# Towards a Wireless Aircraft

## Propagation, Antennas and Radio Standards

David C Hope

PhD

University of York

Electronics

April, 2011

**Abstract**

The replacement of many of an aircraft's wired interconnects with wireless data connections would offer a number of benefits. These include: weight reduction, easier maintenance, faster aircraft design, simpler retrofitting and greater redundancy. However, attempting to do this produces a lot of challenges and risks because data connections on an aircraft must meet very high standards for reliability. This PhD Thesis lays out the main challenges of the propagation environment, explaining that large sections of an aircraft may be highly resonant and explaining the difficulties this presents for wireless receivers.

Modern wireless standards are explored and the question of how suitable they would be for this application is addressed. Two standards are recommended: 802.11n and ZigBee and these are tested in a variety of environments such as a reverberation chamber and avionics bays in a Tornado, to demonstrate the levels of time delay spread needed to prevent a wireless system from functioning. In particular, the importance of the statistical distribution of the channel is emphasised and it is shown that Rayleigh fading of the signal is expected, even over short distances. This has implications for system performance when the channel is time varying, such as in an aircraft wing. A variety of approaches are suggested for reducing the probability of system failure by adding redundancy in various forms to boost link availability.

As well as the detailed study of propagation in an airframe some of the general issues of a wireless aircraft are addressed such as data rates, what types of system could be made wireless and what types of data bus would wireless be replacing. As this work was done under the Flaviir UAV program, some of the collaborative work is presented showing the design of a novel patch antenna written onto the aircraft skin by a genetic algorithm.

# Contents

# List of Figures

# List of Tables

# List of Symbols

$\phi$ Carrier phase

$r(t)$ Received signal as a function of time

$r_i(t)$ Received signal at antenna i in a MIMO system

$h(t)$ Channel impulse response

$h_{ij}(t)$ Channel impulse response between transmitter j and receiver i

$\mathbf{H}(t)$ Channel impulse response matrix

$\eta$ Noise vector of noise powers (in MIMO comms context)

**E** Electric field

**H** Magnetic field

$\mu_0$ Permeability of free space

$\epsilon 0$ Permittivity of free space

$\nabla$ The gradient operator

$\nabla\cdot$ The vector calculus divergence operator

$\nabla\times$ The vector calculus curl operator

$E_{zmnp}^{TM}$ The electric vector z component for the mode with indices a,b,c in the x,y,z variations with time

$k$ The wavenumber (electromagnetics)

$\omega_m np$ The angular frequency of a mode in a three dimensional cavity

$V$ Volume of a cavity

$S$ Surface area of a cavity

$\delta$ Skin depth

$c$ The speed of light in a vacuum

$D_w(f)$ Modal density in a cavity as a function of frequency

$\alpha$ Azimuth angle

$\beta$ Elevation angle

$F(\Omega)$ Plane wave arriving at a point with an angle determined by $\Omega$

$i$ The square root of -1 (when not used as an index)

$f(E)$ Probability density function for random variable E (when used in a statistical context)

$\sigma$ The scale parameter for a Rayleigh distribution (when used in a statistical context)

$Q$ The Q-factor for a resonant mode or average Q-factor of a reverberation chamber

$\lambda$ Wavelength

$\langle X \rangle$ Mean or sample average

$S_{21}$ An S-parameter, also has indices 11,12 and 22

$R(j\omega)$ Fourier transformed received signal

$T_x(j\omega)$ Fourier transformed transmitted signal

$a_n$ Amplitude of the nth component in a multipath signal

$\overline{\tau}$ Mean delay for a received signal

$\sigma^2$ RMS Delay spread (in a communications context)

$P_b(E)$ Probability of a bit being received in error

$Q(x)$ The Q function, a special name for a cumulative Gaussian

$\gamma$ Signal to noise ratio (Communications context)

$\gamma$ Wavenumber, like $k$ but often used instead in materials (Electromagnetics context)

$M_\gamma(s)$ The moment generating function for the random variable $\gamma$

$R$ Antenna separation in the Friis equation

$G$ Antenna gain

$Z_0$ Characteristic impedance

# Acknowledgements

I would like to thank EPRSC and BAE Systems for their part in financing and backing the Flaviir Project without which this work would not have been done.

Thanks also to my supervisor John Dawson for his help and to Mark Panitz of Nottingham University (now at BAE) for always being of assistance during the project.

# List of publications

- Panitz, M.; Christopoulos, C.; Sewell, P.; Hope, D. C.; Dawson, J. F.; Marvin, A. C.; Sato, T.; Fearon, E.; Watkins, K. G.; Dearden, G.; Harley, C. D. & Crowther, W. J., The opportunities and challenges associated with wireless interconnects in aircraft, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aeros, 224(4):459-470, ISSN 0954-4100

- J F Dawson, D C Hope, M Panitz, C Christopoulos, WIRELESS NETWORKS IN VEHICLES, IET seminar on Electromagnetic Propagation in Structures and Buildings, 1-6, London, 9781849190053, December, 2008

- J F Dawson, D C Hope, M Panitz, C Christopoulos, Fly by wireless: Evaluation of the Zigbee radio system for in vehicle connectivity, EMC UK, 93-99, Newbury, October, 2008

- Mark Panitz , Christos Christopoulos, Phillip Sewell, David Hope, John Dawson, Andy Marvin, Modelling wireless communication in highly-multipath low-loss environments, EMC Europe 2008, 709-714, Hamburg, September, 2008

- David Hope, John Dawson and Andy Marvin, Mark Panitz, Christos Christopoulos and Philip Sewell, Assessing the performance of ZigBee in a reverberant environment using a mode stirred chamber, IEEE International Symposium on Electromagnetic Compatibility, Detroit, 978-1-4244-1699-8, August, 2008

# Author's declaration

I am the sole author of this thesis. All the work described in this thesis was my own except where specified. The work was carried out as part of a joint 3 year research project but the work presented here is my own research. The antennas work of Chapter 9 involved very close collaboration with Liverpool university as has been made clear. Also in Chapter A one of the simulations was done by my supervisor and the genetic algorithm based software was provided by a colleague and both these are also explicitly stated in Chapter 9. In the antenna efficiency work of Appendix A some of the measurements were carried out jointly between myself and my supervisor although all the analysis is my own work and I was involved in all the measurements.

# Chapter 1

# Introduction

Over the last ten years the computer and consumer electronics industries have been transformed by a move away from mains powered, cable connected static devices to battery powered, wireless and internet connected mobile devices. A major part of this change has been the rapid adaptation of wireless technology ranging from 3G cellular networks to wireless LAN standards to Bluetooth and even wireless HDMI. Meanwhile, whilst the consumer electronics industry converges towards almost exclusively wireless connectivity the aviation industry continues to use wired, powered devices exclusively. This raises the obvious question: could avionics benefit by adopting the same or similar wireless standards to those seen in the consumer arena.

Our question is a very wide ranging one; there are many different types of avionic system on an aircraft working at many different data rates and requiring many different levels of robustness. This report tries to make a start at tackling the question by looking at a wide range of issues including what systems to make wireless and what the main difficulties would be in making a wireless avionics system suitably robust. Necessarily, given the scope of this topic this report can only focus on a topics and some are ignored. Nonetheless it is hoped that by the end the reader will have a much better understanding of what technologies might work for different applications and what the key challenges will be, providing a springboard to future research. To begin, the advantages of a migration to wireless communication are presented.

## 1.1 Why go wireless

The use of wiring in aircraft presents a number of issues in the areas of weight, cost, safety and maintainability. A very high level of complexity in the wiring can be seen in figures 1.1 and 1.2 It would be desirable if it could be completely or partially replaced by a wire free solution. The

Figure 1.1: Some of the 20km of wiring on the JSF

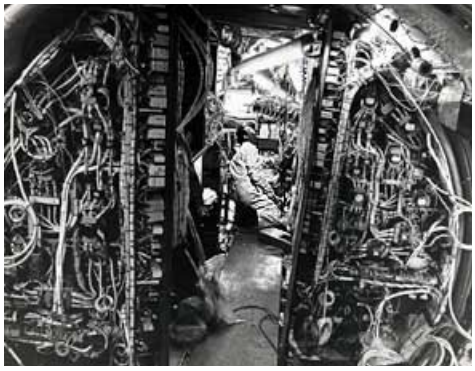

Figure 1.2: Wiring on concorde

main advantages of making an aircraft's wired avionics links wireless are as follows:

**Weight savings** For every wire that is removed, the weight of the aircraft is reduced and thus the engines have to generate less lift and fuel is conserved. For example a Boeing 747 has nearly 200km of wiring[1] adding approximately 1600kg of weight.

**Economic** As well as fuel savings, the cost of cabling is removed and the difficulties in its installation which is time consuming and expensive, as demonstrated by the delays of the Airbus 380 aircraft.[2]

**Wear and arcing** Over time the vibrations and movement in an airframe can lead to damage to cables [3]. This in turn may result in either a link failing or worse a fire breaking out due to arcing. Neither of these issues should arise with completely wireless links.

**Easier maintenance** If new systems must be retrofitted or a device needs replacing it can be done quickly without any changes to the wiring. This again saves money.

**Redundancy** Wireless networks can in principle add more redundancy to an aircraft because rather than using 2 or 3 buses there may be many potential ways to route data without a single point of failure.

One can see that there is a very strong case for moving to a wireless aircraft, provided that regulatory and safety concerns can be resolved. In many different ways costs will be brought down in aircraft design and in operation. At the time of writing the cost of fuel is continuing to rise with an ever increasing demand and instability in oil producing countries which has a direct impact on aircraft operating costs. In the commercial aircraft sector, this can be partially recovered by increased fares for passengers. In the military sector it is just additional cost that can't be avoided. Therefore there is a strong argument for moving quickly towards lighter aircraft that minimises fuel consumption. Part of this will be the move to UAVs but the use of wireless systems that cut down on wiring can play a significant role.

### 1.1.1 Aims

The wireless aircraft topic is a very broad one including areas such as avionics, electromagnetics, propagation, network theory, and general radio and communications technology. There are many potential ways to approach the problem and many different scenarios to design for. Within the scope of a PhD only a limited subset of these can be considered. Therefore at the heart of this investigation are the following aims:

1. Assess the feasibility of a making an aircraft wireless
2. Assuming feasibility, select the technologies best suited to this application
3. Understand the primary difficulties associated with wireless propagation in an airframe
4. Assess the performance of the selected technologies in an aircraft

No assumptions or initial hypothesis were made from the start, because this is a completely new topic. It was intended though, that by the end, some solutions to the problem of making a reliable, robust fly-by-wireless system would have been discovered and that a direction for future research in this field would be much clearer.

### 1.1.2 Document Outline

Throughout this project one particular technology, ZigBee, is used as a basis for understanding the propagation issues in an aircraft. Alongside this other topics will be covered but in less depth due to time constraints. Important areas such as network topologies will not be covered although they were in the original aims when the project began. Instead this thesis

concentrates more on the physical layer and the propagation model because this proved to be a significant body of work in itself, and it is difficult to plan a network architecture without knowing the reliability and data rates of the underlying link.

The document is structured as follows:

**Current aircraft avionics** A brief discussion of the current state of avionics and inter-system communications links, including the problems with the current systems and how a wireless system could overcome those problems. It is necessary to know the requirements of a wireless system before selecting what types to use.

**Radio options** Off the shelf radio systems that might be suitable for the task of flight control are discussed. By using an existing technology the workload is minimised and research can be concentrated on the difficulties specific to aircraft and avionics as opposed to more general communications issues.

**Propagation** Propagation in an aircraft is potentially very different to that in which wireless systems normally function and so investigating this was the major task in this PhD project. The work is divided into 3 smaller sections:

**Propagation theory** The problem of multi-path is discussed from an joint communications and electromagnetics perspective focussing on how multi-path will arise in aircraft and then how it might impact system performance inside an aircraft

**Measurement of propagation** The techniques for measuring the propagation characteristics are presented

**Results and analysis from propagation measurements** Results and analysis from the propagation measurements are provided including single cavity, multiple cavity and noise measurements.

**Antennas** Work done in collaboration with Liverpool University on patch antennas is presented. These are antennas using a novel direct write based manufacturing technique

**Conclusions** To conclude all the work of the previous chapters is drawn together to show what is now known about making an aircraft wireless and what remains to be done.

## 1.2 Research context

All the work described in this PhD Thesis was carried out as part of the BAE/EPSRC Flaviir project and specifically as part of the "Towards a Wireless Aircraft" subproject led by York University and done jointly in collaboration with Liverpool, Manchester and Nottingham Universities. The Flaviir project was setup for the development of novel technologies for UAVs and so the Wireless project had a slight bias towards these too. Wireless technology is especially relevant to UAVs because of the cost pressures on UAVs, which are seen as more disposable than traditional fighter jets and mean moving to wireless can have a big impact. The scope of the sub project was very wide covering topics ranging from antennas to flapless actuation and computational EM modelling. This wireless project was not following any previous BAE/EPSRC project and was not based on or focused around refining previous work. Therefore the aim was to keep the project scope wide, taking an early look at many issues, in the author's case in the areas of propagation, antennas and wireless standards. As a result the chapters that follow do not feature the same level of depth as might be found in a typical PhD thesis but instead cover a much wider range of research areas in an attempt to gain an understanding of where the real work lies in making an aircraft wireless. To begin we shall look at the basics of avionics and the wireless systems that might replace them.

# Chapter 2

# Modern avionics requirements and wireless communications systems

## 2.1 Introduction

To begin, the current state of avionics is considered, because in order to assess what a wireless system must do it is necessary to know what it would be replacing. Avionics can be broken down into different system types and the bus standards that are currently used to communicate data within and between those systems. An avionics system might be a health monitoring system which could be broken down into a set of sensors and processing units communicating at particular data rates on particular bus types. Such a system may have to communicate all its aggregated information to other systems like the cockpit displays.

Having considered the application areas and data requirements, the chapter moves on to the subject of wireless communications systems. The current state of sensor and general wireless networks will be described from the perspective of how they would be applied to an aircraft's systems. There is a particular focus on COTS technology and how current standards based solutions can perform. As will be described later, the use of COTS is an increasing trend in the avionics area because it is more cost effective to make use of a current technology if it is suitable for the job than to develop a new system from scratch.

## 2.2 Avionics

This section presents a very brief overview of modern avionics, looking at what devices constitute an avionics system, what the typical data rates might be, and how data is transferred between the various systems, sensors and actuators, for example the bus standards. Included is a discussion of modern and advanced standards found on aircraft such as the Airbus A380 and the F22 Raptor because these are at the most complex end of the avionics spectrum and present the upper limits on what a wireless avionics system would have to achieve. This is compared to the simpler case of a UAV which is expected to be a major application area for wireless systems.

## 2.3 Wired aircraft avionics

### 2.3.1 Aircraft categories

An appropriate starting point is to consider the various classes of aircraft. For our purposes they are:

- Commercial
- Manned military
- Unmanned aerial vehicles (UAVs)

It will be assumed that fixed wing aircraft are being discussed generally but most of what will follow also applies to helicopters, both civilian and military. This PhD Thesis is the output of work done towards the Flaviir project which was primarily concerned with UAVs. However, it will consider all types of aircraft because the benefits of cutting weight and maintenance with wireless apply to all aircraft and it would be advantageous if a technology could be applied across all aircraft types, not just specific UAV types.

Electronic systems are required in all aspects of aircraft operation and they can be roughly grouped into a number of categories. Starting with civilian aircraft the avionics systems may be divided up into a number of groups [4]. These are the:

**Flight control system** - this is concerned primarily with two objectives. First, to ensure aircraft dynamic stability whilst also ensuring fast response times. For example, to dampen

any oscillations that may make an aircraft difficult to fly or which could be uncomfortable for passengers. This part of control is often referred to as the inner loop control. Outside this is the outer loop which is concerned with keeping to a particular attitude and heading. Part of flight control is about taking pilot input and delivering optimal response but it also relates to automatic pilot functions.

**Flight performance management system** - encompasses functions such as guidance and navigation and also calculation of optimum altitudes and speeds at different points on the route. Partly this feeds data to the pilot, but again it can form part of an autopilot system.

**Engine control** - deals with areas such as maintaining smooth intake for changing conditions, ensuring correct relationship between high and low pressure compressor shafts, reheat control and also temperature monitoring.

**Utilities management system** - this can be divided up into areas such as brakes, hydraulics, fuel management, emergency power, temperature control and fuel gauging.

**Health and usage management** - includes a lot of engine monitoring for variables including speeds, torques and vibration levels but it also includes strain gauges mounted around the aircraft.

**Entertainment systems** - only on commercial aircraft carrying passengers (as opposed to cargo). Although it only applies to one aircraft category it is significant because it is one of the most data hungry applications on a modern aircraft. Modern video allows streaming rates of 250kbps upwards[5] and this is on demand for each passenger, resulting in a worst case scenario of over 100Mbps for a 747 or A380.

When considering military aircraft the key difference is the lack of passenger related systems and the very significant addition of electronic warfare systems [4]. There are a number of features of electronic warfare systems:

**Defensive aids and counter measures** - includes antennas and radar systems to detect hostile threats such as heat seeking missiles. Also included are flares chaff, decoys and other protective measures used in combination with evasive manoeuvres.

**Weapons deployment** - initialisation, guidance information, release from the weapons carriage.

| | | |
|---|---|---|
| Air Superiority | Ground Attack | Strategic Bomber |
| Maritime patrol | Battlefield surveillance | Airborne early warning |
| Electronic warfare | Transport | Refuelling |

Table 2.1: Mission types for military aircraft

**Feeding of weapon** and defensive systems information to the pilot and integration with flight control systems. Included in this are HUD views in the cockpit.

Moir and Seabridge[4] describe a variety of systems architectures for different mission types which are given in table 2.1 Some of these will have weapons systems and their associated buses, others will have more sensors and imaging equipment. Some may have standard displays and some HUDs. The complexity of the flight control systems and mission planning will differ and the speed of computing and performance levels required.

#### UAVs

The final category of aircraft is the Unmanned Aerial Vehicle (UAV). UAVs have particular importance because of their smaller size and the fact they are intended to be manufactured in greater volumes that manned aircraft. Keeping cost and weight down is important and so moving to wireless interconnects can have a significant impact here. Even a large weapon carrying UAV such as the Predator only has a single 4-cylinder engine. In terms of avionics UAVs will generally be simpler than a manned jet and the lack of a pilot can reduce the redundancy[1]. Although there is no cockpit requiring large amounts of information streaming to it, this still needs transmitting to a control station remotely.

One of the themes that was seen through the Flaviir conferences (a yearly event for reporting on all parts of the project) was that a UAV had to be a highly flexible vehicle. At one time it might be used for battlefield surveillance but at other times it might be delivering a weapons payload. For an aircraft such as the Typhoon or F35 there is some flexibility in the missions they carry out but overall it is a situation where a small number of very expensive aircraft do very specific tasks. UAVs are intended to exist in larger quantities with greater flexibility with a given model of UAV being expected to perform a range of the missions types described, over its lifetime. This in particular provides a strong argument for a very modular and ideally wireless aircraft because systems can be easily removed from and placed on the UAV without significant

[1] This is still high as UAVs crashing is unacceptable even if there is no pilot risk

rewiring costs.

There are some fundamentals that can be found in any UAV avionics system. Pastor, Lopex and Royo[6] consider a standard architecture for civilian UAVs and identify the flight control system and the payload/payload controller as the fundamental features on a UAV designed for sensing and video. Included in these are radars, ground to air radio links, sensors actuators, video recording devices, thermal sensing and imaging devices etc. It can be seen that there is still a significant amount of data to be transferred and there are many sensors that might potentially be made wireless. In [7] the significant role of COTS electronics and systems in UAVs is emphasised as well as the likelihood of a UAV requiring video sensing. This is in fitting with the discussion later in this chapter, emphasising the need to use COTS to reduce costs. Dittrich and Johnson[8] discuss similar basic features to Pastor[6] but give particular attention to the navigation system. For any UAV operating at a distance from its ground station it is necessary to have an autopilot system that can navigate reliably without user input (perhaps for some time if a link went down.) Such a system will require links to the flight control system, to other aids such as GPS and also a radio link down to the control station on the ground.

It can be said that despite a UAV being small and cheap relative to conventional aircraft their requirement for autonomy and their use for monitoring and sensing requires a non trivial sensor and avionics system.

**Summary**

This short summary of some typical aircraft types indicates not just the variety of different missions types but also the sizeable number of sub systems on a modern aircraft and thus the high density of communications interconnection that is required. There is no such thing as a typical avionics system with typical data rates and links but such systems can be broken up into the low data rate health monitoring type systems through to the high date rate information streams going to the cockpit or control station.

### 2.3.2 Examples of sub systems and their connections

Having discussed at a high level the different aircraft roles, this subsection now describes two, more concrete examples in order to see the extent to which sensors and electronic monitoring is required in a modern jet.

**Engine control and monitoring**

Figure 2.1 shows the main components in an engine. Reference [9] gives the following as the principle areas to be monitored:

- Speed, temperature and pressure, all with multiple probes
- Exhaust gas temperature and exhaust gas pressure ratio
- Turbine case cooling
- Engine start
- Fuel control
- Turbine blade cooling
- Inlet guide vanes
- Bleed air controls

Thus within an engine there are a large number of probes, all connected to processing units and all of these will require connections to the engine controller. The engine controller itself will need to be connected to other processing units, again with data and power cables. Engine systems are also unique in that they will offer a continually varying propagation environment which might be difficult for a wireless network and may provide an interesting research topic in itself. They provide a good illustration of the level of electronic monitoring in a modern aircraft.

**Fuel tanks and monitoring**

An example of a fuel management system can be seen in figure 2.2 which shows a simplified representation of a 777 system. The left and right tanks have 22 sensor units each and the centre has 15. In total there are 128 signals going to the fuel quantity processing unit. There are also additional signals going to the P310 Standby Power Management panel. Therefore there are a large number of probes and signals and the data rates are low. This is exactly the sort of application where wireless sensor nodes could be very useful.

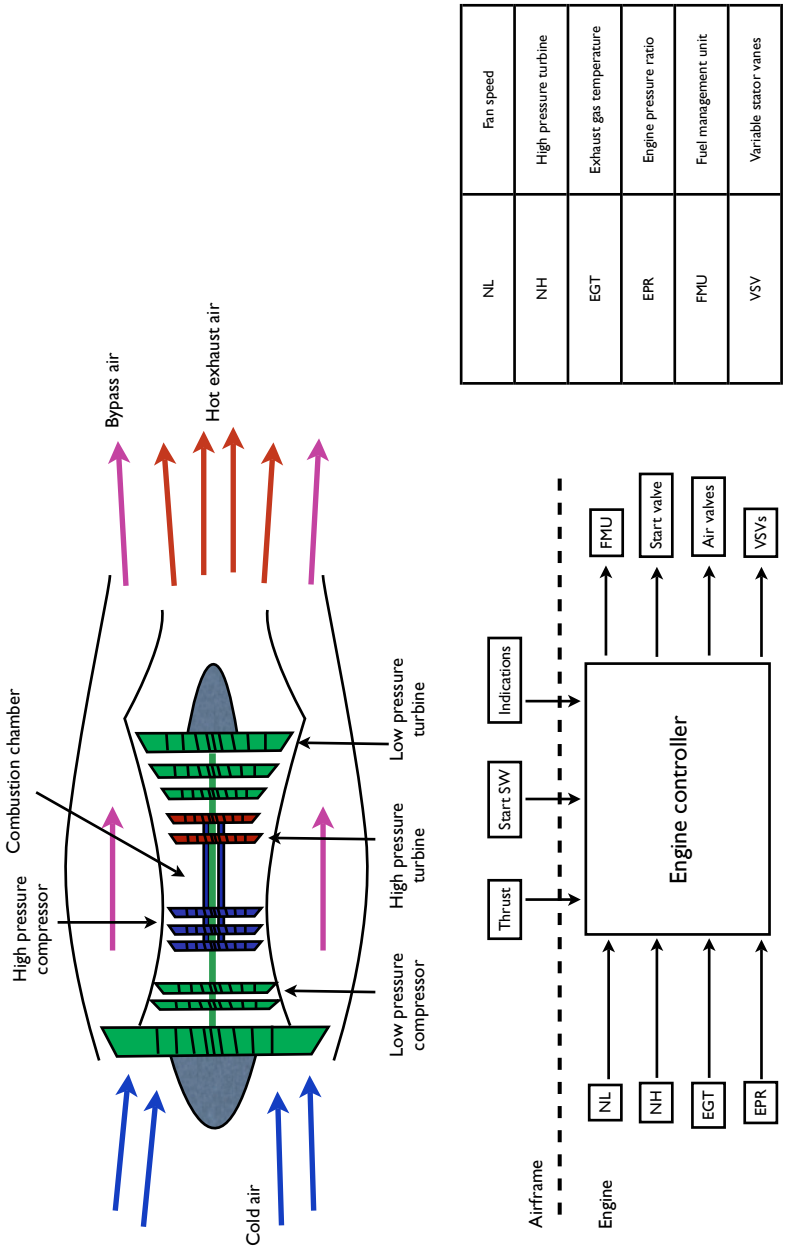


Figure 2.1: Main features of an aircraft engine and its control unit

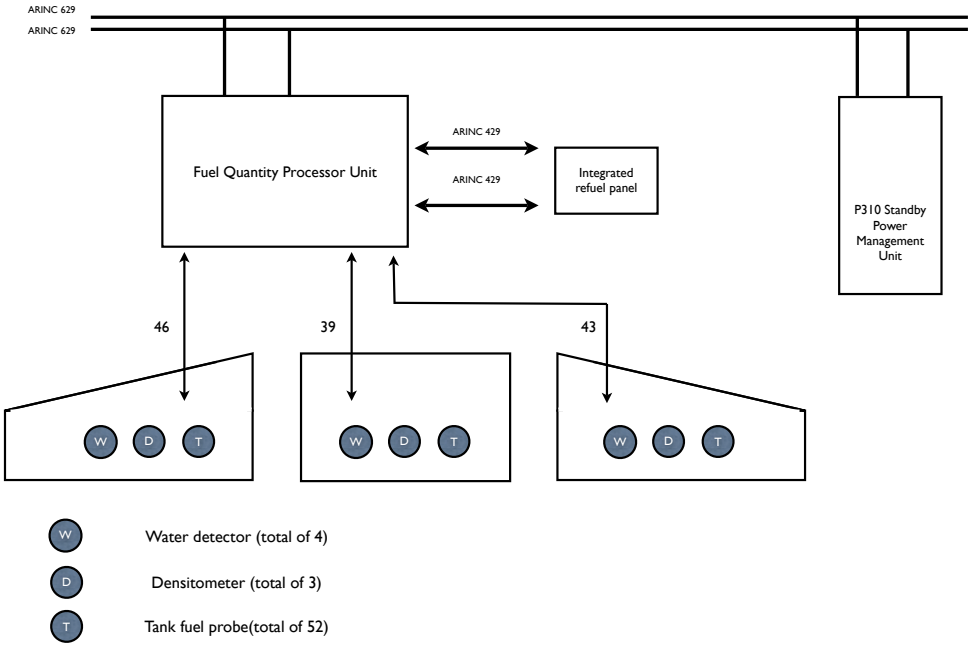


Figure 2.2: Data signals for a Boeing 777 fuel system

## 2.4 Digital link architecture, buses and data rates

Given that the project aim was the replacement of existing wired buses with wireless ones the next area to consider is the buses being used on current aircraft. Digital communications links can be separated into a number of different types[10], as illustrated in figure 2.3:

- Single source, single sink, half or full duplex
- Single source multiple sink (like ARINC 429)
- Multiple source, multiple sink (like ARINC 629, MIL-STD-1553)

The first of these is one we might not expect to find on a modern aircraft and is unheard of in the computer networking arena but it does still exist, for example on Eurofighter. It may be found where all that is required is a simple wire connection between devices, for example a light coming on when a particular system is on. These links are used because they work and are proven to be reliable and there is no need for anything more complicated. The closest

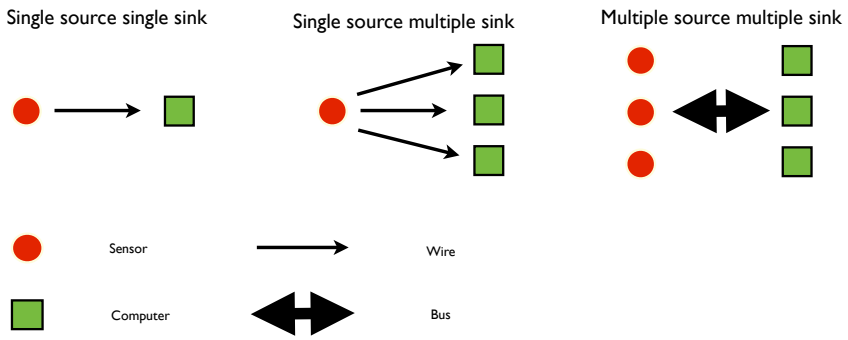


Figure 2.3: Different connection types in an aircraft

type of wireless system would be one in which every point to point connection had its own radio channel but this is a very inefficient use of bandwidth. Therefore in the remainder of this Thesis only bus technologies are considered.

Digital buses are increasingly used on wired aircraft because of their more efficient use of resources and greater flexibility. Multiple source multiple sink is most common in modern communications systems but the single source multiple sink is common because in aircraft there are not 2 computers communicating, or phones or similar but a sensor broadcasting to redundant computer systems.

The following subsections describe a number of interconnection/bus standards in terms of link medium and capability. Further information may be found in [10]

**ARINC 429**

Arinc 429 is a civil standard which acts as half duplex, single source, multiple sink. This leads to a large number of connections being required where there are a large number of nodes. It can be found on the Boeing 767 and Airbus 320. Data can be transmitted between the rates of 12 and 100 kbits per second. Line signalling is done using bipolar return to zero. The weakness of this system is that when new equipment has to be added it is not a simple case of plugging one line to a bus; if both ends of a line need to transmit then extra cables and transmitters are required. With high levels of interconnection this is not practical.

**MIL-STD-1553**

This is much more efficient than the ARINC standard because it is multiple source, multiple sink and when extra systems need to be added to the bus the only change that must be implemented is updating the bus controller software so new devices can easily be added to the bus. It is typically implemented using twisted pair cable but an alternative medium may be used. The data rate is higher than ARINC 429 at 1Mbps but there will be overhead on top of this as any transfers have to travel via the bus controller. It is possible to have a hierarchical architecture whereby multiple devices are linked by a bus interface unit. Remote terminals may also function to connect to multiple sub-systems. It started off as a US military standard but is now used elsewhere; UK helicopter avionics, for example.

**JIAWG**

This refers to an architecture rather than a bus specifically but it includes various fibre optic data buses for communicating data from sensors to processing units, and for transmitting video data to displays. To communicate between the processing racks and other units a high speed data bus is used, which is also fibre optic in a token ring topology and operates at 80Mbps. It was expected that this would be used on the F22 and other aircraft but in fact it has been largely replaced by the Fibre Channel technology described below. However it is given mention here because it gives an idea of the maximum data rates used within recent, very high specification military aircraft and thus what requirements a wireless network must meet.

**Fibre Channel (FC)**

Fibre Channel has it's origins in the computer industry with networking and SCSI disks. After the earlier military standards just described, military specific bus technology stalled in comparison to the computer industry, which has had significant money behind it. Therefore, for the latest very high speed technology, COTS systems are being utilised on aircraft. As the name implies it is often run over fibre optic links although it doesn't have to be. Link speeds are from 1Gbps upwards. High level protocols are available such as FC-AE-1553 which allow FC to communicate easily with existing MIL-STD-1553 standards.The StarFabric implementation of FC was selected as the bus technology of choice for the Joint Strike Fighter (F35)

#### IEEE 1394 (Firewire)

Also used on the F35 (for vehicle management systems) this is the commercial electronics standard for connecting video cameras and other high bandwidth devices to a computer. Data rates are up to 400Mbps in the original variations of the standard.

#### ARINC-664/AFDX

The latest commercial aircraft standard as used by Airbus for the A380 is known as Avionics Full Duplex Switched Ethernet (AFDX)[11] and ARINC-664 is based on this. It works on top of IEEE 802.3, i.e. standard wired Ethernet and wraps it with an additional protocol layer which can guarantee timing constraints and add redundancy. It is made up of switches and end systems and is a network rather than a bus. Over the wire data rates are up to 10Mbps or 100Mbps since it is Ethernet based.

#### Modern integrated systems

Near to the cutting edge of aircraft technology is the F22. It uses an integrated avionics architecture [12] as opposed to federated architecture (where each system is separate with its own processor) so out of the above technologies it is closest to JIAWG. Radar, communications, navigation, identification, electronics warfare and sensor control are all integrated. For such integration it is necessary to not just share information in software but to design hardware in a way that utilises COTS technologies with a common operating system and common buses. For processing 2 Common Integrated Processors (units with many CPUs) are used and each of these has 66 slots of which 19 in one and 22 in the other, are used. Each module is not used to capacity so there is room for growth here as well as the potential for extra modules. Furthermore there is space for an additional CP. Therefore by using this uniform consistent design, the ability to increase processing power or redundancy is made easier. The CIP is used for features such as gun control steering, electronic warfare and communications, navigation and identification. External connections use optical links for high data rates. The F22 also has an intra-flight data link whereby large amounts of information on one aircraft is aggregated and shared with other aircraft. It might be expected that this is a trend that will grow, particularly for groups of UAVs. See [13] for more information on the F22.

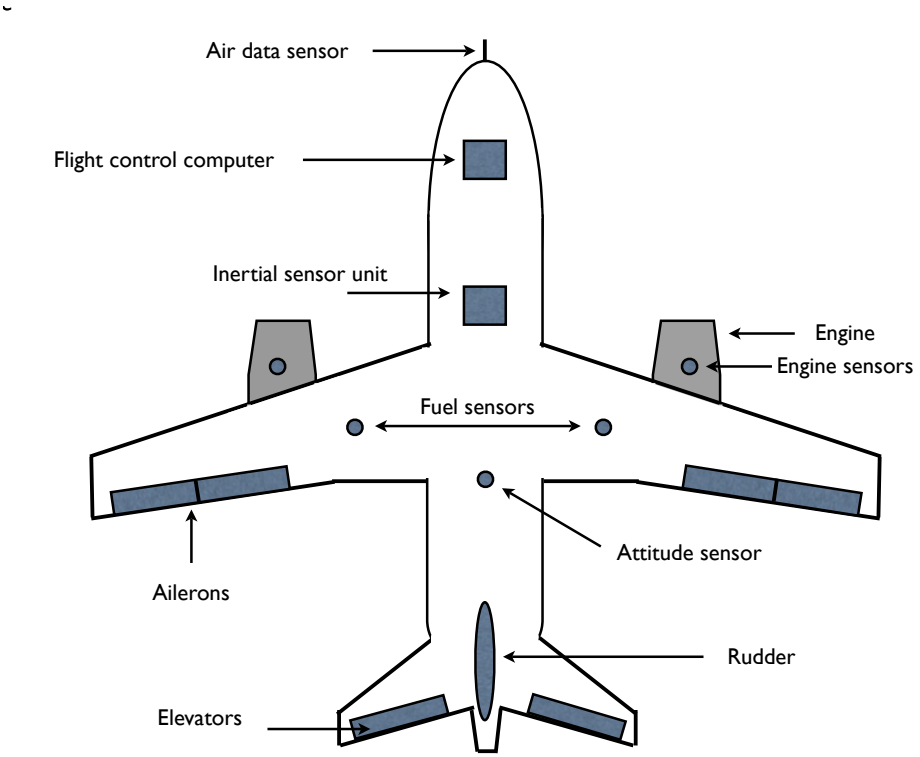


Figure 2.4: Sensors and actuators on an aircraft

### 2.4.1 A flight control case study

To conclude this section on the state current day avionics and data rates the following subsection will look at the requirements for flight control for a small UAV like aircraft, from the perspective of the sensing bandwidth requirements.

Flight control in particular is considered since it is one of the fundamental real time tasks of any avionics system.

#### Sensors and actuators for flight control

An example of the main control surfaces on an aircraft and the communications links required is given in Figure 2.4. Ailerons, elevators and a rudder provide roll, pitch and yaw to the aircraft and the flight control computers send the commands to these control surfaces. It will be assumed no wireless link is needed between computer systems and there is a single computer unit. For larger aircraft there might be a situation closer to that in the Airbus A320, for example, where there are seven computers divided into elevator aileron, spoiler elevator and flight augmentation computers [4] but in a smaller UAV this might not be necessary nor practical. The

sensors may be divided into two groups encompassing air data sensors and inertial sensors. Air data sensors measure the following parameters:

- Indicated airspeed and true airspeed
- Air density and air pressure resulting in barometric height
- Outside air temperature
- Angle of attack
- Mach number
- Rate of change of height

Inertial sensors measure these parameters:

- Roll pitch and yaw rates. These are measured by rate gyros
- Roll, pitch and yaw attitudes measured by attitude sensing gyros
- Accelerations in the x,y and z directions.

This data is used both in the inner control loop for aircraft stabilisation and the outer loop for altitude and heading information. Therefore a flow of data is required from the sensors to flight control computers and from these to the actuators.

Data rates depend on the flight dynamics as these force the sampling rate but also depend on the control system design methodology [2] so there is no absolute number that can fit all aircraft types. For now typical sampling rates in flight control are assumed as being in the order or 30 - 100 Hz and for the purposes of this calculation the upper bound will be taken. This is in line with the figures found in commercially available sensor units; for example the Memsense nIMU gives a sensor bandwidth of 50Hz thus requiring a sampling rate of 100Hz in order to avoid aliasing and Athena's products such as the Microguidestar and Guidestar 311 specify an update rate of 100Hz. These Microstar units are specifically designed for UAVs.

On the nIMU unit, the output from one of the gyros, and for other sensors is 16 bits per sample. Therefore if accelerometers, magnetometers and rate gyros are considered in their 3 orientations this requires a data rate of 14.4 kbits/s. Without jumping too far ahead it should be

[2] All digital Z domain design, or continuous with bi-linear transform etc

noted that the soon to be described ZigBee system is 250kbps. Initially it may appear that this is not a problem for ZigBee but because that is the over the air data rate and not the throughput which may be half this [14] or a lot worse when many units are all trying to send messages at the same time.

The sensors units described above feature inputs for air data information and this would ideally need to be wirelessly transmitted. A Harco PITOT Static Probe is one type of unit that could be used and this has various interfaces including RS232 at the low end requiring 20 kbps although if one calculates the data rates required by three sensors, one for dynamic pressure, one for static and one for temperature and again assuming a sampling rate of 100 Hz and 16 bits leads to a data rate of 4.8 kbps. Adding in the three actuators described earlier required an additional 4.8 kbps. For the engines, sensors would be required to monitor engine speed and temperature and the fuel delivery system may need to monitor rate of flow into the engine. It is also necessary for the flight control system to send demands to the engines. Making the same assumptions as before the engine communications adds 6.4 kbps of required bandwidth. If monitoring sensors were to send data back from both engines this would increase the data rates further. Figure 2.2 summarises this analysis.

| System | Data rate (kbit/s) |
|---|---|
| Pressure sensing | 4.8 |
| Accelerometers and gyros | 14.4 |
| Actuator links | 4.8 |
| Engine links | 6.4 |
| Total | 30.4 |

Table 2.2: Very rough estimates of flight control sensing requirements

What this analysis shows is that data rates over 30kbit/s are needed for the most fundamental sensing needed in flight control, ignoring for example health monitoring and communications such as radar. This is only for a very simplified analysis and it could be a lot higher. Current wireless sensor network standards like ZigBee are capable of meeting these requirements but the reader should be aware that it be easy to reach their limits in control applications.

## 2.5 Radio

Wireless communications technology is the topic of this section that examines those technologies that might be applicable to avionics. These includes the most popular current commercial off the shelf (COTS) solutions, the general trends in modern wireless networks and how they are suited to use in avionics. One point that must be made here before being fully explored in later chapters is that there is a distinct possibility of the aircraft environment being heavily reverberant because avionics bays are often made of metal. Having strong multi-path tolerance built into a given technology is therefore important for a given standard. In fact a large part of this PhD Thesis is concerned with this issue of resonance in aircraft. Firstly the requirements for a wireless avionics system are explored.

### 2.5.1 Wireless requirements

Fly by wireless performance will be heavily dependent on the choice of radio technology at the *physical layer* level. The resistance to interference, the robustness in heavily multi-path environments and the maximum obtainable data rates are all fixed by the choices made at the physical layer for the radio. Furthermore, the limitations of the physical layer level will in turn affect the level of self configurability, the maximum network throughput for a given topology and the number of nodes in any network. Therefore there are many factors to consider in selection of the most appropriate system. It was decided early on to begin with a standard technology for which off the shelf components could be bought. Use of COTS technology is standard now in the state of the art in current wired aircraft technology where Ethernet and Firewire (see section 2.4) are used. COTS technology brings a number of advantages to the development of a wireless aircraft:

- Greatly reduces costs - no need to develop a whole new radio stack
- Results in a greatly shortened development time
- Allows potential problems to be quickly discovered
- Many development tools, and network analysis tools are already available
- Better testing and reliability because standards systems are widely in non aerospace applications

From the research perspective, by using standard technologies one can concentrate on finding out what the main difficulties are in a fly by wireless system rather than getting stuck in the small details of exactly how the radio should be implemented and missing the fundamental. Only if our research into COTS technologies shows that they fail to deliver the required performance should our research divert to custom, new technologies. It's also arguable that COTS can offer superior data rates and performance within a given bandwidth due to the major drive in the consumer electronics arena to increase wireless network data rates.

There is one major disadvantage with COTS based systems for this application, which doesn't occur for wired systems and that is one of radio bands. The most popular band, the 2.4GHz ISM band is only 100MHz wide meaning that a system like IEEE 802.11b has only 3 non overlapping channels. Although non ideal it isn't critical problem because, as is the case with 802.11a and 802.11g, one can use an almost identical protocol stack and modulation scheme but at a different frequency. There will have to be changes in the radio front end but this is a small cost compared to that of developing a radio system from scratch. In any case, some hardware changes will have to be carried out anyway so as to ensure the hardware meets military standards. One should also note that COTS wireless systems are, in the author's opinion, rather more complicated than they need to be, which is the natural result of a standard that has had to meet the requirements and politics of many interested parties. One would expect that a military specced implementation of a COTS system might only used a subset of the complete software specification. IEEE 802.11g and n for example have to be backwards compatible with 802.11b which is unnecessary if all systems are to be 802.11n

The properties that are desirable in any wireless aircraft system are as follows:

**Low cost** There are hundreds of sensors on a large aircraft and so its important to keep the price as low as possible per device

**Low power** Ideally many radios would be battery powered or use energy scavenging

**Scalable** It should be possible to add in new equipment to a system with minimal effort, and a system that can work on a UAV or light aircraft should work on a larger aircraft.

**Secure** It should not be possible to intercept and read or change information being transferred

**Robust against interference and multi-path** Many aircraft equipment bays are metallic and

reverberant. Further, there are many different devices in such bays which may lead to interference. Any system must have strong immunity to these effects

### 2.5.2 Features of current systems

If COTS systems are seen as the best place to begin in the implementation of a wireless aircraft then the first question to ask is what COTS technologies are on the market. This section introduces the most common technologies and describes their main features. All are standards based and offer a different combination of power, data bandwidth and flexibility and so each has advantages and disadvantages for this application. In the relatively new wireless sensor field there are systems available that are non standard but they won't be considered so as to concentrate on technologies with multiple implementations and significant information in the literature. With closed systems it would be possible to evaluate their performance but not to work out why they achieved that performance.

#### ZigBee

ZigBee was finalised in 2005 and it was specifically designed for wireless sensor networks, making it an attractive option for this application. It is based on the IEEE 802.15.4 standard [15] for physical layer and media access control (MAC) and is this considered alongside the ZigBee standard [16] which is built on top of it. It is specifically designed for:

- Low data rate applications - the maximum rate is only 250kb/s because sensor applications generally require much lower rates than, for example, wireless internet.
- Low power consumption - the low data rates and low duty cycles enable them to be battery operated and not require replacement for over a year in many sensing applications. When transmitting, a module would typically be using around 100mW but under 10$\mu$W when sleeping. For a low duty cycle application a high percentage of time would be spent in sleep mode.
- Multiple topologies - can be configured as point to point, tree or mesh which fits well with the different bus configurations given in section 2.4
- Self configuring and repairing - therefore suitable for ad hoc networks and very scalable. For example a network might start with 5 nodes and then this could be increased up to

20 or 30 and the network would adapt itself to this new configuration (provided each node did not have too large a bandwidth requirement). 65000 nodes are supported in a network and multiple channels can be used for separate networks.

- Robust, 128bit AES encryption can be utilised, if data security is required.
- Spread spectrum modulation (DSSS is applied prior to the OQPSK modulation scheme) offering basic resistance to jamming, detection and multi-path.
- The protocol stack is designed to be about one quarter of the size of that of alternative standards such as Bluetooth. This is important for memory considerations in the network nodes.
- The channel access scheme allows for guaranteed timeslots meaning there is support for real time data. For example if the flight control system required roll acceleration every 10ms then it could be given a slot with this periodicity.

This short list illustrates that ZigBee has a number of desirable features that should prove useful in a aircraft wireless sensor system. The low power requirements make it less likely that cabling will still be needed for delivering power to the sensors (it may still be needed for actuators) . There are also some limitations in ZigBee which reduce its suitability for the given application:

- Low data rate. This is one of the reasons that ZigBee is low power. However, whilst the low data rates are sufficient for typical sensor applications (e.g. building climate control systems) it is not suitable for streaming a lot of data to flight control computers or displays.
- Multi-path resistance and error correction is simplistic compared to other standards.
- Low power may prove unnecessary if a technology such a direct writing (see Chapter A) of power lines is utilised or if power needs to be delivered to actuators or even the sensors. It could be better to concentrate on a higher power system that is more robust and less susceptible to interference.

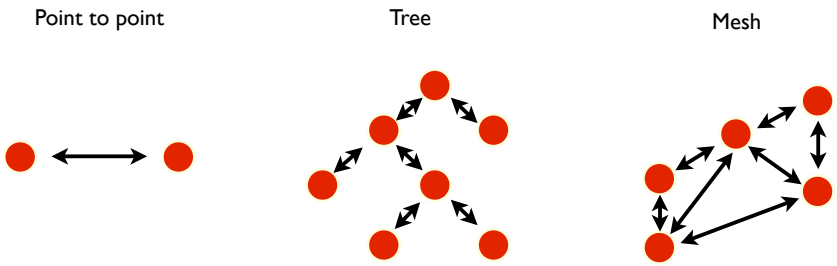


Figure 2.5: ZigBee topologies

**Bluetooth**

Bluetooth technology was a relatively recent standard at the time of writing (2010) and had become established over the last 8 years. It is not developed specifically for sensor networks but it has properties that lend themselves well to that application. It is specifically designed for data transfer between different digital devices such as PDAs and laptops, for data transfer and gaming between mobile phones, cordless mice, keyboards etc. It is also designed to be low power (though not as low as ZigBee). Bluetooth has been standardised as IEEE standard 802.15.1 [17] It has the following technical features:

- FHSS (spread spectrum) making it marginally tougher to detect, jam and adding multi-path delay resistance.
- Transmit powers of 1mW and 100mW so potentially low power meaning there is a possibility of using battery power or energy scavenging.
- Time division multiplexing is used (TDM) which means that devices can be assigned time slots which is important for meeting time constraints in a hard real time system.
- Data rates can be up to 430kbps on the forward and reverse channel for a device, which is higher than that offered by ZigBee. Alternatively data can be transferred at 700Kbps in one direction and 64kbps in the other. Also, as more devices are added the amount of information a given node can transfer to the master is reduced.
- As well as the full operating power mode there are also 3 other reduced power modes: sniff, hold and park which could increase battery life if this proves necessary.
- Security is provided in the form of identification numbers (PINs) and symmetric key encryption (the 128bit SAFER+ cipher) so that there is authentication and data encryption.

There are significant concerns in using Bluetooth, the main one is with regards to how well it scales in larger networks. It works with a concept of "Piconets" in which there can be up to seven slaves and one master where the master controls access to the channel. There can be up to 256 devices in total but only 8 can be active at a point in time. Despite having a number of attractive features Bluetooth appears to be designed primarily for allowing a small number of devices in the vicinity of one another to communicate. It lacks the flexibility of ZigBee to expand as required, forwarding packets as many hops as needed to the destination. Bluetooth 3.0 addresses some of these concerns but this is not yet a proven technology and was only announced towards the end of the research phase of this project. For further information see references [18] and [19] which offer more detailed introductions to Bluetooth.

**802.11 Wireless LAN**

Wireless LAN is the most mature technology of the standards described in this section and is commonly referred to by its IEEE group: 802.11. At the time this project began (2006) 802.11n was being standardised and previous versions of the standard include a, b and g. Later versions are generally more complex and support higher data rates. The main advantage a wireless LAN has in comparison to the other 2 options is very high data rates which go to over 300 Mbps[3]. In order to achieve the higher data rates this system has a higher power consumption than Bluetooth and ZigBee and greater bandwidth. Earlier it was shown that some avionics systems require high data rates so for these this would be an attractive option. The 802.11 standards offer the following features:

- Spread spectrum (either FHSS or DSSS dependent on which 802.11 version is being used) or in 802.11g, OFDM. All of these offer some resilience to the multi-path fading that is likely to be present in an aircraft but the OFDM in the higher date rate standards should be most effective (this is considered in more depth in chapter 8).

- 802.11n uses the latest MIMO theory in order to use multiple propagation paths as independent channels boosting the effective channel capacity. This is also able to improve system robustness under multi-path.

- Carrier sense multiple access with collision avoidance (CDMA/CA) - similar to Ethernet and can be combined with an operating mode known as Point Coordination Function

[3] though practical throughputs can be closer to half the data rate

which allows a form of guaranteed time slots. Thus channel access is analogous to ZigBee and Bluetooth.

- Maximum power of 100mW set by regulations for radio. A wireless LAN card might typically consume up to 1W. This is too high practically for a battery powered device.
- Range of 20m to 300m depending on the environment
- 13 2.4GHz channels available in Europe of which three are non-overlapping providing bandwidth for at least 3 separate networks in an area without interference. More at the 5GHz band.

The two most significant features of the wireless LAN standards for us are the data rate and the power consumption. Wireless LAN is aimed at high power, high data rates in comparison with ZigBee's low power, low data rates. Version n, in particular has the advantage that it is a MIMO system which is used in order to reach very high data rates, with practical throughputs of nearly 100Mbps which better than Ethernet, prior to the gigabit version. For this application however the reason MIMO is potentially useful is that a MIMO communications system makes use of the different multi-path components of a signal. Where an aircraft exhibits significant multi-path a MIMO system would appear to be ideal. Practically however it may not be possible to use multiple antennas when using lots of small sensors. An additional concern is that a MIMO system requires more complex signal processing in order to extract the signal from the separate multi-path components which adds to cost and also the power requirements.

**IEEE 802.15.3**

This is another IEEE wireless standard [20] which is designed for high data rate WPANs. Rates are currently at 50Mbps and are intended to go beyond 100Mbps. Application areas include the streaming of hi-definition video and wireless "Firewire". Data rates this high may be unnecessary for flight control applications where sampling rates are in the tens of hertz but aircraft require the streaming or large amounts of video and data to the cockpit, in a timely manner and so this technology could be very appropriate for that. 802.15.3 has a lot in common with ZigBee in its use of beacons, guaranteed time slots and AES encryption but in terms of network structure it is closer to Bluetooth than ZigBee, being designed for wireless personal area networks WPANs and having a similar concept of piconets and only a small number of devices connected at once. The combination of high power requirements and small networks

makes it less desirable than alternatives and it won't be considered in any depth in the remainder of this document. Nevertheless it could provide a role for one or two of the most data hungry links between flight control racks and the cockpit in conventional aircraft.

### 2.5.3 Swapping current standards for wireless equivalents

At this point both the current state of avionics and existing wireless standards have been described so the next stage is too see how those standards might be applied to replace the wired systems.

In section 2.4 three of the technologies described are: ARINC 429 and MIL 1553 and AFDX. The former is used more in civilian aircraft and offers low speed and high speed modes with data rates of 12-14kbps or 100kbps respectively. A 757 might have 150 ARINC 429 buses linking different systems. The 1553 bus standard is higher performance and offers a speed of 1Mbps. Finally the AFDX system can operate up to 100Mbps

#### Low data rate buses

The ARINC data rates are comparable with those in recent sensor network standards such as ZigBee and so these 429 buses would appear suitable for replacement. One must note that whilst 100 ARINC 429 could be used in an aircraft, modern wireless standards do not present over 100 wireless channels and so unless the networks are well isolated it might not be possible to do a complete replacement. One would have to look at the utilisation of those links. Where signals are only sent once per minute there would be no problem. Where duty cycles are high careful planning is needed to ensure that different channels are used for separate links operating in an area where they would otherwise interfere.

#### Medium data rate buses

Data rates to match the MIL 1553 standard could be provided by any of the wireless LAN standards (IEEE 802.11) but as already stated this would remove at an early stage the possibility of battery power which is a key benefit of wireless sensor networks. This said, even if there is power by wire, wireless still leads to weight reductions and 802.11 wireless LAN can offer these rates. ZigBee does not appear to offer a suitable replacement but again, this is only the case where full utilisation of the link is needed. Where duty cycles are now, ZigBee might also replace MIL 1553.

#### High data rate buses

At the very top level is the 100Mbps AFDX commercial aircraft system and the Fibre Channel system. There are current wireless solutions that can match the AFDX standard but there is nothing currently operating in the gigabit range. One could potentially reach such levels by making alterations to current systems so that a larger number of channels is used.

#### Analysis

Therefore the current IEEE wireless standards do provide equivalent data rates to the different avionics buses. The only question mark is over the issue of number of channels. Many wires can occupy the same space without interfering but this is not true for a wireless system. On the other hand it is more than possible that most devices only need a small part of the available data bandwidth and so can share the channel anyway. Sensors that do nothing but turn on a warning light will not require all their available data bandwidth.

Where multiple buses operating at high data rates are required, one must look at increasing the number of channels and the total bandwidth beyond the 80MHz found in the ISM band or perhaps use the 5GHz U-NII band. The current state of wireless networks is clearly sufficiently advanced to provide the bandwidth needed on an aircraft through a mixture of ZigBee and Wireless LAN based systems.

### 2.5.4 Applicability to future aircraft

It has now been shown that wireless technology is feasible for current avionics systems but what about those of the future that might utilise higher rates. Part of the Flaviir proposal (of which the wireless project was a part) was to look at novel control methods and as an example Manchester University has done some research into control that does not make use of mechanical moving surfaces but instead manipulates the airflow over the wing. Two such control methods are:

**Separation control** Separation of the airflow from the boundary layer on the wing can be controlled by, for example synthetic jet actuators or micro flaps. This is used for improving lift or for manoeuvre effectors.

**Turbulence or boundary layer control** Concerned with reducing turbulent streaks in the boundary layer and thus reducing skin friction drag.

Possible data rates were suggested for these techniques by Manchester University during Flaviir Wireless discussions. For separation control a control system bandwidth of 1 kHz has been cited which would require a sampling rate of 2 kHz. If 10 actuators-sensor pairs are assumed each requiring an accuracy of 16bits this leaves a data rate requirement of 320 kbps. Alternatively if each pair worked on a separate channel then 32 kbps would be needed.

For turbulence control a bandwidth of 100kHz per sensor was suggested (in project presentations) and would require hundreds of sensors, and may require data bandwidths of the order of terabits per second.

Turbulence control is clearly impossible for a wireless system and would require a wired or opical fibre solution. Separation control could not be handled by ZigBee but is within the realms of wireless LAN.

### 2.5.5 Which wireless standards are best suited to avionics applications

Each technology just described in section was shown to have features useful for wireless sensors with most offering guaranteed time slots and encryption for example. The most useful standard appears to be ZigBee because it is specifically designed for wireless sensor networks and because it is much lower power than the others. If sensors on the aircraft could all be battery powered this offers considerably more flexibility in placement and upgradability. Consider, for example the case where extra sensors are needed in an area not previously envisaged. In the worst case power lines need installing, in a better case direct write[4], with new lines tapped off generic power sources distributed over the aircraft. In the ideal case a sensor or a new actuator can be placed where it is needed and work straight off. The other advantage of ZigBee in comparison to its rivals is its re-configurable nature and simple expandability of its networks. Unlike Bluetooth for example it is not designed for 8 devices but could handle a network many times larger, of up to 65536 nodes.

ZigBee should be able to cope with the data requirements needed for some basic sensors on an aircraft with standard moving surface controls. Unfortunately, the reality is that most modern aircraft use buses that run in the tens of megabits. ZigBee is a long way removed from this. For modern navigation, weapons and other systems, high data bandwidths are needed and additional networks, either wired or wireless would be needed alongside ZigBee. Further-

[4]Direct Write (DW) power might be added is a method of writing conductive tracks and other structures directly onto a surface using for example conductive inks cured by laser.

more it was also shown that ZigBee would not be sufficient in terms of data rates for the more novel control techniques being researched as part of the Flaviir project by BAE and universities such as Manchester and that wireless LAN standards are better suited to this. Despite this, ruling out ZigBee would rule out battery powered sensors and the option of removing all wiring, whether data or power is the best option. ZigBee cannot provide data heavy trunk links round an aircraft but will have sufficient data rates for health and monitoring applications where duty cycles are low. In the area of flight control it may still be possible to have a number of ZigBee sensors in a small network that feed data into a more powerful processor which contains either a wired link or a IEEE 802.11n link to a computer. One can imagine an architecture whereby ZigBee removes the vast majority of sensor and actuator links feeding information into a small number of other high data rate links. This technology therefore has great potential for reducing aircraft wiring significantly, even if some wires must remain and is more sensible than trying to replace every tiny sensor with an expensive 802.11 transceiver which is not realistic.

In light of this, the initial assessment is that ZigBee is the most useful technology in the move to wireless aircraft. Therefore it is a ZigBee based radio unit will be utilised for some real world performance measurements in later chapters. Some 802.11 measurements are also provided but the focus is on ZigBee and this is described in greater depth first.

## 2.6 Zigbee radios and the 802.15.4 specification

As ZigBee was been chosen as a standard to work with this section provides a more detailed look at the ZigBee standard, in particular at its physical layer which is of crucial importance in understanding the effect of multi-path on the system.

### 2.6.1 Physical layer compromises

Here ZigBee's physical layer, is presented in terms of its modulation technique and data rates and specifications.

By understanding the modulation scheme it is possible to see how it will interact with the channel and it may be possible to estimate bit error rates and maximum achievable data rates. Ultimately the combination of radio channel and modulation scheme will decide the data bandwidth available to the different sensors and actuators in the aircraft before any consideration of network topology and access scheme. Its important to note that ZigBee and IEEE

802.15.4 have a multiple access scheme and packet structure which put limits on performance not necessarily inherent in the channel. This will be true of all technologies. This section also presents a summary of general issues in modulation schemes and of Direct Sequence Spread Spectrum (DSSS).

Zigbee/IEEE 802.15.4 modulation, like wireless LAN can be broken down into two distinct areas:

- The modulation type used, for example amplitude shift keying (ASK) or phase shift keying (PSK)
- The spread spectrum technique used, such as DSSS, OFDM or FHSS

For modulation schemes there is always a trade-off between the rate at which information is conveyed to the receiver and the signal to noise level which may be tolerated - attempting to transmit too much information in too short a time can lead to a high error rate so there is a limit on the capacity of a channel such as Shannon's limit for the simple case of Gaussian noise. For each scheme there are bandwidth constraints and a tradeoff between the amount of bandwidth consumed and the data that can be transmitted. There is also a difference in the energy required to transmit a given quantity of information for different modulation schemes and there are significant issues in receiver complexity for different modulation types. Questions that must be asked include: Are low noise amplifiers needed, do the transmit amplifiers need to be linear or non linear, is a lot of signal processing required, should they have a high or low intermediate frequency and what is the size requirement on the receiver? Given all these tradeoffs different applications will find different modulation and coding schemes most appropriate.

### 2.6.2 Modulation

Offset Quadrature Phase Shift Keying (OQPSK) with half sine shaping is the modulation type used in IEEE 802.15.4. It corresponds to a type of continuous Frequency Shift Keying (FSK) called MSK. It is widely used for its constant envelope properties (see figure 2.6 and 2.7) and can be demodulated non-coherently which is useful for multi-path channels where accurate phase information may be difficult to obtain.

The bit rate at 2.4GHz is 250kb/s which is the bit rate in the radio, not throughput at the transport layer. When the bit stream enters the radio Direct Sequence Spread Spectrum (DSSS)

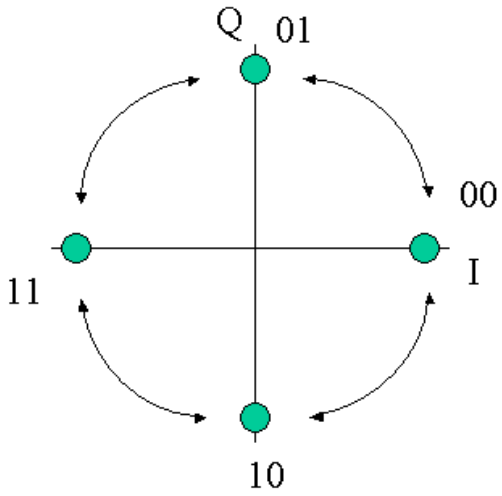


Figure 2.6: IQ diagram showing the constant amplitude transitions for MSK

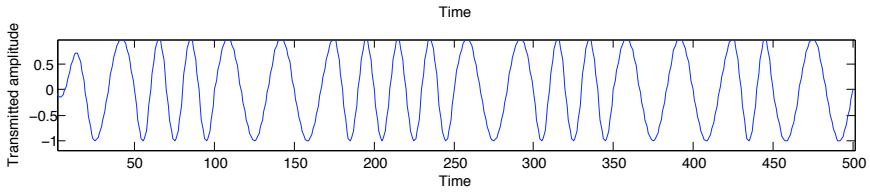


Figure 2.7: Constant envelope of MSK

is applied with spreading factor 8. The DSSS works by using 4 bits of the data stream to select one of 16 different 32bit sequences. At this point therefore there is a chip rate of 250kb/s * 8 chips, 2000kchips/s. This can then be separated so that chips 0,2,4,6 etc are used for the In Phase (I) channel and chips 1,3,5 etc on the Quadrature Phase (Q) channel (assuming an IQ modulator). These are then offset by half the duration of a chip on an I or Q channel and then shaped with a half sine pulse. At this point the data can then be modulated onto the carrier.

A 250kb/s bit rate is converted to a 62.5ksymbols/s transmitted rate. The symbol duration is 1/62.5 which is $1.6 * 10^5$. There are 32 chips per symbol or 8 chips per bit and the chip duration is 0.5$\mu$s. This information will be useful when looking at the frequency responses of various communications channels when the time it takes for a pulse to decay can be compared to the Zigbee symbol time. If the times are comparable then there may be a high level of inter-symbol interference for example.

DSSS is just one form of spread spectrum, another being Frequency Hopping Spread Spectrum (FHSS) where the frequency of transmission "hops" according to a pre decided pattern.

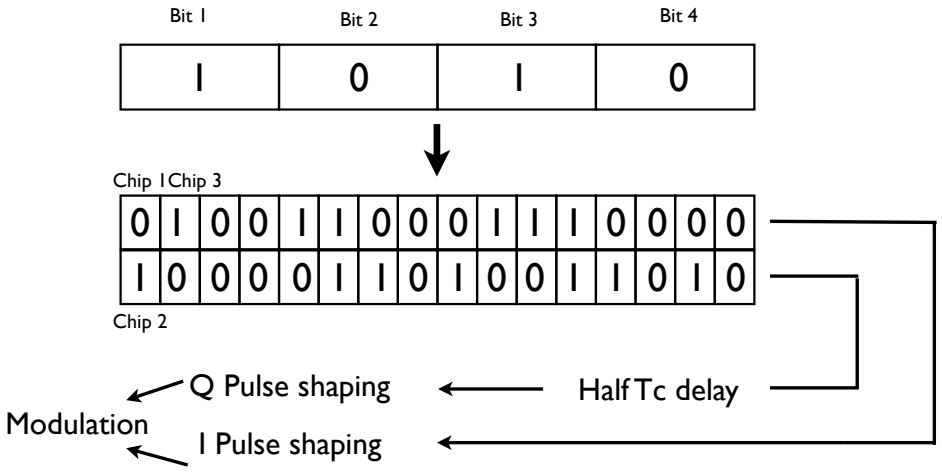


Figure 2.8: DSSS for ZigBee

Common to all systems is the spreading out of the signal power over a wider bandwidth meaning that narrowband noise will have less of an effect on the spread signal than the non spread one. In DSSS a single bit is replaced by a bit stream at a higher rate. Certain codes are used which define what a 1 or a 0 (or the case of ZigBee 4 bits) will get mapped to. Often it works by multiplying a bit stream of 1 or -1 by the code. The choice of code is dependent on the application. For example when used in a multiple access system each user will get a particular code and other user's codes should look like noise. In this instance low cross correlation between codes is important. Low autocorrelation is the priority in single user schemes where a delayed version of a signal should be rejected. The use of spread spectrum offers the following benefits:

1. It makes a signal harder to jam and harder to intercept if the power is spread over a wider bandwidth than a deliberate interferer.

2. It makes a system more resistant in multi-path environments where there might be frequency selective fading. Although the signal may be heavily attenuated over a part of its spectrum other parts will be at a higher level and overall signal power will be manageable in more situations than with a narrowband signal.

3. Delay spread effects are reduced by correlation with a sequence rather than a single bit

The first of these is not relevant for WLAN and ZigBee as the fractional bandwidth of these systems, even once spread is low. Note that in 802.15.4 spread spectrum is not used for multiple access as with a CDMA system.

### 2.6.3 Data link and multiple access

Assuming that the physical layer is working and a stream of bits can be sent over a channel and then decoded back into a stream of bits at the receiver the next issue is that of multiple access and the data link layer. It cannot be assumed that data successfully gets from a transmitter to a receiver. Even assuming that there is a perfect channel, most systems will have many users all trying to send data. Therefore a set of rules is needed so that packets get to the required destination, without error and that should multiple users wish to send data at the same time, this can be dealt with. ZigBee offers various approaches to these issues and works in two different modes:

- Beaconed
- Non- beaconed

It is easiest to consider the latter initially as it is the simpler of the modes. Each device wishing to send data broadcasts it when they are ready. However before transmitting they will check the channel to ensure that no one else is transmitting. It is known as CSMA/CD. This is not as reliable as in a cable system. It may be that two nodes are trying to transmit to the same node which is positioned in the middle of them but are themselves are too far apart to sense a signal from the other.

For beaconed mode a device must be assigned the coordinator role. Only a certain type of device known as an FFD may take this role. Such devices are not permitted to sleep and require power in order to carry out this role. Time is divided up into beacons and the beacons are divided up into slots, some of which are for CSMA/CD and some of which are distinct slots assigned by the coordinator to the devices. This assignment of slots is broadcast in a preamble at the start of a frame. In the measurements that follow in later chapters, the non beaconed configuration was used because the measurements were only testing with 2 radios and so there would be no need to schedule transmissions.

### 2.6.4 Sensor interfacing on ZigBee radios

In order to be useful a ZigBee radio must be connected to a data source and/or sink. In the later measurement chapters a ZigBee solution by Jennic[21] is utilised which is typical of many ZigBee products. Their product, like others has a radio integrated into a microcontroller. The most important implication of this is that a lot of interconnectivity is available including:

- SPI bus
- I2C bus
- DA converters
- General purpose IO

With such interconnectivity it is straightforward to connect to a sensor or actuator as indicated in 2.9. Once a radio is connected to a bus, additional sensors are trivial to add. Further, the processing power available allows basic processing of the data before transmission or perhaps the implementation of a standard such as MIL-STD-1553 so that a pair of wireless sensors could integrate with existing systems.

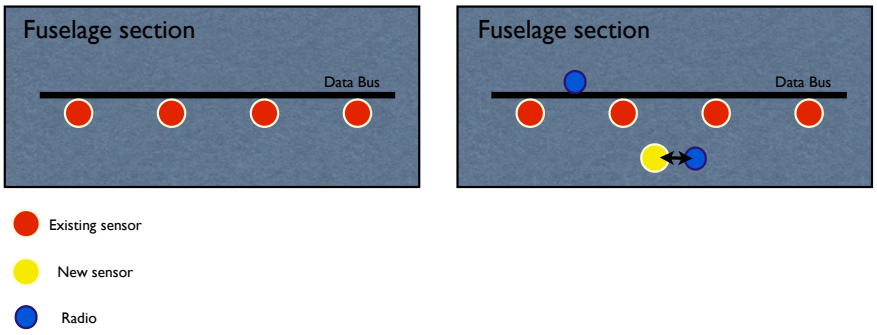


Figure 2.9: Addition of a new sensor into an existing wired system using ZigBee

In terms of size a ZigBee module from Jennic which includes an antennas and PCB has dimensions of 18mm x 30mm and the antenna can be a flat chip antenna. Therefore these are not units that take up significant space and they should not pose significant integration problems.

## 2.7 802.11 standards

Chapter 8 looks at the performance of 802.11 in a reverberant environment so the relevant details of the physical layer are described here.

### 2.7.1 Protocol layers

As expressed in the introduction the physical layer and the channel were researched rather than higher level protocols which are a large topic in themselves. Nonetheless it is necessary to consider the very basics of the protocols in order to know what the physical layer must do.

An 802.11 network can be of 2 types: independent or infrastructure. In the former two or more stations (a laptop and a phone for example) will communicate directly with one another. For the latter, which is by far the most common there is an access point and all stations that wish to communicate must do so via the access point in two hops. The access point provides a number of management functions, including the basic ability to join a network, coordination of encryption and bridging to wired networks like Ethernet and ATM based networks. Transmission is divided up into a concept of frames of which there are various types, examples being:

**Beacon** Periodically sent by an access point advertising the networks existence and capabilities

**Data** Includes the actual user application layer data being sent

**Clear to send** A response to ready to send

**Ready to send** Checking if it is ok to send a packet

The medium access control layer controls the access to the radio channel and can work in one of two ways: for the first all communication is done with the permission of an access point which implements an access scheme called the point coordination function (PCF). This asks stations if they have anything to send and tells them when they can send data. This particular scheme is not very common. The second and more common scheme is carrier sense multiple access with collision avoidance (CSMA/CA) whereby if no ongoing transmissions are sensed then a station will try to transmit its data frame. If there is a collision then it will wait and try again, with a randomised exponential back off utilised to deal with the busiest mediums. Where a threshold in the frame length is exceeded then RTS and CTS signals will be used to try and reserve the channel for a sufficient length of time. This is not the same guarantee one would find in a wired system because there could be interference in the same band or stations unaware of one another.

Other notable features of the MAC are the use of a cyclic redundancy check to ensure that the data has been correctly received and the ability to retransmit lost or damaged packets, up to a pre decided number of times. This is not found in wired systems where retransmissions are done at the transport layer and is needed because of the much greater risk of losing packets in a wireless network. Note that this is not a reliable protocol, it just attempts to give a frame/packet a better chance of getting to its destination than it would is 802.11 was literally a wireless Ethernet.

From the application perspective, what happens (in very simplified terms) in an 802.11 LAN is that the data the application wishes to send is placed into TCP or UDP packets that are then turned into ip packets and then sent to the 802.11 medium access control layer. This places local area network addresses on the packets, the CRC bits and other information before checking the medium and when ready sending the completed frame to the Physical Layer Convergence Procedure (PLCP). The PLCP constructs the final frame to be sent out by the Physical Medium Dependent (PMD) layer which means the actual radio hardware and software. Once received by the other end an ACK will be sent and the transmission is complete or the packet may be resent.

### 2.7.2 802.11b

Of the 3 standards described the first to be ratified and the closest to ZigBee is IEEE 802.11b. Each of the standards differ in the PLCP/PCM layers, or put in plainer english, in terms of their frame headers and radio standards. 802.11b's PLCP receives MAC frames and adds to them a 6 field header that includes a header CRC, synchronisation bits and a start of frame identifier. The synchronisation part, the preamble is 128 bits long for low data rates but can be 56 on higher data rates. It can be seen from this there are a number of ways a frame can be lost or received in error. Besides data bits being in error, if the synchronisation procedure fails then a receiver will receive no data and if the header is in error, then a frame can also be lost. Synchronisation is a particular concern in heavily multi-path environments where bit boundaries are blurred.

802.11b uses Direct Sequence Spread Spectrum DSSS and offers a number of data rates (1Mbps, 2Mbps, 5.5Mbps and 11Mbps). All these *bit* rates use an 11Mbps *chip* rate and fill a 22MHz band. Note that this means there are only 3 distinct non overlapping channels available.

Differential phase shift keying is used for modulation for the 1 and 2Mbps rates. Firstly the signal is mapped to a Barker sequence of length 11 (a 0 gives its inverse) and then the 11Mcps chips determine the carrier phase. When working at 1Mbps data is sent with one of two phases changes, 0 degrees, or 180 degrees where a 1 provides a phase change and a -1 no change. When using 2Mbps 4 phases are used so that two bits can be sent simultaneously and again this is differential (DQPSK) except where the additional phase can be 90, 180 or 270 degrees. The Barker sequence is used because of it's tolerance to multipath.

The fact it is differential is emphasised because in a reverberant environment the phase may be difficult to identify without a complex clock recovery system because of a number of paths arriving at different times and hence with different phases. This is even more the case in a time varying environment where synchronisation is done at the beginning of a symbol and then phase is tracked. For DPSK, provided the phase does not change over a single symbol then it should be straightforward to identify the following symbol. There is a price to pay for this compared to synchronous systems which is slightly low BER performance in the presence of Gaussian noise. For our purposes, where multi-path is a problem, one would expect good performance.

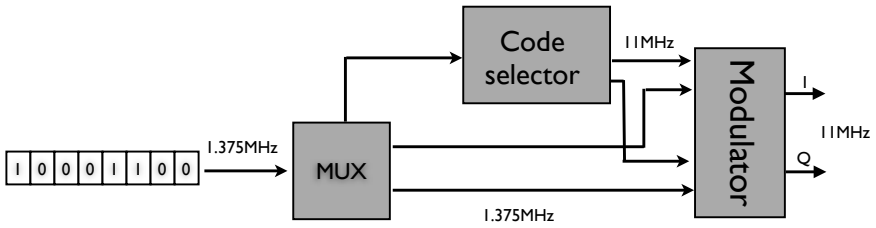


Figure 2.10: Modulation using CCK

The system for 11Mbits/s is more complex. It may be illustrated as shown in figure 2.10. Bits are treated in groups of 8, clocked at 1.375MHz and these groups of 8 map to 8 complex symbols (of length 8 chips). Bits are used as a group of 2 and a group of six. The first two of these generate a phase through differential phase shift keying encoding. The other 3 pairs produce phases through QPSK encoding (or it can viewed it as picking a complex code out of 64 in the diagram) However this is not where the encoding ends, rather these 4 phases are put

into the following formula for the 8 chips:

$$c = \left\{ e^{j(\phi_1+\phi_2+\phi_3+\phi_4)}, e^{j(\phi_1+\phi_3+\phi_4)}, e^{j(\phi_1+\phi_2+\phi_4)}, e^{j(\phi_1+\phi_4)}, e^{j(\phi_1+\phi_2+\phi_3)}, e^{j(\phi_1+\phi_3)}, e^{j(\phi_1+\phi_2)}, e^{j(\phi_1)} \right\} \tag{2.1}$$

The 5MBits/s system uses the same encoding formula. However, whilst phase $\phi_1$ is created from DPSK encoding of bits 0 and 1, $\phi_2$, $\phi_3$ and $\phi_4$ come from bits 3 and 4. In fact $\phi 3$ is always 0 for the 5.5Mbps scheme. At the modulation output for the 8 bits in there are 8 complex chips leaving at 11MChips/s leaving the modulator and an 11Mbps data rate.

The 802.11b modulation is complex in terms of the relationship between the data and the chips so to work out how performance will vary with data rate is not intuitively obvious. This system is only used for the data portion of the packet/frame. In fact for the header the synchronisation will always be 1Mpbs DPSK and the remainder either 1 or 2Mbps DPSK. With various rates over one packet it is impossible to say theoretically what the chance of losing data is, the only way to know is to simulate or test a real system. Nonetheless it would be expected to show resilience to delay in the tens of nanoseconds as one will find in office and home environments. The combination of the CCK code with something like a Rake receiver architecture allows it to deal with low delay spreads.

### 2.7.3 802.11g

The IEEE 802.11g standard works very differently to the b standard and utilises a technology known as OFDM which is a replacement rather than enhancement of the DSSS system. It has a slightly different PLCP header and very different modulation scheme. 802.11g uses a form of Frequency Division Multiplexing that maximises bandwidth by putting signals as close as possible together whilst remaining orthogonal. In the frequency domain there are responses that overlap but with the peak of one being at the 0 of its neighbours. This is characteristic of their orthogonality, and means that at the receiver one can recover each frequency independently by, for example, multiplying by each frequency and integrating over an integer number of periods (i.e. the definition of orthogonality). OFDM can perform very well in a multi-path environment because rather than sending one symbol at a time in quick succession it is sending many symbols in parallel but with a longer symbol time. It is more vulnerable to any frequency distortion such as that from receiver non linearities or from a time varying channel but less vulnerable to multi-path time distortion.

There are some disadvantage of using OFDM that arise because of its implementation.

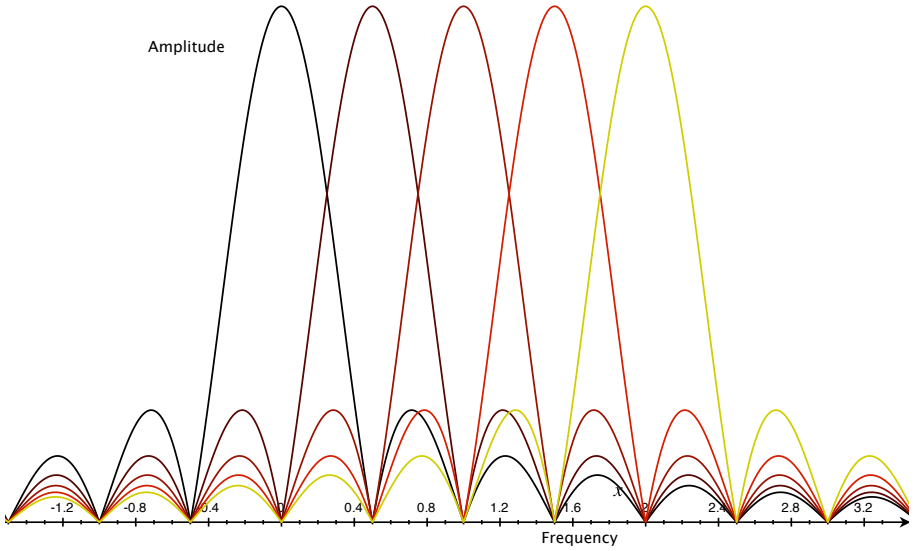


Figure 2.11: Overlapping carriers in OFDM

Firstly, there is a complicated signal in the time domain. One of the strengths of a system such as MSK is the constant signal envelope which allows non linear amplifiers to be used. In OFDM one must use a linear amplifier and these are less efficient and more costly than non linear amplifiers. This is one reason why earlier on it was suggested that a system like ZigBee is better suited to wireless sensor networks in terms of cost and power consumption. A linear amp will not allow batteries to last as long as non linear equivalents. In terms of the implementation of OFDM there are other concerns. OFDM is not implemented not by a large number of oscillators at the different frequencies all mixing separately, but by DSP techniques whereby the spectrum is constructed digitally and then an IFFT is applied to the data to produce the time domain signal. At the receiver an FFT may be applied and then the result can be processed. In 801.11g there are also a number of other coding steps performed upon the data which allow for error recovery and ensure that a rapid noise burst need not lose all data. Bits are scrambled and rearranged to allow original data packet recovery. The need for FFTs and coding means that DSP chips must be utilised which also increases the power and cost requirements of each radio.

OFDM has a feature for coping with multi-path known as the cyclic prefix. This is the time extension of an OFDM symbol placed on the front of the symbol. As a result of adding a prefix, if there is any delay spread, then it should affect the prefix, not the data containing part of the transmitted signal. Clearly this needs to keep the spectrum the same and not affect orthogo-

nality. Therefore each frequency, which is just a sinusoid, is extended back in time rather than having a period of no transmission or anything else that would change the spectrum. In fact this is often done after the IFFT by copying the end of the symbol and placing it on the front which is equivalent due to the linearity of the signal (it is just a sum of the constituent frequencies). The length of this prefix will be significant because if it is longer than the maximum delay spread, then one would not expect to get any errors, at least in a noise free environment. For a real environment a prefix of 4x the maximum delay spread is recommended [22]. In the design of 802.11, 200ns was expected to be the maximum delay spread and so the cyclic prefix (called the guard interval in 802.11) was made to be 800ns. One might at first expect then that an 802.11g system would work well when delay spread is less than 200ns and start to become unreliable as it approaches 800nS. The total symbol time for 802.11g is 4$\mu$S. With such a long symbol time and the prefix one would expect an OFDM system to have no errors when in a low Q environment such as the Tornado of chapter 7. On the other hand, delays of more than 1$\mu$s can occur in a reverberation chamber and even a long cyclic prefix and symbol time as found in 802.11g would not be enough to guarantee reliable performance in an unloaded or lightly loaded chamber.

