



**University of
Nottingham**

UK | CHINA | MALAYSIA

Robust, High-Speed Intra-Vehicular Communication Networks for the Vehicles of Tomorrow

Submitted August 2020, to obtain the degree of
Doctor of Philosophy
in
Electrical and Electronic Engineering.

**Michael James Basford
4275926**

Supervised by Dr. S. Greedy & Prof. M. Sumner

George Green Institute for Electromagnetic Research - GGIEMR
Department of Electrical and Electronic Engineering

Signature _____

Date ____ / ____ / ____

The University of Nottingham, University Park, UK

Abstract

Cars of the 21st century are increasingly reliant upon computers for the most basic of applications. Whether it is safety-critical, such as the transfer of a braking command, or the more banal adjustment of a music playback device, cars have not escaped the relentless tide of digital dominance. As the trend for a higher degree of automation in everyday cars continues, this phenomenon shows little sign of abating soon.

However progress must not come about at the expense of safety. As the metal bars that once linked pedals to brakes were strong to ensure reliable operation, so too must we ensure data networks transfer signals reliably. Traditional means of ensuring the integrity of data on networks utilise shielded conductors, adding both mass and cost to vehicles.

This work presents an overview of contemporary intra-vehicular communication networks and presents the following new and novel findings:

1. A comprehensive survey of the intra-vehicular electromagnetic environment of a fully autonomous, electric vehicle.
2. The discovery of a complex relationship between the bit error ratio and time-domain interference which can be used to design networks with greater data throughput and resilience to interference.
3. The replication of the intra-vehicular environment for the development and verification of a wireless intra-vehicular communication system.

These findings challenge some existing practices regarding data integrity design methods as well as investigate the feasibility of removing wired communication links from intra-vehicular networks.

Publications

C. Smartt, M. J. Basford, S. Greedy, D. W. P. Thomas and M. Sumner, "EMC-oriented multi-conductor equivalent circuit cable models for spice, including transfer impedance coupling and incident field excitation," 2017 International Conference on Electromagnetics in Advanced Applications (ICEAA), Verona, 2017, pp. 1283-1287.

M. J. Basford, C. Smartt, D. W. P. Thomas and S. Greedy, "On the disruption of wired serial communication links by time domain interference," 2018 IEEE International Symposium on Electromagnetic Compatibility and 2018 IEEE Asia-Pacific Symposium on Electromagnetic Compatibility (EMC/APEMC), Singapore, 2018, pp. 183-186.

S. Greedy, C. Smartt, M. J. Basford and D. W. P. Thomas, "Open source cable models for EMI simulations," in IEEE Electromagnetic Compatibility Magazine, vol. 7, no. 3, pp. 69-81, 3rd Quarter 2018.

M. J. Basford, A. E. Pena-Quintal, S. Greedy, M. Sumner, D. W. P. Thomas, "Open Source Bit Error Rate FPGA-based Tester for Serial Communications", accepted for publication in 2020 IEEE Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), Sydney, 2020, in press.

A. E. Pena-Quintal, M. J. Basford, K. Niewiadomski, S. Greedy, M. Sumner, D. W. P. Thomas, "Serial Communication Links Modelling under Radiated EMI and its Impact on Sampling Errors in the Physical Layer", accepted for publication in 2020 EMC Europe, Rome, 2020, in press.

Acknowledgements

I would like to thank all those who have helped me over the last four years. The support and motivation of my supervisors, Prof. M. Sumner and Dr. S. Greedy, has encouraged me throughout my PhD.

I would also like to thank F. Munir and Prof. M. Jeon who facilitated an enlightening research trip to the Machine Learning and Vision Group's autonomous car at the Gwangju Institute of Science and Technology.

Finally, my deepest gratitude to my family for their encouragement in my pursuits. I am grateful to my sister for her understanding and patience. I am forever indebted to my parents for their constant support, generosity and for inspiring me to become the best engineer I can be.

List of Figures

1.1	Historical COMEX Valuations for the Price of Raw Copper, 2000-2020 [6].	3
2.1	The interconnected relationships between space and time, as shown on a sinusoidal waveform.	11
2.2	The four primary means of EM coupling between a source and victim.	13
2.3	A capacitor with an AC source across its plates [16].	15
2.4	Models of digital waveforms	18
2.5	Conversion from time domain model of a data waveform to frequency domain model of the spectral contents [18].	20
2.6	The uses of shielding [23].	21
2.7	Maximum aperture diameters for varying frequencies.	23
2.8	A shielded, multi-conductor cable for use in vehicles [26].	23
2.9	Basic blocks of every wireless network [23].	25
2.10	Analogue modulated network, based on [23].	25
2.11	Digital modulated network, based on [23].	25
2.12	AM in the time domain, for 10 Hz carrier and 1 Hz data.	26
2.13	AM in the frequency domain, for 10 Hz carrier and 1 Hz data.	27
2.14	FM in the time domain, for 50 Hz carrier, 1 Hz data and 5 Hz of frequency deviation.	28

2.15	FM in the frequency domain, for 50 Hz carrier, 1 Hz data and 5 Hz of frequency deviation.	28
2.16	QPSK modulation	30
2.17	QPSK bit encoding on the complex plane.	31
2.18	Frequency/time map of FHSS under Bluetooth standards. Each carrier is separated by 1 MHz, accommodating 79 channels in the band. Hops between f_c occur approximately 1600 times per second [34]	32
3.1	TTL NAND gate. [27]	39
3.2	Graphical representation of TTL family levels.	40
3.3	A single LVTTTL UART frame.	41
3.4	A single LVTTTL UART frame.	43
3.5	CAN bus layout [63].	47
3.6	Effect of noise on data transmission.	48
3.7	The three segments of a FlexRay frame [70].	50
3.8	Detail of a FlexRay header.	50
4.1	Photo of the Kia Soul EV 2017 at the MLV, GIST, Republic of Korea.	53
4.2	Test locations around the Kia Soul EV 2017. Orange represents sensor locations, green represents computer locations, and red represents power supply locations. Edited from [79].	57
4.3	Typical setup for radiated measurements.	58
4.4	Example spectra to determine the net vehicular emissions in a noisy EM environment, as per CISPR 11 Annex C.	61
4.5	0.15 - 30 MHz spectra, measured with a 0.5 m monopole.	62
4.6	902 - 928 MHz spectra, measured with a 110 mm monopole.	64
4.7	902 - 928 MHz spectra, measured with a 110 mm monopole, whilst the car is moving.	65

4.8	2.3 - 2.5 GHz spectra, measured with a 110 mm monopole. . .	66
4.9	Difference in amplitude of measured signals, of varying distance to the emissions source.	67
4.10	Ball probe construction.	68
4.11	Measured electric and magnetic field emissions at the passenger seat.	69
4.12	Measured electric and magnetic field emissions in the engine compartment.	70
4.13	Measured electric and magnetic field emissions in the boot. . .	71
4.14	Setup for time domain measurements on the CAN bus. . . .	73
4.15	CAN bus voltages with only the standard vehicle in operation.	74
4.16	CAN bus voltages during initialisation of autonomous equipment.	74
4.17	Severe induced common and differential-mode noise.	76
4.18	Induced voltage on an idle CAN bus.	76
4.19	Expanded view of the region highlighted in red from Figure 4.18.	77
5.1	Hardware configuration of wired protocols under test.	83
5.2	RF layout for synthesising interfering waveforms.	84
5.3	Diagram of GTEM construction.	85
5.4	Communication network measurement setup.	87
5.5	Complete experimental setup for GTEM measurements. . . .	88
5.6	Time-domain measurement on the conductor of a 5 V TTL compatible communication network.	90
5.7	Effect of increases number of Fourier components of a square wave.	92
5.8	Frequency-domain sample of interference signal.	93

5.9	Effect of increasing the number of Fourier terms on TTL and CAN systems in a GTEM cell.	94
5.10	Side view of wiring loom (not to scale).	96
5.11	Above view of wiring loom (not to scale).	96
5.12	Experimental setup for measuring the BER of a network exposed to time-domain EMI.	98
5.13	Photograph of the interior of the anechoic chamber for crosstalk measurements.	99
5.14	Heat map of BER for systems with 1 M Ω impedance.	100
5.15	Linear plot of BER against number of Fourier terms. Note, lines are naturally sorted by interference amplitude.	100
5.16	Logarithm plot of BER against number of Fourier terms, showing detail of BER for lower interference voltages.	101
5.17	SACAMOS software for creating the loom model.	104
5.18	SACAMOS-generated model incorporated into LTspice.	104
5.19	Comparison of the simulated and measured complex transfer functions.	105
5.20	Comparison of the measured and simulated induced interference waveforms.	107
5.21	Comparison of measured and simulated induced interference waveforms; simulation input from a measured transfer function.	108
5.22	Simulated probability of bit errors. Interfering waveforms amplitudes are represented with each individual line.	109
5.23	Simulated probability of bit errors with a 60 ns sampling period. Interfering waveforms amplitudes are represented with each individual line.	110
5.24	A comparison of induced interference waveforms.	111
5.25	The effect of induced voltages on data, from transient LTspice simulations.	114

5.26	Relationship between number of Fourier terms and time above a threshold voltage.	115
6.1	Interconnections between fundamental BERT components.	121
6.2	Internal layout of a typical FPGA [112].	124
6.3	High-level layout of OSBERT HDL modules, with annotated interconnections. Modules synthesised by Quartus prime are shaded blue, whilst modules unique to OSBERT are shaded orange.	125
6.4	Schematic symbol for a D type flip flop.	126
6.5	Delay module implementation.	127
6.6	The design of the custom soft-CPU in OSBERT, based upon a NIOS II core. Note that several IO ports have been removed from the layout for clarity.	128
6.7	Detailed layout of BERT HDL module.	129
6.8	Connections of the LFSR in OSBERT.	130
6.9	Operating example of OSBERT through Termit terminal software. Commands sent to OSBERT are coloured blue, the responses from OSBERT in green.	132
6.10	Single transmit channel front end of OSBERT.	133
6.11	Example of metastability of a ball.	134
6.12	Excerpt from the schematics of OSBERT, showing the input comparator.	135
6.13	An operational OSBERT unit.	136
6.14	A differential BERT module for OSBERT.	138
7.1	X310 software defined radio unit.	146
7.2	QPSK transmitter in Simulink.	146
7.3	Bit generator module detail.	147
7.4	QPSK receiver in Simulink.	148

7.5	QPSK receiver module in Simulink.	149
7.6	Theoretical number of modes for the RC. Red lines indicate the value at 915 MHz.	153
7.7	Working volume in the GGIEMR's RC.	154
7.8	Location of test points in the Fiat Panda 2009. Edited from [140].	159
7.9	Experiment layout for measuring S-parameters around the car.	160
7.10	Photograph of the experiment around the vehicle.	160
7.11	Complex S_{21} between the driver's console and the boot with a horn antenna. Red line indicates 915 MHz.	161
7.12	Raw K-factor values for DBH. Red line indicates 915 MHz. . .	162
7.13	Selected K-factor values from around the vehicle.	163
7.14	Difference in K-factor in the RC for two test configurations.	165
7.15	Smoothed K-factors in the RC for two test configurations presented in Figure 7.14.	166
7.16	Comparison of RC and car K-factors. Red line indicates 915 MHz.	167
7.17	Transmission spectra for varying Tx gain for $K = 0.093$. . .	168
7.18	BER as a function of time for 1 Mbps transmission through the DSH channel.	169
7.19	Transmission spectra for varying Tx gain for $K = 0.026$. . .	171
7.20	BER as a function of time for 1 Mbps transmission through the DBH channel.	172
7.21	Transmission spectra for varying Tx gain for $K = 0.026$ at a reduced data rate.	173
7.22	BER as a function of time for 400 kbps transmission through the DBH channel.	174
7.23	Constellation diagram with sample-ordered colour grading. .	175

A.1	Electric and magnetic field emissions in the boot.	204
A.2	Electric and magnetic field emissions at the radio antenna. .	204
A.3	Electric and magnetic field emissions at the driver's console.	205
A.4	Electric and magnetic field emissions in the wheel well. . . .	205
A.5	Electric and magnetic field emissions at the rear passenger's location.	206
A.6	Electric and magnetic field emissions in engine compartment.	206
B.1	Colour map of BER for systems with $50\ \Omega$ impedance.	208
B.2	Linear plot of BER against number of Fourier terms. Note, lines are naturally sorted by interference amplitude.	208
B.3	Logarithm plot of BER against number of Fourier terms, showing detail of BER for lower interference voltages.	209
C.1	Probability of interference with a 1 ns sample period. Uses a measured transfer function.	212
C.2	Probability of interference with a 6 ns sample period. Uses a measured transfer function.	212
C.3	Probability of interference with a 10 ns sample period. Uses a measured transfer function.	213
D.1	OSBERT schematic page 1.	216
D.2	OSBERT schematic page 2.	217
D.3	OSBERT schematic page 3.	218
D.4	OSBERT schematic page 4.	219
D.5	OSBERT schematic page 5.	220
D.6	OSBERT schematic page 6.	221
D.7	PCB of OSBERT, top	222
D.8	PCB of OSBERT, bottom.	223
F.1	Complex S_{21} for DBH channel.	230

F.2	Complex S_{21} for DBM channel.	230
F.3	Complex S_{21} for DEH channel.	231
F.4	Complex S_{21} for DEM channel.	231
F.5	Complex S_{21} for DSH channel.	232
F.6	Complex S_{21} for DSM channel.	232
F.7	Complex S_{21} for EBH channel.	233
F.8	Complex S_{21} for EBM channel.	233

List of Tables

2.1	The four EM coupling mechanisms.	13
2.2	List of symbols for the integral forms of Maxwell's Equations	14
2.3	Sources of EMI	17
3.1	Notable digital modules within cars.	35
3.2	The OSI model [49].	37
3.3	Comparison of 5V TTL and LVTTTL operating parameters .	39
3.4	Comparison between intra-vehicular communication protocols.	51
4.1	Test locations on the Kia Soul EV 2017.	56
4.2	Antenna choices for the experimental regime.	59
4.3	Minimum scan time	60
4.4	Measuring Instrument Bandwidth	60
5.1	BERs for a square wave interference source	89
5.2	Harmonic verification of interfering waveform.	93
5.3	Samples above threshold from Figure 5.24.	112
6.1	Fundamental BERT Components	121
6.2	Comparison of commercially available BERTs	123
6.3	Truth table for D type flip flop.	126
6.4	OSBERT commands	131
6.5	Comparison of OSBERT capabilities	139
7.1	Selection of Wireless Networking Equipment	145

7.2	Fiat Panda 2009 test locations.	159
7.3	K-factor values at 915 MHz for selected channels.	164
7.4	Measured and RC K-factor values	166
E.1	USRP QPSK transmitter structure	226
E.2	USRP QPSK receiver structure	227

List of Abbreviations

AC	Alternating Current
AGC	Automatic Gain Control
AM	Analogue/Amplitude Modulation
BER	Bit Error Rate
BERT	Bit Error Ratio Tester
BPS	Bits per Second
CAN	Controller Area Network
CANH	CAN High
CANL	CAN Low
CISPR	Comité International Spécial des Perturbations Radio
CLB	Configurable Logic Block
CMOS	Complementary Metal-Oxide Semiconductor
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
DM	Digital Modulation

DUT Device Under Test

E Electric (Field)

ECU Engine Control Unit

EM Electromagnetic

EMC Electromagnetic Compatibility

EMF Electromotive Force

EMI Electromagnetic Interference

EV Electric Vehicle

EVM Error Vector Magnitude

FHSS Frequency Hopping Spread Spectrum

FM Frequency Modulation

FPGA Field Programmable Gate Array

GGIEMR George Green Institute for Electromagnetic Research

GIST Gwangju Institute of Science and Technology

GPIB General-Purpose Interface Bus

GPS Global Positioning System

GTEM Gigahertz Transverse Electromagnetic

H Magnetic (Field)

HDL Hardware Description Language

HF High Frequency

IC Integrated Circuit

ID Identification

IO Input/Output

ISM Industrial, Scientific and Medical

ISO International Organisation for Standardisation

LFSR Linear-Feedback Shift Register

LIDAR Light Detection and Ranging

LOS Line of Sight

LVTTL Low Voltage TTL

MATLAB Matrix Laboratory

MLV Machine Learning and Vision

NA Not Applicable

NAS Network-Attached Storage

OS Open Source

OSBERT Open Source Bit Error Ratio Tester

OSI Open Systems Interconnection

PC Personal Computer

PCB Printed Circuit Board

PLL Phase Locked Loop

PWL Piece-Wise Linear

QPSK Quadrature Phase Shift Keying

RADAR Radio Detection and Ranging

RAM Random Access Memory

RC Reverberation Chamber

RF Radio Frequency

RS232 Recommended Standard 232

S Scattering

SACAMOS State-of-the-Art Cable Models for Spice

SAE Society of Automotive Engineers

SDR Software Defined Radio

SE Shielding Effectiveness

SNR Signal to Noise Ratio

SPICE Simulation Program with Integrated Circuit Emphasis

TTL Transistor-Transistor Logic

UART Universal Asynchronous Receiver/Transmitter

UHF Ultra-High Frequency

UK United Kingdom

US United States (of America)

USRP Universal Software Radio Peripheral

VNA Vector Network Analyser

XNOR Exclusive-Not-Or

XOR Exclusive-Or

Contents

Abstract	i
Publications	iii
Acknowledgements	v
List of Figures	vii
List of Tables	xv
Glossary	xvii
1 Introduction	1
1.1 The Networks of the Cars of the Future	1
1.2 Research Objectives	4
1.3 Novel Contributions	6
1.4 Thesis Structure	7
2 Electromagnetic Interference and Compatibility in Cars	10
2.1 Sources of EMI	16
2.2 EMI Mitigation for Wired Networks	17
2.3 EMI Mitigation for Wireless Networks	24
2.4 Chapter Conclusion	33
3 Contemporary Intra-Vehicular Communications	34

3.1	Digital Modules and Networks in Cars	34
3.2	The OSI Model	37
3.3	Wired Automotive Communication Protocols	38
3.3.1	TTL UART Networks	38
3.3.2	RS232 Networks	43
3.3.3	Controller Area Networks	44
3.3.4	FlexRay Networks - the Answer?	49
3.4	Chapter Conclusion	51
4	The Vehicular Electromagnetic Environment	52
4.1	Experimental Background	53
4.2	Frequency Domain Measurements	61
4.3	Time Domain Measurements	72
4.4	EMC Recommendations	79
4.5	Chapter Conclusion	80
5	Radiated Interference on Wired Networks	81
5.1	Inducing Bit Errors on Wired Networks under Time Domain EMI	83
5.2	Induced Errors on a Level 1 System	95
5.3	Modelling Distinct BER Regions	102
5.4	Corroboration of Time-Domain Effects	113
5.5	Chapter Conclusion	116
6	Design of an Open-Source Bit Error Ratio Tester	119
6.1	Fundamental BERT Topology and Function	120
6.2	OSBERT - The Design	124
6.2.1	Firmware and HDL Design	125
6.2.2	Hardware Design	132
6.3	OSBERT - A Qualitative Analysis	136

6.4	Chapter Conclusion	140
7	Networks in a Wireless Domain	142
7.1	Creating Flexible Wireless Data Links	143
7.2	Replication of an Automotive Environment	151
7.3	Wireless Transmission Performance	168
7.4	Chapter Conclusion	176
8	Summary and Reflections	179
8.1	Contributions and reflections	180
8.2	Future Work	183
	Bibliography	184
	Appendices	203
A	Electromagnetic Measurements Around a Kia Soul EV	203
B	Additional Study of BER Variations	207
C	Simulated BER Plots Using Measured Transfer Function	211
D	OS BERT Schematics	215
E	USRP QPSK Parameters	225
F	Measurements for K-Factor Replication	229

Chapter 1

Introduction

This chapter presents an introduction to the primary premises behind the work exhibited in this thesis. It briefly discusses how intra-vehicular networks will need to adapt to the demanding electromagnetic (EM) environment of the cars of the future, before Chapter 1 then details the research objectives of this work, before concluding with a description of the layout of this thesis.

1.1 The Networks of the Cars of the Future

Contemporary vehicles are entirely reliant on the consistent and reliable operation of embedded systems and their digital communication networks to function. This is, perhaps, best indicated by the increase size of the embedded systems market, which is projected to reach a value close to 250 billion USD by 2023 [1], with over 20% of the United States (US) market being dominated by the automotive industry. This fast-growing market segment is pushed predominantly by an increasing demand for more sophisticated infotainment systems, but by 2026 this growth is anticipated to be driven by an vast increase in the number of consumer-owned autonomous vehicles. It is anticipated that by the year 2025, less than 1% of all newly

registered vehicles in the UK, as a proportion of all registrations, will have no amount of automation at all [2] - or a so-called level 0 of automation as defined by the Society of Automotive Engineers (SAE) standard J3016 [3]. Indeed, [2] forecasts that 40% of newly registered vehicles in the UK will, by 2026, have a level 3 degree of automation or higher - that is to say that the automated driving system will be the primary monitor of the driving environment and performs the entire dynamic driving task.

These autonomous cars, on each level of automation other than zero, will need a large set of sensors to take in environmental data, which will then be sent along intra-vehicular networks to be processed by on-board computers. Control decisions, such as to accelerate, break, or turn left or right, will then need to be relayed to the actuators which perform these physical operations on the car's controls directly. Since the adoption of intra-vehicular networks in the early 1980s, these networks have been wired [4], and these have served their purpose of ensuring a reliable transmission of data in noisy EM environments for nearly four decades, using techniques outlined in Chapter 2. However, these interference-reduction techniques often rely upon slowing the data rate of networks or introducing a greater amount of metallic, EM shielding. With these heavily shielded wired networks, connection looms typically exceed 30kg in even a small car, such as the Subaru Plug-in Stella [5]. Clearly there are massive economic incentives to reduce or even remove this mass altogether, along the entire lifespan of a given vehicle.

Considering the cost of the raw materials alone, 30 kg of copper is valued at approximately \$165 USD by the Commodity Exchange Group at the time of writing[6]. Historical data suggests the raw price of copper can fluctuate greatly depending on the economic climate and accessibility, but maintains a steady rising value over time as shown in Figure 1.1. When most companies focus their cost reduction efforts on reducing the price they are required to

pay for materials [7] or find that material prices are a major cost driver in the pricing of a product [8], it follows that over the manufacturing of over 95 million cars produced each year across the world [9], enormous reductions in the cost to purchase the raw materials for the manufacturing of each vehicle would occur.

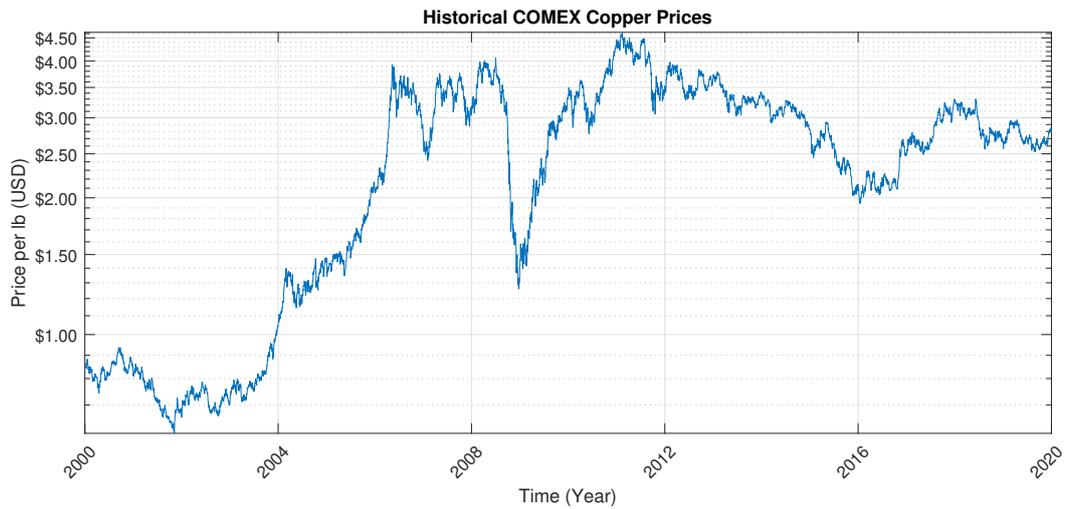


Figure 1.1: Historical COMEX Valuations for the Price of Raw Copper, 2000-2020 [6].

The potential benefits of removing the 30 kg mass of each communication network wiring loom from a vehicle would also be felt by the consumer. The results of a the study shown in [10] indicate that for a small car, a 140 lb (63 kg) reduction in mass resulted in a 0.9 mile-per-gallon increase in fuel efficiency while emulating city driving. This increase in fuel efficiency increased to 1.2 miles-per-gallon when under highway driving conditions. Whilst this mass reduction is nearly twice the 30 kg mass of a typical communications wiring loom, as given by [5], it is not unreasonable to expect an increase in the fuel efficiency of around 0.4 miles-per-gallon. As vehicle lifespans reach close to 200,000 miles on average [11], this would provide around an extra 2000 miles in range for the same amount of fuel, assuming a 38.9 mile-per-gallon efficiency given as a baseline by [10]. In financial terms, assuming the current price of petrol remains constant at 125 pence

per litre, this would equate to around £300 of savings for the consumer. From an environmental perspective this saving results in a reduction of 675 kg of CO₂ emissions from the vehicle over a 200,000 mile life.

Even as countries such as China, France, and the UK are proposing the phasing-out of fossil-fueled vehicles over the next 30 years [12, 13, 14], the same benefits from the reduction of mass offered by a wireless, intra-vehicular communications network would be applicable to electric vehicles (EV). An increase in range, greater fuel economy (albeit from an electric point of view), and reduced chemical emissions from the generation of electricity from non-renewable sources would all be felt. With these benefits, it seems vital that work should be undertaken to investigate the feasibility of such a system.

1.2 Research Objectives

The initial objective of this work is to investigate the feasibility of implementing a wireless data network inside of a car, with the intention of removing at least part of existing wired communication links. As outlined in Section 1.1, this would benefit manufacturers, consumers and the environment by saving costs in the materials, construction and operation of the car. With safety being critical in such a system, it is vital that such a system must prove to be reliable and operate regardless of the sources of EMI encountered in the vehicles of the future. If possible, the performance of this wireless system would be experimentally verified, with comparisons against wired counterparts.

It is also important to realise that the EM environment in which this wireless network should operate will likely differ from that of contemporary vehicles, due to the increased quantity of high frequency and high power devices found in electric and autonomous vehicles. As such, attention should

be paid to understanding what the EM environment will look like inside such a vehicle, and understand where especially polluted portions of the RF spectrum exist in both the frequency and spatial domains.

During the progress of this work, however, an interesting relationship between the shape of an interfering waveform and the errors it generated on a wired, victim network was observed. As a result of this, a new objective was decided upon in order to better understand the newly-observed phenomenon which had not been recognised in prior literature. This would have implications not only on the approach to EMC design on wired networks in cars, but indeed on potentially every wired data network suspected of both emitting, and being victim to, EMI. In order to perform this work to a high standard, it was realised that the use of specialised communications instrumentation would be required. Such equipment was initially unattainable by the author, and so an open-source instrument was to be designed - capable of meeting the requirements for the experiments to understand the wired EMI phenomena. By making such equipment open-source, it would provide the wider academic community with easy access to low-cost instrumentation which may otherwise prove inaccessible due to cost, or equipment flexibility.

The primary objectives of this thesis can be listed as follows:

1. To develop an understanding of the impact an increase in the electronic equipment, necessary for electric and autonomous cars, will have on the EM environment inside and around future cars.
2. To examine the effects of time-domain interference on wired communication links, looking at response of contemporary communication standards to interference, as well as develop an understanding of the influence of EMI on a system as a whole.
3. To develop an open-source means of testing the errors on communi-

cation networks under the influence of EMI.

4. To discover the feasibility of implementing a robust wireless communication link for intra-vehicular communications.

1.3 Novel Contributions

The work in this thesis has resulted in three primary novel contributions to knowledge. The first is the comprehensive EM survey performed in the intra-vehicular environment of a fully autonomous, electrically-driven car. These measurements identified numerous sources of EM interference across multiple frequency bands which were later used to shape the specifications of a proposed wireless intra-vehicular communications system. Furthermore, the findings from this EM survey identified catastrophic levels of time-domain interference on a digital communications network in the car, highlighting the need to further investigate the effect of time-domain interference on the data integrity of wired systems.

The second novel contribution is the identification of the highly complex relationship between the bit error ratio of a system and the rate of change of a time-domain interference source. It was found that contemporary advice to reduce the rate of change of a time-domain source as much as possible to reduce interference may hold true when considering the interference in a purely analogue domain, but that such rules-of-thumb do not necessarily hold true when the conversion of the analogue representation of the data bit to the digital domain takes place. As such, greater system efficiencies are achievable with a lower bit error ratio when sources are designed with a more holistic definition of interference. This work then provides a means of modelling the relationship between time-domain EMI and the integrity of digital data. It is worthwhile to consider that these findings have the potential to impact the design methodology of digital systems in general,

rather than being limited purely to the intra-vehicular case.

The third novel contribution in this this thesis is the measurement of wireless propagation inside the intra-vehicular environment, along with the development and verification of a wireless communication system. Measurements were made of the environment inside a typical car, and were used to recreate the intra-vehicular environment inside a reverberation chamber. A wireless network was created using off-the-shelf components and was found to operate reliably with data rates equal to the upper limits of popular, contemporary wired communication protocols. Such findings indicate that it is indeed practical to implement wireless communication networks for intra-vehicular communications, but limitations of lab instrumentation means additional work is required to determine how widespread the integration of wireless systems may become.

1.4 Thesis Structure

This thesis is divided into 8 distinct chapters, each describing the research procedures undertaken to meet the listed objectives. They cover a broad range of technical topics, ranging between developing an understanding of EMC and contemporary wired communication protocols from existing work to developing hardware to simulations. The work in each chapter, whilst distinct in itself, is closely linked to the work presented in other chapters - often influencing experimental procedures and results. As a result, this section lists each chapter, along with a brief explanation of the topics within.

Chapter 2 leads from this introduction in Chapter 1 to briefly provide a background to the topic of EMI as a whole, before specifically focusing on the impact of EMI within cars. This leads to an examination of some of the most popular means of mitigating EMI and ensuring EMC within data networks for both wired and wireless cases.

Chapter 3 performs a review of some contemporary intra-vehicular communication networks, allowing for an appreciation of the demands required by such networks to operate in a reliable manner. Due to the nature of data networks for intra-vehicular communication being wired, this naturally focuses entirely on wired cases only.

Chapter 4 starts presenting the first of the new and novel findings which have resulted from the work over the course of this project. EM measurements from around an electric, autonomous vehicle in various configurations to understand the sorts of EMI to be expected from vehicles of the future. Due to the specific car in question being retrofitted with equipment to allow autonomous driving, the study can also compare the EM spectra for both human and machine-controlled cars. This results in data which can be used to describe the EM environment in both the near and mid-term future.

Chapter 5 examines how EMI affects wired communications, with particular emphasis on radiated emissions in the time-domain. The examination of the effects of interference using the time-domain rather than the frequency-domain, is slowly being adopted by a greater number of EMC standards and the examination of it in this chapter engenders the possibility for influence on future standards. This chapter also examines both standard-defined communication protocols and abstracted data transmission methods to reveal an, as of yet, undocumented relationship between the form of time-domain interference and the induced errors when considering the system as a whole.

Chapter 6 discusses the design of a new, open-source bit-error-rate tester. This instrument is vital to achieve measurements which describe the real influence EMI has on the most important part of any digital communications system - the transfer from analogue to digital domain during bit determination. Most bit-error-rate tester systems today are very expensive and have specifications which may exceed the needs of fundamental research, such as

that present in Chapter 5. By designing an open-source tester, a versatile instrument may be made more accessible to the wider academic community.

Chapter 7 describes the technology and methods involved to generate a working wireless data connection, which is tested inside of a environment made to replicate that found inside a real car. The experimental findings in this chapter are then followed with a discussion on how they might lead to fully-functional networks with reliable operation in a car.

Chapter 8 closes the thesis with a conclusion covering the novel work presented. It reflects on these new contributions to the field of EMI and discusses how future work may provide a deeper understanding of the explored concepts.

Chapter 2

Electromagnetic Interference and Compatibility in Cars

The fundamental purpose of both wired and wireless communication links in any application is to transfer data from one system to another. In an ideal scenario, the data which is transmitted is the same as the data which is decoded at a receiver, however due to the various forms of electromagnetic interference (EMI) which exist, this is not always possible for every system in every environment. Indeed EMI, whether it is from natural or man-made sources, is the bane of a communication network engineer's life - generally having a greater impact on the reliability and data throughput of a network than other design challenges. It is unsurprising that many tomes have been written on techniques for reducing the EMI emissions themselves, or their ability to affect a victim network. Despite the complexities and intricacies within the detail of EMI (and its engineering twin; electromagnetic compatibility (EMC)) itself, it can be largely broken into quite digestible portions, described quite concisely. This chapter provides a contextual background to the most important EMI-specific concepts, covered in this thesis.

For any periodic waveform with a frequency f Hz, the period of the waveform T (seconds) is given by Equation (2.1) :

$$T = \frac{1}{f} \quad (2.1)$$

If the velocity of the wave is known, then the distance this wave has travelled, during this time period T , can be calculated. This distance is known as the wavelength. Within most cases of EMI, this velocity is taken as the velocity at which EM waves propagate through free space, or c , which is approximately equal to $3 \times 10^8 \text{ ms}^{-1}$. Therefore the wavelength, λ (metres) is given by Equation (2.2):

$$\lambda = \frac{c}{f} \quad (2.2)$$

This fundamental relationship between the period of a sinusoidal waveform and its wavelength is illustrated in Figure 2.1.

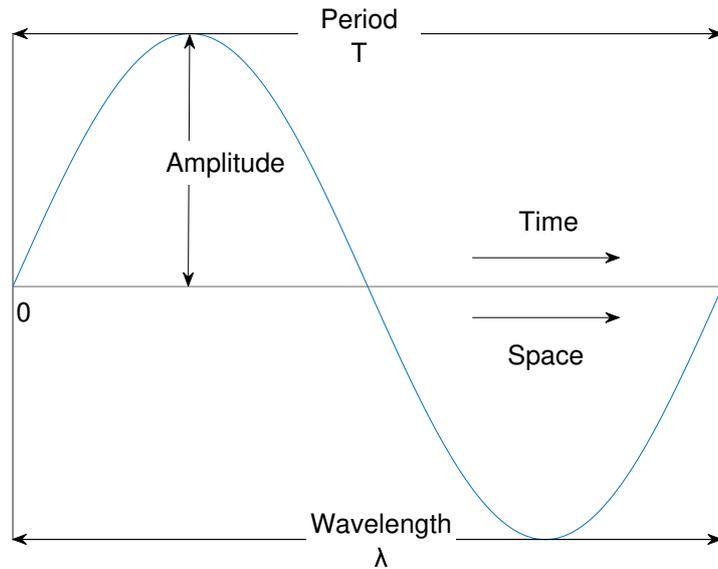


Figure 2.1: The interconnected relationships between space and time, as shown on a sinusoidal waveform.

This highly interconnected relationship between the spacial and temporal dimensions of EM waves indicates that, in certain situations, it may be important to consider how the impact of a particular EM phenomena, often heavily frequency dependent itself, will vary over a interrelated spatial

distance. This is particularly pertinent for cases of radiated EM phenomena, which this thesis details heavily.

EMI itself is a concept which requires three critical components to arise. The first, a source, may emit electrical, magnetic or electromagnetic fields either inadvertently or purposefully. These typically arise from man-made devices, such as radios, microwave ovens and power converters, though natural sources such as lightning can be severely detrimental to the operation of many sensitive devices. The second critical component for EMI to arise is a coupling mechanism, or rather, how emissions from a source can travel and interfere with the operation of the third critical component - the victim. Whilst there are many specific means for a source to couple to a victim, they can all be broken into four fundamental mechanisms, described in Table 2.1.

Figure 2.2 provides a diagrammatic model of the four fundamental coupling paths (or modes) between a source and a victim. Whilst this provides a clear overview of the main forms of coupling, there are many possible combinations of coupling which arise in the real world. For example, the conduction mode is shown as a single conductor but in reality this may consist of multiple power and data connections. Conduction may also be split into common-mode, where a shared interference is induced relative to a system's ground, and differential-mode, where different interference is experienced by different conductors relative to a system's ground. Interference modes may also occur which are dependent on other interference sources, such as a radiated emissions derived from current travelling through a conducted coupling path. In real-world testing scenarios, the measurement of EMI sources may be limited to measuring, at most, the four fundamental coupling modes as the number of all possible coupling modes is too great to fully capture.

As this stage, it is worth noting the influence of the Maxwell's equations on the coupling mechanisms. Maxwell's equations serve as the foundation

Table 2.1: The four EM coupling mechanisms.

Coupling Mechanism	Description
Radiated	Most commonly encountered when source and victim are spatial and/or electrically separated by large distances. EM waves propagate from the source, acting as an antenna of sorts, through space before coupling onto the victim. Analogous to a radio transmitter and receiver.
Conducted	Occurs when an electrical connection between source and victim exists and allows EMI to be electrically conducted from source to victim, such as on power supply connections or data buses.
Capacitive	Changing electric fields between the conductors in the source and victim induce an unwanted change in voltage or current in the victim through displacement currents.
Inductive \ Magnetic	Changing currents in the source creates a changing magnetic field, as per Ampere's Law. This, in turn, induces changes in voltage on the victim through electromagnetic induction, otherwise known as Faraday's Law.

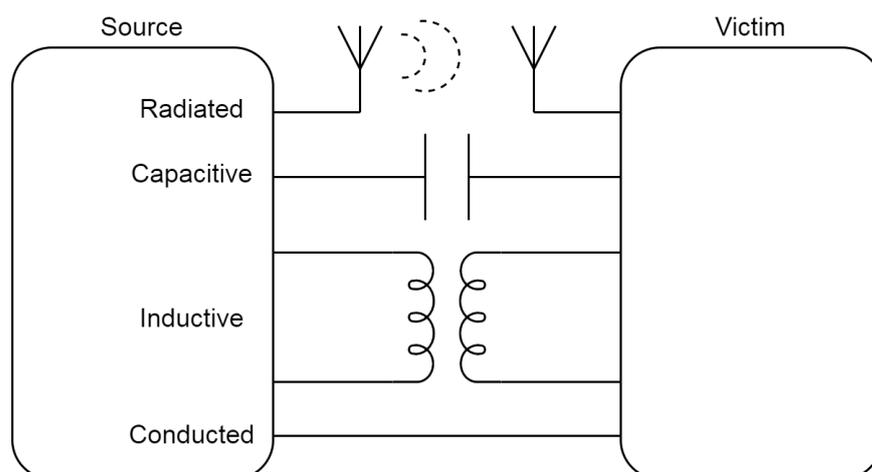


Figure 2.2: The four primary means of EM coupling between a source and victim.

for classical electromagnetism and are fundamental to the accurate and analytical understanding of EM behaviour. They are given as follows by Equations (2.3)–(2.6)[15] - the symbol representations are given by Table 2.2.

$$\text{Gauss' Law: } \oint_{\partial\Omega} \vec{E} \cdot d\vec{A} = \frac{Q_{enc}}{\epsilon_0} \quad (2.3)$$

$$\text{Gauss' Law for Magnetism: } \oint_{\partial\Omega} \vec{B} \cdot d\vec{A} = 0 \quad (2.4)$$

$$\text{Maxwell-Faraday Equation: } \oint_{\partial\Sigma} \vec{E} \cdot d\vec{s} = -\frac{d\phi_B}{dt} \quad (2.5)$$

$$\text{Ampère's Law: } \oint_{\partial\Sigma} \vec{B} \cdot d\vec{s} = \mu_0\epsilon_0 \frac{d\phi_E}{dt} + \mu_0 i_{enc} \quad (2.6)$$

Table 2.2: List of symbols for the integral forms of Maxwell's Equations

Symbol	Description
Ω	Volume with closed boundary surface $\partial\Omega$
Σ	Surface with closed boundary curve $\partial\Sigma$
\oint	Line integral around some boundary
\vec{E}	Vector electric field
\vec{B}	Vector magnetic field
\vec{A}	Vector area
\vec{s}	Vector surface
q_{enc}	Enclosed charge
ϵ_0	Permittivity of free space
ϕ_B	Enclosed magnetic field
ϕ_E	Enclosed electric field
t	Time
μ_0	Permeability of free space
i_{enc}	Enclosed current

Of the equations given by (2.3)–(2.6), the Maxwell-Faraday Equation and Ampère's Law are of most interest when determining the effect of EMI on a system. Consider Ampère's Law. If there is no changing electric field then $\frac{d\phi_E}{dt} = 0$. Therefore Equation (2.6) becomes Equation (2.7), known as the classical form of Ampère's Law.

$$\oint_{\partial\Sigma} \vec{B} \cdot d\vec{s} = \mu_0 i_{enc}. \quad (2.7)$$

It is self evident from the equation that a change in the enclosed free current (the current through the wire) i_{enc} results in direct change in the magnetic field vector, \vec{B} . Any changes arising in \vec{B} from a non-zero changing electric field are the contributions of the so-called displacement current, $\frac{\partial \vec{D}}{\partial t}$, where \vec{D} is the electric displacement field. As described in [16], the existence of displacement current is apparent when considering the case of an alternating current (AC) applied across a capacitor, as shown in Figure 2.3 where S and S' are two surfaces bound to the same closed path.