As with the b standard, the bit rate is variable and can drop down when bit error ratios are high. Bit rates for 802.11g are as follows: 6, 9, 12, 18, 24, 36, 48, 54Mbps. There are 52 sub carriers. Of these 4 are pilot carriers which monitor ICI and path loss. The carrier spacing is 0.3125MHz. The pilot carriers are -21, -7, 7, 21. The modulation techniques used are BPSK, QPSK, 16-QAM and 64-QAM. Two coding rates are used of 1/2 and 3/4. The symbol rate per carrier is 250000 symbols per second. Which provides a period of 4 $\mu$s per symbol. To see how the particular data rate permutations are formed, consider BPSK. There are 48 carriers conveying a bit per symbol at a rate of 250000 bits per second. This gives a rate of 12Mbps. However a convolutional coding may be applied to boost error correction. The rate of this code may be either 1/2 or 3/4. This is where the 6 or 9 Mbps rates come from. In fact the 12 Mbps rate used is the 24 over the air rate with a rate 1/2 convolutional code.

The coding can help in a frequency selective channel in addition to the OFDM. A scrambler is applied to the bits before a convolutional code is used to provide forward error correction and then an interleaver is used in the mapping of bits to OFDM symbols. At the receiver the signal is typically decoded using a maximum likelihood Viterbi decoder. The way this is implemented means that if one carrier is not working, perhaps due to significant fading or delay

spread then its adjacent bits will be sent on carriers with a large separation. Bits will be lost but they will not be adjacent bits in the original stream because the likelihood is that adjacent carriers and adjacent bits in time are most likely to be lost. The coding will be able to tolerate a particular number of consecutive errors within a set of bits, dependent on the coding rate and by ensuring errors are evenly distributed the coding system is given the best chance of working. Where there are deep fades or some frequencies with particularly high Qs this should allow data to be recovered.

Finally, the transmit sequence for 802.11g is as follows: a clear channel assessment is used to decide whether to transmit or not. Following this the PLCP preamble is sent which has a short, then a long training sequence. The former of these is more for automatic gain control then the second for timing recovery. The longer sequence has a guard period because clearly it would be useless to add a guard to the data if the system was never able to synchronise in the first place. Then the signal field is sent followed by the data which is coded. Following transmission an acknowledgement is sent which is sent at a maximum of 24Mbps.

### 2.7.4 802.11n

In order to further boost data rates beyond the 54Mbps found in 802.11g the two most obvious options might be to increase bandwidth or to increase constellation size. The first option is not possible at 2.4GHz because bandwidth is dictated by regulatory bodies and the second isn't because increasing constellation size might make a system excessively susceptible to noise. Therefore the approach taken is to make use of multiple transmission paths and systems which make use of this are called MIMO systems because they use multiple transmitting and/or receiving antennas. What is a problem in some of the following chapters is being used to increase the data rate.

In fact there are two ways to utilise multiple antennas. The simplest way is to counter fading. Where there is constructive and destructive interference some spatial points will have a strong signal and others will have a weak signal. Also, some points may exhibit greater delay spread than others. With multiple antennas, at a given time one antenna will be found to provide a better signal and PER than the other and so this can be selected, improving performance in a fading channel.

One may go further than this and actually make use of the cause of the fading to increase the maximum data rate. Where there is fading there are typically different independent paths

from transmitter to receiver as shown in figure 2.12 Rather than sending identical data streams

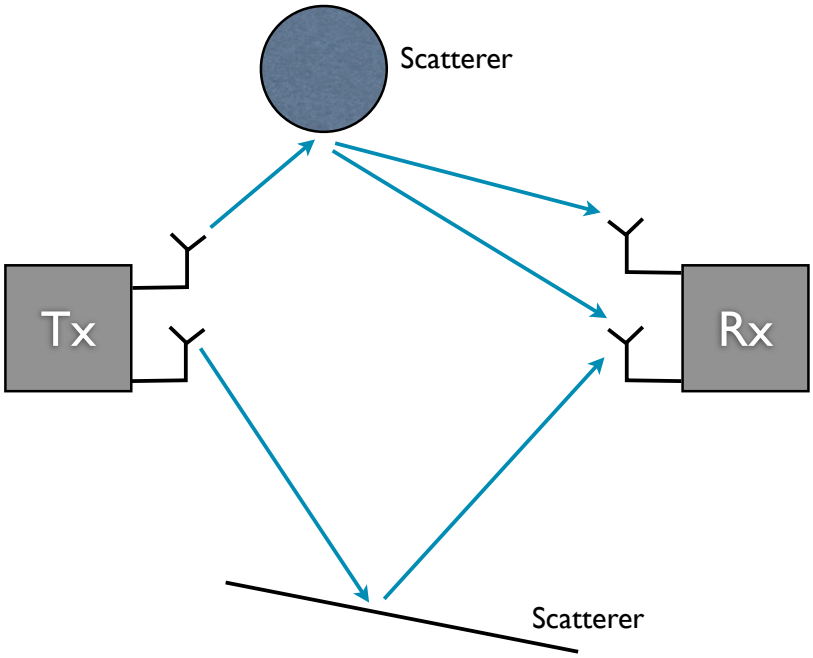


Figure 2.12: Multiple transmission paths due to scattering

from each antenna simultaneously one can send different data streams. Where there is only one path these will obviously interfere. However, where there are multiple paths then each receiving antenna will receive the signal from both transmitting antennas but in different ratios. The easiest case would be when transmitters and receivers work in distinct pairs but it is not a problem if both receiving antennas receive a signal from both transmitting antennas so long as the paths differ. One can then mathematically recover the individual data streams the data rate has doubled.

To make this clearer we may view a MIMO system in a more mathematical sense. When considering a single input single output (SISO) system then we had in the most general case a channel $h(\tau;t)$ where the channel has a time response at a given point in time but for which this response is also changing with time. Typically we imagine that over a symbol the channel is static and so the received signal is written as:

$$r(t)=\int_{-\infty}^{\infty} h(\tau)x(t-\tau)d\tau \tag{2.2}$$

which is the convolution of channel and signal. It is simple to extend this idea to a MIMO system by treating the MIMO system as multiple single channel responses between all the per-

mutations of transmit and receive antennas. All one needs to do is to add a couple of indices to the convolution equation:

$$r_i(t) = \int_{-\infty}^{\infty} h_{ij}(\tau)\, s_j(t-\tau)\, d\tau \tag{2.3}$$

and then one can build up a matrix of channel responses for each combination of antennas so that:

$$\mathbf{r}(t) = \mathbf{H}(t) * \mathbf{s}(t) \tag{2.4}$$

$$\mathbf{H}(t) = \begin{bmatrix} h_{11}(t) & h_{12}(t) & \cdots & h_{1N_T}(t) \\ h_{21}(t) & h_{22}(t) & \cdots & h_{2N_T}(t) \\ \vdots & \vdots & & \vdots \\ h_{N_R1}(t) & h_{N_R2}(t) & h_{N_R3}(t) & h_{N_RN_T}(t) \end{bmatrix} \tag{2.5}$$

This may further be simplified by assuming that rather than having lots of filters or frequency selective channels, instead we have a constant attenuation for each path. By doing this we shall be left with a matrix of constants rather than functions and it will be simple to apply the tools of linear algebra to the problem of solving:

$$\mathbf{y} = \mathbf{Hs} + \eta \tag{2.6}$$

to get the signal given the channel and the additive noise. The flat channel assumption is not unreasonable in OFDM whereby we have already divided the frequency selective channel into many flat channels and if it does not hold we shall other problems, namely the issue of not being able to decode the OFDM symbols.

802.11n is designed primarily as a way of improving data rates rather than improving robustness in a multi-path environment. It does improve robustness in the situation whereby some antenna positions are in deep fades and others are not but it does not increase the symbol time or cyclic prefix, with the modulation being almost the same as 802.11g. In fact 802.11g can use multiple receive antennas to keep the active data rate as close to the 54Mbps maximum by boosting signal to noise.

There are many data rates available for an 802.11 system and they are dependent on 5 parameters:

**Coding rate** 1/2, 2/3,3/4, 5/6

**Cyclic prefix** 400ns or 800ns

**Channel width** 20MHz or 40MHz

**Modulation type** BPSK, QPSK, 16QAM, 64QAM

**Number of spatial streams** 1, 2, 3, 4

The lowest data rate is 6.5Mbps with 1 spatial stream, an 800ns guard interval, BPSK, a 20MHz channel and a coding rate of 0.5. At the opposite extreme is 600Mbps which is 4 spatial streams, a 400ns guard interval, 64 QAM, 40MHz channel and a coding rate of 5/6. If the number of spatial streams is dropped to 2 then the maximum data rate is 300Mbps.

It is seen here that the 802.11n cyclic prefix / guard interval can be halved in order to boost the data rate and this would be expected to reduce the tolerance to reverberant environments. Other differences in 802.11n include shorter headers and an optional low density parity check codes as an alternative to convolutional codes.

## 2.8 Powerline communications

For the sake of completeness, and because it is something that will be in many readers minds, the next small section describes powerline communications. This is as an alternative to the radio systems that can offer high data rates and remove half the wiring. In fact where a radio or sensor will need power, it removes an equal amount of wiring weight.

### 2.8.1 Introduction

Over the last 5 years there has been significant growth in home networking which has been driven by a number of factors:

- Demand for multiplayer gaming
- Sharing of broadband internet
- Music, film and other file sharing

Traditionally home networking would involve a number of wires and some form of switch but this is problematic as all but the most recent homes lack the standard category 5 cabling required for a local area network. As an alternative wireless networks have become widely used and the standards have already been described. However, another technology has been developed making use of a house's wiring system. It is standard for every room in a home to have

a number of mains sockets meaning a complete infrastructure is already present throughout a home. Whilst mains cables are far from being an ideal communications channel there has been success in this area already and it is now seen as a mature technology[23].

Given that one of the main aims of using wireless networks in an aircraft is to reduce weight then removing communications cabling and only having power lines would appear to be a good way to make an aircraft lighter. Although removing power lines as well would be the best scenario it may turn out that battery power is not practical for a lot of systems, especially where actuator, radar and other systems are involved or where high duty cycles are required. In such circumstances powerline communications is as attractive an option.

In the following subsection the Homeplug system is described which is a commercial powerline standard and provides a good introduction to the technology. For more general information on powerline communications one good place to look is volume 16 of the International Journal of Communications Systems published in 2003 which features a brief introduction [23] and then 7 papers covering various aspects of powerline communications technology including an introduction to the Homeplug standard [24] which is summarised in the next subsection. Other papers here consider areas such as EMC considerations and alternative multiple access schemes including a system similar to TDMA where a central base station schedules channel access [25]. The use of multi carrier systems (as used in Homeplug) is considered and its performance is reported.

### 2.8.2 Homeplug

Homeplug is an alliance of companies who have created various power line standards including the HomePlug 1.0 standard and is sponsored by companies such as Linksys, Intel, LG, Motorola and many others. The alliance includes companies who deal in shaping the actual technology, those who build hardware products, software products and those delivering service and content. There are two standards, one is the Homeplug 1.0 standard and another the newer Homeplug AV Standard designed for higher data rates which is backwards compatible with 1.0. This report will only consider the older 1.0 standard because the later AV standard was not available until part of the way into this research.

**Physical layer**

When transmitting data over a mains system there are a number of issues that make it far more difficult than transmission over dedicated cabling and these arise from the high bandwidth needed for high data rate communications (not a lot of data can be transferred with a 50Hz centre frequency). Mains cable is designed to work at 50Hz and can be pushed up to 400Hz without problems. Beyond this it performs badly as a communications channel. Some of the potential problems are given in [26]: the terminating impedance will vary with both frequency and with time due to a varying load pattern. In a home environment attenuations will be of the order 20 - 60dB [24]. There will be a variety of conductor types and differing types of connection between them. Noise will come from lights, power supplies and motors as well as the interference from induced radio signals. Shielding can't be expected.

In order to deal with these difficulties an OFDM technique is used as with wireless LAN and reported in [24] as having the following advantages for this application:

- Copes well with time dispersion
- Minimizes inband narrowband interference
- Scalable to high data rates
- Flexible and adaptive with different modulation techniques available for different carriers, so for example, if one sub band was more noisy than others it could transmit bit less bits per symbol.
- Complex equalisation is not required

Various types of modulation are used: DBPSK and two variations of DQPSK and there is channel estimation in order to avoid the worst sub carriers and select the best modulation type. Physical layer throughput can vary from 15Mb/s to 1Mb/s dependent on the modulation employed. The band used goes from 4.49 to 20.7MHz. As the band is divided up 84 carriers should be available but 8 are masked to avoid interfering with amateur radio. The modulation and processing is quite complex and before the analogue front end involves the following:

1. Data scrambling - often used to remove any DC component and to ensure there are sufficient transitions for synchronization.
2. Reed Solomon-encoding - a code that is effective at correcting errors appearing in bursts.

3. Convolution encoding - involves a code sequence instead of a block and provides a high gain, i.e. good BER for a given signal to noise ratio

4. Puncturing - a method of increasing the rate without changing the encoder structure by removing bits, for example certain parity bits[27]

5. Interleaving - prevents burst of errors from affecting a continuous stream of bits

The radio front end has the following:

1. Constellation mapping block - turn the signal into points on constellation diagram

2. IFFT block - transform to time domain

3. Preamble block - automatic gain control and synchonization information

4. CP block (cyclic prefix) - prevents intersymbol interference at the expense of using power to transmit data with no new information

5. Raised cosine (RC) block - filtering used to prevent intersymbol interference

Finally, note that the medium access control layer makes use of CDMA/CA whereby inference of a collision leads to a back off. This is similar to Ethernet, but as in wireless LAN, performance is reduced because a lost packet caused by channel attenuation is incorrectly interpreted as a collision. A type of encryption described as "DES in cipher block chaining mode" [24] is used in order to offer privacy and security.

It is clear from this is that the Homeplug system used is considerably more complex than that found in ZigBee having a number of similarities with wireless LAN and ADSL and like them requiring significant power and processing power. The additional power constraints will not be a problem here by definition of the technology but cost may be. The transceiver complexity will increase costs per unit which may start to cause a problem when moving to hundreds of transceivers on an aircraft. Its worth remembering that one of the initial objectives of this work is to reduce costs.

**Performance**

In [24] a number of simulations of the homeplug standard were carried out and then "real world" tests were done in a Florida home. In simulation TCP throughput was found to exist

around 5.9-6.1 Mbps and for UDP it was at 7.5 to 8 Mbps. Performance in the residential environment ranged from 1.6 - 5.2 Mbps over ranges 2 feet to 70. Such throughputs are comparable with wireless LAN (at least as good as the earlier b variant) and considerably higher than those for Zigbee where throughput is in the 10s of Mbs. Differences were found between forward and reverse channels but connectivity was 100% throughout the home. Therefore it can be seen to be a robust and high bandwidth technology. This particular paper looked exclusively as the 1.0 standard and it should be noted that the technology has improved since this point in 2003.

The latest system, HomeplugAV is described in [28]. This system operates in the range 2 - 28MHz and has a number of changes compared to the 1.0 standard including the use of a Turbo code to get closer to the channel capacity. There are also many more carriers in the AV standard at over 900 and encryption is 128 AES as found in 802.15.4. For this system Homeplug claims a rate over the physical channel of 200Mbps and an information rate of 150Mbp which is comparable to the 802.11n wireless LAN standard.

**Opportunities on an aircraft**

The Homeplug white papers refers to the technical difficulties which are involved in communicating over power lines. These are:

- Different conductor types
- Random terminating impedances
- Interference from lights, power supplies and many other devices

If an existing aircraft was to be used then with communications being moved to the powerlines for some devices then all these problems would continue to apply. However if instead a UAV was being developed from scratch then the power lines could be developed and standardised in the knowledge that they would be used for communications making such a system much more effective with very high data rates.

### 2.8.3 Summary of powerline communications

From this brief look over powerline communications it can be seen that such systems could be very useful as a way of reducing cabling on aircraft. In those scenarios where power is needed anyway, and where high data rates are required it sounds ideal. However this thesis does not

look any further at powerline and no further research was carried out into it. The reasons for this are as follows:

- Homeplug is an expensive technology that would be hard to use in volume, even more so than high performance wireless
- It is potentially even harder to certify than wireless because performance very much dependent on the loads on the powerline and the many types of cable used (though this is less of an issue for future designs)
- Access to real aircraft is required to carry out measurements and was not available

Powerline technology could definitely have a role to play on aircraft but given that wireless can offer simpler solutions that are no harder to certify and with no connections required (potentially) these are preferred.

## 2.9 Previous research in fly-by-wireless

Fly-by-wireless aircraft is an extremely broad area that can cover many different types of avionics and communications technologies. It is also a very recent subject with little previous work available. Therefore it is not possible to provide a major literature review for this area but what research does exist, specifically in the area of fly-by-wireless aircraft is presented in this section. The subject of fly by wireless aircraft is a relatively new one and in fact wireless sensor networks themselves are a new technology with the popular ZigBee standard [16] only being ratified in December 2004 and the recent standard found in most modules, ZigBee 2006 being proposed in December of that year. Before ZigBee, other custom systems were used such as Mica Motes [29] and Bluetooth was used for sensor applications

#### The first fly by wireless demonstrator

Coelho and Macelo looked at "A fly-by-wireless UAV platform based on a flexible and distributed architecture"[30] and present a very broad overview of a demonstrator platform discussing everything from flight control to ground station issues. In the area of avionics they describe a distributed architecture where processing is distributed on small nodes found near individual sensors. This differs from the trend towards using a modular avionics architecture

with complex central processing as might be found in an F22 (see section 2.2). Each of the distributed nodes communicated by Bluetooth. The demonstrator aircraft was a very small and simple one with total wingspan of 4.8m and a 2.9m long fuselage. The paper states that the system has "one processing unit for the Bluetooth (BT) piconet Master; one flight controller unit; one data logger and earth link; and one embedded vision system (EVS)." The wireless system has the following slaves in the Bluetooth piconet:

1. GPS
2. Engine, flap, aileron
3. Engine, flap, aileron
4. Rudder and elevator
5. IMU
6. Altimeter
7. Pressure probe

This is a very basic setup for the simplest of UAVs. The engines are actually electric motors, there are very few sensors compared to a real aircraft and data rates are very low. However this may not be too far removed from reality for the very cheapest and slowest of the UAVS that may be used for civilian monitoring applications. By using Bluetooth, network expandability is severely limited which is one reason why it won't be pursued in this project. But their paper does mention some of the advantages that are shared by ZigBee, the technology used in this Flaviir project- namely the weight losses and flexibility afforded by using a wireless system.

The distributed architecture is one of the more interesting aspects of this because in older aircraft designs (see section 2.2) there is a configuration whereby each system has its own processing units and the aim has been to go more central in recent years with replaceable racks. However, in the older designs the avionics would be customised for each system. The scenario described in this paper involves the same chips and radios being used in each system just as would happen in the avionics racks. Disadvantages may still exist however as smaller units may have less processing power than the rack based boards and don't benefit from the redundancy in power supplies. For a small UAV that only needs to do basic processing it might be sufficient.

The architecture in this paper also has the problem in that there is not a system of equal units communicating but rather a master slave configuration with a single point of failure. Even for a small UAV the such an architecture, although simple, would be undesirable.

**The Caneus Workshop**

A lot of relevant technologies were described at the 2007 CANEUS/NASA workshop on "Fly-by-wireless for Aerospace Applications". This workshop considered a number of different topics including:

1. Air vehicle applications
2. Space vehicle applications
3. Vehicle health monitoring and test applications
4. Architectures and networks
5. Wireless instrumentation
6. Passive sensor tags

From these items 1,3,4,5 are the most relevant to the Flaviir work. Content in the presentations and briefings at this workshop was generally at a very high/general level. It was the first conference on this subject and as such mainly consisted of different companies introducing what they do, why they are interested in wireless aircraft and what they think the advantages are rather than the details of complete solutions.

In a briefing by Honeywell it was suggested that wireless is best used as a backup for wired solutions rather than as the only method for data transfer [31]. This could help protect against common mode failures that could affect all wiring at certain points were there critical damage. Reasons cited for wireless only aircraft not taking hold in the near future were: conservatism, the need to still get power to sensors and actuators and EM susceptibility. Furthermore it was suggested that wiring is not a big problem now with the much greater uses of data buses.

In [32] health monitoring systems are singled out as a potentially suitable area for the deployments of wireless aircraft. It again refers to EMC issues and the susceptibility to common mode failure. The complexity of current wiring is singled out as a motivation for moving to wireless.

Securaplane revealed that they are working in the area of wireless sensors for aircraft [33] and plan to fly a demonstrator. They also show how wireless is already being incorporated into aircraft with the example of the wireless emergency lighting system in the Boeing 787.

Bombardier are also looking into this area and referred more generally to reduction of wiring [34], as opposed to just fly-by-wireless. One suggestion made here was to use optical systems both for data and power. They state that systems such as avionics, inertial reference though to brakes and fuel monitoring can require many miles of wiring even in small regional jets. Making systems more integrated is given as an aim and some architectures are shown. Despite the move to more buses they point out that there is still a large amount of wiring between the remote interface units and the buses, that could potentially be removed by wireless systems like RFID. Cabin systems are brought up as an area ideal for trials of wireless systems.

In [35] NASA looked at sensor networks and developed a custom system that sensed temperature, pressure, strain and vibrations. This system was tested in 2004 in flight as well as on the ground and it was found that performance remained the same for their system and they achieved a 20.6kbps data rate with 1.15 retransmissions per packet.

Invocon gave briefings on a number of topics including wireless control [36] and on how to improve reliability for wireless sensor networks in flight applications [37]. In the former they provide some useful information as to what levels of improvement could be gained by removing wiring in aircraft. They state that for a general aviation aircraft, the Cessna 310R there could be a 90 lbs reduction leading to a 10% range increase. For the NASA APEX, 130 lbs out of a total weight of 750 lbs could be saved and for a SH60 military helicopter it would be 627lbs. Some examples are given of cases where wireless has already been used such as a four node system on the shuttle with 3 remote sensor units monitoring temperature and reporting back to a controller unit. Other temperature sensors are also mentioned and microgravity monitoring and vibrations. Data rates on the systems used range from tens of kbits to a 2Mbit wireless LAN. All the systems described made use of spread spectrum to reduce the multi-path effects.

In the area of improving redundancy Invocon refer to dependability which has three components: reliability, availability and redundancy. Issues affecting dependability which are unique to wireless systems are:

- A transmitter that gets stuck on can bring down a system since it is not just using one wire it is occupying a channel

- When a node does not communicate it is impossible to know the cause so health information must be continuously gathered
- Battery failure
- Non-deterministic behaviour for a fully connected system (wired non fully connected systems can be more easily analysed)
- Synchronisation is more difficult with wireless

They emphasise the extra redundancy that a wireless network can add and suggest that the most efficient point to add this is at the node level; it is very expensive to do this within a node. Being able to remove cables to sensors will easy requirements such as temperature and vibration. It is also pointed out that data rates are limited compared to wired systems, data reduction needed

**Gulfstream fly-by-wireless**

In Flight International in 2008[38], there was a report of Gulfstream testing a fly-by-wireless system on its G550 aircraft. This was being used purely as a backup system and the article suggests that no attempt at certification had been made. As was suggested in this project they report that they are using a "direct sequence spread spectrum modulation and coding technology". The technology for this system was created jointly with Invocon, spinning out of the NASA related work described above.

**Cars and other vehicles**

Some of the advantages of making aircraft wireless also apply to cars and so there has been research into wireless sensor networks in road vehicles carried out as well. For example Ahmed, Saraydar and others [39] state that a mid range vehicle will have 50 sensors and contain up to 4km of wiring and weigh up to 40kg. This means that the savings to be gained per vehicle are less than a large aircraft but they also point out because of the number of cars sold this represents a multi billion dollar sensor market where there are large incentives to move to wireless in terms of installation and maintenance costs. They explain that wireless systems would avoid the reengineering associated with Controller Area Network (CAN) systems when new nodes are added. Wireless standards are compared such as Bluetooth and ZigBee and

similar conclusions are reached to those of this chapter although for slightly different reasons: ZigBee is the most suitable standard, Bluetooth is too complex and power expensive, 802.11 is a poor match for the protocol light automotive applications and has too large an overhead through it's protocol stack up to internet protocol (IP).

This paper also offers some useful design data by stating that "the most demanding sensor will require a latency of approximately less than 1 ms with data throughput of 12kbps". When investigating a star topology of a customised 802.15.4 network connected to a CAN network he author's conclude that a latency no better than 15ms can be achieved and so this does not match the most demanding applications. However a lot of this analysis is based around the use of super frames rather than CDMA/CA and so this is only true when a central coordinator is required.

Carmo and Couto [40] also look at wireless sensor networks for automotive applications. As in [39] they describe the need to add an increasing number of sensors to vehicles as being hampered by the practicalities of installing wired CAN connected devices. They also assess ZigBee and Bluetooth as being inappropriate because of Bluetooth's latency and ZigBee's low duty cycles and high latency which is said to be appropriate for the real time requirements of automotive sensors. However they then contradict this by stating that their CMOS devices could monitor variables such as water temperature and oil level and so it would be more reasonable to say that certain real time systems are not easily implemented with ZigBee or Bluetooth but many are. The remainder of the paper concentrate on RF design but it is interesting to note that the final design uses 11.2mW when transmitting and is extremely small with dimensions 1.5mmx1.5mm.

## 2.10 Conclusions

This chapter has examined the current state of avionics from the older more conservative technologies through to the state of the art. The latest in COTs wireless networking systems have been described and how we might apply such technology to replace wired avionics links has been researched. It has been found that out of all the technologies available ZigBee is the most interesting and useful one for aircraft sensor networks and so later chapters will focus on this. ZigBee is not only designed for wireless sensor networks but is one of the only standards based technologies that promises battery power for many nodes. When transitioning to a wireless aircraft ideally the designer doesn't want to be left with power lines if it can be avoided; batter-

ies or energy scavenging is preferable.

It is clear that on paper, all of the current COTS wireless networking standards could replace particular aircraft systems at both low and high date rates. ZigBee and 802.11 appear to be the most suitable technologies because they are built with large networks in mind rather than a small number of point to point links as is found in systems like Bluetooth and they are scalable. There are some complications, particularly for ZigBee links as they stand but also applying to wireless LAN. Earlier it was found that ZigBee could equal the data rate of certain types of existing aircraft bus but that there are often tens of such buses over an aircraft or more, a situation that would get even worse if the system designer wished to make each link redundant. ZigBee has multiple radio channels but it does not have sufficient to match this.

Although this looks like a serious problem it is not a reason not to use COTS technology. Currently all wireless standards work typically at 2.4GHz and some at 5GHz. Wireless avionics could be standardised around ZigBee and 802.11 in everything except the radio band and they might be able to operate in more, separate and wider bands bands, particularly if power levels are low and there was confidence of not interfering with anything outside the aircraft. Existing systems would only need the RF front ends and antennas adapting to the new frequency bands and one could keep the remaining architecture including the DSSS/OFDM systems, MAC layer protocols and above which is where a lot of the complexity is. All one would need to change is the centre frequency and number of channels which is far more affordable than developing a new standard from scratch, even for a large number of UAVs. As explained earlier, going with a COTS solution is the best option if the COTS system can do everything you require. It can even help in certification because one can have much greater confidence in a network stack that is being continually used and tested in many applications in tens of thousands of units as opposed to a stack that exists in a few hundred aircraft.

When starting out there was a hope that a single technology might be suitable for all avionics links. This is not possible when there are the conflicting specifications of low price, battery power and high data rates. Instead it is far more likely that an avionics architecture would involve small ZigBee type radios next to individual sensors communicating to central processing units that might in turn relate data to the cockpit or mission computer via a high data rate system such a 802.11. Figure 2.13 illustrates this. Having said that, the wireless UAV described earlier utilises only Bluetooth and in [41] there is a wired UAV solution makes use of a distributed architecture like the wireless Bluetooth one. The point is that by moving to dis-

tributed architectures one can reduce the load on a each link because we are not transferring vast amounts of data between central processing units. Nevertheless there is no escaping the heavy data rates needed for video transmissions, as an example. Ideally there would be a set of technologies that might be applied consistently to both small UAVs and existing manned aircraft both military and commercial that keeps costs low and avoided the need for power lines where possible. As a result, the combination of a low power low data rate and high power high data rate wireless systems appears ideal.

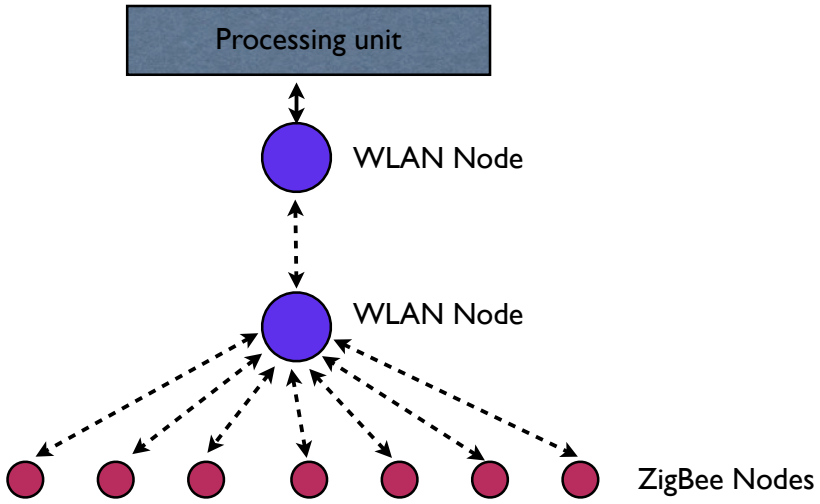


Figure 2.13: An architecture mixing ZigBee and 802.11

Going beyond the simple advantages of moving to wireless avionics that were outlined in the introduction there are many questions as to how much of an aircraft to make wireless, which particular systems to replace and which not to replace. Included in this will be the issue of what data rates are required by different systems and what level of reliability. Clearly it is more straightforward to replace a low data rate non safety critical component linked with a wire by an alternative linked wirelessly, as opposed to a high data rate safety critical function. This is because in general wired standards are not just faster and more reliable but they are also understood and certified. Wireless standards do yet appear close to certification for safety critical functions because research into wireless avionics has only begun to appear in the literature in the last 5 years. When this research phase of the Flaviir Wireless project began there was almost no published literature about wireless aircraft. Since then there have been some developments such as those shown at the Caneus workshop but one can't expect a whole air-

craft to become wireless over a period of 2 or 3 years or even ten years because the topic is so new and the safety requirements so high.

Some of the first applications of wireless on aircraft have been in delivering entertainment on commercial airliners and in the backup of wireless cabin lighting control. In both these applications reliability, robustness and timeliness are unimportant in comparison to the systems on which a passenger's life depends. This is the most likely way for manned wireless aircraft to develop, with a very slow phasing in of wireless backup and non critical functions on new aircraft. On the other hand in the UAV business it might be expected that there would be a much stronger drive for wireless because UAVs are intended to be higher volume and more disposable than very expensive manned jets. Their requirement of being highly adaptable may also drive a move to wireless. It could be expected that UAVs would move to safety critical wireless systems earlier than their manned counterparts. Nonetheless, because moving to wireless has benefits in all types of air vehicle, the discussion is kept open to all aircraft types in later chapters.

Now that the wireless systems that are best suited, on paper, to be adapted to an avionics application have been found the next few chapters reveal how well these systems do actually to perform at the physical layer and if they really can meet the requirements of modern avionics.

# Chapter 3

# Electromagnetics in highly reverberant vehicles

## 3.1 Introduction

At the heart of the development of a wireless aircraft is wireless propagation. An avionics system must be well understood and must be reliable in that sense that no data link should lose availability for sufficiently long to risk loss of control of the aircraft. Loss of internet in an office is inconvenient, but loss of signal in a jet carrying a person or perhaps a weapons payload could be catastrophic. The only way to design a reliable system is through proper understanding of the propagation environment and how that interacts with the wireless systems. To provide a reliable flight control system one requires the high availability of its wireless links during flight and where the propagation environment does not allow this there may be system failure.

Propagation in an airframe is likely to be different to that in which wireless sensor technologies such as ZigBee typically work. Therefore before committing to a technology it is important to consider the propagation environment and whether a given system can function reliably in that environment. In this chapter the electromagnetic environment is analysed from a theoretical perspective. The very basics of resonant cavities are discussed before moving on to reverberation chambers and how these relate to the environment inside an aircraft.

### 3.1.1 Aim

In the preceding chapter the aim was to investigate the radio systems that might be used for a wireless aircraft. Most of the following chapters examine the performance of those systems in

an aircraft like environment and so the aim of this chapter is to provide sufficient background that the reader might understand what to expect in the following chapters and why. Put another way, the aim of this chapter is to connect the wireless systems of Chapter 2 to the results of Chapter 4 onwards.

## 3.2 Electromagnetic basics

Throughout this and other chapters there will be frequent references to propagation in resonant cavities. This is because aircraft can be viewed as resonant cavities of varying dimensions and Q-factors, coupled together. Before explaining and justifying why this is it will aid the later discussion to go through the basics of electromagnetic (EM) waves in cavities.

All treatments of electromagnetics start at Maxwell's equations. When considering fields in free space rather than materials there can be written as:

$$\nabla \cdot \mathbf{E} = 0 \tag{3.1}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{3.2}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{3.3}$$

$$\nabla \times \mathbf{B} = \mu_0 \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} \tag{3.4}$$

and they say that there is no divergence, or outwards movement of field from a point, when there is no charge, and that time varying fields of type magnetic or electric generate fields of the opposite type. **E** is electric field, **B** the magnetic, $\mu_0$ the permeability of free space and $\epsilon_0$ the permittivity of free space.

Of these equations the two curl ones are the most interesting because some simple algebra demonstrates that a changing electric field leads to electromagnetic waves. Take the curl of both sides of either of the curl equations and then rearrange:

$$\nabla \times \nabla \times \mathbf{E} = \nabla \times -\frac{\partial \mathbf{B}}{\partial t} \tag{3.5}$$

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = \frac{\partial}{\partial t}(\nabla \times \mathbf{B}) \tag{3.6}$$

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = -\mu_0 \varepsilon_0 \frac{\partial^2 \mathbf{E}}{\partial t^2} \tag{3.7}$$

$$\nabla^2 \mathbf{E} = \mu_0 \varepsilon_0 \frac{\partial^2 \mathbf{E}}{\partial t^2} \tag{3.8}$$

So here is one form of the wave equation. We could have done exactly the same for the magnetic field or we could have started with the scalar potential $\Phi$ and the vector potential $\mathbf{A}$. There are many ways to manipulate Maxwells equations in order to solve a given problem as easily as possible. At this point one can clearly see the 3 dimensional wave equation and it turns out that in a reflective closed environment, if there is a time varying excitation then there will be standing waves as in acoustics. It is instructive to follow the derivation further to see the form of these standing waves and also the difficulty in solving such equations.

The outline derivation for a closed box shape, rather than any other will be given because that is what we shall be considering in later chapters. An electromagnetic view of an aircraft is examined later but very roughly one can compare avionics bays and even the fuselage to rectangular metal boxes. Although the solutions below are not general and particular to a cuboid, one would obtain very similar results for a closed tube like shape because there would still be a field made up of many standing waves or modes, it is just that their configuration would look different. However, as will be shown, the structure of the standing waves is not important just the statistical nature of the channel that they produce.

What follows assumes a perfectly conducting metal box for the analysis which is important for the application of the boundary conditions. One can begin by using separation of variables for the time and space variables whereby one can express a function of multiple variables as a product of functions of single variables. Clearly when you start with this method there is no guarantee of success but if it isn't valid then the boundary conditions will ensure no solution is found. One takes electric field $\mathbf{E}$ and splits it into two functions of time and position and then applies the wave equation. An equality results that can only work if each side is equal to

a constant, since each is a function of a different variable.

$$\mathbf{E}(\mathbf{r},t) = \mathbf{E}(\mathbf{r})E(t) \tag{3.9}$$

$$\frac{\nabla^2\mathbf{E}(\mathbf{r})}{\mathbf{E}(\mathbf{r})} = \frac{\mu_0\varepsilon_0}{\mathbf{E}(t)}\frac{\partial^2 E(t)}{\partial t^2} \tag{3.10}$$

$$\frac{\nabla^2\mathbf{E}(\mathbf{r})}{\mathbf{E}(\mathbf{r})} = K \tag{3.11}$$

$$\nabla^2\mathbf{E}(\mathbf{r}) + k^2\mathbf{E}(\mathbf{r}) = 0 \tag{3.12}$$

$$\frac{\partial^2 E(t)}{\partial t^2} + {k_1}^2 E(t) = 0 \tag{3.13}$$

(3.14)

In fact although we have separate functions for time and space we need not worry about the time side as we can use a Fourier description of the wave. In fact we could have assumed the following phasor form for the E field's time variation:

$$E = E_0 e^{j\omega t} \tag{3.15}$$

(3.16)

For the spatial part we may again try separation of variables, trying here just to obtain the $E_x$ component of the field which is still a function of 3 variables.

$$\left(\frac{\partial^2 E_x(x,y,z)}{\partial x^2} + \frac{\partial^2 E_x(x,y,z)}{\partial y^2} + \frac{\partial^2 E_x(x,y,z)}{\partial z^2}\right) + {k_r}^2 (E_x) = 0 \tag{3.17}$$

$$\frac{\partial^2 E_x(x)}{\partial x} + {k^2}_x E_x(x) = 0 \tag{3.18}$$

$$\frac{\partial^2 E_x(y)}{\partial y} + {k^2}_y E_x(y) = 0 \tag{3.19}$$

$$\frac{\partial^2 E_x(z)}{\partial z} + {k^2}_z E_x(z) = 0 \tag{3.20}$$

$$E_x(x) = A\cos(xk_x) + B\sin(xk_x) \tag{3.21}$$

$$E_x(y) = C\cos(yk_y) + D\sin yk_y) \tag{3.22}$$

$$E_x(z) = E\cos(zk_z) + F\sin(zk_z) \tag{3.23}$$

(3.24)

$$E_x(x)=A\cos(xk_x) \tag{3.25}$$

$$E_x(y)=D\sin(yk_y) \tag{3.26}$$

$$E_x(z)=F\sin(zk_z) \tag{3.27}$$

$$\tag{3.28}$$

The cosine and sine terms are obtained by inspection. By application of the boundary condition that says tangential fields at a boundary are equal, it was possible to make two of the cosine terms equal to zero. The normal component at the boundary is not necessarily 0, being dependent on the surface charge and permittivity. Thus a sine term is eliminated which would force such another 0 where it did not exist (the same is true for the two tangential B components the curl of which provides the normal E field.). Thus we see that we have a standing wave solution for the field as expected.

There are many other ways to arrive at this. The wave equation could be bypassed entirely applying Maxwells equations directly with the boundary conditions. We may have taken existing knowledge of waveguide or planewave solutions and found a standing wave from the reflections of those waves. Still we note that the solution whilst simple from an EM perspective is still not completely trivial despite this being the simplest geometry one could imagine. When there is a sphere for example then we must use spherical forms of the equations but in these the Vector Laplacian is no longer made up of 3 independent components as here. When we have an environment of partially curved, partially flat surfaces it is near impossible to stitch together an analytical solution and we must use numerical techniques such as FDTD. In fact it does not require many changes to rectangular box to make the field very complex and unpredictable, a property exploited in the reverberation chamber.

Following similar steps to the above one can arrive at the electric field equations found in Hill[42]

$$\begin{aligned}
E^{TM}_{zmnp} &= E_0\sin(\frac{m\pi x}{a})\sin(\frac{n\pi y}{b})\cos(\frac{p\pi z}{c}) \\
E^{TM}_{ymnp} &= \frac{k_y k_z E_0}{k^2_{mnp}-k^2_z}E_0\cos(\frac{m\pi x}{a})\sin(\frac{n\pi y}{b})\sin(\frac{p\pi z}{c}) \\
E^{TM}_{xmnp} &= \frac{k_x k_z E_0}{k^2_{mnp}-k^2_z}\cos(\frac{m\pi x}{a})\sin(\frac{n\pi y}{b})\sin(\frac{p\pi z}{c})
\end{aligned} \tag{3.29}$$

where a,b, and c are the x, y and z dimensions and m,n,p are integers that may be placed into the sines and cosines. There are multiple solutions to this because the m,n, and p are integers and can be nearly (not all zero for example) anything without changing the equation. The

wavenumber/eigenvalues can be shown to fit:

$$\begin{aligned} k_{mnp}^2 &= \left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2 + \left(\frac{p\pi}{c}\right)^2 \\ k_x &= \frac{m\pi}{a},\ k_y = \frac{n\pi}{b},\ k_z = \frac{p\pi}{c} \end{aligned} \tag{3.30}$$

Before progressing to this there are important properties of the rectangular cavity that must be considered. Firstly because this is a resonant environment (energy is trapped inside the box, excluding wall losses and small gaps etc) there will be resonant frequencies associated with the modes:

$$\omega_p = \frac{k_p}{\sqrt{\mu\epsilon}} \tag{3.31}$$

for the one dimensional case and for the 3 dimensional box we have just described:

$$f_{mnp} = \frac{1}{2\sqrt{\mu\epsilon}} \sqrt{\left(\frac{m}{a}\right)^2 + \left(\frac{n}{b}\right)^2 + \left(\frac{p}{c}\right)^2} \tag{3.32}$$

so it can be seen that the box has natural frequencies that can be excited allowing for significant energy storage and thus reverberation.

These resonant frequencies will each have a Q which is dependent on the losses in the cavity. It is explained in Hill[42] that a number of different methods suggest a Q of:

$$Q \cong \frac{3V}{2\delta S} \tag{3.33}$$

where S is inner surface area, V the volume and $\delta$ the skin depth. For the York reverberation chamber (4mx5mx3m) this would suggest a theoretical Q of just under 500000 but this is an order of magnitude higher than one would find in reality due to imperfect joints, apertures and other losses which can bring the Q down very rapidly.

It has been shown then that there are modes in a metal box and one could, if required calculate all of these modes individually from the equations given. However it is often more useful to know roughly the number of modes up to a certain frequency since this affects the statistical behaviour of a box. Knowledge of the mode density is also useful because higher mode densities provide a more rapidly varying channel over a given band. The number of modes up to wavenumber $k$ is defined as

$$Q \cong \frac{k^3 V}{3\pi^2} = \frac{8\pi f^3 V}{3c^3} \tag{3.34}$$

which is an asymptotic approximation and so loses validity when close to the first mode or equivalently the cut off point.

The mode density as a function of frequency is obtained by differentiation with respect to frequency giving:

$$D_w(f) \cong \frac{8\pi f^2 V}{c^3} \tag{3.35}$$

When there are a lot of modes within a band, each with a high Q then there is a very random field in the sense of being uncertain at a given position about signal strength, as will be shown shortly. If there are very few wide modes then there is a very flat frequency response and less randomness. From a communications systems point of view a flat channel is ideal. A very peaky response in the frequency domain implies a long time response and a distorted signal affected by inter-symbol interference. As the channel moves from flat to variable it is very difficult to predict its performance as the signal strength is unpredictable and even the inter-symbol interference can behave differently as an antenna picks up different modes from one orientation and position to the next. One can only talk about performance in a statistical sense and this is best described using a reverberation chamber which is what will be looked at next. Before moving on though it is worth recapping the key points raised in this rapid tour through EM in cavities:

- Maxwells equations dictate a modal structure to fields in closed boxes, tubes and all types of system
- Analytical solution of all but the simplest geometry is impossible
- With the modal structures comes resonant frequencies in the cavity
- As the Q of the peaks increases, so the channel variance increases and the field strength one might find at a given position becomes very uncertain

### 3.2.1 Reverberation Chambers

A reverberation chamber is a type of resonant cavity typically used for EMC measurements. Put simply it is a metal box with a large paddle inside it as shown in figure 3.1 The paddle can be rotated and has non symmetrical geometry ensuring that as it rotates the signal pattern will not repeat. Rotation of the paddle moves the modes by the changing boundary conditions or put another way, angles of reflective surfaces. Typically a reverberation chamber is used to statistically describe EMC performance, e.g. shielding effectiveness or immunity and it performs a similar role to an open area test site (OATS) or anechoic chamber despite being have boundaries that behave differently by reflecting waves.

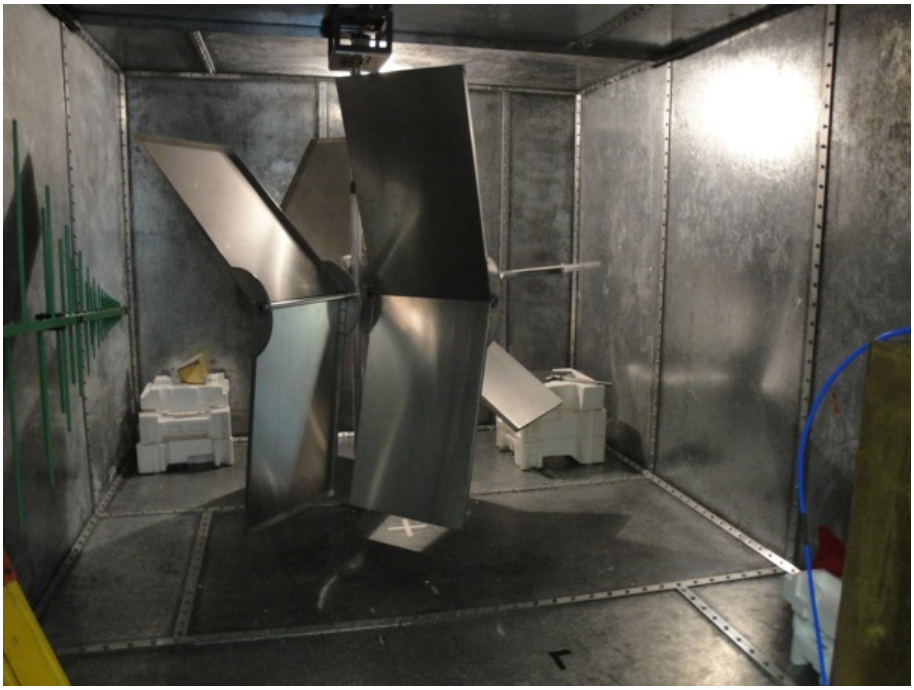

Figure 3.1: The York reverberation chamber

The purpose in using a reverberation chamber here is different to that in electronic system or communications measurements because it is used to simulate reverberant environments in aircraft rather than measure a parameter such as shielding effectiveness. One can see the usefulness of the stirrer if the alternative is considered. In a very reverberant cavity then it is possible (as shown in the next chapter) for the signal to vary by tens of dBs by moving the antennas just a few cms. If we wished to be sure that a system would work reliably then we would have to move the antennas to many different positions for every measurement which would be a very time consuming process. With the stirrer there is no need to move anything. The measurements can be automated so that a computer can take a measurement, rotate the stirrer then take another measurement and the effect is exactly the same.

**Reverberation chamber properties**

There are certain properties of a reverberation chamber that are useful to understand in advance of the upcoming measurement chapters. The notational style used by Hill will be followed and the reader is referred to his papers and recent book [42] for further information as these provide the best summaries of reverberation chambers. In analysis of reverberation chambers it is useful to describe the field in terms of plane waves where we say that the electric

field at a point is the sum of the plane waves all angles[42]:

$$\vec{E}(\vec{r}) = \iint_{4\pi} \vec{F}(\Omega) \exp(i\vec{k} \bullet \vec{r})\, d\Omega \tag{3.36}$$

where $d\Omega$ is the solid angle, $k$ the wavevector and position vector $r$ describing the position in the spherical coordinate system. This plane wave $\vec{F}(\Omega)$ can be decomposed into its components, the $\alpha$ and $\beta$ angles (elevation and azimuth respectively), and then into real and imaginary components so that the $\alpha$ part can be described as:

$$\vec{F}(\Omega) = \alpha F_\alpha(\Omega) + \beta F_\beta(\Omega) \tag{3.37}$$

$$F_\alpha(\Omega) = F_{\alpha r}(\Omega) + iF_{\alpha i}(\Omega) \tag{3.38}$$

The same can be done for the $\beta$ component. We may then look at the ensemble averages of these fields

$$\begin{aligned} &\langle F_\alpha(\Omega)\rangle = \left\langle F_\beta(\Omega)\right\rangle = 0 \\ &\langle F_{\alpha r}(\Omega) F_{\alpha i}(\Omega)\rangle = \left\langle F_{\beta r}(\Omega) F_{\beta i}(\Omega)\right\rangle = 0 \end{aligned} \tag{3.39}$$

These can be assumed from an understanding of what is actually taking place in a reverberation chamber: many different rays are effectively reflecting around the chamber, and because of the stirrer reflecting at different angles they are arriving at measurement point $\vec{r}$ with randomly different phases and polarizations. Measurements also back this up, at least for a well behaved chamber that has many modes at the frequency of interest.

The above lead to the electric and magnetic field also having ensemble averages of 0:

$$\left\langle \vec{E}(\vec{r}) \right\rangle = \left\langle \vec{H}(\vec{r}) \right\rangle = 0 \tag{3.40}$$

It is shown by Hill (by a double integration of $F$ and $F^*$ at two angles) that:

$$\left\langle \left| \vec{E}(\vec{r}) \right|^2 \right\rangle = E_0^2 \tag{3.41}$$

and in terms of components:

$$\begin{aligned} \langle |E_x|^2 \rangle = \left\langle |E_y|^2 \right\rangle = \langle |E_z|^2 \rangle = \frac{E_0^2}{3} \\ \langle |H_x|^2 \rangle = \left\langle |H_y|^2 \right\rangle = \langle |H_z|^2 \rangle = \frac{E_0^2}{3\eta^2} \end{aligned} \tag{3.42}$$

To summarise the results so far: the averages of the fields are all 0, the averages of the power in the electric field in each component are equal and the powers in the 3 magnetic components are equal. This is to be expected from a randomly scattered field in a closed box there is no sense of direction of single direction of propagation.

**Distributions**

It can also be that the real and imaginary components of the field are uncorrelated and each have mean 0. We can denote their variance $\sigma$. Since we are considering a field made up of many identically distributed independent variables in terms of the many different paths then it should be possible to apply the central limit theorem and describe the statistical distribution of the real or imaginary components of the field as Gaussians:

$$f(E_{xr})=\frac{1}{\sqrt{2\pi\sigma}}\exp\left(-\frac{E_{xr}^2}{2\sigma^2}\right) \tag{3.43}$$

The magnitude of the E field is obtained by application of pythagoras' theorem to the complex vector and when there are two components each with a Guassian distribution, the magnitude is distributed with a Rayleigh distribution and their power is exponentially distributed. Similarly when looking at the total field rather than a component we get chi and chi squared distributions for magnitude and power respectively, each with 6 degrees of freedom:

$$\begin{aligned}
f(|E_x|)&=\frac{|E_x|}{\sigma^2}\exp\left(-\frac{|E_x|^2}{2\sigma^2}\right)\\
f(|E_x|^2)&=\frac{1}{2\sigma^2}\exp\left(-\frac{|E_x|^2}{2\sigma^2}\right)\\
f(\left|\vec{\mathbf{E}}\right|)&=\frac{\left|\vec{\mathbf{E}}\right|^5}{8\sigma^6}\exp\left(-\frac{\left|\vec{\mathbf{E}}\right|^2}{2\sigma^2}\right)\\
f(\left|\vec{\mathbf{E}}\right|^2)&=\frac{\left|\vec{\mathbf{E}}\right|^4}{16\sigma^6}\exp\left(-\frac{\left|\vec{\mathbf{E}}\right|^2}{2\sigma^2}\right)
\end{aligned} \tag{3.44}$$

The received current distribution is given by:

$$f(|I|)=\frac{|I|}{\sigma_I^2}exp\left(-\frac{|I|^2}{2\sigma_I^2}\right) \tag{3.45}$$

It is possible to work instead with the received power which is distributed as:

$$f(P_r)=\frac{1}{2\sigma_I^2R_r}exp\left(-\frac{P_r}{2\sigma_I^2R_r}\right) \tag{3.46}$$

Again there is a Rayleigh and Exponential distribution. Why there is the Rayleigh current distribution is easy to see intuitively. The randomness in the E field components results in a random phase and in fact gaussian real and imaginary components for the current. Where this occurs the magnitude will always be Rayleigh distributed, as one can see by integrating a ring around a polar plot.

**Q factor from the statistical ensemble**

Previously the Q-factor of a resonant cavity was expressed in terms of its physical properties. That was for an ideal case where there are no leaky joints, apertures etc. If one wishes to obtain Q from measurements then it is possible to do so by taking the ensemble average of the power received and dividing it by the power transmitted:

$$Q = \frac{16\pi^2 V}{\lambda^3} \frac{\langle P_r \rangle}{P_t} \tag{3.47}$$

$$Q = \frac{16\pi^2 V}{\lambda^3} \langle S_{21}^2 \rangle \tag{3.48}$$

$S_{21}$ here is the forward scattering parameter. This can be viewed as the transfer function of the system, it simply tells us the ratio of input to output. This formula also takes into account the polarization mismatch of the antennas but not the efficiency which must be factored in when converting from the measured $S_{21}$ to the power ratios in the chamber.

**Summary**

At this point it has been shown that reverberation chambers are electromagnetically complex and are best described by their statistics. We have seen the distributions for the field components and field and their powers. In particular, a Rayleigh distribution for the S21 network parameter, which is voltage based, would be expected. It has also been shown that the Q-factor can be found through a measurement of the network parameters averaged over stirrer positions

## 3.3 Frequency and time views of a channel

Now that the main physical characteristics of the reverberation chamber have been discussed the next step is to look in more detail at how to measure it and analyse it in such a way as to gain useful information about how it will affect a communications system.

For all the complexity of the modal structure of the channel, in a communications system what the designer wishes to know is what does the communications channel look like. The path a radio wave travels between a transmitter and receiver is known as the radio or transmission channel and there are two ways one can look at it: what happens to the signal in the time domain, or what happens in the frequency domain. Generally in communications the frequency domain is used the channel is analysed through transmitted sinusoids which get

distorted as they pass through the channel, arriving at the receiver with a different amplitude and phase. The key benefit of the frequency domain is that communications systems have limited bandwidth and by definition it is easy to look at this bandwidth in the frequency domain, from which one can quickly reach a conclusion about the quality of the channel in terms of how flat it is and the attenuation. Measurement equipment such as network analysers work in the frequency domain, sweeping across frequency which also makes it the more attractive domain.

On the other hand we have previously explained that we expect a key difficulty of metal box type channels to be the large energy storage or high Q. The effect of symbols decaying slowly and interfering with one another is easier to visualise in the time domain. Therefore we would expect that we shall move between the two domains as is appropriate using the Fourier Transform.

### 3.3.1 Network analysers

All channel measurements presented in the later chapters were carried out using either a network or a spectrum analyser with the majority using the network analyser. The spectrum analyser is the simpler tool intuitively and provides a frequency domain view of a signal. A time domain signal is passed through narrow filters, envelope detected and the resulting magnitude from each filter is shown on the display. This is done at a number of frequencies to build a frequency domain view of a signal. For example, if measuring a ZigBee signal one would see a plot such as that in figure 3.2 This is an absolute measurement of power and contains no phase information. On the other hand a network analyser is a tool that gives a relative measurement including signal phase and thus generates the input signal itself. As shown in figure 3.3, in any RF system there is an input voltage which is best viewed as a wave rather than static voltage. Part of this will enter the system, part will be reflected and there will be voltages on the input and output terminals that are made up of forward and reverse waves. In the case of the upcoming measurements it is voltages close to the transmit and receive antenna terminals that are measured. A network analyser measures the various ratios of transmitted to reflected voltage waves at the input and output ports and these ratios are referred to as S-parameters. These are favoured over input and output impedances (which would give the same information) because impedance measurements at high frequencies are not practical at high frequencies. S11 is the ratio of incident voltage to port 1 to reflected voltage on port 1 and S21 is the ratio of transmit-

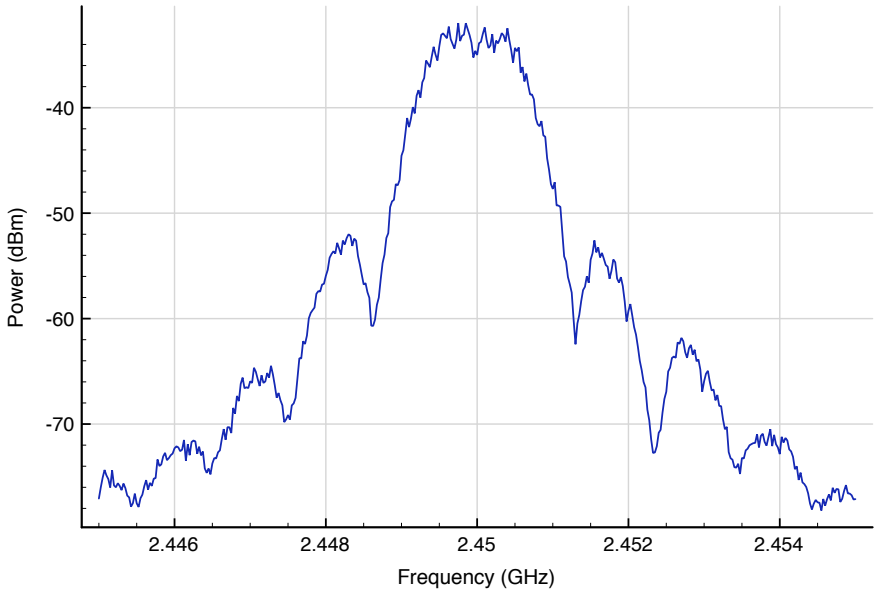


Figure 3.2: Spectrum of a continuous ZigBee signal in anechoic measurement with 100KHz Video and Resolution bandwidths

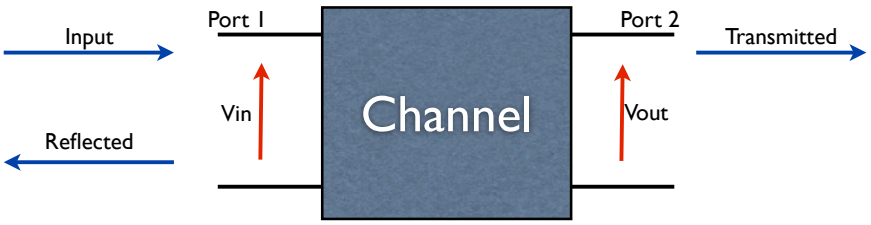


Figure 3.3: Forward and reflected waves into a system

ted voltage out of port 2 to input voltage on port 1, i.e. the transfer function. Like the spectrum analyser this measurement is swept across a frequency range. For both measurements types there are compromises to make: for example how narrow to make the filters since greater resolution in frequency gives a longer sweep time and longer measurement. Also there are options for averaging which are useful in preventing noise from masking the signal.

### 3.3.2 Fourier representation of a channel

Network analyser measurements, as just described give what in linear systems theory is called the impulse response of the channel or at least the band-limited or filtered impulse response.

In other words, if we passed a time domain impulse through the system the output would be the impulse response and if we transformed it to the frequency domain it'd be the same as we measure with the NA. It is possible to do this because a resonant cavity is a linear system. By virtue of its many reflections, any signal, at a given frequency, coming out of an antenna will be delayed and attenuated at the receive antenna but it will still be of the same frequency. We may write the received signal in each domain as:

$$r(t) = \sum_{n\in\mathbb{N}} \alpha_n t_x(t-\tau_n) \tag{3.49}$$

$$R(j\omega) = \sum_{n\in\mathbb{N}} \alpha_n T_x(j\omega) e^{j\tau\omega} \tag{3.50}$$

The first consequence of this is that we can apply the Fourier Transform to the signal to obtain it's time domain representation. With the transformed signal we can see how long it takes for an input to decay which is critical for system performance. If symbols are transmitted every $\mu$s, but it takes 10$\mu$s for each of those to decay, then they will all add and we shall have no idea which symbol we are receiving at the receiver. NA measurements allow us to see when this will occur.

Linear systems theory also makes it very simple to test a communications system design against the channel. Once the transfer function is obtained it is possible to find the output to an arbitrary input by convolution in the time domain, or even simpler, by multiplication in the frequency domain. To see how this might be used, assume an aircraft cavity where the Q is known to be in the region of 700-900 approximately. One might wish to know how this would perform under a variety of modulation schemes and bit rates. Using either a reverberation chamber or some form of simulation we could obtain 100 different channels with the upper Q of 900. If simulating a radio system we could then take each of the modulation schemes and convolve or multiply these with the channel and see whether it produced the correct data in each simulation.

#### Transforming from a band-limited frequency domain signal to the time domain

Practical measurements, will only cover a particular frequency range, rather than going from DC up to the maximum frequency. This is due to the nature of Signal and Network Analysers but also it only makes sense to measure a limited band when a communications system operates in a limited band. This complicates the transformation process from the time to frequency domain. If for example there are 1600 samples at the end of a measurement and these samples

represent 2.4GHz to 2.484 then one can't necessarily blindly apply the Fourier Transform as this is not going to look the same in the time domain as going from 0 to .084GHz.

There are a number of ways to address this:

- Pad the measured spectrum with zeros before transforming
- Use the frequency shift theorem

In the first case assume that 1600 measurements were taken over a frequency range of 160MHz varying from 2.32 - 2.48 GHz. The frequency spacing here $\Delta\omega$ is 100kHz. If zeros were to be added this would require an additional 23200 samples to be prefixed to extend the spectrum to zero frequency. It may be noted that what is happening when transforming in this case is multiplication by zero. This can lead to an unnecessarily slow operation and in fact, surprisingly, the resulting spectrum will still have the same envelope as if we had blindly applied the transform to the initial 1600 measurements.

The frequency shift theorem provides a way to deal with this. It is described in the continuous and discrete cases by:

$$X(\omega-\omega_0) \Leftrightarrow X(\omega)e^{j\omega_0 t} \tag{3.51}$$

$$H(n-i) \Leftrightarrow h(k)e^{j2\pi\frac{ik}{N}} \tag{3.52}$$

The magnitude of a complex number can be described as the product of magnitudes and since the magnitude of the exponential is 1, this shows that the time domain signal will not change the shape of its magnitude response as the frequency shifts.

It is necessary to be careful using the frequency shift theorem because it is normally required that a real signal is present in the time domain and therefore symmetry must be maintained, meaning that as one half of the spectrum shifts right the other shifts left (the negative frequency half). By doing this the original envelope is maintained but it is modulating a cosine wave (formed by the addition of the positive and negative exponential) and it is for this reason that the frequency shift theorem is often known as the modulation property. In fact in order to make use of this property practically, measured data was put into a complex baseband form and then transformed, which will also provide the signal envelope.

Both techniques were used in the course of the described measurements. Initially zero padding was used for its simplicity but when in some cases (where the frequency spacing was small) the computer run time because a problem the second method was used. Each is equally

valid in terms of demonstrating the delay spread in the time domain. For both cases care was taken so that a frequency range was used that could be zero padded, i.e. if we padded back to zero frequency we would have a sample on 0.

**Metrics for describing the channel in the frequency and time domains**

There are certain metrics which can be used to characterise a multi-path communications channel. The most prominent ones are coherence bandwidth and mean delay spread [43]. The former measures how flat the channel is in the frequency domain and the latter the delay in the reflected paths when they arrive at the receiving antenna. In reverberation chambers the typical metric used would be the Q-factor but in communications delay spread is more common and so this will be used. There are already many papers available reporting different error rates for different modulation schemes for different delay spreads. One could use these to predict how ZigBee and any other system might perform in a given environment.

There are in fact two delay metrics used and these are delay spread and mean delay. Mean delay is defined as follows:

$$\bar{\tau} = \frac{\sum_k a_k^2 \tau_k}{\sum_k a_k^2}$$

RMS delay spread, $\sigma$ is

$$\sigma = \sqrt{\bar{\tau^2} - (\bar{\tau})^2}$$

$$\bar{\tau^2} = \frac{\sum_k a_k^2 \tau_k^2}{\sum_k a_k^2}$$

Note that the square of the amplitude is being used because these metrics are being calculated from a power not amplitude profile. With these values calculated one may compare them to the symbol duration to provide a rough estimate of whether there will be a problem or not.

Although these are metrics that have originated from the field of communications they happen to have a close relationship to the Q of a resonant cavity. Each resonant mode has a Q-factor and this corresponds to an exponential decay with time constant $\tau$ where $Q = \tau * 2\pi f$.

If we put this exponential decay into the formula for mean delay we see:

$$
\begin{aligned}
&f^2(t)=Ae^{-\frac{t}{\tau}}\\
&\overline{\chi}=\frac{\int_0^\infty f^2(t)\,t\,dt}{\int_0^\infty f^2(t)\,dt}=\frac{\int_0^\infty Ae^{-\frac{t}{\tau}}\,t\,dt}{\int_0^\infty Ae^{-\frac{t}{\tau}}\,dt}=\frac{\int_0^\infty e^{-\frac{t}{\tau}}\,t\,dt}{\int_0^\infty e^{-\frac{t}{\tau}}\,dt}\\
&\int_0^\infty e^{-\frac{t}{\tau}}\,dt=-\tau\\
&\int_0^\infty e^{-\frac{t}{\tau}}\,t\,dt=\left|(-\tau)^2\left((-\tau)-1\right)e^{-\frac{t}{\tau}}\right|_0^\infty\\
&=-(-\tau)^2=-\tau^2\\
&\overline{\chi}=-\tau^2\div-\tau=\tau
\end{aligned}
\tag{3.53}
$$

Note that $\chi$ has been used here for mean delay to avoid confusion with the time constant.

One can see that the mean delay and the exponential delay associated with the Q are effectively the same thing. Conveniently this means that one can work out a rough Q for a given propagation environment and directly relate it to a standard communications metric. It is particularly prominent in the mobile phone literature, but discussions of wireless LANs (e.g. [22]) also refer to delay spread in terms of what the system was designed for. We also see in Appendix B that this relationship allows for a novel calculation of antenna efficiency.

### 3.3.3 Types of fading

#### Flat and frequency fading

What is considered narrow and what is considered wide depends on the signal bandwidth. In frequency selective fading fluctuations in signal power occur over within bandwidth of the communication system, the alternative being that no fluctuations can be seen over the bandwidth being used. This distinction is significant because all the standard formulas for assessing system performance are valid only for the flat fading scenario. However it is likely that for very high Q-factor in a reverberant environment the fading will not be flat. In this situation it is not possible to predict the error rates through any analytical means. Instead simulation must be used to assess performance or measurements. How flat a channel is can be quantified by examining the autocorrelation function in the frequency domain and the definition of coherence bandwidth in Rappaport [43] uses the ACF stating the coherence bandwidth to be that

over which the correlation function is over 0.5.

**Fast and slow fading**

This refers to the fluctuation in signal strength over time. Given that all communications on the aircraft are fixed relative to one another and communications with the ground are not being considered it might be thought that time variation of the channel does not need to be considered as it might be when assessing a mobile phone system and the fading that occurs as the user moves. Unfortunately this is not necessarily the case. Aircraft wings will continually flex during flight and there will be vibrations caused by the engines so aircraft are not completely rigid. Discussions with Flaviir project partners on the issue of panel vibration suggest that any movement is no more rapid than 10s of Hz and the amplitudes are small and of order 1mm. In a very resonant environment even small changes in geometry can bring about large signal strength and phase variation. For example, in the mock up wing section demonstrator produced by York for the Flaviir roadshow, one could push down on a panel such that it was displaced by 3mm at the centre point and signal strength would change by nearly 10 dBs.

Just as fading caused by multi-path can be characterised as frequency selective or flat, fading caused by a moving channel may be characterised as fast or slow. What this means is amplitude changes may occur during the duration of a symbol (fast) or the channel may appear static over the duration of a symbol (slow) In the former case there will be symbol distortion which can lead to incorrect symbol detection.

### 3.3.4 Calculating performance under fading and noise

Inside aircraft structures and cavities there are likely to be two separate issues leading to errors and the deterioration of a communications link:

- Noise and interference generated by other aircraft systems
- Fading caused by reverberant properties of aircraft bays

An error is defined for the purposes of this section as anytime a symbol is incorrectly decoded as being any symbol other than the transmitted one, for example 00 instead of 01. One may also consider chip or packet errors instead. Certain types of error are easy to predict, others very complicated. Later chapters will deal with performance that is not easily calculated by

simple (or complex) equations but first some examples from the standard books and literature are considered to see what can currently be calculated.

**The simplest case - a flat channel and Gaussian Noise**

Firstly noise and interference will be considered. Typically in communications noise is assumed to be Additive, White and Gaussian which is a good assumption because a lot of the noise in receivers is Gaussian and environment noise before the receiver is also often Gaussian.[1]

It is possible analytically to calculate bit error rates in Gaussian noise for a lot of systems. For example, in a receiver circuit the received signal will undergo an integration over its period followed by sampling. If at the point of sampling a '1' would be seen as $4\mu V$ and a '0' as -$4\mu V$ then clearly in order for an error to occur $4\mu V$ must have been added or subtracted by the noise dependent on whether a 0 or 1 is being sent. At the input to a receiver there will be a particular signal to noise ratio in terms of signal power and noise power. From this one may look at the energy over the symbol period to calculate the likelihood of incorrectly decoding the symbol. The probability of error will involve a cumulative Gaussian and so tables are generally used for analysis. One may write the probability of error in terms of the Q function which is defined as:

$$Q(z) = \int_z^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}\, dx \tag{3.54}$$

where $x$ takes the value of the Gaussian and $z$ is the value we are integrating up to. It gives the probability of a particular value being exceeded for a 0 mean, variance of 1 Gaussian. Note that this is completely unrelated to the earlier Q describing energy storage in resonant systems. The probability of a bit error $P_b(E)$ can be calculated by:

$$P_b(E) = Q\left(\sqrt{\frac{2E_b}{N_0}}\right) \tag{3.55}$$

where $E_b$ is the energy in a bit and $N_0$ is the single sided noise spectral density. Note that even though this is for noise in the simplest case, the calculation is not trivial and requires numerical techniques or tabulated values. This result applies when one assumes perfect synchronisation in the sense that the carrier phase is known precisely and so is only offering a maximum performance. The result will also change if one assumes non-coherent demodulation where no phase information is used for demodulation. In that case the probability of bit

[1] The central limit theorem ensures that if there are lots of things all adding noise the noise will appear Gaussian.

error will instead be given by:

$$P_b(E) = \frac{1}{2} exp\left(-\frac{E_b}{2N_0}\right) \tag{3.56}$$

In this case the form is simpler and can be calculated directly. It should be noted that performance is inferior for the non coherent case.

**Fading Channels**

It has already been mentioned that one can expect a Rayleigh distribution in a reverberation chamber which has the form:

$$f(|E_x|^2) = \frac{|E_x|}{\sigma^2} e^{-\frac{|E_x|^2}{2\sigma^2}} \tag{3.57}$$

This, in fact is the case for a system with no direct path. When there is a direct path there is a different distribution known as the Ricean distribution:

$$p(r) = \begin{cases} \frac{r}{\sigma^2} e^{-\frac{r^2+A^2}{2\sigma^2}} I_0\left(\frac{Ar}{\sigma^2}\right) & for(A \geq 0, r \geq 0) \\ 0 & for(r \leq 0) \end{cases} \tag{3.58}$$

where A is the amplitude of the dominant component, $I_0$ a Bessel function of zero order and $\sigma = \sqrt{\frac{\chi}{2}}$ with $\chi$ being mean power in the random component. Again this is a measure that occurs typically in mobile phone literature but it's also relevant to reverberation chambers and can be seen by looking at a polar plot of measurements data and seeing how the points cluster around the centre of the plot. If they are not centred then the phase is not totally random and there is likely to be a direct path. The Ricean case is very difficult to deal with analytically due to the Bessel function and so it is easier to work with the Rayleigh distribution. The Rayleigh represents a worse case scenario where there is no direct component. Given that the amount of direct component will be different in every environment it is logical to work with the worst case and assume no direct path. Many of the measurements in the following chapters are setup to deliberately avoid a line of sight path.

It is possible to calculate the probability of error for a Rayleigh Fading channel providing that the errors are caused by the modulation of the signal to noise ratio rather than inter-symbol interference. In other words a flat channel is being assumed where a flat channel is one which in the frequency domain appears flat over the bandwidth of the data being sent. If there is a bandwidth greater than the coherence bandwidth then we don't a flat channel and we might expect errors.

Provided there is no inter-symbol interference, the signal to noise ratio is modulated as antennas are moved or as the boundary conditions change in the propagation environment. It is

possible to think in terms of instantaneous SNR $\gamma$ and the probability of a given instantaneous signal to noise ratio and from this the probability of an error. Mathematically we can write the average SNR and the probability of an error as:

$$\overline{\gamma} = \int_0^\infty \gamma p_\gamma(\gamma) d\gamma \tag{3.59}$$

and the same approach is given for PER with

$$P_b(E) = \int_0^\infty P_b(E|\gamma) p_\gamma(\gamma) d\gamma \tag{3.60}$$

where $P_b(E|\gamma)$ might be the Q function just described. This calculation is rarely trivial and can be very hard to calculate for the more complex modulation schemes where coding is employed or where multiple antennas are used. Simon and Alouni [44] show how error probability may be calculated using a unified framework based around moment generating functions. As an example of the technique one can consider the problem of calculating the average SNR $\overline{\gamma}$. This may be rewritten with moment generating functions as:

$$\overline{\gamma} = \left. \frac{dM_\gamma(s)}{ds} \right|_{s=0} \tag{3.61}$$

where

$$M_\gamma(s) = \int_0^\infty P_\gamma(\gamma) e^{s\gamma} d\gamma \tag{3.62}$$

Effectively one would end up performing Laplace transforms to obtain moments from which average SNRs or error rate probabilities can be obtained. There is no attempt to cover the results here as they are generally complex and numerous with complicated derivations for the many different types of modulation including the many types of MIMO system. There is a special case for binary FSK that we shall give here though. The probability of symbol error is:

$$P_s(E) = \frac{1}{1+\overline{\gamma}_s} \tag{3.63}$$

Other variations can be seen in 3.4. For more anything else the reader is directed to [44]. One must again emphasise that all the results in the aforementioned are only for the flat fading case and ignoring equalisation, time varying channels and coding issues. Assessing communications system performance is non trivial and must be backed up by simulation and measurements especially for a safety critical application as described here.

For most digital wireless systems data is sent in packets and performance metrics such as sensitivity are specified in terms of packet error rate. In ZigBee, for example there are a

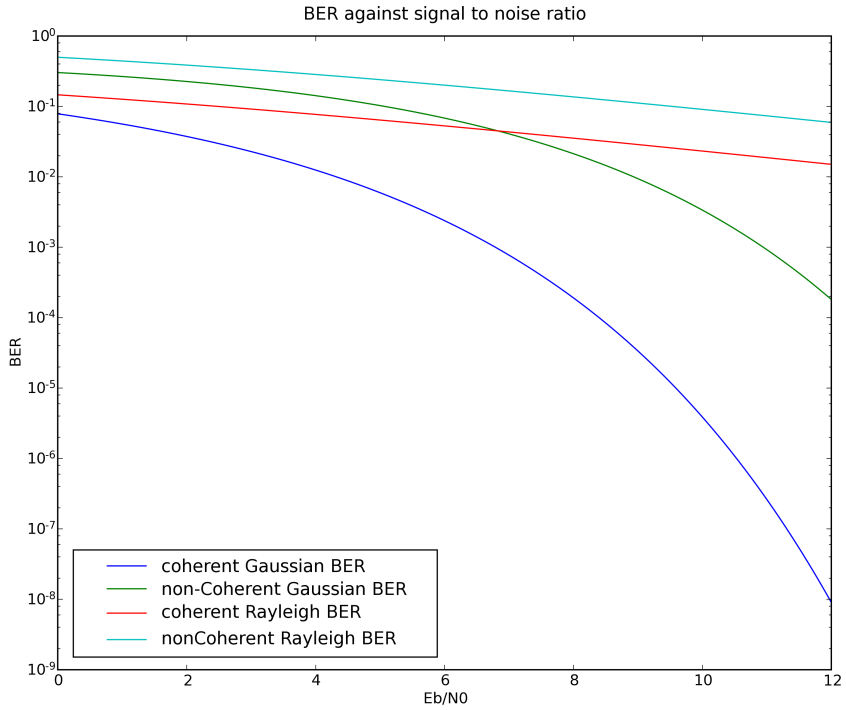


Figure 3.4: BER against average SNR for Gaussian and Rayleigh channels for FSK

maximum of 127 octets in a packet and only one of the 1016 bits needs to be in error to cause a packet error. Packet error can be linked to bit error by considering the probability of all bits in a packet being correctly received and then looking at 1- this quantity. Probability of packet error will be approximately:

$$PER = 1 - (1 - P_b)^N \tag{3.64}$$

where $P_b$ is the probability of a bit in error and equal to the BER. With packets there will be additional complications such as packet synchronisation on the preamble which could fail and so the problem is slightly more complicated.

### Noise Sources

Inside an aircraft the noise may appear less Gaussian than in mobile communications since interference may be concentrated around the particular frequencies other equipment operate at, for example noise from motors and actuators. EMC is a difficult issue in aircraft because they contain a lot of devices from radar to flight control computers to high voltage electrical system. These are all contained within a relatively small space (compared to an industrial or home environment) making interference more likely, and this is even more so on a small

UAV. In [45] the following are given as sources of EMI: current carrying wires and cables, electromechanical components like relays, other switching circuits, oscillators, radio frequency and high power transmitters, lightning, high-tension power lines, electrostatic discharge etc, extra-terrestrial sources such as the sun, cosmic effects, radio stars etc. This also points out that carbon fibre may not always provide the same level of shielding effectiveness that is provided by metals. Lightning in particular can be a serious problem since up to 200KA can be swept through the aircraft.

The other main external noise source is static charging, which is examined in [46] and comes from potentials between clouds, from water friction with water particles and the engine exhaust. When charge builds up there will inevitably be a discharge and this leads of EMI. It is reported that these effects do not generally cause problems for narrowband radio systems and there is greater danger to wide-band systems such as digital electronics and so in a wireless aircraft we might be more concerned here with the microcontrollers than the radios. The paper concentrates in particular in the range of 1 to 100MHz and systems in the GHz bands should not be affected.

Another point of interest is that in [45] and [47] the role of power lines, connectors and data cables is emphasised in terms of their EMI role and the power they can radiate, especially if terminations are faulty or badly installed etc. This is one area of strength in a wireless aircraft in that by removing a lot of long current carrying wires we also reduce many potential radiators of electromagnetic energy. The more of an aircraft one can make completely wireless the more reliable those wireless systems should become. As a final point most of the avionic systems mentioned in [47] that might contribute to EMI problems operate in the hundreds of MHz or around 1GHz. Therefore the effects of existing systems can be reduced by only using bands nearer the ISM band and the 5GHz band used for 802.11. Lightning on the other hand has energy beyond 1GHz [48] and in all these systems one must be aware of potential product terms developed in mixers that could cause out of band interference to enter the band of interest.

Gaussian noise is a particular issue when communicating over a distance, such that a very weak signal may be received, with receiver noise becoming significant. Again this is different for the situation on an aircraft whereby distances may be small and energy high, due to resonance. Therefore, sensitivity would be expected to be far less important as a measure of radio system performance for an avionics system than it is for a satellite or even mobile phone system. On the other hand, in an environment with a lot of electrical equipment the ability to

resist narrowband on board noise sources on the aircraft is likely the most important requirement. Any radio must have either sufficiently wide channels or sufficient channels over a wide range to deal with interference sources.

## 3.4 Propagation in an airframe

**Model of an aircraft**

Near the beginning of the chapter it was stated than an aircraft may be viewed as a set of resonant cavities. Having now discussed what these are and what their main characteristics are they must now be related back to aircraft. In [49] it is stated that "aircraft structures are basically boxes" and this is stated from a structures point of view rather than electromagnetic. A lot of the analysis in that book and in [50] looks at tensions, loads and stresses on box structures so it appears to be a reasonable assumption to treat aircraft as boxes when looking at them from an EM perspective.

If one looks at some of the images of aircraft cross sections found in structures texts such as [49] and [50] then this is also apparent. The most obvious structure is the fuselage and whilst this is not rectangular it is a closed structure not dissimilar to a large rectangular box. In fact when it is split into two halves (cabin and cargo areas in a commercial jet, for example) then it looks even more like two boxes. A large number of aircraft types are shown in the cross sections of the aforementioned texts but all have similar features. We find that the wings may be viewed as a set of smaller boxes, some coupled through ribs and some cut off such as the fuel tank areas. The fuselage will be divided up into many areas, from avionics bays, to the cockpit, cabin area on a commercial jet, spaces for landing gear, the tail section etc.

So it is clear that an aircraft has lots of closed boxes and tubes making up its structure but what about the assertion that these structures would be reverberant. Until recently all commercial and most military aircraft were constructed from a combination of steel and aluminium alloy with the latter being by far the most widely used and steel reserved for certain joints, fasteners, joint plates, the undercarriage etc. Titanium and other metals may also be used. Regardless of the particular type of metal the main point is that these structures are metal and thus have high conductances and reflect electromagnetic waves. Further, large metal boxes are exactly what a reverberation chamber is and so reverberation chambers would appear very suitable for modelling aircraft structures electromagnetically.

There are some complications: more recently carbon composites have been replacing aluminium alloys in aircraft construction. Now carbon fibre is a conductor but it has a low conductivity compared to metals and so we might expect a lower Q-factor in aircraft sections made from carbon fibre. This may not be as big a change as it first appears though because to deal with issues such as lightning strikes it is now common to weave metal into the structure, thus making it a much better conductor. Further, tests by colleagues in the York EMC department have suggested that CFCs may be regarded as very similar to metals for our purposes. The other major complication is that reverberation chambers are typically well sealed and have a very high Q which is desirable for EMC so as to ensure a low correlation between measurements at each stirrer position. Although aircraft are metal boxes, many of the boxes will be leaky and have apertures for rods, hydraulics, cable looms etc. Further they may contain materials that will absorb electromagnetic energy; the miles of wire in aircraft will certainly absorb energy. As a result the Q in a modern aircraft, regardless of its type is likely to be very much lower than a reverberation chamber. We shall look at measured Qs later but it is worth noting that as all the cable is replaced by more and more wireless radios then the Q-factor is likely to rise inside the airframe and we shouldn't assume a Q based on aircraft filled with wiring. We may also lose a lot of the hydraulic links and just have electrical based actuators communicating to a computer wirelessly with little in the space in between. Electric actuation is already on the A380 as a backup and it is likely to appear much faster in the less conservative UAV space. With time, parts of an aircraft may become more empty.

To summarise our abstract view of an aircraft, we shall regard it as a set of coupled and nested boxes of differing sizes. Some of these boxes will be large such as the aircraft fuselage, others may be smaller such as parts of a wing, divided by ribs or perhaps an avionics bay. We do this for a number of reasons:

- It is reasonable according to the aircraft structure literature
- Such structures are easier to analyse than complex geometries
- No aircraft is the same as so there is little to be gained in looking at one detail in depth for a single aircraft

### 3.4.1 Potential challenges for communications systems

The propagation within an aircraft will be significantly different to that outdoors, within a house or an office space: all the sorts of areas where wireless networks have been used extensively up to now. The propagation is likely to have more in common with mobile phone propagation where there are a number of multi-path rays, travelling significant distances relative to the wavelength and with time variation of the channel. This is because of the resonant nature of the environment that we have just discussed. Unlike the mobile phone example the distances are short but the fact that reflections happen many times means that the end effect could turn out to be similar in that an attenuated signal has travelled many times round and has a significant phase shift relative to a direct signal or first received ray. We might find that the inter-symbol interference in a fuselage stops ZigBee or Bluetooth or 802.11 from working correctly.

One could decide that that these issues should be ignored and we should say that wireless aircraft should be designed with a non reflective coating inside equipment bays, wings and other areas but this is not practical for adding wireless equipment to current aircraft and would be an unnecessary expense if if it found to be unnecessary. Also, these issues with resonant environments may apply to other sectors such as shipping, trucks etc and are worth investigation.

### 3.4.2 Previous work in modelling the inside of an airframe

There has already been some research into propagation inside aircraft, although most of this has been looking at commercial aircraft with a view towards safe usage of consumer electronics. For example: ensuring radiation from portable games consoles and similar technology does not couple significantly into avionics bays and the cockpit. Freyer did some work on coupling between different aircraft cavities [51] for this purpose. In this he identifies the use of the reverberation chamber as good way to simulate the environment inside an aircraft and shows how, due to many differences between different cavities, a statistical approach is needed which is the approach taken in this project.

There has been one piece of work looking specifically at wireless sensors in an aircraft which was carried out by Hof and Stancil [52]. This was looking at UAVs, aircraft wings and other highly reverberant environments. In this they report on the close link between Q-factor and delay spread which was something identified early into the Flaviir work although they

don't quantify the link however. Delay spread is plotted up to 9GHz for two different resonant environments. However the Q of these environments is not given nor is any absorber used. The need for a high number of nodes for reliable communication is emphasised and different modulation formats of different complexities are compared according to the SNR needed to maintain a BER of $10^{-4}$. Unfortunately most of the conclusions reached are specific to the two propagation channels used and not very general.

In [53] a reverberation chamber is used as a substitute for an aircraft by lowering the chamber's Q-factor to realistic levels. The aim of this work is to design a system for detection of mobile phone signals but the paper concentrates on the characterisation of the reverberation chamber at low Qs and states that research has found Q-factors in a Boeing 757 aircraft to vary from 140 to 1250 as frequency moves from 300 to 3200MHz. One issue brought up by this paper is how well the statistical qualities of the chamber are maintained when the Q becomes very low.

Johnson, Hatfield and others at the Naval Surface Warfare Centre (NSWC) have produced some significant research into the reverberant nature of aircraft. In one of their reports entitled "Electromagnetic Reverberation Characteristics of a Large Transport Aircraft" [54] they take measurements on a decommissioned Boeing 707-720B and compare those measurements to a simulated avionics bay in a reverberation chamber. It was found that losses were 15dB greater for the avionics bay and 12dB greater for the cockpit than in the chamber setup and that the maximum to minimum field would vary by up to 20dB past 800MHz. This confirms a point made early on in the report that "if a cavity has an adequate mode density and a sufficiently low loss factor, it may be possible to bound the bay reponse under mode stirred conditions". This is effectively what we are doing in a large part of this PhD Thesis, where we aim to bound the communications system performance under mode stirred conditions. Their results suggest that real losses are greater than one might expect from the simulated environment but at the same time there is a modal structure to the fields causing a change in power by order 100 over short distances. They point out that there are a lot of openings to the cockpit and avionics bay with apertures for wiring looms for example and that these as well as other aspects of the surrounding area of the avionics bay will account for differences between the simulated and Boeing 707 results. The good news here is that in using a reverberation chamber setup to simulate an aircraft we may get too negative impression of system performance but where we are trying to bound the reliability of a system this is not a bad thing compared to believing a

system will work that won't.

A second report by the NSWC used the same decommissioned Boeing 707 and carried out further testing on the cockpit and avionics bays and also the passenger cabin's electromagnetic properties. These tests were more extensive, considering coupling into avionics systems of various kinds and the decay times and Q-factors found in these environments. Two types of stimulus were used which led to slightly different results but in terms of the decay time it was of order 100ns for the passenger cabin, around 20-50ns for the cockpit. They found that Qs varied from order 10 to 100 over the frequency range 0.2GHz to 6GHz. Moving to Q-factor they presented Q in two way, one where Q was derived from the time delay and another they called $Q_s s$, the steady state Q derived from the ratio of power in to power out. The steady state results suggest that at 2.4GHz , the standard wireless network frequency range, a Q of anywhere between 50 and 150 could be expected. This differs significantly from the time response derived Qs that suggest a Q in the region of 500 to 1000 in the passenger cabin and 200 to 500 in the cockpit and avionics bay. Why there is such a big disparity is not clear; it will partly be due to an inaccurate volume calculation since the aircraft is not made up of perfect cuboids of one material as we find in the lab. Further, we must remember that the Q and time decay are equivalent for a single peak but this is not necessarily the case when we have many overlapping peaks. The stirrer average of the peaks will match the decay constant but only if we have a lot of samples, if there are not many stirrer positions then we might expect differences when we have a field varying by 20dBs from one position to another. At the very least we can ascertain from this that we should definitely look at environments with Qs under 1000 when calculating how wireless systems might perform when fitted to existing aircraft.

A final point from the NSWC papers is that they find the shielding effectiveness of the bay to be in the region of 30dBs plus and the cockpit and cabin to be lower, in the region of 20dBs. The issue of shielding effectiveness of an aircraft will not be considered. Nevertheless it is of interest if one wants to know the risk of jamming of internal avionics systems by an external signal, or what signal leakage there will be out of the aircraft.

### 3.4.3 Measuring communications systems in reverberation chambers

We have already described how we plan to use a reverberation chamber to analyse the performance of communications systems. The reason that this is being done is because a reverberation chamber can be made to resemble the inside of an aircraft. In fact both before the

start of the Flaviir project and during it there have been a number of publications on the subject of communications measurements in reverberation chambers but using the chamber for a completely different reason: to simulate other communications environments. This is because the multi-path environment that they provide can be used to model the environment faced by mobile phones and any other system where there are reflections. To be more precise we can say that one can replicate Rayleigh and Ricean channels in a reverberation chamber. Recent multiple input multiple output (MIMO) antenna technology actually makes use of the difference paths to increase the data rate and reverberation chambers offer a way to test the full system from the antenna to the signal processing with the added benefit of being able to obtain performance statistics via the stirrer.

At EMC Europe 2006 Carlsson and Kildal presented a summary of work characterising wireless devices in the period 2000-2005 and cited examples such as:

- Antenna efficiency measurements
- Measuring total radiated power from phones
- Antenna diversity
- Capacity of MIMO systems

More recently in the last 2 years they have looked at very similar issues to those presented in this PhD Thesis, examining the relationship between delay spread and coherence bandwidth for different reverberation chamber loadings [55] finding the obvious point that coherence bandwidth will increase as the loading is increased but also showing that the theoretical relationships between delay spread and coherence bandwidth appear to hold.

Chris Holloway and others from NIST have also taken an interest in this area in the last 5 years. At EMC Europe 2006 they presented a paper on the use of reverberation chambers to simulate multi-path environments. They show here that it is possible to set up different K-factors using a reverberation chamber through different antenna separations and angles. A K-factor is the ratio of direct to indirect path and was mentioned earlier because it is important in determining whether a signal has a Ricean or Rayleigh distribution. They have since progressed to measurements similar to those to be described here. In 2008 Holloway published at EMC Europe on the subject of using a reverberation chamber to simulate different power delay profiles demonstrating what we also show which is that it is easy to control the delay

spread with a simple setup of a reverberation chamber and EM wave absorbing foam. They find that with no or very little absorber there is a frequency dependence on the delay profile because the wall loss dominates losses although this is not noticeable for the small bandwidths that are presented here. In 2010 an updated journal paper was published on the same topic but also incorporating BER measurements and some continuously moving stirrer measurements. These are exactly the types of measurements that we have made in York to characterise an aircraft propagation channel. They found irreducible BERs for the empty chamber case of around 8% regardless of signal to noise ratio for a binary PSK system operating below a 1Mbps symbol rate. Changing the signal to noise ratio made no difference suggesting inter-symbol interference was the main issue which matches results that will be presented later here. They also found that a continuously moving system will present a higher error rate although they are uncertain of the cause and doing further work in this area.

Sorrentino and Ferrara have also looked at the use of reverberation chambers for wireless system measurements recently publishing on the simulation of a time moving channel using a reverberation chamber, and in particular showing how one can control the coherence time [56].

There has also been research into the performance of wireless systems in a resonant environment for the purpose of learning about performance in cars and similar vehicles. Flint, Ruddle and May looked at the received signal strength indication (RSSI) for ZigBee in a number of antenna locations in the body shell of a car and found that "standing wave patterns that were set up were very sensitive to position" and that changes in position of only 10mm could make dramatic differences.

Delangre and Lienard have published a number of papers in this area such as [57] which compares a car as an example of a confined environment with a reverberation chamber. This paper looks at the relationship the power delay profile and time constant and the Q-factor and explains that when the Q is low and the field is not uniform the normal relationships don't hold but can still provide an lower bound on the Q-factor. Their measurements of a car and a reverberation chamber loaded with absorber show that that the chamber can provide a very accurate model in terms of the power delay profile and its relationship with the Q-factor.

There are many other examples of similar work and we can say over the last 10 years the use of reverberation chambers for multi-path simulation has increased significantly and it's a widely used tool now. On the other hand the details of the environment and how it affects a

wireless system are less well understood. Consideration of the time response and it's impact is something that has only appeared in the literature at the same time as the research described in the upcoming chapters was ongoing. We certainly know that the multi-path environment in a high Q can lead to high BERs and also that we may accurately control the delay spread via the absorber in a reverberation chamber. Exactly why certain error rates occur is less well understood however and is a topic of ongoing research. In the chapters that follow we shall offer some further insights and results to explain how wireless systems behave in a reverberant environment with the aim of presenting a much more complete picture than we have from the small but interesting pockets of prior research.

## 3.5 Conclusions

The propagation environment has been discussed in this chapter from two perspectives: the electromagnetics perspective and the communications system perspective.

Firstly a channel model for an aircraft has been presented, which is based around coupled resonant boxes of differing Qs and sizes. It has been shown that electromagnetic waves in resonant boxes display a highly modal structure just as could be found for acoustic waves in a violin and this has many effects. In the worse case in the time domain there may be a long exponential decay of energy and in the frequency domain there may be frequency selective fading. In the better case fading will be flat but the signal to noise seen by a communications system will have an uncertainty associated with it, dependent on antenna position or any movement in the airframe. Where there is a long exponential decay that is comparable to the system symbol time it is impossible to say analytically what the PER might be, but one can say that we would not expect a reliable link. Where there is flat fading, it has been shown that the PER can be calculated and formulae for many different modulation schemes exist in the literature.

Electromagnetic and communications theory has been bridged using the mean delay metric which has been shown in this chapter to be directly related to the Q of a resonant mode in a cavity thus providing a useful way to take results found in communications literature and apply them to communications in resonant cavities. Coherence bandwidth has been presented as the frequency domain equivalent of mean delay and we have seen that it will be convenient to move between both time and frequency domains rather than doing all analysis in one or the other.

Returning to the channel model it has been seen that the electromagnetics imply a Rayleigh

or Ricean voltage or current distribution at the receive antenna which is exactly the same as that often found in multi-path communications. The mechanism is different in that rather than a few reflections with big path length differences, in a resonant cavity or reverberation chamber the difference arises from energy storage and any reflections over a short distance. It is interesting that these two very different scenarios can both lead to the same statistics of channel amplitude and similar concerns in terms of signal fading and high delay spreads. The issue will be discussed further in the following chapter but one of the most profound questions, given what we have seen in this chapter, is how do we certify a wireless aircraft. A shielded, properly terminate cable is easy to characterise and always looks the same to transmitter and receiver whereas a resonant wireless channel is the exact opposite. Ignoring noise and interferers, the received signal strength is unknown until the transmitter begins transmitting and can vary by tens of dBs as antennas move or the aircraft structure flexes or vibrates. Where there are many modes, those modes will have different Qs and thus different delay spreads. Our propagation environment is very unpredictable and if there is to be any chance to certify a wireless, passenger carrying aircraft (or any type) the environment must be described statistically so that we can work out chances of failure and design to minimise these.

A channel measurement technique has been presented for use in the future chapters whereby a network analyser will be used to measure the channel transfer response at the equivalent of many antenna positions using a reverberation chamber and a stirrer. This technique allows quick (in relative terms anyway) measurements of a channel, providing the complete statistical description of the channel that we require and it also makes it easy to produce time domain responses. With this technique we should be able to build up a picture of resonant channels at many different loadings, thus obtaining a large results set with which one could predict how an arbitrary communications system might perform in a variety of aircraft bays. The statistical reverberation chamber measurements provide some of the tools we need to meet the necessary reliability constraints.

Finally, we have observed that the areas of communications system testing in reverberation chambers, and of assessing resonant channel propagation are rapidly growing research areas. They have seen a number of papers published in parallel with the work done in Flaviir Wireless and generally indicate what we would expect, that a resonant environment makes for a challenging propagation environment.

# Chapter 4

# BER and channel measurements in a lab environment

## 4.1 Introduction

In the last two chapters the basics of wireless systems were described and then the theory of reverberation chambers. With this the two can be combined and analysed for how communications systems might perform in a reverberant environment. We present these results with two aims:

1. To characterise resonant environments from a communications perspective.
2. To show how a real system performs and to link performance to the channel and find out if resonant environments as challenging as we might expect for standard communications systems

The results presented in this chapter will demonstrate the main contributing factors affecting performance of a wireless aircraft. They will show what types of environment cause difficulties, which are benign and what the dominating propagation effects are.

### 4.1.1 Types of measurements

The aims above require an understanding of both of the electromagnetic and wireless systems theory and require that electromagnetic aspects of the communications channel are correlated with measured results. To do this two type of measurements are necessary:

- Electromagnetic measurements of the communications channel
- Measurements of radio performance, for example packet error rate or bit error rate tests

With the former what will be done is similar to channel sounding in the design of a mobile phone system or wireless LAN system. Just as those systems require design data so does a wireless aircraft system. In fact it is even more important in this situation as the application is safety critical and the environment is not well understood.

For the error rate tests a ZigBee system was used, because as explained in Chapter 2, ZigBee is one of the most suitable COTS technologies for sensor networks on wireless aircraft. Both PER and BER measurements were carried out because the former indicates real system performance and the latter offers information that can be more easily matched to communications system theory.

Key information expected to come out of the measurements was as follows:

- The form of the frequency response for different types of aircraft or vehicle and different bays within those vehicles
- The corresponding performance of a wireless system, such as ZigBee
- Q-factor of those environments
- Other metrics that can be taken from these environments that allow effective prediction of system performance or building of statistical or simulation models.

## 4.2 Measurement environment and use of reverberation chambers

It became clear early into the research phase that gaining access to aircraft for testing purposed would be difficult. Some compromises were found and these are presented in later chapters.

Significantly, even if highly realistic environments were available from the start then their use should not have been the focus of the measurement program. As explained in the introduction, the Flaviir program was a UAV program but the wireless project was to have a more general slant. UAVs themselves are a very new area, and how they will look in 20 years is very uncertain. Already they come in a variety of sizes from the american Predator to the much smaller UAVs that might be used for roles now occupied by police helicopters, for example. Further, if one looks to other types of aircraft or even vehicles, then there is a vast difference

between a passenger jet and a fighter jet in terms of size, in terms of materials inside, and in terms of internal structure and layout. One has a very large, open fuselage section, another has a a smaller narrow fuselage with a set of closed avionics bays. The implication of this is that a measurement program should be quite general, considering a wide range of potential environments.

The other important point which was conveyed in section 3 is that knowledge of the statistics of a resonant environment is crucial. For avionics systems total link failure cannot be tolerated because it could cause the loss of an aircraft. If one was to place a wireless system in a given aircraft and do measurements, both propagation and radio then the conclusion could be that such a system worked well and it should be installed in all models of that aircraft. Such conclusions would be unlikely to pass certification however because they don't address the statistical chance of failure. A system could be functioning on the runway but then under the force of takeoff, some cables or other equipment could move and the system would drop out.

To deal with these issues, there are two main requirements.

- A variety of structures must be considered in the lab
- All measurements must give statistical results where possible

In order to meet these requirements the reverberation (also know as mode stirred) chamber is the best option. It is already a highly resonant environment and the level of reverberation or Q-factor (see Chapter 3) can be controlled by loading it with absorber. This allows the simulation of different environments. It also allows the statistics of the channel to be found with minimal effort. In principle one could move the antennas many times to obtain the statistics but this would be very time consuming and difficult to set up whereas the reverberation chamber automates the process. Moving the stirrer moves the modes so that at a static position a maxima in space can become a minima in the same way as this would occur by moving the antenna itself.

Pictures of the 2 chambers that were used are presented in figures 4.1 and 4.2. Their dimensions are also presented in table 4.1. Other boxes and tubes were also used although these had the disadvantage that they could not be mechanically stirred and so only frequency stirring was available. These will be described as they are introduced. One can learn a lot by using just one reverberation chamber and changing its load. However there are a few characteristics of resonant cavities that change with size which may or may not make a difference

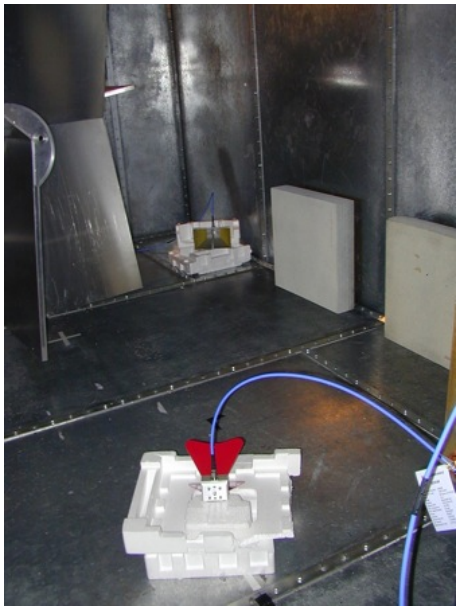

Figure 4.1: Reverberation chamber with 2 horns in a direct path configuration

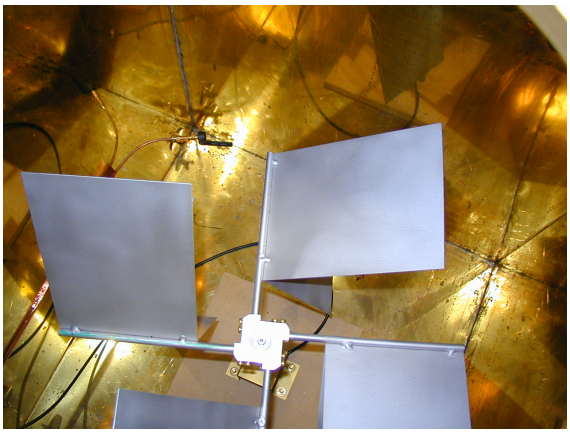

Figure 4.2: Small reverberation chamber

and so it was important to test more than one chamber.

The differences may be described by the following parameters:

**Cut off frequency** This is the lowest frequency at which a wave will be propagated. A larger cavity will support a longer wavelength or lower frequency frequency.

**Mode density** For a given frequency band the mode density will be higher for a larger chamber leading to narrower frequency and autocorrelation peaks.

**Maximum Q** The Q will increase with chamber volume for a equivalent loading

| Chamber | x(m) | y(m) | z(m) |
|---|---|---|---|
| Main | 4.7 | 3.0 | 2.37 |
| Small | 0.6 | 0.7 | 0.8 |

Table 4.1: Dimensions of the large and small reverberation chamber at York

**Coherence bandwidth** A larger chamber will have a smaller coherence bandwidth (see Chapter 2) because higher Qs give narrower bandwidths per mode which leads to more pronounced maximas and minimas across a frequency band.

**Delay** In the time domain the larger chamber will result in a signal with greater delay spread and thus more opportunity for inter-symbol interference.

## 4.3 Experimental setup

To measure the propagation a Vector Network Analyser was used which provides a frequency domain view of the channel through the S21 and S12 parameters. These specify the response over a given frequency range between two ports and provide both amplitude and phase information. This was connected to the chamber as shown in figure 4.3. Inside the chamber 2 antennas were connected to the chamber walls via coaxial cable. Typically 2.4GHz dipole antennas were used but in some incidences these were swapped for horns such as when broadband measurements were required. In all cases the antenna efficiency was taken into account when performing channel calculations. The use of the 2.4GHz dipoles should be assumed unless otherwise specified. The network analyser was calibrated using a full 2 port calibration before all measurements and all cables were terminated with either N-type or SMA connectors.

The setup for ZigBee measurements was the same except that in place of the Network Analyser, two Jennic boards were used such as those shown in figure 4.4 It is important to note that whilst the network analyser was calibrated , no such operation was possible for the radio system and losses of approximately 0.5 to a worst case of 5 dBs were possible dependent on the length of cables required. Therefore the power put into the antennas should not be regarded as just the board transmit power. This point is only significant if considering absolute signal and noise power which is not the case in the measurements presented.

For whichever reverberation chamber was being used the Q-factor was controlled by loading the chamber with AN79 foam which absorbs RF energy. With this absorber the chamber

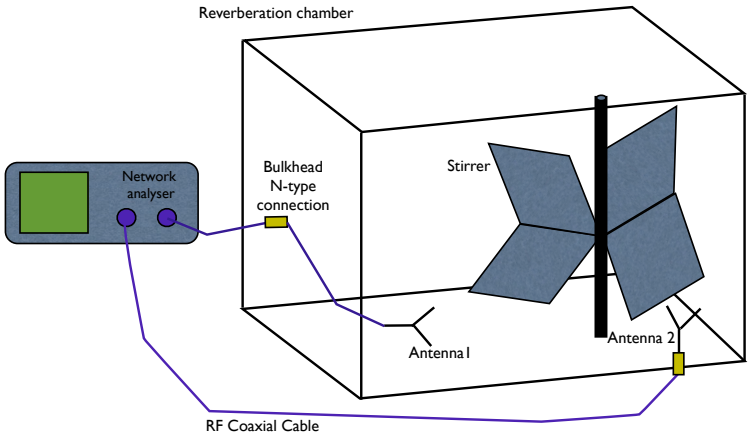


Figure 4.3: Basic reverberation chamber setup with network analyser and two antennas

Figure 4.4: A Jennic ZigBee unit as used for the first measurements

can, over 20MHz for example, be varied from flat to frequency selective and the Q could be varied from 30000 to 300. The standard blocks were of dimensions of 61cm x 61cm and had thickness 11cm. In the main chamber these were used either in this form or in 61cm x 30cm or 60cm x 15cm and placed flat on the floor or against walls. Typically the orientation was not found to affect the results due to the theoretical properties of reverberant cavities when fields are averaged with a stirrer. Therefore the Q is always reported but not the precise amount of absorber and its positions. To recreate the measurements provided in this and later chapters the reader is advised to use chamber dimensions as close as possible to those throughout this thesis and to add absorber until the same Qs are achieved.

### 4.3.1 Procedure

When carrying out a Network Analyser measurement in one of the reverberation chambers the process was as follows:

1. Prepare the measurement by setting the absorbing material and antenna positions and calibrating the network analyser.
2. Log the S parameters
3. Rotate the stirrer
4. Return to 2 until enough positions have been acquired

Typically when this was completed an equivalent measurement would be done using the ZigBee radios in the place of the NA before then changing the chamber setup. For later measurements a more complicated setup was used with an RF switch that allowed the radio measurement to be done with the exact same position as the network analyser measurement allowing for better correlation between the two. This is described in more detail later.

### 4.3.2 BER Measurement

To measure the packet and bit error rate a ZigBee kit from Jennic was used. This was used in two ways: with the Jennic provided PER test software and with custom software using the Jennic Production Test API and which was used will be made clear as needed. The major difference between these is that the former requires node association (see the 802.15.4 standard) and does not allow any BER information to be obtained. The latter does allow access to BER information and doesn't need association before data transmission so it is less realistic but more useful in terms of the data that can be logged. In terms of the ZigBee stack only the physical and medium access control layers were used because the aim was to try to obtain point to point link quality over the channel, not maximum capacity in topologies with large numbers of nodes.

Jennic ZigBee modules were the only ones used throughout the experiments. There is a danger with only using one vendor's technology, in that when issues arise they may be due to a peculiarity in the implementation of that specific card and not a general issue. However, by sticking to one vendor performance can be fairly compared over a variety of propagation channels and it saves significantly on time for measurements and programming. Tests were carried out to make sure that the radio performed as expected in a noise free and noisy environment alongside the resonant environment measurements.

## 4.4 Propagation through a metal cylinder

Rather than beginning with reverberation chamber measurements this section begins by looking briefly at propagation through a metal cylinder rather than a reverberation chamber as this can be described analytically and is simple enough to show the main characteristics of the communications channel. The tube is seen in figure 4.5 and inside the tube are monopole antennas. The frequency and time response for this can be seen in figure 4.6.

Figure 4.5: The metal cyclinder used to demonstrate waveguide effects

If we first consider the frequency response then we see that there is no propagation up until around 2GHz. In fact it is easy to calculate the exact frequency at which this happens using the solutions to the wave equation found in Balanis [58] who shows that the dominant mode of a cylindrical waveguide is at 1.954GHz for the particular radius of this cylinder. Up until this frequency there is no propagation of energy/power and so all that is seen is the analyser noise floor. After 1.95GHz there is wave propagation and as the frequency increases it goes from having some identifiable modes / peaks to a situation where there is a very unpredictable structure to the channel.

There are a few implications that this response may have for wireless communications. Firstly, there is a limit put on any system in terms of where it may operate, and this is deter-

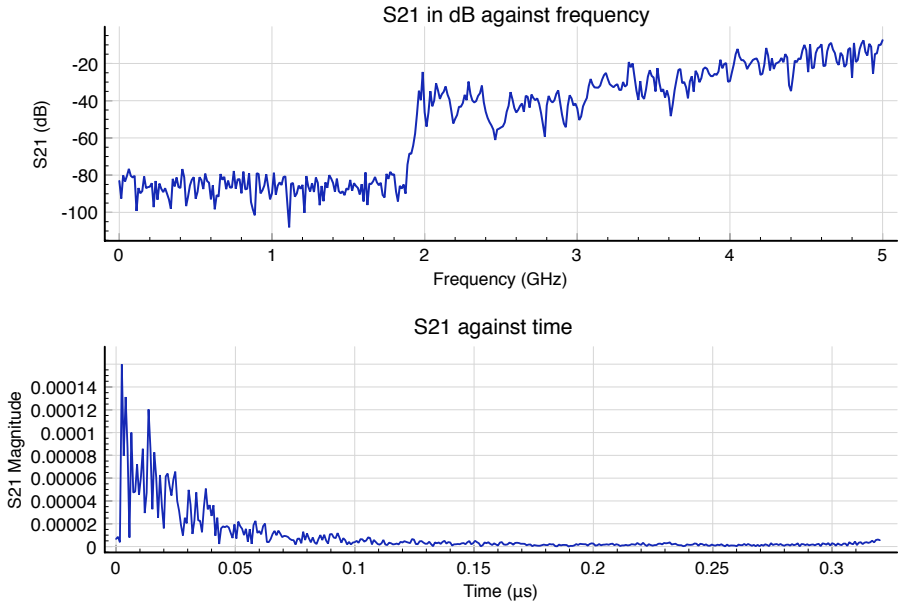


Figure 4.6: Frequency and time response of a metal cylinder

mined by aircraft geometry. The radius of the cylinder (4.5cm) is quite small and it is unlikely that anything so low would be encountered in a modern commercial aircraft. However, one could imagine a fuel tank in a UAV (or a space just behind a fuel tank) where radios may be required to communicate down a narrow tube. This puts restrictions on the radios in terms of the bottom of the frequency bands that could work. Fortunately fly-by-wireless systems are likely to be operating above 2GHz because that is where current standards work and because higher bandwidths may be available for a given centre frequency.

Considering the response after 2GHz then the next problem can be seen: the channel is extremely variable and the signal power may vary by nearly 40dB, i.e. 4 orders of magnitude, in a 1GHz span. This is not insignificant both in terms of what it means for the signal to noise ratio and also in terms of signal distortion. It may be that this variation exists within the bandwidth of an 802.11 channel and there is frequency selective fading. To see this more clearly we can look at a zoomed in part of the response as shown in figure 4.7 which demonstrates a response over 160MHz. There are two peaks in the response and if a WLAN signal, for example was centred over either of these the system performance would be hard to predict due to signal distortion. The remainder of the response on the other hand is relatively flat within the range of a typical wireless signal (which for these purposes is defined as anything from 2 to 25MHz). In the time domain the effect of a peaky/high variance frequency response can be seen and

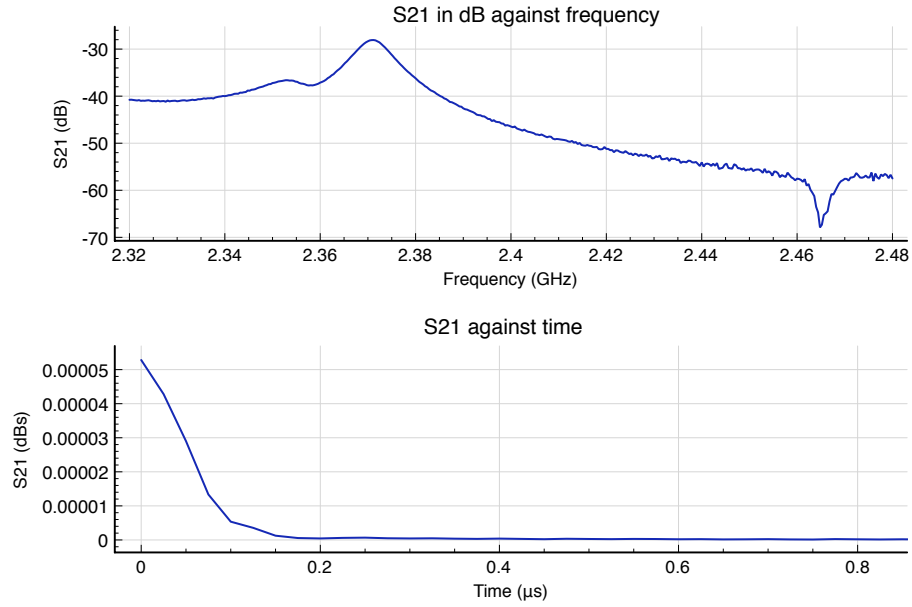


Figure 4.7: Frequency and time response of a metal cylinder

that is the none instantaneous, decaying impulse response. As explained in Chapter 3, the more greater the variation in the frequency response, the longer the exponential decay time leading to potential inter-symbol interference.

To summarise, this simple example shows that when there is a reverberant environment we must always be aware of the waveguide cut off, the variation in signal power and the potential for inter-symbol interference.

## 4.5 Results - Radio PER performance in highly resonant environments

### 4.5.1 A first look at the reverberation chamber channel

The reverberation chamber introduced in section 3.2.1 is a more complicated situation. For the cylinder one could have obtained the measured response by being very precise in the geometry, applying the standard formulae and adding together all the modes. In the reverberation chamber this is not possible as the stirrer paddles in particular complicate the geometry disallowing analytical solutions. It is possible to know the cut off, but how the chamber will look to a pair of antennas is very dependent on their exact position and that of the paddle.

To get a basic understanding of its behaviour, first the chamber's response is observed when it isn't loaded (i.e. no deliberate absorbing surfaces). This demonstrates the worst case

and a situation that might arise in a sealed environment which could arise in aircraft fuel tanks or in other vehicles such as large bulkhead separated sections of ships. Following this additional absorber is added. 4 measurements sets were carried out corresponding to 4 different loadings of AN79: an unloaded chamber, a chamber with 1 small piece of AN79, a chamber with 1 large piece and a chamber with 2 large pieces. To be clear about what this means, in 4.1 the piece shown is classed as large and a piece of half this width is classed as small. It will be seen that not many pieces are needed to being the Q-factor down by an order of magnitude.

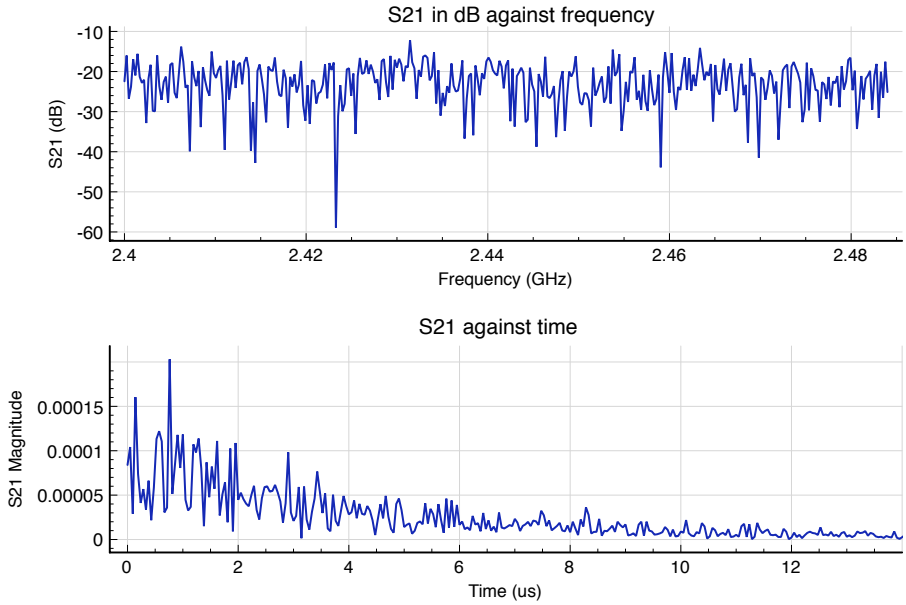


Figure 4.8: Frequency time graphs for chamber with no absorber pieces

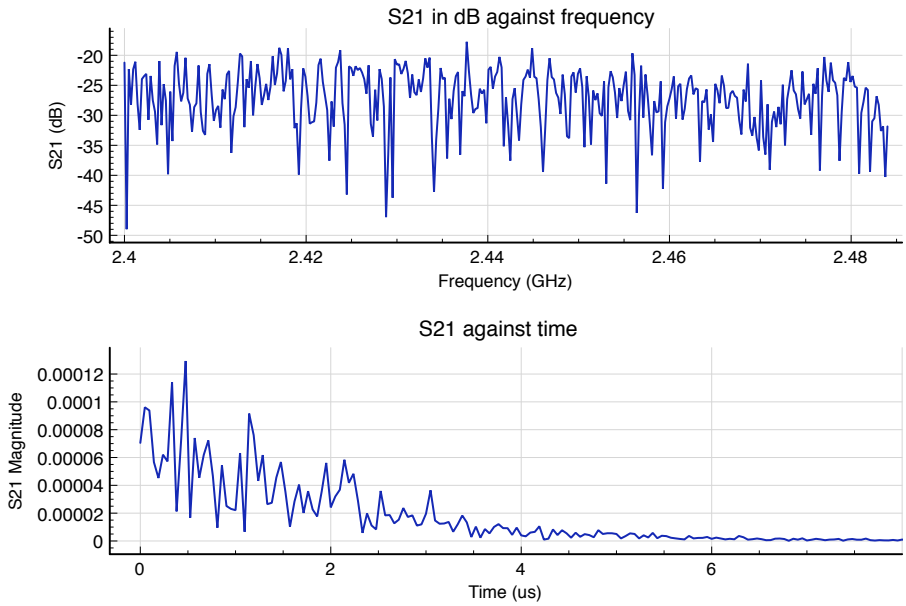


Figure 4.9: Frequency time graphs for chamber with 1 small absorber piece

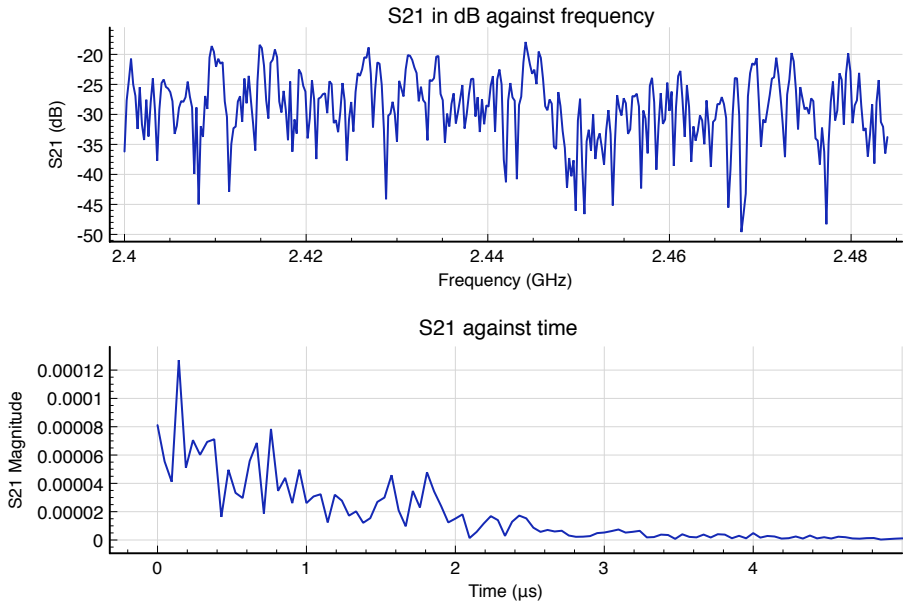


Figure 4.10: Frequency time graph for chamber with 1 large absorber piece

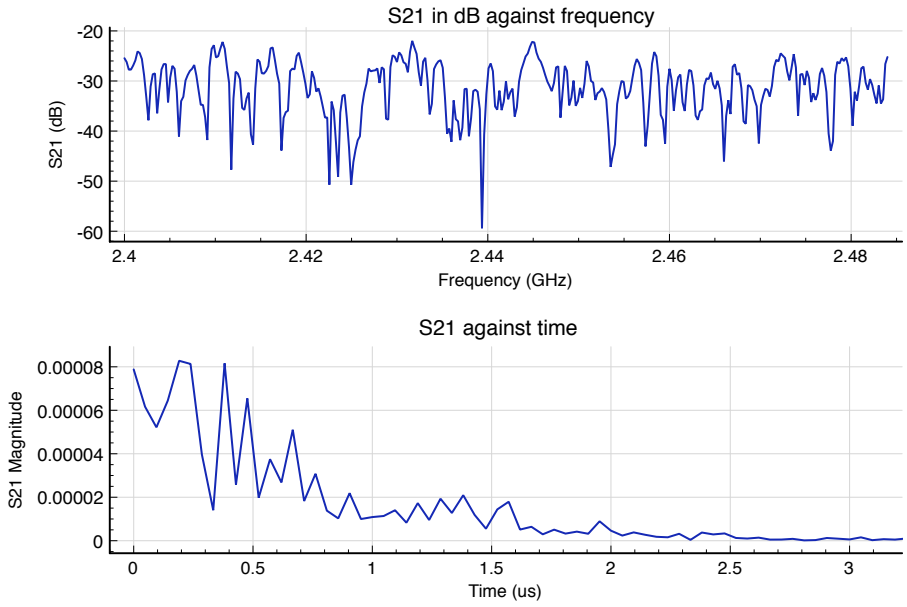


Figure 4.11: Frequency time graph for chamber with 2 large absorber pieces

Time and frequency responses are shown in figures 4.8 to 4.11. These graphs demonstrate that when unloaded the chamber is extremely chaotic and it is impossible to identify individual peaks at this resolution of 0.06MHz between points. Signal power is varying regularly by 25dBs and sometimes by even more. One very sharp null can be seen at nearly 50dBs below the maximum value. Correspondingly, in the time domain the response is not just oscillatory but slow, in that the signal has got to 1/10th the voltage level in 10$\mu$s. This may be compared to the maximum bit rate on an IEEE 802.11 LAN which is 11Mbps with symbol time 0.09$\mu$s. When the decay time of the order of or longer than the symbol time, system performance will be very unpredictable and reliable communications cannot be achieved. Various reasons for this are identified in Chapter 3 but the key point is that the amplitude at the receiver depends not on the current symbol but on the last one, in fact on many previous symbols in this case.

Looking at the other 3 graphs we see a pattern of the channel becoming flatter relative the the empty case but the maximum to minimum deviation remaining similar. If we considered the auto correlation of the frequency channel, as we shall later, then its peak will widen as absorber is added. Note also that the maximum and average values have been reduced which is expected due to the absorber damping the resonances and absorbing energy. On the time domain plots the decay time approximately halves from one plot to the next as the absorber is doubled. This is not of significance except in the sense that it shows only a very small surface area of absorber having a large effect on the chamber resonances. One can imagine an aircraft with a large number of apertures or wiring looms having a low Q-factor.

### 4.5.2 PER and Q-factor

Now that the channel looks like at high Qs has been seen, both in the frequency and time domains, next the topic of how these Qs might affect a wireless system is introduced.

| Absorber (blocks) | Q | PER (%) with ACKs | PER without ACKs |
|---|---|---|---|
| 0 | 32000 | 100.0 | 100.0 |
| 0.5 | 10000 | 97.9 | 100.0 |
| 1 | 7000 | 0 | 0.1 |
| 2 | 4000 | 0 | 0.1 |

Table 4.2: Table of ZigBee PER for a single stirrer position

Table 4.2 shows the performance of a Jennic ZigBee system in a static environment, i.e.

the stirrer and the antennas are in one specific orientation. The Q-factors quoted are for the specific stirrer position at which the BER is recorded. Variation of the Q is shown in the later graphs. In each case 40000 packets were sent between two radios, firstly with acknowledgements (ACKs) and retransmissions turned on and then without. The surprising aspect of the results is that there is no smooth drop off in the performance of the radio. When the Q is 10000 then the system is unusable with 100% PER being recorded but when the Q is 7000 the error rate (allowing for retransmissions) is 0. An interesting question is raised here as to why this suddenly changes.

The results show that caution is needed in assessing system performance. If these results had been obtained by putting antennas in two different aircraft then it'd be possible to conclude that a system won't function in a high Q aircraft (calling a high Q over 10000, medium between 1000 and 10000) but will work in a medium Q aircraft. However it was seen earlier that the channel varied significantly with frequency and it should do the same with stirrer position. Therefore one cannot assume that these results are typical but must consider what happens with antennas in different positions or the stirrer in different positions.

| Absorber | PER (%) with ACKs | PER without ACKs |
|---|---|---|
| 0 | 100.0 | 100.0 |
| 0.5 | 19.4 | 85.8 |
| 1 | 16.8 | 9.3 |
| 2 | 2.2 | 7.9 |

Table 4.3: Table of ZigBee PER for a moving stirrer

Table 4.3 offers a first look at this issue. In this table the results are for a moving stirrer where the rotation rate is 0.03Hz and so sufficiently slow not to affect the channel in terms of frequency modulation. The BERs are much lower here than in the static single position for the measurements with only a small absorber piece but the BERs are higher for measurements with more absorber. Particularly noteworthy is the transformation from 19.4% with ACKs to 85.8 without, especially when compared to the static case. With movement, by the time of a retransmission the channel could have transitioned from a minima to maxima in the signal. In all cases the change in PER indicates that there is variation with stirrer position and therefore antenna position. The case where the moving stirrer PER is higher than the static case can be explained because there is a chance of finding in a single sample, a position with very low PER

relative to the average over stirrer positions. The opposite applies where the particular position is one of a very high error rate. What is seen here is that a wireless system is not robust in a reverberant environment in the sense that performance will fluctuate and is very dependent on antenna position in the environment.

Similar experiments (with a slowly turning stirrer) were done for the small reverberation chamber and are presented in tables 4.4 and 4.5. The volume here is much lower and so the Q for the unloaded setup is an order of magnitude smaller than the large chamber at just under 1000. Note that these Qs are averages rather than for a single position. BER was consistently higher for the more highly damped case, this may be because of the additional attenuation or may be because of the different modal structure and the statistical nature of our results, being within the error bounds.

| Absorber | Q (approx) | PER with ACKs | PER without ACKs |
|---|---|---|---|
| None | 900 | 0.2% | 0.6% |
| 1 piece AN79 | 550 | 0.1% | 0.1% |
| 3 pieces AN79 | 370 | 0.2% | 0.8% |

Table 4.4: Packet error rates (PER) for Jennic ZigBee kit with stirrer in small chamber with varying Q-factor, first run

| Absorber | Q (approx) | PER with ACKs | PER without ACKs |
|---|---|---|---|
| None | 900 | 0.3% | 1.2% |
| 1 piece AN79 | 550 | 0.1% | 0.2% |
| 3 pieces AN79 | 370 | 0.2% | 0.2% |

Table 4.5: Packet error rates for Jennic ZigBee kit with moving stirrer in small chamber with varying Q-factor, second run.

In these small chamber results the average PER is very much lower than for the larger which is encouraging for the potential performance of ZigBee in a UAV where the bays are significantly smaller than the York large chamber. However, with average PERs of over 0.2% performance could not be considered good compared to a wired system, especially considering this is a noise free environment.

On carrying out these experiments it was noticed that for the moving stirrer case the BER tended to shoot up rapidly for half a second then settle again. In other words the average

was not representative of the performance at any given stirrer position. Either the system was working well with no errors or badly with a lot of errors but rarely anything in between. Understanding this behaviour has influenced a lot of the work that follows.

### 4.5.3 Further investigation of radio PER in high Q-factor environments

So far there are two key observations:

- There is a very sharp change between a system working and not working as absorber is added
- There is a very sharp change between a system working and not working as the stirrer moves.

Further information is required to better understand this and so additional measurements are presented in the high Q-factor range that provide further insight. For these a similar setup was used to before where 6 different combinations of absorber blocks were used to control the Q from an average over stirrer positions of 28000 to 3800.

The results are shown in figure 4.6. There is a lot of information in this table but it has been presented like this in order to show a few of the main characteristics of the data. What the table is showing is the packet error rate for different chamber loadings, with each loading a column in the table. The first ten rows are the recorded PER for 10 independent stirrer positions (i.e. the movements are large to ensure independence of signal amplitudes). In the first case we see that all ten positions are unusable in the sense that no packets can be successfully transmitted. Then as the Q halves there are a couple of positions with some successful packet transmissions and once the Q is at a 10th of its maximum then 8 of the ten positions have no packet errors at all. If we look at the average then even the best case in terms of errors we see that the average PER is 10% which would be far too high for any system that needed to utilise the link more than occasionally. However, in this case it is necessary to be aware that an average over stirrer positions does not necessarily apply. In fact there are two cases, in an aircraft, one where wings are flexing and hydraulics are moving and the channel is changing in flight, and the other where the signal in an avionics bay is constant. For the former case the average PER is what a real system might see. For the latter, it is misleading in that the link is likely to be either 100% working or 100% not. In order to understand the implications of the measurements, throughout this chapter both the average PER and the number and percentage of positions

within a certain error range will be referred to.

| Stir position | No absorb | 1 quarter | 1 half piece | 1 half , 1 quarter | 2 halves, 1 quarter | 1 l, 2 h, 1 q |
|---|---|---|---|---|---|---|
| 1 | 100.00% | 0.00% | 100.00% | 82.20% | 0.00% | 0.00% |
| 2 | 100.00% | 0.00% | 26.00% | 2.40% | 0.00% | 0.10% |
| 3 | 100.00% | 100.00% | 19.80% | 100.00% | 0.00% | 0.00% |
| 4 | 100.00% | 22.80% | 100.00% | 0.00% | 0.00% | 0.00% |
| 5 | 100.00% | 0.00% | 0.00% | 100.00% | 0.00% | 0.00% |
| 6 | 100.00% | 100.00% | 100.00% | 0.00% | 0.10% | 0.00% |
| 7 | 100.00% | 100.00% | 0.00% | 59.90% | 100.00% | 0.00% |
| 8 | 100.00% | 100.00% | 100.00% | 0.00% | 0.00% | 0.00% |
| 9 | 100.00% | 100.00% | 0.00% | 33.40% | 100.00% | 100.00% |
| 10 | 100.00% | 100.00% | 100.00% | 0.00% | 0.00% | 0.00% |
| average | 100.0 | 62.28% | 54.58% | 37.79% | 20.01% | 10.01% |
| No. of 100%s | 10 | 6 | 5 | 2 | 2 | 1 |
| No of 0%s | 0 | 3 | 3 | 3 | 7 | 8 |
| Mean delay($\mu$s) | 15.5 | 6.75 | 4.64 | 4.2 | 3 | 2 |
| Delay spread($\mu$s) | 14.9 | 6.65 | 4.8 | 3.9 | 2.95 | 1.95 |
| Q-factor | 28800 | 12600 | 10000 | 7300 | 5500 | 3900 |

Table 4.6: Packet error rates for Jennic ZigBee kit with moving stirrer in small chamber with varying Q-factor, second run.

Figures 4.12 and 4.13 better illustrate this data showing the number of positions with PER > 1% and average PER, respectively. What they show is that when the channel is time varying then performance is unsatisfactory for all but the lowest data rate applications until the Q-factor reaches 4000. In the static case what we see a minimum of a 1 in 10 chance of a system not working. In both these situations, even for non critical applications, this is too high to be widely useable and so we can conclude that *Qs from a few thousand should be avoided for ZigBee based sensor networks.*

The other piece of information presented in table 4.6 is the delay which should be proportional to Q-factor. To check this we plot Q-factor against mean excess delay in figure 4.14. Delay spread was discussed in Chapter 3 as a parameter which gives an indication of the likelihood of inter-symbol interference and it is interesting to check if this corresponds linearly to Q (as the theory suggests) because if it does then we may then move between the two easily as required.

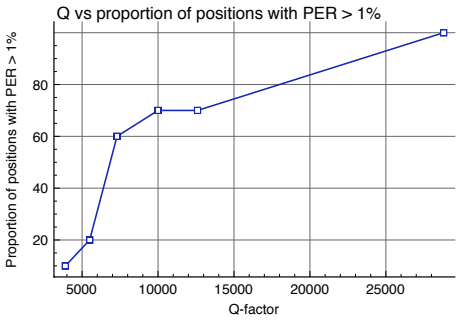


Figure 4.12: % of positions with PER > 1% against Q

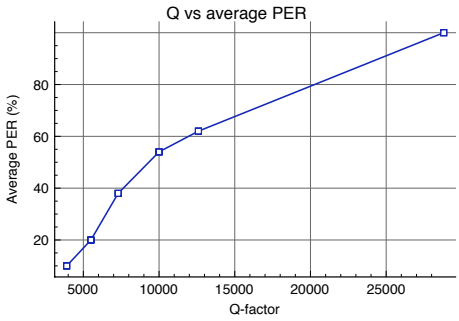


Figure 4.13: Average PER against Q

For ZigBee the bit rate is 250kbps and the chip rate is 2Mbps which means the bit period is 4$\mu$s and chip period is 0.5$\mu$s (although in principle there are two orthogonal offset chip sequences with period 1$\mu$s). The delay spreads presented in 4.14 are at best comparable with this bit period. Rappaport suggests[43] that a link without equalisation may be unreliable unless the delay spread is of the order of $0.2T_b$ or less where $T_b$ is the bit period so it is not surprise that the PER is high. In this situation exactly how high the PER will be is not calculable and requires measurements such as these or simulation, but we would not expect a low PER.

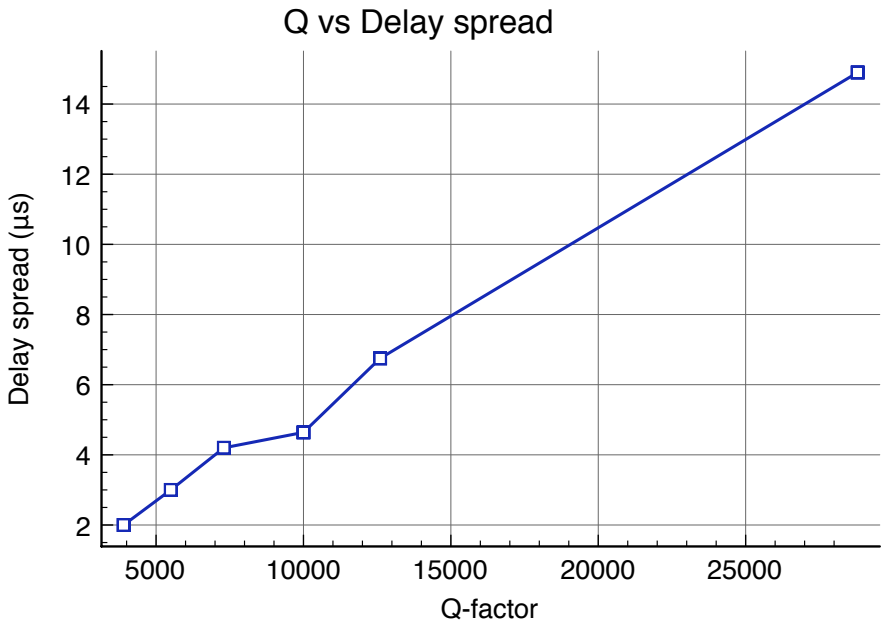


Figure 4.14: Mean excess delay vs Q-factor

**Variation of parameters with stirrer position**

Single delay spread and Q values were presented in the previous tables which correspond to averages over the 10 measured stirrer positions. Both of these (and by implication signal strength) will vary with position as well. Plotting S21 for two frequencies-loading pairs one can see significant differences in signal power from one position to another. This is plotted in figures 4.15 and 4.16 which correspond to the cases where there is no absorber and where there are 2 pieces of absorber in the chamber respectively. 64 independent stirrer positions are used over a complete rotation of the stirrer.

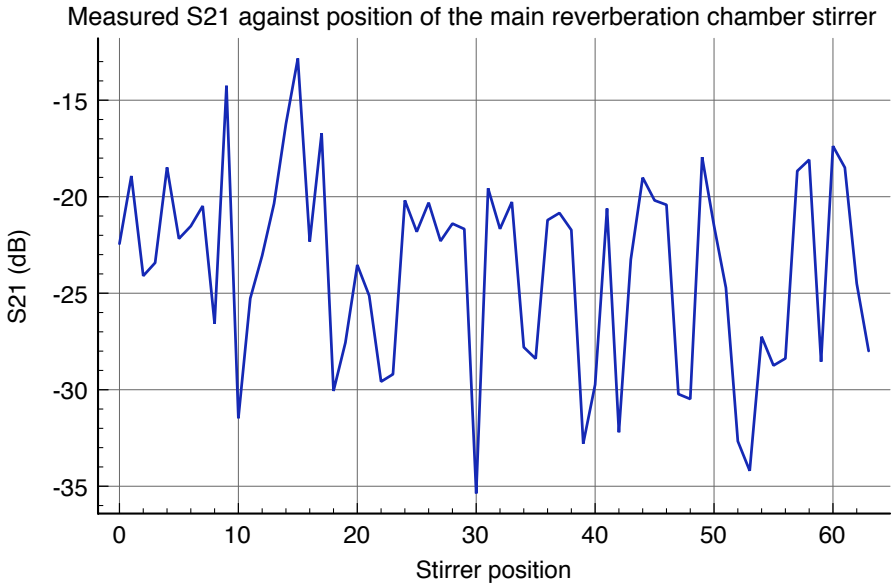


Figure 4.15: Variation of S21 at a single frequency over 64 independent stirrer positions with no absorber

We see a significant variation of up to 20dBs in both cases although there is a very slightly wider range in the higher Q/lower absorption setup. This is not an easy environment for a communications system to operate in. It can be seen that when combined with a weak signal or noise then a wireless system could alternate between working and not working dependent on antenna positions.

Q-factor is calculated from the S21 measurements and so this too can be expected to change significantly with stirrer position and frequency. Firstly we shall consider the variation of Q-factor with frequency. In figures 4.17 and 4.18 we have averaged over all stirrer positions (for each frequency) in order to obtain the Q-factor. For the higher Q case there is a range of 60000

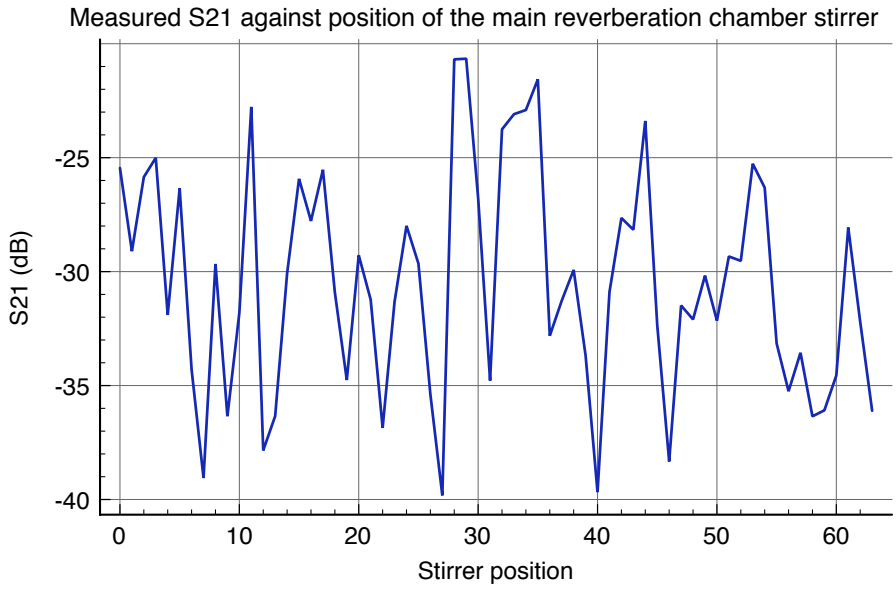


Figure 4.16: Variation of S21 at a single frequency over 64 independent stirrer positions with 2 pieces of absorber

to 25000 and for the lower Q case there is a range of 4000 to 11000 despite averaging over 60 positions and so the minimum Q is just under half the maximum in both cases. Clearly in comparison to the single frequency S21 measurements the range is far smaller but there is still a difference between frequencies. Certain modes are affected more than others by the absorber in the chamber, particularly when there are only one or two pieces of absorbent material. The wavelength is sufficiently small that some modes will have minimal interaction with the absorber and others may not.

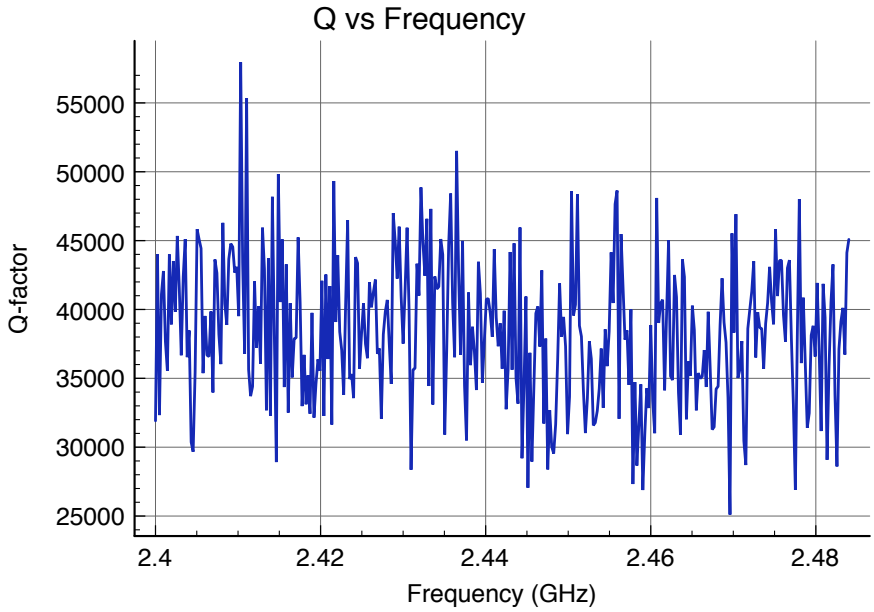


Figure 4.17: Variation of Q-factor with frequency, averaged over 64 independent stirrer positions for an empty chamber

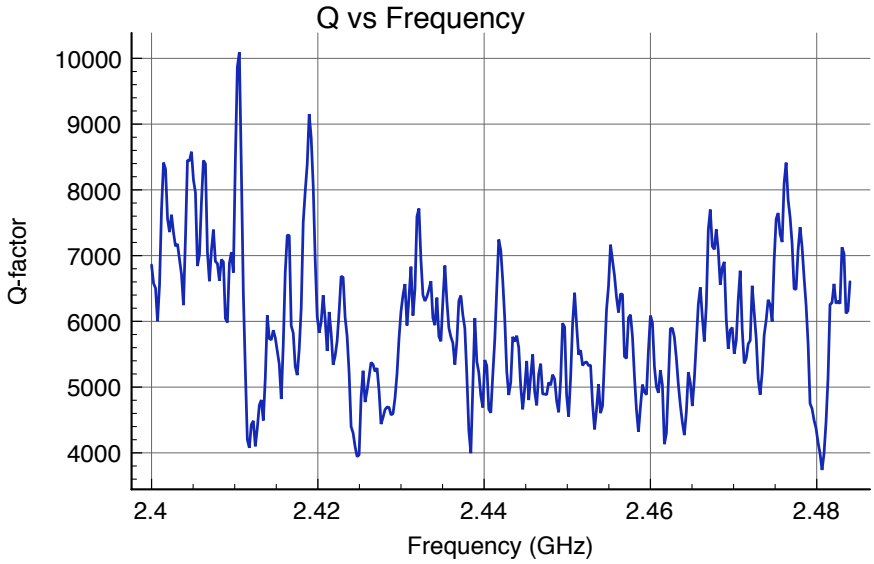


Figure 4.18: Variation of Q-factor with frequency, averaged over 64 independent stirrer positions for a loaded chamber

Variation of Q-factor with stirrer position is shown in figures 4.19 and 4.20. These are both

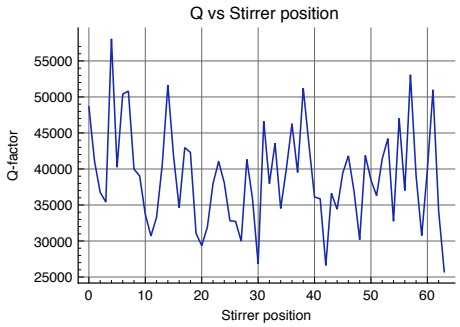


Figure 4.19: Variation of Q-factor over 64 independent stirrer positions, 5MHz bandwidth

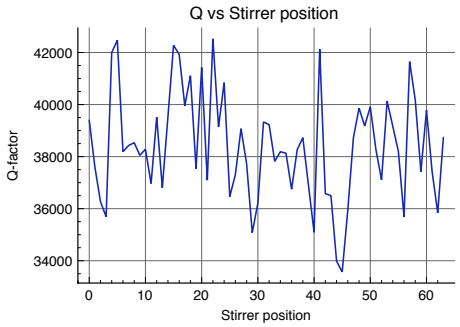


Figure 4.20: Variation of Q-factor over 64 independent stirrer positions, 80MHz bandwidth

for an unloaded chamber but represent two different bandwidths, 5MHz and 80MHz. The maximum to minimum range is 4 times greater in narrow band case compared to the wide band case where movements in Q have a range of 6000 as opposed to 25000. This demonstrates the advantage of a DSSS system, some of the effects of the reverberant channel are reduced by averaging over an increased bandwidth.

To conclude this subsection we present the the channel's stirrer variation another way: in terms of its delay spread. Variation of delay spread with stirrer position is plotted in figures 4.21 and 4.22 which show the delay spread for the case with no absorber and 2 pieces of absorber. These again demonstrate that exponential decay constant shows significant variation across stirrer positions with order of magnitude changes between positions. This could be the difference between a system havin 0% PER and 100%PER. As with the Q measurements, greater absorption results in not just a lower delay spread but also a smaller variance and difference between maximum and minimum levels.

### 4.5.4 The key discoveries so far

At this point we have introduced a number of important issues that may affect wireless systems in aircraft and similar resonant environments:

**Variation of signal power with frequency** This can vary by tens of dB and so certain bands may work whilst others don't. At high Qs there will be frequency selective fading.

**Variation of signal power with antenna/stirrer position** Antenna/stirrer position modulates the signal to noise and so the exact locations of the antennas and the current shape of the


Figure 4.21: Variation of delay spread over 60 independent stirrer positions, 80MHz bandwidth

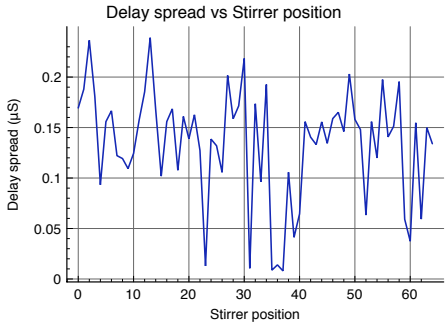


Figure 4.22: Variation of delay spread over 60 independent stirrer positions, 80MHz bandwidth

environment may result in a system that has a working signal to noise ratio or doesn't. We may only talk of the probability of a system working.

**Variation of delay spread and Q with frequency** Delay may cause inter-symbol interference in some bands and not others

**Variation of delay spread and Q with stirrer position** Delay may cause inter-symbol interference at some antenna positions but not others

From what has been shown so far, if the Q-factor is over 1000 there is no way to be certain of not ending up in a deep fade and a high packet error rate.

When looking at a set of error rates we saw that typically the PER is either 0 to 100%. There were few values in the middle, unless a moving stirrer was employed causing an averaging of the SNR and the PER. For the measurements in the last section the signal strength appeared too high for receiver noise to be an issue given there was little absorber and no noise sources. Therefore we might assume that another mechanism is the cause of the errors. This is likely to be inter-symbol interference. We have shown that delays spreads are typically of a similar or even higher magnitude than the bit period. It is quite possible that the change between 0 and 100% PER is due to moving between tolerable and intolerable levels of inter-symbol interference. This is the issue that will be considered next.

## 4.6 A more detailed look at the relationship between Q, delay, signal strength and error rates

### 4.6.1 Experimental setup changes

Given the unusual nature of the errors seen, a more sophisticated measurement setup is required in order to understand the cause of link failure. There is uncertainty about whether signal strength or ISI is dominant in causing errors. To deal with this we must record frequency responses for each individual PER measurement. Further, it would be beneficial to measure BER rather than PER so that one could relate back to the standard formulas found in the literature. To deal with the issue of correlating the channel at a given stirrer position with the BER, two RF switches were employed and incorporated into an automated switching system, that allowed for computer control. The setup was shown in figure 4.23 and essentially it allows a network analyser measurement to be carried out followed by a ZigBee radio measurement before then moving the stirrer and doing both measurements again.

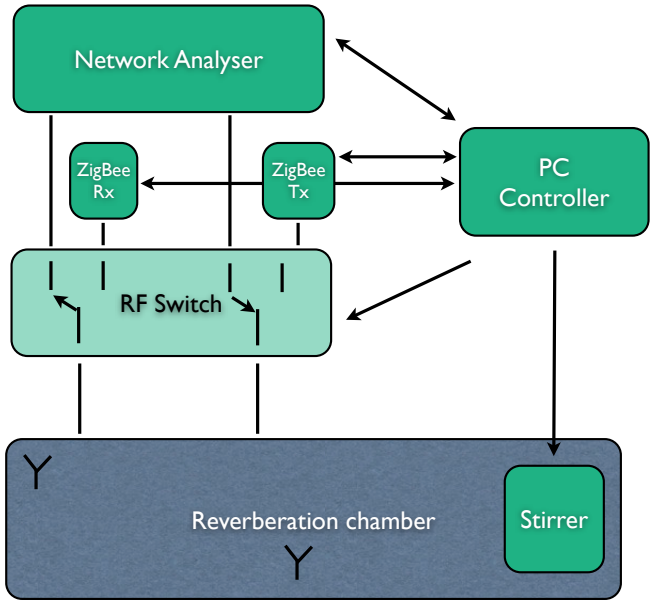


Figure 4.23: ZigBee measurement setup

For the PER measurements a standard Jennic supplied application was used that was triggered manually. Given what has been shown so far, we can see that the error rates are unpredictable, changing dramatically from one position to another. So that this can be better understood a large number of measurements will be needed and the whole system must be

automated. This was done by programming the Jennic modules, using their Production Test API to send and receive packets on the request of a signal from the pc. This has other advantages over the PER test because at no point does either node try and join any network; there is no association. Thus we can see the raw network error rates at the physical layer. When doing the PER tests there were occasions where the Jennic supplied software would not try to send anything if it could not first identify the node it wished to send packets to, meaning PER would be presented as 100% even if the rate would have been lower if the radio had actually transmitted packets. Although this is exactly as would happen in real life, it makes our analysis difficult and so it is preferable to carry out measurements with association turned off.

The final point to make is that often BERs of 100% are reported in the following sections which for a pure BER measurement would not make sense. However, this is still a packet based system and so if part of the header has errors, or synchronisation has failed after the preamble, or automatic gain control hasn't adapted then we shall not get a packet and the error will be reported as 100%.

### 4.6.2 Results

In the results that follow 2 ZigBee(ISM) band dipole antennas were used inside the reverberation chamber. Note that these were rather inefficient, worse than their data sheets suggested and so Qs must be taken as a guide and are not completely accurate. AN79 absorber was used as before to control the chamber loading. The number of stirrer steps between measurements was selected in such a way as to be independent and was adapted as required to ensure this. Per stirrer position 1000 packets were sent which is low for a PER measurement and would not be high enough to capture low bit error rates. However, Gaussian noise errors are well understood and the aim was not to see Gaussian noise but to capture the pattern seen before of alternating 0 or 100% PER, for which 1000 packets should be sufficient. Each packet contained 104 bytes of data. The absorber was used to vary the Q from approximately 40000 down to 1000 and for each stirrer position the S21 parameter was measured alongside the PER measurement.

Table 4.7 has the chamber setups, in terms of number of absorber blocks and associated average Q. Alongside these are plotted average BER results and their breakdown in terms of numbers of positions with BER greater than 1% and 50%. 100 stirrer positions were used for each of these measurements. The Qs and BERs in this table were calculated as averages over

the available stirrer positions and demonstrate the dramatic variation of the error rate as the Q drops. We see that the average BER goes from a high of 90% to around 1% as the Q drops off towards 1000. Significantly we also see the number of positions where the signal is not working go from 97 out of 100 to 1 or 2 out of 100.

| Absorber | Q | BER | No. BER > 0.1 | No. BER > 0.5 | No PER > 0.01 |
|---|---|---|---|---|---|
| 0 | 45773 | 0.908 | 97 | 97 | 100 |
| 1 | 10566 | 0.16 | 39 | 39 | 50 |
| 3 | 4646 | 0.0000425 | 0 | 0 | 0 |
| 5 | 2540 | 0.039 | 1 | 1 | 0 |
| 6 | 1776 | 0.0003 | 2 | 0 | 2 |
| 8 | 1388 | 0.0189 | 2 | 2 | 2 |
| 10 | 1345 | 0.000 | 0 | 0 | 0 |

Table 4.7: Bit error rates for ZigBee (1 is 100% errors) and their breakdown including no. of stirrer positions with BERs and PERs below a threshold

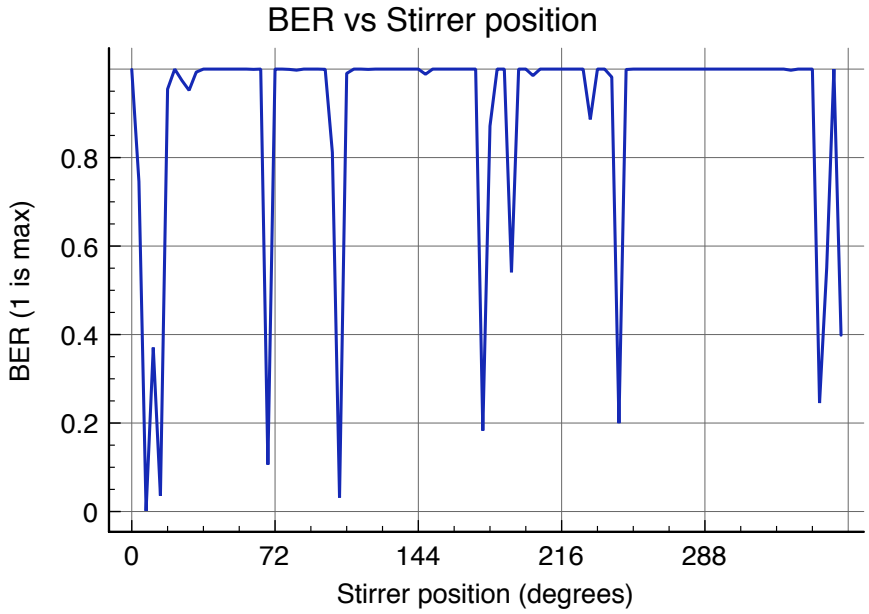


Figure 4.24: BER against stirrer position for Q of 45000

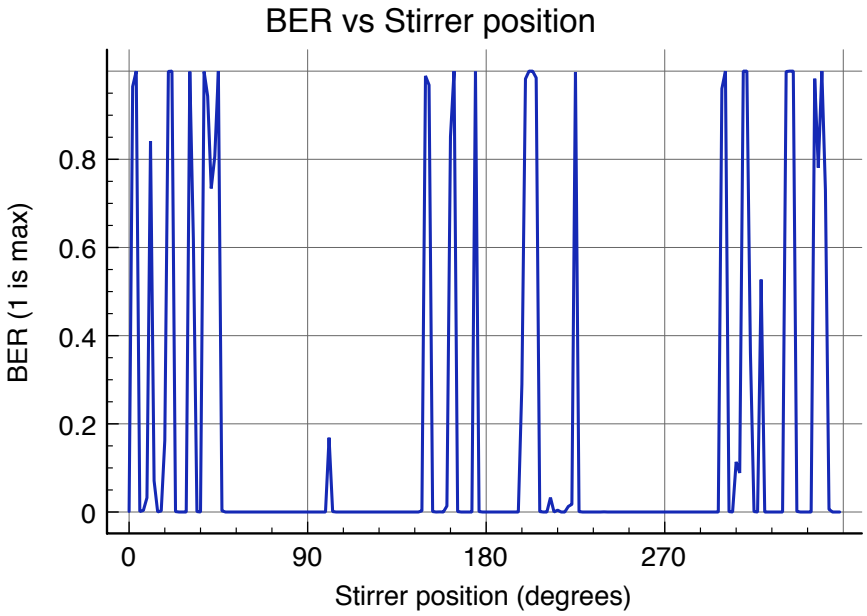


Figure 4.25: BER against stirrer position for Q of 10566

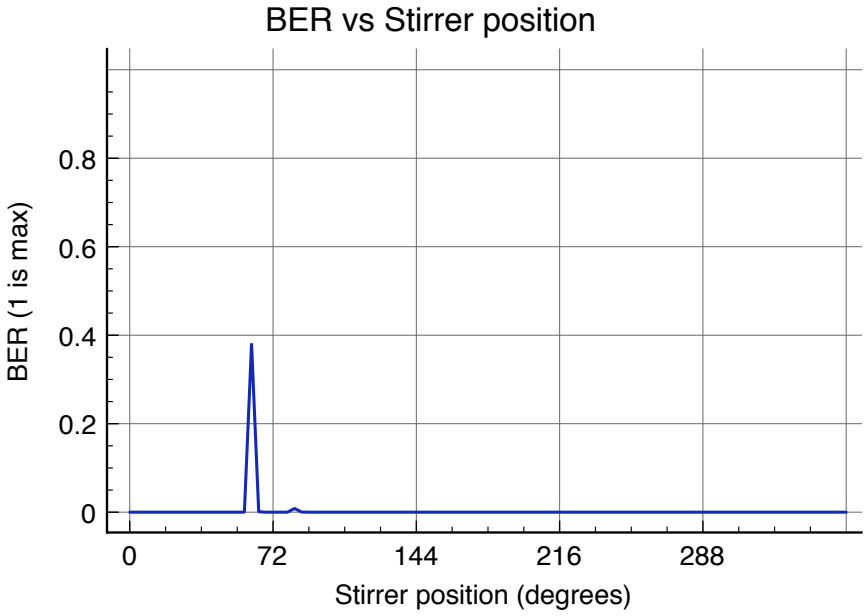


Figure 4.26: BER against stirrer position for Q of 2540

Due to the on-off pattern of the system it is useful to see the plot of BER against stirrer position and so figures 4.24 to 4.26 have BER plotted against stirrer position for three different Qs. A pattern is seen of BER being either 0 or 1 with only a few points in between. The system is either working well or not working at all dependent on the stirrer position, it is not a gradual change. Again we see a dramatic change from the high Q measurements to the lower ones.

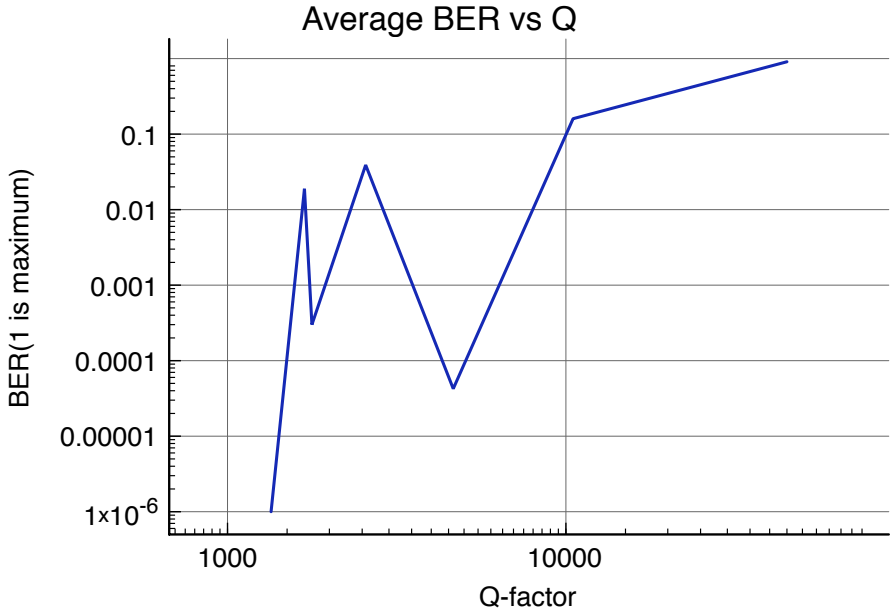


Figure 4.27: BER against Q-factor

Figure 4.27 plots the average BER and shows that whilst it drops dramatically with some absorber compared to none, after that it is quite variable and this is because over the 100 stirrer positions we are only finding 1 to 2 positions with 100% BER. In total there are 6400 positions and so with between 1 or 2 high error rate positions in 100 positions we could easily miss these in a particular measurement set. It would be impossible to consider all positions for all Qs due to measurement time but some single steps measurements were carried out. This pattern of a select few stirrer positions causing problems does mean we can't take the result for the Q of 1345 (0 BER) as certain. All we can say is that the BER in an environment of Q 1345 is likely to have less than 2 in 100 independent positions with BER greater than 0.5.