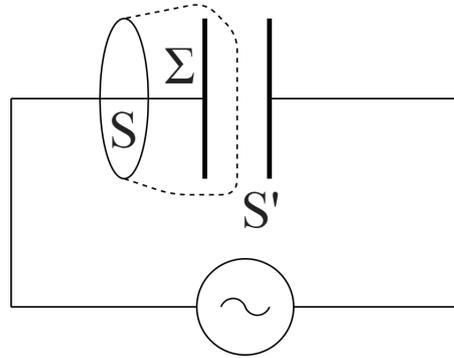


Figure 2.3: A capacitor with an AC source across its plates [16].

As the magnetic field through the two surfaces must be equal, it follows that the magnetic field between the plates should be equal to the field outside of the plates. Therefore, any displacement current between the plates must equal the current flowing through the wires - current is therefore not only the movement of electric charges, but the time-variance of an electric field. It is this realisation that forms the foundation for the capacitive coupling mechanism to function.

This same mechanism from Ampère's Law also describes half of the process of inductive coupling. In the mechanism time-varying electric currents produce varying magnetic fields, but instead of the fields coupling directly onto the victim as in capacitive coupling, the time-varying magnetic fields

induce voltages on the victim through electromagnetic coupling, described by Faraday's Law. This is presented in Equation (2.8).

$$\mathcal{E} = -\frac{d\phi_B}{dt}. \quad (2.8)$$

Where \mathcal{E} is the electromotive force (EMF), in Volts. Maxwell expanded on this by stating that a time-varying magnetic field always accompanies a varying (in both a spatial and temporal sense) electric field, as in Equation (2.5).

2.1 Sources of EMI

Whilst at first glance, it may seem that all sources of EMI are man-made, this is far from the truth. Indeed, the earliest form of ensuring EMC was perhaps the invention of lightning conductors in the 18th century to protect buildings from one of Earth's most powerful EM phenomena [17]. The emissions from such natural phenomena are outside the control of a design engineer and provide a clear distinction between design philosophies for ensuring EMC. In systems where EMI is affecting a victim system, it may suite a designer engineer to produce a *preventative* solution whereby the source is redesigned in some manner to reduce the emissions from itself. However, the *corrective* design philosophy of EMC seeks to better design the victim to protect it from, or correct, the effect of EMI. Of course, when natural sources of EMI are considered, it is the latter aesthetic which must be adopted - no person may redesign a lightning bolt! Table 2.3 details a few of the most common sources of EMI, both natural and not.

Table 2.3 indicates many sources of EMI which are liable to affect the data network of a car. From industrial tools such as welders, to mobile phones being used by passengers the EM environment of a car is a quite polluted. The following subsections discuss contemporary means of preventing and

Table 2.3: Sources of EMI

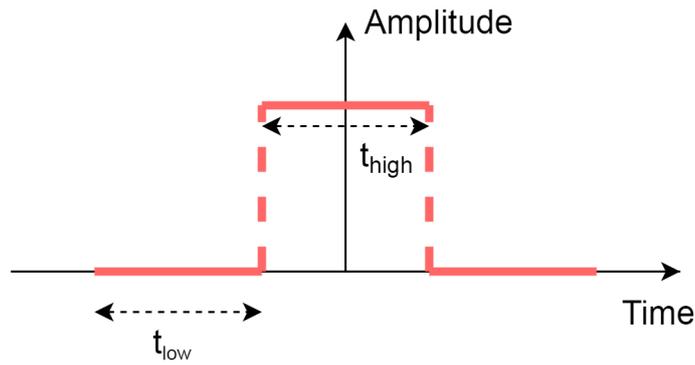
Man-made Sources of EMI			
Industrial	Radio & Communications	Energy	Consumer
Welders Lasers Compressors RF Heaters	RADAR Mobile Phones Microwave Links Spark Gap Transmitters	Motors Power Converters Power Lines Generators Substations Heaters	Air Conditioning Lighting Fluorescent & LED Lighting Human Generated Static
Natural Sources of EMI			
Lightning	Interstellar	Solar	Precipitation Static

correcting the effects of EMI in both wired and wireless networks.

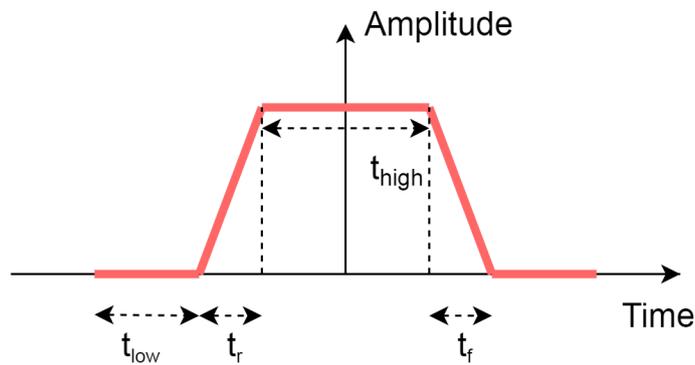
2.2 EMI Mitigation for Wired Networks

Wired data networks are typically susceptible to a mixture of all four coupling mechanisms, but can also be a potential source of EMI if they are not properly designed. This is partially due to the inherent characteristics of the digital domain - it aims to operate at two discrete voltage levels - only ever settling at high and low states, of duration t_{high} and t_{low} - with the absolute minimum time between the two data states, as shown in Figure 2.4a. Here the amplitude of the data is at two discrete levels with an instantaneous transition between states, represented with vertical, dashed lines. However, a real world system is better approximated by the waveform in Figure 2.4b, with non-zero transition periods labelled as t_r and t_f . Any sampling of this data waveform which occurs during the rise and fall transitions (t_r and t_f) will fall outside the receiver's ability to assign a digital state to the analogue voltage - therefore resulting in an unspecified system response. These unspecified voltages during t_r and t_f are known as undetermined states.

Generally speaking, however, the ratio of either t_r or t_f (more often than



(a) An ideal digital waveform, with instantaneous transitions between high and low states



(b) A more realistic model of a digital waveform with non-instantaneous transitions

Figure 2.4: Models of digital waveforms

not, almost nearly identical in duration) to either t_{high} or t_{low} (which again are nearly always equal) is deliberately designed to be low. Figure 2.4 graphically illustrates that for a given t_{high} and t_{low} period, the longer the duration of t_r and t_f , the longer the overall duration of the data bit, as given by (2.9), assuming half of the duration of t_r and t_f are associated with the duration of the particular bit t_{bit} .

$$t_{bit} = t_{high} + \frac{t_r + t_f}{2} \quad (2.9)$$

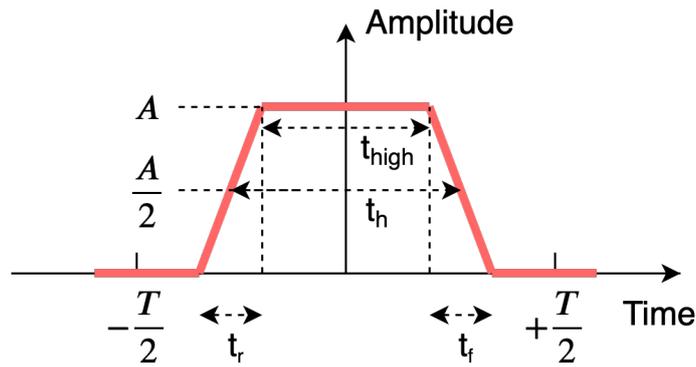
As the maximum data throughput, or data rate, of a wired system can be calculated by Equation 2.10, it then follows that the maximum data (or bit) rate (R_b) of a wired system is inversely affected by t_r and t_f .

$$R_b = \frac{1}{t_{bit}} \quad (2.10)$$

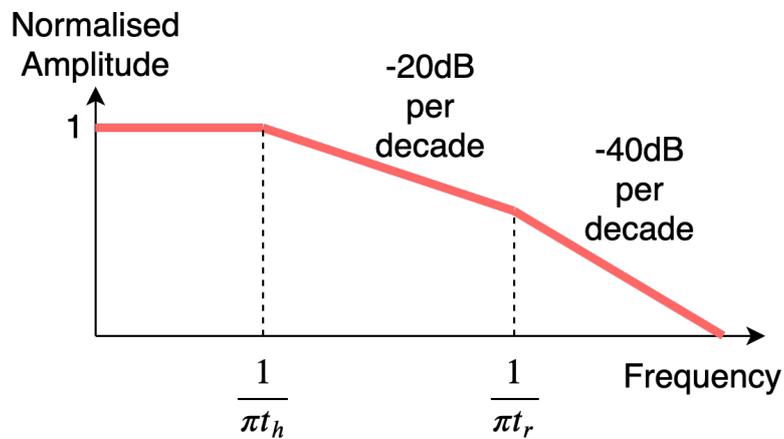
In other terms, a longer duration of t_{bit} results in a smaller bit rate R_b . As systems may have a lower limit of t_{high} and t_{low} to ensure sufficient time to sample the bit in a known-determinable state, then the duration of t_r and t_f become the only variables which may be decreased to realise a higher data rate. Another name for the period t_r and t_f is the slope length.

Whilst, from a data rate perspective, it may seem entirely beneficial to strive towards reducing the duration of the periods t_r and t_f (in other words, to increase the slope $\frac{dV}{dt}$ for the same given amplitude), in reality a shortening of t_r and t_f is actually detrimental from an EMC point of view. The fast transitions between states in a wired digital communications system can create a great deal of high-frequency emissions in the EM spectrum, which inevitably increases the potential for EMI. Figure 2.5 shows a simplified means of converting a time-domain waveform representation into an emission spectrum in the frequency domain, using a three-straight-line approximation from [18]. Figure 2.5a describes, in greater detail, the trapezoidal approximation shown in Figure 2.4b for a data bit with the addition of amplitude A , mid-point amplitude $\frac{A}{2}$, half-amplitude period t_h and the signal period T .

The result of the three straight line approximate for the spectral content is shown graphically in Figure 2.5b. At frequencies above $\frac{1}{\pi t_h}$ the asymptotic spectral contribution is approximately inversely proportional to the frequency (-20dB per decade) - this contribution becomes inversely proportional to the square of the frequency (-40dB per decade) at frequencies above $\frac{1}{\pi t_r}$. This demonstrates that reducing the transition periods in around a data bit increases the spectral content of the signal, increasing the potential for EMI to arise.



(a) A simplified model of a data bit with non-instantaneous transitions, where A is the amplitude in Volts and the half-amplitude period t_h .



(b) Three straight line approximation model's approximation of spectral content.

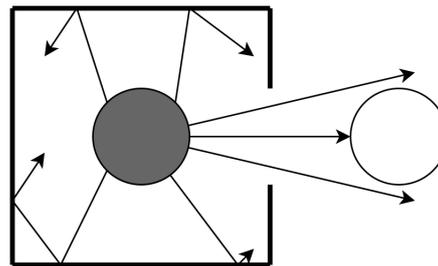
Figure 2.5: Conversion from time domain model of a data waveform to frequency domain model of the spectral contents [18].

Therefore, whilst it is possible for data throughput to be increased by increase the slope rate, it can be demonstrated that this also increases the potential for the network to act as a source of EMI. As a result, a common preventative technique utilised in the design of networks is known as slope control. By deliberately decreasing the $\frac{dV}{dt}$ of the slope, the amount of emissions decreases. At present, this translates into design procedures advocating a balance between decreasing the slew rate of the slope as much as possible to reduce EMI, whilst also maintaining a slew rate high enough to ensure data throughput requirements are met - indeed some examples of

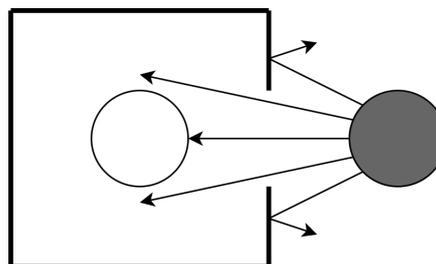
slew control seem to attempt to slow transitions by such a degree so as to eliminate EMI harmonics altogether [19], which may come at a significant reduction in maximum data throughput.

The technique of increasing slope periods is also gaining traction in power switching applications to improve both the quality of power delivered to a system and to reduce the potential for EMI [20, 21, 22], asserting the benefits of lower slope $\frac{dV}{dt}$, albeit in a power switching application.

A common corrective method for the protection of a victim of EMI is through the use of EM screens or shields. This can be described as the use of a protective layer, shielding the victim conductor from the emissions from an EM source. The effect of shielding a conductor is reciprocal, in the sense that it can also be used to contain any EM radiation from the conductor too, demonstrated with the concept of EM leakage in Figure 2.6.



(a) Shielding to prevent EMI leakage.



(b) Shielding to provide EM immunity

Figure 2.6: The uses of shielding [23].

In an ideal case, such shields would consist of thick, solid metal completely enclosing the conductor. Thick, in this instance, is generally taken to

describe a dimension which is greater than ten skin depths - skin depth being the depth below a surface where the AC current density J falls to $\frac{1}{e}$ [24], as given by (2.11).

$$\delta = \sqrt{\frac{2\rho}{\omega\mu}} \sqrt{\sqrt{1 + (\rho\mu\epsilon)^2} + \rho\omega\epsilon} \quad (2.11)$$

Where ρ is the resistivity of the conductor, ω is the angular frequency of the current, μ is the permeability of the conductor and ϵ is the permittivity of the conductor.

Shields with a thickness of multiple skin depths may provide a shielding effectiveness (SE) above 200dB [23]. Shielding effectiveness is a measure of the effectiveness of a shield to attenuated EM radiation from or on a conductor as shown in Figure 2.6. Considering the case with of a single aperture in a shield, the SE may be obtained from (2.12).

$$SE = 20 \log_{10} \left(\frac{\lambda}{2d} \right) \quad (2.12)$$

where d is the maximum diagonal dimension of the aperture. As shown in [23], this means that the maximum permissible aperture d can be calculated for given frequency, as in (2.13).

$$d = \left(\frac{c}{2f \cdot 10^{\frac{SE}{20}}} \right) \quad (2.13)$$

An illustrative example of the shielding effectiveness of various aperture diameters can be seen Figure 2.7, which plots the maximum allowable aperture diameter for a SE of 70 dB between 1 MHz and 1 GHz. Note that as the frequency increases the change in the maximum permissible aperture decreases at a logarithmically lower rate. This can be advantage in situations when using shields to protect equipment from EMI at lower frequencies where a small change in aperture results in large increase in

the maximum frequency for a given SE. However this relationship may also prove a disadvantage at higher frequencies, where the variation of small apertures resulting from, for example, manufacturing imperfections results in a large decrease of the maximum frequency for a given SE.

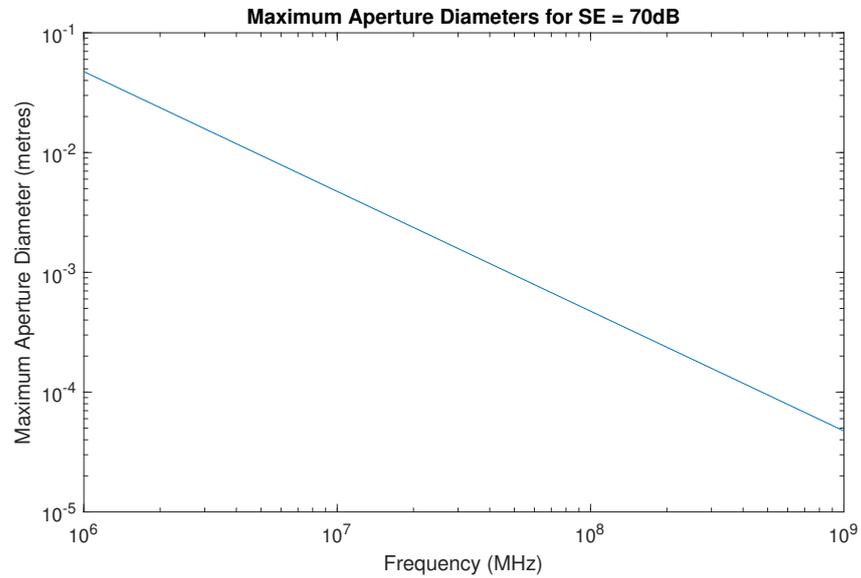


Figure 2.7: Maximum aperture diameters for varying frequencies.

Whilst this applies only to the case of a solid shield with a single aperture, it has been shown by [25] that the addition of multiple apertures only serves to decrease the SE. In the context of this thesis, considering that real-world network cables require the ability to flex and move in order to be installed properly into cars, the shielding on them is woven into a braid which inevitably contains many small apertures, as shown in Figure 2.8.

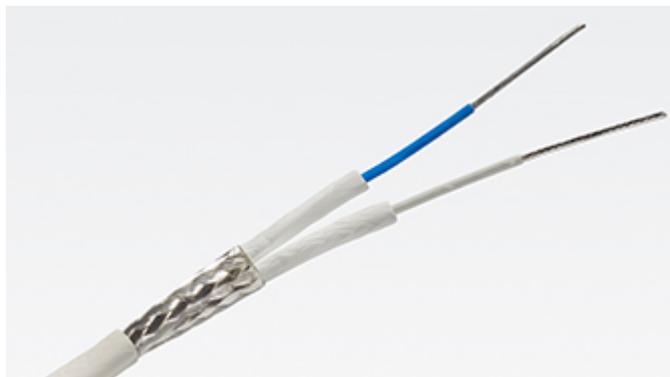


Figure 2.8: A shielded, multi-conductor cable for use in vehicles [26].

These apertures, whilst small and consistent in dimensions in a straight cable will distort and vary in size as soon as any flex occurs on the cable. Importantly from an EMI perspective, this will result in wider apertures forming on the outer, extended radius of the cable. As demonstrated, this reduces the upper limit frequency for a given SE, potentially exposing the shielding the inner conductors to harmful EMI. Whilst it is possible to utilise braids with a higher density weave and a higher density of metallic fibres to reduce the change in aperture dimension when flexed, this will also increase the flexibility, cost, and mass of the cable - less than ideal for cost-conscious manufactures and users of cars.

2.3 EMI Mitigation for Wireless Networks

As one of the aims of this thesis is to produce a wireless data link for use in cars, it is also worth examining some of the most prevalent techniques for reducing the effect of EMI on wireless networks. As it is the very nature of a wireless network to have no physical electrical connection in the transmission channel (considering only the channel between transmit and receive antenna), out of the four fundamental coupling mechanisms in Table 2.1 it is only radiated modes which affect wireless networks. Despite this single mode of interference, the use of multiple modulation techniques allows for a variety of EMI mitigation strategies to exist, which are described and analysed in this section.

All wireless systems can be thought of as a system of five blocks, as shown in Figure 2.9 [23]. These can be divided into two categories based upon how they convert, or modulate, a data signal with a higher frequency carrier wave. These are analogue and digital modulation (AM and DM, respectively). AM works by modulating a continuous data signal (Figure 2.10, whereas DM modulates an analogue carrier with a discrete

signal. AM schemes tend to be more susceptible to relatively small amounts of interference, but are more readily designed. On the other hand, DM schemes are more resilient to the same amounts of interference due to the ability to arbitrate between only one of two digital states, though DM may fail completely at a certain level of interference where some limited AM performance is maintained. This increased resilience to interference (and therefore EMI) comes at the cost of greater system complexity, as shown in Figure 2.11. Any analogue data source must be first converted to digital information through a source encoder, before it is processed by a channel encoder which permits the receiver to detect bit errors. This digitally encoded information is then processed by a digital modulator before it is transmitted through the channel and the process reversed to reconstruct data as close as possible to the original source.

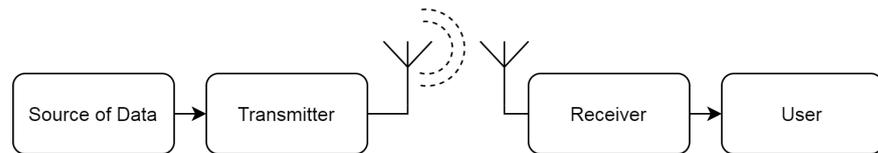


Figure 2.9: Basic blocks of every wireless network [23].

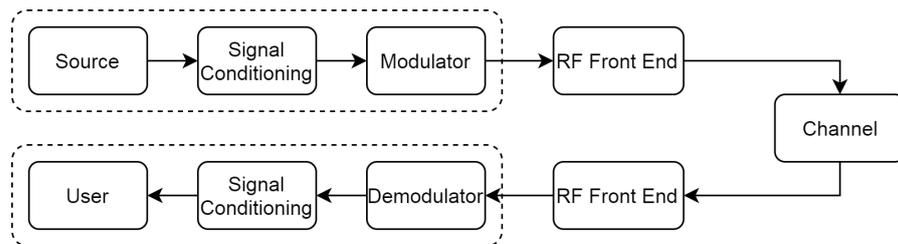


Figure 2.10: Analogue modulated network, based on [23].

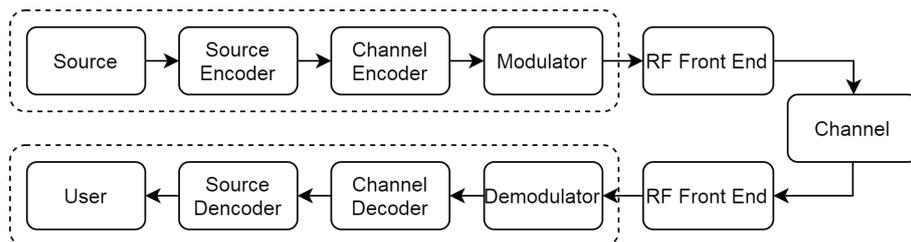


Figure 2.11: Digital modulated network, based on [23].

Examples of analogue modulation schemes include Amplitude Modulation (AM) and Frequency Modulation (FM). Amplitude modulation consists of varying the amplitude of the higher frequency carrier wave. Consider a simple signal $\cos \omega_m t$ which is modulated with a carrier wave $\cos \omega_c t$ as in (2.14).

$$V(t) = (1 + m \cos \omega_m t) \cos \omega_c t \quad (2.14)$$

Where $V(t)$ is the modulated signal, and m is the modulation index - the ratio of the limits of the modulated to unmodulated carrier wave. Expanding the product of (2.14), as per [27] yields (2.15).

$$V(t) = \cos \omega_c t + \frac{1}{2}m \cos (\omega_c + \omega_m)t + \frac{1}{2}m \cos (\omega_c - \omega_m)t \quad (2.15)$$

From this, it can be deduced that the AM waveform consists of three peaks in the frequency domain. This modulation is illustrated in both time and frequency domains by Figure 2.12 and Figure 2.13, respectively.

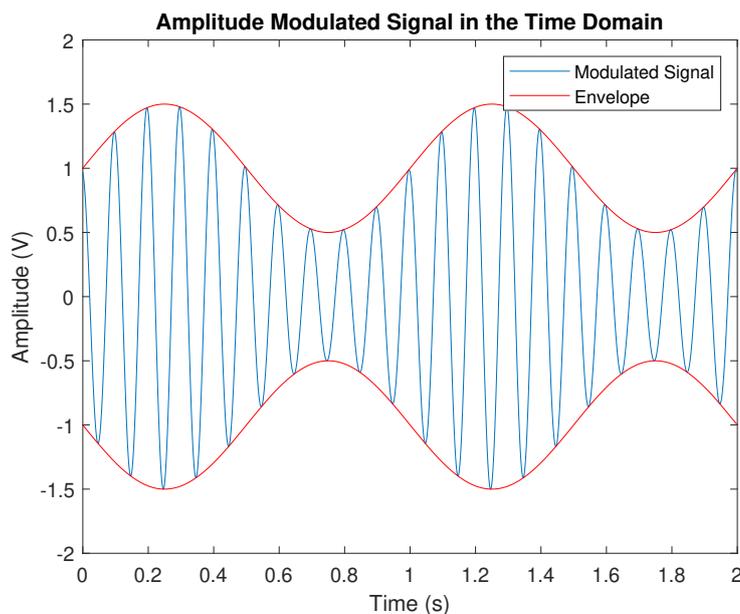


Figure 2.12: AM in the time domain, for 10 Hz carrier and 1 Hz data.

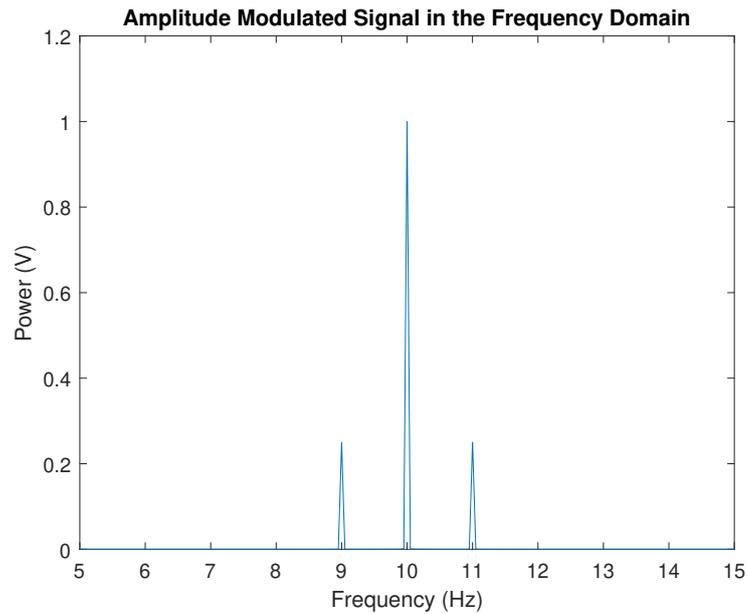


Figure 2.13: AM in the frequency domain, for 10 Hz carrier and 1 Hz data.

In Figure 2.12 the envelope of the transmitted signal (in blue) represents the data from the user, drawn in red. It can be seen that in this case a carrier frequency of 10 Hz is modulated with a frequency of 1 Hz. This is confirmed by the spectrum in Figure 2.13, with a large carrier component at 10 Hz, and two sidebands at ± 1 Hz. Whilst AM is a relatively simple modulation method it is, as mentioned before, highly susceptible to EMI as there is no way for any AM demodulator to arbitrate between noise added to the waveform during transmission in the channel, and the underlying signal. Indeed, an AM modulator detects EMI and the underlying signal in equal proportion. Furthermore, the frequency spectrum of an AM signal shows that the vast majority of the transmitted waveform's power is in the carrier - being highly inefficient and requiring high power for reliable data transfer it is an unsuitable candidate for use in a high bandwidth, high reliability system - such as that required for a car.

Another form of Analogue Modulation is Frequency Modulation (FM). In this case, the frequency of the carrier is modulated by the continuous data signal, as in (2.18).

$$V(t) = \cos(\omega_c t + k \int f(t) dt) \quad (2.16)$$

Where $k \int f(t) dt$ is the frequency deviation. Figure 2.14 shows the time domain representation of a FM scheme with a 10 Hz carrier, a 1 Hz data waveform and a frequency deviation of 5 Hz, with its frequency domain representation in Figure 2.15.

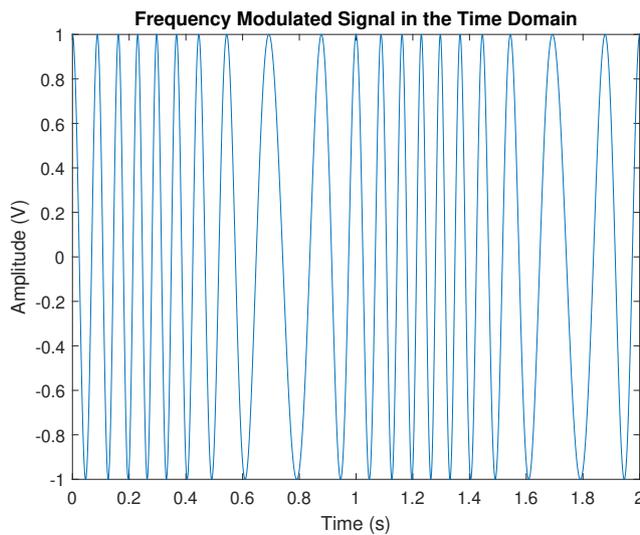


Figure 2.14: FM in the time domain, for 50 Hz carrier, 1 Hz data and 5 Hz of frequency deviation.

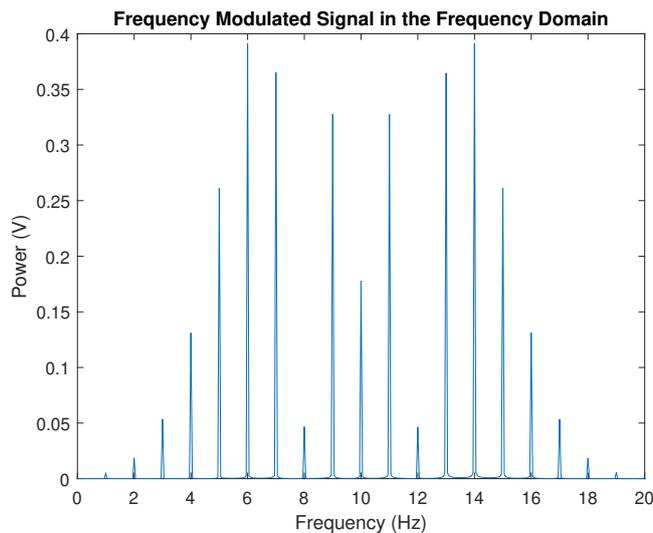


Figure 2.15: FM in the frequency domain, for 50 Hz carrier, 1 Hz data and 5 Hz of frequency deviation.

Note that there are many more sidebands in the FM spectrum, separated by the frequency of the transmitted data f_m of 1 Hz, with a total spread much further than the AM case. Whilst not ideal for minimising the use of the band, it can be clearly seen that only a very small portion of the total power is in the carrier component, f_c , at 10 Hz, resulting in a higher power efficiency as nearly all transmitted power is derived from the data source and not the carrier - which contains no information at all. FM also provides a higher resilience to interference when compared to AM, indeed FM demodulators contain an amplitude limiting function before detection [27], so amplitude variations caused by EMI do not manifest in the demodulated signal.

As described previously, the other primary form of modulation is digital modulation (DM). Such modulation schemes are generally far better suited to the transfer of data in the form of a bit stream, given that there is no need to convert between the discrete digital domain and the continuous analogue domain. Indeed, multiple DM methods intrinsically perform symbol encoding of an incoming bit stream, such as in the case of Quadrature Phase Shift Keying (QPSK), which permits a higher data rate when compared to simple binary forms of modulation, which would be readily achievable through analogue modulation methods of on or off.

QPSK works by dividing an incoming digital bit stream (d_{in}) into two distinct data streams named in-phase (d_I) and quadrature (d_q) which can be described as in (2.17).

$$\begin{aligned} d_{in}(t) &= d_0, d_1, d_2, d_3, d_4, d_5, \dots \\ d_I(t) &= d_0, d_2, d_4, \dots \\ d_q(t) &= d_1, d_3, d_5, \dots \end{aligned} \tag{2.17}$$

Now d_I and d_q are amplitude modulated onto the cosine and sine functions of a carrier wave as in (2.18).

$$V(t) = \frac{1}{\sqrt{2}}d_I(t) \cos \left[2\pi f_0 t + \frac{\pi}{4} \right] + \frac{1}{\sqrt{2}}d_Q(t) \sin \left[2\pi f_0 t + \frac{\pi}{4} \right] \quad (2.18)$$

This process is represented diagrammatically in Figure 2.16. Notice that the modulation and subsequent transmission of d_I and d_q occur simultaneously. This permits a data rate twice that of a binary modulation scheme, without an increase in the occupied channel bandwidth [28].

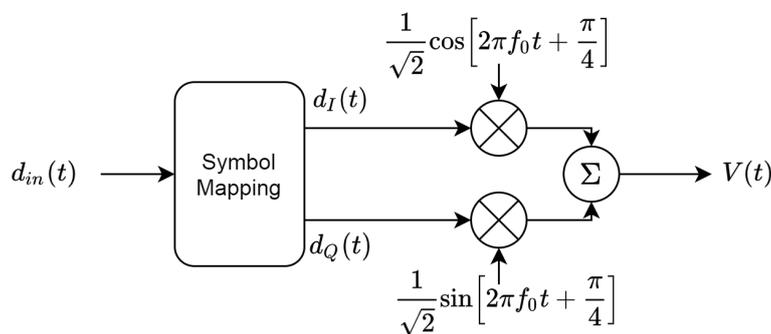


Figure 2.16: QPSK modulation

Whilst it is possible to visualise QPSK modulation in the time and frequency domains, it is more useful to inspect the performance of QPSK systems through signal constellations - a means of mapping the phase shift and amplitude of the signal on the complex plane as illustrated in Figure 2.17.

However, QPSK is not without its disadvantages. When considering the demodulation of the signal, a large enough disambiguity in the phase at the receiver can manifest in complete loss of the data stream. The ability of any QPSK receiver to determine the correct frequency and phase of the received signal generally relies upon the response rate of the phase locked loops PLL used to generate, for example, the demodulation carrier [29]. Whilst far more complex than the analogue modulation examples, the use

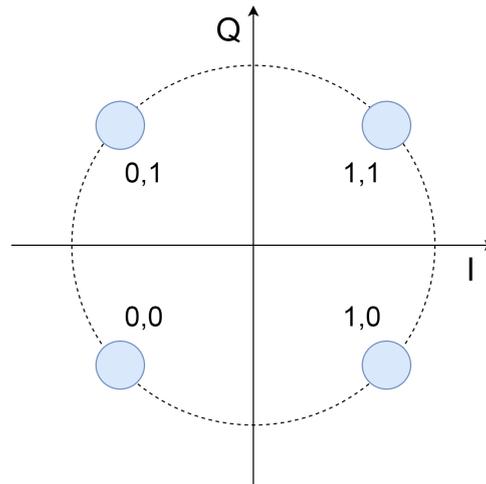


Figure 2.17: QPSK bit encoding on the complex plane.

of QPSK (and in general, all digital modulation schemes) permits a higher data rate as well as less potential for interference from EMI due to the purely phase modulated structure of the signal being transmitted through the channel. Furthermore, with the advent of high volumes of processing power in ever decreasing physical size and cost, the technical complexities of creating highly reliable QPSK networks are continuously being reduced. For this reason, such modulation schemes may be considered ideal for adoption in intra-vehicular networks.

As mentioned, the predominant means of preventative wireless EMC design is to utilise appropriate modulation schemes to inhibit the effect of EMI summing with the amplitude of a transmitted signal through the channel. However, a technique which does not rely upon a modulation scheme alone is known as Frequency Hopping Spread Spectrum (FHSS). [30] considers that for a data signal given by (2.19), the resulting FHSS signal is given by (2.20).

$$d(t) = Aa(t) \cos [2\pi f_c t + \theta(t) + \phi] \quad (2.19)$$

$$V(t) = Aa(t) \cos [2\pi f(t) + \theta(t) + \phi] \quad (2.20)$$

Where $f(t)$ describes the carrier frequency as some function of time, which is known by the transmitter and receiver. Real world communication protocols which use FHSS, such as IEEE 802.11, Bluetooth, and Zigbee ([31, 32, 33], respectively), all use pseudo-random functions to determine the carrier frequency at a certain time, t . Therefore, data is transported across multiple carrier frequencies rather than sitting on a single frequency as shown in Figure 2.18.

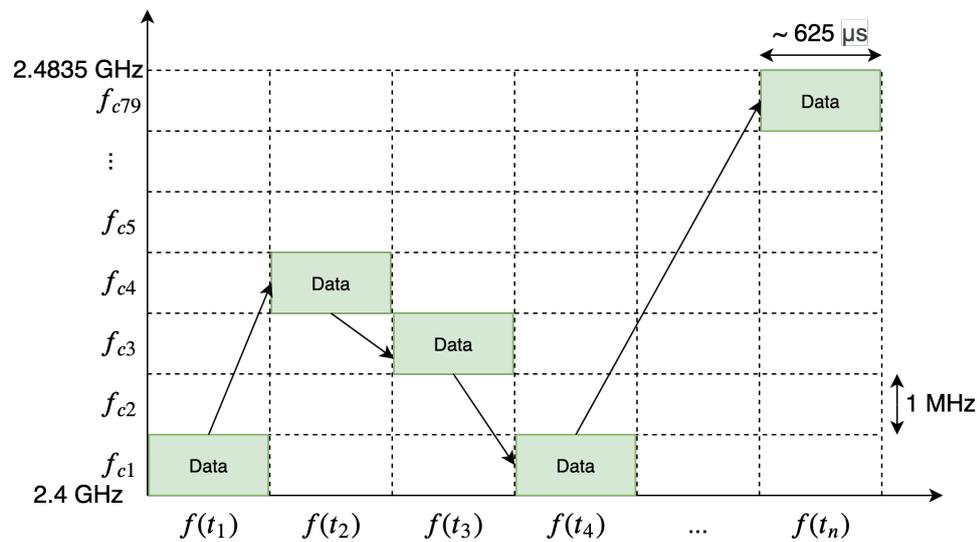


Figure 2.18: Frequency/time map of FHSS under Bluetooth standards. Each carrier is separated by 1 MHz, accommodating 79 channels in the band. Hops between f_c occur approximately 1600 times per second [34]

This can be extremely useful in communication systems which may encounter a wide variety of users operating spatially close, as the function determining the f_c can be used to spread out data transmission across the spectrum such that no two users continuously interfere with each another. Whilst FHSS has clear advantages for preventing interference on a particular transmission band (from other band users and EMI sources simultaneously), it does not offer any fundamental protection from a noisy channel, which is where the choice of data modulation becomes vital once more.

2.4 Chapter Conclusion

This chapter has given a basic background to the concepts and mechanisms which cause EMI and allow EMI to couple from sources to victims. As this thesis focuses on wired networks with a view towards changing towards wireless networks, methodologies for both preventing EM emissions and protecting potential victims on both wired and wireless networks was described. In the case of wired networks, existing EMC design methodologies for wired networks may be effective against preventing the coupling of EMI, but comes at the cost of either data rate or additional shielding (and therefore adding cost and mass to cars). The inherent advantages of advanced DM for wireless systems, such as QPSK, coupled with a reduction in cost and physical size of computational power, means that pursuing these more complex modulation schemes may provide a convenient and reliable means of creating future intra-vehicular data networks.

Chapter 3

Contemporary Intra-Vehicular Communications

The car manufacturer Volkswagen introduced the first car with a fully computer controlled module for the use of improving the reliability, efficiency, and easy of use of their engines. This Engine Control Unit (ECU), introduced in 1968 was to mark a fundamental shift in the way humans were to interact with vehicles [35]. Since then, the number of electronic modules, vital to the operation of cars has increased greatly since their introduction by car manufacturers in the 1980s [36, 37, 38, 39], and as processing power, over time, has become cheaper and more compact, so too has the number of non-safety-critical embedded devices - often used in personal entertainment systems or environmental control.

3.1 Digital Modules and Networks in Cars

In order to better understand the roles digital units perform in cars, Table 3.1 provides a brief description of a selected number of the most prevalent electronic devices found within cars over the past 50 years, along with the date of their first introduction on commercially available cars.

Table 3.1: Notable digital modules within cars.

Device	Date	Function	Description
Engine Control Unit	1968	Vital	Set air/fuel mixture, ignition timing and idle speed of an engine by comparing sensor readings to achieve optimum engine performance.
Auto-Linear Skid Control	1969	Safety	Sensors read the wheel speed - if about to lock (resulting in a skid), the breaks were momentarily released.
Trip Computer	1978	Vital	Kept track of metrics such as mileage, fuel consumption. Also alerted users to maintenance requirements.
Anti-Lock Braking System	1978	Safety	Improved version of the auto-linear skid control, with four independent control channels.
Emission Control Module	1980	Vital	Reduced pollutants, such as nitrogen oxide by redirecting exhaust gases to the engine, reducing combustion temperature and pollutants.
Touchscreen Infotainment	1986	Comfort	Allowed for buttonless control of climate and radio.
PowerTrain Electronic Control	1991	Vital	Controlled the greater drivetrain, including gearbox. Improved car reliability and traction control.
Dual Zone Climate Control	1992	Comfort	Allows for the control of multiple temperatures within the same vehicle.
Inbuilt Satnav	1995	Comfort	The GuideStar system offered increased user convenience for navigation between destinations.
Bluetooth	2003	Comfort	Allowed for the connection of user devices with the car's stereo.
Variable Force Airbag Deployment	2006	Safety	Varied the size and force of airbags, depending on the severity of the crash. Reduced likelihood of injury from airbag deployment.
Lidar	2020	Vital	Detects the presence of objects within 200m of the vehicle for obstacle detection and automated driving.
Inertial Measurement Unit	NA	Vital	Offers a more robust means of tracking a vehicle's location when compared to external navigation aids.
Neural Networking	NA	Vital	Dedicated hardware for utilising the decision-making capability of neural nets for navigating unmapped routes.

Note that many early embedded units are vitally important for easing the use of complex mechanical systems, and for performing safety-critical functions, such as the ECU [40, 41, 42, 43, 44, 45, 46, 47, 48]. Over time, a greater number of modules tend to be included for entertainment and user comfort.

However, this trend is likely to change in the next decade, as the immensely volume of work on electric-driven and autonomous vehicles drives a change on the basic functionality of new cars in the marketplace. In an engineering renaissance (of sorts), newly incorporated computing modules will be required to ensure the safe operation of high-power systems and safety-critical driving devices. This too, is tentatively summarised in Table 3.1, with descriptions of digital modules which are either already in use or are anticipated to be in cars either commercially available within the next decade.