Figure 4.28 shows a measurement not included in the above table where the BER is 0.07 and the Q is 6000. This was a single step measurement done over all 6400 stirrer positions to demonstrate the complete pattern. It was found that there were 400 and 600 positions with a BER at or over 0.5 and 0.1 respectively. This indicates that even at this high Q it is possible

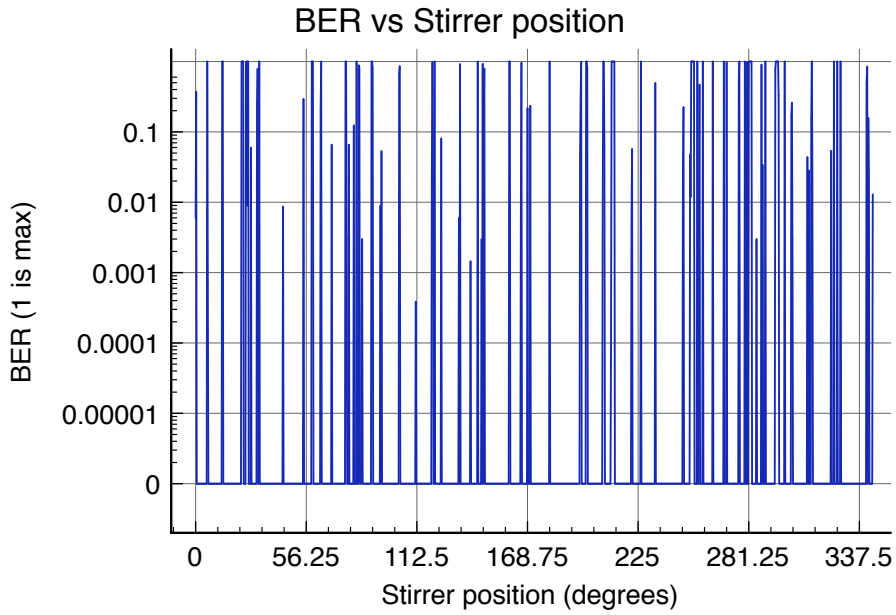


Figure 4.28: BER against stirrer position for Q of 6000 for all stirrer positions

we could theoretically pick up not a single position with a high BER despite doing 100 measurements and we must be aware that there is always an uncertainty associated with resonant environments. In order to deal with this problem some single step measurements was carried out to verify our results for independent stirrer positions.

Although we are presenting BER results we should not forget the implications for the PER. At the Q of 4000 there are no recorded positions with high error rates and every position has BER less than 1%. However, for ZigBee the packet size is 128 bytes maximum, of which 104 can be the payload. Therefore we have 832 bits and we may be concerned about the chance of a packet being in error. This is shown in the far column and we see that it is not dramatically different to the bit errors. This is because of the pattern seen before whereby the errors tend to be 100 or 0% at the different stirrer positions. Where it may have more of an effect is in the case where the channel is varying with time. To send a packet of 128 bytes takes 0.004 seconds. If the stirrer is rotating at a rate of 100 steps per second then it will not have moved a single step in the time that a packet is sent but if the rate of change is higher then the system could pass through working and non working mode configurations as a packet was transmitted.

### 4.6.3 Low Qs

To complete our analysis of BERs under different Q-factors we now look at the lower Qs in more detail. We have seen for these that a very small percentage of positions will have a high PER and so we analyse these with many more stirrer positions than before, 1000 per chamber configuration and with greater chamber loadings. The measurement procedure was exactly the same as in the previous section meaning that BER and channel response measurements were done for every stirrer position. In table 4.8 the results are presented for Q-factors ranging approximately from 600 to 4000. As before the chamber configuration is presented in terms of the number of AN79 blocks used.

| Absorber | Q | BER | Number > 0.1 | Number > 0.5 | Number of packets > 0.01 |
|---|---|---|---|---|---|
| 2 | 3948 | 0.1037 | 143 | 103 | 162 |
| 4 | 2648 | 0.04388 | 55 | 44 | 60 |
| 6 | 2012 | 0.0106 | 10 | 8 | 11 |
| 8 | 1428 | 0.00759 | 13 | 7 | 14 |
| 10 | 1120 | 0.007 | 10 | 6 | 11 |
| 12 | 708 | 0.0050 | 7 | 5 | 7 |
| 14 | 616 | 0.009699 | 11 | 10 | 11 |

Table 4.8: Packet error rates for Jennic ZigBee kit with moving stirrer in small chamber with varying Q-factor, second run. Mean delay varies from 0.3$\mu$s to 0.04$\mu$s

Even with the lowest Q we find some positions with >50% error rate because of the much greater sample size. The way that BER drops with increased loading is similar to before but we now have much more accurate figures per configuration. There appears to be a limit on how low the BER can be made as the Q drops below 1000, something that was also noted in the literature study in Chapter 3. A likely explanation for this is the experimental setup whereby blocks of absorbing material are used that may end up attenuating some modes significantly whilst leaving others intact, especially because most absorber was against the walls rather than the floor and was on the ceiling. It is possible that real world performance may be better where attenuation is more evenly distributed throughout the propagation structure. No real world tests on structures of this size were carried out however, so this is purely speculative. All we can say at the end of this measurement set is that for a resonant cavity of size comparable to the York reverberation chamber, one should not expect a radio system to work consistently with a single antenna pair, even for high loadings. Our results show that movement of the

antennas or geometry can stop a ZigBee system working, even if it is only 1 in 100 positions of the antennas.

### Relating signal power and delay spread to BER

We now have a good idea of the bit error rates that one will obtain for a variety of Q-factors for a commercial Q-PSK/MSK system. Our results have indicated that where delay spread is within an order of magnitude of the symbol time, we can expect errors. Still this does not totally explain the pattern seen of a system either working or not working at all with small changes in stirrer deflection.

Therefore, the next area to consider is the way in which bit error rates relate to the channel. We might expect that we can relate the errors to either the signal to noise ratio or the delay for a given stirrer position. Clearly in a standing wave scenario signal strength will change with geometry or antenna position. Given that for now we are ignoring all noise except for that intrinsic to the receiving system, the variation of the $S_{21}$ parameter(the channel impulse reponse) fully describes the variation of signal to noise.

A change in signal to noise from one position to another is not the only parameter affecting system performance. Already we have seen that Q-factor varies with stirrer position, as we pick up different standing waves that may have more or less contact with absorbing surfaces, and this means a varying delay spread too. To try and understand the causes of errors we shall look at a set of graphs where error rate has been overlaid on a channel parameter in order to observe correlations.

In figures 4.29 and 4.30 we plot the power with BER and mean delay with BER respectively against stirrer position. In this measurement most positions have a BER of 100%. Both graphs are from the same measurement set but with slightly different scales. In both cases we can see some correlation between the in band power measurement and the BER and the delay and BER but it is not clear cut. To take an example we might look at the first pair of BER

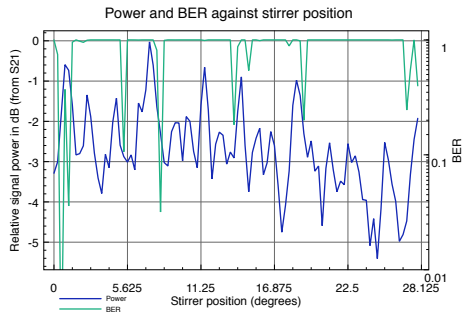


Figure 4.29: Single radio channel power and BER plots for no absorber

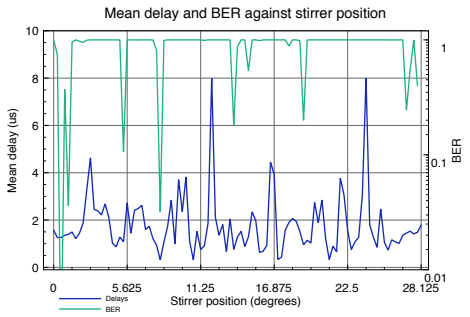


Figure 4.30: Mean delay and BER plots for no absorber

spikes. These occur at a point of high power but low delay spread. On the other hand the second and third distinct spikes occur at low BERs and low delays. Often a high power and high delay would be expected to go together because a high received power could be the result of a strong resonant peak. In other cases a more direct path might have a lot more energy due to the shorter distance but also it will have a lower delay spread because of the dominance of the early component. Nonetheless, when compared with the standard mobile propagation environments we see resonant environments can be different because there can be correlation between signal strength and delay.

One could observe that all drops in the error rate occur when the delay spread is in or close to a local maximum. However it is not neccessarily in the absolute minimums nor is it possible to put a threshold on the signal where if the delay is below that threshold there won't be errors. All we can say is that for this highly loaded chamber there does appear to be a much clearer correlation between delay and errors than is seen for power variation.

Plots 4.31 to 4.40 show the same information but for increasing absorber loadings and thus lower Qs. These graphs appear to show the correlation increasing between signal power and BER and delay spread appearing to have less effect as Q increases. It's important to note that these graphs are not on the same scale and has each been rescaled so that you can easily see the relationship between errors and the channel. One reason why the delay will become less important is because it becomes small relative to the symbol time. Also, when we have a much lower Q the the accuracy of the delay spread calculation will be reduced and the concept of delay spread will be less well defined (for example which is the initial peak is not always clear).

We can also observe that the low BER positions are associated with low power for the loaded chamber (unlike the loaded case). However, we would not expect to be on the limit of the receiver's sensitivity here even though there are some deep fades of over 20dB from a point where there is significant attenuation already from the absorber. It would appear that we have a complex situation where there are a combination of factors causing high error rates. Delay spread or attenuation by themselves are not sufficient to stop a radio from working but the combination of the two can cause errors because a heavily attenuated signal will already be affected more by receiver noise and the delayed previous symbol. There might be a vector addition of phase (for example) that gives a very small value which combined with a deep fade, leads to a high PER.

At very high Qs the errors are very clearly dominated by the delay spread but then as this

drops towards 1000 and less then we move to a situation whereby the error process is more complex and due to both energy and delay considerations. Nevertheless in both situations we still don't see a gradual change in error rate but spikes of error and non error producing positions. The resonant nature of the channel, even when heavily damped means unpredictable performance.

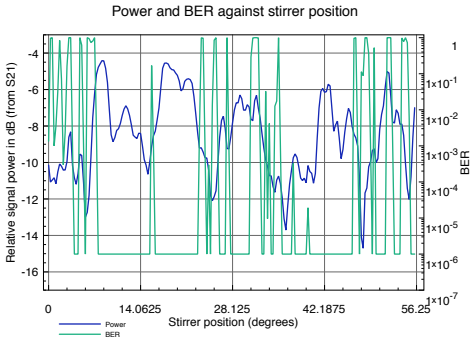


Figure 4.31: Single radio channel power and BER plots for one piece of absorber

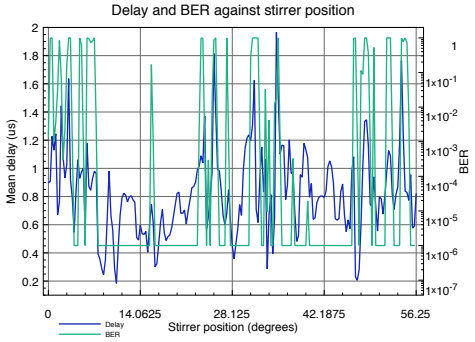


Figure 4.32: Mean delay and BER plots for one piece absorber

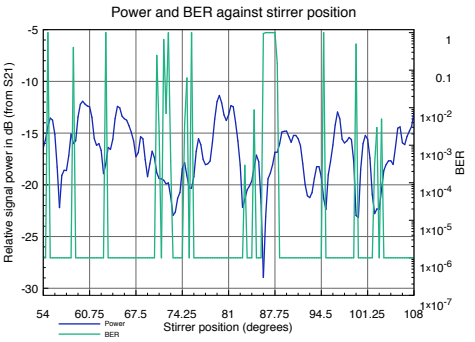


Figure 4.33: Single radio channel power and BER plots for 2 pieces of absorber

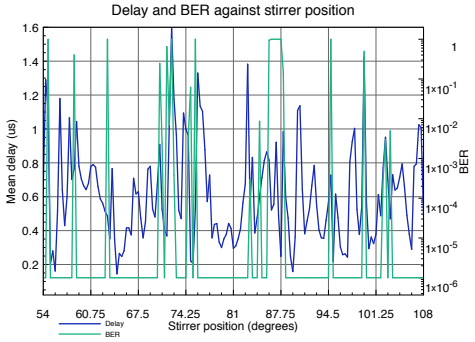


Figure 4.34: Mean delay and BER plots for 2 pieces of absorber

## 4.7 Conclusion

Through this chapter we have started to characterise reverberant environments in such as way as to understand how they might affect wireless systems. We have started out by looking at a few of the basic problems in a reverberant environment such as waveguide cut offs, long time responses and frequency selective channels. Following this we have examined the channel properties and performance of a ZigBee radio in two reverberation chambers with Qs varying

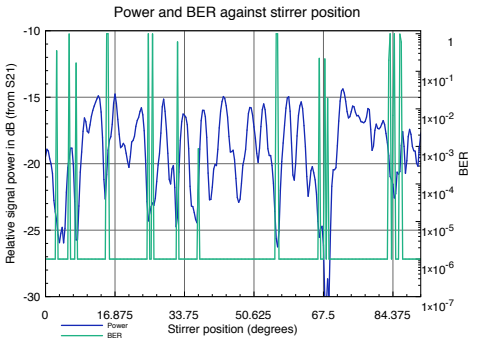


Figure 4.35: Single radio channel power and BER plots for 4 pieces of absorber

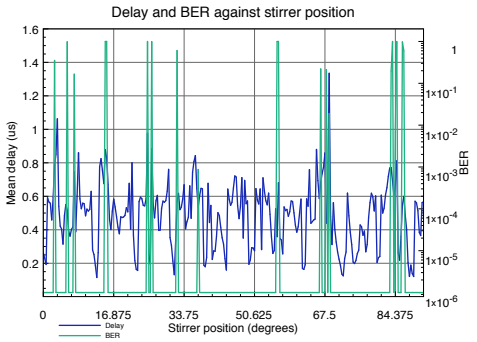


Figure 4.36: Mean delay and BER plots for 4 pieces of absorber

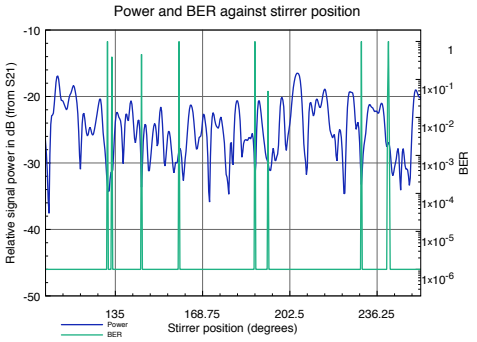


Figure 4.37: Single radio channel power and BER plots for 10 pieces of absorber

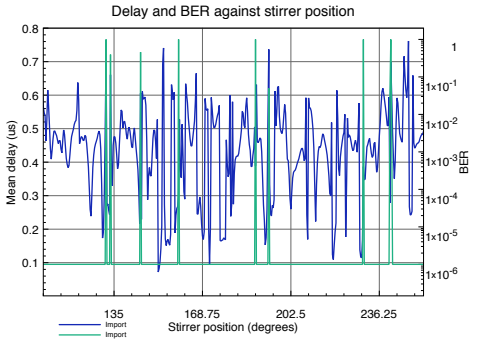


Figure 4.38: Mean delay and BER plots for 10 pieces of absorber

from just under 40000 to just over 400. We have looked at the variation of the S21 parameter with stirrer position and the variation of Q-factor with stirrer position. Although abstract in concept one must not forget that this is exactly the channel we could see in a large metal transport container, or nearly empty aircraft fuselage. These structures will not have a deliberate stirrer but they may experience changes in structure. Perhaps an antenna could be moved by violent vibrations in turbulence or on landing. Or an antenna might be in a wing that is continually flexing or be near some sort of rod being moved for actuation. Our results show that when this happens for Qs in the aforementioned range the performance of the wireless system will vary and in a dramatic way.

Of particular note is the way that stirrer positions with high BERs can occur even when the average Q appears to be relatively low, being under 1000. Such Qs suggest a $\tau$ of approximately 0.1$\mu$s which is an order of magnitude less than the symbol time for ZigBee and so it should not cause problems. The main problem is that it is an average Q and we have seen that both Q

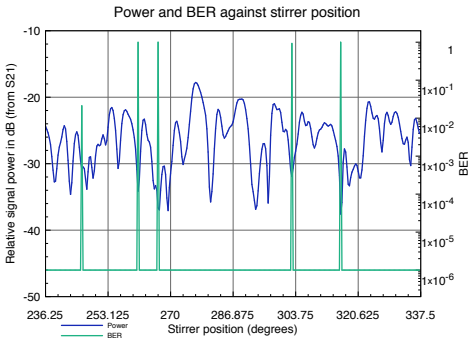


Figure 4.39: Single radio channel power and BER plots for 12 pieces of absorber

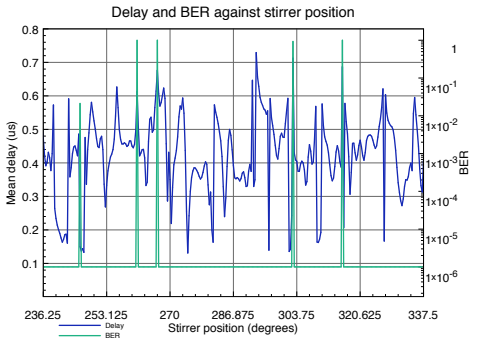


Figure 4.40: Mean delay and BER plots for 12 pieces of absorber

and mean delay can change by an order of magnitude, which would be sufficient to approach the ZigBee symbol time. A point also worth noting is that the average delay is also comparable with the chip time for ZigBee which might be the cause of some symbol errors. Spread spectrum allows the effect of a heavy fade in a narrow band to be negated but the shorter chip time introduces its own problems. If the delay is large and sufficient chips are randomised then the receiver correlation which decode the wrong symbols despite its autocorrelation properties. One would expect an OFDM system to perform better in this regard. The high PERs in some positions in the main chamber may also be a side effect of the way in which the Q was controlled, where rather than having a consistent absorption over the entire chamber, blocks of absorber were used. It is possible that this meant some field configurations were heavily attenuated and others not although this is only speculation. If that was the case then in a real aircraft there would be less problems because the absorption mechanism would be different.

In Chapter 2 we identified ZigBee as the ideal technology, in principle, for a wireless aircraft and so the key question is whether it could work in this environment without any modification. The analysis here gives a mixed picture. For very high Qs, i.e. over 5000 it is not suitable for flight control and if one wishes to analyse predictable links between pairs of radios it is not because there is a probability of the system failing of 1 in 100. However there are obviously ways in which to improve performance and those are to use diversity and redundancy. By moving to 2 or 3 antennas rather than 1 we can reduce the probability of failure to $10^{-4}$ or $10^{-6}$. We might also combine antenna redundancy with with other forms of redundancy such as allowing re-routing a packet via other radios.

For example, consider the setup where there is a direct link from an actuator to a micro-

processor system and that link fails during flight. This could be serious, but there might be other radios being used for other less critical applications transmitting nearby. If these had a lot of spare capacity they could take over routing of critical packets restoring flight control functionality. In fact such a setup is supported by ZigBee's mesh network configurations. In an even better solution additional logic might be added to move between frequencies upon failure. Therefore the situation is not as bad as it first looks; the resonant environment is very bad at guaranteeing a PER at a particular frequency and antenna positions but between the problem positions are many that work that we could take advantage of.

What we have seen so far is that sealed metallic boxes are highly reverberant and this leads to a communications channel that isn't flat over the width of a ZigBee (or Wifi which uses the same range) channel. This leads to inter-symbol interference and high average BERs for the ZigBee standard, until the Q drops to order 1000. We can observe that it is often a case of a system either working or not working, rather than a gradual performance drop off as one might expect from Gaussian noise. Also it has been shown that the causes of the errors are varied and at higher Qs delay spread is dominant, at low Qs received power has a much larger effect. Next we shall look at more complicated scenarios with multiple cavities, measurements in a real aircraft and the introduction of noise.

# Chapter 5

# Coupled cavities

## 5.1 Introduction

Sometimes in an aircraft the radio system designer will only be interested in sending data through one part of a wing or the fuselage or a single avionics bay. In other cases data must propagate from one end of an aircraft to another or through the entire length of a wing or from a central point into a number of avionics bays, perhaps through apertures of limited size. Furthermore, there may be unwanted coupling between sections of the aircraft that ideally would be working on independent networks but are sharing the same frequency band. Where there are very few and narrow coupling points the literature suggests that MIMO system performance will drop.

In light of all this it is necessary to consider what happens when different resonant cavities are coupled to one another. The electromagnetics literature has covered in detail the shielding effects and attenuation introduced by slots and apertures but there is much less information about what this coupling means for a radio system, in terms of the field statistics and the delay spread and what might be expected in terms of received signal strength. Such areas of interest are the principal topic of this section.

### 5.1.1 Possible scenarios

As already alluded to there are a number of possible scenarios that feature coupled resonant cavities. Reference [50] shows that there is no such thing as a typical aircraft but there are certain features that occur in most. For example, in an aircraft wing one would expect a ribs and spa structure which would effectively separate the wing into a number of separate cavities

with apertures just large enough to allow rods, cables or hydraulics through. In terms of the fuselage this too may be sub divided with bulkheads providing structural integrity and there may be a number of different smaller bays within these larger spaces. If the apertures separating the different areas were large then one might intuitively expect that the system would look similar to a single cavity and as the aperture got smaller then there would be a gradual move to a multiple cavity type system.

There are many possible combinations of cavities and apertures and it is not practical to measure or simulate all the possible permutations. The best option is to take some very simple scenarios and examine those in order to understand what might happen. This has been done in a number of ways. To begin the main reverberation chamber is used which has already been discussed in detail a single box is added inside this which may be likened to the situation where there is a small box or bay in a larger fuselage. Following a small reverberation chamber is used in a nested configuration. This has the advantage that it can be stirred, providing statistical information about what happens to the signal in coupled cavities. It should also have a much larger mode density than the small boxes.

### 5.1.2 Theory of coupled cavities

The main reason for considering coupled cavities, besides the fact that they are regularly encountered in aircraft, is that they might be expected to present different statistics compared to the case of a single cavity. This would mean a different probability of the signal dropping below a given level and thus a different probability of a packet error. It may also be found that the delay spread is different in a coupled cavity environment compared to a single cavity one and this needs to be established. Where there are narrow coupling points it would also be expected greater attenuation than in the single cavity situation.

To begin a very basic model is presented for the coupled system which is shown in figure 5.1 In this there are on the left a set of parallel resonant circuits each similar but slightly different component values. Each branch will have a transfer function of the form:

$$Y_i(s) = \frac{I_i(s)}{V(s)} = \frac{s}{L_i(s^2 + \frac{R_i}{L_i}s + \frac{1}{L_i C_i})} \tag{5.1}$$

If this were rewritten in standard second order system notation, replacing the Ls,C,and Rs with damping factors and resonant frequency $\omega_0$ then it would be exactly the same as a mode in the reverberation chamber. The total response is obtained by summing these branches and

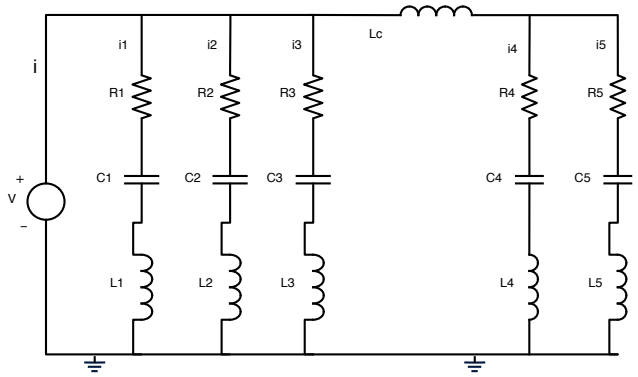


Figure 5.1: A simple model for coupled resonant cavities

obviously in reverberation chamber there are many more branches than seen here. [1] When a second cavity is added, the same circuit is added, only this time the Rs,Ls and Cs may be significantly different producing different Qs. Linking these together we have shown an inductor in the diagram which makes the assumption there is only one aperture or coupling point. It's an inductor because we might expect that when the wavelength is very much bigger than the aperture no energy will be propagated but when it is of comparable size or smaller the there will be little impedance and energy can easily flow between the two cavities. When this happens the cavity might almost look like a single cavity with energy bouncing between the two and relative energies being dependent on the 2 Q-factors. This type of situation has seen recently researched by Hill and Holloway and like a lot of areas described in these chapters has been been given a brief treatment in [42] which includes a power balance description.

Clearly our model here is overly simplistic in a number of ways. Firstly an aperture does not look just like an inductor, and may in fact have it's own resonances and may only couple particular polarizations. It many ways it could closer resemble an antenna with a radiation pattern. Secondly, the assumption that there is only one aperture might be simple to keep in the lab but it won't be the case in an aircraft which will have many couple points. Despite these problems this model will be used for now as it provides an intuitive view of the problem and can help to understand the statistics of a two cavity system.

In section 3 it was seen that the distribution of electric field followed a Rayleigh distribution and the distribution of power an exponential distribution. The first question to ask then is what distribution one might expect in a multiple cavity system. What will happen is de-

[1] In fact a model such as this was discussed with, and used by Nottingham university who were the project partners on the electromagnetics part of the Flaviir Wireless project.

pendent upon the coupling mechanism: we may have a single coupling point or a number of these. Initially, consider the case of a single coupling point, this may be a small opening (relative to the wavelength) or it could even be antennas in each of the cavities with wires between them through the cavity walls. As has just been discussed, what there is in essence is a linear system transfer function between transmitter and coupling point in cavity A. Then, on the other side is another linear system function between the coupling point in cavity B and the receiver. For each of the two systems, at a given frequency there is a random variable describing the attenuation and the total received signal is a joint distribution.

Rayleigh distributions are assumed for the S21 parameter within a single cavity and denote this $f(a)$ and $f(b)$ for the two cavities. Thus one can write

$$f(a) = \frac{a}{\sigma_1^2} e^{\frac{-a^2}{2\sigma_1^2}} f(b) = \frac{b}{\sigma_2^2} e^{\frac{-b^2}{2\sigma_2^2}} \tag{5.2}$$

for the density functions of random variables A and B taking on the values a and b. $\sigma$ defines the distribution scale. To obtain the new coupling distribution a new random variable U = A*B can be defined and we shall be interested in the probability:

$$F(u) = P(U \leq u) = \int_{x=0}^{x=\infty} \int_{y=0}^{U/x} \frac{ab}{\sigma_1^2 \sigma_2^2} e^{\frac{-a^2}{2\sigma_1^2} + \frac{-b^2}{2\sigma_2^2}} \tag{5.3}$$

Fortunately this does not require solving as the solution has been derived in the literature and is given, for example in [59] as:

$$f(u) = \frac{u}{\sigma^2} K_0(u/\sigma) \tag{5.4}$$

This distribution and a histogram of the results are plotted in figure 5.2

If there are more than one coupled cavity then there will be a multiplication of 3 or 4 or more Rayleigh variables leading to distributions shown in figures 5.3 and 5.4 where the number of displayed bins have changed to try and show the behaviour more clearly. It is found that the distribution becomes far more skewed to the left but with the small possibility of quite extreme values as shown by the long tails. The significance of this from a communications systems perspective is that one should not be very cautious about testing a real world system with a small data set and assuming it is well understood. When sampling the single Rayleigh channel it has already been seen that a particular number of samples is needed to bound the PER and SNR and to get a representative histogram for the underlying distribution. For the cascaded cases most of the area is concentrated in the first few bins, and far more extreme attenuations (and relative gains) may occur and so one must be aware of this.

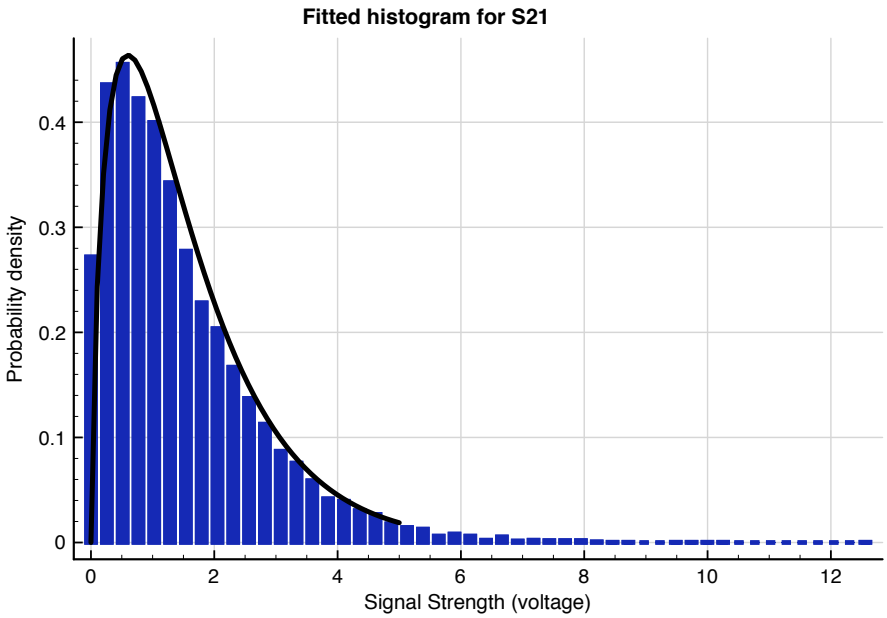


Figure 5.2: Distribution function shape for a double Rayleigh histogram

In many cases there will not be a single propagation path through a narrow aperture but a variety of paths through various apertures. Where this happens there is a sum of fields and the distribution wil revert back to the Rayleigh distribution. Note that this won't become Gaussian because there is a vector addition, leading to the Rayleigh.

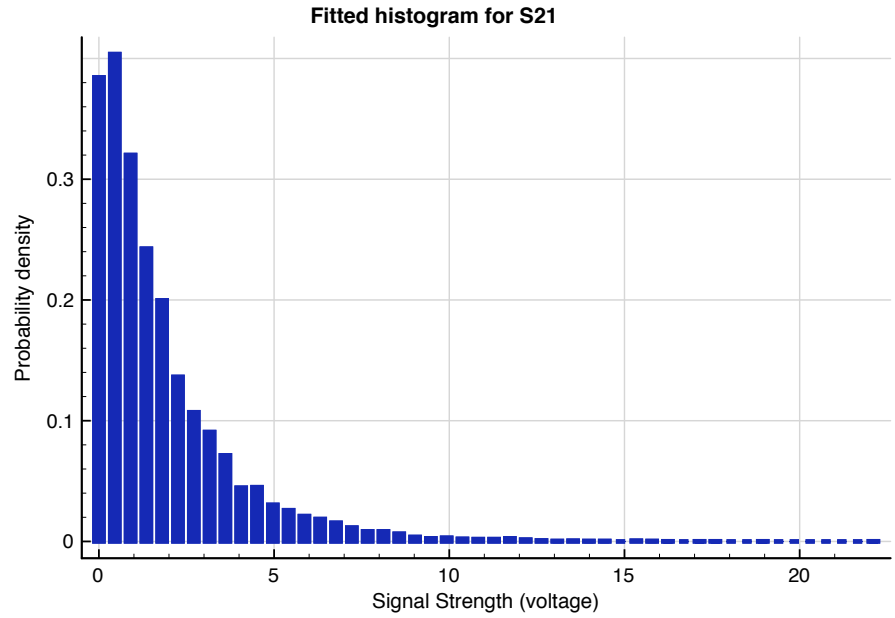


Figure 5.3: A 3 Rayleigh cascaded distribution

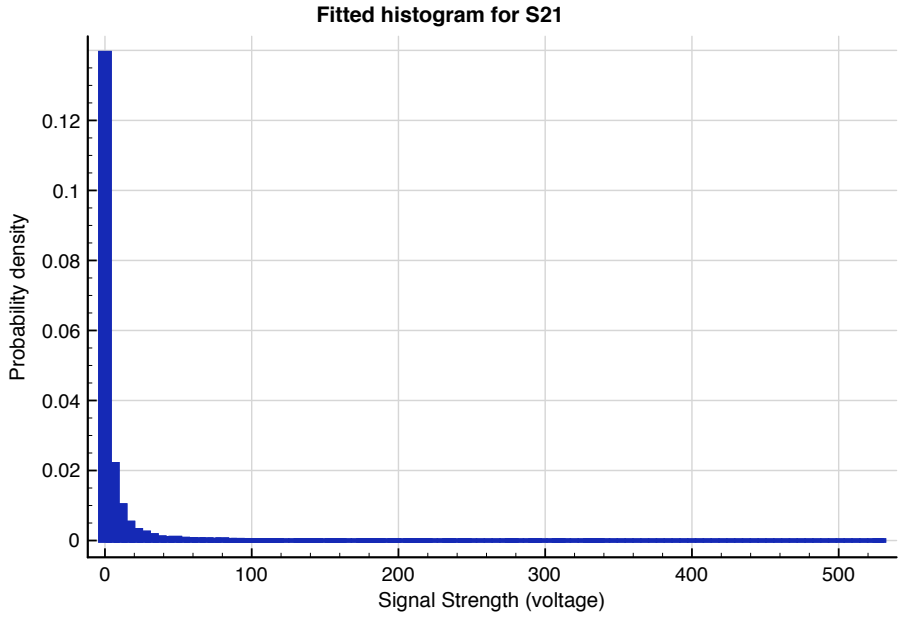


Figure 5.4: An 8 Rayleigh cascaded distribution

Alongside the amplitude statistics we should also consider the time domain signal. In a reverberant environment there is a decay of energy approximately proportional to the Q-factor of that environment. If there are two cavities then both will have their own time delays. When two paths are combined via a "pinch point" there are cascaded transfer functions and these correspond to a convolution operation in the time domain. Intuitively it might be expected that the delay to increase when two reverberant environments are cascaded because a signal at a pinch point can be approximated by a train of weighted impulses, with each of those impulses being spread out by the impulse response of the second cavity. Figure 5.5 demonstrates

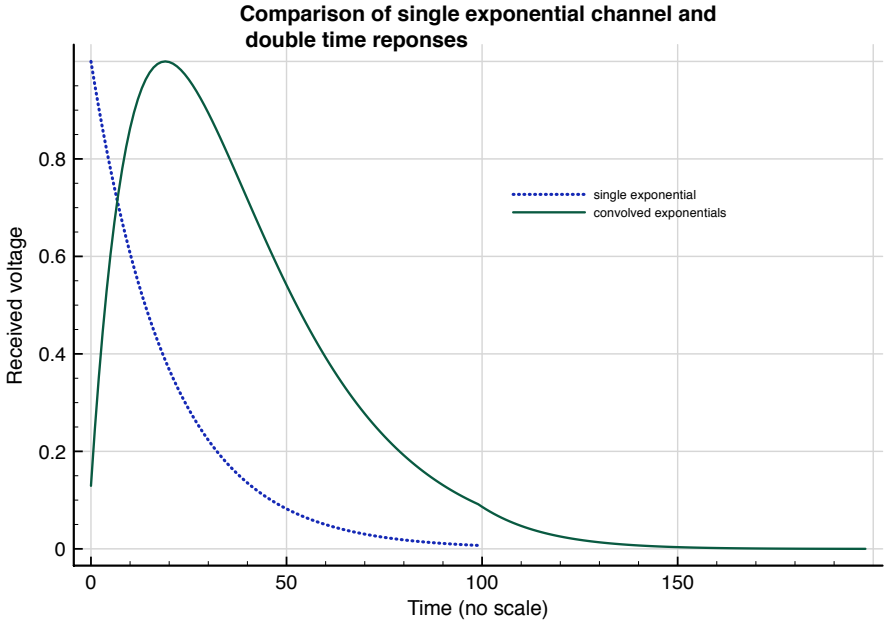


Figure 5.5: Time decays for a single cavity and multiple cavity scenarios

the change in time response for the case where there is a single cavity and the case where we have 2 cavities with identical time constant $\tau$. Analytically this can be obtained via convolution:

$$\int_0^t exp(-\frac{t}{T})exp(-\frac{t-\tau}{T})d\tau = texp-\frac{t}{T} \tag{5.5}$$

where T has been used as the time constant to avoid the clash with $\tau$ typically used in the convolution integral. There is still the same exponential decay constant and there won't be large, order of magnitude changes in delay spread but the time response can look slightly different in the way it increases before decaying. This should not have a significant effect on a communications system except where the delay spread is very close to the limit of causing inter-symbol interference.

### 5.1.3 Small boxes in a reverberation chamber

Next some experiments are presented on coupling into small boxes to explore further the issue of cavity coupling before moving onto a more interesting setup with the small reverberation chamber. Pictures of the two boxes are given in Figures 5.6 and 5.7 and their dimensions are 270x270x270 and 170x170x170 respectively. With the boxes open and a wavelength signifi-

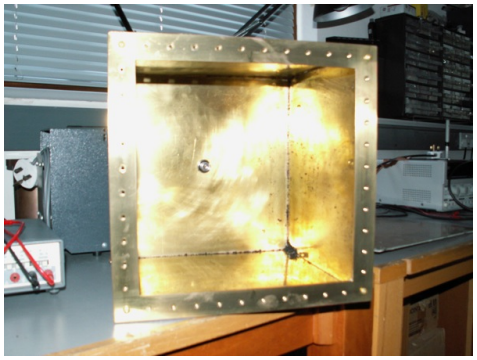

Figure 5.6: Large brass box

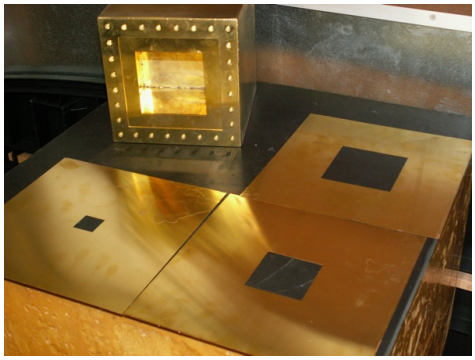

Figure 5.7: Small brass box and some apertures

cantly higher than the opening one would expect a channel no different to placing antennas directly into the reverberation chamber. Therefore a selection of apertures were placed on the open end of the boxes and were secured with screws of a separation just less than the minimum aperture size, so as to avoid the gaps between them leaking as much energy as the aperture. Aperture sizes are given in figure 5.1

| Aperture | Width(mm) | Height(mm) |
|---|---|---|
| Large box 1 | 240 | 240 |
| Large box 2 | 100 | 100 |
| Large box 3 | 70 | 70 |
| Large box 4 | 30 | 30 |
| Large box 5 | 20 | 80 |
| Small box 1 | 145 | 145 |
| Small box 2 | 100 | 100 |
| Small box 3 | 70 | 70 |

Table 5.1: Aperture sizes for boxes of size 270x270mm and 170x170mm

For all these measurements the chamber was loaded with 6 pieces of absorber. As usual one antenna was behind the stirrer to avoid a direct path, the other was inside the box, at the

opposite end of the box to the opening. The largest aperture size in the table actually corresponded to the case where there was no aperture and one end of the box was open. The number of stirrer positions was varied according to the measurement (due to the measurement time) meaning that less positions were used when error rates were found immediately to be very high and on each measurement the stirrer moved 0.3degrees. 400 packets were sent per stirrer position. Results for these set ups are presented in table 5.2

| Setup | AverageS21 | Average BER | % Positions > 1% BER | % Positions > 50% BER | % Positions > 1% PER |
|---|---|---|---|---|---|
| Ext | 264 | 0.0228 | 0.78 | 0.52 | 0.78 |
| L 240 | 268 | 0.0118 | 0.39 | 0.39 | 0.79 |
| L 100 | 80 | 0.01985 | 1.17 | 0.78 | 1.18 |
| L 100 | 64 | 0.0228 | 2.09 | 1.39 | 2.78 |
| L 70 | 8.54 | 0.0322 | 1.25 | 0.94 | 1.56 |
| L 30 | 0.03 | 0.326 | 24.5 | 19.01 | 27.6 |
| L 30 | 0.04 | 0.464 | 29.6 | 27.8 | 31.3 |
| L 20 x 80 | 1.56 | 0.0476 | 1.6 | 1.36 | 1.6 |
| S 145 | 210 | 0.0099 | 0.32 | 0.32 | 0.65 |
| S 70 | 172 | 0.027 | 1.3 | 1 | 1.3 |

Table 5.2: BER and different aperture sizes for boxes in the main reverberation chamber

Rather than tabulating Q a relative measurement is given called average S21 and why this is is explained later on. It can be seen that the BER performance gets worse as the aperture reduces in size which would certainly be expected since a narrow aperture will cause attenuation once its dimensions make it, in effect, into a waveguide cut off below the operating wavelength. In fact there is little change for the larger box until an aperture dimension of 30mm which is $\lambda/4$. At double this there was no difference in errors compared to larger apertures, but at $\lambda/4$ there are 30% of positions with a BER of 0.5 and greater. This is not necessarily the only effect relevant to the BER increase. Potentially there are a number which include a change of in delay spread, a change in the channel statistics and in expected signal strength may all lead to a different BER relative to the single cavity situation. Now these are examined in more detail.

By comparing the delays and signal strengths of the empty chamber case, with the situation with antennas in the boxes it should be possible to see what effects are dominant. In Figure 5.8 the variation of Q in the chamber is shown which demonstrates that although typically the Q is around 400 there are some positions where it increases to around 3200. These are relatively low compared to some of the Qs considered in Chapter 4 and so the relatively small

number of positions with a high BER is not surprising and it can be assumed that high Q is not a cause of errors with no aperture.

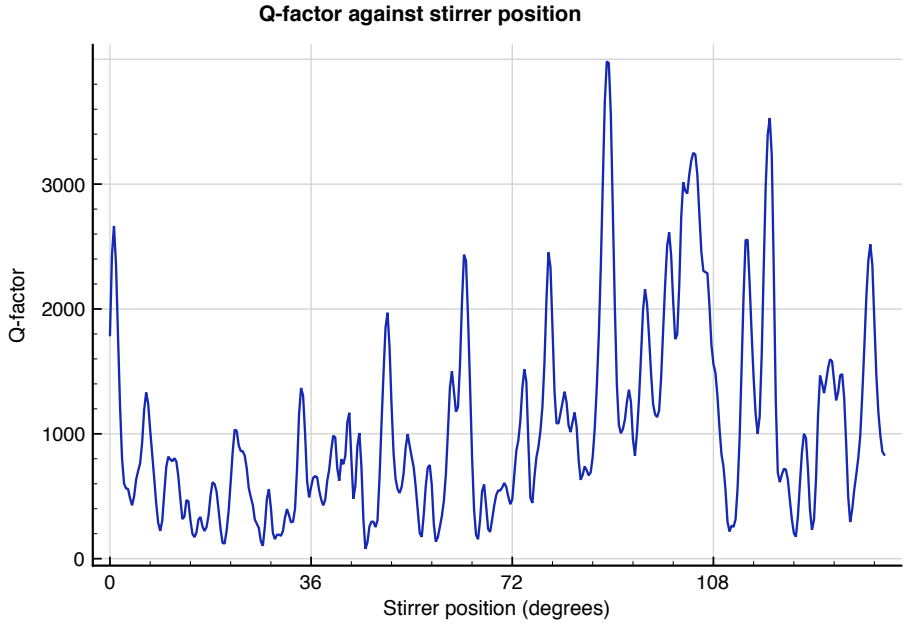


Figure 5.8: Variation of Q-factor for reverb chamber with 6 absorber pieces

Moving to a two cavity scenario it doesn't make complete sense to talk about the Q-factor because both cavities will have their own Q. However it is still possible take an average of S21 measurements squared which is what the Q measurement actually is along with a coefficient that takes account of the volume and frequency and antenna efficiency. On the other hand the time delay is still a relevant quantity and won't be affected in the same way as the Q-calculation by the aperture attenuation. Figure 5.9 plots the average (over frequency and stirrer position) of S21 against the average (over stirrer positions) of the delay spread and what it demonstrates is that there is no change in delay spread for the different aperture sizes despite large changes in the ratio of received to transmitted energy. There is actually a 10% variation in delay spread but this is small both in comparison to the differences one sees by varying loadings and also in comparison to the changes seen at different stirrer positions.

Before continuing there are two points worth raising: firstly that 5.9 demonstrates the usefulness of the delay spread metric as opposed to frequency domain based metrics. Secondly that it shows that when there is a large mismatch in Qs (in this case due to volume differences rather than absorber), one cannot expect convolution of the two time domain responses to make any difference. In the situation where one Q is significantly larger than another it is the

reverberant effects of the larger cavity that are more significant.

Note that the larger reverberation chamber has an unloaded Q of between 30000 and 40000. The larger of the two boxes might be expected therefore to have a Q-factor of less than 100. Noting that Q is also $\frac{f_0}{\delta\omega}$ any modes will have a bandwidth of 24MHz which is far larger than the ZigBee band and also larger than that for the WLAN bands.

It is possible to calculate the mode density by

$$d_n = \frac{dN_{modes}}{df} = \frac{8\pi abcf^2}{c_0^3} - \frac{(a+b+c)}{c_0} \tag{5.6}$$

which gives the number of modes in a ZigBee channel as 0.32. In fact across all the ZigBee bands one can expect 9 modes and so within a given band there will only be one very wide mode. The cut off frequency for the larger box is 0.7GHz and for the smaller box is 1.2GHz so the measurements were well above cut off but with a low mode density and wide modes.

The main parameter of interest then is the level of attenuation through the aperture. If there is a low signal level close to the smaller box and there is only an aperture of 1/4 of the wavelength then there may be high error rates. On the other hand, with an aperture of half a wavelength or more there is strong coupling into the box and the signal strength variation in the main chamber and the box are the same. Note that in the case of the thin slot coupling is strong as the width is greater than $\lambda/2$. Polarization should not be a concern in a reverberant environment such as this.

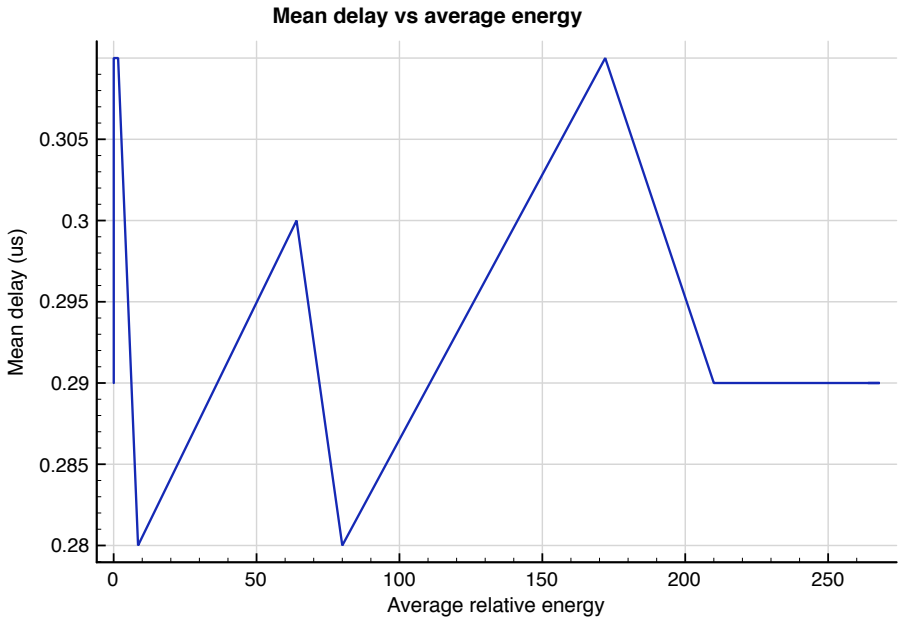


Figure 5.9: Average S21 squared (normalised) against mean delay

There remains one final point to consider here and that is the channel statistics. A histogram of S21 (received to transmitted voltage) is plotted in Figure 5.10 for the case of just the reverberation chamber and for the chamber-box system with the 70x70mm aperture in Figure 5.11. There appears to be a Rayleigh for the chamber but not for the box case despite expectations because of the box's low modal density and Q. Figure 5.12 explains the discrepancy by showing the average S21 with frequency (it is presented as a Q calculation but as before this does not have real meaning) which is far from flat.

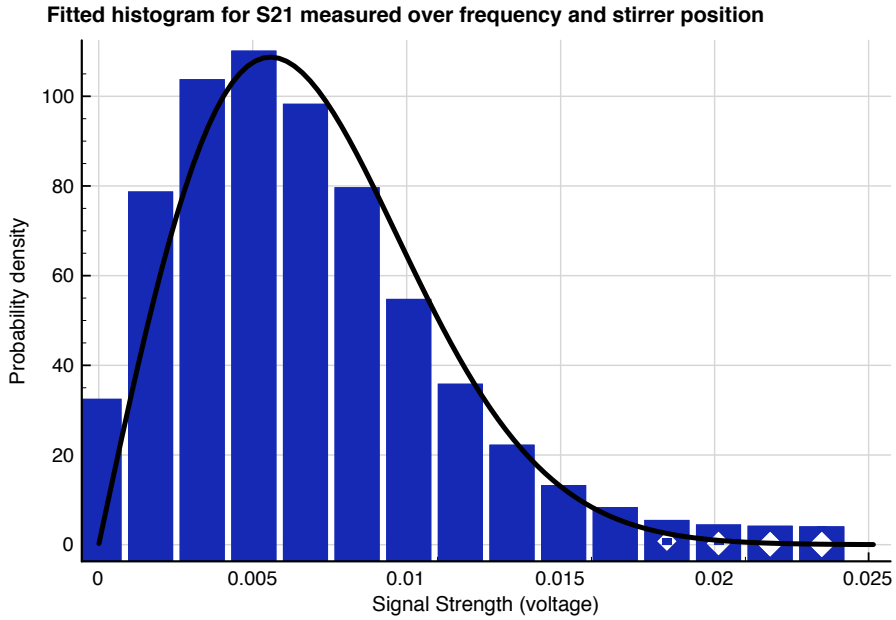


Figure 5.10: Histogram of received field for external antennas

In the frequency domain there is a very clear resonance associated with approximately 2.466GHz and 2.461GHz, with transmitted energy being an order of magnitude lower at all other frequencies. If the same graph is plotted for the aperture of size 3x3 then what is seen is a much sharper resonance at the frequency corresponding to a wavelength that is 4x the aperture width and height. The coupling inductor actually looks like a resonant filter and distorts the response.

Through this small example some of the complexities of coupled cavities have been illustrated. As far as an aircraft is concerned, it is not realistic to expect the precise coupling function to be understood. Typically there would not be just one aperture between two sections and that aperture would not be a perfectly square or rectangular shape. Exactly what size and shape they are though can't be known but there are two points to be made:

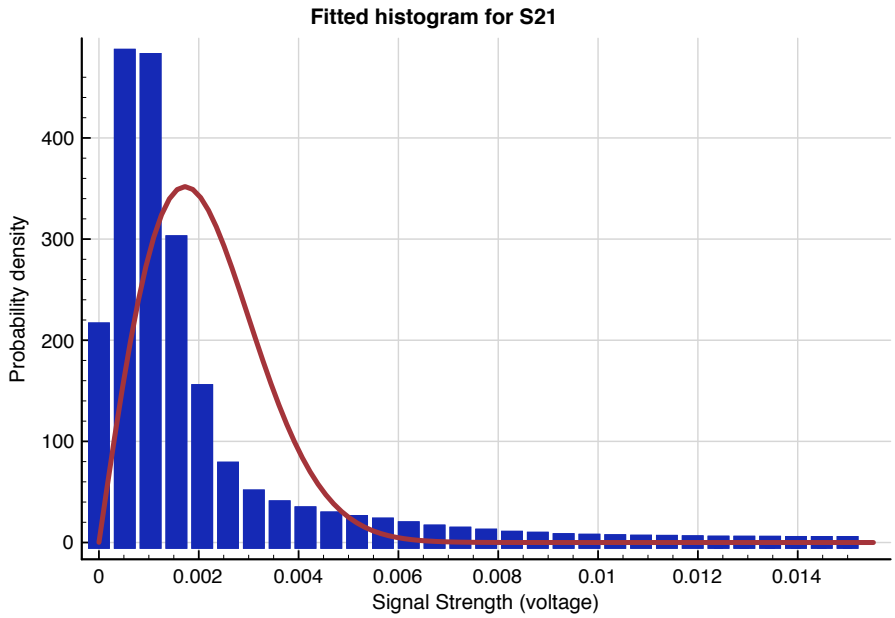


Figure 5.11: Histogram of received field for internal antenna

1. Wherever any two sections must be linked and the apertures are of a size less than half a wavelength, significant attenuation is to be expected and compensate for this.

2. Where all coupling is done over an electrically small aperture the statistics can be skewed and care must be taken in making assumptions about the probability distribution of the signal strength

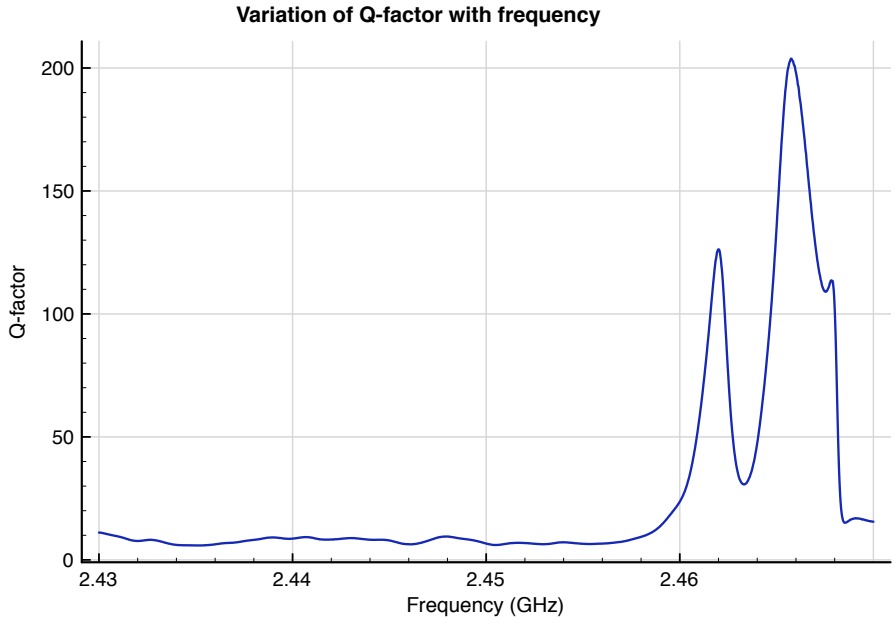


Figure 5.12: Variation of Q-factor against frequency for 70x70mm aperture

## 5.2 Nested chamber

This section presents measurements and results for a slightly more complex situation than the previous because the small boxes are replaced with another reverberation chamber. This means that both cavities are complex multiple moded environments that can be stirred. The small chamber was in a nested arrangement set inside the main chamber and this is shown in figure 5.13. The inside of the chamber is shown in figure 5.14 where the stirrer and some absorber can be seen. Of particular importance is the use of the small stirrer to generate

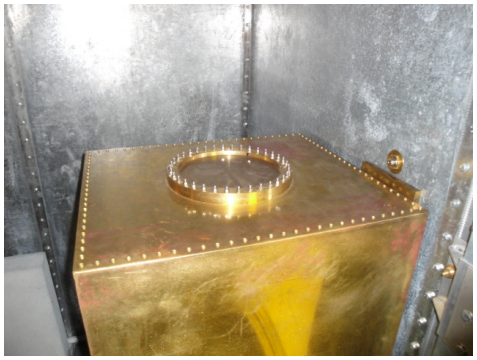

Figure 5.13: Small reverberation chamber in larger chamber

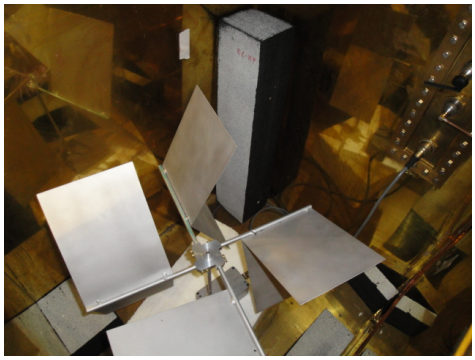

Figure 5.14: Paddle and absorber inside smaller chamber

the complete statistics of the system. This wasn't done for the small boxes and would not be of great use because of their low Q-factors. When there are two reverberation chambers then performance may vary both with the position of an antenna in chamber A and with the position of an antenna in chamber B. If one did not stir the small chamber then for small apertures it would be as though there were two static antennas, or a constant attenuation in the small chamber meaning that one would not see a double Rayleigh distribution.

### 5.2.1 Experimental setup

The nested chamber was placed on the side of the main/larger chamber, as shown in Figure 5.13 so that both receive and transmit cables could be brought into the large chamber at the same point from the outside. Placement of the antenna in the outer chamber was behind the paddle as before which helped reduce any direct path and allows for maximum scattering by the paddle onto the absorbing surfaces. In the inner chamber the antenna was placed as shown in figure 5.15, connected to the wall but at sufficient distance that it received a random

signal (otherwise the wall boundary conditions restrict the received field) By using this setup

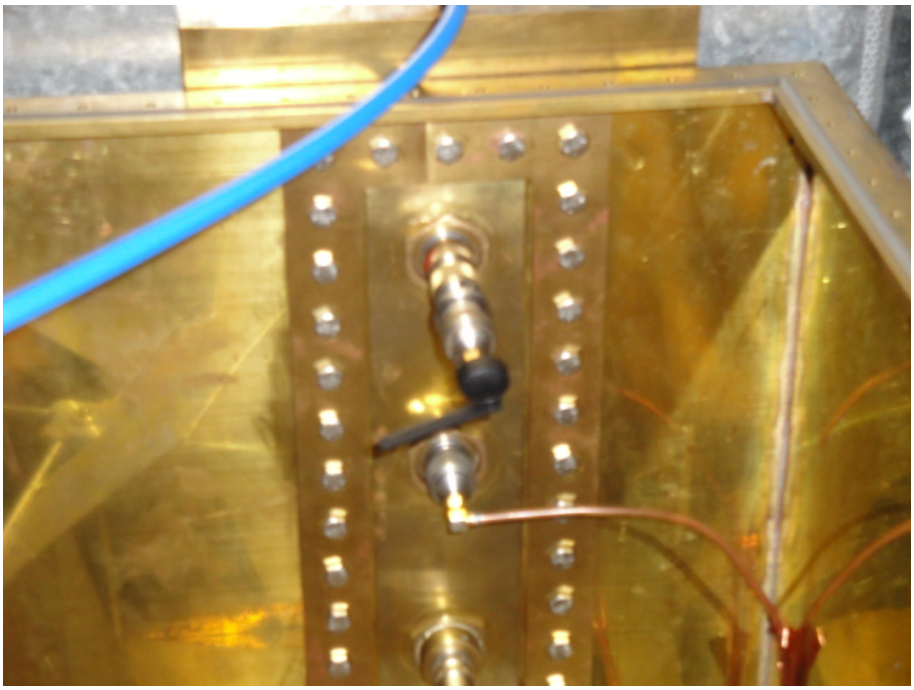

Figure 5.15: Antenna placement in small chamber)

it possible to get a statistical description of a a system that has two separately fading environments because now one can stir both resonant environments separately. The larger of the two chambers was at a higher Q when empty and so required absorber to bring it to a level close to the smaller chamber. This was done so as to avoid a situation where one of the two cavities dominated the channel response, as happened with the large reverberation chamber and the much smaller boxes. In fact two configurations were used, one a higher Q configuration with 6 pieces of absorber in the larger chamber, none in the smaller and the other with 9 pieces of large absorber in the larger and 3 in the smaller.

As in the previous experiments coupling into the smaller chamber was controlled by using a selection of different apertures. Aperture size was controlled by using thin brass circular sheets, each with different sized aperture in the centre. These were clamped under a metal ring screwed down to the box's lid with tightly spaced screws.

The various configurations of aperture size and Q/absorber are presented in figure 5.3. As before measurements were done with a static stirrer at a number of discrete positions and both ZigBee BER and network analyser channel response measurements were done for each position, utilising an computer controlled RF switch to alternate between them.

| Number | Aperture(cm) | Small Chamber Q | Large Chamber Q |
|---|---|---|---|
| 1 | 26, diameter | 4500 | 1000 |
| 2 | 6x6, square | 4500 | 1000 |
| 3 | 4x4, square | 4500 | 1000 |
| 4 | 3x3, square | 4500 | 1000 |
| 5 | 1x1, square | 4500 | 1000 |
| 6 | 0x0, square | 4500 | 1000 |
| 7 | 26 diameter | 400 | 550 |
| 8 | 6x6, square | 400 | 550 |
| 9 | 4x4, square | 400 | 550 |
| 10 | 3x3, square | 400 | 550 |
| 11 | 1x1, square | 400 | 550 |

Table 5.3: Aperture sizes and BERs for coupled cavity measurements

### 5.2.2 Results

First the bit error rates are presented for the ZigBee radios in figure 5.16 which shows the BER for both the higher and lower Qs at the different aperture sizes.

The BER shown in the figure is the average over the stirrer position, that which might be expected for a system in a slowly continuously stirred or time varying environment. A value of 1 means 100% packet loss. In the best cases there is a large aperture size and average BER is 2.8% which is as expected from previous measurements. When the gap comes down to 6x6cm there appears to be no effect but once we pass this point is passed (which is approximately $\lambda/2$) then there is an increase in the bit error rate. For the case of the higher absorber it is not possible to send any packets successfully, for the lower case it might be but the system is unusable.

Interestingly, for a 1x1cm aperture the BER is much higher at the low Q setup. On the other hand for the 3x3cm aperture the BER is higher for the higher Q setup. It appears there are at least two mechanisms causing BER, absolute signal power and also inter-symbol interference. This is reinforced by table 5.4 which provides information about the breakdown of the errors. For most aperture sizes the percentage of stirrer positions with a PER > 1% is about 1% before dramatically increasing to be every position. Clearly there has been no increase in the Q, but the signal strength has decreased with the size of the aperture. In both cases there is a limit on

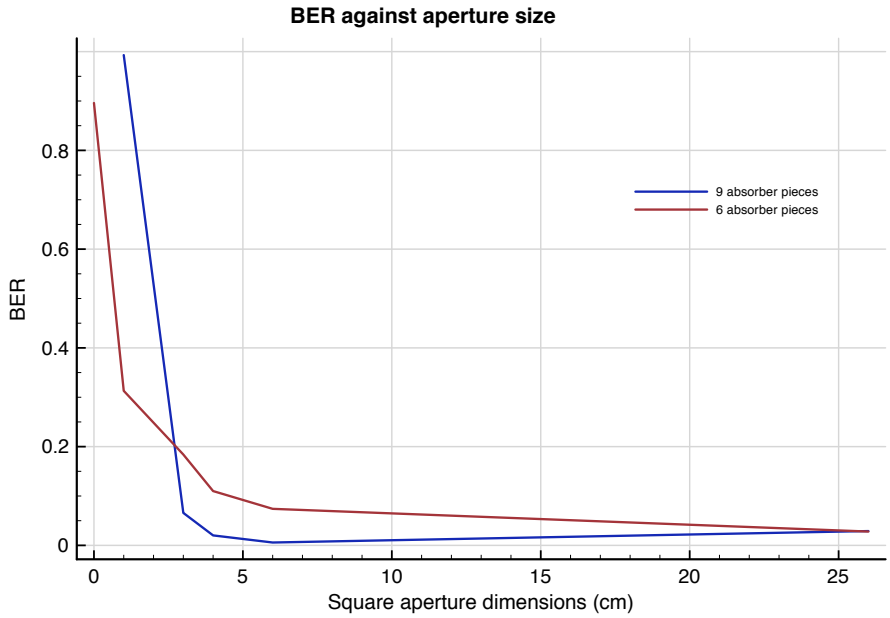


Figure 5.16: Bit error rates against aperture size for the higher Q setup

the sizes of the apertures after which PER rapidly increases. Further, we must note that this is a noiseless environment and that in a real world scenario a system could not tolerate the same level a attenuation.

| Setup | Q | BER | NoPositions > 0.1% | No Positions > 0.5% | No Positions > 0.1% |
|---|---|---|---|---|---|
| 26 diameter | H | 0.028 | 4 | 3 | 4 |
| 6x6 square | H | 0.074 | 1 | 7 | 11 |
| 4x4 | H | 0.109 | 16 | 11 | 19 |
| 3x3 | H | 0.184 | 24 | 18 | 27 |
| 1x1 | H | 0.313 | 40 | 32 | 49 |
| 0x0 | H | 0.896 | 95 | 92 | 100 |
| 26 diameter | L | 0.029 | 3 | 3 | 3 |
| 6x6 | L | 0.0057 | 1 | 1 | 1 |
| 4x4 | L | 0.0201 | 3 | 1 | 3 |
| 3x3 | L | 0.065 | 1 | 0 | 1 |
| 1x1 | L | 0.99 | 100 | 100 | 100 |

Table 5.4: Aperture sizes and BERs for coupled cavity measurements, final column is PER

By looking at the delay spreads in figure 5.17 the delay spread does increase as was suggested at the beginning of this chapter and that the increase is far more prevalent for the case of higher Q channel. In fact the delay for the low Q case, even with a 1cm aperture is still less than that in the open, high Q setup. Therefore it is unlikely that any changes in error rates at low aperture sizes are due to the delay. The network analyser results show that the signal is very close to the noise floor and exists at a level somewhere below -85dBs (after cable calibration). With a transmit power on the Jennic of 1dBm and a sensitivity of -96dBm, it is therefore likely that the combination of absorbing material and aperture attenuation is what causes the highest error rates.

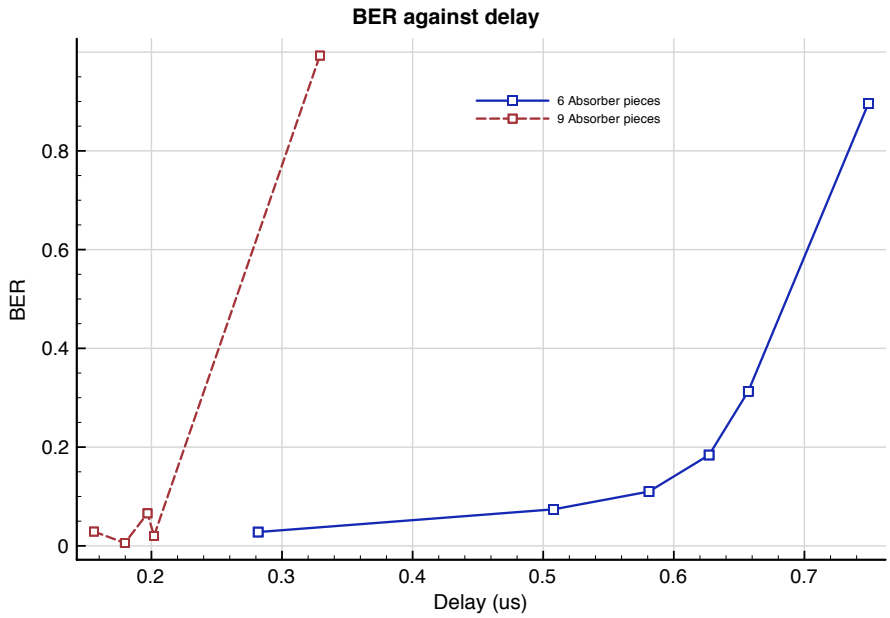


Figure 5.17: Bit error rates against time delay, longer delays correspond to smaller apertures

In the high Q cases the delay spread is ranging from 0.3 to 0.75$\mu$s and IEEE802.15.4 has a bit time of 4$\mu$s and a chip time of 1$\mu$s. This is heading into the range where a delay spread could have an effect. It would be useful to compare the delay with the errors for some of these cases to see if the maximum bit error positions correlate with maximum delay positions. We shall take the case of the 4x4 aperture first in the higher Q environment and in calculating the delay only the ZigBee band is used.

Next overlays are given of signal power and errors.Both are plotted on dB scales and the maximum BER is 1 for reference. The signal power is not from a signal analyser measurement but comes from integrating over S21 and has a variation of approximately 5-10dB. In figure

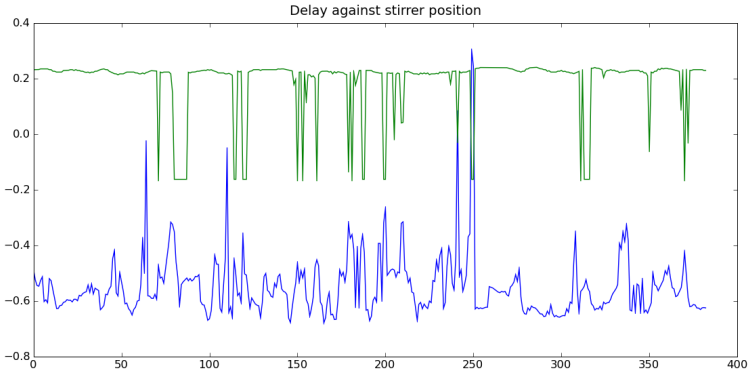


Figure 5.18: Bit error rates overlaid with delays (4x4 aperture)


Figure 5.19: Bit error rates overlaid with channel power (4x4 aperture)

5.19 and in figure 5.21 for the 4x4cm and 3x3cm apertures respectively it can be seen that there is some correlation between the signal power and the error rates although in different ways for the 4x4cm and 3x3cm aperture. For the former the minimum BER often corresponds to a low power whereas for the latter case this is sometimes the case but there are also incidences of high powers and low BERs coinciding. This may be because high powers can suggest high Qs and delay spread as well as low attenuations and so sometimes the high power corresponds to modes with energy storage, sometimes low power is due to serious attenuation.

The same is also true for delay spread which is shown in Figures 5.18 and 5.20. A number of medium and high delay peaks can be seen that line up with high BERs. For example in figure 5.18 there are two very obvious matches around stirrer position 250 and a couple of less obvious matches just at approximate position 120. It may be that other positions also corresponded to the high BERs. However, in order to try and match the errors and delays the

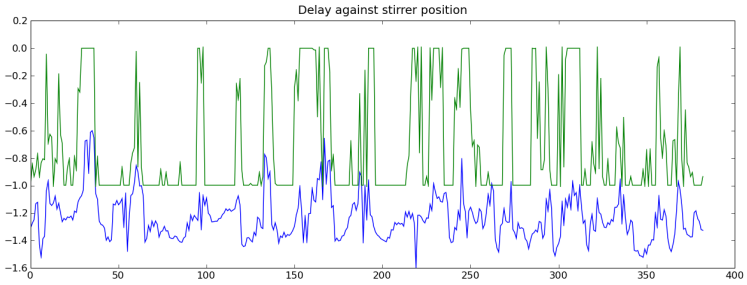


Figure 5.20: Bit error rates overlaid with delays (3x3 aperture)

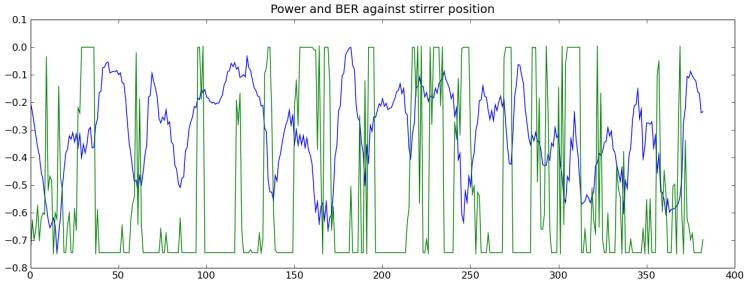


Figure 5.21: Bit error rates overlaid with channel power (3x3 aperture)

delays had to be calculated in the relevant band and the use of a narrow 6MHz band vastly reduces resolution in the time domain because it means the signal is being convolved with a much wider time response. The issue may be seen in figure 5.22. Therefore these results are not precise and it would not be expected that all delays would match up with all errors even if the delay was the only cause. In fact what is seen from these graphs is that when operating in an environment that has significant attenuation but also significant delay spread then both contribute to the failure of a communications system through inter-symbol interference and low signal strength. An additional point worth mentioning is that both these factors may be interconnected in the context of synchronisation, where if a system is coherent but the signal is heavily attenuated and with phase shifted I and Q channels then carrier and symbol synchronisation will be challenging.

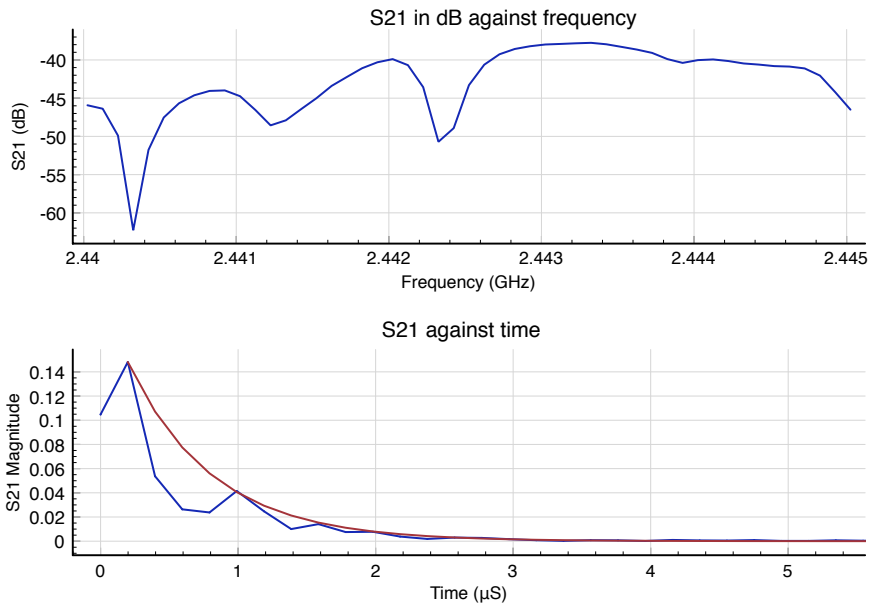


Figure 5.22: The limited number of points available calculating a time response decay constant

### Histograms

So that the probability of a failure may be calculated the statistics of the channel are required. We shall also look briefly at the form of the delay statistics because it has been seen again in these measurements that the delays are an important cause of bit errors and knowing the delay statistics could be useful in understanding how different wireless systems might perform. As before, initially we shall look at the statistics in the chamber when it is empty; the voltage histogram, along with a Rayleigh fit line is presented in figure 5.23

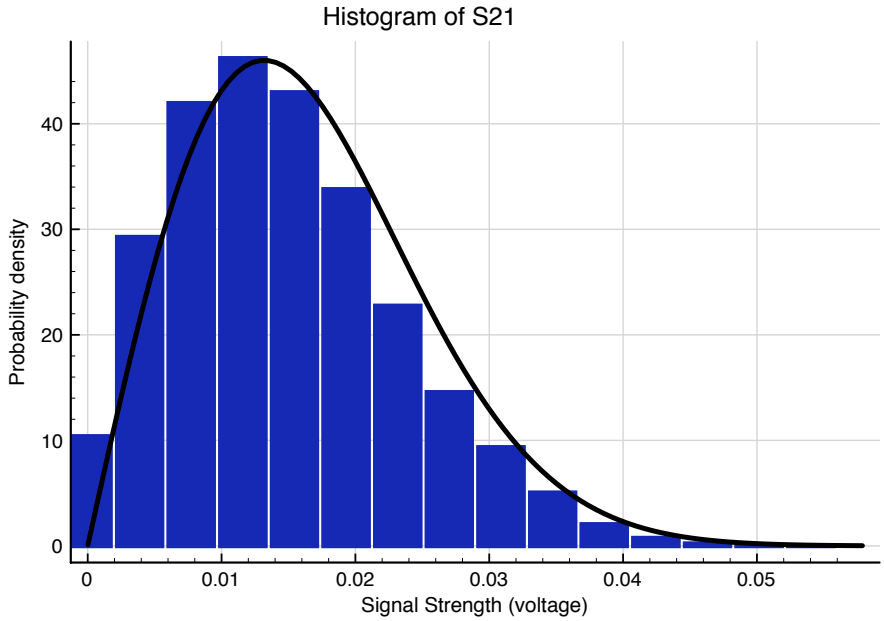


Figure 5.23: Histogram for signal between two antennas, one in the main reverb chamber, one in an open nested chamber

In this case, without any need for tests of fit it is very clearly Rayleigh which is what one would expect. Next it is plotted for the 4x4 and the 3x3 apertures in figures 5.24 and 5.25 with the latter looking like a much weaker Rayleigh fit than the previous two. By the 3x3 aperture the distribution is no longer clearly Rayleigh and the statistics appear to be moving to a double Rayleigh distribution. Finally the distributions for the 3x3 and 1x1 apertures are plotted in figures 5.26 and 5.27. The distributions closely resemble the double Rayleigh distributions discussed at the beginning of the chapter. There doesn't appear to be any problem here with the box/aperture combination leading to distinct resonances in and distorting the statistics.

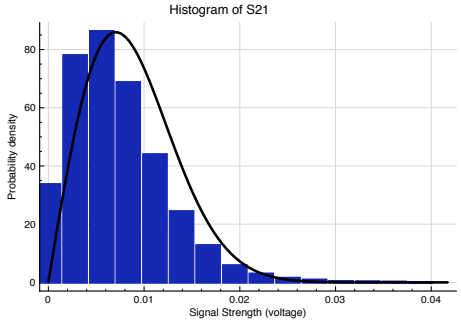


Figure 5.24: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (6x6 aperture)

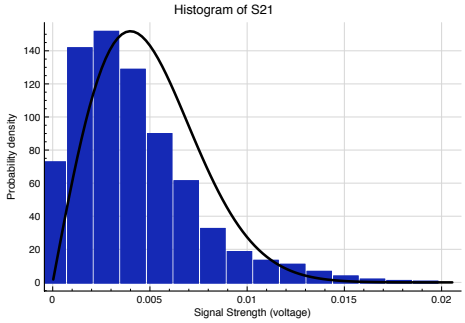


Figure 5.25: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (4x4 aperture)

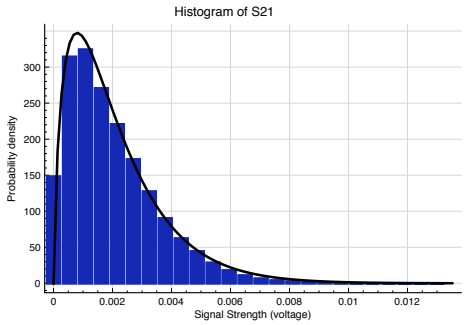


Figure 5.26: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (3x3 aperture)

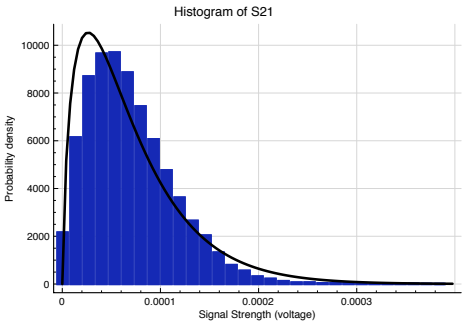


Figure 5.27: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (1x1 aperture)

### Delay Histograms

Next time histograms are presented giving the distribution of the delay spreads. It is interesting to see how these look so one might have an understanding of the range of values that can occur. Histograms for the open aperture, 3x3, 2x2 and 1x1 aperture are provided in figures 5.28 to 5.31 corresponding to the higher Q setup. These are nearly triangular in shape with slight tails on either side. Their distribution function is uncertain but it can identified that 95% of the delays are contained within a deviation of approximately 60 to 66% of the mean offering some limits on the delays that might be encountered over many stirrer positions.

It was noted in these and other measurements (not shown) that there appeared to be a slight skew in the results, particularly for the smaller apertures when compared with a single chamber, in the sense that they distribution was more heavily weighted to smaller delays. Why

this might be is not entirely clear and would require further research into the delay statistics.

Exactly what the distribution function is not certain but there are a few ways one could come up with an approximation. One might approach the analysis from the perspective of set of rays arriving at the receiver where each ray has a number of reflections and an amplitude dependent on the path length. Once a decision has been made about the distribution of path lengths we could calculate the mean delay as a function of this. Alternatively one could consider the Q-factor of the chamber and its distribution since Q is closely proportional to $\tau$.

The Q per stirrer position is calculate by taking the average of the power over stirrer positions (see chapter 3):

$$Q=\frac{16\pi^2 V}{\lambda^3}\left\langle S_{21}^2\right\rangle \tag{5.7}$$

where we are just specifying an exponential distribution in the form of 3. Now we know that distribution of power is:

$$f(P_r)=\frac{1}{2\sigma_I^2 R_r}exp\left(-\frac{P_r}{2\sigma_I^2 R_r}\right) \tag{5.8}$$

meaning that when we calculate the Q we are doing

$$Q=\sum_{i=1}^{N}\frac{1}{2\sigma_I^2 R_r}exp\left(-\frac{P_r i}{2\sigma_I^2 R_r}\right) \tag{5.9}$$

Perhaps then these results show a distribution approaching a Gaussian given there are a sum of (partially) independent identically distributed random variables. Generating a Gaussian on a computer with 200 to 300 samples also gives a histogram very similar to the figures shown adding further weight to the likelihood this is Gaussian. If this is the case though then any skew in the graphs is more likely to be a fault of the measurement or delay spread calculation. Generation of a double Gaussian suggests that this is also just a Gaussian and has no skew. For the cases of low Qs or the small chamber there will be wider bandwidth modes and less modes so some of the assumptions shouldn't hold as well but similar shapes were observed. Clearly the argument presented here is far from rigorous; there was not a great deal of time to look at this and so these delay statistics are presented purely as an introduction for the interested.

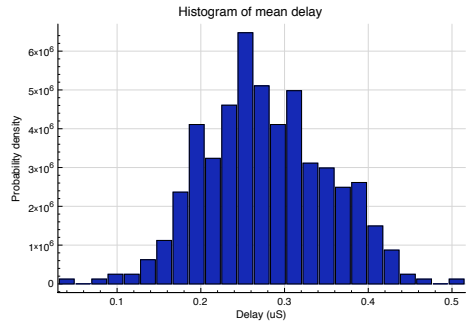


Figure 5.28: Histogram for signal between two antennas, one in the main reverb chamber, one in an open nested chamber (open aperture)

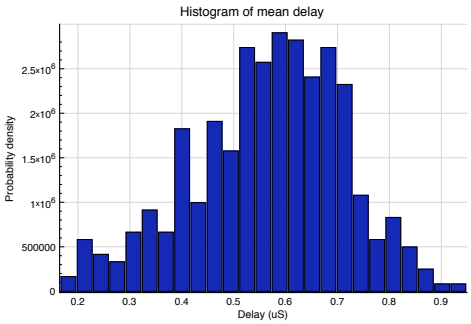


Figure 5.29: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (3x3 aperture)

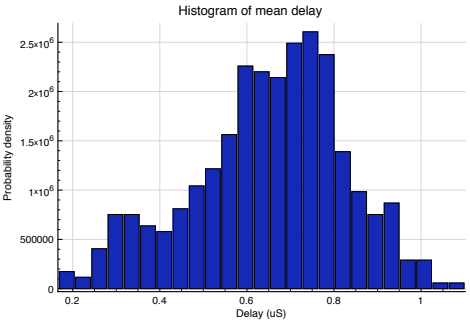


Figure 5.30: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (2x2 aperture)

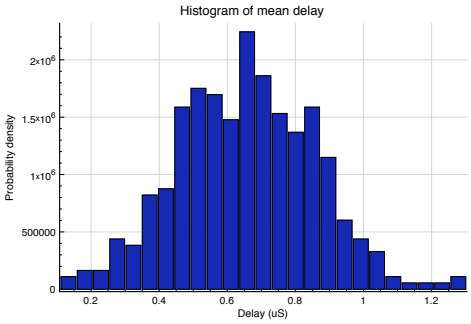


Figure 5.31: Histogram for signal between two antennas, one in the main reverb chamber, one in a nested chamber (1x1 aperture)

## 5.3 Wing

The focus of this PhD thesis is on aircraft and UAVs, not metal boxes and although more theoretical shapes lend themselves well to simulation and analysis this next section describes a more aircraft like measurement environment. In order to better predict the performance on a small UAV a metal structure was specially constructed that loosely resembled a section of a wing. This was done partly in the absence of a real system to do certain tests on but also to provide a new geometry not dissimilar to what could be found in an aircraft. An image of the wing is provided in 5.32 This is 3m in span and approximately 1m in chord and so is much smaller

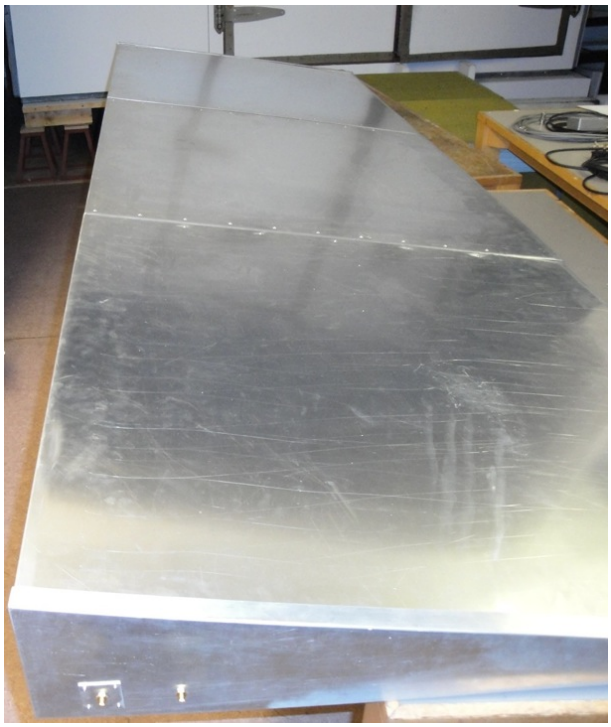

Figure 5.32: York wing like structure

than a section of a military fighter jet aircraft but could be found in an autonomous vehicle, for example. It might represent a cavity behind a fuel tank. The construction was from aluminium and it was made in 3 parts so that it might be divided into a number of cavities. Holes were placed in the top to allow stirrers to be placed in the wing just as they have been used in the coupled boxes but in this case with the stepper motors on the outside. These stirrers were necessarily much smaller than the main chamber and the small chamber stirrers and are shown in figures 5.33 and 5.34. Nonetheless, they were of comparable size relative to the wing sections with the only problem being the lack of a third stirrer in the middle section, where the

other two will have had little effect. The wing was not sealed and there were small openings of up to 1mm at the joints and at the ends. In the construction there bolts and screws had 3cm between them. Unlike the reverberation chamber and the boxes this was a far more leaky cavity like the sort one might find in a real aircraft where there are many openings between different sections. For example, a real aircraft avionics bays would require fans and openings for cooling, and across the wing there might be hydraulic systems. The empty Q of the wing was of order 1000 calculated very approximately from S21 measurements and its volume and also from time response measurements.

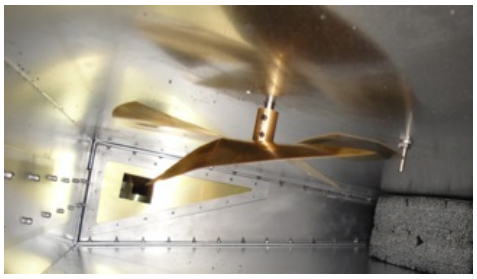

Figure 5.33: Left side wing stirrer and aperture

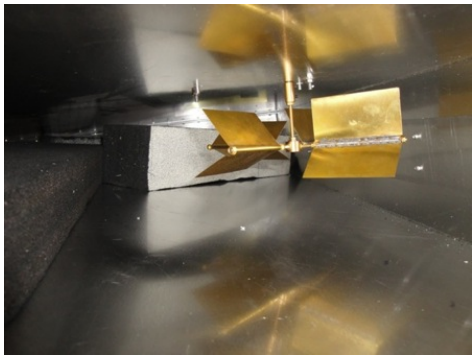

Figure 5.34: Right stirrer in wing section

## 5.4 Results

The 3 aircraft sections were divided up into 2, using metal sheets as with the boxes, but this time with a mixture of square and rectangular apertures. 80 stirrer positions were used which was 40 rotations of stirrers in the outer cavities. Stirrer rotations were alternated and the middle section was not stirred. Antennas were placed at each end and so sometimes there was a direct path if there was no absorber to block this. The measurement configurations and their associate BERs are given in figure 5.5.

There appears to be very little variation in terms of BER between the different channels, certainly nothing statistically significant. Examining the delay spreads also given in figure 5.5 then again there is very little difference here between each case; it would appear as though the cavity is always functioning as a single cavity rather than a few. Our wing environment would appear to be much more benign an environment that the large reverberation chamber.

The case of the horizontal slot demonstrates double the BER and more positions in error

| Setup (cm) | Delay | BER | no > 0.1 | no > 0.5 | no > 0.1 |
|---|---|---|---|---|---|
| 25x10 | 0.29 | 0.0125 | 1 | 1 | 1 |
| 25x10 | 0.27 | 0.01 | 1 | 1 | 1 |
| 7x5 | 0.3 | 0.001 | 2 | 0 | 2 |
| 5x4 | 0.26 | 0.025 | 2 | 2 | 2 |
| 3x3 | 0.32 | 0.018 | 2 | 1 | 2 |
| 8x0.5 | 0.32 | 0.06 | 6 | 5 | 6 |
| 0x0 | 0.32 | 0.025 | 2 | 2 | 2 |

Table 5.5: Aperture sizes and BERs for wing cavity measurements including number of positions out of 80 with a BER above a threshold. The final column is for PER.

which is much greater than the others and so this example will be looked at in more detail. A plot of the BER is given in figure 5.35 showing the most of the errors come in a group in the stirrer positions around 70. There is no problem with signal energy in this instance, in fact for all these wing measurements the energy was high (in the ZigBee scale it is typically between 25 and 35).

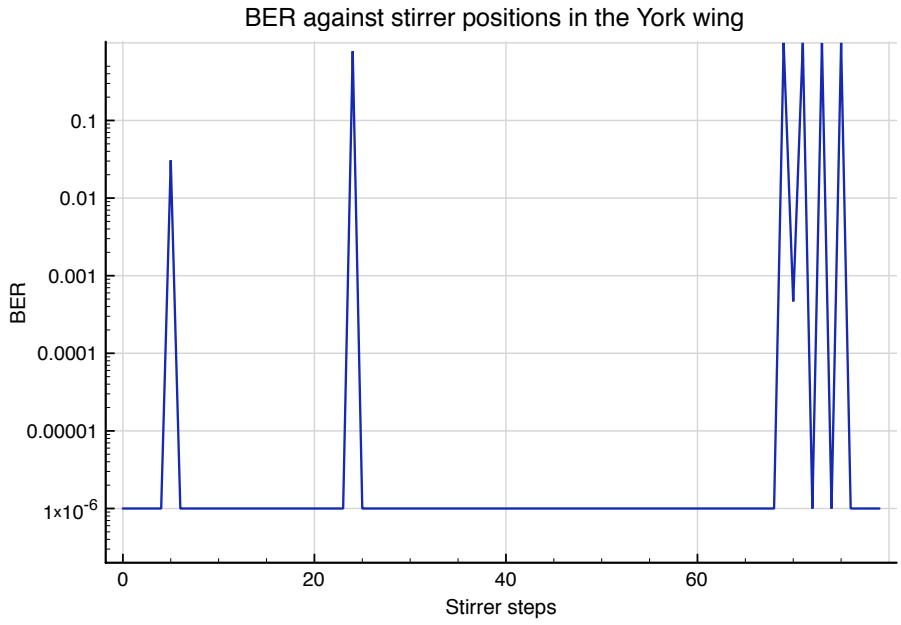


Figure 5.35: Variation of BER with stirrer position

As done earlier the BER is plotted against the delay spread in figure 5.36 where it is clear that the delays are the primary cause of the high BERs because the peaks correlate. However

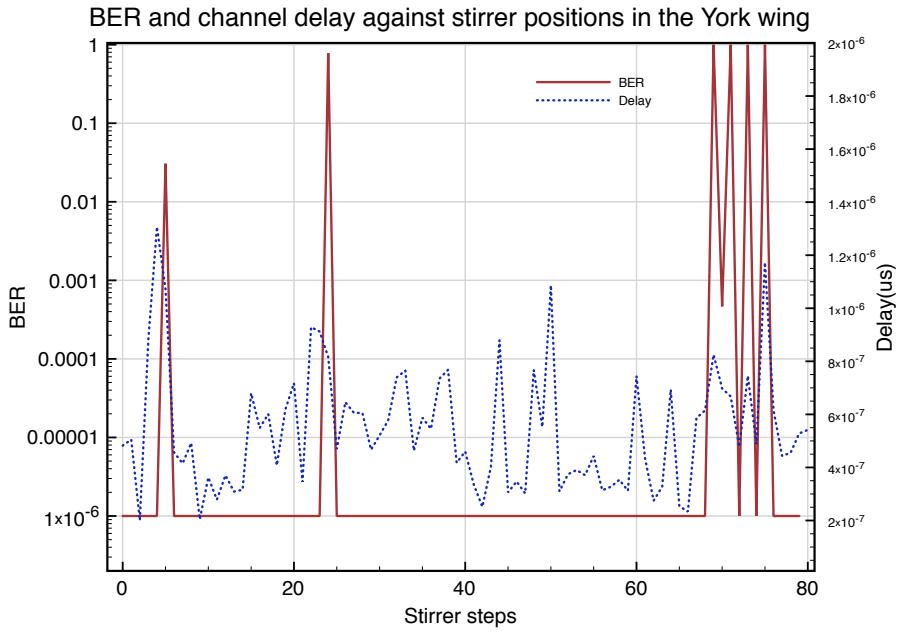


Figure 5.36: BER overlaid with delay spread

there are also higher peaks with no delay and there is some unpredictability in this. As before part of this comes from the fact that the delay spread calculation is not totally accurate and what will happen when delay spread is high is unpredictable.

### Statistics in the wing

Moving back to the topic of channel statistics histograms were also plotted for the coupling between the wing cavities and 4 of these are plotted in figures 5.37 to 5.40 What this shows us

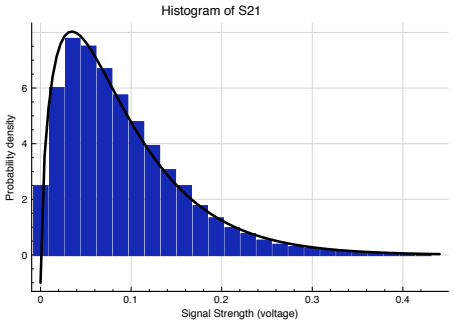


Figure 5.37: Histogram for no aperture (25x10)

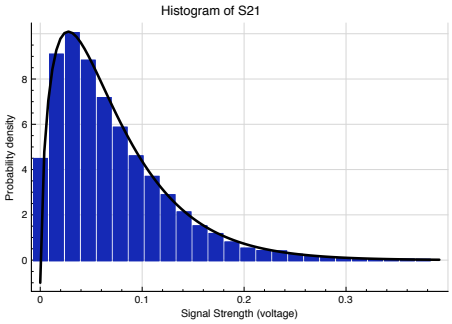


Figure 5.38: Histogram for aperture 4x5

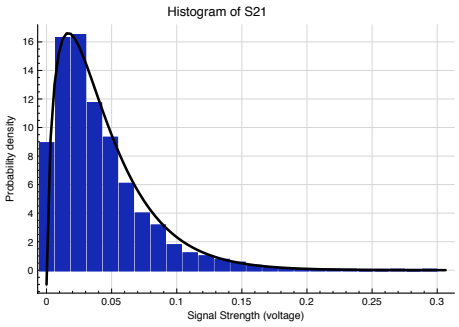


Figure 5.39: Histogram for aperture (0.5x8)

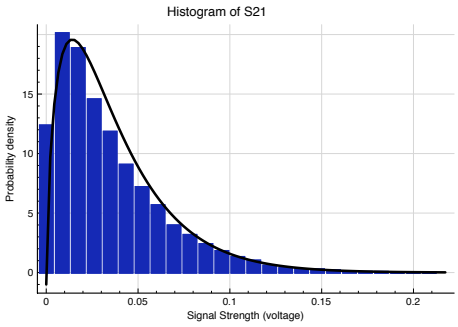


Figure 5.40: Histogram for filled aperture (solid plate)

is that the reason for the errors and delays being so consistent (i.e. not changing with aperture size) is not that the channel was acting as one cavity but that it was acting as multiple cavities in all cases. It appears in this case that the 25x10 aperture size is small enough to allow the individual sections to act as separate cavities leading to double Rayleigh distributions for all of the apertures. This is likely due to some polariazation in the E-field which has not been seen in previously measurements which would also explain the dramatic change for the aperture only 0.5cm in height. The other interesting point is that even with a solid sheet same distribution is found and if signal strength were plotted it would be found to be high(around 20 on the ZigBee scale) which is explained by the point made earlier: the wing and its dividers are not sealed in an electromagnetic sense because gaps due to inaccurate manufacturing (and space between

screws) allow significant propagation between sections. The fact it has not reverted to Rayleigh suggests only one or two of these gaps was dominant.

#### 3 cavity measurements

Additional measurements were carried out where 2 apertures were used rather than 1 and where absorber was varied. Both stirrers in each end of the wing were turned a full rotation over 40 positions each, 80 in total alternating between the two stirrers. Two aperture sizes were tried and were applied to both points in the wing so there was always 2 5x4 or 2 3x3cm apertures. In this case the aim was simply to show how the BER might vary with 2 variables, Q-factor and aperture size. These measurements are summarised in table 5.6.

What is found is that it is that the absorber leads to a PER typically lower than before with 1 exception where it reaches a 9% average. There is no obvious reason why this result has occurred and it further underlines the unusual results one can find in a random reverberant channel. In 5 of the 8 measurements the PER was 0 at all stirrer positions including the non direct case which might be expected to have poor performance compared to the other measurements. Looking at the differences between measurements with no absorber and a lot of absorber no difference is seen with both having incidences of more than 5 positions with greater than 50% BER. It should be remembered that one reason for less difference between measurements with high and low absorber configurations is that the cavity was already very leaky compared to the reverberation chambers and so absorber had far less of an effect. An observation made as these measurements were carried out was that the absorber configuration appeared more important than in the reverberation chamber and the results were less repeatable. Due to the taper there is not an even power distribution throughout the structure as is found in the reverberation chambers. The statistics are not shown here and add nothing new. They are half way between a Rayleigh and Double Rayleigh. There is nothing that is obviously triple Rayleigh but this is likely because was no stirrer in the central cavity which would be needed to observe the triple Rayleigh statistics.

## 5.5 Conclusions

This chapter started with a basic model of coupled cavities and has shown circumstances where it can be and can't be applied to BER and channel response measurements for small

| Setup | BER | no > 0.1 | no > 0.5 | no P > 0.1 |
|---|---|---|---|---|
| 5x4, HighAbs, Direct | 0.0 | 0 | 0 | 0 |
| 5x4, HighAbs, Non Direct | 0.000856 | 0 | 0 | 0 |
| 5x4, MedAbs, Direct | 0.0 | 0 | 0 | 0 |
| 5x4, NoAbs, Direct | 0.062 | 5 | 5 | 5 |
| 5x4, NoAbs, Direct | 0.0 | 0 | 0 | 0 |
| 3x3, HighAbs, Direct | 0.069 | 7 | 6 | 12 |
| 3x3, MedAbs, Direct | 0.000036 | 0 | 0 | 0 |
| 3x3, NoAbs, Direct | 0.09 | 8 | 7 | 8 |

Table 5.6: Aperture sizes and BERs for wing measurements for Qs of approximately 1000, 600 and 300. Final 3 columns show [BER, BER, PER].

boxes, a double reverberation chamber system and a stirred wing structure.

It was found that when there are two cavities and a dominant coupling point, they produce a double Rayleigh distribution for the S21 parameter, which is as expected. One can only see this when both cavities are stirred as was done for the small/large chamber system and the the wing. These statistics, as with the single Rayleigh case, can be interpreted in one of two ways. They provide the probability of finding a transfer function between two antennas, placed in two random positions, in two aircraft bays or they provide the distribution that will be sampled in a slow time varying environment like two cavities in a flexing wing. This distribution can cause even more difficulties than Rayleigh because it appears to show slightly more kurtosis as values less evenly spread and so it may produce more unpredictable results. When designing a communications system this is the opposite of what is wanted, where consistent performance would be ideal.

As well as affecting the statistics the apertures also cause attenuation and we have seen that this has minimal effect until the aperture becomes much smaller than half the wavelength. It was found that the signal only attenuates to a point where BER goes up significantly where the aperture size is between 1cm and 3cm in each dimension at 2.4GHz. The reality is that in modern aircraft such as the Eurofighter or UAV demonstrators seen by the author, many apertures will be larger than this and allow cables, rods etc though the structure but there will be exceptions and we need to be aware that even though distances are short in these environments there can be high levels of attenuation that lead to poor performance of a ZigBee system.

Moving to the channel delay, it was found that a two cavity system, where each cavity has the same Q, the system will have a greater delay than a single cavity system of the same combined volume although only where the system really looks like two cavities with one coupling point as was the case in the well controlled two reverberation chamber measurements. In a typical aircraft, this effect might again be less noticeable and again and although an interesting effect, the impact on PER is small compared with modest changes in Q in either cavity. Whilst not conclusive we have also shown that the delay statistics may have a Gaussian distribution.

PER and channel differences have been seen between the different coupled cavity types. In the wing it was seen that some of the reverberation chamber assumptions of chapter 3 no longer hold and there is field polarization in the direction of propagation, as shown by the differences between the horizontal slot and other apertures. We have also observed that very small boxes do not act as reverberation chambers because they do not have sufficient modes and that they may, in conjunction with an aperture generate a frequency response that is not flat on average. In the smaller cavities we have also observed much lower PERs than were seen in the previous chapter in a large reverberation chamber. This is interesting because it means that the PERs in an aircraft should not be as bad as assumed earlier. There remains the same problem of a small number of stirrer positions producing a high PER but the probability of finding these positions is a lower.

# Chapter 6

# Introducing noise and time variance to the channel

## 6.1 Introduction

Noise in communications system is well understood and the performance of systems based around ZigBee or wireless LAN is typically specified in terms of the minimum signal to noise at the receiver that can yield a particular BER or PER. For this reason the effect of a particular noise power on ZigBee or alternative systems has not been considered in this report. However, there is one aspect of noise that does require attention and that is the fading of the noise in a resonant environment. If there is no interfering source in an aircraft then the Gaussian noise in the receiver is the limiting factor. Electronics, actuators, deliberate jamming and other sources may produce additional noise. Where this is the case one might expect that the noise will fade just as the signal does. If this happens then there won't be just Rayleigh fading where the signal to noise ratio is modulated as one unit but something more complicated with two interacting Rayleigh distributions.

### 6.1.1 Aim

Given that the signal to noise will vary in a noisy, resonant environment, in a way that is unknown, our aim is to describe the SNR statistically. This was done experimentally by considering different levels of injected noise at differing Q-factors to cover the variety of scenarios one might find in different parts of different aircraft. Further, measurements were carried out of real system performance using ZigBee, just as in previous chapters. This will provide a picture

of what the statistics mean practically for a real system.

## 6.2 Measurements

To find out how signal to noise is affected in a resonant environment a similar setup was employed to that in earlier chapters, using the same aluminium wing described in the previous chapter. This had a Q of order 1000 (some delay spread measurements suggested a maximum of 4000) when empty and the use of absorbing foam could bring this down to order 100. As with any reverberation chamber measurement designed to show statistics, measurement time can be long and so a limited subset of possible measurements had to be selected. Therefore 3 different absorber levels were used (3 Q-factors) and the wing was left as a single cavity. Whilst it would have been interesting to see what happened with signal and noise going through 2 narrow apertures or the signal through 1, the noise 2 and other combinations, time did not allow this. The noise powers to be used were not fixed from the beginning of the experiment as these were varied for the ZigBee measurements to show the more interesting regions of performance. We are not so interested here in exactly what noise power causes a specific BER, rather we are interested in the range of powers over which high BERs can occur taking into account fading and time varying fading. The exact setup used for the measurements is described in the following subsection.

#### Measurement setup

The experimental setup can be seen in figure 6.1 and is summarised by the block structure in figure 6.2. Two small stirrer motors are used at either end of the wing to maximise the stirring all the way through the wing. Noise is injected into the wing at the same end as the transmit antenna which allows it to undergo fading with the signal and it is kept within half a wavelengths distance from the transmit antenna to allow separate fading and independent paths. The noise generator is broadband and fed into an amplifier. So as to avoid overloading the amplifier whilst allowing a high enough noise power, a 2.45GHz bandpass filter is employed which reduces out of band power before the noise signal is fed into the amplifier. In order to vary noise power the attenuation is varied because the amplifier is of fixed power. Measurements were under the control of a laptop computer and proceeded as follows:

1. Start continuous transmit of ZigBee packets

Figure 6.1: Setup for noise measurements

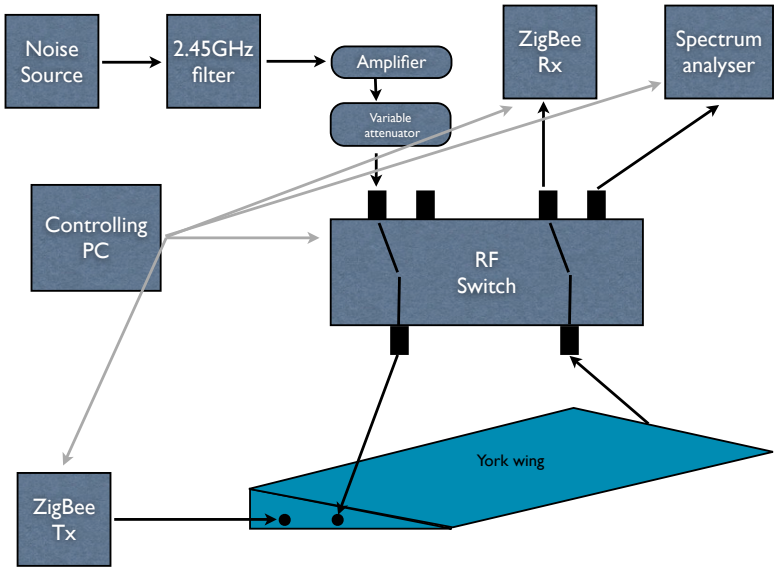


Figure 6.2: Block diagram of setup for noise measurements

2. Measure the spectrum using the spectrum analyser
3. Change the RF switch so that noise enters the wing
4. Measure the noise power on the spectrum analyser
5. Change the RF switch and transmit ZigBee packets to the ZigBee receiver
6. Move a stirrer (alternating between the 2) and return to step 1

These steps were repeated 80 times before changing the attenuation and repeating again. Care was taken to correctly measure noise, so that correct power densities were obtained with compensation for filter shapes and bandwidths. For the power in the ZigBee signal the measurement was not of a continuous signal, but a pulsed one. Therefore an averaging mode was used on the analyser and a constant (determined experimentally) was added to the magnitude of the spectrum during analysis as it was found that the power loss due to the duty cycle was constant.

### 6.2.1 Measurement results - correlation functions

To begin, consider the effect of the multi-path on the noise for a low Q configuration. Figure 6.3 is a plot of received noise in the wing over complete alternating rotations of two stirrers in either end. Noise power is measured in a band of 2MHz which could be considered a flat bandwidth and can be converted to a PSD. This is also approximately the width of the main lobe for an IEEE 802.15.4 signal spectrum. The Q in the wing is of order 1000 and so the noise fades just as S21 did in the NA measurements. The amplitude variations are very large and the noise is bounded between -120dB and -97dB which is equivalent to a power variation of 200. In the case of an aircraft, this would mean a major increase (or decrease) in the noise seen by or coupled into a device when the wing flexed or cargo was added or actuators moved. Alongside this plot is figure 6.4 showing the signal to noise ratio (SNR). This differs even more dramatically with a 45dB variation occurring as the stirrer is rotated.

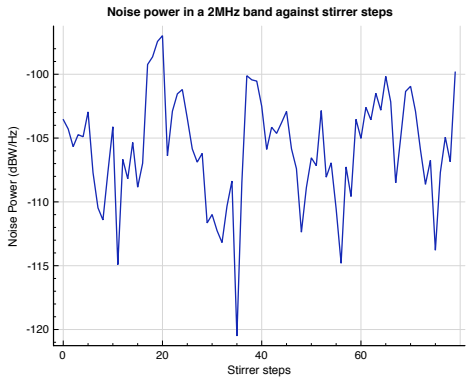


Figure 6.3: Variation of noise power with stirrer position for a Q of approximately 2000

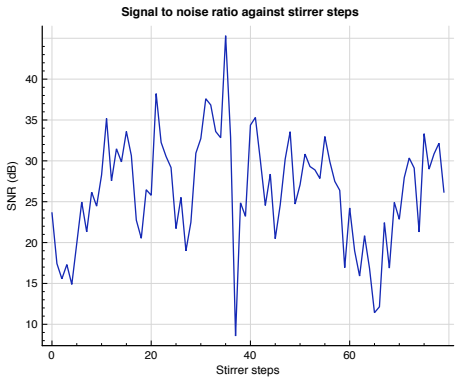


Figure 6.4: Variation of signal to noise with stirrer position for a Q of approximately 2000

An effective way to demonstrate the double effect of the channel resonance, is shown in

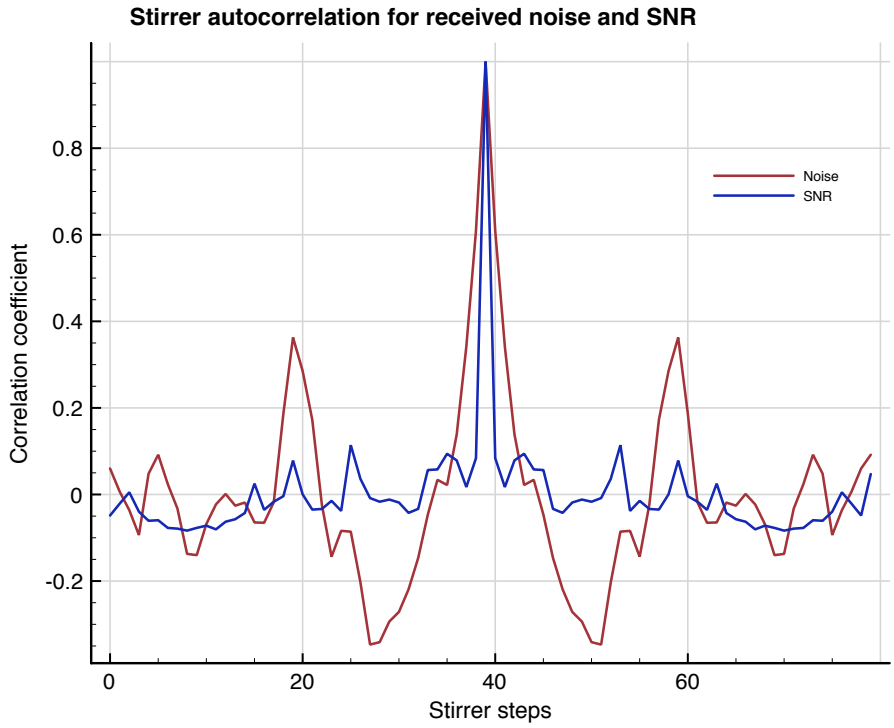


Figure 6.5: Auto correlation for noise and signal to noise for a Q of approximately 2000

figure 6.5 which plots the autocorrelation function for the noise power and the SNR overlaid on the same axis [1]. One can see that for both noise and SNR there is a very sharp peak at the central point which corresponds to no offset and this peak only lasts 2 or 3 offsets on each side which is expected given the stirrer step separations were intended to avoid highly correlated measurements. Comparing the two peaks it can be seen that the noise peak is about half the width of the signal to noise peak. What is happening here is that noise and signal are fading independently so there is no reason why as the noise increases the signal might not decrease. The antennas were deliberately separated to ensure that they emitted independent signals. What is seen here could be likened to an aircraft where a noise source from a high voltage switch is nowhere near close enough to a ZigBee transmitter to produce noise coherent with the ZigBee signal.

The following figures (6.6 through to 6.11) show equivalent results with additional absorber loading used to bring the Q down to approximately 600-700 and 300-400. The results are similar, in that there is a greater variation in signal to noise compared to just noise and a narrower peak for the autocorrelation function of the SNR. However, rather than the absorber reducing

[1] Plotting signal power and SNR would show the same effect

the effect as might be expected it appears to accentuate it. Although the noise autocorrelation function has a widening peak, the SNR peak hardly widens and so the disparity becomes much greater. As the frequency response flattens then one would expect the SNR peak to widen too. Instead it stays constant width, which can be explained by looking at the plots of just signal or noise because these still show close to independent movement between measurements, even for the lowest Q case. Even if the overall trend is upwards over a portion of the noise graph, there are still small variations which don't line up with the small movements in the signal power. Although the more dramatic changes in power of over 20dB are rarer for the more loaded case, they do still occur and we end up with a large maximum to minimum signal to noise difference of 40dB.

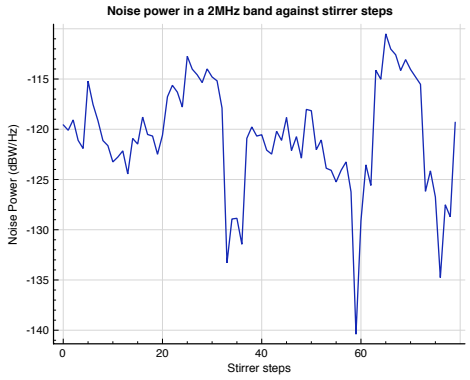


Figure 6.6: Variation of noise power with stirrer position for a Q of approximately 700

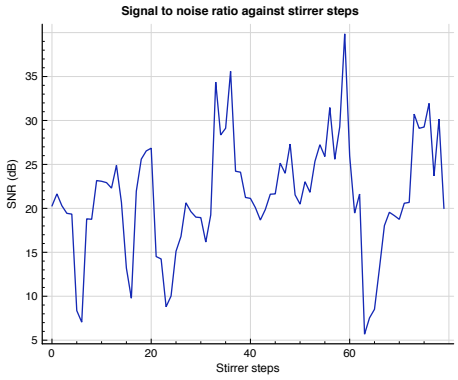


Figure 6.7: Variation of signal to noise with stirrer position for a Q of approximately 700

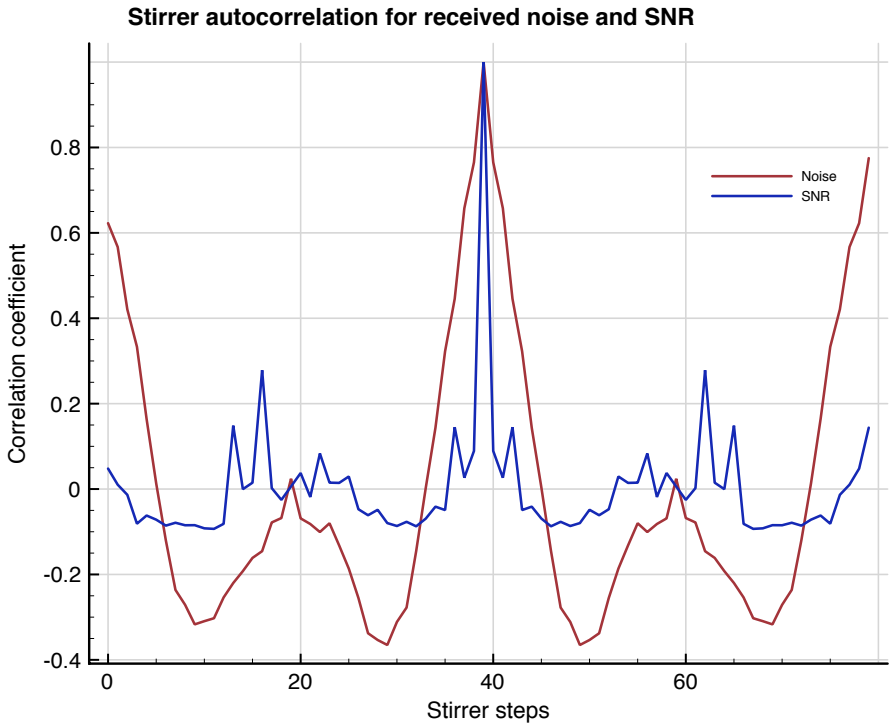


Figure 6.8: Auto correlation for noise and signal to noise for a Q of approximately 700

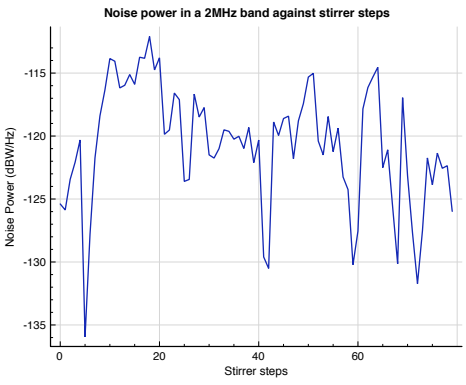


Figure 6.9: Variation of noise power with stirrer position for a Q of approximately 300

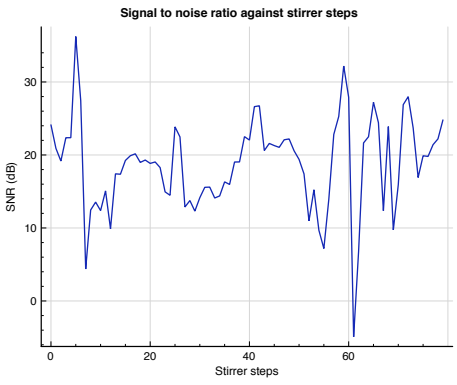


Figure 6.10: Variation of signal to noise with stirrer position for a Q of approximately 300

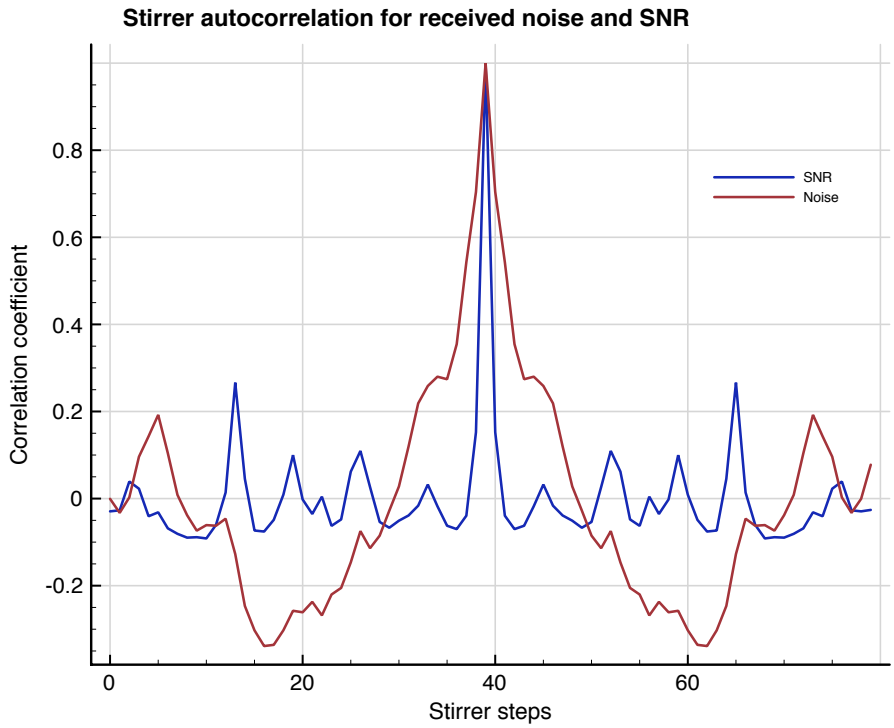


Figure 6.11: Auto correlation for noise and signal to noise for a Q of approximately 300

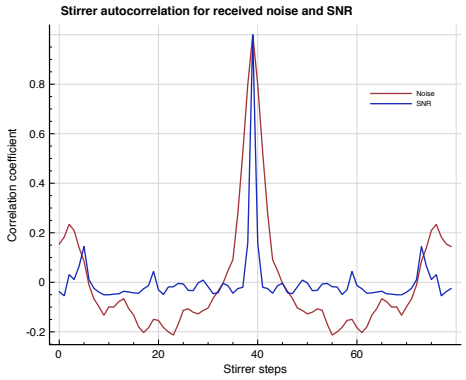


Figure 6.12: Auto correlation for noise and signal to noise for a Q of approximately 700

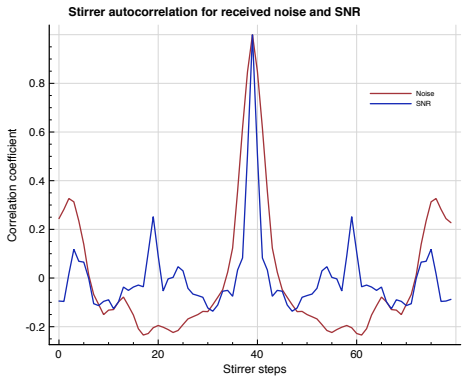


Figure 6.13: Auto correlation for noise and signal to noise for a Q of approximately 700 with 13dB attenuation

### 6.2.2 Results - distributions

The next section looks at the distribution of the signal to noise ratio as compared to just signal or noise. Given the results of the last section a standard Rayleigh voltage / Ricean power distribution would not be expected; both the signal and noise fade independently so one would expect to see a new distribution. To see if this happens histograms have been plotted in figures 6.18 to 6.25 corresponding to the lowest Q-factor out of the 3 main noise measurement setups. Histograms are given for signal power, noise power and signal to noise are given each time and the signal and the noise graphs would be expected to have the same exponential shape each time. All histograms are for linear power and, for the SNR graphs, 2 histograms have been plotted, one with 15 bins, one with 50, with the latter being there to show more detail, given almost all data is concentrated in the first couple of bins when only 15 are used. 3 noise powers have been selected spanning the range that was used (i.e. highest, middle and lowest attenuations) and so there are 3 different sets of histograms. Noise power was controlled with RF attenuators and the higher the attenuation quoted the lower the noise power. All were above receiver noise on reception and are used here just to provide a range of examples.

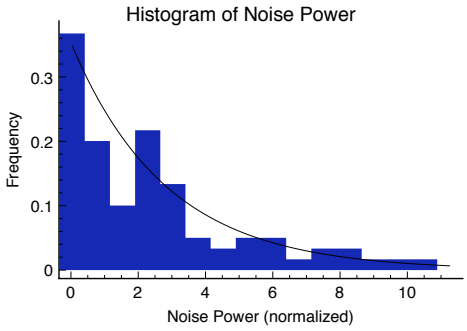


Figure 6.14: Histogram of noise power (16dB noise Attenuator)

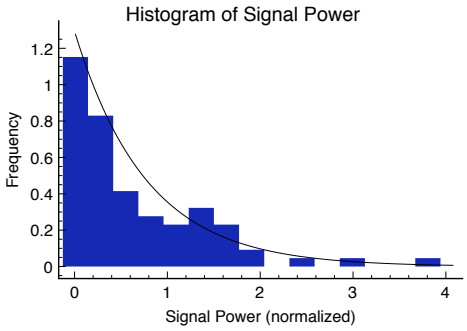


Figure 6.15: Histogram of signal power

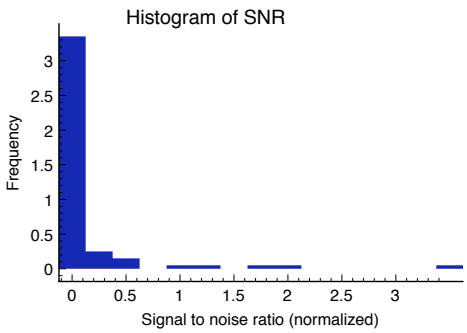


Figure 6.16: Histogram of Signal to noise ratio (16dB attenuation)

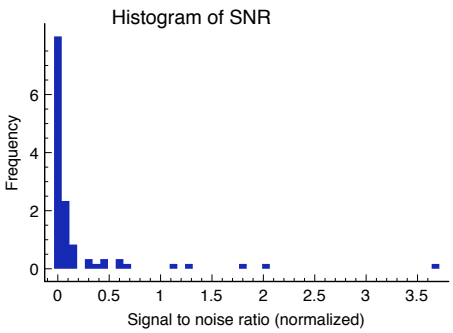


Figure 6.17: Histogram of Signal to noise ratio (16dB attenuation)

Before considering the SNR the form of the individual distributions is presented for the wing with our loadings. From the fitted lines we see that the signal histograms are likely to have an underlying exponential distribution which is as would be expected theoretically for an ideal reverberation chamber. In fact in this case the measurement is using a cavity of much lower Q than a reverberation chamber, loaded to bring its Q further down and with a far lower mode density but the power distribution is still exponential. There was a risk that this would not be the case with a low Q in a non uniform shape and the fact we are seeing this is useful in the sense it suggests we can assume a Rayeleigh/Exponential distribution in a wide range of bays and aircraft sections.

Moving on to the distribution of the signal to noise ratio, the plots show that this looks different, as expected. In figure 6.16 the SNR is almost entirely in one bin. Increasing the number of bins to 50 for the same data leads to the plot of figure 6.17. Since most of the 80 samples are contained in the first bin when using only 15 bins, it makes sense to increase the number significantly. There is now a distribution that looks similar to an exponential function but much steeper and with a few outliers.

So as to be sure of this shape the same graphs are plotted for 2 more cases where there is greater attenuation and less attenuation of the noise source and these are in plots 6.18 to 6.25, thus also allowing us to see the consistency of the distribution over different relative signal and noise levels.

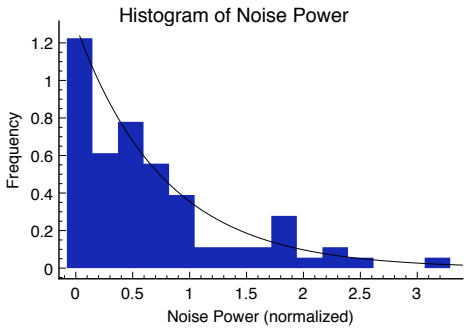


Figure 6.18: Histogram of noise power (22dB attenuation)

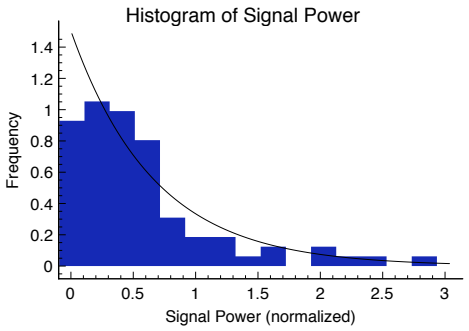


Figure 6.19: Histogram of signal power

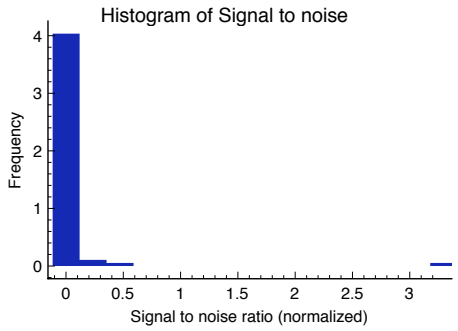


Figure 6.20: Histogram of SNR ratio (22dB attenuation)

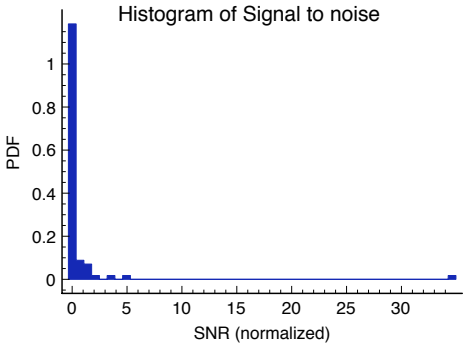


Figure 6.21: Histogram of SNR ratio (22dB attenuation)

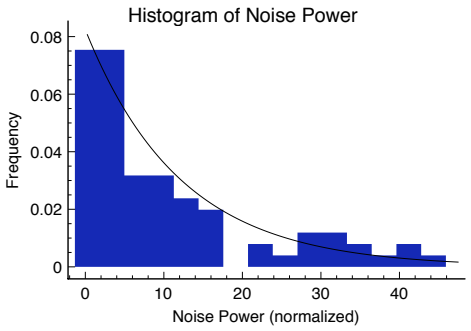


Figure 6.22: Histogram of noise power (11dB Noise Attenuator)

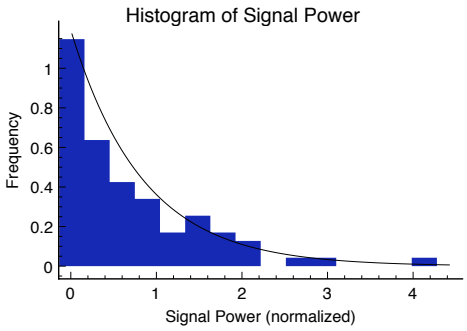


Figure 6.23: Histogram of signal power

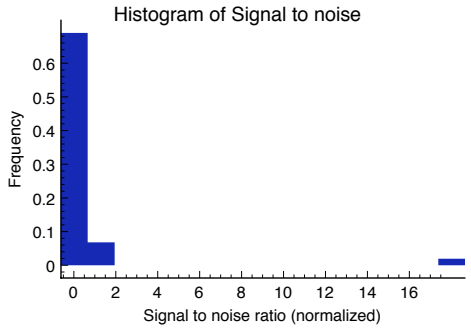


Figure 6.24: Histogram of signal to noise ratio(11dB attenuation)

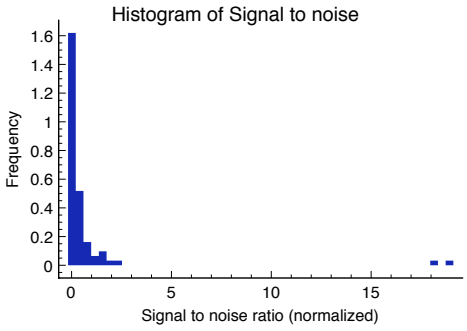


Figure 6.25: Histogram of SNR (50 bins)(11dB attenuation)

In all these figures the same curves can be seen for the signal, noise and SNR with the latter always having a very sharp fall off from the first 2 or 3 bins. One might expect that as the signal power becomes much greater than the noise then the statistics would be dominated by the signal. This is not seen here because we are looking purely at a ratio but, if there were other noise sources that the noise faded into then the distribution might return to a standard exponential. As far as performance is concerned then this dual fading effect is most important when there is a signal and a noise signal of similar magnitudes or within an order or two of, not when the signal will always be greater than the noise, or when the noise is below noise in the receiver electronics.

The results given so far demonstrate the performance of a system in a low Q environment where the Q is under 700. For completeness two cases where the wing has no absorber are considered. In this case the results should be closer to the theoretical as the channel will be less correlated between stirrer movements. Figures 6.26 to 6.29 show the noise power and SNR for two noise levels, so as to further demonstrate the reproducibility of the distribution.

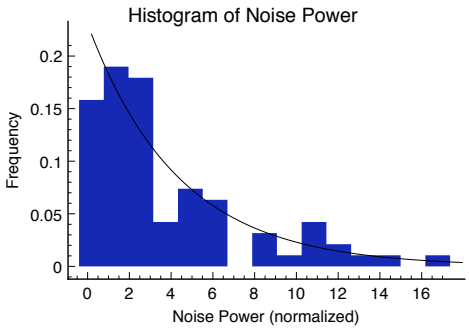


Figure 6.26: Histogram of noise power (16dB attenuation, no absorber)

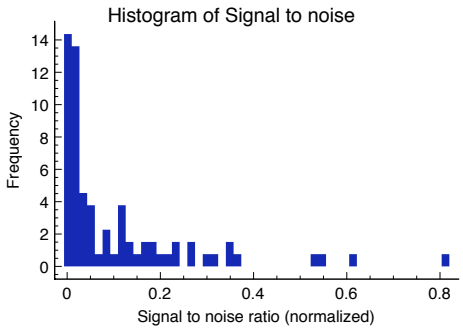


Figure 6.27: Histogram of SNR (15bins)(16dB attenuation, no absorber)

It would be expected that the SNR histogram would be similar to the low Q case and it is. For these two plots there are more data points to the right of the main peak than before which is likely due to the lower correlation between stirrer positions, allowing a higher number of independent samples in a rotation.

From a communications systems perspective the interesting part of the distributions is the very long tails. Although here they show extreme values of noise relative to signal, the same would be true if we plotted a noise to signal ratio. Where there is a coincidence of a very high noise value and a very low signal value then the measurement may occasionally be out on the

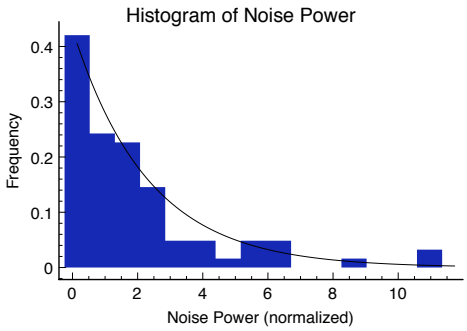


Figure 6.28: Histogram of signal to noise ratio(18dB attenuation, no absorber)

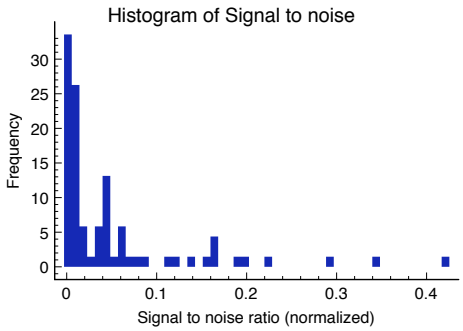


Figure 6.29: Histogram of SNR (50 bins)(18dB attenuation, no absorber)

ends of these tails. The danger is that a few measurements in an aircraft when static will only identify points in the first, highly populated bin. Unfortunately, it has been seen that this ratio may change up to a maximum of 45dB. The only way out of this problem is to ensure our band of operation has no in-band interferers, or at least no in band interference within 40dB of the signal power that is undergoing independent fading.

**Derivation of distribution**

With the shape of the SNR distribution known the next step is to try and derive it. It is possible to analytically derive the new distribution by considering how to combine the two exponential distributions. We shall work with these rather than Rayleighs because noise is described in terms of power and because exponentials are simpler to manipulate. Consider initially two exponential distribution density functions:

$$f(x) = \lambda_1 e^{\lambda x} \tag{6.1}$$

$$f(y) = \lambda_2 e^{\lambda y} \tag{6.2}$$

What is required here is the distribution of signal to noise and so an additional variable called $SNR$ is created where the aim is to find $f(SNR)$. Before working with the function of the random variables it is necessary to look at the joint distribution. One can say that the fading of the noise will not affect the fading of the signal. This section is deliberately looking at this case and in the experiments a distance between antennas was chosen to provide independence of the noise and signal. Therefore we can say

$$f(x,y) = \lambda^2 e^{-\lambda(x+y)} \tag{6.3}$$

where we have made $\lambda 1 = \lambda 2$ which is reasonable since the both noise and signal are using the same channel. With the joint distribution we can try to obtain the distribution of the signal to noise function of $x$ and $y$:

$$f(SNR) = f\left(\frac{x}{y}\right) \tag{6.4}$$

To do this we shall try and find the cumulative distribution and then differentiate. A cumulative distribution may be found by integration over a given region. If the new variable (i.e. signal to noise in this case) is denoted SNR then we wish to find the distribution $F(u) = P(SNR < u)$.

If one plotted this they would see that we need to integrate $x$ from 0 to $\infty$ and $y$ from 0 to $x/u$. We may solve this integral as follows:

$$F(u) = \int_0^\infty \int_0^{u/n} f(x,y)\,dx,dy \tag{6.5}$$

$$= \int_0^\infty \int_0^{u/n} \lambda^2 e^{-\lambda(x+y)}\,dy\,dx \tag{6.6}$$

$$= \int_0^\infty -\lambda e^{-\lambda(x+x/u)} + \lambda e^{-\lambda x}\,dx \tag{6.7}$$

$$= 1 - \frac{u}{u+1} \tag{6.8}$$

This can be differentiated to give:

$$f(u) = \frac{1}{u+1} - \frac{u}{(u+1)^2} \tag{6.9}$$

Plotting this and overlaying it on a histogram formed on a computer by using two perfect exponential distributions shows that it is indeed the correct fit. 8000 samples were produced and 100 histogram bins were used for 6.31 where the items in all but the first bin are too few to see in this plot. Figure 6.30 splits the first bin of 6.31 into bin to show the actual shape and the fit. What is seen first of all here is that the there is a very rapid drop off in probability of SNR and then a relatively slow drop off which means that quite extreme values are possible. In fact it is not possible to calculate an expected value for the signal to noise as the integral does not converge and grows logarithmically with the limit of the dependent variable, signal to noise.

### 6.2.3 BER

Through the previous sections the theoretical aspects of noise in a reverberant environment have been discussed. It remains to be show how this affects the BER for a ZigBee system.

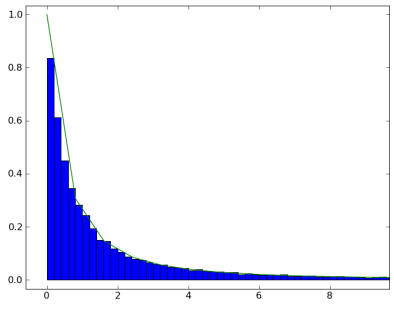


Figure 6.30: Computer generated histogram of signal to noise ratio zoomed in on first bin

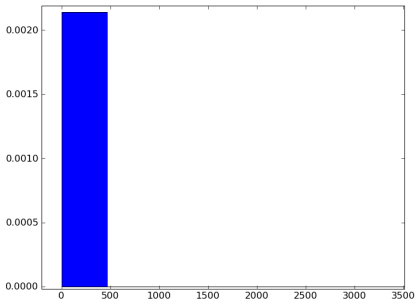


Figure 6.31: Computer generated histogram of signal to noise ratio

Figure 6.32 shows how ZigBee performs under noise in an anechoic environment first , in order to show us how the receivers would perform under direct injection of noise. The theoretical graph for MSK is also shown here which will be different because of various factors, the most important of which is that a packet based system was used.

Figure 6.33 shows the equivalent for the wing, plotting average BER for 3 different Qs. For these the shape of the drop off in BER is similar to the anechoic case but there is a lot more variation because of the reverberant environment. The BER and SNR given is the average over stirrer positions in the wing and this shows that a much greater average $Eb/N0$ is needed to drop the BER compared to the direct injection case. Intuitively one would expect a higher signal to noise is needed because the extreme values of SNR where noise is very much higher will take the BER 0.00001 to 0.1. Figure 6.32 reminds us that at the transition point from high to low BER the change in SNR is only a few dB. Therefore, with the tens of dB of variation for the reverberant channel it is easy to move back to a BER of 0.1 on certain stirrer positions and this BER will be mask the much lower BER at other stirrer positions.

The nature of the errors is explored in figures 6.22 to 6.37 which show the number of positions at which the BER is greater than 0.1 and is 0.5 and greater for the range of $Eb/N0$s. The first thing to note is that the lines for BER>0.1 and 0.5 are nearly overlaid on one another which tells us that when errors do occur they are typically very high providing a probability no better than chance of recovering the correct data. As with the noiseless case, there is often a 0% or 100% pattern for BER performance. In all the plots the lower noise levels are providing 8 or more stirrer positions with a BER reaching 50% which is 10% of all positions measured. This is for an average SNR that would give a BER of less than $10^{-6}$ if there was only Gaussian noise in a

static environment. Previously, when looking at coupled cavities it was found that the number of positions was sometimes 0 and at most 2 with 50% BER and so this must be a consequence of the noise in combination with fading. The change in absorber does not appear to have an effect with the same pattern seen in each graph. Over the stirrer rotations both signal and noise are changing independently with the same variation in ratios and different absolute levels.

The other notable feature of these graphs is that compared to a traditional graph of SNR/BER the equivalent drops in BERs are over wider SNR ranges - there is no reduction of BER of 2 or 3 orders of magnitude over just a few dB. Finally, note that in two cases there are 2 positions with BER at 50% for an SNR of 30dB which may or may not be due to the reverberant effects alone.

Whilst this is only one system in one environment, it does tell us that noise at any given power, in a reverberant environment, has a significantly worse impact on system performance than noise in an anechoic environment.

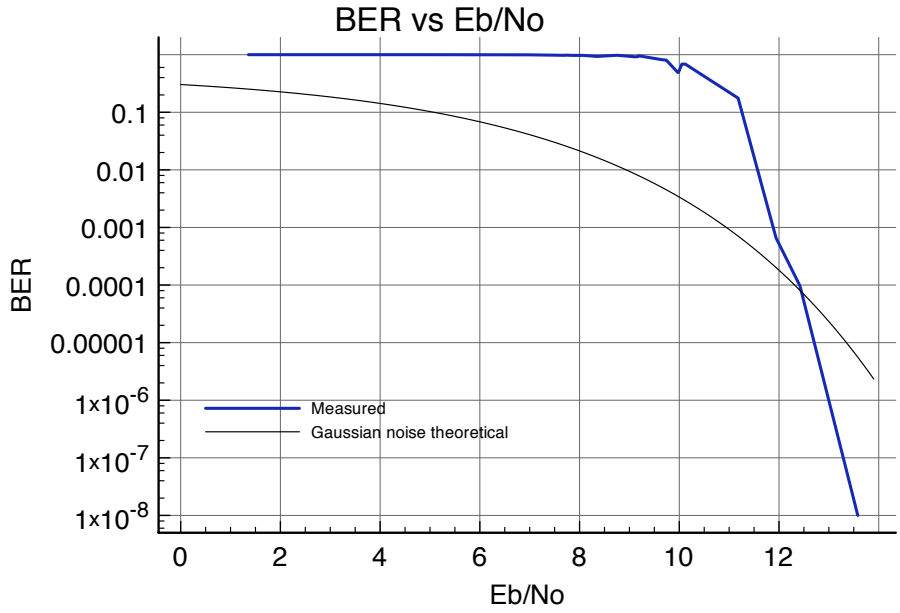


Figure 6.32: BERs against signal to noise ratio for ZigBee in an anechoic chamber

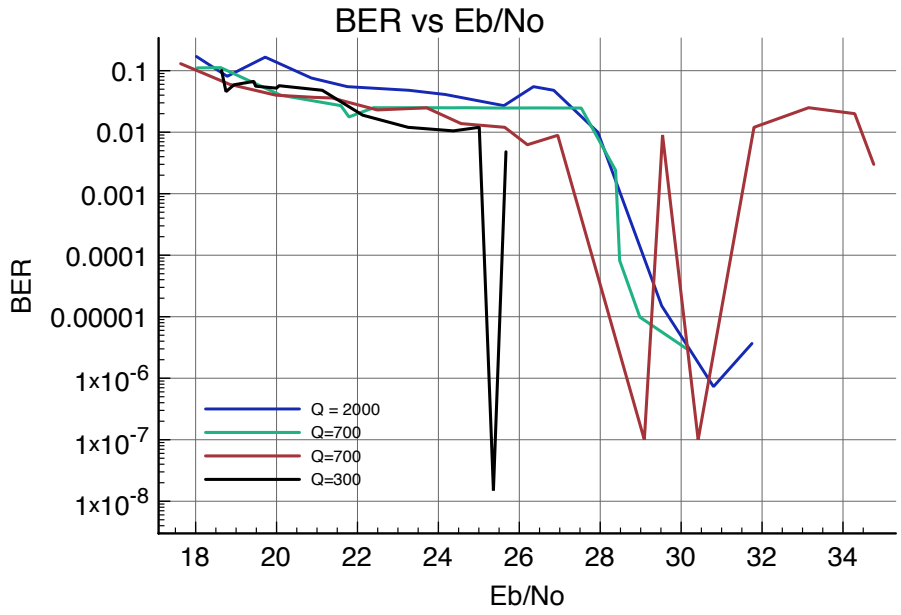


Figure 6.33: BERs against signal to noise ratio for 4 different wing configurations and 3 significantly different Q-factors

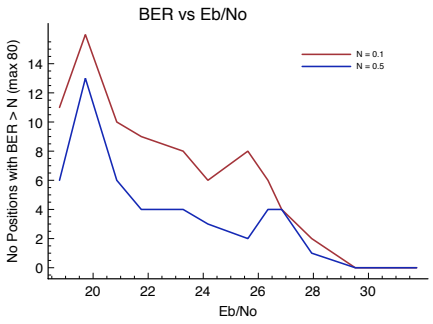


Figure 6.34: Number of high error positions against SNR for a low Q ( 300)

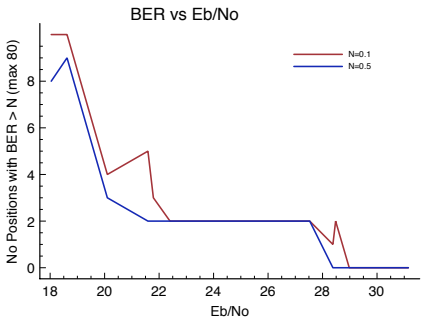


Figure 6.35: Number of high error positions against SNR for a medium Q( 700)

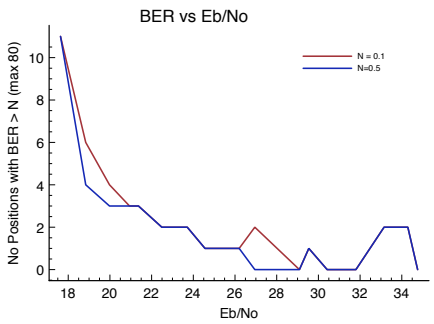


Figure 6.36: Number of high error positions against SNR for a medium Q( 700)

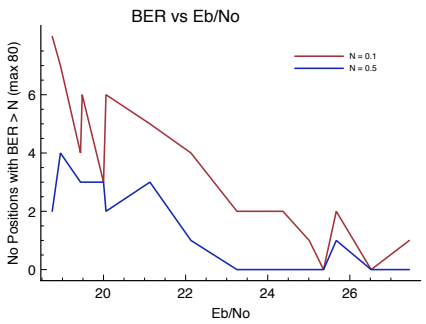


Figure 6.37: Number of high error positions against SNR for a high Q( 2000)

## 6.3 Time variance

Alongside noise, the other factor which may degrade performance is time variance. In any communications textbook the issue of fast and slow fading is discussed, for example see [60]. Up to this point slow fading signals have been considered and where the transition from flat frequency to frequency selective channels occur. We have imagined that as symbols are transmitted then the channel stays constant with a given phase and magnitude response over the time of a symbol. It might be expected that this would be the typical channel behaviour for an aircraft, but it nonetheless warrants a closer look to see what happens for a moving channel. In a modern avionics bay there might be metal fans to circulate air, wings will flex during a flight, engines and flight will cause vibration and any sensors that must transmit through the engine will see a rapid time variation.

The question then comes up as to whether this movement affects performance and whether the channel is fast or slow fading. Fast fading is the case where Ts » Tc, i.e. where the symbol time is significantly higher than the channel coherence time. In other words the amplitude and phase of the channel are changing over a symbol period which causes two problems:

- Symbol distortion
- Doppler spread

These can be almost the same thing, in the sense that time distortion is seen as doppler shift in the frequency domain. Rappaport [43] shows graphs demonstrating that doppler shift can put a limit on how low BER can go as noise drops. The well known problem of a mobile phone on a train moving away from a transmitter as a symbol is sent is not our main concern because in a reverberation chamber, for any mode that is delayed, another is advanced so the situation is subtly different. In a resonant cavity there is a highly frequency selective channel fading quickly because there are large variations in phase and amplitude between stirrer positions. When the stirrer is moving quickly there will be a large $d\phi/dt$ distorting the signal. In such circumstances it is impossible to predict BER rates analytically (just as for frequency selective only channels) and all one can say is that performance will degrade. In high end systems such as ADSL or wireless LAN one can navigate these problems with coding and processing.

One of the motivations for looking at the case of the moving stirrer was some results seen early on into the first year of measurements. ZigBee were done manually rather than under computer control which was time consuming so it was hoped that the time average could

be used in place of the average over a number of stirrer positions. Upon trying this it was repeatedly found that the two did not match up leading to the discovery that the movement, although not fast was affecting the wireless receiver performance.

### 6.3.1 Results

First of all we examine the case of the main reverberation chamber for a very high Q channel before moving to low Q channels. Figure 6.38 plots the time averaged BER over 10000 packets for a range of stirrer speeds. The chamber was loaded with only one piece of AN79 absorber and so the Q was in the region of 10000 which is as high as well shall look at. The graph clearly shows the movement having an effect on the error rate because as stirrer speed increases so does the average error rate. What is unexpected is the shape shown here whereby the BER rises dramatically from close to the per stirrer average before rolling off and appearing to approach 0.3 asymptotically. It was not possible to go beyond the maximum speed shown on the graph due to the mechanics of the stirrer motor. Graphs for some higher loadings are plotted in figures 6.39 which show the error rate going up with speed for the lowest of the loadings (2 absorber pieces) but no BER variation with speed for lower Q-factors. The variation with speed for the aluminium wing is plotted in figure 6.40 and also shows that there is no increase in BER with speed for this structure either..

These are positive results for the system designer because they suggest that within the range of Q-factors likely to be found in a UAV or military aircraft there will be no problem with time variance and there is no need to worry about it. Most bays in an aircraft with the exception of the fuselage in a transport or passenger aircraft won't have a Q as high as 10000 even when unloaded because of the relationship between volume and Q-factor.

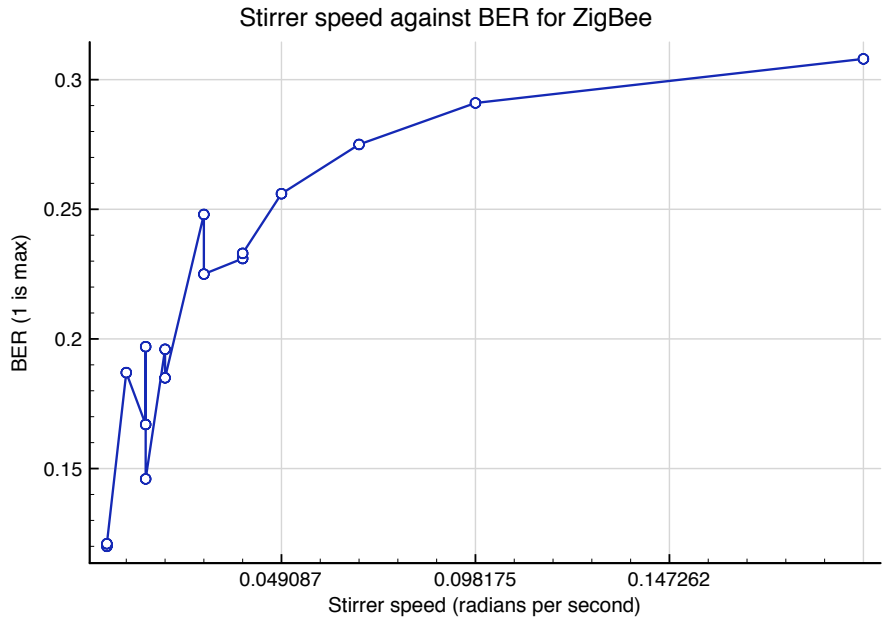


Figure 6.38: Variation of ZigBee BER in the York reverberation chamber with stirrer speed for a high Q (over 10000) configuration

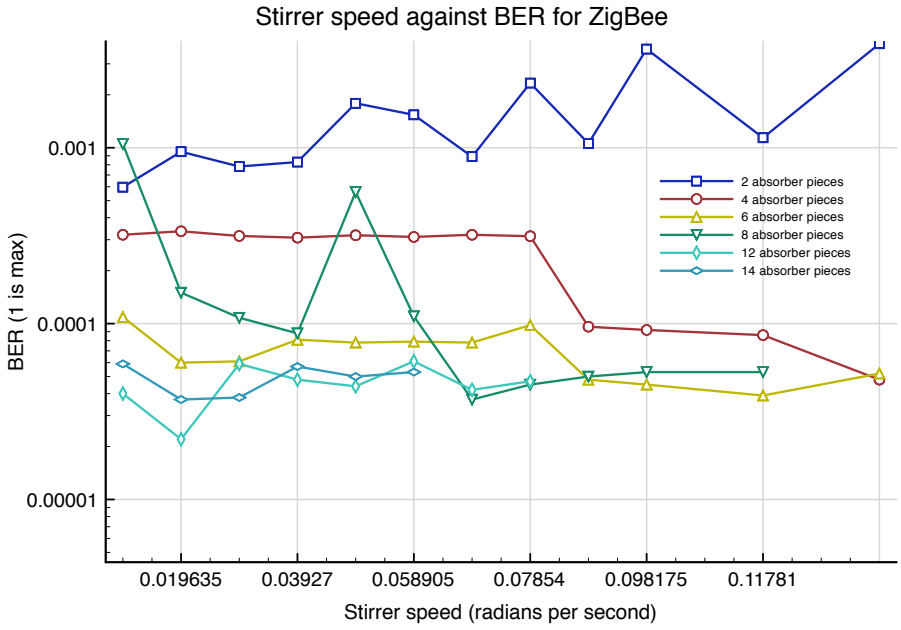


Figure 6.39: Variation of ZigBee BER in the York reverberation chamber with stirrer speed for different chamber loadings

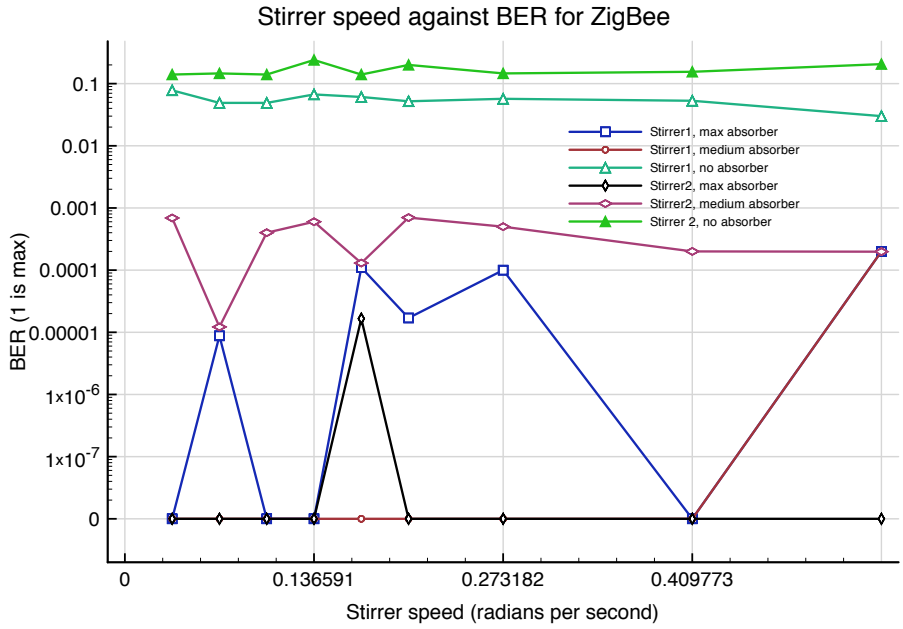


Figure 6.40: ZigBee performance in the York wing with a moving stirrer and varying loadings

In order to understand why there is no effect at the lower Qs, look at the phase response in figures 6.42 6.41 which show the phase against stirrer angle at single frequencies in the wing like structure. In the first of these graphs there appear to be dramatic phase shifts over tiny stirrer movements although a lot of these are actually just phase wrapping. In this particular plot it is hard to identify all the step to step changes as a lot of the measurements are uncorrelated but it does show well the rapid variation. It turns out that the phase change has an upper limit of approximately 1 radian of phase in 0.01 radians of stirrer change. Therefore even at the maximum speed of nearly 1 rotation per second there is a rate of change of phase of 628 radians per second. The symbol time for ZigBee is 4$\mu$s and so over this time period one could expect a phase change of 0.0024 radians. Therefore a much faster rotation of the stirrer would be required to have any effect. Even panels in a avionics bay were moving at 100Hz with sufficient displacement to replicate the effect of the stirrer then the phase change would only be 0.2 radians over a symbol. Therefore one would imagine that the rate of change of the environment, would have to be approaching 1kHz to have a serious effect. For the low Q case 6.41 most of the phase transitions are a lot slower although there are still occasional rapid movements between positions as seen at about 3.9 radians.

Figure 6.43 shows the variation in phase for the main reverberation chamber with no absorber. It turns out in this case that for a rotation speed of 60rpm the rate of change of phase would be 4000 radians per second. This is 6 times more rapid than for the empty wing but it is not enough to cause a problem over the symbol time of a ZigBee system or even a WLAN type system. However, having said this the situation it not necessarily quite as simple as suggested because of the nature of the reverberation chamber. For a large chamber the ZigBee signal will not be contained within a single mode but a set of modes all of which vary in phase independently and precisely what effect this will have is unclear and requires simulation. All that can be said is that from a receiver's perspective it should lead to a more rapidly varying detected phase between symbols. It is unlikely to be a serious effect in any case because the delay spread will cause packet loss regardless of the phase change, at the highest Qs.

Figure 6.44 shows a phase response when the chamber is lightly loaded. This particular snapshot is not typical of the phase response but may indicate a reason for the convergence seen in figure 6.38. There is not a regular phase movement but there are patches of high and low variance in phase. It could be that as the speed increases a little the fast variation changes lead to errors as the stirrer passes them but then a much greater rotation speed would be

needed to produce errors on the low variance transitions.