Such a high number of devices cannot function in a collaborative manner without the presence of a data network ensuring the cohesive management of all the data collected and processed by the systems within a car. As the number of new embedded control modules in cars has increased, so has the number of intra-vehicular networks connecting them. Initially, these networks were based on pre-existing communication networks standards for non-automotive applications, but as an increasing reliance on digital control modules arose, purposely designed networking standards followed. Until this point, all EMC design methodologies in this thesis have been based on hardware - or the physical layer, as set out by the Open Systems Interconnection OSI model. The following subsection briefly outlines the (OSI) model in order to allow for abstraction between hardware level and firmware level (which are often unique to the protocol) EMC design techniques.

3.2 The OSI Model

The OSI model, created by the International Organisation for Standardisation (ISO), describes the function of the hardware and software components which may make up any communications system, by abstraction of the layers involved in transferring data from one user to another through a digital communications network. In total, seven layers are defined by the model [49], which is presented in Table 3.2.

Table 3.2: The OSI model [49].

Layer		Description
Host	7	Application High level application control of data network
	6	Presentation Exchange of data between networking and application
	5	Session Manages communication sessions between devices
	4	Transport Multiplexing and segmentation of data networks
Media	3	Network Unit addressing and other networking control
	2	Data Link Data framing and checksum methodologies
	1	Physical The raw bitstream - analogue level bit representation

As mentioned, all the EMC design techniques discussed thusfar have focused on methodologies affecting only the physical layer (layer 1) of a system. Generally, many such techniques may be applied to a given communications protocol, and many wired communication protocols may share similar physical layer design techniques. However, the biggest difference between the standardised protocols mainly exists between the data link and networking layers - at least from an EMC standpoint. Therefore, a selection of the most influential wired communication network protocols, used in an automotive environment, are discussed below. In all these cases, there exists no EMC design features which are represented in layer 4 or

higher - as such, these layers are not covered in this work, though more information is available through the appropriate references.

3.3 Wired Automotive Communication Protocols

The following subsections provide a brief description of each of four prevalent communication protocols in the automotive sector today [50, 51, 52, 53, 54, 55], detailing typical application connections, open system interconnect model (OSI) descriptions, methods for detecting the effects of EMI, and data throughput. These protocols are, in order; transistor-transistor logic universal asynchronous receiver/transmitter (TTL UART), Recommended Standard 232 (RS232), Controller Area Network (CAN), and Flexray.

3.3.1 TTL UART Networks

Strictly speaking, TTL is not a networking protocol *per se*, rather it is one of the two most popular digital logic families (the other being complementary metal-oxide-semiconductor (CMOS)). The naming structure describes the component performing the logic arbitration (the first transistor) and the output current sourcing (the second transistor) [56] as per the example in Figure 3.1, sourced from [27].

The first stage in Figure 3.1 shows two diodes on each input A and B - the anodes of which are connected to a transistor inverter (Q_1) which provides the input logic detection. If either voltage at A or B is held low, then Q_1 is turned on. The output, Q, is then forced low by the saturated Q_4 and the switching off of the Darlington pair made by Q_3 and Q_4 . However, if one or both of A and B are low, then Q_1 is held off, forcing Q_4 off and the Darlington pair to turn on. This in turn pulls the output high.

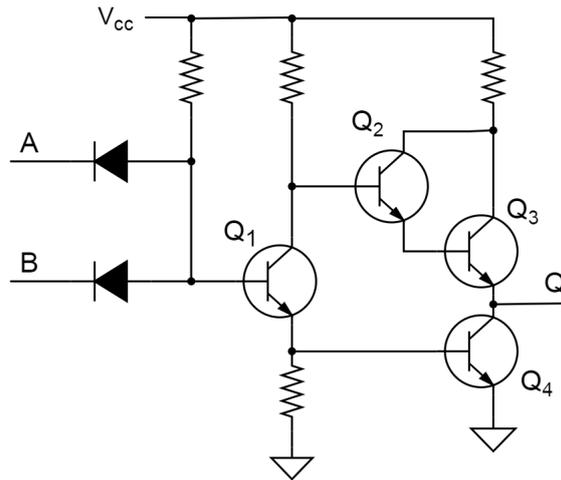


Figure 3.1: TTL NAND gate. [27]

Note that the voltage levels affecting the entire logical operation in the gate are primarily affected by the forward biasing voltage of the diodes. These fundamental, component level specifications determine the actual logic levels for a particular family of devices. The two most common variants of TTL are 5V TTL and 3.3V (or low-voltage) TTL (LVTTTL).

Table 3.3: Comparison of 5V TTL and LVTTTL operating parameters

Characteristic	Name	5V TTL	LVTTTL
Input High Voltage	V_{IH}	2V	2V
Input Low Voltage	V_{IL}	0.8V	0.8V
Output High Voltage	V_{OH}	2.4V	2.2V
Output Low Voltage	V_{OL}	0.4V	0.55V
Propagation Delay	t_{pd}	22ns	4.1ns

The specific operating voltages, and their names, are described in Table 3.3 [57]. Note that for each family, the maximum output low voltages are specified so as to be lower than the maximum input low voltage. Similarly, the minimum output high voltage is higher than the minimum input high voltage. This is to ensure that the voltage of any transmitted bit is acceptable for the input of a device from the same family. This is to

avoid transmission of bits in the non-defined operation region, which is highlighted in Figure 3.2. This difference between the limit of an output and input voltage is known as the gate noise margin. In other words, the gate noise margin is the minimum amplitude change an interference source must induce on a victim to cause incorrect data transfer.

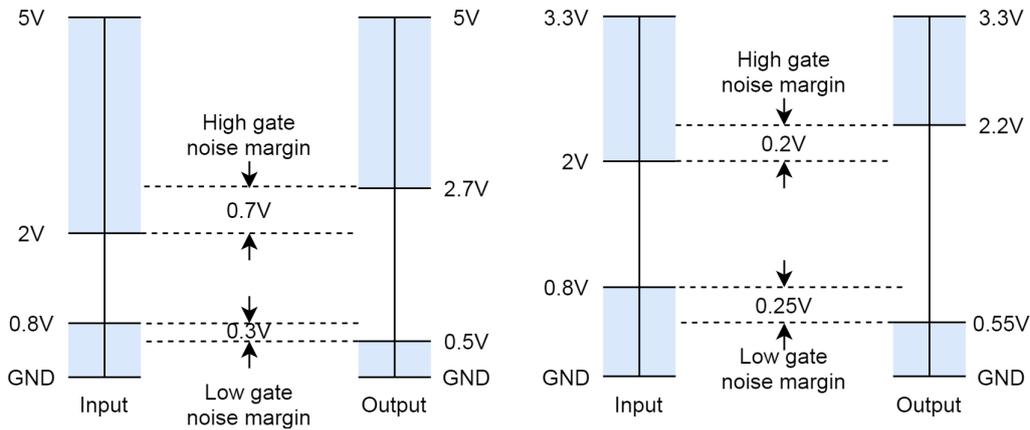


Figure 3.2: Graphical representation of TTL family levels.

This is a common technique for preventing unspecified behaviour arising from coupled EMI on every form of wired communications network - it is not unique to TTL families. Note that whilst Table 3.3 shows that LVTTTL has a faster propagation period (t_{pd}) than 5V TTL (allowing for a higher upper data throughput limit), it also has smaller gate noise margin. This means that at a physical level, it is more susceptible to induced voltage changes from EMI.

It is important to note that TTL in itself is not a communications protocol - it merely describes the operating voltages for communication. However, these voltages can be used to specify the operating conditions of a universal asynchronous receiver-transmitter (UART).

A UART sequentially transmits bytes (eight data bits) of information in sequential frames [58]. There is no clock signal accompanying UART data - hence the reference to asynchronous communication in the protocol's naming. Each serial frame, modelled as layer 2 in the OSI model, contains a

single byte of data, along with a start bit, one or two stop bits, and a parity bit. An example of a UART frame (with LVTTTL logic levels) is presented in Figure 3.3.



Figure 3.3: A single LVTTTL UART frame.

In its idle state, the UART is held in a high state, in this case 3.3 V. This is a remnant of telegraphy standards in which an idle line was held high to indicate that the networking equipment was functional [59]. A start bit, indicating the start of a frame, is then held low for one bit. The eight data bits, from least significant to most significant, is then transmitted. Next, a parity bit is sometimes transmitted. This parity bit is a means of detecting interference, through EMC or other means, on the transmitted data. It does this by performing an exclusive-or operation (XOR) on the eight data bits, and transmitting the result in the parity bit.

For example, if the sequence 01101011 is sent, the XOR of this set is 1 ($0 + 1 + 1 + 0 + 1 + 0 + 1 + 1 \pmod{2} = 1$). Under even parity, a 1 is then transmitted as the parity bit - in odd parity this would be a 0. The receiver can then perform the same XOR operation on the received data and compare it with the received parity bit. If they do not match, then the frame is discarded. Whilst this may suit networks under low levels of EMI, it may not be practical in systems which must operate in challenging environments - such as the automotive. If an even number of data bits are altered at the receiver, the the parity will still match - enabling incorrect data to be passed through. Furthermore, the parity bit itself may also be altered in the channel - having only one bit for checking the safe exchange of data may be a significant risk in high noise environments. One or two stop bits then complete the transmission.

It can be seen that for every eight bits of data sent through a UART network, up to 12 bits must be sent. This requirement to transmit more bits than just the data exists across all networks in order to facilitate efficient transmission and allow for the addressing of devices in protocols with layer 3 characteristics. It can be useful to measure this extra transmission requirement as a bit efficiency - that is, the ratio of data bits to transmitted bits in a given frame as in (3.1).

$$b_{eff} = \frac{b_{data}}{b_{frame}} \quad (3.1)$$

In the case of TTL UART, this can range from 0.8 (with no parity check and one stop bit) to 0.66 (with parity check and two stop bits).

TTL UART has no specified upper or lower data limit [58], though it can be used to transfer data between 2400 bits per second (BPS) and 3 Mbps. However, difficulties in ensuring the proper decoding of such a simple (on all three OSI layers) protocol at high data rates mean that in practice data rates exceeding 128 kbps are extremely rare.

In summary, TTL UART is a protocol which is simple to understand and implement. This makes it ideal for a quick implementation in many simple electronic systems. However, this simplicity also results in a protocol which cannot inherently address different network components, and which can only detect the influence of EMI - not correct it. Indeed, the primitive use of parity checking is easily misled by multiple bit errors. As a result, TTL UART networks are primarily used for very short distances between a pair of devices - such as in entertainment systems.

It should be noted that the use of 'TTL' in many contemporary systems is typically used to describe the operating voltage levels of a system, rather than provide a strict description of the fabrication process of the device.

3.3.2 RS232 Networks

RS232 networks are functionally identical to the TTL UART counterparts in layers 2 and 3 of the OSI model. However, as indicated by Figure 3.4 the differences between the two protocols lies in the analogue representation of bits in the physical layer. Note that the RS232 standard does not specify precise voltage levels [60], rather a range of voltages which may represent a 1 or 0. In this case, the lowest voltages are chosen, which may facilitate the highest possible data data rate (limited by the $30\text{ V}/\mu\text{s}$ slew rate specification).

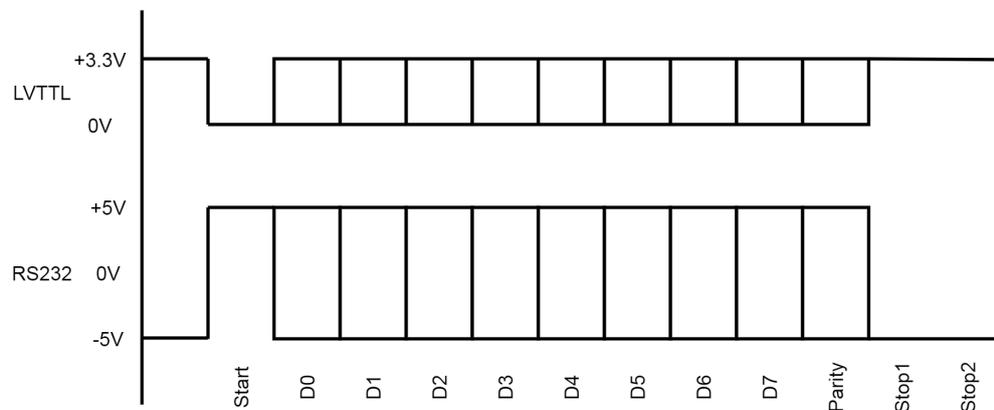


Figure 3.4: A single LVTTL UART frame.

Note that the high voltage level now represents a logic 0 and a low logic level represents a logic 1 - the inverse of the TTL UART case. Furthermore, the voltage difference between high and low is also much greater, at 10V. This increase in complete noise error margin, means a much larger induced voltage swing is required to alter the state of a single bit in an RS232 system, compared to a TTL UART - increasing its resilience to EMI. Furthermore, the RS232 standard also specifies the separation of the common ground (reference for all data voltages) and a protective ground - a shield held at the equipment's frame ground - thus providing effective shielding throughout the entire transmission channel.

As a result, RS232 is well suited for use in networks which need to

operate in noisy environments, but the layer 2 and 3 limitations of TTL UART still apply. As a result, RS232 is almost entirely obsolete in today's automotive networks, as the hardware is prohibitively large and complex for short connections, but is greatly surpassed in capability by the purpose-designed, automotive protocols which follow.

3.3.3 Controller Area Networks

The first network protocol design for use in automotive environments specifically is the Controller Area Network (CAN). Since 1983, Robert Bosch GmbH has developed the CAN specification for use in industrial and automotive environments [61]. This specifies only the data link layer within the Open Systems Interconnection (OSI) model. The physical layer, was not formally standardised until the separation of ISO standard 11898 in 1993. ISO 11898-1 [62] took the place of [61] in determining the data layer, whilst ISO 11898-2 [63] would define the physical layer for high speed CAN operation. For completeness ISO 11898-3 [64] was released in 2006, which covers the physical layer for low-speed CAN. It should be noted that despite the emergence of [63] and [64], the CAN specification itself does not contain any information regarding the physical layer. It is therefore common practice within industry to state that a CAN component is fully or partially compliant with a relevant physical layer standard, as seen in the technical document [64].

[61] and [63] describes the CAN system as a multi-master bus network. This means that multiple, addressable nodes exist within a CAN arrangement, each of which can initiate a transfer of data. This permits an exchange of data throughout the network without any centralised control[65]. The use of a multi-master bus means there is no single failure point for the network. Should the master (or transmitter/receiver) on other data networks be

damaged due to EMI or other causes, the entire network cannot function as there is no master capable of routing data. Such an event may lead to disaster in a safety-critical application, such as connecting a pressure sensor on a pedal to a break actuator in a car. The CAN, however, will continue to function despite any number of non-relevant nodes being removed or damaged in the network.

Data transmission within CAN is performed using a technique similar to synchronous transmission where clocks, internal to each unit, are employed to maintain synchronised reading and writing to and from a bus. Every CAN node constantly samples the network, if it is found to be in contention with another node, priority is given to node with the lowest identification (ID) number.

It should be noted that CAN IDs can be divided into two groups, part A of the CAN specification describes an 11-bit ID network, whilst part B describes an extended-frame, 29-bit ID. The 29-bit ID frame is referred to as CAN B, and is for lower speed operation, at around 125 kbps, vs the 11-bit ID CAN which can achieve 1 Mbps. The specification states that a CAN B device must not affect the operation of a CAN A network. To identify the type of CAN, an Identifier Extension (IDE) bit is used. If this bit is 1 then the message is in CAN B format, and 0 for CAN A. Work in this PhD is been focused on CAN A, as it is generally expected that with the need for an increase in data transfer within cars (for example, to permit driver-less operation) high data rates would be required. As such, only the frame for CAN A will be examined, though the general layout is very similar for CAN B.

A standard CAN A frame consists of 55 bits. A 1-bit start-of-frame is sent first, this is always a 0. Following this is a 12-bit arbitration field, which begins with the 11-bit node ID and ends with a 1-bit Remote Transmission Request (RTR). RTR will be a 0 for standard data frames,

where a node independently decides to transmit data and 1 for a remote request, where another node requests information from a particular node. After the arbitration field is a 6-bit control field. This starts with the IDE bit, to determine if the frame is in the standard or extended frame mode, indicated by a 0 or 1 respectively. Next follows a reserved bit which must be held at 0. Finally, in the control field is a 4-bit number which denotes the number of bytes of data within the frame. This is known as the Data Length Code (DLC) and may range between zero to eight bytes of data. After the control field is between 0-64 bits of data in the data field, the length of which is determined by the aforementioned DLC. After the data field is the 16-bit Cyclic Redundancy Check (CRC) field, consisting of a 15-bit, predetermined number which is used to detect the quality of the transmission. This number is compared by all receiving nodes to determine if the data was corrupted whilst travelling around the network. This number is followed by a 1-bit delimiter, which is used as the CRC field does not use bit-stuffing (discussed later on this page). After the CRC field there is a 1-bit acknowledge (ACK) slot, where the transmitter sends a 1 and receivers a 0 to confirm the message has been sent. This is followed by another 1-bit delimiter, and the CAN frame ends with a 7-bit end-of-frame (EOF), which must consist of all 1s.

CAN also makes use of bit stuffing for the detection of EMI. This technique involves the transmitter detecting the state of each bit as it is being transmitted. If five bits have been transmitted with the same value, either 1 or 0, then an extra bit is added and inverted. Normal transmission then resumes. However, in fields such as CRC, bit stuffing may appear like the effect of EMI to a transmitter, so it is not used. As such, a delimiter bit is used as a means of re-synchronising the network, should a long list of identical bits be transmitted. However, this does come at the cost of a much lowered bit efficiency, which maximises at 0.57 and can reach as low

as 0.15 when little data is transmitted in a frame [66].

So far, this section has examined the multitude of unique layer 2 and 3 methods for detecting EMI in a CAN system. CAN systems also have a unique physical layer implementation, which also fundamentally protects it from the effects of EMI. As an example of a standard for CAN implementation, [63] was referenced. This standard, first published in 2003 and revised in 2016, details the acceptable boundaries for a Medium Access Unit – in effect describing part of a suitable physical layer for CAN operation. Figure 3.5 is taken from [63] and shows the physical layout of a CAN bus.

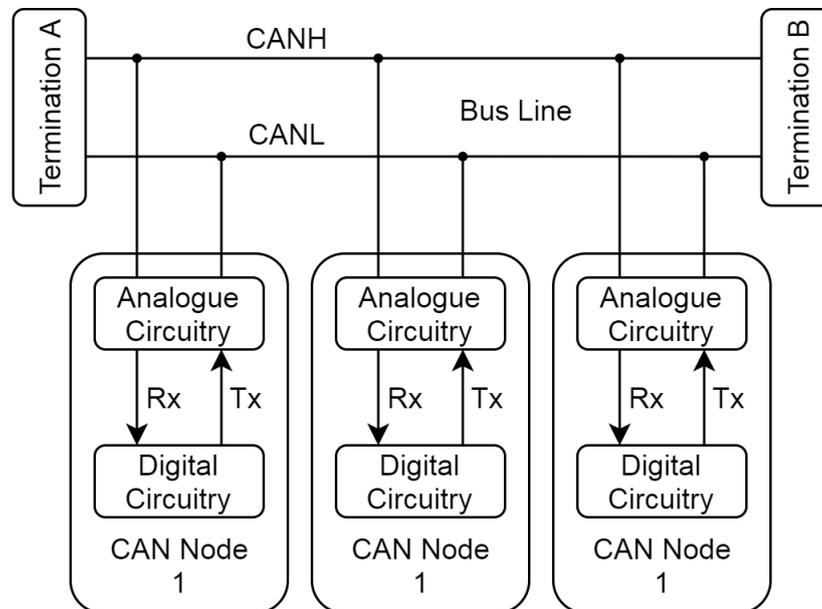


Figure 3.5: CAN bus layout [63].

It can be seen that the CAN bus itself consists of two wires; CAN high (CANH) and CAN low (CANL) configured for a differential-type signalling method. This involves representing each bit as two voltage levels. In the case of CAN, when transmitting a 1, both voltages of CANH and CANL are very similar, at around 2.5 V. In this condition, it is said that the bus is in the recessive state. When transmitting a 0, CANH goes to 3.5V and CANL goes to 1.5V - the dominant state. [63] gives the acceptable noise error margin for each state. For example, whilst the bus is in the recessive

state (in other words, transmitting a digital 1), the differential bus voltage must be between -120 mV and 12 mV. Any deviation from this, or other criteria, would result in disqualification from the standard and would in the eyes of this thesis, count as a failure to counteract EMI.

From the same standard, parameters regarding the physical construction of the bus are discussed. For example, the terminations shown in figure 1 must be $120\ \Omega$. This allows for a fast change in bit state due to a relatively low impedance, however the differential signalling is used primarily for its higher resilience to common mode noise. Take, for instance, the case shown in Figure 3.6.