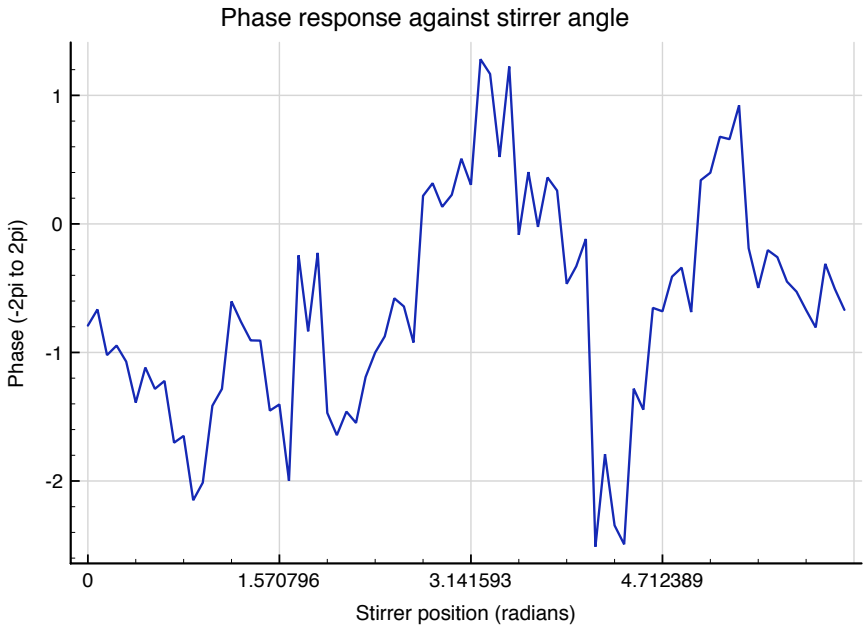


Figure 6.41: Phase response at a single frequency in the wing for a low Q configuration

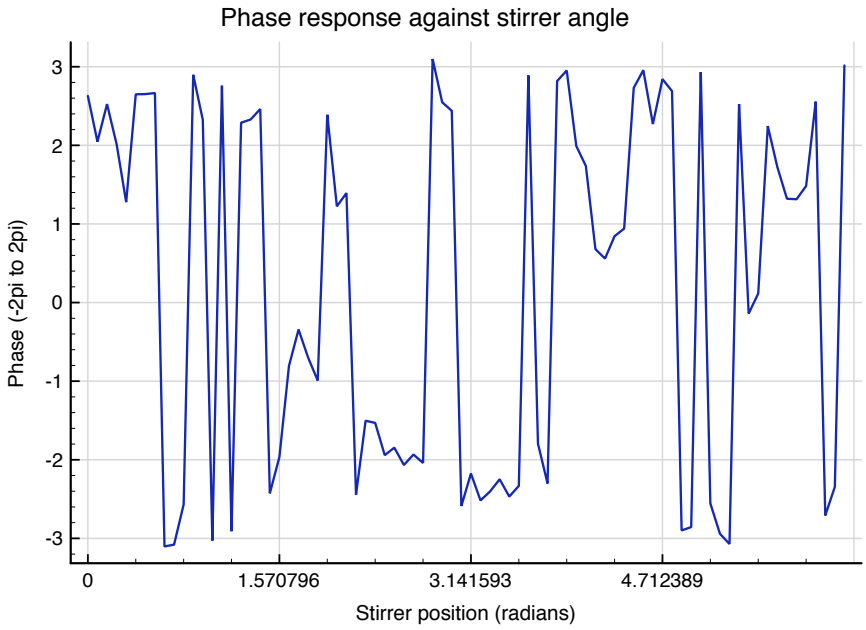


Figure 6.42: Phase response for a single frequency in the wing for a high Q configuration

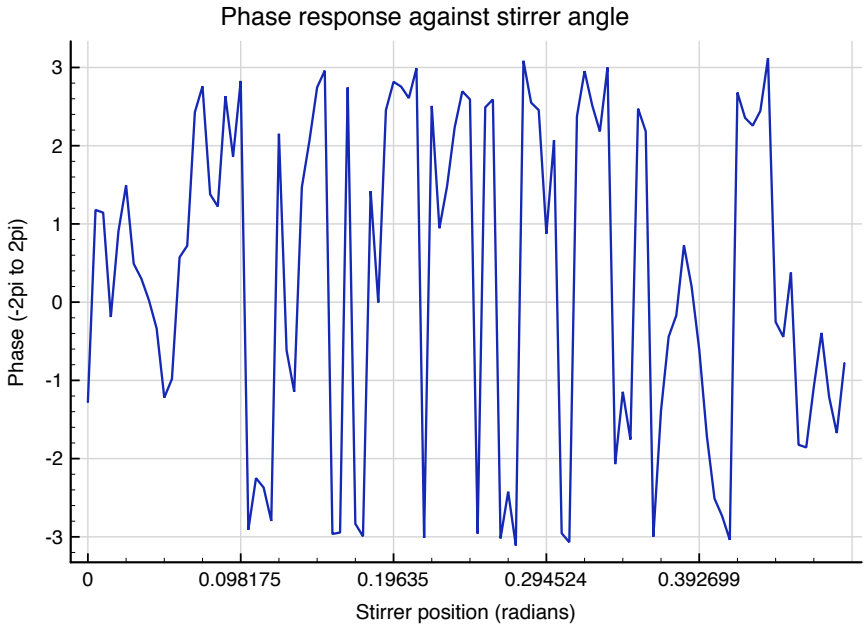


Figure 6.43: Phase response at a single frequency for the unloaded York reverberation chamber

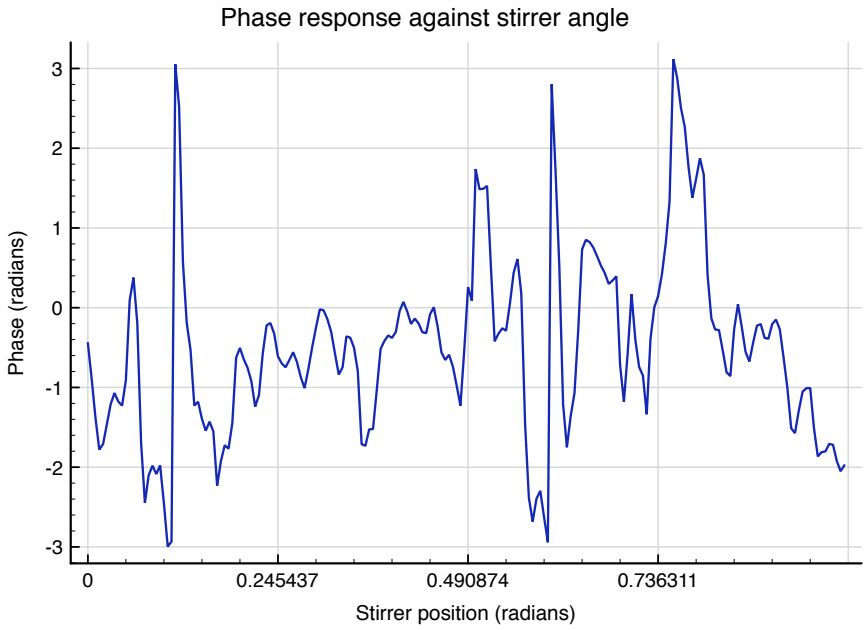


Figure 6.44: Phase response at a single frequency for lightly loaded (Q 10000) York reverberation chamber

### 6.3.2 The receive process

To understand why the moving stirrer appeared to increase the BER at high Qs a slightly lower level view of the demodulation was required than the packet level view from ZigBee radio. Therefore a signal analyser (an Agilent N902A) was used that could decode MSK signals (and many other types). This gave some insight not just into the case of a moving stirrer but also of the static errors. It was used as a simple replacement for one of the two Jennic radios taking over the receiver role. A picture of the unit and one of its display is shown in figures 6.45 and 6.46. In this case the IQ diagram is seen and an eye diagram where the signal is perfectly synchronised and with very little noise. There was only limited time with this unit and so it was tried in a variety of configurations in the reverberation chamber with both a moving and static stirrer.

When used at low Qs generally the display looked as shown and typically the eye did not close as the stirrer rotated although occasionally there problems like those shown in figure 6.47. However what was seen regularly was an eye and an IQ like those shown in figures 6.48 and 6.49. On the device screen when the display showed a plot like these (albeit varying rapidly with time) it was accompanied with a message about losing synchronisation, a problem which is clear to see from the two figures where symbols are all at completed different positions and phases in relation to one another. It was observed that a lack of synchronisation as shown occurred at some static positions and not others. Further, it was observed visually that this situation occurred for a greater percentage of time when the stirrer was moving than the percentage of static samples it occurred for. Unfortunately it was not possible to find a way to quantify this or show it rigorously but what was seen suggests the following:

- That the reason for often losing whole packets under high delay spread is because of a synchronisation failure
- That for very high Q factors, synchronisation is more likely to fail as stirrer speed, or channel time variance increases in rate

In chapter 4 the observation was made that most errors involved whole packet loss and thus 100% BER rather than just losing bits. There was a pattern of 0 or 100% errors. What the signal analyser showed was that this total loss is not just because of parts of the header being lost, but because of an inability to synchronise symbol timings and phase. The Jennic ZigBee system uses non coherence demodulation and so does not need to be carrier locked in phase

but it nevertheless requires to be symbol locked. To aid this process a 32bit preamble is transmitted at the start of a data frame (a packet at the physical layer) which allows both the symbol synchronisation and the automatic gain control (AGC) loops sufficient lock time. It would appear that for high delay spreads, (which will cause inter-symbol interference on each of these 32 bits) one or both of AGC and Synchronisation is failing and causing packet loss. When the stirrer was moving this was occurring disproportionately more frequently which shows why the packet error rate increased - the synchronisation and or gain loops were unable to track the rapid variation in the signal phase and amplitude.

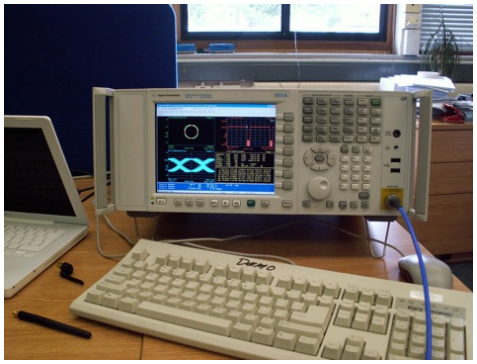

Figure 6.45: Agilent N902 signal analyser

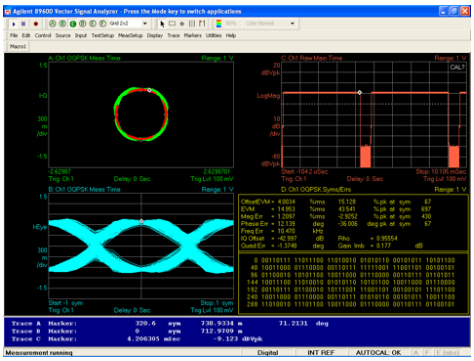

Figure 6.46: Signal analyser display with ZigBee signal

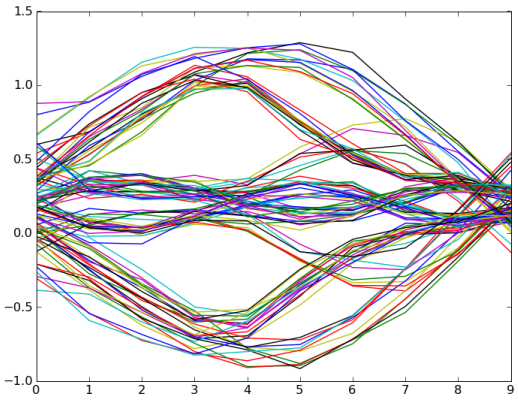


Figure 6.47: Eye diagram for ZigBee in the York reverberation chamber where there are errors in demodulation

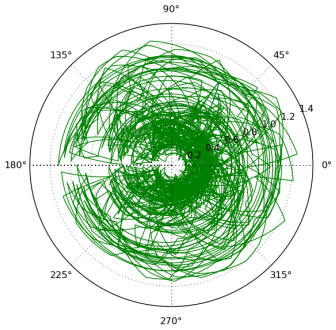


Figure 6.48: An IQ diagram for ZigBee after losing synchronisation

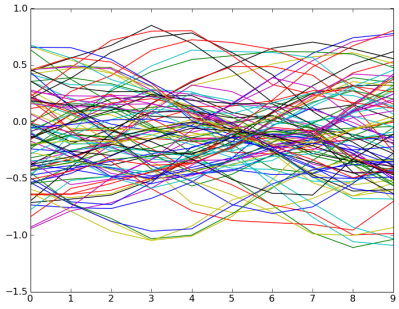


Figure 6.49: An eye diagram for ZigBee after losing synchronisation

## 6.4 Conclusion

This chapter has investigated two complications that can make a reverberant channel even more hostile for a wireless radio:

- Noise
- Time variance

The investigation into noise has shown some interesting effects in a reverberant environment. In many traditional radio communications areas the problem has been one of ensuring that at the receiver the signal is sufficiently strong that it can be detected under Gaussian noise, a problem sometimes complicated by Rayleigh, Ricean and other forms of fading. Now what we see is that noise sources are likely to undergo exactly the same statistical fading in a resonant environment. When designing a wireless system in an aircraft it will not be enough to say that a particular noise source will generate NdB of noise power in band at the position of a receive antenna. This chapter has shown that the situation is not that simple and that the noise with have maximas and minimas at different positions just as a signal does.

Statistically this gives rise to a new distribution for signal to noise ratio distribution that has been plotted and derived. These plots show that the distribution has long tails with extreme values that have a small chance of occurring but will nonetheless occur over hundreds of stirrer positions. Where wireless will be used for reliable safety critical systems the unlikely chance of finding these points cannot be ignored because even if they occurred only in 1 in 10000 flights that would be too often. As with the standard Raleigh channel, these problems are best countered through the use of redundancy to further reduce the probability of high BERs.

It has been shown that average BER is significantly increased over the theoretical case for a given average Eb/No as a direct result of the changed distribution, both at low and high Qs. Where there will be extreme values of noise to signal, then the BER will reach very high values. The average BER will be important in a slow or fast time varying environment where a receive antenna may be going in and out of a high noise, low signal position every second, or 10 second or some other period. For some applications this is not a problem and a retransmission could take place, in others a short period without transmissions would not be acceptable. These are all issues for the system designer and at a higher level of the protocol stack than is considered in this thesis here but they are important issues to be aware of in order to make a fly-by-wireless aircraft robust.

For the moving channel, the results have indicated that time variance should not be a problem in practice for a ZigBee type system. The symbol time may have to be two orders of magnitude less in order for a time varying channel to have an effect or the time variance must be much more rapid. There was a PER increase with stirrer rotation rate for an empty and lightly loaded large reverberation chamber but once Q was below 5000 the problem went away and so we would not expect problems in an aircraft. From a purely theoretical perspective there could be some interesting research in finding out how the many different modes at different Qs affect a receiver system that is trying to track phase. Our work on the time varying channel has also revealed more about the error pattern seen in the prior chapters indicating that the reason high delay spreads are causing whole packets, rather than bits to be lost is that synchronisation is never established following the frame preamble.

# Chapter 7

# ZigBee measurements in an aircraft

## 7.1 Introduction

The majority of the work described in the previous chapters has been based around measurements carried out in the laboratory. This chapter presents the results obtained from measurements on an aircraft and see how these fit in with the scenarios set out in the other chapters.

Early into the 3 years of the Flaviir project a plan was formulated to test ZigBee systems on a BAE UAV demonstration platform which would have provided an ideal way to test in a realistic environment with all the complications that can bring. Unfortunately this fell through due to outside pressures and an alternative test platform was required. The final decision was to perform tests on a Tornado GR4.

This is an aircraft that was introduced in 1979 and was manufactured up until 1998. It is not a UAV type aircraft in any way but a twin engine combat aircraft. Clearly the age has to be a concern in the sense that results from airframe measurements should be applicable to a modern jet and ideally UAVs. The latter requirement is impossible to fulfil here as no UAVs were made available. However, it is to be expected that all aircraft will have a combination of bays and aperture coupling points and the standard features of hydraulics, electrical systems and cabling. Although the size of a combat aircraft will be significantly greater than a surveillance UAV, for example, there will be common characteristics.

In comparing the Tornado to a modern jet then there should be even more similarities. The author was able to see the inside of a Typhoon during manufacture and there was a lot in common with the Tornado. To put it in simple terms, both are metallic and feature a cockpit, a fuselage which is divided into a number of separate but coupled bays, fuel tanks, wiring looms,

metal rods and hydraulics. In other words, although the avionics technology has become more sophisticated the actual aircraft structure has the same features.

### 7.1.1 Aims

The primary aims of the work presented in this chapter were:

- To obtain as many channel response measurements as possible on a real aircraft so that the models could be compared to those used in the laboratory and used to improve them in future.
- To observe the performance of the ZigBee system in an aircraft to see whether the environment poses a hostile or benign channel.

As mentioned in the introduction, the original plan for carrying out real aircraft measurements was for flight testing on a UAV demonstrator. With this came some secondary aims, which were to learn about:

- Practicalities of placing radios in a real airframe
- Design of electronics to fulfil air safety requirements
- Performance evaluation in the presence of in flight noise (e.g. electrical systems)

The last of these was not possible in the revised (Tornado) measurement plan but 1 and 2 were. 1 and 2, whilst barely touched on in the rest of this PhD thesis are worth mentioning because they were not trivial. The test equipment was going to be on a flying aircraft, operational in flight and therefore there was a need to design a robust airworthy system. Part of this is shown in figure 7.1. Any system to go on a real aircraft must deal with issues of pressure, temperature, robustness to vibration and damage. All these features may restrict the initial ideal of a very portable light system and although the system shown is more extreme in its robustness than is practically needed, it does illustrate a point not made so far: the major weight savings and flexibility that would be provided by a wireless system have been taken for granted. However, where there are power cables (as in the figure) or large batteries or processing units alongside the radio then suddenly each chip can end up a lot larger than expected and finding space for these in the tight confines of an aircraft can be difficult. New and novel architectures and manufacturing are needed to maximise the wireless advantage. In the antennas appendix

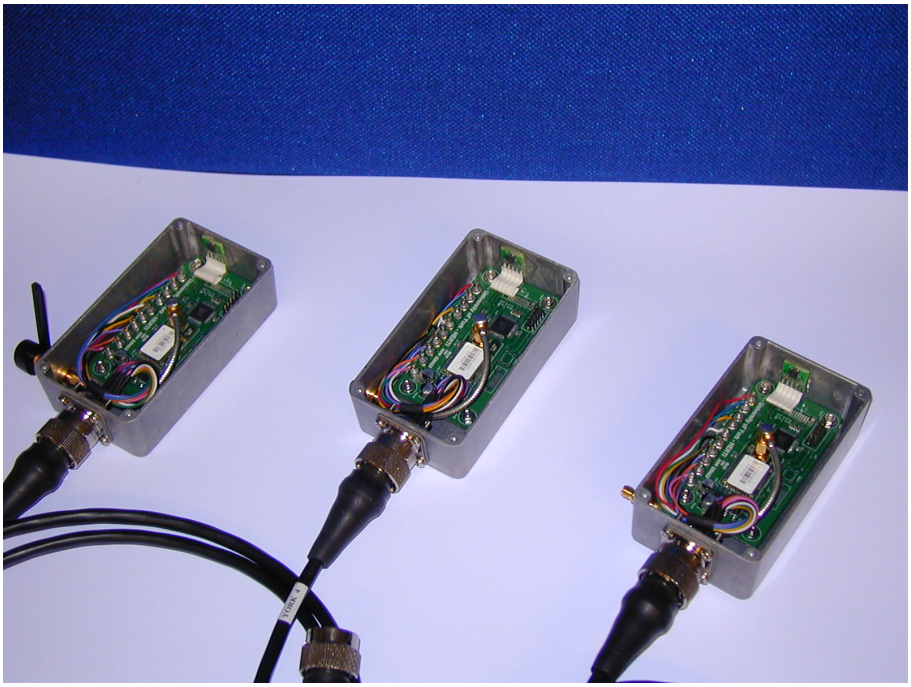

Figure 7.1: Overview of the measurement and aircraft

there is an example of this with direct write a technology that allows the writing of both novel antennas, sensors and power buses allowing a number of devices in a constrained space.

## 7.2 Measurements

### 7.2.1 Measurement procedure

The measurement procedure is the same as in previous chapters so as to allow easy comparison of results. 2 stirrers were employed to stir the aircraft bays just as the reverberation chamber or aluminium wing structure. By using these a statistical description of the aircraft could be obtained - something not found in any of the literature. A network analyser was used to measure the frequency response and an FFT was employed to transform this into a time domain power delay profile. A ZigBee test system was employed to measure bit error rates around the aircraft for a sample typical sensor network system and for every channel response measurement a BER measurement was carried out at the same antenna and stirrer positions.

### 7.2.2 Measurement configurations

The measurements were broken up into a number of types and the results that follow are arranged to match this. Two bandwidth ranges were used: one from 2.43 to 2.47GHz and another from 500MHz to 3GHz. The measurements vary from a single cavity to multiple cavities between different sections of the aircraft. Stirrers were used in all cases and typically there would be one stirrer per antenna containing cavity. In other words any intermediate cavities were not stirred.

A summary of the setups that will be analysed is given in figures 7.1. Other scenarios were also measured but these provide a good representation of what was found. Some measurements were done with both antennas in one avionics bay and others were done with one antenna in one bay, one antenna in another. Where two bays were used 2 stirrers were employed as specified in the table. Intermediate bays were not stirred as only two small stirrers were available. All measurements were very time consuming because of the stirring and because an older network analyser was used[1] which limited the number of configurations that could be measured. In some cases there were no intermediate bays and pairs of bays could be adjacent along the length of the aircraft or opposite on each side. So that both electromagnetic and ZigBee measurements could be carried out an RF switch was used allowing them to be done in succession on each movement of a stirrer. This was exactly the same as the laboratory setup in previous chapters.

To get a better understanding of what is meant by a bay and the separation between them as well refer to figures 7.2 and 7.3 which also show the complete test setup. Very roughly, a bay might be 1m across and 0.5m deep but some were smaller and some larger than this. Further, some were filled with many cable looms, some less so and some with large apertures and others with smaller apertures. The main principle of the measurement plan was to obtain as large a range as possible of measurements in the given time, so that the extremes of electromagnetic behaviour in the aircraft could be understood as well as the typical behaviour.

One of the challenges was that it wasn't possible to drill holes into the airframe and place bulkhead connectors in-between, as was done in the lab to avoid any uncontrolled leakage of energy. A way had to be found to get in both RF coaxial cable and power cables for driving the stirrer. This was done any way possible dependent on the measurement. In some cases such as that shown in figure 7.3, a direct route was available avoiding the need to make any additional

[1] this was used to due the risks of working outside

| Number | Band (GHz) | No of cavities | Separation | Other notes |
|---|---|---|---|---|
| 1 | 2.4 | 1 | NA | Open bay |
| 2 | 2.4 | 1 | NA | Closed bay |
| 3 | 2.4 | 1 | NA | Closed bay |
| 4 | 2.4 | 2 | 1 bay (adjacent) | 2 stirrers |
| 5 | 2.4 | 2 | 1 bay (opposite) | 2 stirrers |
| 6 | 2.4 | 2 | 2 bays | 2 stirrers |
| 7 | 2.4 | 2 | 2 (diagonal) | 2 stirrers |
| 8 | 2.4 | 2 | 2+ (top to bottom) | 2 stirrers |
| 9 | 2.4 | 2 | 3 | 2 stirrers |
| 10 | 2.4 | 2 | 3 | Big separation |
| 11 | 0.5-3 | 1 | NA | Single stirrer |
| 12 | 0.5-3 | 2 | 1 (opposite) | 2 stirrers |
| 13 | 0.5-3 | 2 | 1 (diagonal) | 2 stirrers |
| 14 | 0.5-3 | 2 | 2 (same side) | 2 stirrers |

Table 7.1: Antennas, Q-factors and the time constants associated with the measurement

openings. Where this could not be done there was usually an internal route available and so a neighbouring cavity was used to get the cable in (perhaps one with an opening as just described). In most cases an direct or indirect route in was found that did not involve leaving a bay partially open.

Returning to table 7.1, this shows that there are situations with antennas very close together and situations where the antennas are apart and there are multiple cavities between the stirred cavities. It might be expected that there would be more attenuation at greater distances but also different statistics, such as multiple Rayleigh. A large number of images are provided throughout this chapter to illustrate the propagation environment for each measurement.

## 7.3 Results

The results are divided into single and multiple cavity results and begin with network analyser only measurements to build a picture of the channel in terms of how flat it is, what its auto-correlation looks like, whether there are any direct path components and what attenuations

Figure 7.2: Overview of the measurement and aircraft

might occur though the aircraft.

### 7.3.1 Single cavity channel measurements

In the first case measurements were taken for a single cavity and a single stirrer. The cavity is shown in figures 7.4 and 7.5 and the measurements numbers referred to here are 1 and 2 for the case where the cavity is open as in the image and when it is closed. For measurement 3 a very similar setup was used but with a different cavity. It is tricky to see in the images but the antennas were positioned at arbitrary angles and orientations tied onto the wiring looms but kept a quarter to a half wavelength away where possible. The cables were brought in through apertures from the outside. Clearly this is a leaky cavity compared to the sealed reverberation chambers considered in Chapter 5. There were numerous thick wiring looms that would be expected to absorb RF energy in the same way that the AN79 absorber did in the reverberation chamber.

The frequency,time and correlation responses are plotted in figures 7.6 and 7.7 for the bay in figure 7.4 with its door closed. The frequency and time response given is for the first of the 80 stirrer positions and the correlation plot is averaged over all 80 positions. Another measurement was done in a separate single cavity, the results for which are shown in figure 7.8 and 7.9

Figure 7.3: One route into the aircraft

corresponding to measurement 3.

Figure 7.4: Stirrer and one antenna for the first single bay measurement

Figure 7.5: The other antenna for the first single bay measurement

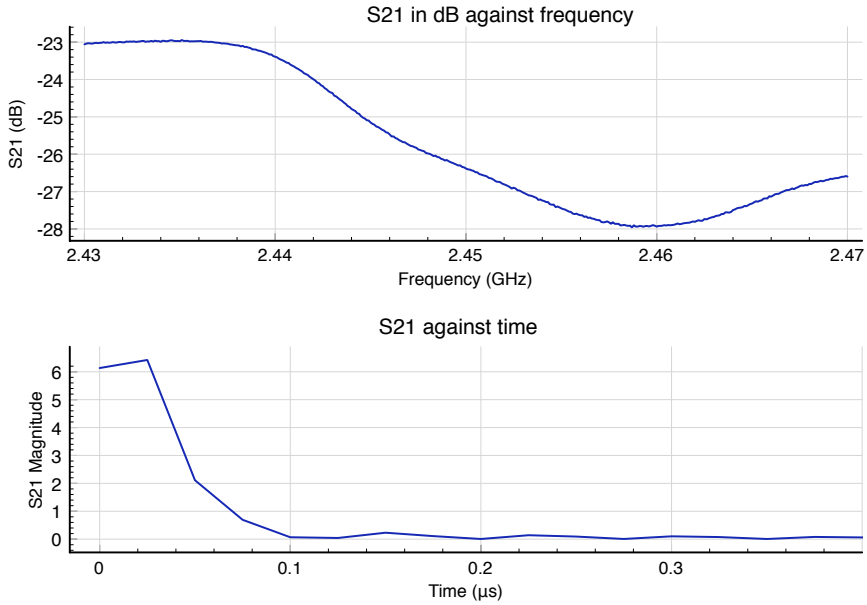


Figure 7.6: A single cavity frequency-time response (measurement 2)

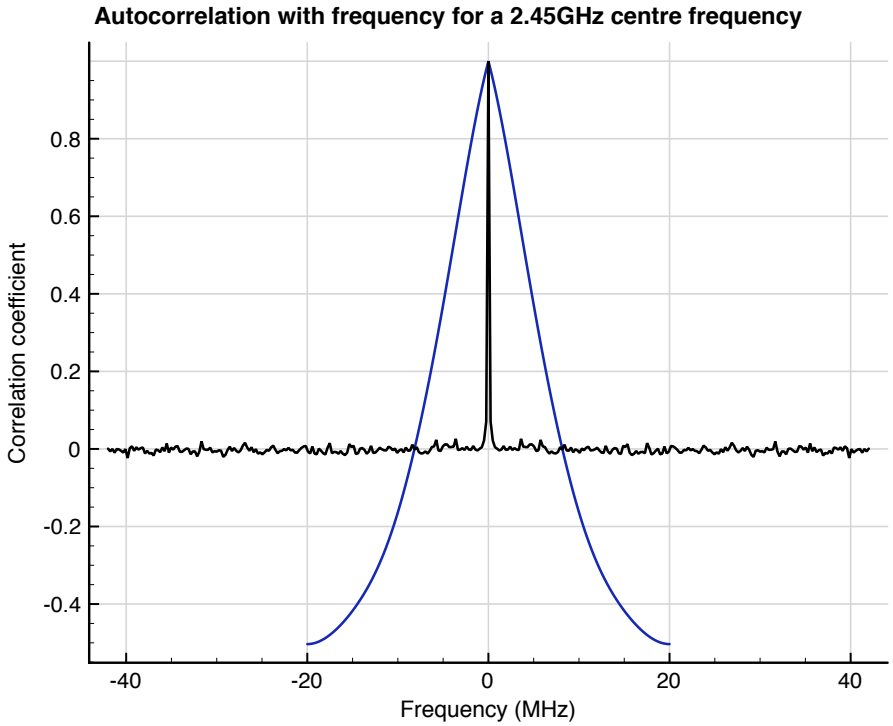


Figure 7.7: A single cavity correlation response (measurement 2)

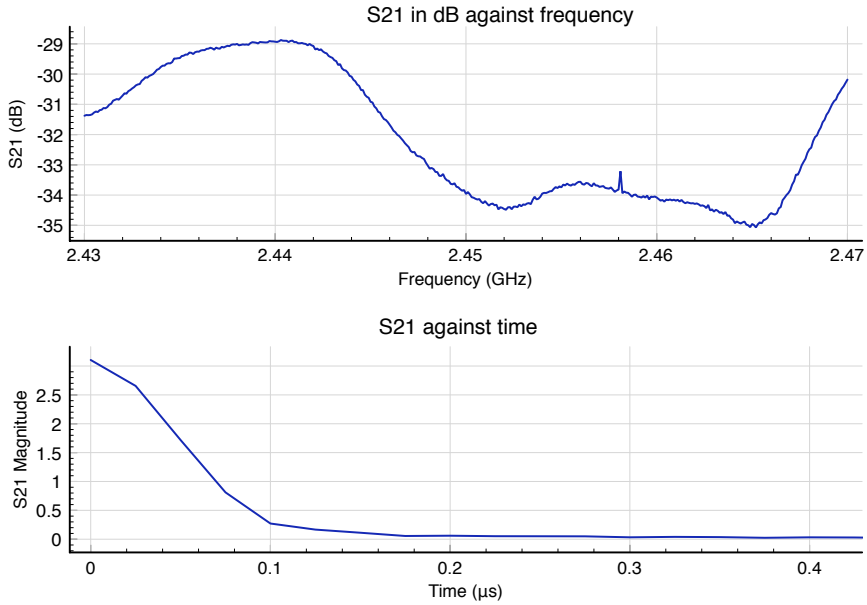


Figure 7.8: A single cavity frequency-time response (measurement 3)

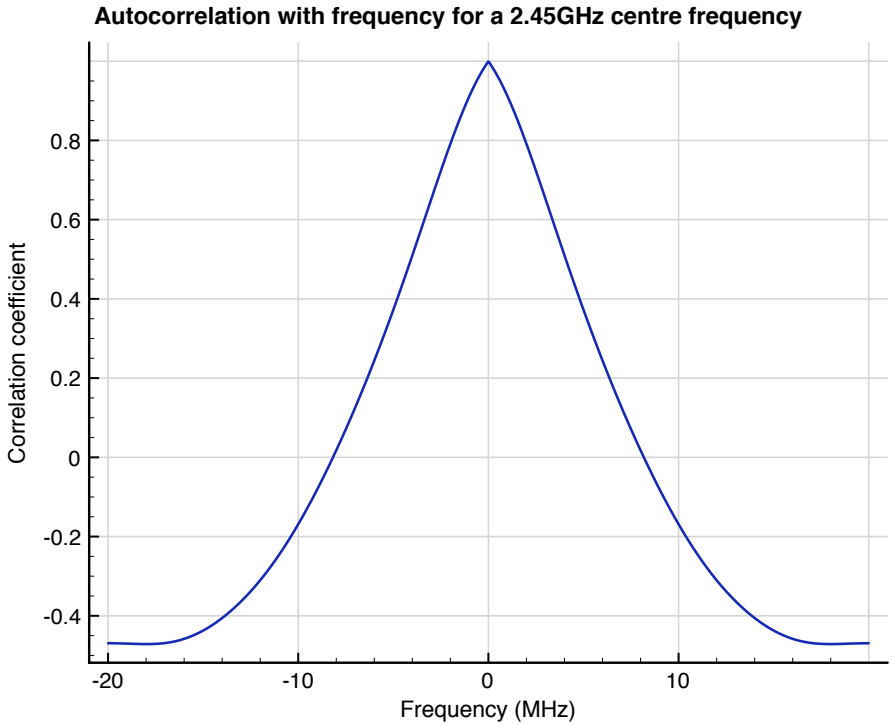


Figure 7.9: A single cavity correlation response (measurement 3)

Looking at the frequency/time plots first of all, the response is much flatter than before with a variation of only 6dB which compares with 20dB+ in the reverberation chamber and the wing, even where absorber is present. Over the range of a ZigBee or wireless LAN channel it is flat. One would expect from this that the channel would not provide a difficult communications environment. There is a slight difference in absolute level between the two bays used in measurements 2 and 3 but this is to be expected as the bays were different sizes with different apertures and containing different wiring looms. Overall the responses are very similar.

Moving onto the time domain, the time response is very rapid compared to the reverberation chamber. When the time constant is 0.05$\mu$s and under the Q is no greater than 700 which is very much at the lowest level of the lab measurements described in earlier chapters. It was found before that low delays like this were usually benign to ZigBee systems (and should be for almost any system). In fact it is difficult to see the response, as the band's width limits the resolution and it may not be easily seen over the filter response.

Therefore, in this section we shall also show (for the 2.4GHz measurements) the correlation coefficient in the frequency domain which is inversely related to the time constant [43]. The corresponding correlation plots for measurements 2 and 3 are presented in 7.7 and 7.9 respectively. It is the average autocorrelation over frequency in the sense that the correlation has been calculated for each stirrer position and then the resulting correlations have been averaged per frequency offset. For the first of these the empty reverberation chamber is overlaid for comparison. The difference in correlation here is very clear to see. In fact the correlation is at one half its maximum with a bandwidth one hundredth that of an empty reverberation chamber. Similarly the delay spread is in the order of a few of tens of nanoseconds whereas in a reverberation chamber it is in microseconds. Although this is not to say anything about the absolute accuracy of the correlation coefficient in it's relation to delay spread it shows that an approximate relationship fits here just as it does for more traditional multi-path channels described in communications literature.

**An open single bay measurement**

In figures 7.10 and 7.11 measurement 2 is repeated but this time with the bay door open as a point of comparison to see how much energy storage and resonance is actually exhibited by the closed cavity. It is seen that the frequency domain is again quite flat with under 8dBs variation and the time response is near instantaneous, with little more than the FFT filter re-

sponse. Although the shape in the frequency response looks similar to the closed case in that it is not totally flat, note that it is 20dB down on the magnitude for the closed case (over the lower frequencies) and so although the ratio of the peak to peak is slightly greater the absolute differences are much smaller. The central maxima is centred close to 2.45GHz and we may be seeing the antenna response rather than cavity resonances. This measurement shows that due to the signal strength being 20dB down with an open door, the cavity does store energy and has resonances, even if they are lower in energy in comparison to the lab. It is also possible to see the dangers in the correlation coefficient plot and time response once the time period becomes very small because factors other than resonance become far more important.

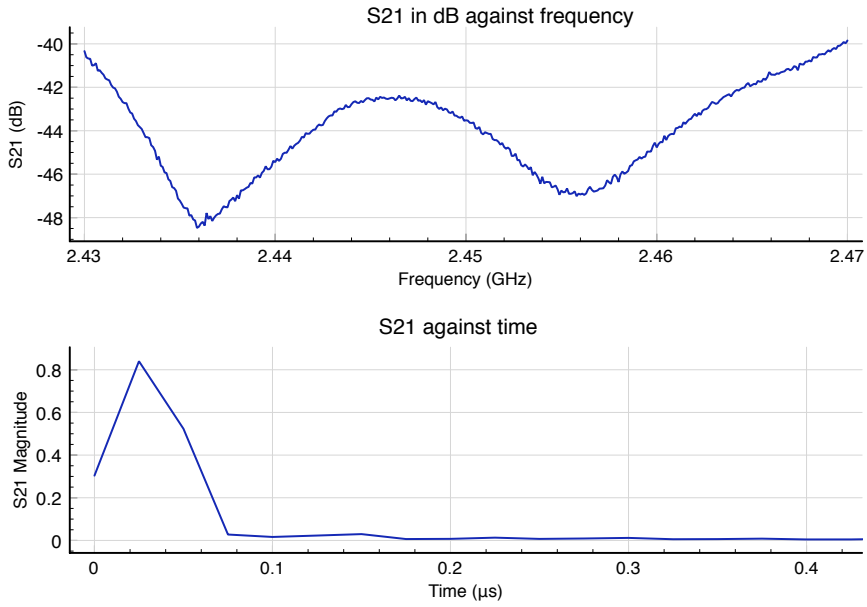


Figure 7.10: Frequency time graphs for a single open cavity (measurement 1)

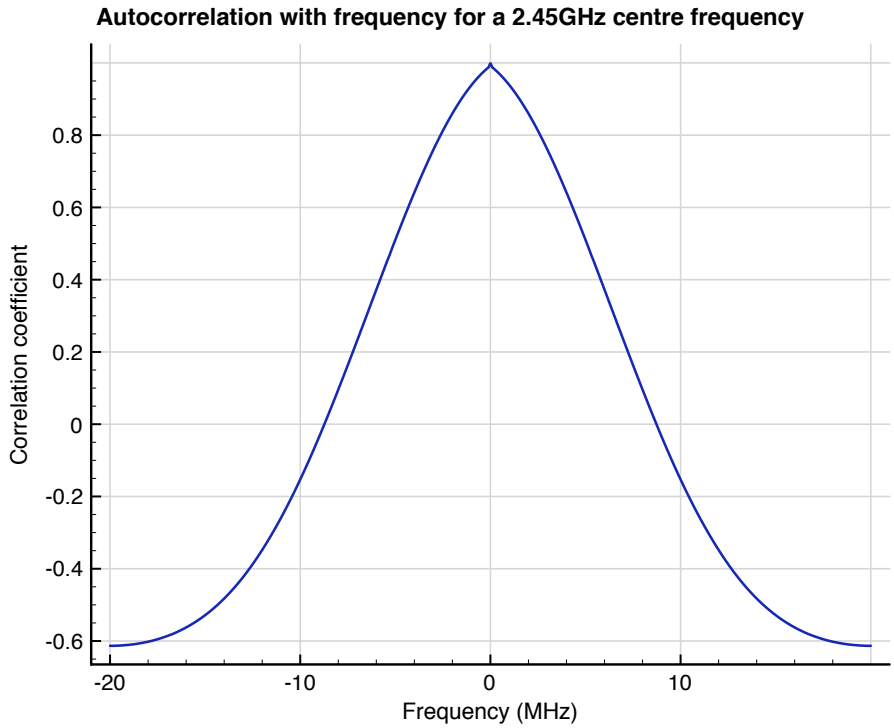


Figure 7.11: Frequency correlation for a single open cavity (measurement 1)

### A broadband response

Next the broadband response of the cavity are show for the cavity used in measurements 1 and 2 because it shows if their response is consistent over a wide frequency range. Time and frequency single cavity broadband responses for measurement 11 are shown in figure 7.12 which is the same cavity and setup as measurement 2. One cannot read too much into the

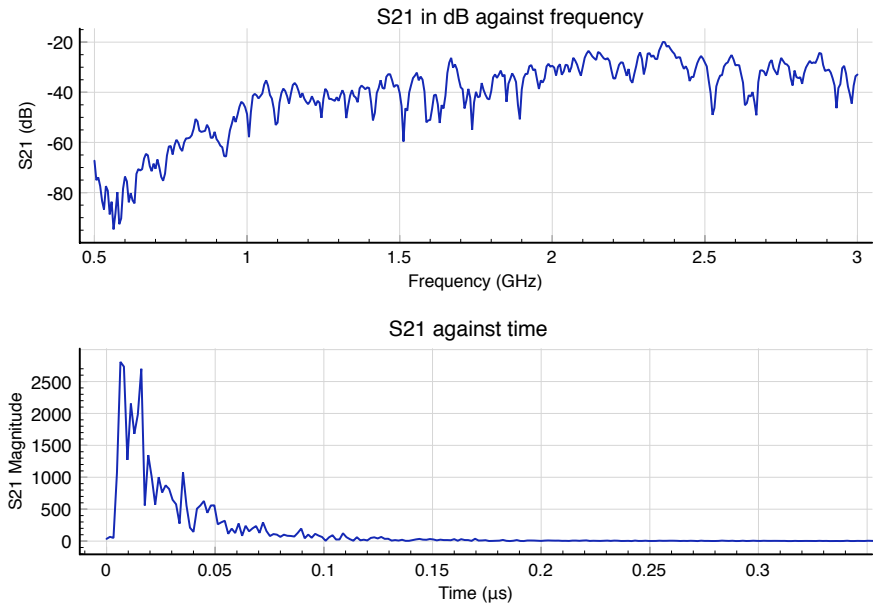


Figure 7.12: Broadband frequency-time graphs for the first single bay measurement ( same setup as measurement 2)

increase in S21 from 0.5GHz to 1GHz because the antenna used was not designed for broadband measurements. Nevertheless, although being designed for 1GHz it is effective in exciting the cavities past 1GHz which is interesting in itself. The frequency response pattern looks the same from 1GHz to 3GHz which means that our 2.4GHz centred analysis should apply just as well to other bands.

Turning to the time domain, with this much greater bandwidth the resolution in the time domain has been enhanced but the decay time has not changed dramatically from the narrow band measurement having still decayed completely by 0.1$\mu$S. However it does appear to decay a little quicker which suggests the narrowband measurement, distorted by filtering is a slight over estimate of delay. It is worth observing that for the aircraft environment the use of wideband delay spread measurements is more useful than in higher Q measurements because and can provide an upper limit on delay spread.

**Conclusions so far**

From the first two cavities presented the following conclusions can be drawn:

- A military aircraft is packed heavily with wire and has many openings, so absorbs and leaks a lot more energy than a near empty metal box as used in the laboratory
- The wireless channel appears flat for ZigBee and 802.11 bandwidths
- The delay spread is so low that it can't be accurately reported
- As a communications channel it should be far more benign than the laboratory cavities

### 7.3.2 Multiple bay setups

#### Adjacent bays

Following on from the single cavity case this section moves to two cavities. The configuration in the previous section with two antennas close to each other in a high absorption environment could be considered an ideal communications channel in that there is low delay spread and a direct, low loss path. For multiple cavities there is no direct path, there will be an increase in attenuation and possibly an increase in delay spread dependent on the number of coupling points (see Chapter 5). The two cavities used in the first of the two bay/cavity measurements are shown in figures 7.13 and 7.14. These were adjacent bays that were individually stirred. Cables were placed in through existing apertures and the bays were closed during the measurement. The apertures between the two bays were close to being electrically large.

Figure 7.13: The first of two cavities used in measurement 4

Figure 7.14: The second of two cavities used in measurement 4

Plots of the frequency and time response measurements are presented in figure 7.15 and the correlation in figure 7.16 As before there is very little sign of delay spread in the time response but this time and the correlation plot drops to a value of 0 within 10MHz again which is comparable to the earlier measurements. Where this measurement is different is in the attenuation of the signal (a loss of 60dBs in the worst case) and in the maximum to minimum S21 variation of 12dBs which is higher than before. Although only one stirrer position is shown in the plots these were true across positions. It was seen in Chapter 5 that these differences can be due to two things:

- The frequency dependence of the apertures through which coupling happens.
- The combination of the two cavities in a double Rayleigh type system.

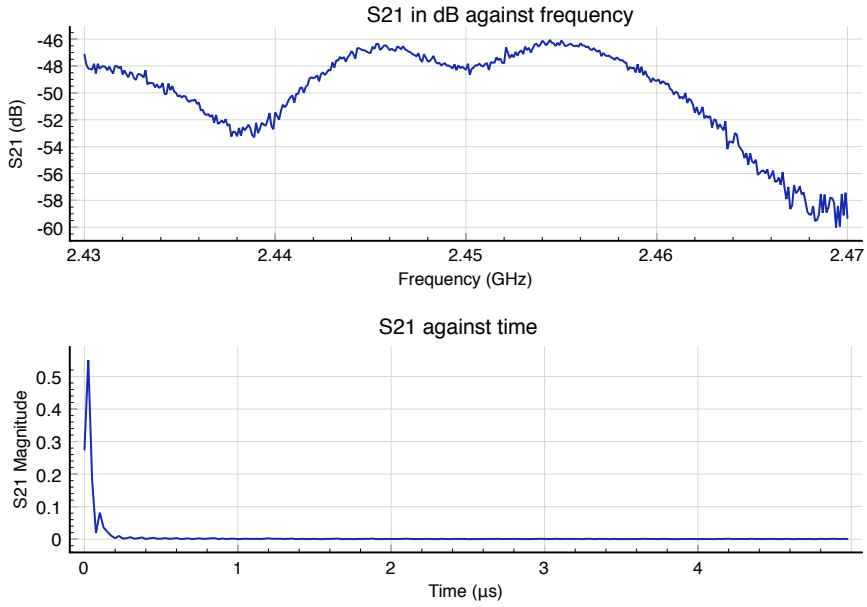


Figure 7.15: Frequency and time response for transmission between 2 adjacent cavities (measurement 4)

After looking at these measurements and the aperture sizes it appeared more likely that the greater variation was due to one of the cavities being larger than in the previous measurements.

If wireless systems are to be placed in existing aircraft it appears that signal strength may be as significant an issue as delay spread. Within a single cavity a lot of absorbing material is beneficial in that it avoids resonances but when the direct path from antenna to antenna is cut off then there is very heavy signal attenuation. This is not an issue in a low noise environment but it could be an issue when dealing with an SNR close to the sensitivity limit. Further it demonstrates a problem with the obvious technique for reducing delay spread: introducing absorber. When there is a highly resonant environment, adding absorption will reduce delay but simultaneously decrease signal to noise, unless we compensate with an increase in output power for the radios; something that is not desirable on low power systems. Within the single cavity environment high absorption is ok, but when there is a need to communicate down the length of an aircraft it may be a problem.

A correlation for a similar system is shown in figure 7.17. The difference here is that the path is across the aircraft's width rather than down its length. In this instance there were were smaller coupling points, without the largish clear openings for cables that one might find for

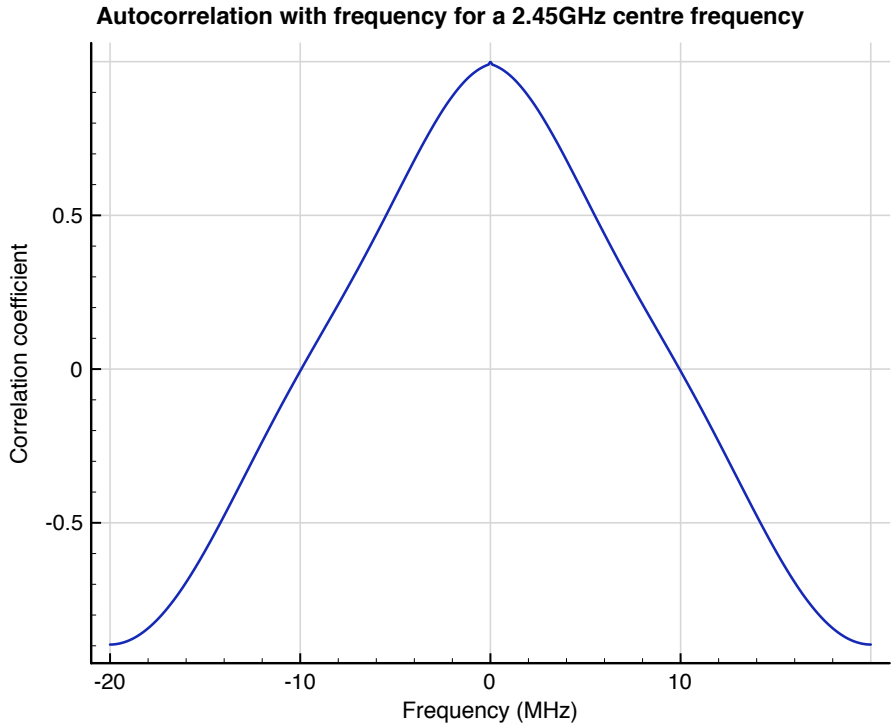


Figure 7.16: Frequency autocorrelation for 2 adjacent cavity transmission (measurement 4)

cables going down a length of the aircraft. The correlation plot is 50% narrower this time than the single cavity case.

When carrying out the measurement, it was noted that there was more of a direct path through a larger aperture for the measurements down the aircrafts length rather than across its width. To test for this the polar plots are provided for all points over all stirrer positions in figures 7.18 and 7.19 It is seen here that neither plot is centrally based but that figure 7.18 is more skewed and has one or more indirect paths

**Separated bays**

The next measurement examines what happens when the separation distance increases so that the 2 cavities are no longer adjacent. In this case the antenna in figure 7.14 has moved to the bay shown in figure 7.20 For this measurement there are still only two stirrers and so only the bays with the antennas are being stirred meaning that intermediate bays will be nearly static. Results are presented in figures 7.21 and 7.22. Attenuation has increased further and so significant noise and be seen with the signal nulls approaching the measurement noise floor of approximately -85dBs. The correlation width lies between that of the adjacent and opposite

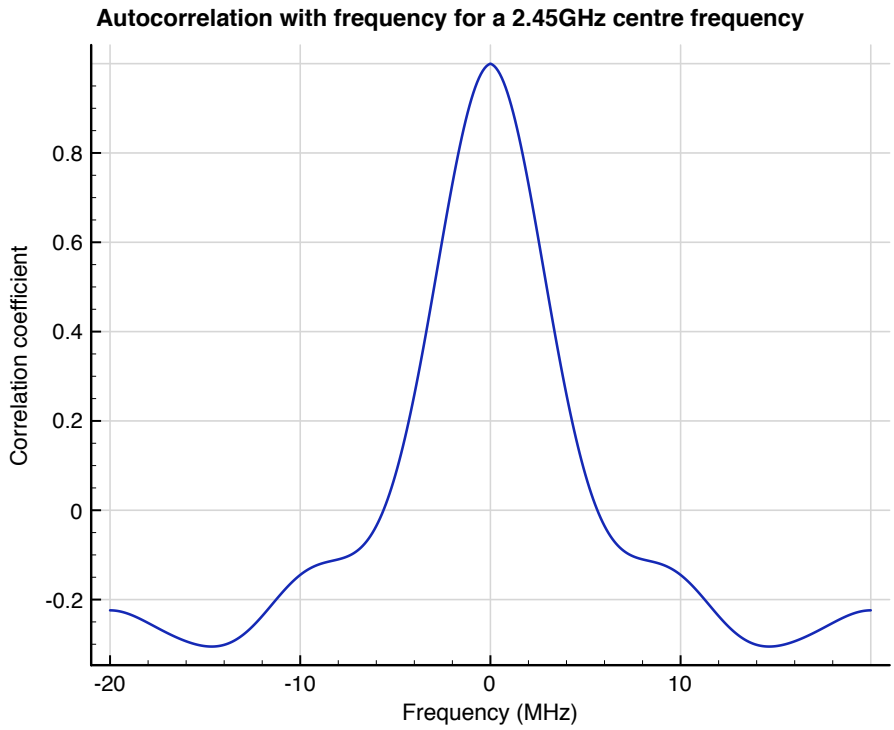


Figure 7.17: Frequency correlation for 2 adjacent cavity transmission (measurement 5)

two cavity measurements, if ignoring the noise spike at the top of the plot.

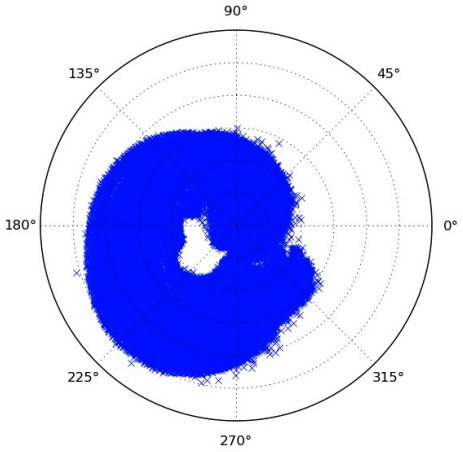


Figure 7.18: Polar plot for 4

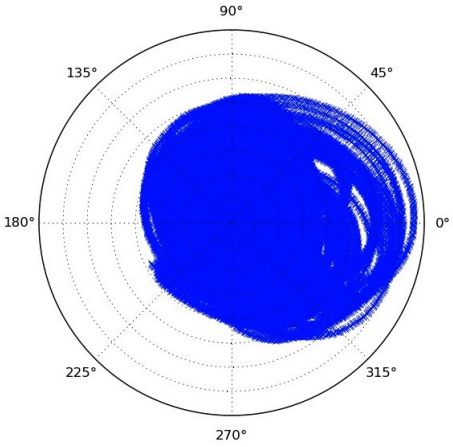


Figure 7.19: Polar plot for 5

Figure 7.20: The bay used for the more distant antenna in measurement 6

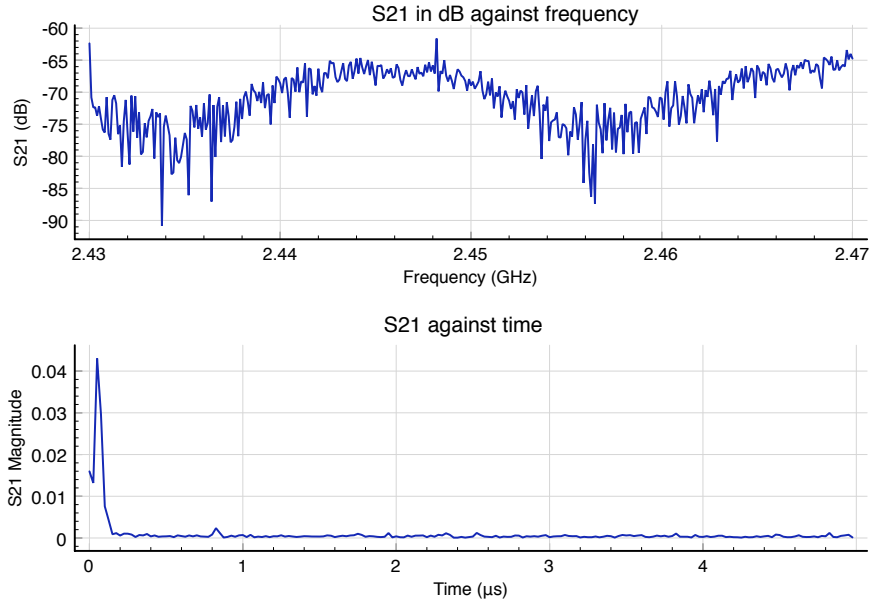


Figure 7.21: Frequency and time response between 2 bays with a separation of 1 bay (measurement 6)

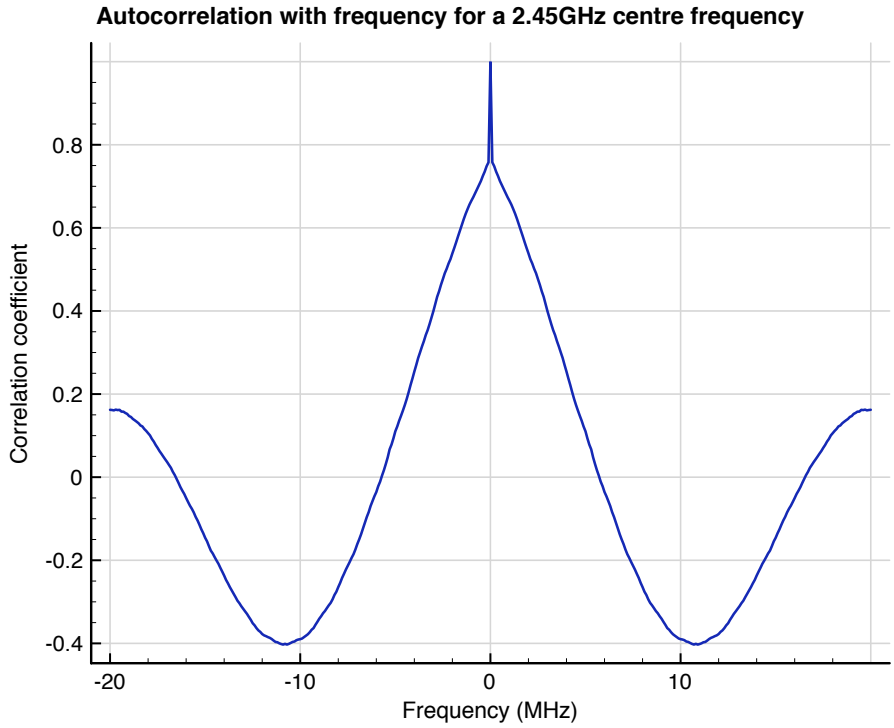


Figure 7.22: Frequency autocorrelation between 2 bays with a separation of 1 bay (measurement 6)

### Further scenarios

A range of other setups were measured going between different aircraft sections: top to bottom, side to side, diagonally, etc. These were done because it's would be useful to know if the coupling between any two sections is reasonably predictable in terms of frequency response or totally different for any 2 measurements. One could expect some consistency because where there are a large number of coupling points and lots of absorbing structures then statistically similar results would be expected from the ensemble. The first setup is a diagonal coupling across the aircraft using the two cavities shown in figures 7.23 and 7.24

Figure 7.23: Diagonal coupling, bay 1 (measurement 7)

Figure 7.24: Diagonal coupling bay 2 (measurement 7)

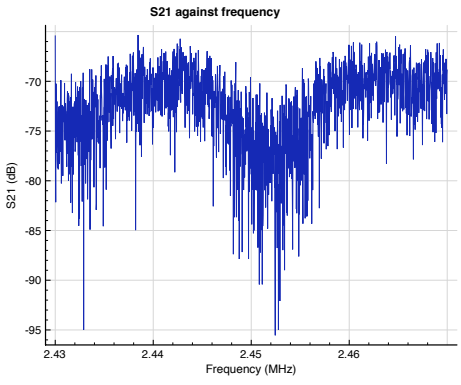


Figure 7.25: Frequency response for diagonal coupling (measurement 7)

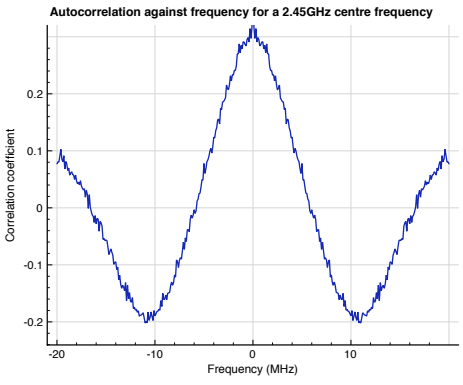


Figure 7.26: Autocorrelation response for diagonal coupling (measurement 7)

There is a problem with the noise on this measurement as with some of the other measurements at large separations and it would have benefitted from going to the maximum power on the network analyser and a longer measurement time. Nevertheless one can still see the main

shape of the frequency response in figure 7.25. The correlation plot is clearer, although the peak is not at 1 again here because of noise meaning top of the plot in the figure should be treated as 1 when comparing half or quarter value bandwidths. The null bandwidth (which is straightforward to compare) is exactly the same as the 1 bay separation measurement as is the maximum to minimum signal variation which is approximately 15dBs across stirrer measurements, 3 times the variation in the single cavity case.

In figures 7.27 to 7.30 the measurement went from the top to the bottom of the aircraft instead, utilising openings close to the cockpit and again obtain similar results in terms of correlation width and maximum to minimum signal strength. Although the delay spread here is clearly very low, due to wide auto correlation, signal to noise is also low.

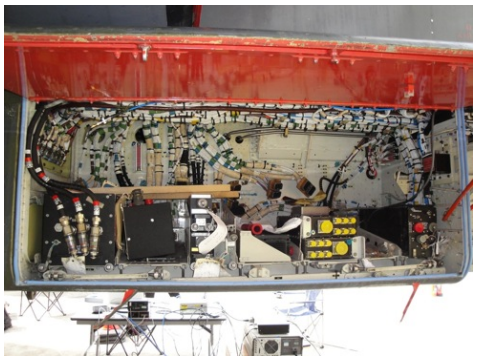

Figure 7.27: Bay in the top of the aircraft (measurement 8)

Figure 7.28: Lower bay for top to bottom coupling

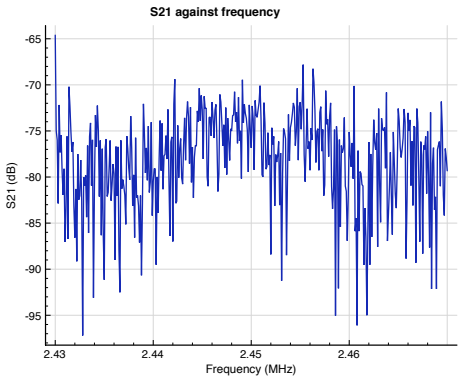


Figure 7.29: Frequency response for top to bottom coupling (measurement 8)

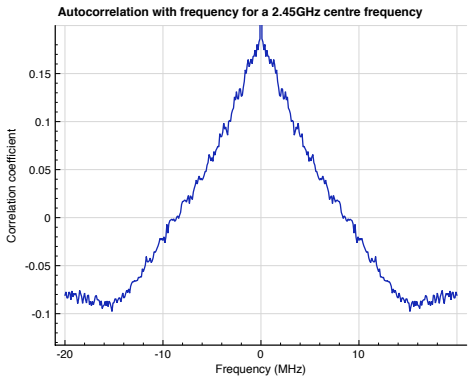


Figure 7.30: Autocorrelation response for top to bottom coupling (measurement 8)

Finally a diagonal measurement is shown in figures 7.31 to 7.34, where the cavities are on

opposite sides of the aircraft, and at difference distances down it, with a 2 cavity length separation. This configuration illustrates the attenuation over a greater distance than before as well the effects of coupling through multiple cavities. For this measurement one can see that the noise floor is lower than before due to setup changes. The channel has greater variance than before and this is reflected in a correlation function nearly half that of the last 2 measurements and a decay that is visually almost twice as long. Measurement 9 was a very similar setup and so won't be shown here although the BER is presented later.

Figure 7.31: Cavity 1 for final diagonalmeasurement (10)


Figure 7.32: Cavity 2 for final diagonal measurement (10)

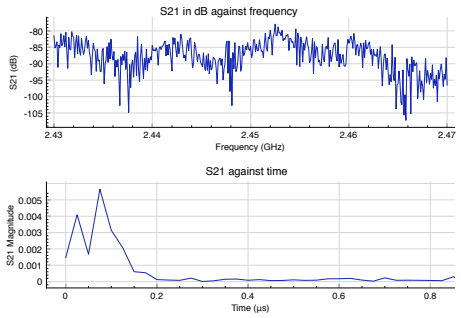


Figure 7.33: Frequency response for final narrowband measurement (10)

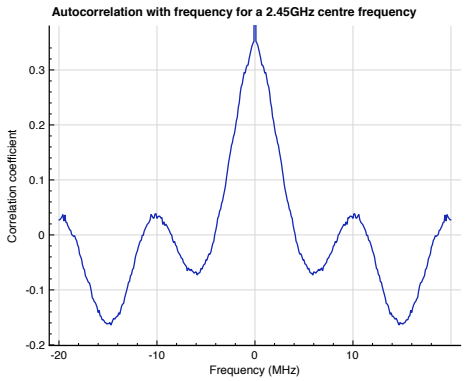


Figure 7.34: Audo correlation response for final narrowband measurement (10)

**Broadband responses for multiple bays**

3 more broadband responses are plotted in figures 7.35 to 7.37 which show the case of adjacent bays on opposite sides of the aircraft, opposite sides at a different length and a separation of 2 bays down one side. The diagonal one in this case had much lower signal strength than the others because there was no direct path as exists on a single side. Note that this was using a dipole antenna here that won't be efficient outside the 2.4GHz range. This affects the absolute levels recorded but doesn't stop one from seeing the localised patterns over the plots. Also, in a small cavity like this the concept of antenna efficiency is different to free space or larger cavities - all we need to do it excite the cavity rather than create waves efficiently. One implication of this is that wide bands could be used in an aircraft without having to worry about large antenna dimensions.

The results show a maximum S21 of -40dBs and a minimum S21 of -90dBs although not for the same measurement setup. Given that this is for transmission over no more than 3 or 4 metres it is a very dramatic difference. For the tornado there are not the large signal differences due to resonance that were seen in single cavities in the laboratory but equally large variation is seen because of the absorption and coupling losses between equipment bays and the implications of this are the same if considering flat fading - the expected signal strength is highly random. Fortunately, there was little frequency selective fading in the 2.4GHz measurements and this is backed by the broadband time responses that suggest a range of 0.03 - 0.08$\mu$s.

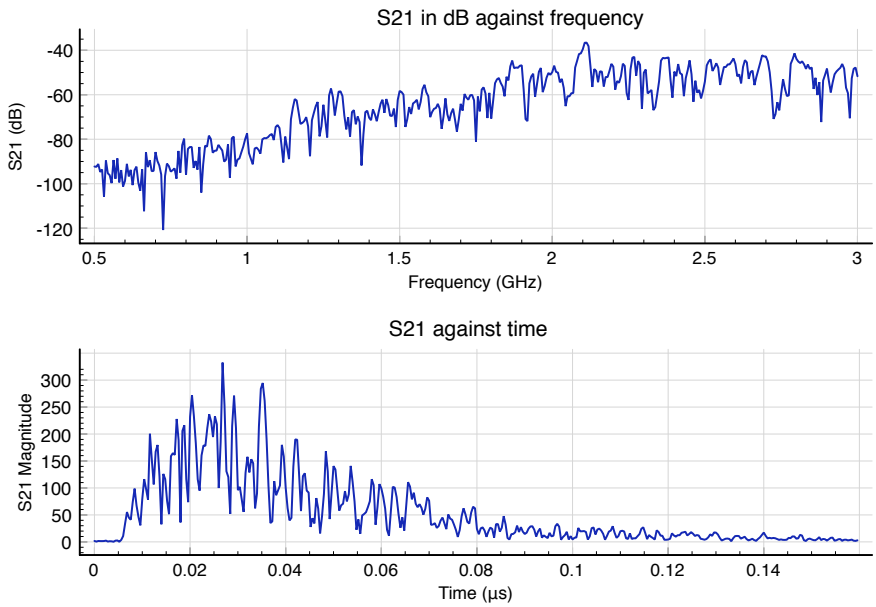


Figure 7.35: Broadband for multiple bays facing on opposite sides of the aircraft (measurement 12, same setup as 5)

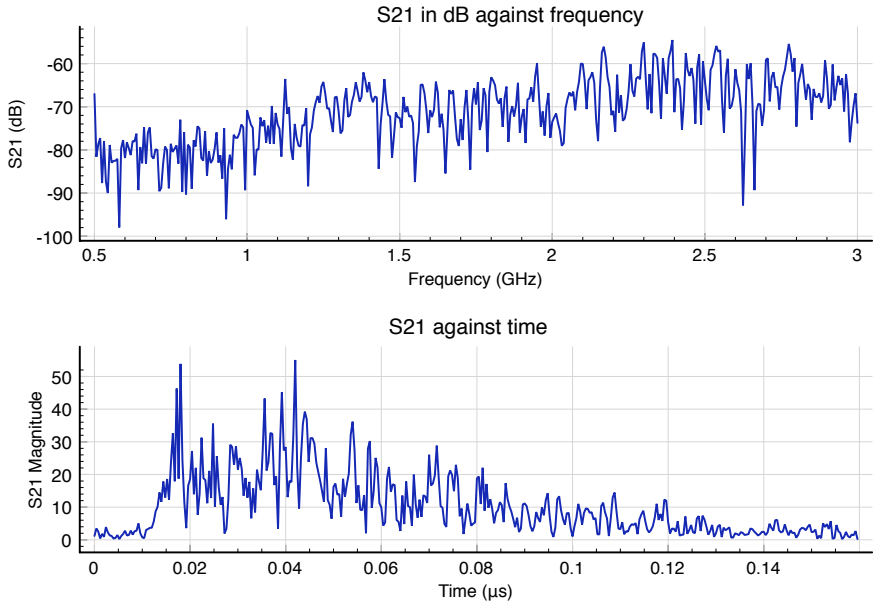


Figure 7.36: Broadband frequency time response for diagonally opposed bays (measurement 13)

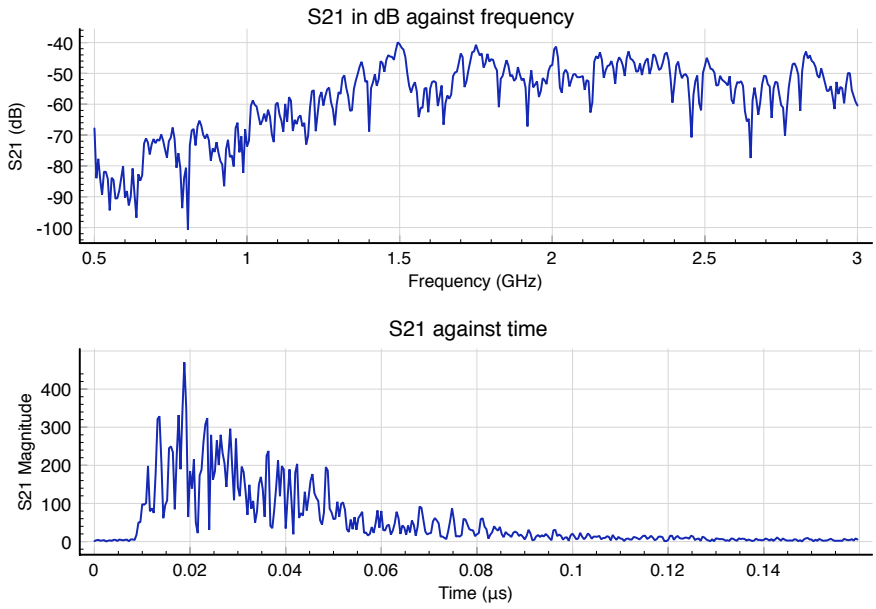


Figure 7.37: Broadband frequency time graphs for bays with an intermediate bay (measurement 14)

## 7.4 BER Results

So far the form of the frequency and time responses in an existing aircraft have been considered and it has been seen that it is a relatively benign propagation environment where there are large amount of existing wiring looms and apertures to allow the absorption of energy and damp resonances. The time response and autocorrelation change slightly between the measurements but the difference are tiny compared to those seen in previous chapters and for all measurements the time response is quick and should not cause errors in ZigBee or 802.11 systems.

To verify that the environment is consistently benign to wireless system it is necessary to record how ZigBee performed in the aircraft. Table 7.2 summarises the results.

| Measurement number | Average BER | No. > 0.01 | No. > 0.05 |
|---|---|---|---|
| 1 | 0 | 0 | 0 |
| 2 | 0 | 0 | 0 |
| 3 | 0.00005 | 0 | 0 |
| 4 | 0 | 0 | 0 |
| 5 | 0 | 0 | 0 |
| 6 | 0.00051 | 1 | 0 |
| 7 | 0.06 | 4 | 4 |
| 8 | 0 | 0 | 0 |
| 9 | 0.0002 | 0 | 0 |
| 10 | 0.73 | 32 | 27 |

Table 7.2: Average BER and number of positions (from 80) where the BER is over a limit. 1 is max BER

It is seen that for 5 of the ten measurement configurations the average BER is 0 and for 3 configurations it is less than 1 in 1000. Only in 2 configurations was the BER high in more than one stirrer configuration. Both of these cases had very high average attenuation and so the stirring was dropping the signal to the edge of the receiver sensitivity limit. For the situation with 27 high BER positions many of the errors were due to packets not being received at all rather than part of the payload being in error. These results are not surprising given the IEEE 802.15.4 specification of a 85dBm sensitivity requirement or better for 1% PER because we have seen that the highest attenuations are in the region of 80+ dBs with a transmit power of a

2.5dBm for a ZigBee radio.

It appears that there are no issue with errors due to delay spread for the ZigBee radio in the Tornado but that it is difficult to transmit down the fuselage of an aircraft where even over 2 or 3 metres there is high attenuation. For this type of environment the problem is a more traditional one of Gaussian noise but with the complication of the signal transitioning over time from a error free peak into nulls below the receiver sensitivity. What this implies is that if one wished to retrofit wireless systems onto existing aircraft the primary concern might be one of boosting the signal at regular intervals rather than one of trying to deal with delay spread. This boosting can be done by routing between nearby radios using the standard ZigBee protocol.

## 7.5 Statistics and interesting relationships

To conclude this section a few other features of the responses are presented including the statistics of the environment and how delay spread varies with other parameters.

#### Statistics

Two histograms, one for a single bay measurements and one for a 2 bay measurements are shown in figures 7.38 and 7.39. Both appear to be of Rayleigh distribution as can be seen from the fitted curves. Although the channels were highly correlated over the 2.4GHz band in comparison to some of the lab cavities, which might have changed the distribution, the results indicate there was still sufficient variation with frequency and stirrer position to create a Rayleigh distribution. Thus it is possible to continue to use the Rayleigh model of the channel used in previous chapters.

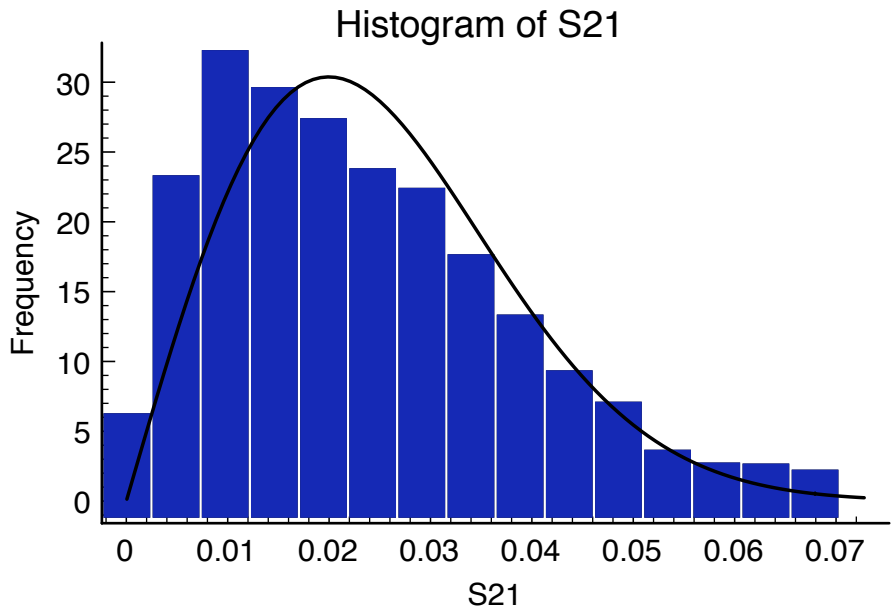


Figure 7.38: Histogram of S21 in measurement 2 (single cavity)

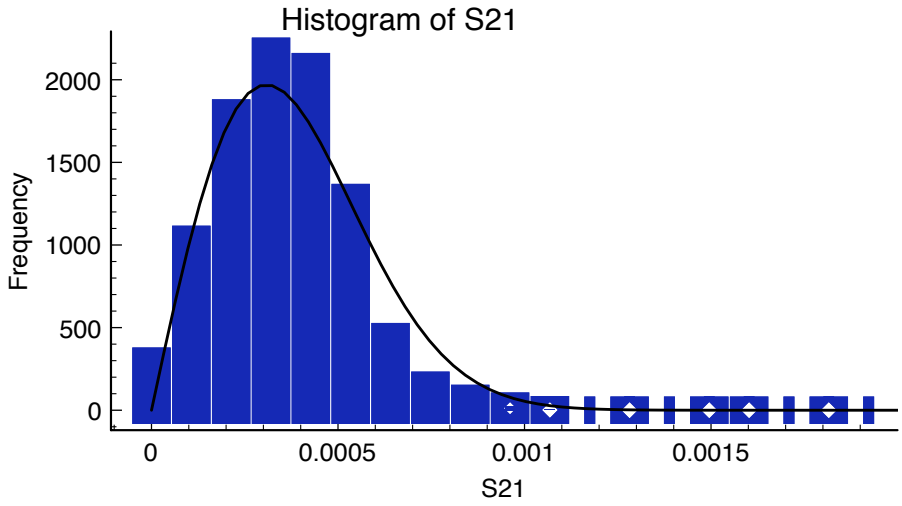


Figure 7.39: Histogram of S21 in measurement 6 (2 cavities with an intermediate cavity)

### Delay spread and coherence bandwidth

As a brief aside we shall compare the time constant and the correlation bandwidth over the 2.4GHz measurements in order to see how they relate and whether either one might be used as a substitute for the other. Guidance in the literature (e.g. Rappaport [43]) suggests that the relationship may be of the form

$$\tau = \frac{1}{KB_c} \tag{7.1}$$

where coherence bandwidth $B_C$ is very loosely defined as the bandwidth over which the signal changes and is adjusted by a constant dependent on whether the calculation is half the maximum correlation, 1/10th etc. For a frequency correlation over 0.5 Rappaport suggests a K of 5.

The results for delay spread and average coherence bandwidth are presented in figure 7.40 where the bandwidth is that over which the correlation coefficient halves, taking care to avoid artificial peaks caused by noise. As has been seen, accurate delay spread calculation was not possible but the figures here have been averaged over stirrer positions and were all calculated the same way so they may be compared even if in absolute terms they have large error bars.

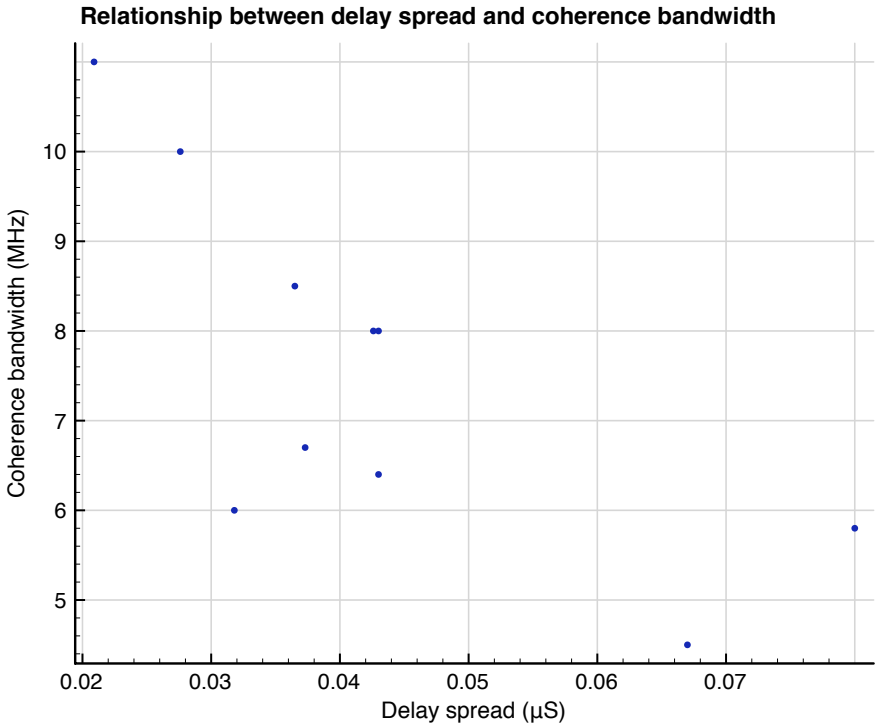


Figure 7.40: Coherence bandwidth against delay spread

One can't infer too much from this plot because the sample size is small and it would be hard to do any meaningful analysis of variance. Visually though it is possible to see that the two delay spreads under 0.03$\mu$s have a higher correlation bandwidth than those greater than 0.3 and those with delay less than 0.05$\mu$s have a higher coherence bandwidth than those with a greater delay spread. These results do at least match the theory which is reinforced because the formula above says that a delay of 0.02$\mu$s should correspond to a coherence bandwidth of 10MHz (for correlation of >0.5) which is exactly what the figure shows.

For the engineer trying to characterise a relatively low Q environment in a narrow bandwidth, average coherence bandwidth is a more practical metic than delay spread because it is easier to see than a short time response and easier to calculate automatically. That it appears to give exactly the same information is useful.

### 7.5.1 Delay and distance

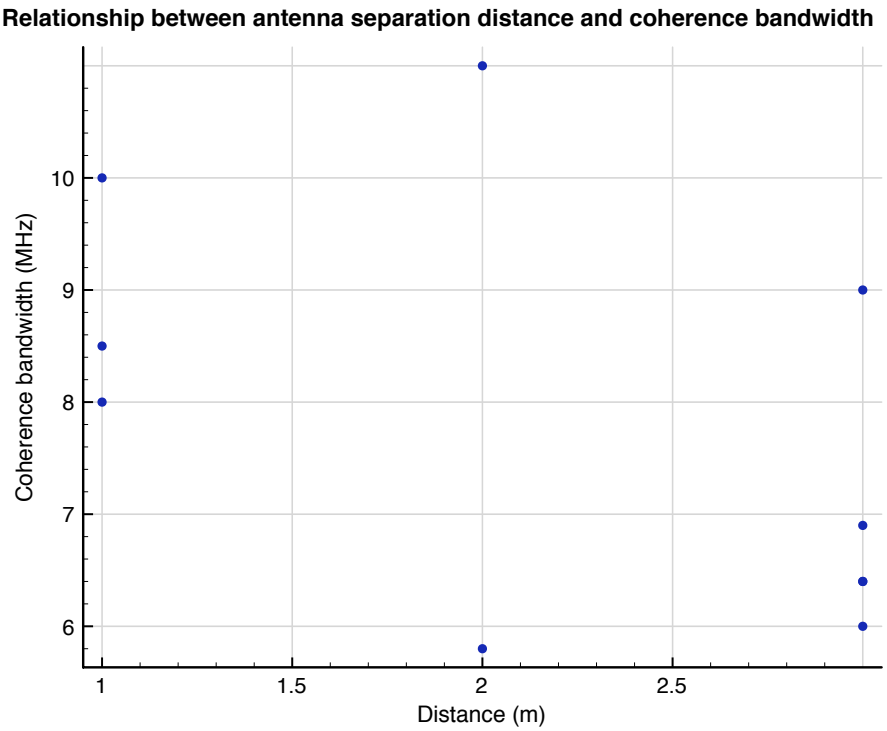


Figure 7.41: Variation in coherence bandwidth with distance (very approximately)

Finally, in figure 7.41 the variation in the time constant around the aircraft is plotted, grouped according to the antenna separation. The three levels of separation are:

1. Same cavity ( 1m separation or less)
2. Adjacent cavities ( 1-2.5m separation)
3. Distant cavities (>2.5m separation)

There is some visual evidence of a difference here, suggesting a longer response (following the initial impulse) with greater separation. It does fit with expectations that with multiple resonant cavities, the time response can increase, and also that the correlation will become narrower when there are fading effects from more than one cavity. However, the cavities in the Tornado were almost non reverberant and so the effect here is more likely to be because of the large number of potential paths through different apertures as the separation increases.

## 7.6 Conclusions

At the beginning of the wireless aircraft project it was predicted that reverberant effects could be the main cause of packet loss for wireless sensor networks. Previous chapters have presented large numbers of channel and BER related measurements for these environments to gain a better understand when and how these errors occur as well as providing useful design information for future networks. However what we have seen here, in the only measurements carried out on a real aircraft is that actually the reverberant effects are far less significant than expected. For half the measurement setups the BER was 0% and where there was a high error rate it was due to a low signal to noise level rather than inter-symbol interference.

With signal to noise being the dominant bit loss mechanism there is an unexpected problem in retrofitting a wireless system to an existing jet. One can see that it might not be possible to wireless connect two cavities down the full length of the fuselage in a reliable fashion which is a problem because the aim of using wireless is to replace long and heavy wiring looms. There are however a number of ways that this can be overcome:

- Antennas could be deliberately placed on either end of dividing bulkheads with coaxial joining them in between, thus providing stronger coupling than provided by small apertures
- The ZigBee stack could be utilised to route the 802.15.4 packets between bays rather than trying to use two node links

If measures such as these are taken then reliable propagation in an airframe like the Tornado should be relatively easy to accomplish, given the low level of delay spread and no frequency selective fading.

The measurements in the single bays suggested some energy storage but to call it reverberant is probably the wrong word because the maximum to minimum changes at 2.4GHz were so much smaller than in the previous chapters and because the time response was very rapid. Where the time response increased slightly it may simply have been down to distance and multiple paths. When antennas are 3m+ apart there will be a propagation time of 0.01$\mu$s and some path lengths may be double this leading to higher delay spread. From an analysis perspective, this chapter has demonstrated the clear relationship between delay spread and coherence bandwidth. It has been shown that the latter is especially useful where the time resolution is low and easily affected by the filtering inherent in the FFT operation. Where wide bandwidth sweeps can be done it has been shown that this can place a useful bound on the maximum delay spread.

All of this raises the question as to whether our laboratory setup is invalid and the most important areas to analyse would be noise sources in aircraft and attenuation. The stepping of the stirrer did change shape and magnitude of the channel response between measurements but all changes were far smaller than in the laboratory just as the packet errors were far fewer. Undoubtedly the results suggest that more time should be given to testing systems at low Qs and that if wireless components were to be retrofitted into a Tornado there would be no problems with reverberant effects. However, it should be remembered that the Tornado is an older aircraft with a lot of wiring looms throughout its bays. If these were to be removed and replaced with wireless radios then the Q of the aircraft would increase and so the more wireless connections replace wires, the more difficult the wireless environment could become. A similar effect may occur as actuation changes; if the hydraulics are all replaced by electric actuators or even flapless flight then some of the apertures will disappear between cavities and the Q-factors of the cavities may increase. Also there are some parts of the aircraft that have to be sealed for fuel or air and leakage may be much lower in these; unfortunately it was not possible to test any such bays on the Tornado. In conclusion, the testing on the Tornado has given a lot of useful information about the electromagnetic responses of existing wired aircraft but we should not design wireless systems assuming future aircraft will be the same.

# Chapter 8

# Alternative radio technologies - the IEEE 802.11 standards

## 8.1 Introduction

Throughout the previous chapters ZigBee has been used to illustrate the difficulty in transmitting data in a resonant environment. Now the scope of this research will be expanded to look at the IEEE 802.11 technologies, often referred to as Wireless LAN or Wifi. There are 3 main variants that we shall test with the postfixes b, g and n. These offer different modulation techniques and so may react differently to a resonant channel. They all present the same interface from the MAC protocol layer upwards. This chapter will concentrate on the physical layer and only consider higher level issues as they affect the physical layer.

Up to this point Wireless LAN has been ignored because the more energy efficient and widely applicable ZigBee standard has been the focus. However, carrying out wireless LAN measurements offers three benefits:

- It allows assessment of how the higher data rate systems would perform in an aircraft
- It provides a comparison point for ZigBee. By providing diversity in the systems tested it is possible to be more certain that results seen previously were not due to peculiarities of the particular hardware used.
- WLAN has been suggested in Chapter 2 as forming part of a two tier network architecture with Zigbee. Therefore it is important to know how it performs

**Benefits**

Wireless LAN (IEEE 802.11 in particular) technologies include a variety of features designed to boost wireless data rates to levels whereby real time video streaming, gaming etc are possible and a large number of users can be supported. They are used throughout the home, office and public spaces. Therefore these are mature technologies at the cutting edge of performance and could provide a basis for high performance wireless aircraft systems just as Firewire and Ethernet do for wired systems. Importantly, between them they cover a number of modulation and coding techniques and so, as a measurement platform they have relevance beyond Wireless LAN, offering an insight into the benefits of MIMO, OFDM and other modern communications techniques. Results are divided by the types b,g, and n as each has some unique features to further boost data rates.

### 8.1.1 Aims

The primary aim of this chapter is to check the performance of high specification DSSS,OFDM and MIMO systems in a reverberant environment. This will provide a useful contrast to the simpler 802.15.4 radios seen earlier. The secondary aim of this chapter is to provide an overview of some of the radio technologies available in COTS systems at the physical layer, as they relate to reverberant environments.

## 8.2 Measurements and results

Chapter 2 discussed the basics of some existing wireless LAN systems and found that they include a wide variety of technologies including ODFM, multiple antennas and advanced forward error correction coding. It is beyond the scope (and the available equipment) of this project to take each one of these and test or simulate it in isolation in different resonant environments. Instead in this chapter measurements are taken of a COTS 802.11 system, similar to those done with ZigBee in order to find out how systems with much more advanced modulation and coding than ZigBee can perform, providing an upper limit on performance if power requirements and radio cost are no object.

Figure 8.1 shows the measurement setup and there is an access point (AP in the figure) and a station communicating. For the wireless base station a Senao/Engenius 9710 router was used and for the end device a Senao/Engenius 9701 usb device with Senao being chosen be-

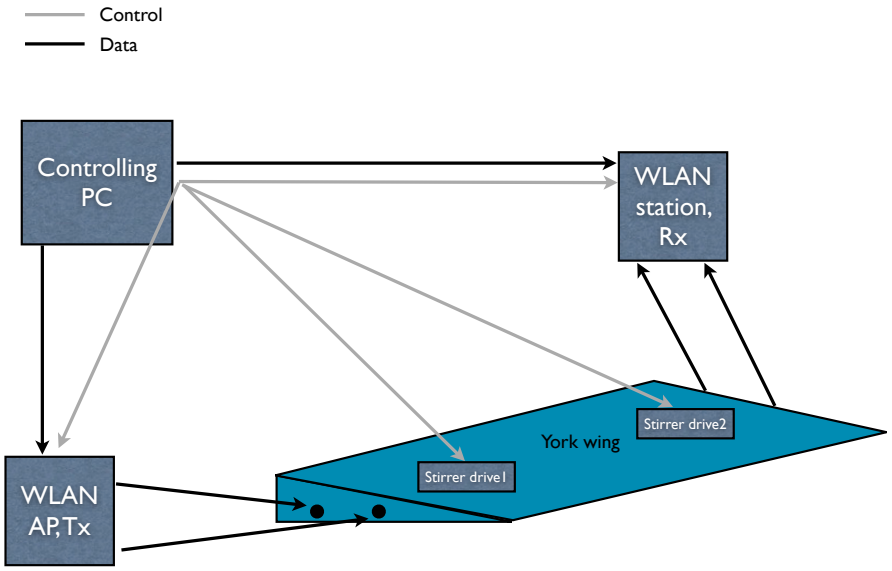


Figure 8.1: Wireless LAN measurement setup

cause they have USB stations with external SMA antenna connectos, as required to get into the wing through bulkhead SMA connectors. There was no third station as is typically the case in a WLAN network, i.e. there was no setup with a station to station transmission meaning there were no 2nd hops. This was done for the purpose of simplicity in order to see how a given link looked at any time. One might imagine that a more realistic scenario might involve a base station and then two users of the network that need to communicate such as an actuator and a sensor with some distance between them. In this instance there are two possible points of failure at a given time as in a resonant environment each path's fading is statistically independent. Such a scenario, although realistic would not allow a close comparison to the ZigBee measurements seen earlier. In order to link to the laptop, the station used a USB2 connection and for the access point to the laptop gigabit Ethernet was used ensuring that the links to the laptop were not limiting factors.

There is a direct link option in WLAN which could be employed and this is known as an ad hoc network in IEEE802.11. Unfortunately this is not widely supported and is only required for 802.11b. Therefore it was decided that this route would not be pursued and the access point-station combination was chosen. Although the diagram only shows 1 computer, actually this was split into two to avoid issues of the operating system trying to bypass the network, although logically the setup was as shown. Python scripts controlled the whole process through

UDP sockets. The system worked by a controlling computer sending a message over ethernet to a second computer telling it to send a set of packets. This second computer sent the packets over ethernet to the wireless router which then sent the packets wirelessly back to the controlling computer. 2000 packets were sent per stirrer positions and each packet contained 1024 bytes and the number successfully received were recorded. This is not sufficient packets to see noise problems but was found to be sufficient to see inter-symbol interference based errors. It was not possible to control the setup in such a fine grained way as the ZigBee one as there was no option presented by the device or its drivers to change the MAC retransmissions.

Clearly this was only a single system and other wireless routers might behave differently. Time restrictions meant that only this could be used. However any system must conform to the physical layer specification and have the same symbol time, guard time (for b and g), coding and modulation which is all critical to multi-path performance.

All wifi tests were carried out on the wing like structure which is at the medium to low end of Qs considered in previous chapters. From 7 it would be expected that Qs on a real aircraft would be lower than those in York's main reverberation chamber and so it was felt that the aluminium wing like section would provide a good balance between medium and low Q. There wasn't the time to carry out performance measurements on such a wide range or setups as for ZigBee, not helped by the fact that there is not a single data rate for Wireless LAN, but there are in fact a wide variety of transmission standards to test. Finally, alongside the wireless LAN measurements network analyser (NA) measurements were used to check the delay spread for each configuration.

## 8.3 No loading

The results are divided by the loading in the wing and the version of 802.11 used starting with no loading and 802.11b.

#### 802.11b Measurements

BER against stirrer position was measured for all 4 data rates and the results are plotted in figures 8.2 to 8.5 As with ZigBee errors appear to come in spikes rather than a gradual change as the stirrer rotates. There is no difference between the 2Mbps and the 1Mbps system. On this particular run the lower data rate system has performed worse but this was within the

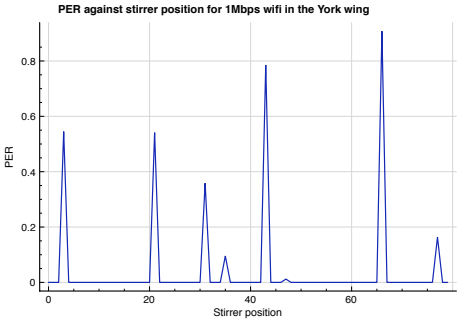


Figure 8.2: 802.11b PER in the empty wing at a data rate of 1Mbps

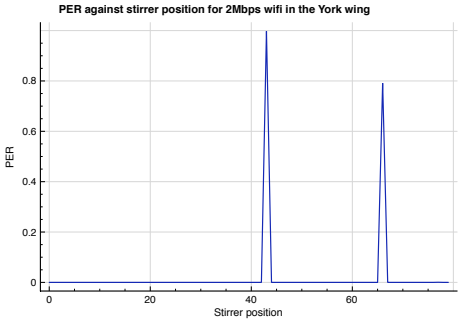


Figure 8.3: 802.11b PER in the empty wing at a data rate of 2Mbps

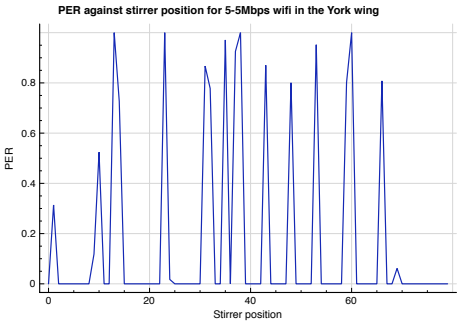


Figure 8.4: 802.11b PER in the empty wing at a data rate of 5-5Mbps

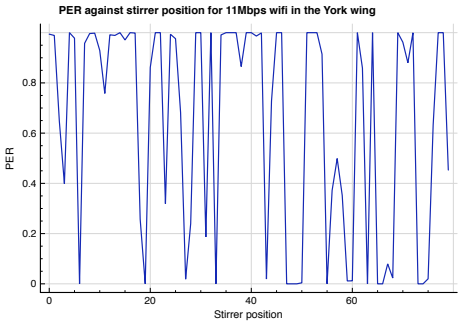


Figure 8.5: 802.11b PER in the empty wing at a data rate of 11Mbps

error bound for the measurement. The stirrer was not always at the exact same positions and any deformity in the wing would lead to large PER changes. When moving to the higher data rate system the difference is far more pronounced. At 11Mbps there are a majority of stirrer positions where the system does not work and at 5.5Mbps there are still nearly 1/4 of positions where the error rate is greater than 0.5%. In all cases the error rates are higher than for ZigBee where typically only 1 in 80 positions in the wing were found where the PER was over 0.1. At 1Mbps the performance would be a concern for a high date rate communications link between the cockpit and flight computers and at the higher rates the system would be unusable.

The corresponding delay spreads for this measurement were approximately 0.26$\mu$s which suggests a Q of approximately 1000. As reported earlier, the chip period for 802.11b is 0.09$\mu$s and the bit period is 0.09$\mu$s for the 11Mpbs case. For the 1Mbps case the bit period is $10^{-}6$, which is 4 times longer than the delay spread. Although all 4 transmission rates have the same chip rate the bit period is most important because when there are less possible bits per se-

quence of chips the detector's job of deciding the most likely bit given the received sequence is simpler. There is a difference between the 2Mbps and 5.5Mbps rate in terms of the DSSS scheme used but this is less likely to be the cause of the large change in PER than the move to a bit period that is less than the delay spread.

### 802.11g Measurements

Next the equivalent measurements for 802.11g are presented, again for a range of data rates. The results are shown in figures 8.6 to 8.9. At the low data rates this performs much better than 802.11b with no positions of PER approaching 0.5 and only 3 positions with a PER close to 0.1 for 6Mbps. It is slightly worse for 9Mbps but this is still much better than the 11Mbps 802.11b rate that was receiving no packets in half the stirrer positions. Once the data rate increases to 18Mbps the PER pattern is similar to 802.11b at 11Mbps alternating between working or having 100% PER. The symbol time is the same for all of these as is the guard interval and the number of carriers so none of these are leading to the higher error rates. Modulation is where they differ because 6 and 9Mbps use BPSK, 18Mbps uses QPSK and 24Mbps uses 16QAM. The guard interval is 800ns for this modulation, designed to be 4x a maximum delay spread of 200ns. For the wing section the delay spread was slightly higher than this on average with some stirrer positions reaching 400ns delays. As with the 802.11b measurements it would appear that the system can cope with distinguishing between 2 possible symbols or sequences but when this moves to 8 or 16 the high delay spread has an effect and error rates rapidly increase.

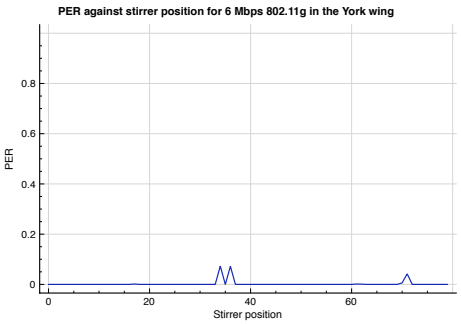


Figure 8.6: 802.11g PER in the empty wing at a data rate of 6Mbps

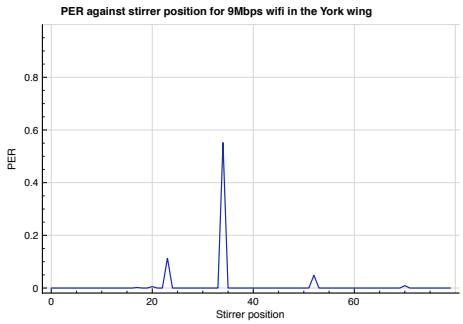


Figure 8.7: 802.11g PER in the empty wing at a data rate of 9Mbps

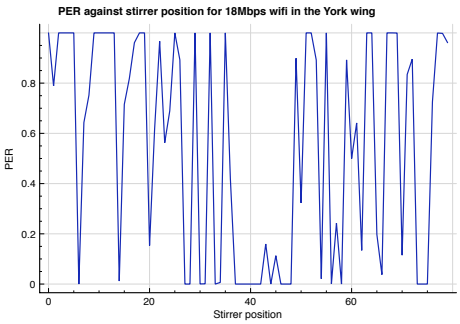


Figure 8.8: 802.11g PER in the empty wing at a data rate of 18Mbps

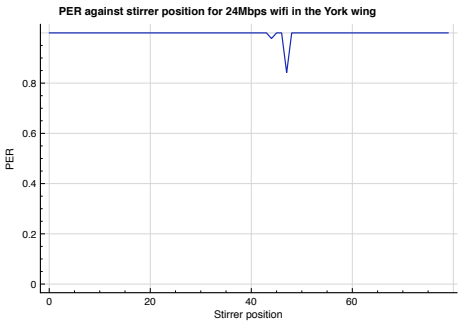


Figure 8.9: 802.11g PER in the empty wing at a data rate of 24Mbps

### 802.11n SISO

To complete the measurement set in the empty wing, measurements for 802.11n are presented. These are only using one transmit antenna and one receive antenna. The 820.11g pattern is repeated which is to be expected since the modulation is the same.

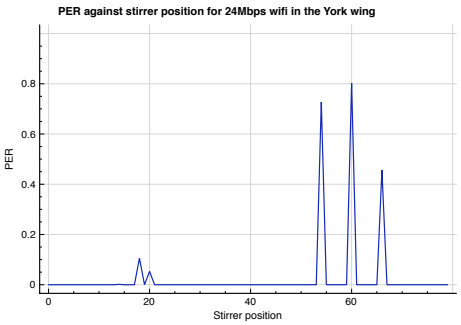


Figure 8.10: 802.11n performance in the empty wing at a data rate of 6.5Mbps

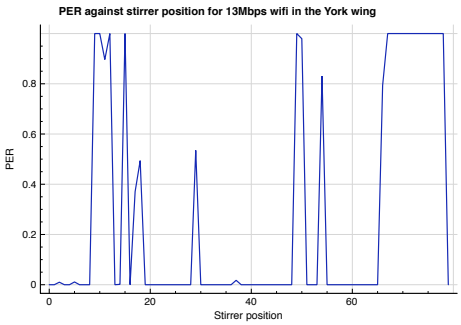


Figure 8.11: 802.11n performance in the empty wing at a data rate of 13Mbps

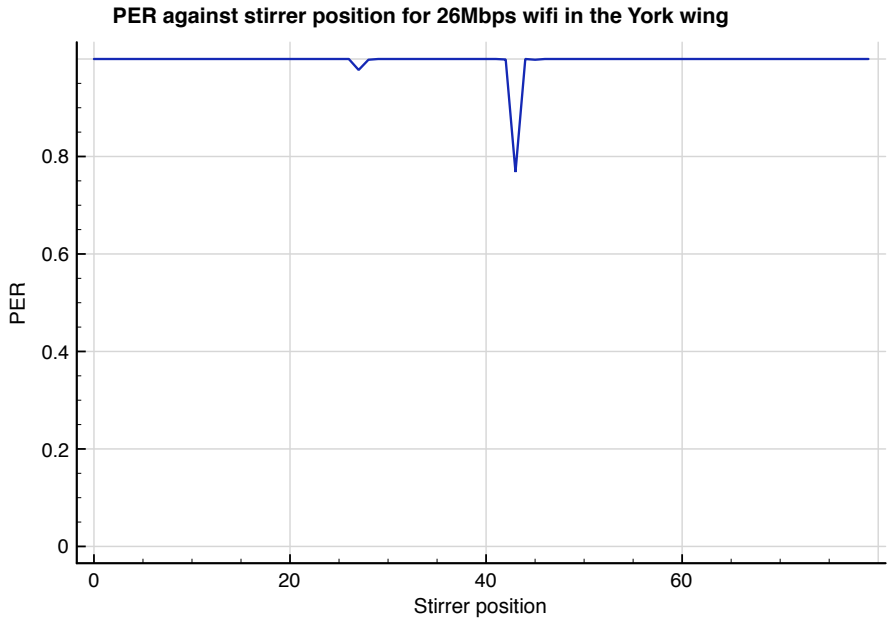


Figure 8.12: 802.11n performance in the empty wing at a data rate of 26Mbps

### 8.3.1 Results for a loaded wing

Further measurements were carried out with a medium level of absorber (relative to the maximum used in previous measurements) and an average delay over stirrer positions of 0.08$\mu$s.

These are presented here:

### 802.11b

The situation is rather better now that the delay spread is 1/3 what it was before. It is now possible to send packets at the maximum data rate of 11Mbps without errors for 80% of the stirrer positions. This is starting to approach a position whereby a reliable system could be created by adding redundancy. Considering the fact that the delay spread is of the order of the symbol time the link is surprisingly reliable.

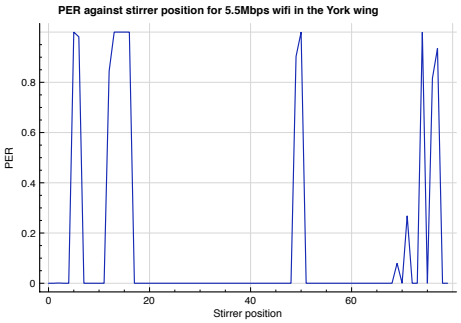


Figure 8.13: 802.11b performance in the empty wing at a data rate of 5.5Mbps

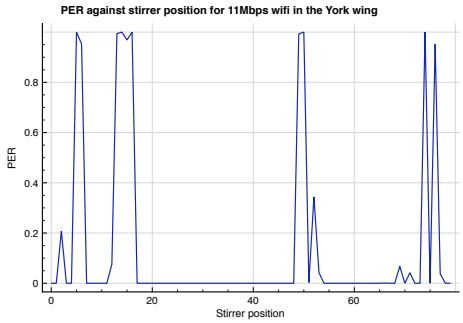


Figure 8.14: 802.11b performance in the empty wing at a data rate of 11Mbps

### 802.11g

For the g system there are no positions with any packet errors now until the 18Mbps link speed, which is the transition to QPSK. Now that the delay is less than 1/8 of the guard interval it becomes possible to communicate reliably using OFDM at data rates over 10Mbps. However it is again not possible to go beyond this without the risk of finding a position where the system does not work. The top rates from 36Mbps to 54Mbps are unusable.

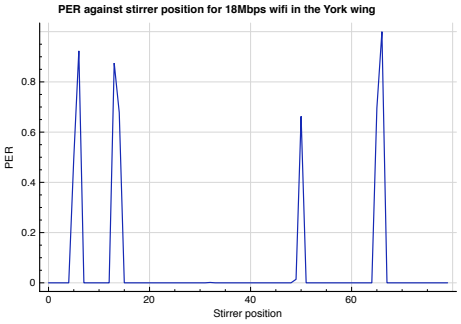


Figure 8.15: 802.11g performance in the empty wing at a data rate of 18Mbps

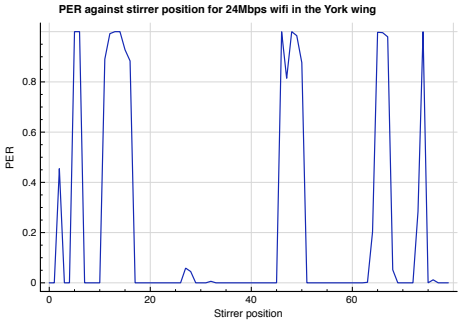


Figure 8.16: 802.11g performance in the empty wing at a data rate of 24Mbps

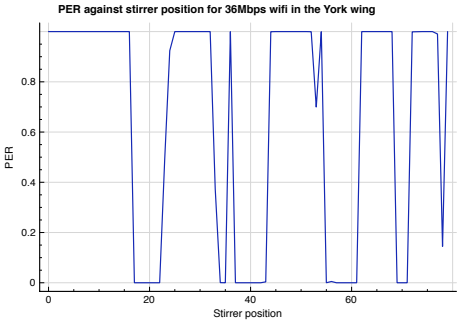


Figure 8.17: 802.11g performance in the empty wing at a data rate of 36Mbps

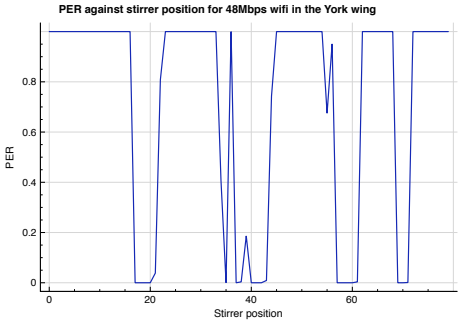


Figure 8.18: 802.11g performance in the empty wing at a data rate of 48Mbps

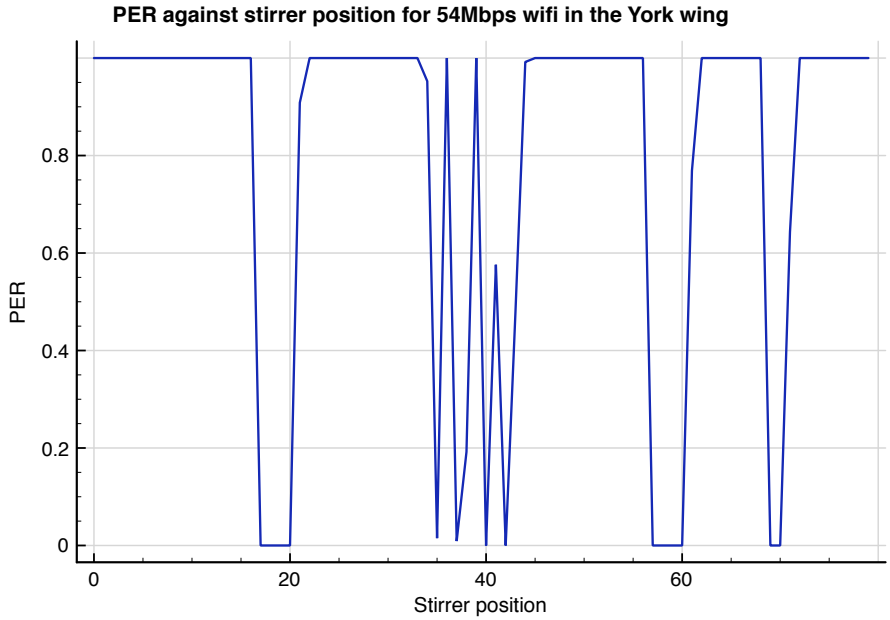


Figure 8.19: 802.11g performance in the empty wing at a data rate of 54Mbps

### 802.11n

The PER pattern for 802.11n was almost identical to that in 802.11g as shown by this example at 26Mbps which is very close in pattern and average PER to the case of 24Mbps g.

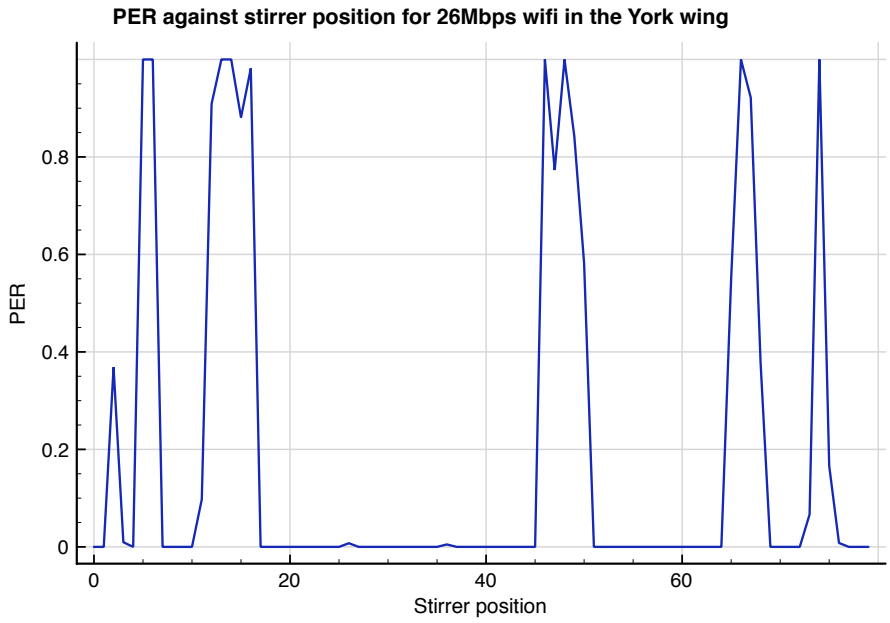


Figure 8.20: 802.11n performance in the empty wing at a data rate of 26Mbps

### Summary

Figure 8.21 summarises the measurements shown so far by plotting average PERs against data rates. It shows that after 10Mbps the PER is relatively lower for the lower Q but it also shows that average PER is high for both cases. Although the delay spread is below that for which 802.11g was designed it is nonetheless unreliable once the modulation scheme moves to QPSK.

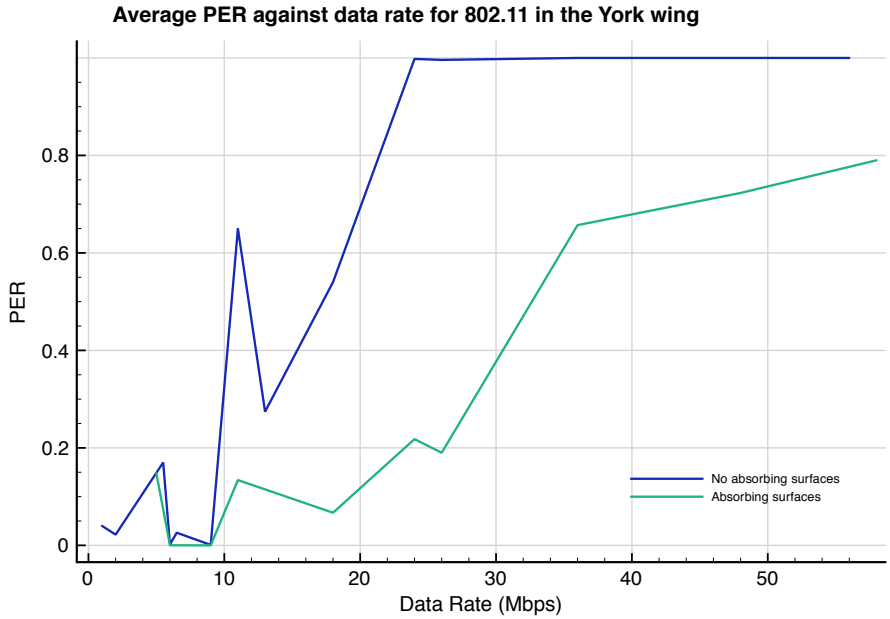


Figure 8.21: Summary plots for all 3 802.11 standards of the stirrer average PER for an unloaded and loaded wing, with delay spread 0.26$\mu$s and 0.08$\mu$s

### 8.3.2 Multiple antennas

Multiple antennas systems should exhibit better data rates and reliability by adding diversity to the receiver and by allowing additional data streams over the different possible transmitter to receiver paths. Next the PERs are presented for the b,g and n systems with multiple antennas. Note that for the b and g standards the extra transmit antennas are both transmitting the same data so as to give the receiver a better chance of decoding to the correct bits. For these measurements use a setup of an unloaded structure, without AN79 absorber.

The results for 802.11b are shown in figures 8.22 to 8.24. Comparing this to figures 8.2 to 8.5 then a clear improvement can be seen, particularly at 5.5Mbps where there are half as many positions with PERs over 0.5.

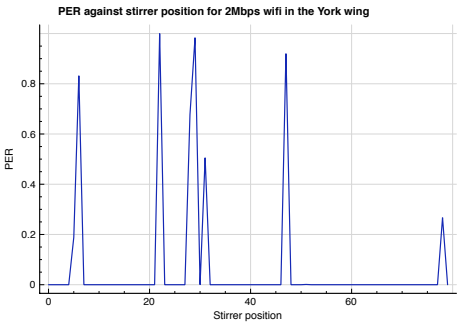


Figure 8.22: 802.11b performance in the empty wing at a data rate of 2Mbps with antenna diversity

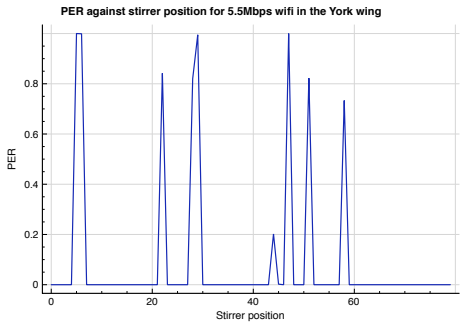


Figure 8.23: 802.11b performance in the empty wing at a data rate of 5-5Mbps with antenna diversity

Moving on to a g setup we see in figures 8.25 and 8.26 that at the lower data rate there is very little difference between the MIMO and non MIMO system for the 9Mbps rate but a very significant improvement for the higher date rate where the system was unusable previously for half the stirrer positions and is now unusual for 7 positions out of 80.

At the data rate of 36Mbps there is no position at which it was possible to send data with neither antenna proving successful in picking up a useable signal and so this is not shown. MIMO diversity can't offer any improvement when all combinations of the antennas and stirrer are unusable.

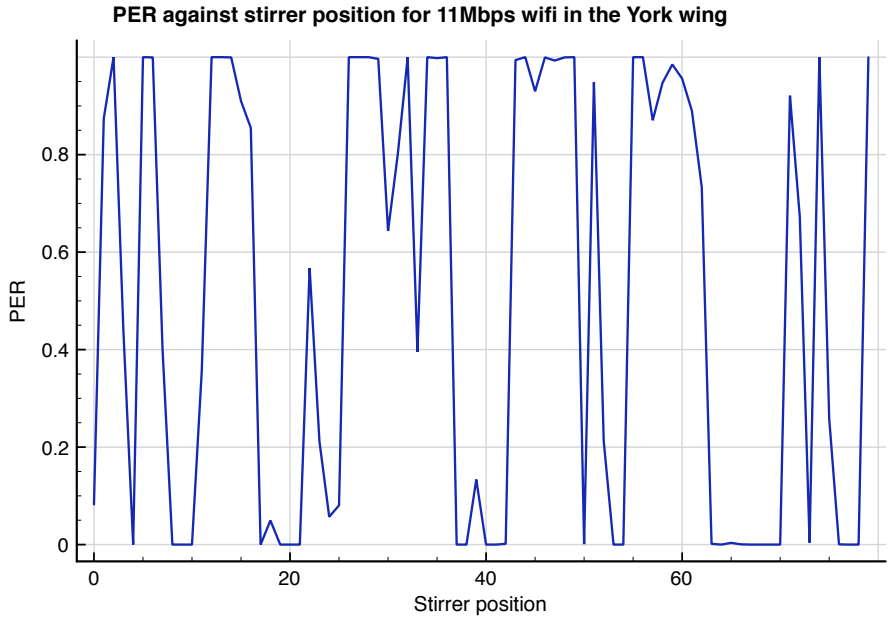


Figure 8.24: 802.11b performance in the empty wing at a data rate of 11Mbps with antenna diversity

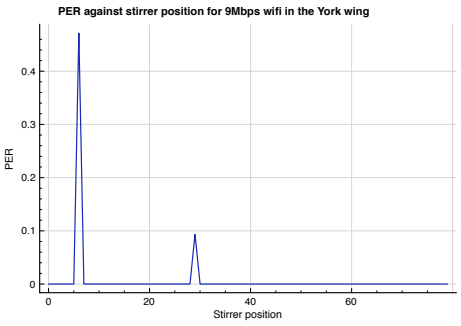


Figure 8.25: 802.11g PER in the empty wing at a data rate of 9Mbps with antenna diversity

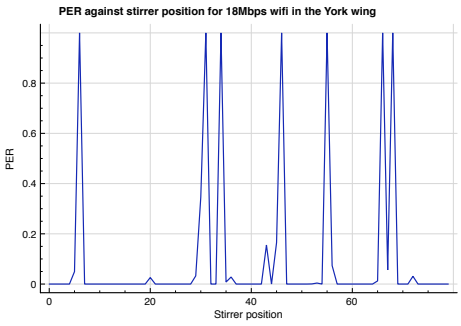


Figure 8.26: 802.11g PER in the empty wing at a data rate of 18Mbps with antenna diversity

### 802.11n

For n we have compared two types of diversity, both operating at the 26Mbps rate. One is operating through diversity and the other through increasing capacity over the two routes as described in the introduction and the plots are shown in figures 8.27 and 8.28 respectively. It can be seen that although many positions of the stirrer result in 100% PER 1/4 of positions allow error free transmission for the case of identical streams whereas for the dual streams the PER is 100% as seen previously.

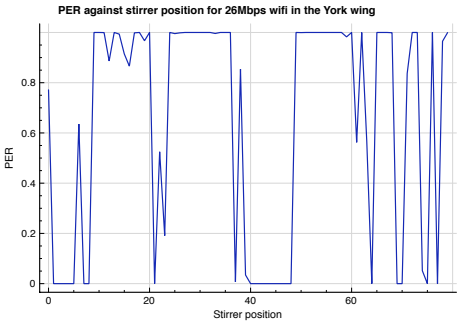


Figure 8.27: 802.11n performance in the empty wing at a data rate of 26Mbps with antenna diversity

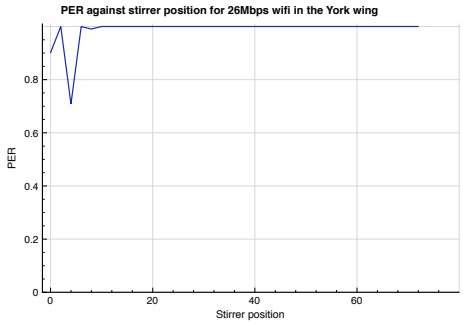


Figure 8.28: 802.11n performance in the empty wing at a data rate of 26Mbps with 2 data streams

## 8.4 Performance with a moving stirrer

When looking at complications for ZigBee measurements were presented with a continually moving stirrer. It was found for the particular ZigBee system tested that a moving channel could severely degrade performance such that average PER and BER was higher than that seen averaging over a number of static positions suggesting the shifting phase was adding to the error rate. Therefore it would be interesting to see what happens for wireless LAN as the speed varies.

3 configurations were used to check the effect of a moving stirrer. The moderate loading configuration was used because in this configuration errors were occurring but only at some positions meaning there would be a better chance in seeing an increase in errors. The results can be seen in figure 8.29. There is no sign here of PER increasing with stirrer rotation rate. There is a variation with speed but it is only the uncertainty of the measurement. The pat-

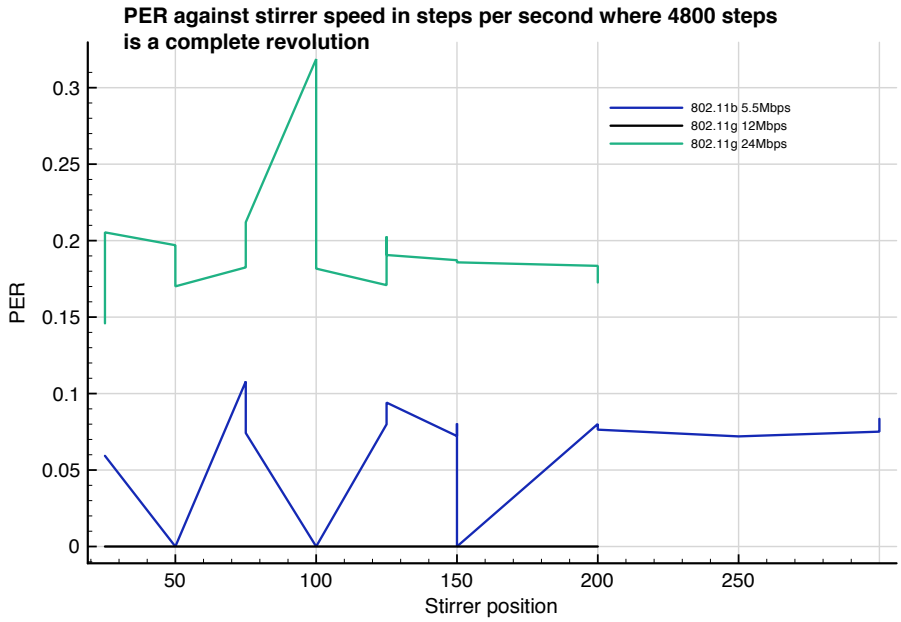


Figure 8.29: 802.11 performance in a time varying channel

tern is reminiscent of that seen for ZigBee at lower Qs which shows that for all Qs that will be encountered realistically, a time varying channel won't introduce any effects beyond altering the signal to noise ratio per symbol. For the slower measurements the stirrer did not rotate a complete revolution which may explain the oscillatory nature of one of the results sets.

## 8.5 Conclusion

This chapter has demonstrated the performance of IEEE 802.11 wireless LAN standards in a reverberant, aircraft cavity like environment, to find out whether the additional processing and different modulation systems result in different behaviour to ZigBee.

The performance of the 802.11b standard was as expected. This did not perform well in both the loaded and unloaded wing measurements with positions of 100% PER even for the lowest data rate. For this standard the mean delay time was high compared to the symbol time and there is no special protection against delay spread except through the correlation properties of the chip sequence resulting in a high number of positions have an unusable link after 2Mbps.

Perhaps more surprising was the performance of 802.11g which managed to work at the maximum rate of 802.11b without errors but then introduced positions of 100% PER from

18Mbps. There did not appear to be a problem with synchronisation failing at the higher data rates since higher and lower rates use the same rate and sequence for synchronisation. The data rate is increased by adding more points to the constellation and/or changing the coding rate so it appears the problem lies here. What is most likely is that the delay spread in the wing was causing some bits to be received in error and these were recovered by the FEC for BPSK. When the constellation got larger with QPSK, then QAM, this lead to many more symbol errors and therefore bit errors, beyond the number that the FEC could cope with. All the modulation schemes used in 802.11g (unlike ZigBee which is an FSK/PSK hybrid) require accurate phase information and it may be that even with the guard interval the delay in the wing was preventing correct decoding of signal phase.

It was seen that when multiple antennas were used there was a definite improvement in PER performance, particularly at the higher data rates. Why it did not eliminate all packet errors at the low rates is not clear but to understand this would require further knowledge of how the signals from the two antennas are combined. If the diversity is designed to work with a weighting according to signal strength it will have limited effect where there is a delay spread problem. The n standard was found to bring few benefits despite being based around multiple antennas and reflections. It can only bring high data rates if the delay spread is low and the OFDM demodulation is successful. Where this is not happening then the data rates cannot increase. Worse is the comparison between 26Mbps rates where with maximal ratio combining (2 identical data streams) was compared against spatial division multiplexing with two 13Mbps data streams and the former performed better. The reason for this is unclear but it is probably due to the coding strategy and the way data is spread over the two streams. If one antenna receives a high delay spread and can demodulate no data then the whole system will fail. The data rate can only be increased when the system demonstrates multiple paths, but low delay spread.

Overall, 802.11 has been shown to offer data rates much higher than ZigBee, being able to achieve 10Mbps reliably with OFDM. At the same time the data rates achieved are far below what was expected. In Chapter 2 it was suggested as a replacement for the 100Mbps cables in aircraft that form the network backbone. In real aircraft we may see better than the results of this chapter because delay will likely be lower as Chapter 7 indicated. If we want to go to higher data rates for delay spreads of 0.26 and 0.1 $\mu$s then changes to the standard may be needed such as using an even greater guard time. One could also tailor the way diversity is

used in order to obtain the signal least affected by time delay rather than combining signals to maximise signal to noise.

# Chapter 9

# Conclusion

When the 3 years of the Flaviir wireless project started there was almost no literature in the area of fly-by-wireless aircraft and for many people the first reaction to the idea of wireless aircraft was skepticism. Mobile devices can often demonstrate unreliability in the form of mobile phones dropping calls or Bluetooth connections that perform poorly. The idea of making all the components in an aircraft wireless when the consequences of failure would not be a failed download but a crashed plane is thus difficult to justify. However, the benefits of moving to wireless are significant and so one has to look at it objectively and try to ascertain whether or not wireless systems in aircraft can be made sufficiently reliable, given the safety critical requirements. Perhaps the largest obstacle in the path to a wireless aircraft is skepticism as to whether it can be sufficiently safe and so this had to be addressed.

There are a great many areas that one could look at in the design of a wireless aircraft but this PhD Thesis has focused mainly on that of propagation and reliability. It could have focused on topologies that would fit with avionics architectures or on resistance to jamming or lightning protection. All of these are interesting areas but they are perhaps for the next stage of research. The fist question to address is that posed above: if extreme factors are excluded can an aircraft's avionics be reliably made wireless. In this section these issues will be addressed.

## 9.1 Practical restrictions

It would be reasonable to say that if one wanted to communicate with absolute reliability over a given link in an aircraft then one could. Broadband radios and antennas could be used with very powerful amplifiers to send out heavily coded signals almost guaranteeing that the signal

would be wider band and more powerful than any interference, and that the data signal would be so heavily spread that any frequency selective effects of the channel would be averaged out. Unfortunately, this is not practical. As was explained in Chapter 2 there are numerous sensors all over an aircraft from control loop sensors like those for air speed, sensors to monitor fuel, sensors to check that flaps positions, to check landing gear and to check temperatures and other components. Then there are also all the high data rate links that transfer data between imaging devices, computers, radar etc. As the number of devices that one could make wireless goes up, so does the need to keep the cost of those devices low.

Further adding to the cost pressure is a desire that whatever solutions are found should be applicable not just to 100 expensive fighter jets but that they be applicable to UAVs of various shapes and sizes. This leads to another requirement, that any system be scalable from a small UAV to a large commercial passenger carrying or transport jet. One could of course use different standards for different aircraft types but if this was not necessary it would be better not to do so. As far as frequency bands are concerned then it is likely that special bands could be released for the use in aircraft, given that a radio at 30,000 feet transmitting at 1mW is not a serious interference threat. However, this may not be desirable for the smallest UAV, being operated close to the ground which might have to utilise one of the existing free bands.

When all these factors are put together the clear conclusion is that it makes sense to investigate off the shelf solutions, COTS systems. As was explained in Chapter 2 this is already done in wired avionics so it makes sense to do it in wireless. Issues of scalability, reliable transmission, low power have already been addressed here and because they are widely used, the unit prices are much lower than they would be if only used in certain types of aircraft. .

The other big technical question mark was over the use of power lines or batteries. As has been seen, if power lines are used then designing a wireless aircraft is made simpler because higher transmission powers may be used at wider bandwidths and complex coding can make a link more reliable. At the same time, not using batteries is removing half the advantage of going to wireless. Related to this was the question of data rates because at the moment a battery powered device simply cannot deliver the bandwidth required to replace the high data rate applications in an aircraft. It was decided in the early phases of the research that battery powered technology should be the focus to deliver maximum weight savings, installation savings, and ease of maintenance. There are many more low data rate links on an aircraft than high so this makes sense. The issue of replacing batteries has largely been ignored not because it isn't im-

portant but because it may not be important in the near future. Growth in phones, tablets and all manner of mobile devices is likely to solve some of these issues for us by the time wireless aircraft become reality. For the ZigBee radios used in the tests in the prior chapters, battery power could last years on low duty cycle applications.

Out of all the available technologies it has been shown that ZigBee is the best suited to the fly-by-wireless application. The use of ZigBee has in turn forced a requirement of a dual technology system in order to accommodate high data rate transfer where required. 802.11g and n were decided on as being the most suitable standards for this because of their high data rates and built in resilience to multi-path delay. One of the most important reasons for recommending that 802.11 and ZigBee are best suited to this application is their inherent ability to scale. Some of the other wireless standards are designed for devices in close proximity and don't do this so well. Part of the benefit of these COTS solutions is their inherent ability to have new nodes join a network and to re-route data frames when links fail.

Given that ZigBee is regarded as the best way forward for creating a fly-by-wireless aircraft our original question thus narrows slightly to one of can ZigBee be made reliable and the only way to address this is through finding out what the wireless aircraft channel looks like and how ZigBee performs in it.

## 9.2 The wireless channel

Trying to understand just what the channel looked like (in terms of time and frequency responses) and how it affected wireless systems became the major focus of research once questions of radio systems were dealt with. Aircraft must be reliable and be predictably so with calculable percentage chances of failure and of time to failure. There was an expectation early on that an aircraft could be a hostile environment, not just because of noise sources on the aircraft but because of its structure, with many metal boxes that would be resonant.

### 9.2.1 Tornado measurements

Late into the project measurements were done on a Tornado aircraft and these suggested that a different problem, that of high attenuation over the aircraft length. There was very little information in the literature on aircraft Qs and no one appeared to have attempted to describe an aircraft statistically using electromagnetic stirrers so it was impossible to be certain of the

channel.

In the Tornado measurements the aircraft looked non resonant and it was easy to see why, by looking at what was inside the bays. Most were packed with electrical cables and had many apertures. Two issues are raised: can ZigBee deal with these problems and is it a mistake to have devoted so much time to resonant environments. On the former question ZigBee was shown to perform very well except over the very longest and most attenuated distances tested where the PER reached levels that were too high to be usable for this application. One would therefore conclude that ZigBee could be used on an aircraft with these characteristics because it works without any errors for shorter distances and the routing capabilities in the standard to could be used to take packets further distances. The modelling of the routing, and proving link guarantees in the advent of a failure for a given traffic load are non trivial. In principle though, from a physical layer perspective it could be done. The channel measurements backed up the PER measurements showing a flat channel with a high coherence bandwidth in all configurations of measurement.

Returning to the other question about whether one should research higher Q cavities given the results found on the Tornado, the answer has to be yes. The Tornado was full of leaky apertures and wiring looms and the direct result of making an aircraft wireless is that these looms will go and the bay Q-factor will increase. There are other factors to consider as well such as the types of bay used in these measurements which were generally side bays with a lot of equipment in, not fuel tanks, air tight fuselages or other such structures. It might also be that many bays are deliberately partially sealed electromagnetically in the future so as to avoid deliberate interference that might jam a radio system. Overall the Tornado measurements have revealed a lot about propagation and reliability on modern jets but one must not read too much into them and it would have been beneficial to do tests on a UAV without wires for comparison. Without this, all that is known is that the Q can higher than measured on the Tornado and the system designer should know the limits within which ZigBee can be used.

## 9.3 Higher Q environments

A wide variety of measurements have been reported, in various cavity types, including a large reverberation chamber, a small reverberation chamber and a structure custom made to resemble half of an aircraft wing. The very highest Q measurements are beyond those one would ever find in practice on an aircraft. For example York's empty reverberation chamber has a Q

of over 30000 which would never be encountered in a military aircraft or UAV because even empty bays would be too small to reach a Q like this. The empty wing was only 2000 at its maximum although it was a lot leakier than the large and small reverberation chambers which was part of the reason. It is possible to imagine a few structures though that could hit the highest Q, empty sealed transports containers being an example. In this instance it is reasonable to say that ZigBee and all the IEEE standards would not work at all and if one wanted to use a wireless system then absorber would have to be added deliberately.

As the Q-factor was dropped from 10000 down to 1000 there was a dramatic change in the performance of ZigBee which went from not functioning at many stirrer positions to functioning at any. Again it is worth remembering that this is not an abstract concept, it is exactly what anyone would see happening as antennas were moved to different locations or as the environment changed due to vibrations or flexing. Under 1000 and particularly by a Q-factor of 500 it was found that ZigBee worked reliably for almost all stirrer positions in all structure types but that occasionally it could drop to 100% PER. This is undoubtedly the largest worry with using wireless systems in a reverberant or even semi reverberant environment; an installer might be sure that a system is reliable because on installation and test it works with 0% BER and 0% PER even with all noise sources turned on but at a later time it could switch instantly to 100% PER. How to deal with this will be dealt with shortly.

The nature of the 0 to 100% PER is itself interesting because it is quite different to the Gaussian noise situation. It could be difficult to deal with in practice because a link could not detect a gradual worsening and make changes, a system would instantaneously fail. Some of the results in chapters 4 and 5 have indicated that these errors are closely correlated with delay spread at the higher Qs. This was backed up by work done with a signal analyser reported in chapter 6 which indicated that it was not just a case of the header or part of a packet being corrupted but the synchronisation process failing because of the delay spread and thus causing a whole frame and packet to be lost. For the lower Qs the exact packet loss process was harder to understand but appeared to be a combination of delay spread and attenuation leading to errors. As was expected, the relationship between the symbol time and multi-path delay was clear to see. Once the average delay was under a tenth of the symbol time the percentage then the average BER dropped rapidly although at individual positions the Q and delay could vary by almost an order of magnitude. Similarly once the Q was under 5000 then the channel moved from frequency selective to flat. It was convenient that it was possible to show that the

Q-factor exponential decay corresponds exactly to the mean delay calculation value found in the communications literature. It means that if one knows the Q roughly one can estimate performance from the communications literature, in which the performance of many systems in different delays may be found. Although the delay mechanism is different for resonant cavities the effects of inter-symbol interference and frequency selective fading remain.

Additional complications can be added by bay or cavity coupling, noise and time variance. The first of these has been shown to make minimal different in terms of delay spread. It theoretically increases it but in practice it was found not to have much impact, especially when there were multiple coupling points. The main danger through cavity coupling appeared to be signal strength loss, especially over short distances in the wing if the aperture was of the wrong polarization. Chapter 5 also showed how there might be a double Rayleigh distribution in theory where there is only one coupling point and this is something system designers must be aware of when making outage probability calculations. Double Rayleighs were not seen however on the Tornado which was low Q anyway and had many coupling points between any two bays. This is a good result from the perspective of MIMO systems where MIMO is used to to boost capacity. Coupling through a single cavity causes singularities in the channel matrix and limits total capacity since all paths must go through one path at the coupling point but this would not occur on the Tornado.

Time variance was shown to be harmful at very high Qs but was not an issue in the reverberation chamber for Q-factors under 5000 and for the aluminium wing like structure and so it should not be a concern except where there are fast turbines spinning or alternative wireless systems are used with a symbol times order os magnitude less than ZigBee. Noise was seen to be more of a problem because it can fade independent of the signal in a reverberant environment. A new distribution was derived for this which showed long tails representing unlikely but possible occurrences of very low and very high signal to noise ratios. Where one might expect a constant noise and a signal variation of 20dBs they must now expect SNR variations of 40dBs with only small movements of the environment or antennas. This makes even more difficult the task of guaranteeing that a system will work reliably.

In chapter 8 the various 802.11 wireless LAN standards were investigated because they were seen as a possible companion to ZigBee for the high data rate communications links in an aircraft and because they illustrate a lot of additional technologies. The aim was to show that wireless LAN could offer much better performance than ZigBee and be more reliable. Un-

fortunately the performance was not as good as anticipated. Firstly, the 802.11b standard was not able to get past 2Mbps without becoming unreliable in the sense that over 5% of stirrer measurements positions were high BER even for the loaded wing scenario. This is still 8 times the ZigBee data rate but much better was expected. For OFDM and 802.11g higher data rates were achieved with 6 or 9Mbps looking usable before an 18Mbps rate caused 1 in 20 positions to have a 50%+ BER. Ideally one would wish to find 0 or 1 such positions if there is to be a chance to use redundancy to reduce the chances of link failure to improbable. The n standard appeared to suffer the same problems because it's modulation is the same and in fact some of its data rates use a shorter guard interval which is not desirable in this environment. Diversity was found to improve performance but again not as much as hoped particularly at lower data rates. Why this is was not clear but it is likely to be related to the diversity algorithm taking account of signal strength over error rate, or perhaps due to large unintended correlation in the receiver. It should be noted that the symbol time per carrier in 802.11a,g and n is actually the same as for ZigBee but without the benefit of spreading but with the advantage of many narrower channels. From the channel measurements one could expect many of these channels to be correlated and thus to all fail together. If this happened the error correction would be of no use, as we'd likely have many channels unsynchronised and with all bits in error. Wireless LAN then looks useable for anything up to 10Mbps but not to be suitable as a 100Mbps cable replacement when used in an off the shelf configuration in an environment with a Q approaching 1000 or above. However channel measurements suggest it could perform a lot better if diversity was correctly applied to counter fading caused by delay.

The statistics of the transmitted signals in cavities were found to be Rayleigh as expected for a reverberation chamber. These were found to hold for the wing and the Tornado although predictably the correlation with adjacent frequency and small distances was a lot higher. Therefore we can say that when designing a wireless aircraft one should assume this for their model of the signal in the absence of noise, obviously being wary of the points made above. In the measurements shown in previous chapters the power levels varied greatly from system to system with S21 going from -20dB to -85dB and so although the likely distribution is known, each situation will need its parameter determining via experiment and/or simulation so that the probability of an outage can be calculated.

## 9.4 Improving reliability

Given that in resonant parts of an aircraft, total packet loss can be expected at certain antenna positions or at certain times in a time varying environment the next question is one of how to prevent that or make up for it without making large changes the complete radio hardware and protocol stack. There are a number of solutions to this problem all of which involve adding redundancy:

**Multiple antennas** It has been shown that certain stirrer positions create a poor channel between two antennas and other's don't. If both the receiver and the transmitter used a second antenna the probability of an outage $P(O)$ would no longer be $P(O)$ but $P^4(O)$ which if the original probability was 1 in 100, as seen in many measurements in the wing, would drop the probability to the very improbable. This does however assume no correlation between antennas and might not be practical because even if tiny chip antennas are used, the requirement to keep them a half wavelength apart at least will take up a lot of space per sensor. A better solution can be seen in figure 9.1 which takes account of the fact that sensors are not likely to be communicating with other sensors but are more likely to be alerting a control computer or processing electronics or their readings. The likelihood is that this central device will be aggregating information from many sources and be a lot larger than the individual sensor units with radios. Therefore it is logical to have many antennas on the larger units and one on the sensors because the benefits are the same, the number of independent paths between two points has been increased. How this might be implemented is uncertain. The protocol stack could

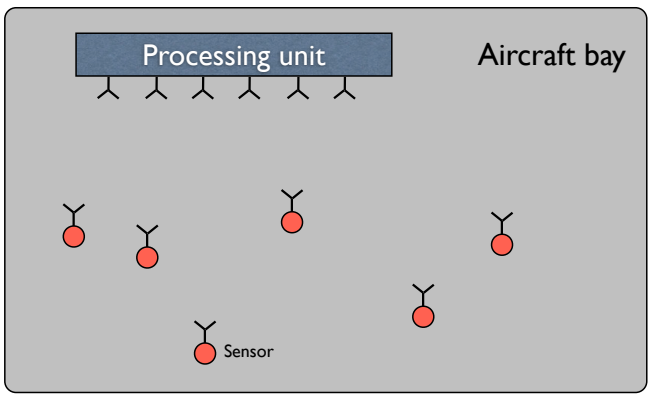


Figure 9.1: A selection of sensors all communicating with a processing array.

be slightly altered as could existing module designs to accommodate the extra antennas and choose between them according to delay in the received signal or some other metric. Alternatively some form of smart antenna and beamforming technology could be employed. This is already allowed for in 802.11n as described in [61] which explains that 802.11n does not specify a method for beamforming but typically a weighting is given to each antenna based on a singular value decomposition of the channel. Such a technique could be used in a reverberant environment but might need adapting to cope with high delays and not jus signal strengths in fades.

For ZigBee these techniques make less sense because there are not one or two devices continually communicating but instead there may be many antennas being read from. Perhaps a higher level approach would be better whereby multiple modules are used at the processing end with radios being enabled and disabled according the which sensor was being read from or sent a message. Radios could be sent messages every second or two via every antenna until one was found that gave an error free response.

The other way to deal with this would be moving antennas, for example dipoles that can rotate to minimise error rates. Whatever is used the system designer must add some type of receive diversity so that links are not lost and must design this for error rate minimisation as opposed to channel power.

**Multiple routes** A complex mesh network already has build in redundancy. In 802.15.4, if an access point is used, called a PAN coordinator then it works in a similar way to wireless LAN whereby all end points communicat with the coordinator. If this was to stop working then all links would break and there would be system failure. However in alternative configurations the higher ZigBee topology allows routing of data packets meaning that if one coordinator went down another node could be used to route packets. So long as there were enough coordinators around an aircraft with spare data bandwidth then a link failure need not be serious. This is very similar to the triple redundancy found on avionics systems today but even more so because there could be many many different paths between end points. There are a few difficulties here: coordinators must always be on and so must be powered, they need to be on the same frequency as the failed nodes and the traffic modelling would be complex because it would not be straightforward to prove that a reliable route would be available as other links failed. In principle though,

given that our channel is not reliable, it is sensible to take advantage of the inherent routing capabilities of the ZigBee (and internet protocol).

**Multiple frequencies** Just as there is spatial variation in signal strength and delay there is also frequency variation and so this too could be better exploited. Neither ZigBee nor WLAN explicitly support dynamic frequency change in a network but if combined with spatial diversity it could prove a powerful way to improve reliability. The problem is that at 2.4GHz there is limited bandwidth with only 85Hz and so channels are selected and they kept (although for the 5GHz band this is much higher). If ZigBee and 802.11 were adjusted to work over much greater bands with high level protocol changes to adjust frequencies could have a big effect.

## 9.5 Suggestions for the design process for wireless aircraft

Given the prior discussion, the final subject to address is how everything that has been found should be applied in the design of future wireless aircraft. Our requirement of a robust and reliable wireless aircraft must be included from the beginning of any design process and the propagation environment must be understood. As an example of how this might be incorporated into aircraft design some very general steps are given that address this problem:. From an avionics perspective one could do the following:

- Determine the sensors, actuators, radar and imaging electronics and the processing units that will be required and where they will be placed
- Perform stirred channel measurements between all pairs of linked bays to determine Q-factors and delay spreads
- Select a combination of ZigBee and WLAN systems to meet the data rate, physical space and reliability requirements, e.g. to make the probability of a link failure between an actuator and flight control computer less than $10^9$, given steps 1 and 2.
- Also select a sufficiently high number of frequencies to deliver reliable performance
- On a prototype, post manufacture and fitting, run stirrer BER tests on the aircraft. For example, place an stirrer in every aircraft bay and step each one individually each time exercising all links. Do so with all noise generating systems functioning.

- Where reliability tests have not been met add additional robustness in terms of frequency distributions, additional antennas, more PCFs to relay signals and other redundancy
- Run reliability tests again

Using stirrers in validation is an unusual approach but such an approach provides a route to certification, unlike the case of placing Bluetooth or otherwise on an aircraft and showing it can fly. The key in aircraft is to show that wireless systems will work in all circumstances and what this PhD Thesis has demonstrated is that one cannot be sure of this, particularly where the channel is resonant. One must add as much redundancy as possible into the links at the design stage and one should test the finished aircraft using electromagnetic stirring in order to prove true reliability over the long term.

It is most likely that these issues will be explored first on small UAVs rather than manned jets because UAVs will be an easier certification environment and allow technical experience and knowledge to grow in a less risky way both in terms of financial costs and safety requirements. Wireless may be used on manned jets for additional redundancy over the existing wired systems. Where wireless is used for this, the above suggestion of analysing systems in terms of stirrer BER measurements would also have value because this will allow the collection of data to aid the design of future aircraft.

## 9.6 Future research

To conclude, we look at the areas where the most research still needs to be carried out. At the physical layer there is still research to be done into the exact mechanisms of the link failure in resonant environments. Previous chapters have demonstrated the end results in terms of packet loss, and bit error rates but exactly how synchronisation is lost or automatic gain control fails are not fully understood. These types of issues are best solved using computer simulation, taking real channel measurements as found in this report and incorporating them into synchronisation algorithms to see what happens over a bit and over a packet. Some work was started on this during the Flaviir wireless project but unfortunately there was not time to complete it. From a theoretical perspective it would be useful to know how other modulation types compared to the ones shown here. There are many different types of receiver architecture, of carrier recovery, of FSK, PSK, spreading algorithms, and of automatic gain control. In

this project the focus was on finding a fast and cost effective route to a fly-by-wireless aircraft and so standard technologies were chosen but it would be interesting to see how others performed.

Lightning protection and deliberate jamming are two other areas that require significant work. The former is as much a pure EMC issue as it is a communications one and it could be imagined that the existing research in this field could be incorporated into a model of the wireless channel. Deliberate interference from anything up to EMP strikes must be given a lot of time for the obvious reason that it is rather easier in principe to interfere with a wireless network than a shielded cable. It may be that this has implications for airframe design in terms of aperture sizes on the aircraft or in the power levels that have to be used for the wireless radios.

The largest area of future work is likely to be at the network layer and in terms of architecting the wireless system so as to make up for the problems that we have identified at the physical layer. There are complex issues in terms of data traffic modelling, frequency organisation, antenna configurations and how sub systems are separated. Being able to prove that complex multi-node systems can guarantee reliability in this environment will not be simple. In the far future, as redundancy becomes cheap we may wish to move away from dedicated avionics racks to almost entirely distributed processing, self healing and all sorts of multi agent systems topics and this is also a huge topic in itself.

The previous chapters of this PhD thesis have explained the best technologies to use, demonstrated how well they work and provided significant design data in terms of distributions, ranges of Q-factors, delay spreads and how these interact with radio receivers. The next step is to take this and transition from a point to point link to a system with multiple nodes on real demonstrator aircraft so as gain a fuller picture of how a fly-by-wireless system will fit together.

# Appendix A

# Antennas for wireless aircraft

In the introduction of Chapter 1 it was explained that the work described in this PhD Thesis was carried out in collaboration with 3 other universities and one of the most biggest areas of collaboration was that of antennas for which work was done jointly with Liverpool University. All wireless systems require transmitting and receiving antennas and Liverpool's Direct Write (DW) expertise and facilities provide a new way to fabricate these for aircraft. The design and application of some novel DW antennas will be dealt with in this chapter. For detailed information about their fabrication see the work of Liverpool's laser group including articles such as [62]

## A.1 Design of antennas for direct write

### A.1.1 Direct write

Direct write is a process for creating structures by depositing functional materials onto a surface [63]. In the Flaviir project this was done by the deposition of silver inks via a syringe with the inks then being cured using an oven or a laser. This technique can be used to create complex, multi-layer 3d structures over many passes as shown in [64] multiple passes. As well as being used for antennas DW could be used for applications such as for producing small flow control actuators or sensors or for writing power buses over a surface. It may prove useful in writing and linking a number of systems on an aircraft and so it is logical to try and use it for wireless antennas as well.

### A.1.2 Direct Write antennas

An antenna typically used for wireless applications such as mobile phones is the planar antenna which provides advantages in terms of size, shape and profile in comparison to a dipole or monopole antenna. It can be as simple as a conducting patch of metal placed on a dielectric with a ground plane beneath and has many variants such as the PIFA, low profile printed monopole and various aperture coupled antenna types although in this chapter we only cover the simplest variety. Reference [65] covers the more exotic varieties. Patch antennas are also cheap to design and manufacture since even at 2GHz and above they can be made on FR4 or with other cheap materials. The only problem with these antennas comes with the issue of how to fit them onto an aircraft both for new aircraft and retrofitting. If these antennas are to be used in an aircraft it would be imagined that they'd be attached to the same PCB or enclosure as the wireless electronics or if separate they would be attached to the aircraft. Sometimes this might not be straightforward, especially if it compromised the structure and if we wished to place antennas on the outside surface of the aircraft the profile of even the narrowest patch antennas might be too high from an aerodynamics point of view. With DW all that is needed is access to the surface onto which we can draw a device; there is no need for screws, drilling holes, clamps etc.

The idea in principle is that given the surface of an aircraft wing or any other internal or external part of an aircraft it should be possible to lay down a conducting surface, then a dielectric and then a second conducting surface forming the antenna and that this might be done both inside and outside the aircraft. Being able to do this is dependent on being able to deposit the ink in a finely controlled way, so that it leaves a smooth surface and accurate geometry. This requires suitable inks with the correct conductivities and is also very much dependent on the curing process. There are a number of variables involved here when doing laser curing such as the wavelength to use, the laser power and the number of passes. These will all vary according to the substrate that is being used for the deposition and so DW poses a lot of manufacturing challenges which were tackled by Liverpool University and which should be published in the PhD Thesis of Taku Sato. In the following sections we largely ignore these and concentrate on the electromagnetic aspects and antenna design.

## A.2 Design of a simple test patch

At the beginning of the Flaviir project and still there was very little literature on the subject of DW antennas and so the first task was one of proof of principle. In order to show that the DW technique could write a patch a starting point was chosen of a rectangular patch antenna on FR4 dielectric like that shown in figure A.1. A centre frequency of 2.442GHz was selected because the aim was to design antennas that could work with the ZigBee and 802.11 systems and also because this would mean we could compare it against the 2.442GHz antennas we were already using. This is as simple a design as one could find and as will be explained later, it was not going to provide a high bandwidth compared to antennas like the Jennic provided monopoles or fractal based chip antennas. We were only testing the actual patch deposition because FR4 dielectric was used and this was for three reasons: Liverpool had no experience depositing dielectrics, to do so would be very time consuming and most importantly we wanted to have a simple comparison between chemical etching of copper and direct writing of silver inks.

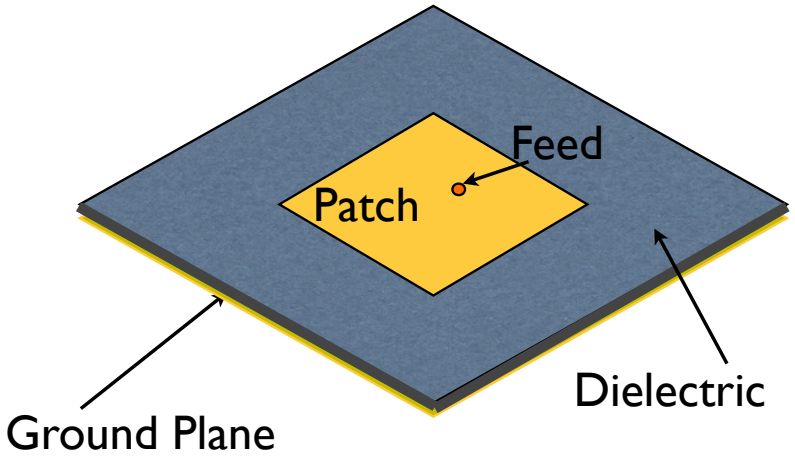


Figure A.1: A simple patch antenna

### A.2.1 Gridded antennas

The process for making direct write antennas consisted of a computer controlled syringe moving up and down in adjacent lines depositing the ink and the same when laser curing. This was initially a slow process (curing was at approximately 1mm/s [66]) and so it was desirable to cut down on the number of lines. Even with a faster process it is best to make the fabrication as

quick a process as possible. Therefore it was suggested that it might be better to fabricate the patch as a grid structure rather than a completely filled structure to better match the manufacturing technique meaning that the syringe and laser would need to do many less lines across the antenna. What sort of effect this might have on the antenna electromagnetically is not obvious. The expectation was that if the outer dimensions remained the same one would expect the same resonant frequency as the full patch but it was likely that as the spacing in the grid got bigger this would eventually change. Simulations were carried out at York by John Dawson to test gridded patch designs against full patches and see if our expectations were correct.

In order to test DW antennas against a reference equivalent antennas were produced using a PCB etching process using copper on a 1.6mm FR4 PCB laminate with a relative permittivity of 4.05 at 2400 MHz. Simulation and design was carried out assuming these parameters. The antennas were designed for a centre frequency of 2.442 MHz selected for being in the band of many wireless products and standards. To design the antennas standard equations were used from the antennas literature, e.g. Balanis's Antennas [67] or Carver's paper on patch antennas [68]. Effectively the method involves working out the effective length including fringing fields, calculating the velocity inside the patch and calculating the width and length where the effective length specifies resonance. Following this the feed point may be selected. It can approximated by equations based on a transmission line theory where the feed offset is selected to minimise the susceptance at the input. Simulation may also be employed to select the feed point precisely by trying different positions and observing the return loss.

For the 2.442GHz antenna a patch was designed with a length of 30.11 mm, width of 38.63 mm with the feed point 9.54mm along the length. John Dawson simulated this using an in-house FDTD code and the model was simulated on 0.4 and 0.2 mm meshed grids which gives some error in the patch dimensions. For the 0.4 mm grid the length was 30.0mm (-0.37%) and the width was 38.8 mm (+0.45%) and for the 0.2 mm grid the length was 30.2mm (+0.29%) and the width was 38.6 mm (-0.6%). The patch is shown in figure A.2 Simulation results for the gridded and solid patches with both FDTD mesh sizes is provided in figures A.3 and A.4 which show the reflection coefficient and the radiated power respectively.

Two points are apparent from these:

- The plots confirm it's possible to produce a just usable ZigBee antenna using 1.6mm FR4
- Gridding the antenna has little effect on performance

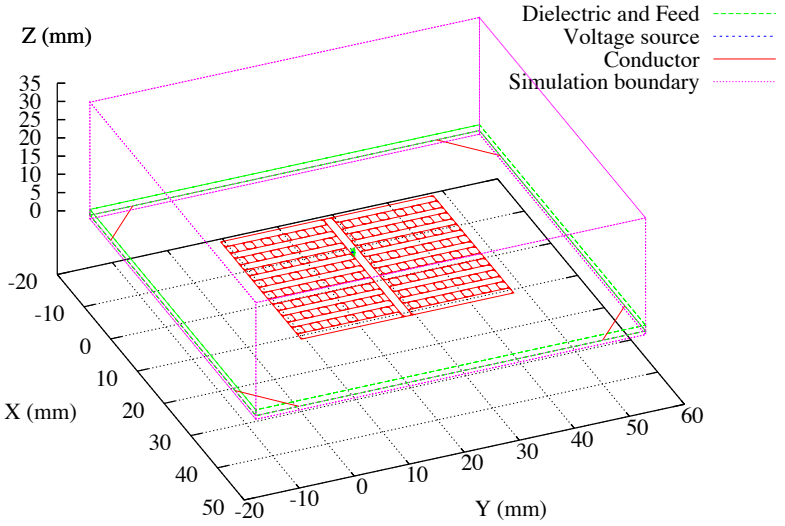


Figure A.2: The Zigbee patch antenna from a grid of lines (FDTD model on 0.2mm grid), plan view showing model dimensions (filled rectangles have diagonal lines on corners)

To produce the antennas may not be trivial though as we see the antennas only have a 10dB bandwidth of 50MHz for the S11 parameter which means their useful frequency of operation is smaller than the ISM band but larger than a single ZigBee or 802.11 channel. There is also a shift in centre frequency for the gridded patch compared to the solid one and none of the antennas has the correct centre frequency. This is a general problem with patch antennas, in that they are narrow band and any small change in dimensions or dielectric or other parameter may move the antenna operation out of the required band.

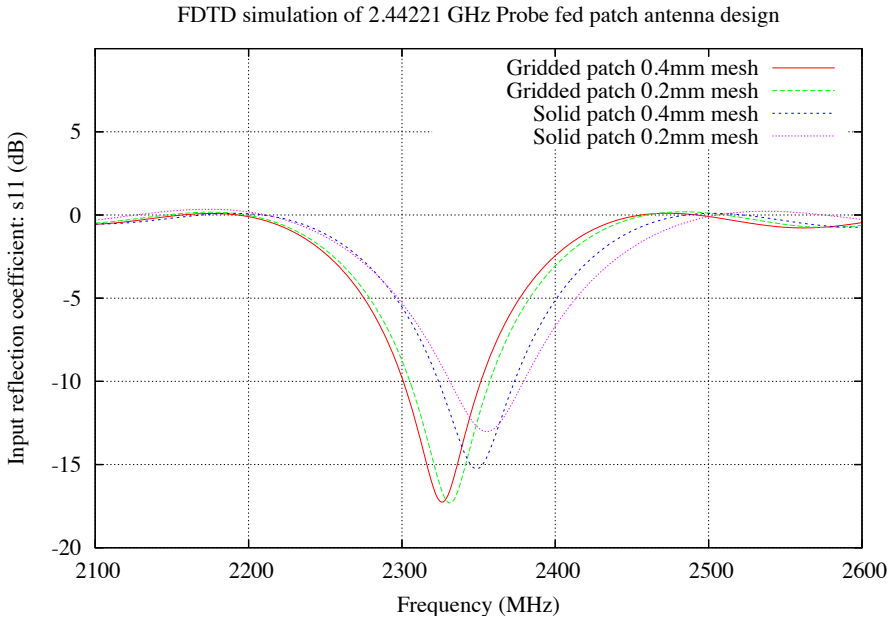


Figure A.3: The simulated input reflection coefficient of the Zigbee patch antennas

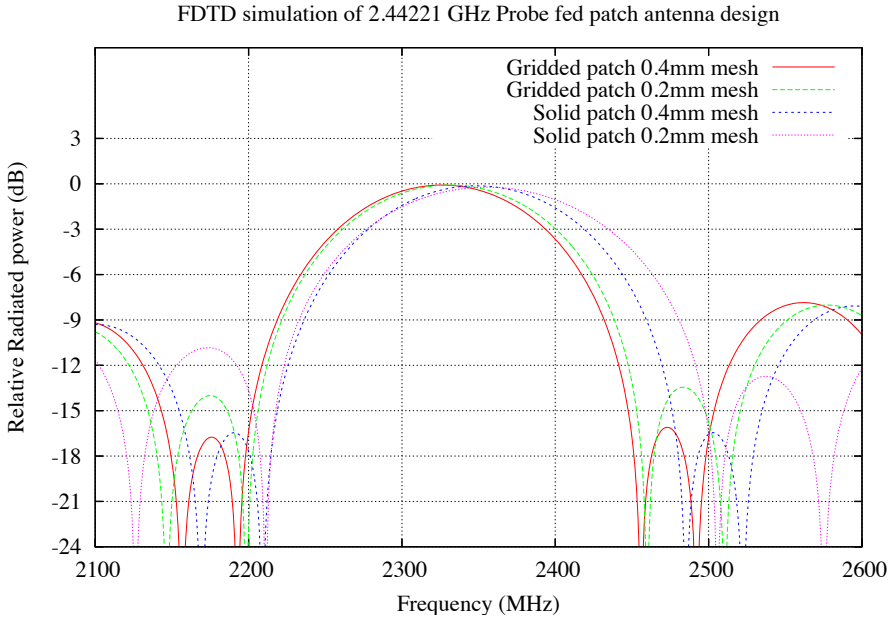


Figure A.4: The simulated radiated power of the Zigbee patch antennas

### A.2.2 Fabrication

Having found that the design should work (albeit not exactly at the right frequency) antennas were fabricated for testing. Direct write and etched antennas were used, both on FR4. As explained before, although in the longer term it was expected that a dielectric might be direct written, for now FR4 was used so the performance of the silver inks could be checked and the basic DW technique verified. The two fabrication methods were:

- Despostion of ink followed by oven curing
- Copper etching on a FR4 PCB

Laser curing was not used in this case but the technique of depositing ink followed by oven curing would allow us to test the inks and is a faster process.

The objectives were to compare the performance of the gridded to non gridded patch and to provide a reference to which the ink deposited antennas could be compared with the hope there would be no difference between the full PCB, gridded PCB and DW antenna. Full PCB patches had been made by colleagues at York previously at other frequencies and the technique is very well controlled so there was confidence in this method providing a good reference. Figure A.5 shows the PCB patch and gridded patch antennas side by side.