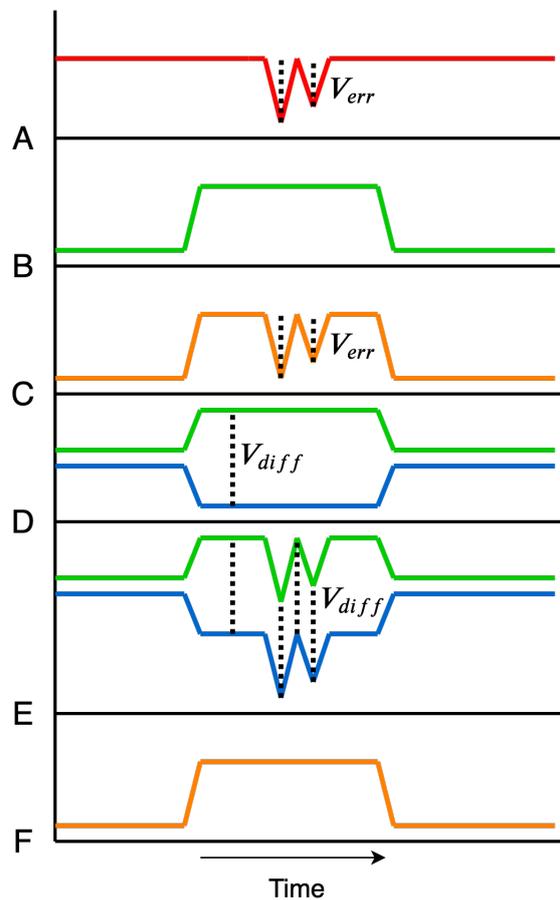


Figure 3.6: Effect of noise on data transmission.

Here an induced voltage swing V_{err} (A) affects the transmission of a regular, single-ended data bit (B) and a differential data bit (D). The noise is summed with B to create what the receiver sees on from a single ended

network (C) and a differential network (E) The result of the difference of E is shown in F. The induced voltage creates two substantial troughs in the waveforms of both the single-ended and differential data. In the case of (C) this may result in the receiver erroneously miss-sampling the data bit as a 0. However, the noise is coupled onto both high and low lines of the differential signal in D. The difference between the high and low lines remains the same V_{diff} throughout, as shown in F. Therefore, the data output matches what was transmitted in B.

As a result of its robustness to EMI in layers 1, 2 and 3 of the OSI model, a data rate suitable for contemporary cars, and network module addressing, CAN has cemented itself as a firm standard network in the automotive sector for networking complete vehicles. However, as the requirement for higher data throughput arises, the 1 Mbps capability is becoming ever more restrictive. As such, a need exists for new networks to fulfil higher bandwidth networking requirements.

3.3.4 FlexRay Networks - the Answer?

In order to facilitate higher network speeds on wired networks, the FlexRay Consortium developed the FlexRay standard [67], which is currently set by ISO 17458-1 [68] and ISO 17458-5 [69].

On the physical layer, FlexRay consists of one or two unshielded differential mode channels, allowing for two independent communication modes versus the single mode of CAN [70]. Indeed, FlexRay is specified to provide up to 10 Mbps data rates when in operation on each channel - 20 Mbps total. Furthermore, FlexRay nodes are configurable in star and bus topologies - ideal for creating highly customised networks for reducing weight in vehicles.

The FlexRay framing implementation also includes addressing and error detection, with a far higher amount of protection than CAN. As shown in

Figure 3.7, each frame can be broken into three components.

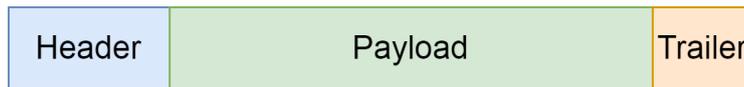


Figure 3.7: The three segments of a FlexRay frame [70].

The header contains a fixed 40 bits. The first 5 describe the type of frame being sent, which is followed by an 11 bit frame ID. A 7 bit payload length indicator is then sent, describing the size of the payload which may be up to 2032 bits long. This followed by a CRC, specific to the header. The header finishes with a cycle count which counts each communication cycle. This header is shown in Figure 3.8.

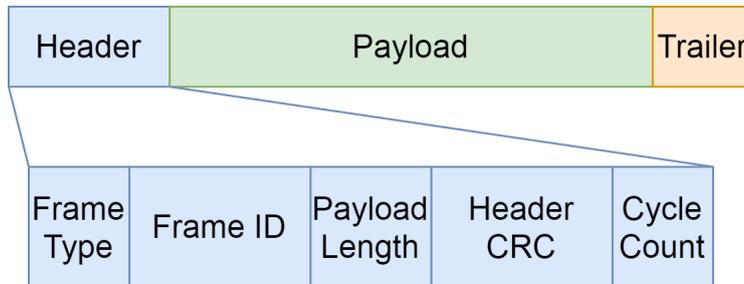


Figure 3.8: Detail of a FlexRay header.

The payload from Figure 3.7 consists of between 0-254 bytes (0-2032 bits). The trailer is made from three CRC for an extended examination of the susceptibility of the channel to EMI - three times as many for an entire CAN frame. This provides a great deal of layer 2 detection of EMI whilst the massive payload length means a bit efficiency of 0.97 is possible - up to 6.5 times as great as CAN.

However, adoption of FlexRay has been slow in the automotive sector and non-existent in other markets [71]. It is both costly and complex to implement [72]. Furthermore, the specific tailoring to automotive applications has resulted in little demand for FlexRay in other markets. This reduces demand for tools and development hardware, which only serves to increase adoption costs of FlexRay. Indeed, some sources, such as [72], qual-

itatively asses Ethernet to be better suited to adoption for high bandwidth wired networks in cars. However due to the complexity and cost of not only creating custom Ethernet networks, but also the cost associated with instrumentation capable of properly measuring the ultra-high bandwidth capabilities of Ethernet networks (up to 10 Gbps by IEEE 802.3ak-2004 [73]). The cost-per-node and network complexity of Ethernet has limited its adoption as a fix-all solution for intra-vehicular networks [74]. As a result, Ethernet is not examined further in this thesis. Instead, networks with bandwidths up to tens of Mbps are investigated as relatively high bandwidth in the context of intra-vehicular communications.

3.4 Chapter Conclusion

An introduction to the OSI model has been presented in this chapter to highlight the differences in approaches to EMI detection and mitigation amongst a selection of wired communication protocols, used for intra-vehicular communications. These protocols have been discussed, and a simple qualitative analysis of each can be seen in Table 3.4.

Table 3.4: Comparison between intra-vehicular communication protocols.

Name	Cost	Complexity	Bandwidth	Fault Tolerance	Adoption
TTL UART	Low	Low	3 Mbps*	Low	High
RS232	Low	Low	20 kbps	Medium	Low
CAN	Medium	Medium	1 Mbps	High	High
FlexRay	High	High	20 Mbps	High	Low

*Formally unspecified, but is an extreme upper limit provided by [58].

The results of this analysis suggest that an alternative wireless network must not only include a robust communication channel, as per existing standards, but quantitatively must at least offer 1 Mbps data rates to compete with the popular CAN protocol.

Chapter 4

The Vehicular Electromagnetic Environment

The process of determining the suitability of any communications system, wired or wireless, cannot take place without first understanding the EM environment in which it is expected to function. Part of this task requires measuring the RF emissions within a vehicle and using this to derive realistic interference sources. These source models can then be used with relative ease in laboratory experiments for testing the resilience of networks, as well as determining appropriate emission standards. Previous work, such as [75] describe the EM spectrum across multiple frequencies inside an vehicle, though at over 21 years of age at the time of writing this thesis, the vehicle itself is dated and no longer representative of contemporary vehicles - let alone the vehicles of the next decade.

This chapter provides a new EM survey on the EM environment inside a Kia Soul EV 2017, as shown in Figure 4.1 A commercially available, electric vehicle it allows for the measurement of emissions from a car with an electric drive system. As discussed in the Introduction, this is the power drive for cars in the future - not the internal combustion-engined vehicle as present in [75]. This particular car was also retrofitted with equipment and sensors

to perform autonomous driving, as part of the Machine Learning and Vision (MLV) group at the Gwangju Institute of Science and Technology (GIST), Korea. With a number of standard communication network protocols on-board (such as TTL UART, RS232 and CAN), along with much higher bandwidth networks and devices (such as servers, LIDARs and thermal cameras) this vehicle would allow for the acquisition of new data on the EM environment around a vehicle laden with high bandwidth networks and equipment - unlike vehicles before it.



Figure 4.1: Photo of the Kia Soul EV 2017 at the MLV, GIST, Republic of Korea.

4.1 Experimental Background

A large number of standards for measuring the EM environment exist. These standards cover a wide range of frequency bands, in different environments at different power levels. The measurement of EMI in the automotive environment usually requires the measurement of fields with a spectrum analyser and antenna outside the vehicle as per CISPR 25, CISPR 11 and SAE J551 [76, 77, 78]. However, as [75] highlights, taking near field measurements within the tight confinements of the interior of a vehicle makes it impractical to measure the emissions from a specific device

under test (DUT) when compared to measuring a single unit (i.e. the car) from a distance of several meters or more.

For example, CISPR 11 states that measurements in the 9 kHz to 1 GHz region should be performed “at a close distance but not less than 3m”, whilst in the 1 GHz to 18 GHz band the distance between the DUT and receiving antenna is fixed at 3 metres. CISPR 25 is marginally better, with a separation of 1 metre. However, both of these measurements are clearly impractical within the tight spacial constraints when measuring the interior environment of a car measured to be only 1.8 meters in width. This is a result of the scope of standards such as CISPR 11 and CISPR 25 primarily covering the *inter*-vehicular communication cases. However, the methods presented in [75] offer a means of measuring the near field environment within a car by simply recording the field components at precise locations within the vehicle. Whilst such an approach may not provide a pass/fail qualification for DUTs (due to the highly varying field strength within the near-field environment), it does provide a means of creating numerical data sets for the variation of the EM environment within a vehicle as well as field strengths which may be compared with standards such as CISPR 25 - if proper consideration of the difference in measurement procedure is minded.

Another variation from the procedures presented in the popular automotive standard CISPR 25 is the lack of any ground plane and anechoic facility for measurements. However, this does not necessarily reduce the importance of the captured data if context is considered. In the case of the ground plane, it was unfeasible to create a ground plane large enough for the car, especially when travelling distances of 800 metres or more. The absence of a ground plane, as specified in CISPR 25, would only serve to reduce the measured field strength as the non electrically reflective asphalt road surface would reflect far less EM energy towards a receiving antenna than a smooth, conductive metal sheet.

The unfeasibly large size of the required anechoic chamber meant that no attenuation of external signals was available during measurement. However, Annex C of CISPR 11 describes a means of subtracting a background signal from a measured emission source as per (4.1).

$$E_g^{1.1} = E_t^{1.1} - E_s^{1.1} \quad (4.1)$$

Where E_g is the electromagnetic radiation disturbance, E_t is the measured value of the electric field strength, and E_s is the background electric field strength - all terms in units of $\mu\text{V}/\text{m}$. The powers of 1.1 are mostly likely an experimentally-derived means of scaling the resultant source. Annex C of CISPR 25 advises that whilst it is advisable to use (4.1) when it is not possible to avoid the effect of transmitters, it should only be used when all sources are stable. By repeating background measurements for E_s over the minimum scan time dictated by CISPR 25, it was observed that the remote nature of the test track for the car provided a relatively quiet EM environment, with all sources being stable in transmission. Therefore this approach from CISPR 25 was deemed suitable for reliable measurements of the emissions from devices within the car.

The highly varying nature of a near-field emission across a spatial dimension means that a number of test locations would be required to determine the EM environment of the car - simply taking one measurement within the vehicle would not represent another point only a few tens of centimetres away. Therefore, a list of six test locations was compiled, as per Table 4.1. These were chosen to meet three criteria. The first was to reflect measurements taken in previous work, such as [75]. The second criteria was to examine how EM spectra could vary over short distances in an automotive case. The third was to chose locations which would be classed as functionally important to the operation of the car, reflecting the observations made in

Chapter 3 that the future of intra-vehicular communications networks lies in vehicle functionality, rather than for more trivial purposes.

Table 4.1: Test locations on the Kia Soul EV 2017.

Number	Location	Rationale
1	Car Radio Antenna	A common location for a basic RF application. Measurement location in [75].
2	Rear Left Boot	Inside compartment shared with main server and power supply.
3	Driver Display Unit	Important location for relaying car performance metrics to a driver.
4	Rear Right Passenger	Location of passenger devices such as mobile phone or laptop.
5	Rear Left Wheel Well	For examining the change in EM environment with minimal separation from server (30 cm).
6	Under Bonnet	Same compartment as high power motor drive unit. High levels of EMI anticipated.

The corresponding location numbers are shown on an orthographic drawing of the Kia Soul in Figure 4.2. The locations of large modifications to the original Kia Soul are also shown, with colour coding to indicate whether the alteration is for sensors, computing and network, or power applications. A large amount of equipment was added to the boot of the car, primarily consisting of processing equipment. More specifically, a server consisting of two Intel i7 processors, two Titan X graphics cards, and 32GB of RAM was bundled with an Odroid XU4 (NAS) device and a WiFi router. The boot also contained a 2500W DC converter for powering all retrofitted equipment.

The front and rear bumpers of the car were fitted with Delphi ESR 2.5 RADAR modules. These devices produce high power, but highly directional EM beams (between ± 45 degrees and ± 10 degrees).

A number of sensors, vital for the operation of the vehicle were mounted roof racks. These included a Velodyne HDL-32E LIDAR for high resolution mapping of the environment, four Logitech C920 Webcams and a FLIR T420 thermal camera for low light conditions, and a Novatel SPAN-CPT Global Positioning System (GPS).

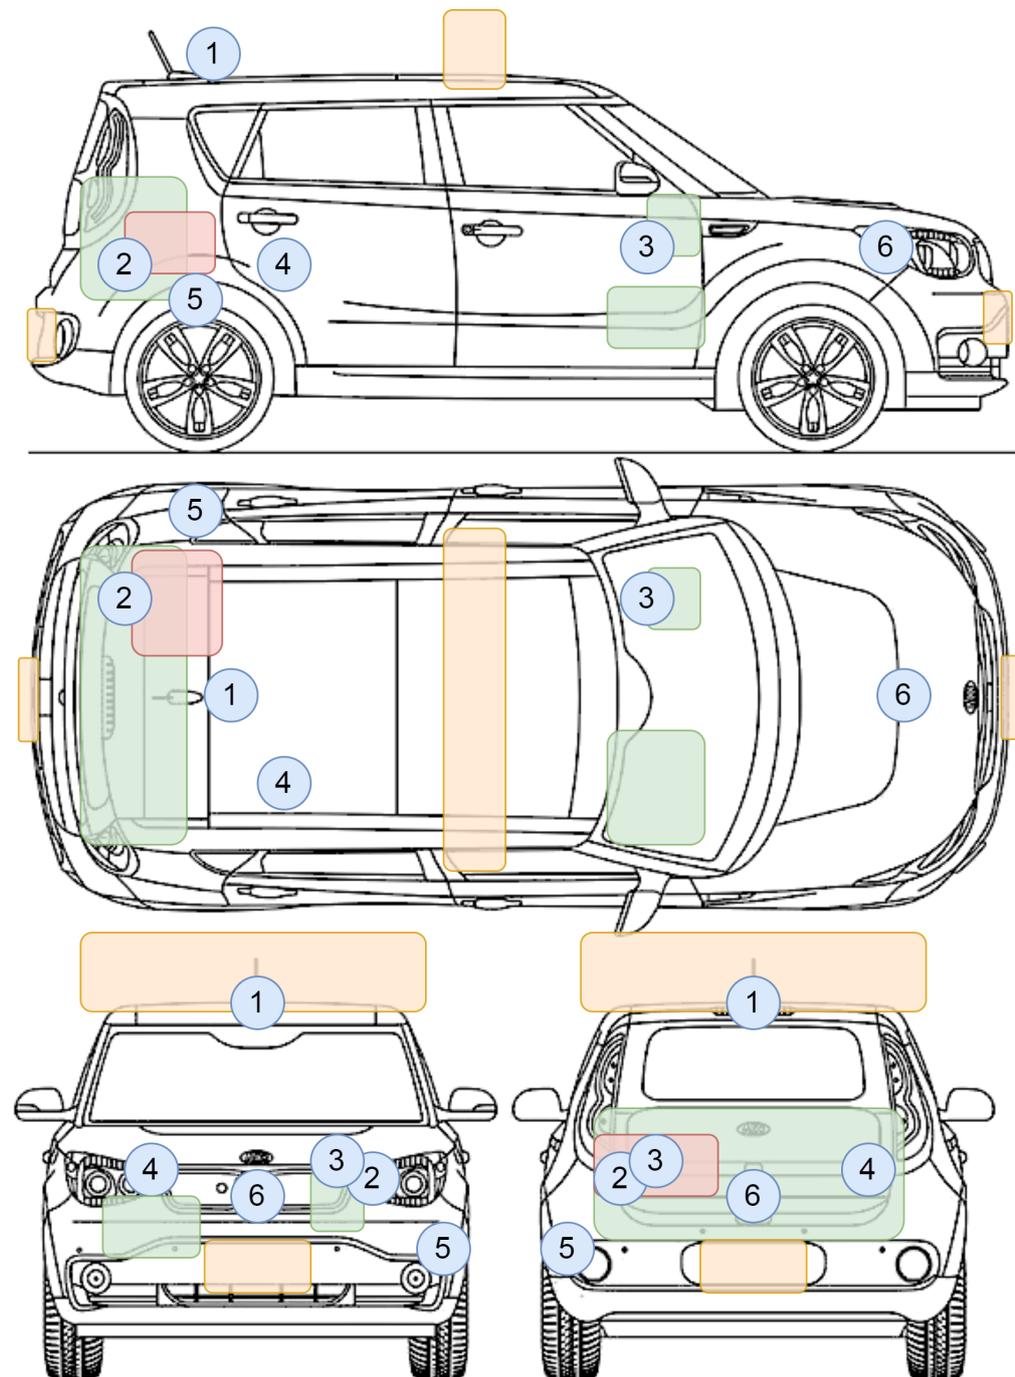


Figure 4.2: Test locations around the Kia Soul EV 2017. Orange represents sensor locations, green represents computer locations, and red represents power supply locations. Edited from [79].

At the front console, two smaller subsystems were installed for initialising the operation of the autonomous driving capability through a personal computer (PC) in the passenger's footwell. This was directly interfaced to a PolySync Drive Kit which connects to the vehicle's CAN for full control over all vehicle operations.

The large volume, and high density of power and computing equipment in the boot meant several test locations were selected for this region, primarily locations 1, 2, 4 and 5. Location 3 was chosen for its position at the driver's interface and the PolySync Drive Kit. Finally, location 6 was chosen for its close proximity to the 375 V, 3-phase AC synchronous motor [80] - initially thought to be a potential source for a great deal of EMI.

The experimental setup consisted of a Keysight FieldFox N9912A portable spectrum analyser and a series of antennas and near field probes. Whilst CISPR 25 permits the use of a preamplifier, it was noted during initial measurements that this was not required in order to detect the sometimes significant levels of EMI present. In order to maintain a consistent setup throughout the measurement regime, and to facilitate quick comparisons during the time-limited measurement schedule, no preamplifier was used. A typical setup is as shown in Figure 4.3.

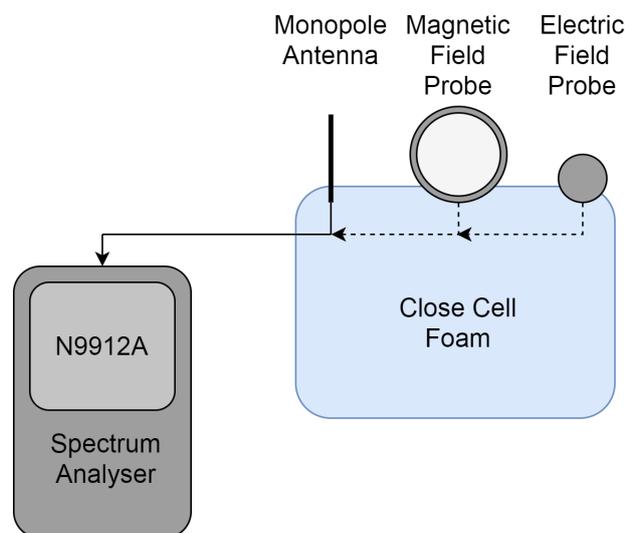


Figure 4.3: Typical setup for radiated measurements.

In this case the hardware for the experiment is very simple. One of the antennas from Table 4.2 is directly attached to spectrum analyser. For certain antenna and measurement locations, antennas were held in place with a combination of closed cell foam and tape - both transparent to EM waves at all frequencies of interest and with less influence on measurements than a human operator.

Table 4.2: Antenna choices for the experimental regime.

Probe	Measures	Rationale
0.50 Metre Monopole	Electric and Magnetic Fields	Similar antenna to that used in CISPR 25. Simple design, with equal radiation pattern at the same plane. A useful base measurement.
0.11 Metre Monopole	Electric and Magnetic Fields	Antenna which better matches a real implementation antenna. Smaller than the previous monopole, therefore easier to position.
Model 904 Ball Probe	Electric Field	3.6cm ball probe for isolating electric fields. Best for measuring high impedance sources, such as long network cables.
Model 901 Loop Probe	Magnetic Field	6cm loop probe for isolating magnetic fields. Better for detecting low impedance sources, such as high current supplies.