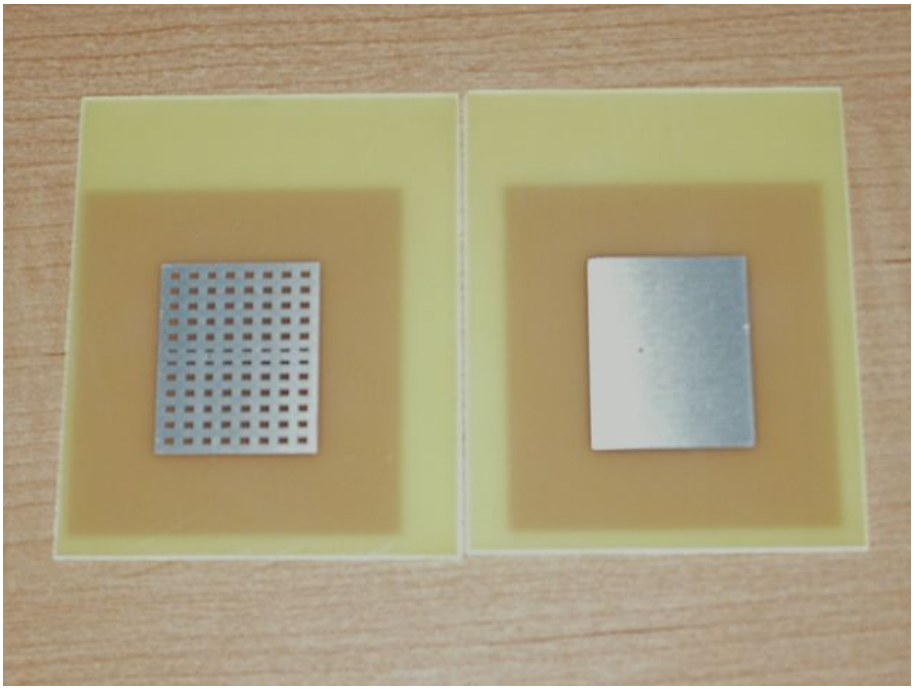

Figure A.5: A grid and a full patch 2.4GHz antenna side by side

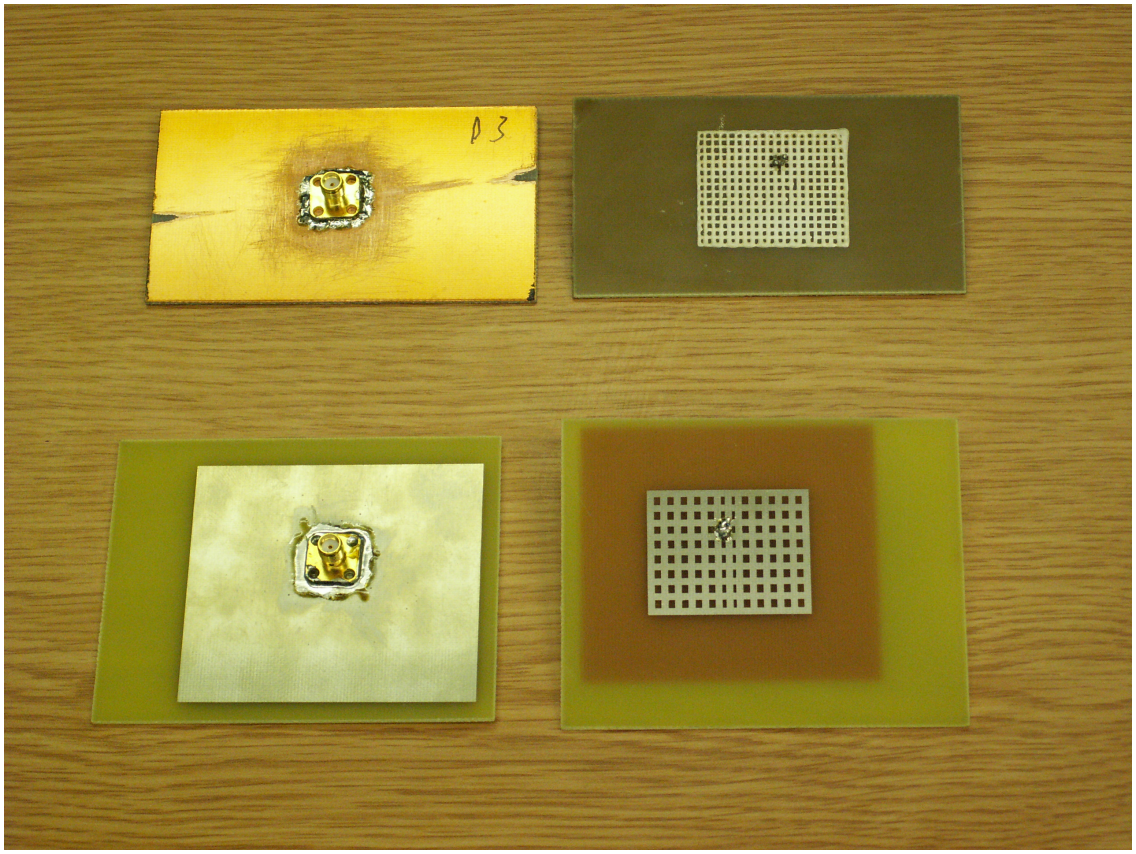

Figure A.6: York and Liverpool patch antennas where the top two antennas are from Liverpool ones and the bottom two York

**Testing of basic antennas**

The PCB gridded patch and DW gridded patch can be seen in figure A.6 and it shows that the grid sizes were different here although according to the simulation done prior this should not have made any difference. This occurred because in manufacturing it proved easier to make more thin lines than less thicker lines with DW. One can also see from the images that the DW process at this stage was not very precise and the grid geometry varied by 1-2mm over the PCB process which would be expected to cause some frequency shift given the resonant frequency is dependent on antenna dimensions. The antennas were tested by measuring the S11 parameter using a network analyser and no power measurements were carried out at this point, just tests to see their frequency response.

The measured performance is plotted in figure A.7. It is presented in terms of the antenna's S11 parameter which provides the reflection coefficient. The lower the value the better as more energy will be transmitted.

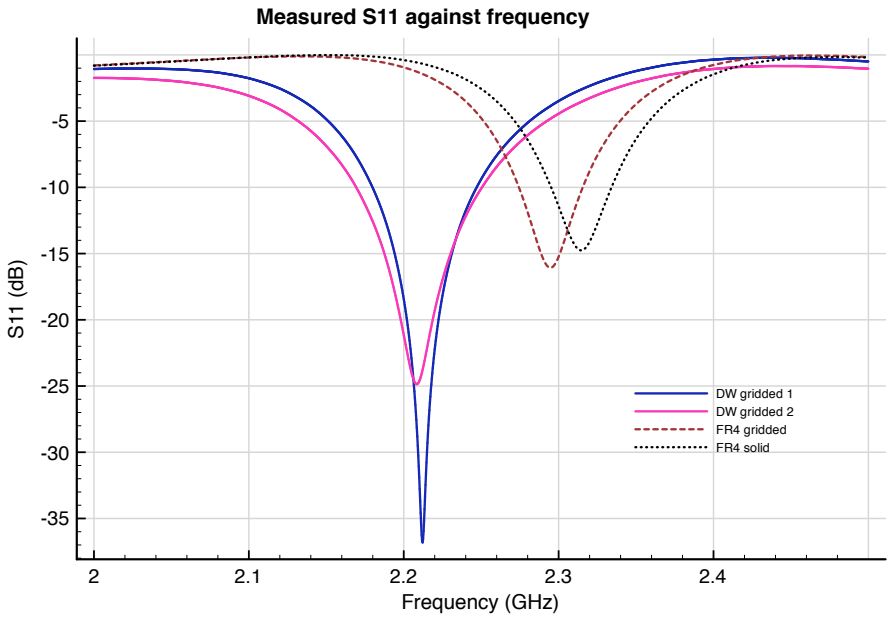


Figure A.7: S11 for FR4 and DW solid and gridded patch antennas

The first thing we can see here is that the FR4 gridded and solid patches are very closely matched in terms of bandwidth performance although there is a slight frequency shift between them. We also see that there is little difference in -10dB bandwidth performance between the FR4 and the DW antennas although the latter is wider. The centre frequencies are shifted relative for the DW relative to the FR4 ones but this was to be expected given that the dimensions were not exactly the same being a few mm out. We should also note that FR4 does not guarantee a permittivity and this could vary also. Furthermore the Liverpool antennas drop to a lower value of S11 when at their resonant frequency by 15dBs. and so surprisingly the DW antennas actually have the better response. Neither antenna type is suitable for ZigBee since they have bandwidths that are too small. In all cases the -10dB bandwidth is between 50MHz and 70MHz whereas the band in which the channels can work for ZigBee is over 85MHz. This was expected however and so it is not a fault with the fabrication of any of the antennas. Why the DW antennas performed slightly better is not known but one possibility is that because the profile of the ink is quite high the conductivity was increased relative to the very thin 0.06mm copper track. This is only speculation though and the import point to take from this measurement is that the DW gridded antenna is able to produce an S11 response as good as the conventional fabrication method, albeit requiring some tuning to obtain the correct frequency.

At this point we have shown that a planar antenna can be fabricated with direct write in

a grid pattern having been designed using the standard equations for planar antennas. There were many issues associated with the direct write process but these are not addressed here. Although this could be regarded as a success these FR4 patch antennas were not ideal for an aircraft ZigBee application because they were large compared to commercially available chip antennas and they were low bandwidth. In order to address these problems two significant changes were made in:

1. Dielectric
2. Design methodology

### A.2.3 Genetically evolved alumina antennas

To make the antennas smaller a new material was selected: alumina. It has a dielectric constant of close to ten which is more than double that of typical FR4. This reduces the effective wavelength in the material (by a factor approximate to the square root of the change) and means that we can build smaller antennas. The other proposed change was to try designing the antenna with a genetic algorithm (GA). By using genetic algorithms one is not constrained to the performance limits inherence in the standard design equations and it might be possible to increase bandwidth. A researcher at York University, R Woodhouse developed various genetic algorithms for the design of patch antennas, looking in particular at patches for mobile phones including dual band patches. This had some success in the design of GSM frequency antennas and so it was decided to take this software and try and apply it to 2.4GHz antennas which was not unreasonable because the software made no assumptions about frequency. The following section offers a brief introduction to genetic algorithms and how they apply to antenna design. It is kept necessarily brief as it was primarily a tool in the work, not the focus.

#### Application of genetic algorithms to antenna design

In brief, genetic algorithms work by taking a set of random designs or solutions to a problem, testing them and then selecting a subset or mixture of those designs to go forwards with. If, in a very simple case we imagine we have a string of numbers like 0,1,0,0,0,1,0,0 then if we wanted to evolve a string with the highest binary value then we might proceed as follows:

1. Generate 6 random strings of 0s and 1s
2. Take each one and calculate its numeric value

3. Select 2 randomly but with bias so that those with higher value are more likely to be selected

4. Mix those two somehow. Perhaps take half the bits from one and half from the other

5. Repeat 3 and 4 six times till we have 6 new strings

6. Mutate some of the strings according to a mutation rate, anything from all bits in all strings to a small chance of one bit in one string

7. Now do 2 and check if we have the string 1,1,1,1,1,1,1 and if we do stop, if not continue

This is a genetic algorithm at its simplest but there are many variations. For example, one might make not use cross over at all and only use mutation. Rather than allowing all chromosomes in the population to be crossed over or mutated we might ensure we always keep a copy of the fittest individual in a population. For example the $\lambda$+1 strategy works by finding the fittest (in this case the one with the highest value) individual and replicating it 6 times and then mutating 5 of the 6. It can take many generations (tens to tens of thousands dependent on the problem) to produce a useful individual in some problems and the parameters used can have a big effect, such as the population size and the mutation rate.

Before going any further it is important to note that although a genetic algorithm typically uses a string of 1s and 0s we often talk of a genotype phenotype mapping which describes how the 0s and 1s map to whatever we are trying to design. Also, when 1s and 0s are not used then it is often given a different name such as genetic programming although the principle is always the same.

For the antenna designs two options were available:

- A simple genetic algorithms with mutation and crossover
- Cartesian genetic programming

Whichever method was used the principle was as follows: a maximum width and height on a PCB is selected. The rectangle this forms is divided up into a square grid. Each square on the grid could either be copper plated or not. We might also allow variation in the dimensions of the grid within a constraint of a maximum size. It is easy therefore to see how a genetic algorithm might be applied to this. If we take a string of bits and assign each bit to a square on our grid, each bit's state of 1 or 0 can correspond to copper or no copper.

For the case of Cartesian Genetic Programming (CGP) the situation is slightly more complicated. CGP is a system where the genotype is a grid of functions and an input and an output. The functions can be anything but will typically have one or more inputs and only one output, typically a bit or an integer or a floating point number. In the way this was applied to patch antennas the functions described the layout of the copper and the input and output were not important except for how they connected. Functions might be turn left, move 2 squares or lay down copper. It may be thought of as resembling a 'logo'[1] program. Because it is a genetic algorithm, there are a number of programs and the most successful ones are kept and get mixed with other programs. CGP is different to standard genetic algorithms in that there is a dependency between elements through their links, and because in a crossover one might obtain a big chunk of useful functionality from one unit and a big chunk from another (in this case big sections of antenna placement functionality, perhaps a section that acts well at one frequency in one and a section that acts well at another frequency in another).

It was suggested by the author of the software that the standard genetic algorithm had generally performed better for this application and so that was used. The procedure was as follows:

- Generate 60 random grids of a pre specified size, stored in a bit string form
- Convert all 60 into FDTD form and simulate
- Assess all 60 for how wide their -10dB bandwidth is and how close their centre frequency is to the one specified, producing a fitness number
- Perform cross over 60 times picking pairs randomly from the 60 but with a higher probability of selecting fitter bit strings
- Mutate all the bit strings with a probability of each bit changing of 5%
- Repeat 2-5 forty times and then select the fittest from the final generation

**Fabrication and performance - FR4**

Before trying any novel antennas on new material the technique was tested first on FR4. The shape of the evolved antennas is shown in figure A.8 and the S11 parameter is plotted in figure A.9 for both measured S11 in an anechoic chamber and FDTD simulation. This shows that the

[1] http://en.wikipedia.org/wiki/Logo_(programming_language)

bandwidth is double what it was before for both the measured and the simulated responses which is a major success. It would appear that this has been achieved by having two resonant frequencies that overlap, as indicated by the two minimas. Unfortunately the bandwidth is shifted slightly from the simulated response and would need tuning. This is most likely due to an incorrect assumption of the FR4 dielectric constant in the FDTD simulations because the geometry in the etching process is very accurate.

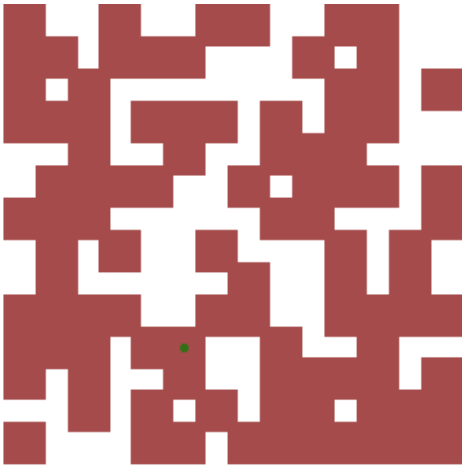

Figure A.8: Shape of a genetically evolved antenna for 2.4GHz.

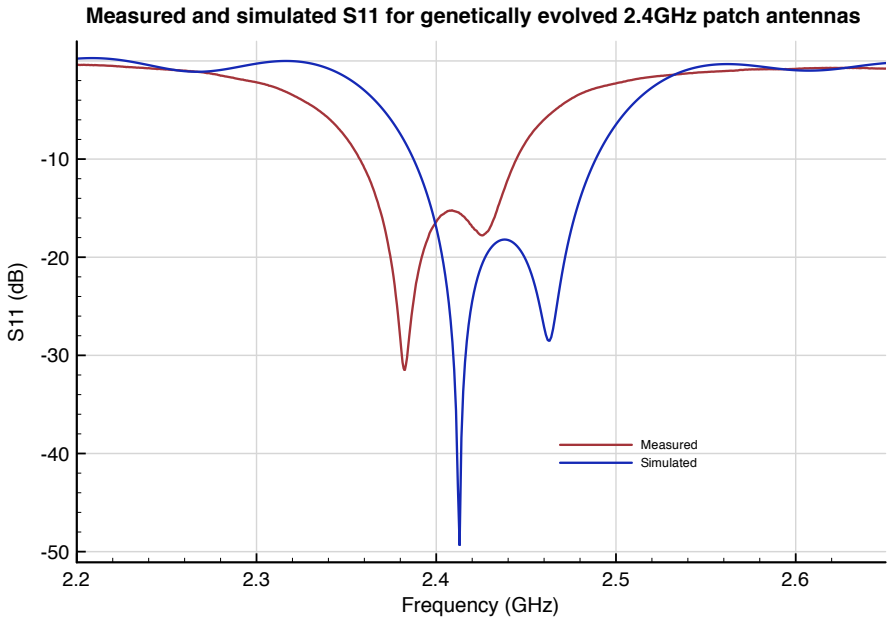


Figure A.9: S11 parameter for a genetically evolved patch antenna, 45mmx45mm

**Fabrication and performance - Alumina**

After proving the GA design technique for FR4 at 2.4GHz the next stage was to design an alumina antenna and find out if that also could provide the required bandwidth but with a smaller footprint. A number of designs were evolved in exactly the same way as for FR4 and two of the designs are shown in figures A.11 and A.12. Their dimensions were 20mm x 20mm which compares favourably with the original FR4 patches of 28.6mm x 30.2mm and the evolved version at 45mm x 45mm. They are also lower profile at heights of 5mm rather than 1.6mm. Their S11 performance is shown in figure A.13 and it shows that the bandwidth is not as high as for FR4 and but both are wide enough for the the 2.4GHz ISM band for -10dBs which is what it was designed for [2]. One of the two gave 40MHz or margin which is valuable because it reduces the problems of a small offset in frequency. However in this case there were significant offsets and this was more likely to be due to manufacturing tolerance problems than the substrate which was better controlled than for FR4. Overall the performance was much better than expected given the small size of the antennas.

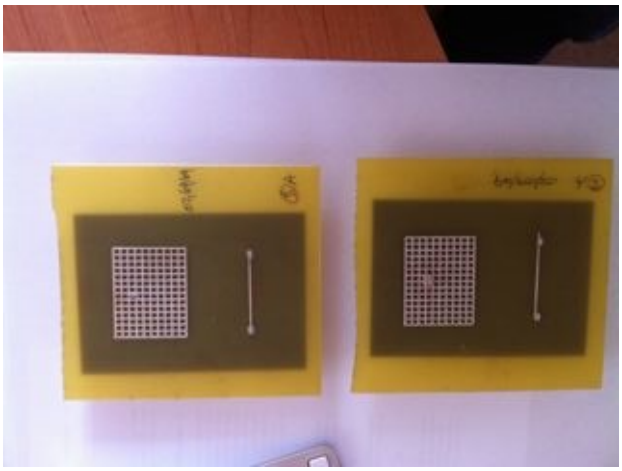

Figure A.10: A genetically evolved alumina antenna

Given that the S11 was sufficiently wide bandwidth for our requirements, no more time was spent trying to improve this. These GA alumina antennas were tested in power measurements in the anechoic chamber to find out how they compared to horns and to an FR4 patch made using etching which has been used as the reference. 2 other measurements were done at the

[2] There is no special reason for concentrating on 10dBs in particular but this number was suggested by colleagues as an appropriate level of performance to aim for in the band of interest, it means that at least 90% of the signal power is going to the antenna or the radiated wave.

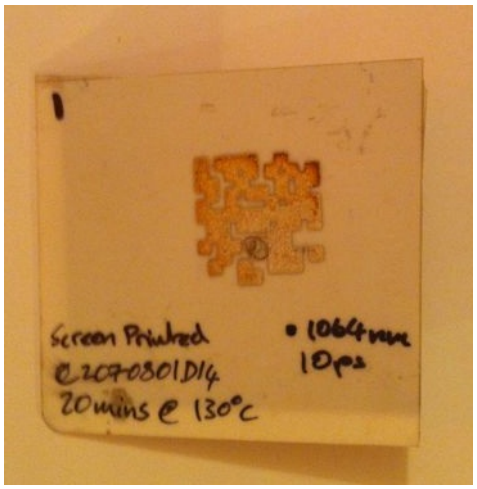


Figure A.11: A genetically evolved alumina antenna

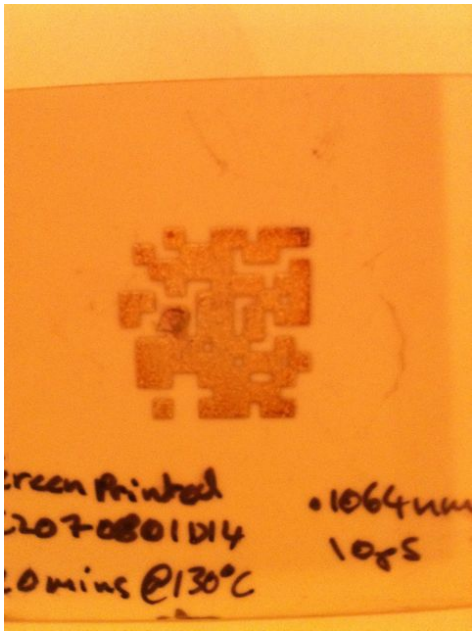

Figure A.12: A second genetically evolved alumina antenna

same time for direct write but done to a textbook design and gridded (similar to those shown at the beginning of this chapter but with more accurate geometry, see figure A.10) and these are presented for comparison. The gridded DW FR4 performance is plotted in figure A.14 and the GA alumina performance is plotted in figure A.15. Also on each of these plots is the solid FR4 patch and a horn antenna measurements which has its theoretical line plotted over it in order to show the level of inaccuracy in the measurement. For all the anechoic measurements we shall only plot the central measurements in the centre of the main beam for each antenna and for the GA antennas their orientation was that that transmitted maximum power to the horn as found experimentally because they don't have an obvious polarization.

For the case of the FR4 measurement the DW antennas were again frequency shifted from 2.4GHz, even more so than the solid patch. Their radiated power was almost as good being 3 and 5dBs down on the solid patch for the two antennas. That there was this difference was not surprising given that these DW antennas had much larger spaces in their grid than any of the other antennas and because they were being compared to a solid copper patch. Overall this was a positive result verifying that the DW technique can produce antennas just as effective as the chemical etching PCB process.

The alumina GA antennas performed less well as we can see in figure A.15 where both

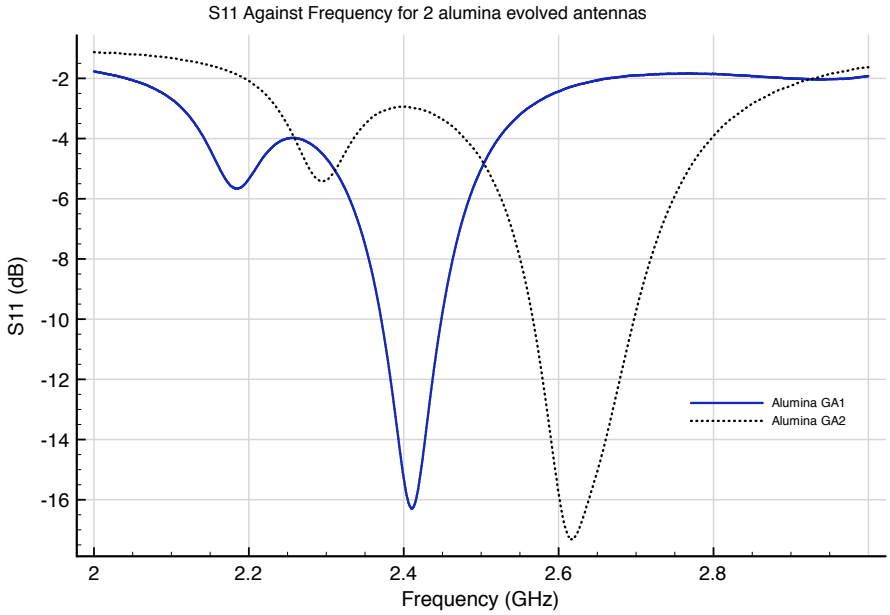


Figure A.13: S11 for two alumina antennas

antennas are 10dBs down on the maximum power of the FR4 copper antenna which is in turn about 5dBs below the horn. In this case one of the antennas got close to the centre of the ISM band whereas the other was some way off but this is just a function of the problems in manufacturing. Although only 10% of the power of the FR4 patch is being radiated this is not a bad result still because the distances inside an aircraft do not require a very strong signal. An antenna like this could be good enough, and the smaller size might be preferable to a better performing but larger antenna.

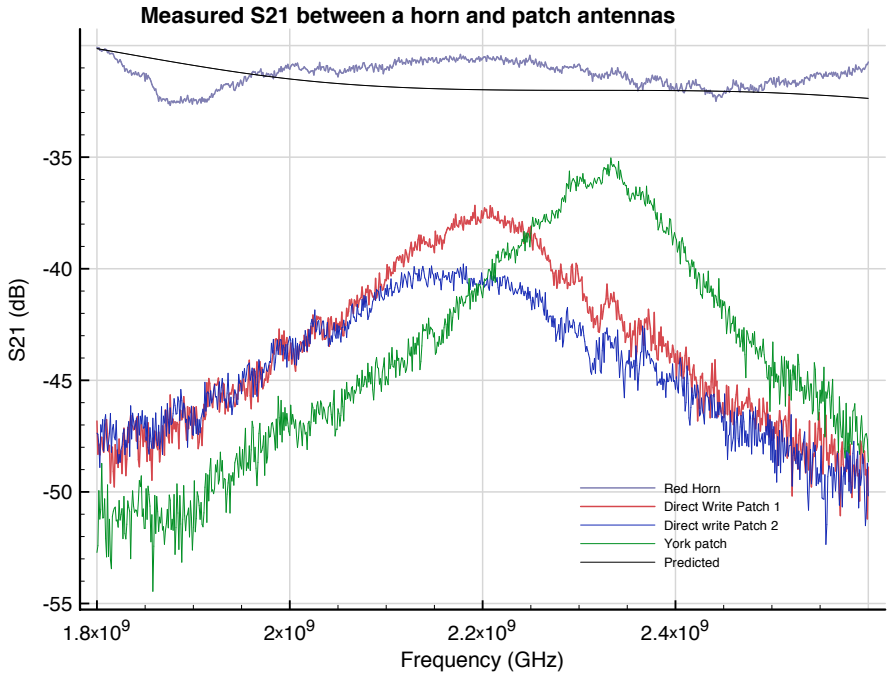


Figure A.14: S21 for two FR4 etched and DW antennas and a horn

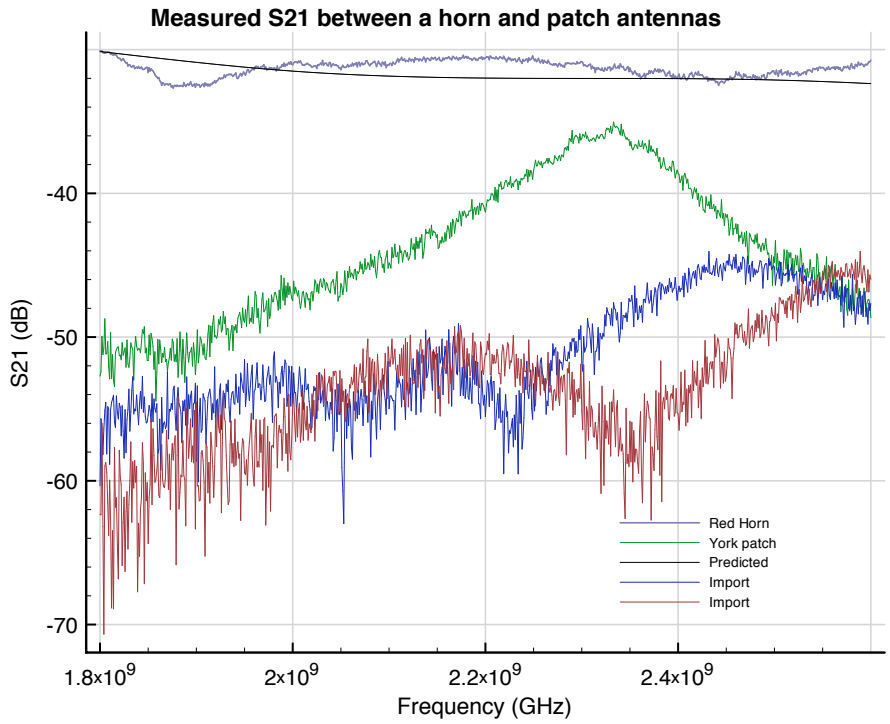


Figure A.15: S21 for two alumina genetically evolved antennas

## A.3 Fabrication on a foam wing

Once antennas had been successfully evolved for the ZigBee band the next step was to look at designing an antenna for a more interesting and practical example. So far direct write was being applied to FR4 and alumina substrates and there is no reason why chemical etching could not be employed to make such an antenna. There was a strong interest in looking at a scenario where the dielectric and ground plane was written using DW as well as the patch. This was not practical during the course of the Flaviir project as the DW technique had not reached a level whereby it could produce a consistent and thick ground plane as required. However an interesting possibility was identified whereby part of the aircraft structure could be used as the dielectric and so this was pursued.

### A.3.1 Aim

The aim was to apply the GA design method with direct write to make a work ZigBee antenna demonstrator

### A.3.2 Design

As part of a flaviir project wireless technologies demonstrator a ZigBee system was to be used in a foam wing transmitting with DW antennas. Typically such an expanded foam material would be used for model aircraft but they are also applicable to small light weight UAVs. The scenario envisaged was that of an antenna at each end of a wing communicating (e.g. a small computer at the fuselage end communicating with an engine or flap.) In our example we have the antennas communicating via the outside of the wing, but there is no reason why they could not be written onto the inside during construction of a wing it just would not make a very visually useful demonstrator.

The specification given for the antenna was very simple in that it should operate in the IEEE 802.15.4 ISM band, which is from 2.4 to 2.485GHz, and be capable of communicating over a distance of approximately 0.5m. This was made less straightforward in that the antennas were to be on the outside of the wing and the main beam of a standard patch antenna is normal to the surface and so the gain would only be -10dBi [69] for a standard patch. For an evolved patch it may be different, and will almost certainly be less uniform. When inside a conducting

cavity, this is unimportant however. The Friis transmission formula is:

$$\frac{P_{rec}}{P_t} = G_t G_r * \left(\frac{\lambda}{4\pi R}\right)^2 \tag{A.1}$$

From this, assuming we want an S21 of greater than -80dB for a ZigBee system to work (from the typically sensitivity fo a ZigBee system and allowing for an additional 10dB of noise) the path loss is 0.00199 and so the gains of two identical antennas must be greater than $2.78*10^{-3}$ transmitting in the direction across the wing. Of course for a real system the wings would be a lot longer and we would need to do significantly better than this. There was no choice in the dielectric used which was to be an unspecified polystyrene type blue foam of 10cm thickness.

**Measuring dielectric constant**

To design the antenna it was necessary to know the dielectric constant of the foam and so a sample was requested. The foam was also checked for losses, to ensure that it could potentially work as a dielectric, and this was done using an airgap microstrip as shown in figure A.16 whereby the air gap could be partially filled with the foam and the S21 parameters for the two cases compared. In measuring S21 with foam and without no difference was seen in terms

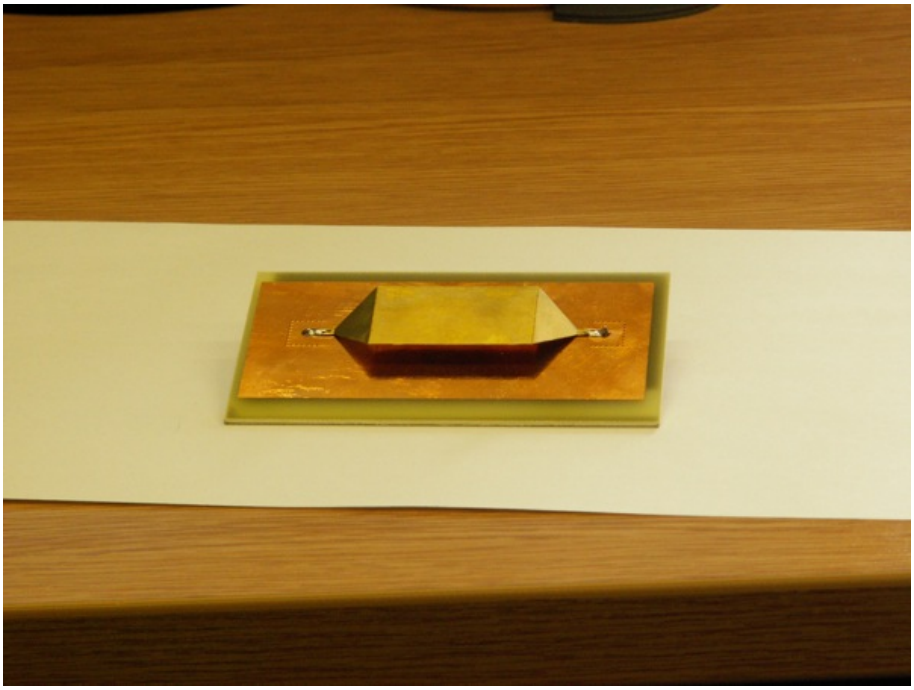

Figure A.16: An airgap microstrip for material conductivity and dielectric constant measurement

of losses. Following this it was necessary to measure the dielectric constant and this had to be

accurately measured. Were the dielectric constant not accurate the patch might be centred at the wrong frequency causing the ZigBee system not to function, a problem seen in most of the FR4 and alumina antennas. The actual wing surface was to have a thin resin coating on top of the foam and so this had to be measured with the correct height of foam. Therefore a sample of foam was requested with a ground plane underneath and a metal strip on top so as to also form a microstrip. Two of these were provided both with some problems that made accurate measurements difficult. One was of the wrong height and the other had some problems in the way the microstrip (made from copper tape) connected to SMA connectors at each end. Unfortunately there was no time to refine these so they had to be used as they were.

There are some approximate ways in which the dielectric constant can be calculated from S11,S22 and S21 and all 3 were used for checks on each other. S11 and S22 plots can be used to find the positions of the resonant peaks using which one can predict the dielectric constant using the fact that the resonances are dependent on the electrical length of the strip. For the S21 measurement we may look at the delay between the two ends which gives us the velocity in the medium and we may relate the velocity to the dielectric constant. Doing these suggested an $\epsilon$ of between 1.65 and 1.7. To be totally accurate was difficult because of the poor fabrication of the line and the large mismatch between the coaxial and the microstrip. With such uncertainty it is not possible to get a dielectric constant with perfect accuracy and there was a worry that when the the final antenna was fabricated this may lead to a large shift in centre frequency from the desired point. Whilst such an issue is not critical from an academic point of view it was a worry for the demonstration aspect of the work where the radio band was constrained.

To obtain a more accurate (and frequency dependent) estimate of the dielectric constant, an optimization method was also used to calculate the dielectric constant with the hope that this in combination with the other estimates would give more certainty as to its value.

For this method one can begin by writing down an expression for the impedance of a microstrip transmission line:

$$\frac{Z_0}{2\pi\sqrt{(1+\epsilon_r)}} ln\left(1+\frac{4h}{\omega_{eff}}\left(\frac{14+\frac{8}{\epsilon_r}}{11}\frac{4h}{\omega_{eff}}+\sqrt{\left(\frac{14+\frac{8}{\epsilon_r}}{11}\frac{4h}{\omega_{eff}}\right)^2+\pi^2\frac{1+\frac{1}{\epsilon_r}}{2}}\right)\right) \tag{A.2}$$

which is clearly dependent on the paramter $\epsilon_r$ which we wish to find. There may be some variation of $\epsilon_r$ with frequency and so we may represent this by:

$$\epsilon_0 \left( 1 + \frac{\epsilon_r - 1}{1 + \frac{j\omega}{\omega_r}} \right) \tag{A.3}$$

whereby we have assumed a relatively crude first order model (suggested by the PhD supervisor, John Dawson) for the electric susceptibility to take account of the fact that polarisation becomes more difficult at high frequencies. Such a model assumes a simple mass spring type system in terms of moving the electrons (on average) in the direction of the field shifted slightly from the centre of the atoms. It may be found in either [58] or [70] The propagation of the voltage wave may be described by

$$Ve^{-\gamma z} \tag{A.4}$$

where $\gamma$ is the direction of propagation. We can add to this the reflection coefficients at each end which are dependent on the line impedance. as given above and we can calculate what the loss and phase change should be between two ends of the line which is what S21 provides. $\gamma$ may be described by:

$$\gamma^2 = -\omega^2 \mu \epsilon' + \omega^2 \mu \epsilon'' \gamma = \sqrt{(j * w * \mu_0) * (j * w * \epsilon)} \tag{A.5}$$

With this information we may write use an optimization algorithm such as non linear least squares to fit the measured S21 to the analytical propagation equation. A fit produced using the Python 'numpy' least squares implementation is shown in figure A.17

Given that this fit and the earlier estimates suggest that $\epsilon_r$ is 1.7, this is the value that was used in the next stage of the antenna design.

**Antenna design**

A traditional textbook patch design was tried and then a GA. It was expected that the GA would be better but the other was trialled purely to see if it would work at all. To check this an FDTD simulation was done for a standard rectangular patch and it performed badly with the S11 parameter not showing any resonance around 2.4GHz. For the foam wing antenna bandwidth was viewed as very important because the measured dielectric constant had an uncertainty and so the wider the bandwidth the better, in order to negate the effects of a shift in centre frequency. We have seen that a GA can deliver more bandwidth. Also, the standard design equations for a patch antenna found in the literature and books such as [67] will not work for the thickness of foam being used. $\lambda/2$ is 62.5mm approximately but the antenna height

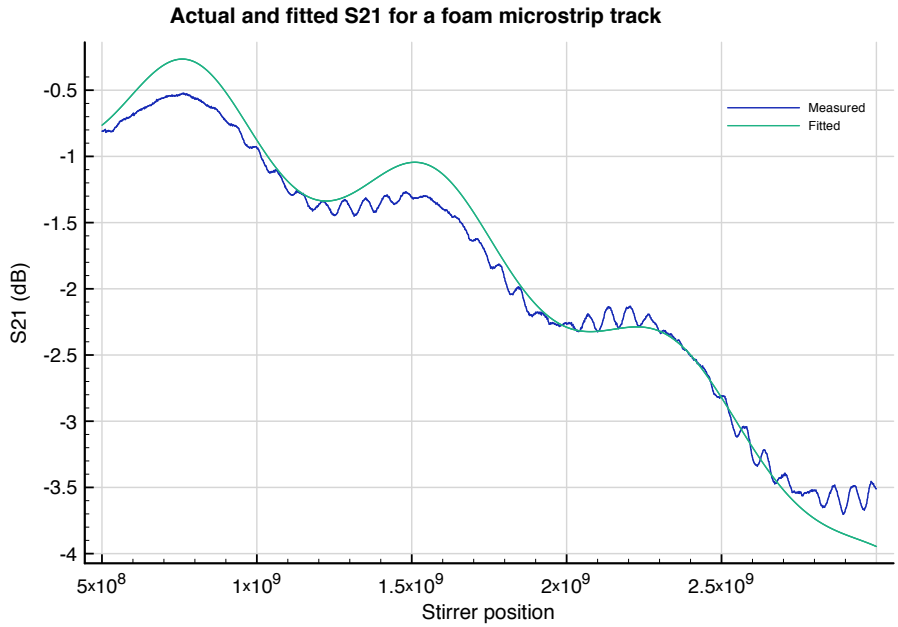


Figure A.17: The measured and simulated performance of the microstrip on blue foam when $\epsilon_r$ is 1.7 at 2.4GHz

was to be 11mm. The literature suggests a height very much smaller than the patch dimensions because otherwise the standard models that make assumptions such as ignoring fringing fields, do not apply and we found this to be the case. Therefore the evolutionary approach was the only option worth pursuing because it has the advantage of not making any assumptions about dimensions, frequencies or otherwise - it will just try random designs, simulate them and optimize them until it finds a good design.

4 runs were carried out (each taking in the region of 18 hours) and the best result had the S11 parameter shown in figure A.18. This had dimensions of 39mm x 39mm. Power was not considered at this point, the most important requirement was to get an antenna that was resonant with low return loss at the 2.442 centre frequency and with sufficient bandwidth. As the figure shows this was well within the specification providing greater -10dBs bandwidth than the ZigBee band. The antenna design was passed onto partners at Manchester and Liverpool university who fabricated the wing section and the antennas, respectively. Note that two antennas were fabricated for two ZigBee devices to communicate. Upon testing with the radios connected it was found that data could be reliably sent between them, using a test application that moved an LED bar graph as air was blown over a pressure sensor on one end of the wing.

The fabricated antennas can be seen in figure A.19 and A.20 and the full wing section is

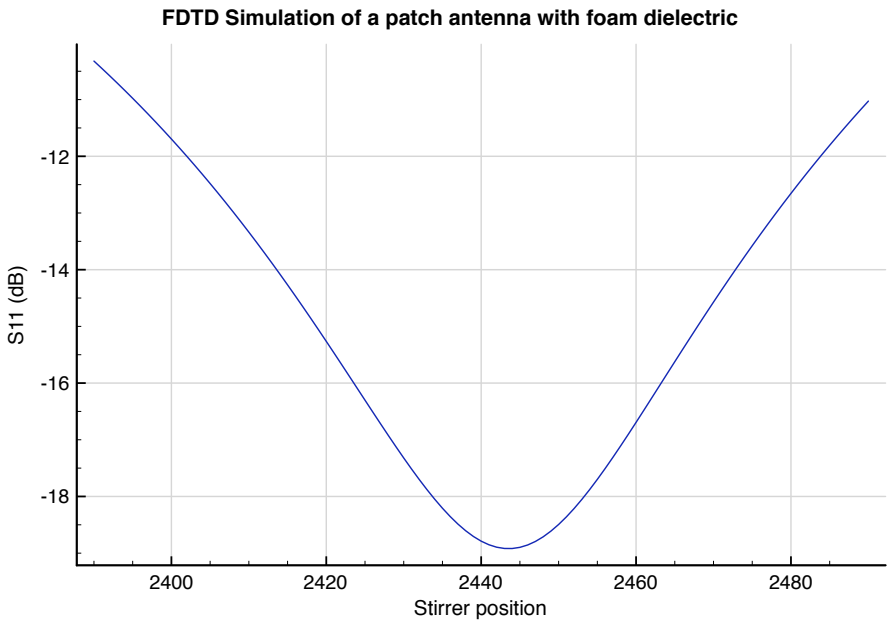


Figure A.18: FDTD simulation of the foam substrate antenna

shown in figure A.21 along with the feed in figure A.22

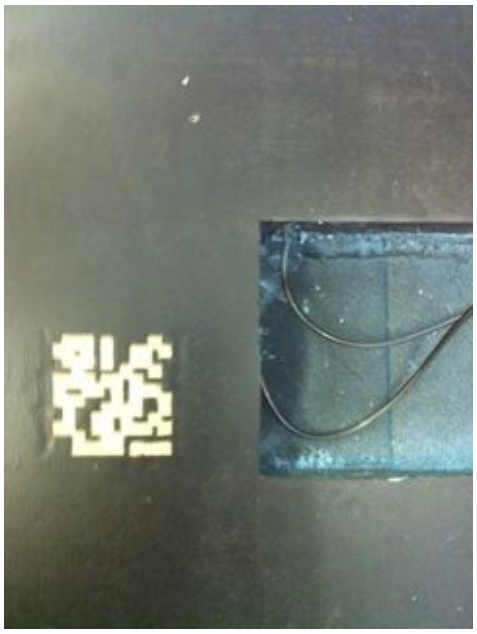

Figure A.19: Wing patch antenna

Figure A.20: WingPatchAntenna

Figure A.21: The wing before measurement in the Anechoic chamber

Figure A.22: WingPatchAntenna

In order to measure these antennas an anechoic chamber was used for the reflection coefficient and a reverberation chamber for efficiency. Directivity was ignored because it was not in the specification of the antenna. The S11 is plotted in figure A.23 Overall the S11 performance

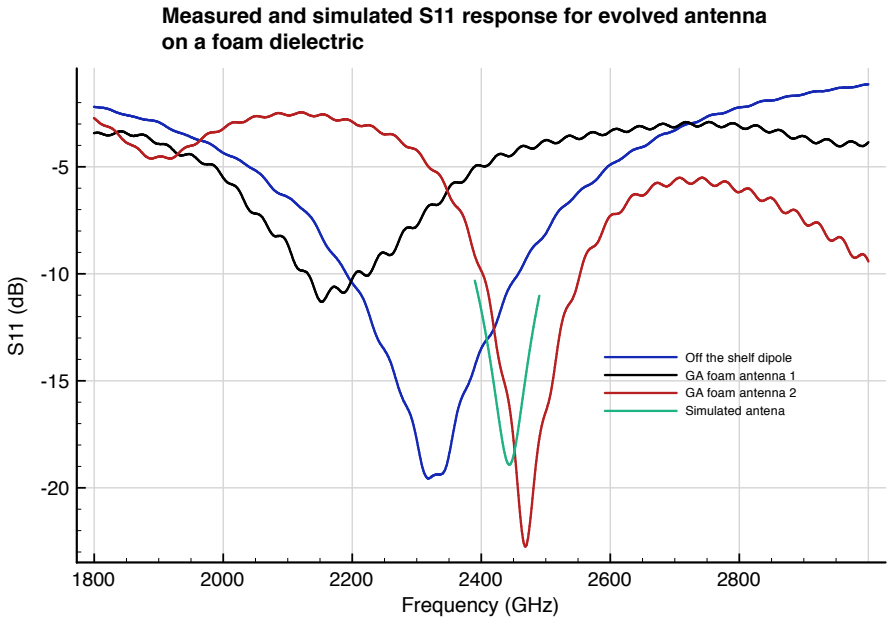


Figure A.23: Measured responses of GA foam antennas

of one of the antennas was excellent matching up very closely to the simulation whereas the other was a long way off. It is likely that the other had a problem in the way the SMA connector or ground plane was attached given that the patch itself was identical. The S11 bandwidth is not quite as high as for the dipole (one of the ones used in measurements in previous chapters) but it would be good enough for use with a ZigBee system.

The wing was placed in the reverberation chamber to measure efficiency, using a technique described more thoroughly in Appendix B. First two horn antennas of known efficiency were placed in the chamber and their response (through the S21 parameter) was measured over 100 stirrer positions. Following this one antenna was swapped out and replaced by one of the wing antennas. A metal sheet was placed over the other antenna to minimise its energy absorption. For both the horns and the patch the Q for the chamber was averaged over position and the frequency range 2.4GHz to 2.45GHz. Q is obtained by a relationship of the form:

$$Q \propto \frac{< S21 >}{\eta_1 \eta_2} \tag{A.6}$$

The Q of the chamber is not changing so we can work out the difference in antenna efficiency

between the two measurements. It was found that the average efficiency was 1/3 that of the horn. Assuming an efficiency of 0.9 for the horn (from the data sheet) this gives us an efficiency of 0.3 for the patch. Although this could not be regarded as a high efficiency the performance is good considering it is a new untested antenna design, the fabrication was poor and the patch was using silver inks rather than copper. Overall, given this was a first attempt at building the antenna into the foam of a wing, it proved the principle very effectively.

## A.4 Conclusions

In this chapter the problem of designing an antenna for direct write has been discussed and some successes have been presented. DW allows for low profile antenna designs written directly onto a surface so it could be very useful on aircraft for retrofitting antennas, placing on surfaces where holes cannot be drilled and it may be good for the outside of an aircraft for aerodynamic reasons. At the same time it has been seen that it presents challenges. The techniques available meant that patches were the easiest antenna type to manufacture but that they had to be made as a grid rather than a solid patch. One of the interesting results shown in this chapter is that it is possible to use a grid in place of a complete patch with almost no performance impact. As well as this we have seen that there is little performance penalty in using a DW antenna rather than an etched copper one which was a positive result as there had been worries over the silver inks used in the DW process not being sufficiently conductive. One of the FR4 gridded antennas was only 3dBs below the maximum power of the a solid copper grid and its bandwidth and S11 response were just as good.

From an electromagnetics perspective the use of an elementary patch design brought difficulties in terms of generating sufficient bandwidth for the full ZigBee band and in keeping the centre frequency in the correct place. These problems were seen throughout the simulations and measurements and were due to various factors, including the inaccuracy in the manufacture of some of the early DW antennas to an uncertainty in the dielectric constant of FR4. If these types of antenna are to be used for fly-by-wireless then the manufacturing process will need finer control and this something that Liverpool University improved considerably over the course of the project. We have seen though that novel antenna design can also counter this problem by expanding the bandwidth of the antenna, allowing it to tolerate a greater offset of the centre frequency. By utilising an evolutionary algorithm the bandwidth is expanded and the antenna becomes useful: firstly it won't be adversely affected by small offsets but also it opens up the possibility (regulatory bodies allowing) of widening the range of frequencies available which is something we might wish to do to boost reliability.

The final antenna design was particularly novel because it actually used the aircraft structure as the dielectric. Again, there is no way using conventional techniques that one could design an antenna for this application and but by utilising the optimisation of a GA it has been possible. We should note that all these antennas were designed with very little experimentation in parameters of the GA and with only 18 hour runs. Much better performance yet could

be found by tweaking the algorithm and increasing the run time to allow more generations in the GA. There is one downside to using these GA patch antennas and that is the unpredictable radiation patters they produce which we have not discussed here because they could not be controlled as part of the GA design. This could be added to the parameters of the GA and included in the assessment of each antenna's "fitness". However, it was not regarded as a problem for this application because distances are generally small and typically these antennas would be used within a resonant environment. Once an antenna has excited a bay or resonant cavity then the modes of the cavity will be what determines the signal strength at a receive antenna, not the directivity of the antenna. In summary, we have demonstrated that a 2.4GHz antenna may be written directly onto the surface of an aircraft using direct write, when we combine direct write with a genetic algorithm design method, enabling greater flexibility in the placement of wireless radios in an aircraft.

# Appendix B

# A fast way to measure antenna efficiency

## B.1 Introduction

In this appendix a novel method for calculating antenna efficiency is described. The traditional approach to measuring an antennas efficiency can be a time consuming process since to obtain efficiency very accurately the complete field must be measured for a sphere around the antenna in free space or in an anechoic chamber and better ways to measure this are valuable. Two known methods will be described and then the new method. Before going any further we should be clear about what is meant by efficiency here, as it can be used to refer to ohmic losses relative to the radiation efficiency or may refer to all losses which means reflections should be included. Here we refer to all losses one could separate them by measuring S11 and removing that term.

## B.2 Relative Q

A lot of the work in previous chapters refers to Q-factor of reverberation chambers which is dependent on the volume of a cavity, the antenna efficiency, and the losses in the cavity. Q can be described in terms of energy and power dissipated as:

$$Q = \frac{2\pi f U}{P_d} \tag{B.1}$$

where $P_d$ is power dissipated, $f$ is frequency and U is energy stored. It can be measured in the reverberation chamber as:

$$Q = \frac{16\pi^2 V}{\lambda^3} \frac{< P_r >}{\eta_{tx}\eta_{rx}P_t} \tag{B.2}$$

where $\lambda$ is wavelength, $\eta$ is efficiency,$P_t$ is power transmitted, $P_r$ is power received and $V$ is volume Therefore in order to calculate efficiency the average power could be measured for 2 antennas for which efficiencies are known, then 1 can be swapped out and replaced with an antenna with an unknown efficiency. The change in the average field $< S21 >$ is proportional to the change in efficiency $\eta$ for the antenna that is swapped out. This is a standard technique for calculating efficiencies and it is illustrated in table B.1 showing the result of a applying equation B.2 for reverberation chamber measurements for 4 different antenna types.

| Antenna type | Q | Efficiency |
|---|---|---|
| Horn | 40182 | 0.9 |
| York full patch | 12364 | 0.27 |
| York gridded patch | 17963 | 0.40 |
| York ZigBee COTS dipole | 32871 | 0.74 |

Table B.1: Antennas, Q-factors and the antenna efficiencies calculated from those Q factors

If the efficiency of the horn is assumed to be 0.9 then by looking at the power ratios, for example 40182 to 32871 we can approximate the efficiency of the ZigBee antenna to 0.74. This is lower than the efficiency of 0.8 specified by the data sheet for the ZigBee antennas and this may be due to the fact that the horn antenna's efficiency is not known precisely.

Looking at the two patch antennas efficiencies appear to be very low because the average power is a lot lower so we are able to clearly see difference in antenna performance. Complications exist in the measurement, where for example thee bandwidths are narrow or not aligned, requiring recalibration between measurements for the band of interest. Overall though this is a simple measurement that takes no account of patterns and gains. Reverberation chamber antenna measurements can be rapidly setup and automated making them a lot more practical than anechoic measurements.

## B.3 Peaks

There is another method that can be considered. Q-factor may also be written as:

$$Q = \frac{f_0}{BW} \tag{B.3}$$

which says that for the individual modes in the chamber we may be able to get the Q by looking at the 3dB bandwidth of the modes. The modes can be modelled as second order systems (i.e. 2 poles) and we may fit such a function to the measured modes in order to automate the collection of Q for a number of modes. Now the modal structure of the chamber is a function of the chamber itself not the antenna. So although the Q according to the power calculation will vary from one antenna to another the frequency domain view will not change in terms of radios of frequency to bandwidth and we can recover the antenna efficiency by equating equations B.2 to B.3

In principle the method should work but it is flawed in practice.. Consider graphs B.1 and B.2. We see here that when we have a very reverberant chamber then there are a lot of very

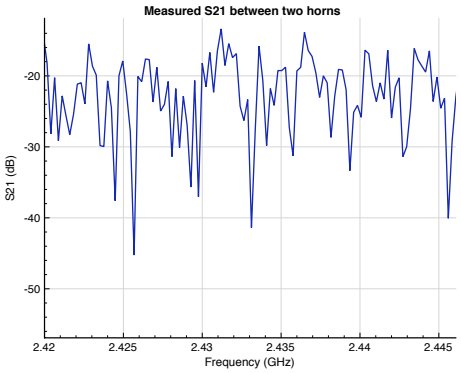


Figure B.1: A high Q chamber response

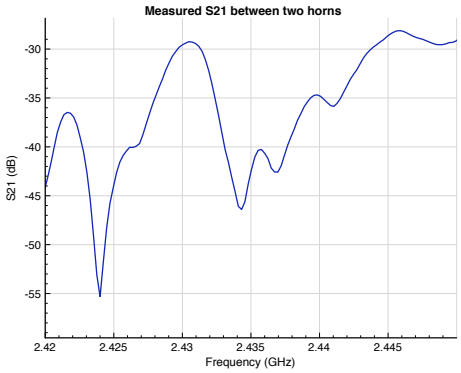


Figure B.2: A low Q response

narrow modes and they are hard to pick out accurately at the resolution shown here. For a lot of network analysers 1601 points is the standard for a measurement which does not provide a lot of resolution over a useful band, (for example the ISM band for the case of WLAN and ZigBee). This can be rectified by doing a number of sweeps over narrow frequency bands to increase the resolution. However, there are nevertheless issues in terms of the resolution of the analyser itself in terms of its filters and the issue of overlapping modes within the chamber response itself. When the peaks are made wider, with additional damping in the chamber then the overlapping is increased between peaks and it becomes harder still to obtain the Q

accurately. Therefore this method is hard to use because individual peaks cannot be resolved.

## B.4 Using the time response

The work on the time responses of reverberation chambers reported in the previous chapters offered some inspiration for such a technique. Before looking at that, An alternative is available though and this is the time response of the chamber. Each mode has it's own Q and time response which is hard to obtain but we may alternatively get a composite time constant $\tau$ for the chamber by considering the Fourier Transform of the time response over a band as was done when looking at delay spreads. When we look at the average Q then we have the expression given in terms of $S21$ and this may be rewritten and solved as shown to obtain a time constant: Equation B.2 may be rewritten as:

$$Q\frac{dE}{dt} = \omega E \tag{B.4}$$

where we have rewritten energy dissipated per cycle $P_t$ as a differential multiplied by $f$ and have treated the energy stored as a starting energy because to see the decay we must imagine that the energy being added to counter the energy being dissipated has been removed. This differential equation can be solved as follows:

$$\frac{dE}{dt} = \frac{1}{Q}\omega E \tag{B.5}$$

$$E = Ae^{-\frac{\omega}{Q}t} + B \tag{B.6}$$

$$E = Ue^{-\frac{t}{\tau}} \tag{B.7}$$

$$\tau = \frac{Q}{\omega} \tag{B.8}$$

We see therefore that the time decay is related to the chamber Q and the frequency of operation. For a given frequency or small band of frequencies we can make the time constant from the measured decay equal to the time taken from $Q/\omega$.

The method is as follows:

- Place two identical antennas in the reverberation chamber
- Measure S21 over a frequency band and over a set of stirrer positions
- Obtain the Q from the average of S21 using equation B.2 with $\eta$ set to 1 and then turn it into $\tau$ as above.

- Transform the frequency responses into time responses and calculate the mean delay for each response and average them
- Equate the two Qs. Divide the frequency based Q by the time based Q and take the square root. The result is the efficiency

This gives one number but by using lots of small bands one could see how efficiency changed with frequency

In order to test the technique a set of antennas were selected and placed in the reverberation chamber. These included 2 horn antennas and 2 log periodic antennas. The test setup can be seen in figures B.3 and B.4

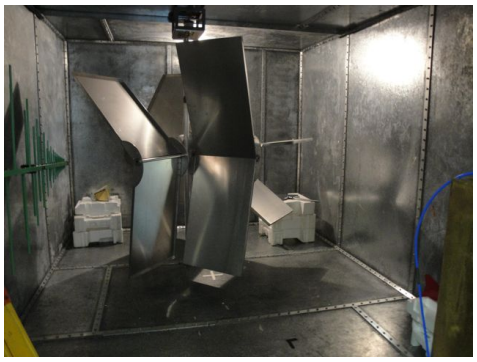

Figure B.3: Reverberation chamber setup

Figure B.4: One of the log periodic antennas

Measurements were carried out between most combinations of antenna at 3 different frequency bands:

- 2.4 - 2.48GHz
- 1-1.08GHz
- 500 - 580MHz

The narrow bands used are 80MHz each with 1601 points thus covering the majority of the ISM band for the 2.4GHz case and not being so big that a Q change might be seen over the range. Plots of stirrer averaged S21 are given in figure B.5 for horn to horn and log periodic to log periodic: Here we see that the horns start to work after 1GHz whereas the log periodic antennas are more effective at frequencies under 1GHz. It is clear that these antennas have significantly varying responses between the measured bands. On the other hand we would expect the time constant measurements to be constant regardless of antenna in a given band.

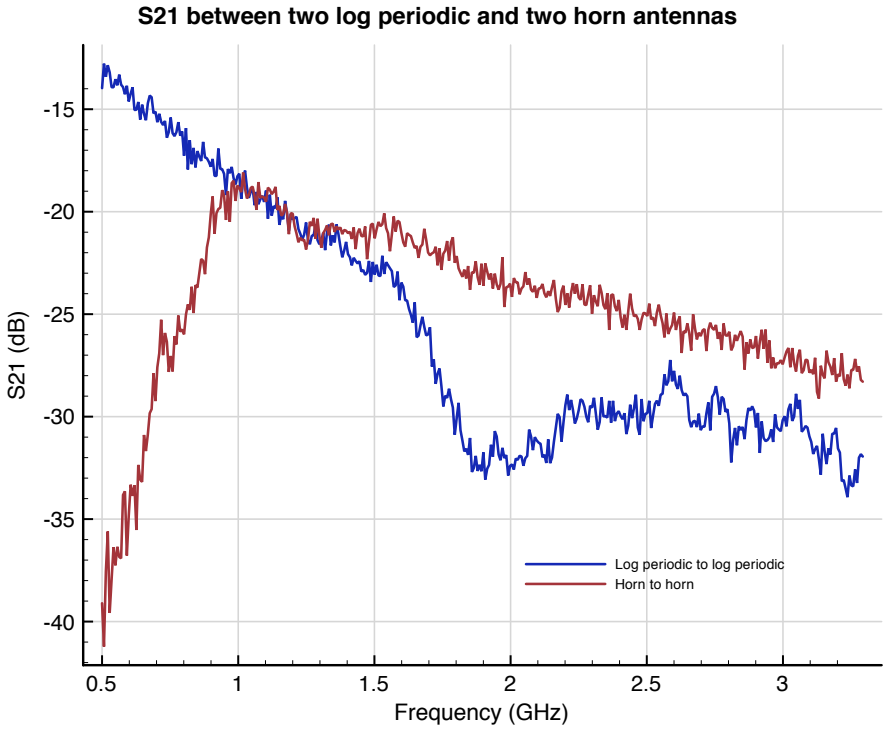


Figure B.5: S21 between log periodic and horn antennas compared

The Qs and time responses for the 1GHz band are provided in B.2 and we would expect from looking at figure B.5 the efficiencies should be very similar.

We see here that the measured Qs are very close as would be expected and that the time constant varies between around 1.15 $\mu$s by approximately +/- 0.5. The differences in the Q-factor and tau do not appear to be correlated. If we were to say that $tau$ is 1.15 for the chamber then this gives a Q of 7200 and an efficiency of 0.72 for a measured Q of 3700. This is lower than we might have expected from the antenna data sheets but not implausible.

A more interesting case is shown in table B.3 the Qs and $\tau$s are plotted for 2.4GHz and 500MHz. This time the Qs have all changed dramatically but the time constant stays the same. This is because the antenna efficiencies are now different for different antennas with the horns having high efficiency but standard ZigBee antennas a slightly lower efficiency. In fact this measurement revealed a problem with one of the ZigBee antennas which has a much lower efficiency than another of the same type, suggesting damage. The horn ZigBee combined efficiency at 2.4GHz is exactly as would be predicted from the data sheets. For the 500MHz measurements we see that the horns are not yet working efficiency and the combined efficiency for an LPDA and a horn is very low.

| Antenna 1 | Antenna 2 | Q | $\tau$ | sqrt($\eta_1\eta_2$) |
|---|---|---|---|---|
| Green LPDA | Grey LPDA | 3589 | 1.15 | 0.69 |
| Green LPDA | Red horn | 3705 | 1.19 | 0.69 |
| Green LPDA | Silver Horn | 3817 | 1.10 | 0.73 |
| Grey LPDA | Red Horn | 3659 | 1.07 | 0.72 |
| Grey LPDA | Silver Horn | 3594 | 1.21 | 0.67 |
| Red Horn | Silver Horn | 3609 | 1.15 | 0.69 |
| Silver Horn | Silver Horn | 3713 | 1.16 | 0.70 |

Table B.2: Antennas, Q-factors and the time constants associated with the measurement at 1GHz

In figure B.6 we see the relationship plotted graphically for results in both the 1GHz and 2.4GHz bands. What is very clear to see here is that there is no relationship between measured Q and measured time constant. Although time constant varies in the reverberation chamber with a standard deviation of 3.5% of the mean it is nothing like the order of magnitude Q variation with different antennas and frequencies.

These results clearly show that so long as one knows the time constant of a chamber at a particular frequency then it is possible to calculate efficiency (albeit with a slight question mark over absolute accuracy) with two identical antennas.

It is hoped in the near future to validate these efficiencies against anechoic or OATS measurements and to see if the same thing can be done with an S11 measurement because this would mean only one antenna was needed.

| 2400-2480MHz | | | | |
|---|---|---|---|---|
| Antenna 1 | Antenna 2 | Q | $\tau$ | sqrt($\eta_1\eta_2$) |
| Green LPDA | Red Horn | 8754 | 1.1 | 0.72 |
| Red Horn | Silver horn | 11682 | 1.18 | 0.8 |
| Silver Horn | Silver Horn | 11074 | 1.15 | 0.79 |
| Silver Horn | ZigBee | 9851 | 1.14 | 0.75 |
| ZigBee | ZigBee | 2083 | 1.11 | 0.34 |
| 500-580MHz | | | | |
| Green LPDA | Grey LPDA | 1721 | 1.06 | 0.69 |
| Green LPDA | Red Horn | 156 | 1.11 | 0.2 |

Table B.3: Antennas, Q-factors and the time constants associated with the measurement at 2.4GHz and 500MHz

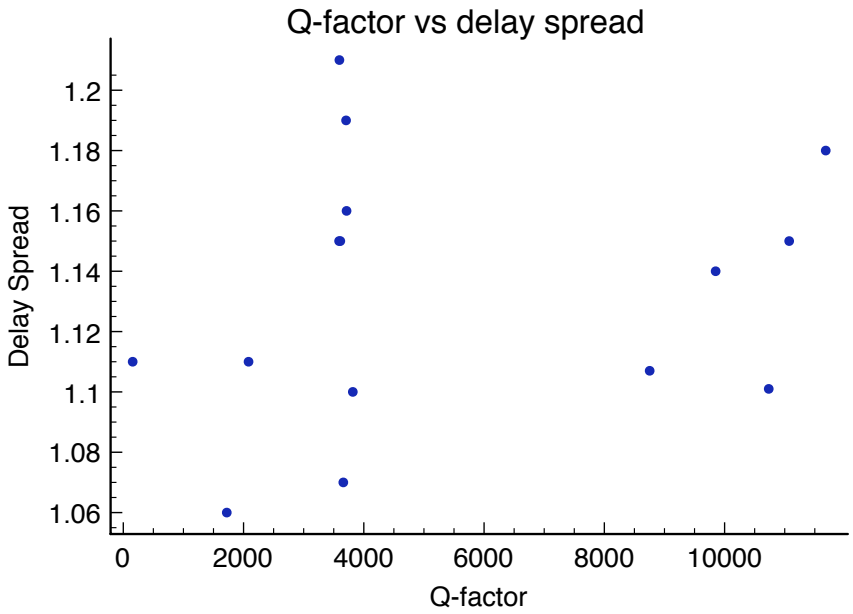


Figure B.6: A high Q chamber response

# Appendix C

# Jennic architecture

Since measurements of errors have been carried out with a Jennic development kit it is necessary to know what sort of receiver the modules use since the errors may be particular to the way Jennic works including the noise figure and whether there is carrier recover for example. The Jennic wireless kit that has been used in ZigBee tests uses a low-IF architecture and is structured as shown in figure C.1:

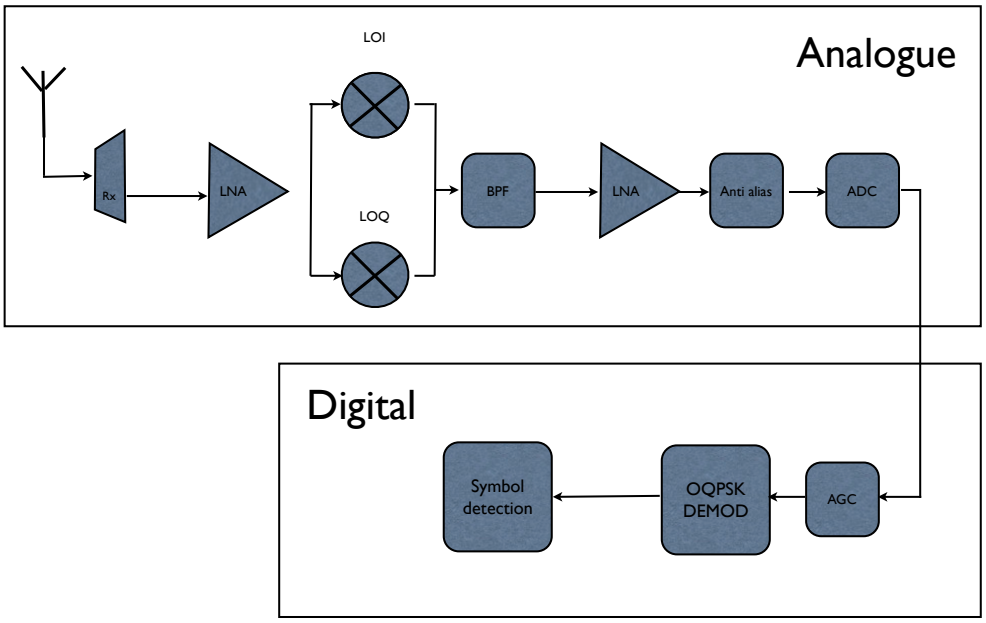


Figure C.1: Overview of Jennic's 5139 chipset receiver section

As this shows it is a low IF design whereby the signal is down converted in I and Q sections before being recombined and filtered in order to provide a certain amount of channel selec-

tion. Following this there are some digitally controlled variable gain amplifiers followed by an anti alias filter in preparation for the digitisation of the signal. From this point everything is done in the digital domain, the matched filtering followed by envelope detection and then symbol detection. After contacting Jennic it was found that the demodulation is non coherent and so no synchronisation is needed.

Jennic was contacted in order to find out more details about the architecture and how it might perform in a multipath environment. They provided the following information:

- Receiver noise figure at room temp is 10dB.
- For signal powers less than -100dBm (approx) the demodulator is not be enabled to conserve power
- If demod does lock onto a packet - it will then remain enabled regardless of the signal level.
- The demod in the 5139 is non-coherent, so there is no carrier-recovery/synchronisation to be performed.
- Demodulation requires detection of a minimum of 2 successive preamble symbols (1 preamble byte) followed by the SOF delimiter to acquire packet lock.
- Packet errors are predominantly due to noise.
- Multi-path signals would have to be strong relative to the wanted signal and be delayed by a few hundred ns (e.g. >100m equivalent path) to cause a problem

The exact method of correlation is not specified and more information here might be useful if any simulations need to be performed.

# Glossary

**3G** Third generation mobile phone standards. 1

**AFDX** Avionics Full Duplex Switched Ethernet . 13

**BER** Bit error rate (or ratio), ratio of bits in error to bits correctly decoded. 64

**Bluetooth** A frequency hopped packet radio service often used for wireless mice, keyboards and phones. 1

**CFC** Carbon Fibre Composite. 57

**COTS** Commercial off the shelf. 13

**Ethernet** A LAN standard specifying the physical layer and medium access control. 16

**Firewire** A high speed serial data bus standard often used for digital video. 16

**HDMI** High Definition Multimedia Interface, a high data rate digital sound and video standard. 1

**HUD** Head-up display. 7

**JIAWG** Joint Integrate Avionics Working Group. 13

**LAN** Local Area Network. 1

**MIMO** Multiple in multiple out - refers here to multiple antenna systems. 20, 21

**MSK** Minimum Shift Keying. 25

**PER** Packet error rate (or ratio), ratio of dat packets in error to packets correctly decoded. 64

**SCSI** Small Computer System Interface. 12

**TDMA** Time division multiple access, giving different users different time slows to transmit. 29

**UAV** Unmanned Aerial Vehicle. 5, 7

**WPAN** Wireless personal area network. 21

# References


[1] W. Cinibulk. Aircraft electrical wire, wire manufacturers perspective. http://www.mitrecaasd.org/atsrac/FAA_PI-Engineer_Workshop/2001/aircraft_electrical_wire.pdf, September 2010.

[2] D. Shannon. Airbus reveals latest a380 delay will cost 2.8 billion euros in pre-tax earnings 2006-2010. http://www.flightglobal.com/articles/2006/10/04/209671/airbus-reveals-latest-a380-delay-will-cost-eads-2.8-billion-in-pre-tax-earnings.html, September 2010.

[3] A. Paterson. Aircraft electrical fire and the case for the virgin electrical bus. http://www.vision.net.au/~apaterson/aviation/electrical_fire.htm, September 2010.

[4] D. H. Middleton, editor. *Avionic Systems*. Longman Scientific and Technical, 1989.

[5] Focus on the mpeg-4 fast motion codec. http://www.tomshardware.com/reviews/mpeg,293-4.html, September 2010.

[6] E. Pastor, J. Lopez, and P. Royo. A hardware/software architecture for uav payload and mission control. In *25th Digital Avionics Systems Conference, 2006 IEEE/AIAA*, pages 1 –8, 614 2006.

[7] D. Uhlig, K. Bhamidipati, and N. Neogi. Safety and reliability within uav construction. In *25th Digital Avionics Systems Conference, 2006 IEEE/AIAA*, pages 1 –9, october 2006.

[8] J.S. Dittrich and E.N. Johnson. Multi-sensor navigation system for an autonomous helicopter. In *Digital Avionics Systems Conference, 2002. Proceedings. The 21st*, volume 2, pages 8C1–1 – 8C1–19 vol.2, 2002.

[9] I. Moir and A. Seabridge. *Civil Avionics Systems*. Wiley, 2006. ISBN 978-0-470-02929-9.

[10] I. Moir and A. Seabridge. *Military Avionics Systems.* Wiley, 2006.

[11] What is adfx. http://www.afdx.com/, September 2010.

[12] F22 raptor avionics. http://www.globalsecurity.org/military/systems/aircraft/f-22-avionics.htm, September 2010.

[13] Raytheon products page. http://www.raytheon.com/products/f22_cip/, September 2010.

[14] B. Latre, P. D. Mil, I. Moerman, B. Dhoedt, and P. Demeester. Throughput and delay analysis of unslotted IEEE 802.15.4. *Journal of Networks*, 1(1):20 – 28, 2006.

[15] IEEE Std 802.15.4-2003. http://www.ieee802.org/15/pub/TG4.html, September 2010.

[16] Zigbee Specification. http://www.zigbee.org/en/index.asp, September 2010.

[17] IEEE 802.15 wpan. http://www.ieee802.org/15/pub/TG1.html, September 2010.

[18] Bluetooth tutorial. http://www.tutorial-reports.com/wireless/bluetooth?PHPSESSID=00d6207d7dbc0f42779e979ed7b28789, September 2010.

[19] K. Hole. Bluetooth overview. http://www.kjhole.com/Standards/BT/BTdownloads.html, September 2010.

[20] IEEE 802.15 WPAN task group 3 home page. http://www.ieee802.org/15/pub/TG3.html, September 2010.

[21] Jennic website. http://www.Jennic.co.uk.

[22] M. Gast. *801.11 Wireless Networks, 2nd Edition.* O'Reilly, 2008.

[23] A. Vinck F. Pavlidou, H. Latchman. Powerline communications and applications. *International Journal of Communications Systems*, (16):357 – 361, 2003.

[24] M. Lee, R. Newman, H. Latchman, S. Katar, and L. Yonge. Homeplug 1.0 powerline communication lans - protocol description and performance results. *International Journal of Communications Systems*, (16):447 – 473, 2003.

[25] G. Bianchi and G. Conigliaro I. Tinnirello. Design and performance evaluation of an hybrid reservationpolling mac protocol for power-line communications. *International Journal of Communications Systems*, (16):427 – 445, 2003.

[26] Homeplug 1.0 technology white paper. http://www.homeplug.org/products/whitepapers/HP_1.0_TechnicalWhitePaper_FINAL.pdf, September 2010.

[27] I. Glover and P. Grant. *Digital Communications*. Pearson Education, 2004. ISBN 0-130-89399-4.

[28] Homeplug 1.0 av white paper. http://www.homeplug.org/products/whitepapers/HPAV-White-Paper_050818.pdf, September 2010.

[29] Moteworks product introduction. http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MoteWorks_OEM_Edition.pdf, September 2010.

[30] Paulo Carvalhal, Cristina Santos, Manuel Ferreira, Luis Silva, and Jose A. Afonso. Design and development of a fly-by-wireless uav platform, 2009.

[31] W. Pond. Considerations for a wireless primary flight control system in a commercial airline. *Caneus conference*, March 2007.

[32] J. Daniel. Fly-by-wirless airbus end-user viewpoint. *Caneus conference*, March 2007.

[33] H. Blair. Securaplane wireless overview. *Caneus conference*, March 2007.

[34] T. Smith. Aircraft wiring reduction. *Caneus conference*, March 2007.

[35] E. Valencia, J. Perotti, and R. Birr. Real-time wireless data aquisition. rf health node and sensornet. *Caneus conference*, March 2007.

[36] K. Champaigne. Techniques for improving reliablility of wireless sensor networks in flight applications. *Caneus conference*, March 2007.

[37] K. Kiefer. Real-world experience in wireless instrumentation and control systems. *Caneus conference*, March 2007.

[38] J Croft. Gulfstream fly-by-wireless trials to continue. http://www.flightglobal.com/articles/2008/10/31/318170/gulfstream-fly-by-wireless-trials-to-continue.html, September 2010.

[39] M. Ahmed, C.U. Saraydar, T. ElBatt, Jijun Yin, T. Talty, and M. Ames. Intra-vehicular wireless networks. In *Globecom Workshops, 2007 IEEE*, pages 1 –9, nov. 2007.

[40] J.P. Carmo, P.M. Mendes, C. Couto, and J.H. Correia. A 2.4-ghz cmos short-range wireless-sensor-network interface for automotive applications. *Industrial Electronics, IEEE Transactions on*, 57(5):1764 –1771, may 2010.

[41] E. Pastor, J. Lopez, and P. Royo. An embedded architecture for mission control of unmanned aerial vehicles. *Digital Systems Design, Euromicro Symposium on*, 0:554–560, 2006.

[42] David A. Hill. *Electromagnetic Fields in Cavities: Deterministic and Statistical Theories*. Wiley-IEEE Press, 2009.

[43] T. S. Rappaport. *Wireless Communications Principles and Practice*. Prentice Hall, 1996. ISBN 0-13-042232-0.

[44] M. Simon and M. Alouni. *Digital Communication over Fading Channels*. Wiley Interscience, 2005. ISBN 0-471-64953-8.

[45] A.K. Shukla and S. Nirmala. Emi/emc for military aircraft and its challenges. In *Electro-Magnetic Interference and Compatibility (INCEMIC), 2006 Proceedings of the 9th International Conference on*, pages 101 –107, feb. 2006.

[46] J.E. Nanevicz. Static charging and its effects on avionic systems. *Electromagnetic Compatibility, IEEE Transactions on*, EMC-24(2):203 –209, May. 1982.

[47] M.S. Rao. EMI/EMC effects on EW receiver systems of military aircraft. In *Electromagnetic Interference Compatibility, 2008. INCEMIC 2008. 10th International Conference on*, pages 63 –67, Nov. 2008.

[48] D. Vine. Review of measurements of the RF spectrum of radiation from lightning. Technical report, NASA Technical Memorandum 87788, March 1986.

[49] J. Cutler. *Understanding Aircraft Structures*. Wiley-Blackwell, 2006.

[50] T.H.G. Megson. *Aircraft Structures for Engineering Students*. Elsevier Aerospace Engineering, 2007.

[51] G J Freyer, D M Johnson, M O Hatfield, and T A Loughry. Cavity to cavity coupling measurements in commercial aircraft and the implications for on-board operation of per-

sonal electronic devices. Technical report, Naval Surface Warefare Center, Dahlgren, Virginia, USA.

[52] J. P. Hof and D. D. Stancil. Wireless sensors in reverberant enclosures: characterizing a new radio channel. *Vehicular Technology Conference Proceedings*, (3):1747 – 1750, 2005.

[53] S Greco and M Sarto. Low q reverberation chamber to reproduce aircraft-like em environment. *EMC 2006, Barcelona proceedings*, pages 502 – 507, 2006.

[54] D M Johnson, M O Hatfield, M B SLocum, T A Loughry, A R Ondrejka, R T Johnk, and G J Freyer. Phase II demonstration test of the electromagnetic reverberation characteristics of a large transport aircraft. Technical Report NWSCDD/TR-97/84, Naval Surface Warefare Center, Dahlgren, Virginia, USA, September 1997.

[55] Xiaoming Chen, P.-S. Kildal, C. Orlenius, and J. Carlsson. Channel sounding of loaded reverberation chamber for over-the-air testing of wireless devices: Coherence bandwidth versus average mode bandwidth and delay spread. *Antennas and Wireless Propagation Letters, IEEE*, 8:678 –681, 2009.

[56] A. Sorrentino, G. Ferrara, and M. Migliaccio. On the coherence time control of a continuous mode stirred reverberating chamber. *Antennas and Propagation, IEEE Transactions on*, 57(10):3372 –3374, oct 2009.

[57] O. Delangre, P. De Doncker, S. Van Roy, M. Lienard, and P. Degauque. Characterization of a confined environment based on acoustic and reverberation chamber theory. comparison with the case of a car. In *Communications and Vehicular Technology in the Benelux, 2007 14th IEEE Symposium on*, pages 1 –4, nov. 2007.

[58] C. A. Balanis. *Advanced Engineering Electromagnetics*. Wiley, 1989.

[59] J. Solo, H. M. El-Sallabi, and P. Vainikainen. The distribution of the product of independent rayleigh random variables. *IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION*, 54(2):639 – 643, 2006.

[60] J. G. Proakis M. Salehi. *Digital Communications International Edition*. McGraw-Hill, 2005.

[61] Eldad Perahia and Robert Stacey. *Next Generation Wireless LANs*. Cambridge University Press, 2008. ISBN 978-0-521-88584-3.

[62] T. Sato, E. Fearon, C. Curran, K. G Watkins, G. Dearden, and D. Eckford. Laser-assisted direct write for aerospace applications. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 224(4):519–526, 2010.

[63] Direct write. `http://www.baesystems.com/ProductsServices/ss_tes_atc_adv_mat_sm_mat_dw.html`.

[64] J. A. Lewis. Direct ink writing of 3d functional materials. *Advanced Functional Materials*, (16):2193–2204, 2006.

[65] K. L. Wong, editor. *Planar Antennas for Wireless Communications*. Wiley-Blackwell, 2003.

[66] T. Sato, G. Deardon, and E. Fearon. Towards a wireless aircraft - october 2007 progress report. Technical report, Lairdside Laser Engineering Centre, University of Liverpool, University of Liverpool, October 2007.

[67] C.A. Balanis. *Antenna theory, analysis and design*. wiley, 1982.

[68] K. Carver and J. Mink. Microstrip antenna technology. *Antennas and Propagation, IEEE Transactions on*, 29(1):2 – 24, January 1981.

[69] P. S. Nakar. Phd thesis - design of a compact microstrip patch antenna for use in wireless/cellular devices. Technical report, Florida State University, April 2004.

[70] D. J. Griffiths. *Introduction to Electrodynamics*. Pearson, 2008. ISBN 0-13-919960-8.