CISPR 25 states that any receiving element may be used, such as the near field probes, providing they have the proper correction factor applied to them. [81] describes the correction factor as the ratio of the incident electric field to the voltage induced at the antenna's terminals, as per (4.2).

$$A_f = \frac{E_i}{V_r} \quad (4.2)$$

Where A_f is the antenna factor in 1/m, E_i is the incident electric field in dB μ V/m and V_r is the induced voltage at the receiver in dB μ V. By carefully selecting the logarithmic scale for voltage, A_f can be simply summed with the measured signal to calculate the incident field.

Frequencies for measurement can be divided into three bands. The first, of 0.15 - 30 MHz, is derived from the results shown in [75], where only emissions from a CAN bus in the high frequency (HF) were shown. By performing measurements in this band, it was anticipated that a comparison of the EM environment could be made with a vehicle from the late 1990s. The second two bands, of 902 - 928 MHz and 2.4 -2.5 GHz are both ultra-high frequency (UHF) bands reserved specifically for industrial, scientific and medical (ISM) applications. Whilst the 902-928 MHz band is available only in Region 2 (as laid out by the International Telecommunication Union in [82]), it is an unlicensed band in the Republic of Korea. These bands are available for unlicensed wireless communication development, providing they do not interfere with other devices and meet certain power and duration requirements. As such, they would make ideal bands for the development of an intra-vehicular network, as examining them may prove useful.

With the measurement bands decided, the final consideration for this setup was the bandwidths and sweep durations for the spectrum analyser. Fortunately, such detail was presented in CISPR 25, as in Table 4.3 and Table 4.4. These parameters were set appropriately for each of the following measurements.

Table 4.3: Minimum scan time

Band	Peak Detection	Quasi-Peak Detection
9 - 150 kHz	NA	NA
0.15 - 30 MHz	100 ms/MHz	200 s/MHz
30 - 1000 MHz	1 / 100 ms/MHz	20 s/MHz

Table 4.4: Measuring Instrument Bandwidth

Band		Broadband or Quasi-Peak	Narrowband or Average
0.15 - 30 MHz		9 kHz	9 kHz
30 - 1000 MHz	FM	120 kHz	120 kHz
	Mobile	120 kHz	9 kHz

Generally speaking, unless specified otherwise, the minimum limits for peak detection measurements for narrowband signals was observed due to time limitations.

4.2 Frequency Domain Measurements

Figure 4.4 demonstrates the method of subtracting background noise from (4.1) in CISPR 11. Measurements were taken in the boot.

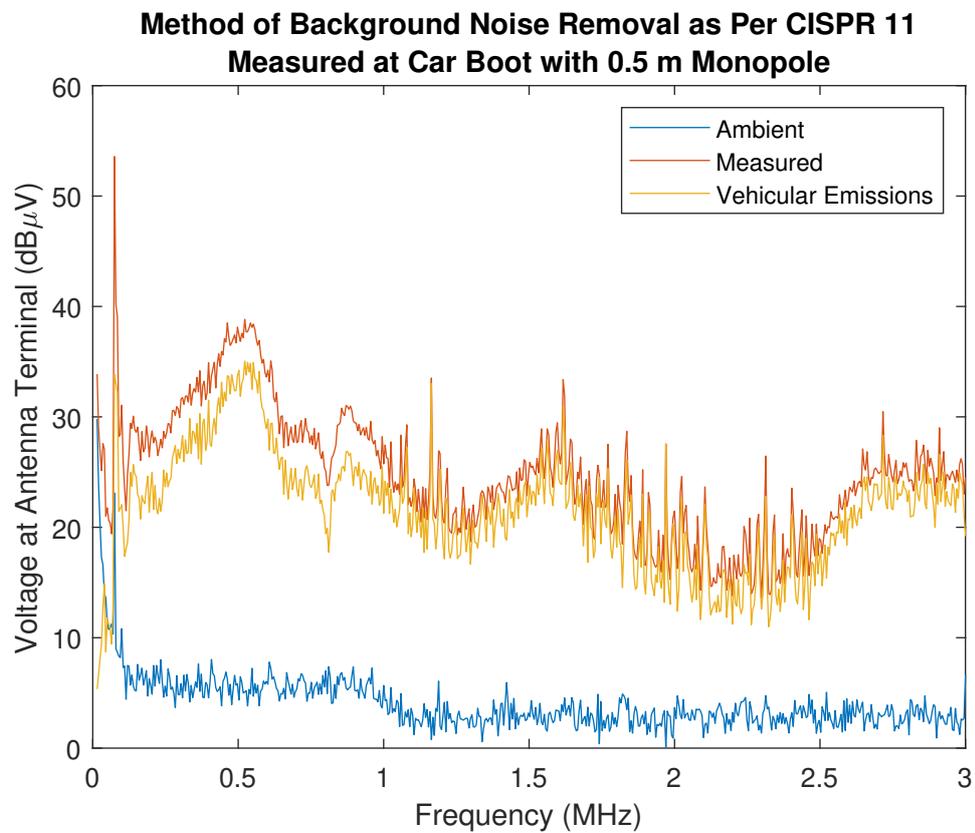


Figure 4.4: Example spectra to determine the net vehicular emissions in a noisy EM environment, as per CISPR 11 Annex C.

The large peak in the measured spectra at 0.1 MHz when the vehicle is switched on (in orange) may be interpreted as an emission of almost 54 dBμV, if it were not for the background measurement (in blue). This measurement indicates a steady background radio source, which contributes 23 dBμV to the measured signal. Subtracting this background contribution

from the measured spectrum results in the actual vehicular emission spectrum (yellow), which indicates the actual voltage induced at the antenna's terminals is $34 \text{ dB}\mu\text{V}$ - above the CISPR 25 limits, but far lower than might be assumed without accounting for background emissions. As this chapter reports only on the emissions from the vehicle itself, all subsequent spectra are the result of subtracting background spectra from the measured.

The spectra captured in the 0.15 - 30 MHz band with a 0.5 m monopole antenna clearly demonstrates the high spatial variability of the EM field in close regions of the car. Emissions from the boot (number 1 in Figure 4.2) dominate the other measurement regions in terms of amplitude, greatly exceeding the limits set out in CISPR 25 of $22 \text{ dB}\mu\text{V}$ for broadband continuous peak sources. Even for short duration broadband sources, the peak at 0.5 MHz exceeds the $28 \text{ dB}\mu\text{V}$ threshold by nearly $7 \text{ dB}\mu\text{V}$, as show in Figure 4.5.

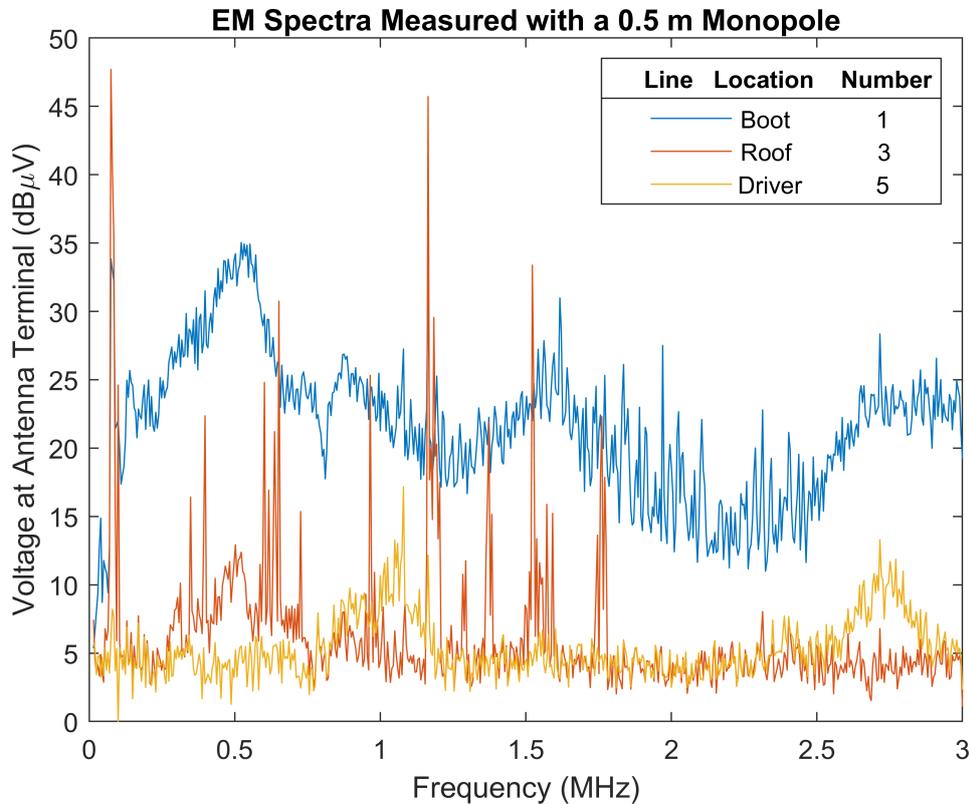


Figure 4.5: 0.15 - 30 MHz spectra, measured with a 0.5 m monopole.

By switching on and off differing combinations of the equipment in the boot, it was determined that the majority of the emissions in this region were the result of the power supply, a strong indicator that regardless of the advanced autonomous technology present in the vehicle, the utilisation of properly designed power supplies is also extremely important to meet emission standards.

It is possible to note that examples of field leakage, directly attributable to the emissions from the boot were observed in particular frequency regions in the driver and radio test locations (numbers 3 and 5 on Figure 4.2). Whilst both driver and radio locations attenuated the field strength to base levels, different frequency regions were observed at each point, despite being around the same distance from the source (0.5 m). Figure 4.5 shows that emission peaks from the boot (in blue) line up nearly perfectly with those recorded in the driver's seat (in yellow), though attenuated by at least 10 dB, in the case of the 10.8 MHz emission and 15 dB in the 27.2 MHz peak. A single example of EM leakage from the boot exists in the measurement at the roof (orange), suggesting a greater attenuation of emissions due to the shielding effect of the car's body. The attenuation of 22 dB at 5 MHz suggests that the shielding effectiveness of the car's body is greater than that of the rear boot bulkhead.

The spectrum of the radio location contains far more peaks than either the boot or driver's spectra, with multiple examples of peaks exceeding the amplitude of those in the boot and driver locations. This indicates that an EM emissions source is likely closer to the radio test point than either the boot or driver, or another emissions source has a channel to the measurement location with a lower shielding effectiveness than the channel to boot and driver locations. In this case, it was determined that the emission peaks were caused by the autonomous sensors fitted to the roof. With no means of attenuating the line of sight between the noisy sensors

and the radio test location, very high amplitude peaks were observed. This suggests it may be possible to determine the location of particular emission sources through making multiple frequency domain measurements in differing locations. The amplitude of any peaks may then be compared in order to triangulate the source. This is, however, beyond the scope of this work, though members of GIST expressed interest in using their machine learning capabilities to explore this further.

The 902 - 928 MHz band was measured using a 110 mm monopole antenna, with the results from all six test locations shown in Figure 4.6.

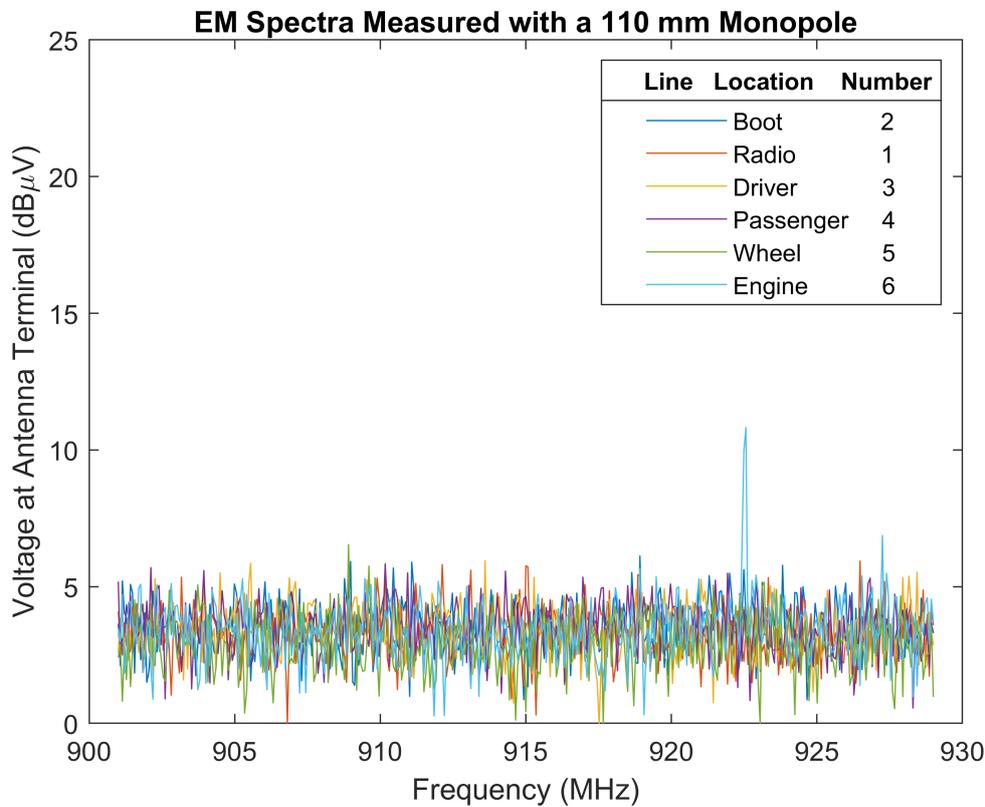


Figure 4.6: 902 - 928 MHz spectra, measured with a 110 mm monopole.

The 110 mm antenna was chosen for this band as a better approximation of a practical antenna in a UHF network, when compared to the 0.5 m monopole. The spectra throughout the vehicle was much quieter than in the 0.15 - 30 MHz band, with only one noticeable peak observed at 922 MHz in the engine compartment. This peak, a high-order harmonic from power

switching as determined by systematically turning vehicle components on and off, is completely attenuated by the engine bulkhead and does not appear in other measurements. The low level of activity in this band indicates that the 902 - 928 MHz band may well be suited to an intra-vehicular network, as no onboard equipment utilises this frequency range.

The smaller dimensions of the 110 mm antenna allowed for fixing the antenna in place securely during travel. As such, the opportunity was taken to perform the same measurements during travel, as shown in Figure 4.7. Note, when comparing to the levels in Figure 4.5, that the y axis is at half of the scale.

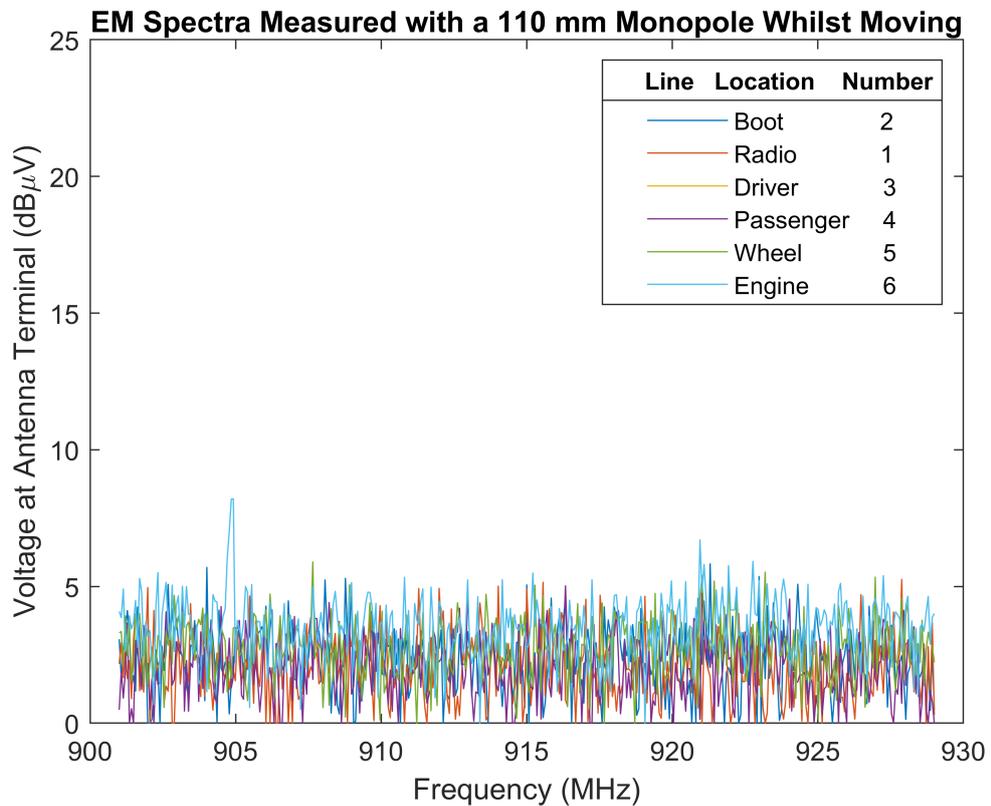


Figure 4.7: 902 - 928 MHz spectra, measured with a 110 mm monopole, whilst the car is moving.

Generally, all the spectra from all test locations do not change whilst the vehicle is moving - though with one notable exception. Again the engine compartment measurement produces one distinct peak, but it occurs at a

frequency of 905 MHz - a 15 MHz frequency shift from its location in Figure 4.6. The large variation of this peak may make wireless intra-vehicular communication more difficult in this environment, though techniques such as FHSS would mitigate the effect of any variable frequency emission by hopping between differing carrier frequencies, as described in Chapter 2.

Whilst the 2.3 - 2.5 GHz band exceeds the accommodated frequency range in CISPR 25, measurements were also performed in this band to assess the suitability for a wireless, intra-vehicular network. The results of the measurements when the car was stationary are shown in Figure 4.8.

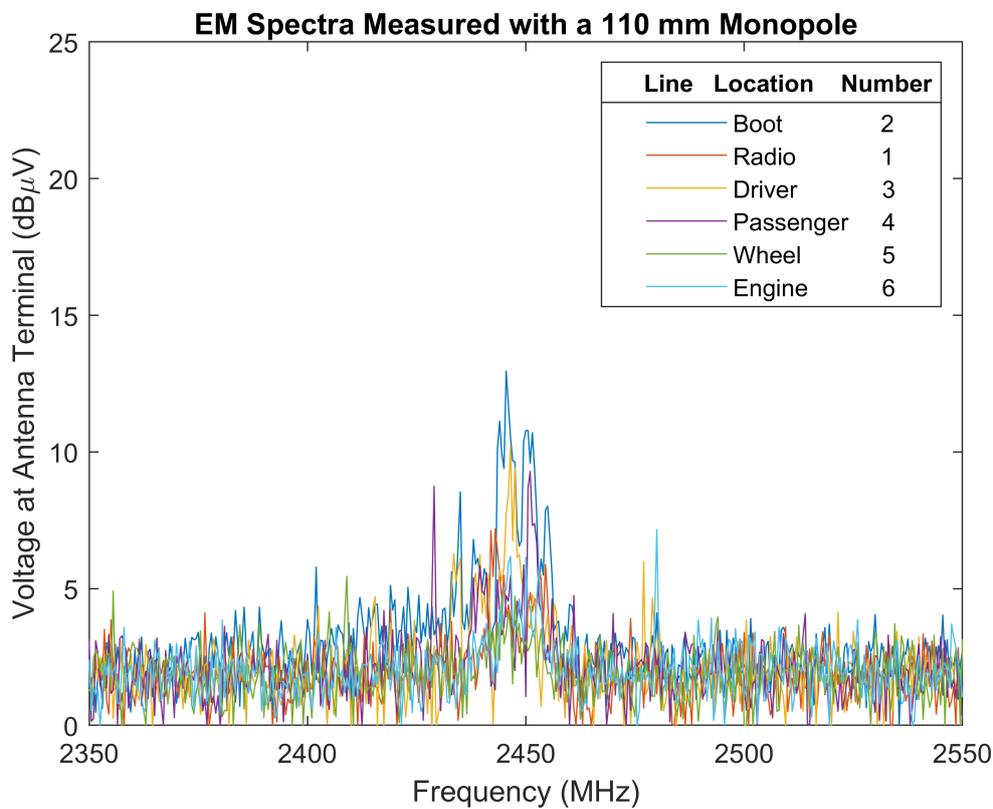


Figure 4.8: 2.3 - 2.5 GHz spectra, measured with a 110 mm monopole.

In this band, the spectra are far more occupied by equipment in the vehicle - primarily 2.45 GHz emissions from the wireless router. Extrapolating the CISPR 25 limits for 1 GHz to these spectra, it can be seen that none come close to exceeding the 28 dB μ V limit. Further to the observations made about the relative positioning of emission sources made regarding

the 0.15 - 30 MHz emission sources, a similar idea can be seen in the 2.4 GHz band. The closest measurement location to the main source of these emissions (in the boot) clearly stands higher than all other sources, whilst moving further through the vehicle's cabin through the passenger and driver measurement points the signal decreases in strength. By the time the signal reaches the engine compartment, it is nearly entirely attenuated.

However this band also highlights that in the complex vehicular environment, a decrease in field strength is not only dependent on the distance between source and measurement location. The plot in Figure 4.9 shows the received signal strength from measurements performed in the rear wheel well. Emission sources in the boot, separated by around 30 cm, are far lower in amplitude than the same signal measured at a distance of 1.5 m at the driver's console.

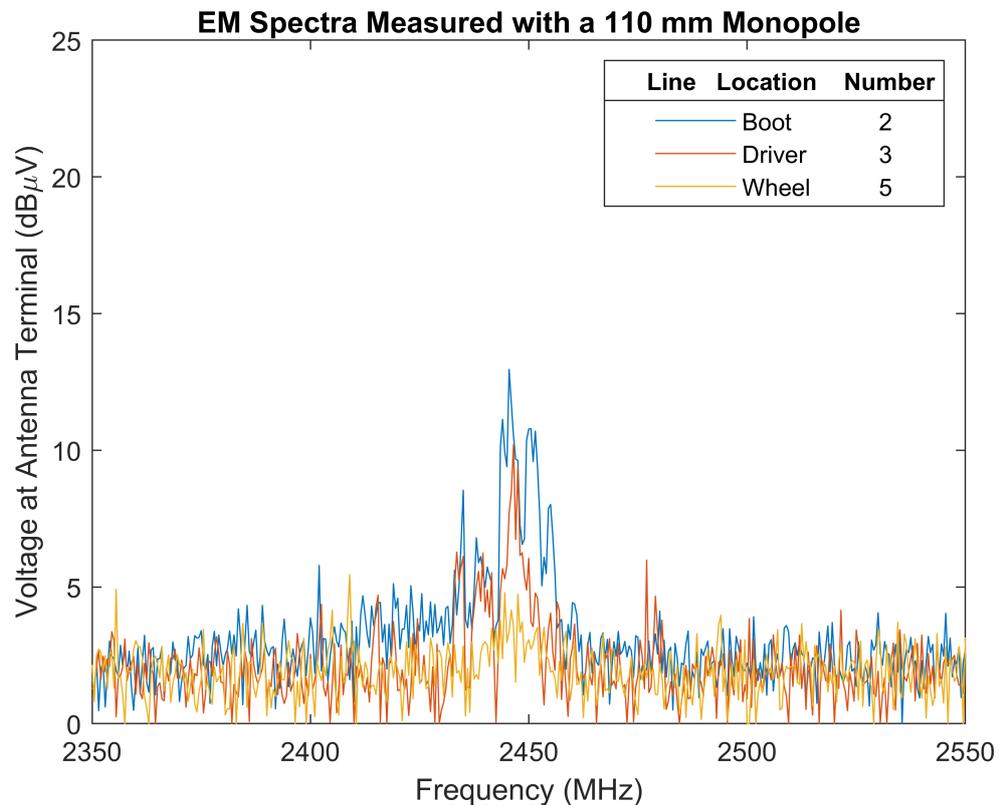


Figure 4.9: Difference in amplitude of measured signals, of varying distance to the emissions source.

This is primarily due to the large number of metal objects (boot bulkhead, car body, chassis, brakes, etc) between the boot and wheel well. Whilst these measurements indicate that the 2.4 GHz ISM band is already quite congested within the vehicle itself, the results do indicate the need for caution when assuming the operational capability of wireless channels in a complex environment such as a car. Whilst tests may indicate that a certain node to node communication is reliable, this does not guarantee the capability of another node combination to function - even if they are closer, spatially, by an order of magnitude or more!

Measurements were also conducted using near field probes, to isolate the electric (E) and magnetic (H) fields at each test location for comparison with results from [75]. A model 904 ball probe was chosen to measure the E-field at the prescribed test locations. The ball probe was chosen as it is more sensitive to the E field than stub probes, at the cost of refined location detection [81]. Effectively a metal ball on the end of a $50\ \Omega$ termination, the absence of a closed loop within the probe results in no current flow - eliminating any contribution from the H-field, as per Figure 4.10.

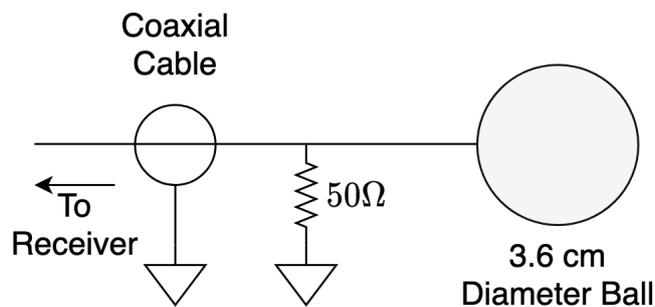


Figure 4.10: Ball probe construction.

A model 901 H-field probe was used to measure the magnetic field. Construction consists of a loop of coaxial cable with the centre conductor shorted to the shield at the base of the loop. A notch is cut in the shield at the midpoint of the loop, which rejects the E-field due to the balanced shielding of the conductor.

Figure 4.11 shows the E and H-fields at the passenger seat in the 0.15 - 30 MHz band.

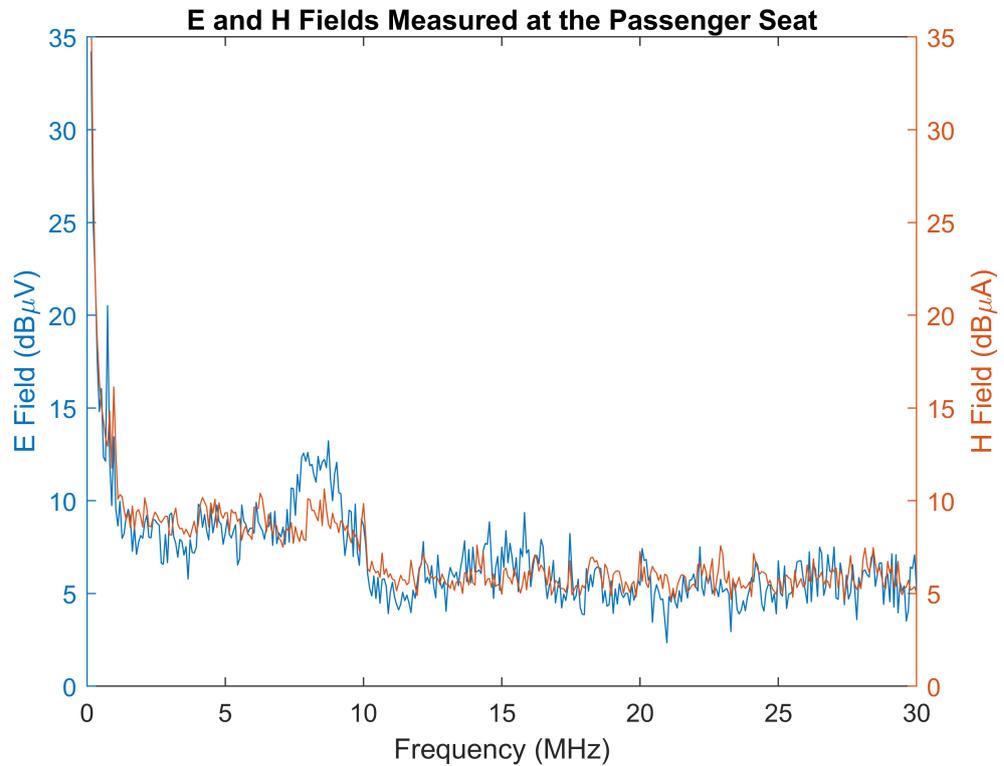


Figure 4.11: Measured electric and magnetic field emissions at the passenger seat.

When compared to results from the passenger test location in the car from [19], it is apparent that both the E and H-fields are significantly lower than that from the vehicle made in the late 1990s. The Noise floor is lower by around 15 dB in both measurements, and there are no harmonics from any communications network. Data from the 1990s car showed harmonics across the 10 - 25 MHz region in the E-field, and across 10 - 20 MHz in the H-field. However, these have either been very effectively eliminated as an emissions source, perhaps through the use of refined slope control implementations, or are better shielded from radiating into the wider environment.

It is possible to note the contribution from the power inverter at 9 MHz, which is quite noticeable in the E-field, with a peak amplitude of 14 dBµV, but much smaller in the H-field at only 2 - 3 dB above the noise floor.

The measurements from the engine compartment are presented in Figure 4.12. Again the noise floor is much lower in both E and H-field measurements when compared to the measurements from the 1990s, and no obviously distinct communication harmonics are observable.

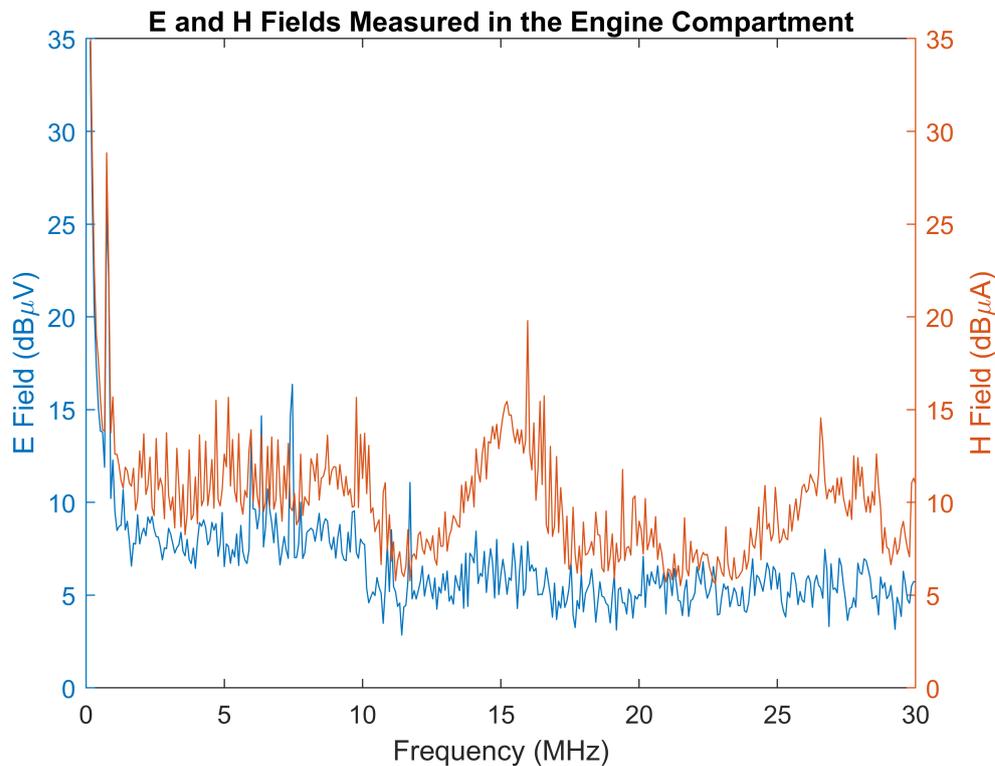


Figure 4.12: Measured electric and magnetic field emissions in the engine compartment.

In this plot, the major contributions are nearly entirely from the H-field, particularly around 15 MHz and 25 MHz and above. Evidence of emissions from a communications system is present in the high peak at 1 MHz of around $28 \text{ dB}\mu\text{A}$ - a strong indicator of a CAN network at maximum bandwidth, which does match the maximum peak amplitude measured in the vehicle from the 1990s. However it is important to note that if this peak is indeed the result of a CAN emission, it is not a wide-band emission and less likely to interfere with other systems. The lack of harmonics again suggests that improvements have been made in reducing the emissions from such communication systems over the past two decades, though the wider

emission peaks in the H-field suggest that some emissions from the electric drive may become more problematic unless they are properly designed.

In addition to the two measurements taken from the engine and passenger locations, other field measurements were taken around the vehicle for reference. One such example of particular interest is that from the boot of the vehicle. In previous measurements, the boot has proved to be noisier than other regions of the vehicle and the E and H-field measurements again confirm this is the case, as shown in Figure 4.13.

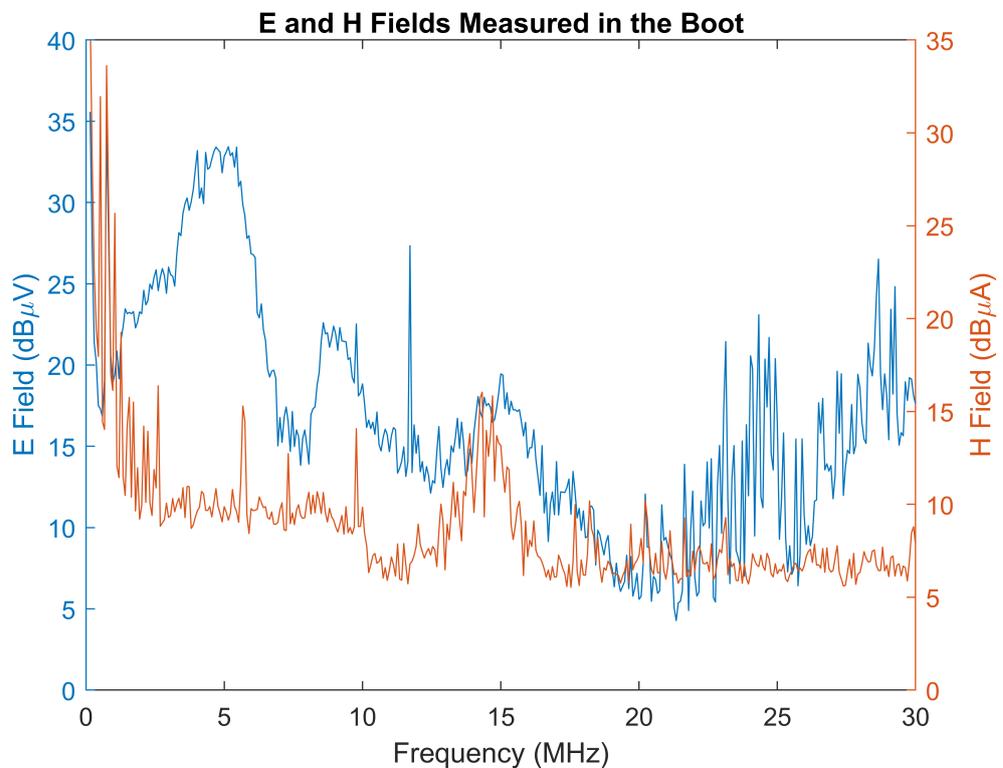


Figure 4.13: Measured electric and magnetic field emissions in the boot.

Due to the large amount of high power and high bandwidth communication equipment in the boot, a variety of emission types can be noted in both the E and H-field. Large wide-band emissions present in the E-field are not present in the the H field at 5 and 10 MHz, though one such emission at 15 MHz is visible in both. These independent measurements of the electric and magnetic fields reinforce the measurements taken of the EM spectra

earlier in this section, and again demonstrate that over a short distance the precise nature of the emissions varies greatly in the complex automotive environment. For completeness, all electric and magnetic field plots are available in Appendix A.

4.3 Time Domain Measurements

The team working on the Kia described some very intermittent failures of the car to react to objects which were fully detected by the autonomous driving sensors and machine learning algorithms, suggesting a problem between the output of the autonomous driving computer and the car's driving system. The vehicle had several modules attached by the MLV group to the internal CAN bus, including a module for controlling the car's control inputs. It was decided to examine the CAN bus network in the time domain, using an oscilloscope, in order to try to detect any interference on the transmitted and received CAN information, and to try to corroborate this interference with the previous measurements, performed in the frequency domain.

The experimental setup for this work is shown in Figure 4.14. Two channels of a GW Instek GDS-2102E, 100 MHz, 1 GS/s oscilloscope were connected through calibrated x10 attenuation probes (minimising capacitive loading) to the CANH and CANL stub lines. These stubs were, in turn, connected to the internal CAN loom. The stub lines, created by the students of the MLV consisted of two free wires, which were not physically connected in any way. The lack of any impedance controlled, or twisted pair conductors would likely allow easy coupling of differential mode noise. Furthermore, neither the stub, nor the equipment on the node were shielded, allowing both common and differential mode noise to couple onto the CAN wiring loom.

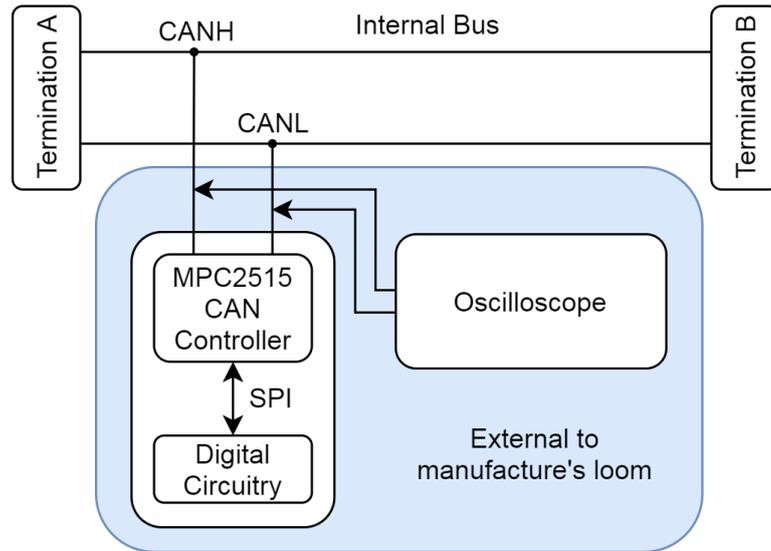


Figure 4.14: Setup for time domain measurements on the CAN bus.

Several measurements would be taken, to first determine the normal operating condition of the CAN network, with only the car switched on, and no autonomous driving sensors or computing equipment, indicating whether or not the additions to the CAN bus were detrimental to the normal operation of the car in a typical EM environment. This requires the assumption that the internal loom can operate in a typical EM environment - as the car manufacturer Kia is well known for producing vehicles which meet international standards, this was deemed a reasonable assumption.

An example of the CAN bus under normal operation, with only the standard vehicle switched on, can be seen in Figure 4.15. Both the dominant and recessive states of the CAN data bits can be clearly resolved at the correct voltage levels described in Chapter 3. Two independent coupled noise frequencies can also be observed at around 40 kHz and 350 kHz - though both are relatively small levels of common-mode noise. This common mode noise would be easily rejected by the MCP2515 CAN controller, used on these additional stubs, and easily falls within the maximum permitted voltage for the device inputs of 6.3V ($V_{SS} + 1$ in the datasheet [83], where the supply voltage V_{SS} is 5 V.).

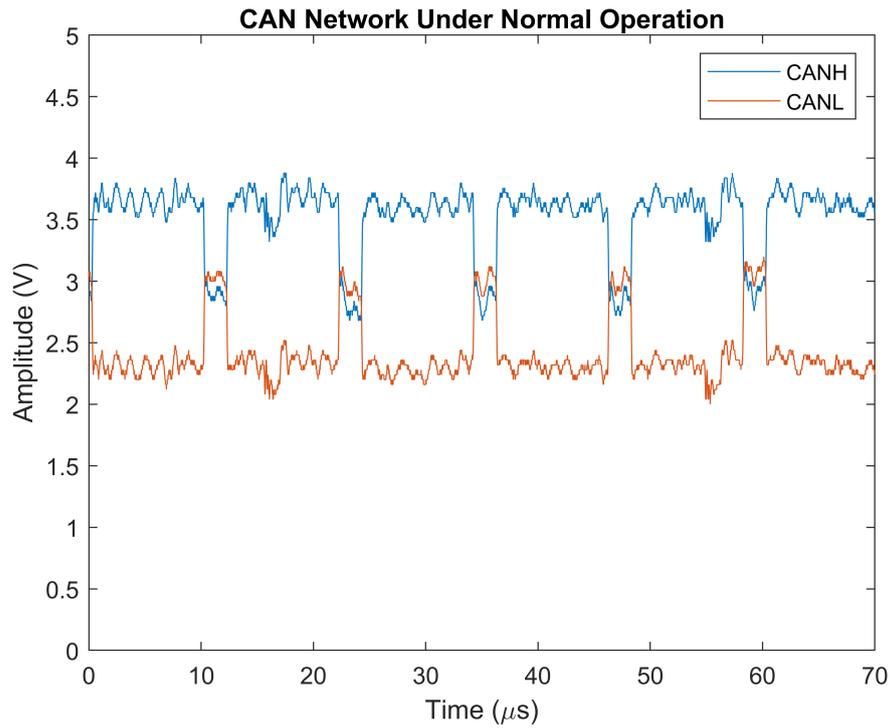


Figure 4.15: CAN bus voltages with only the standard vehicle in operation.

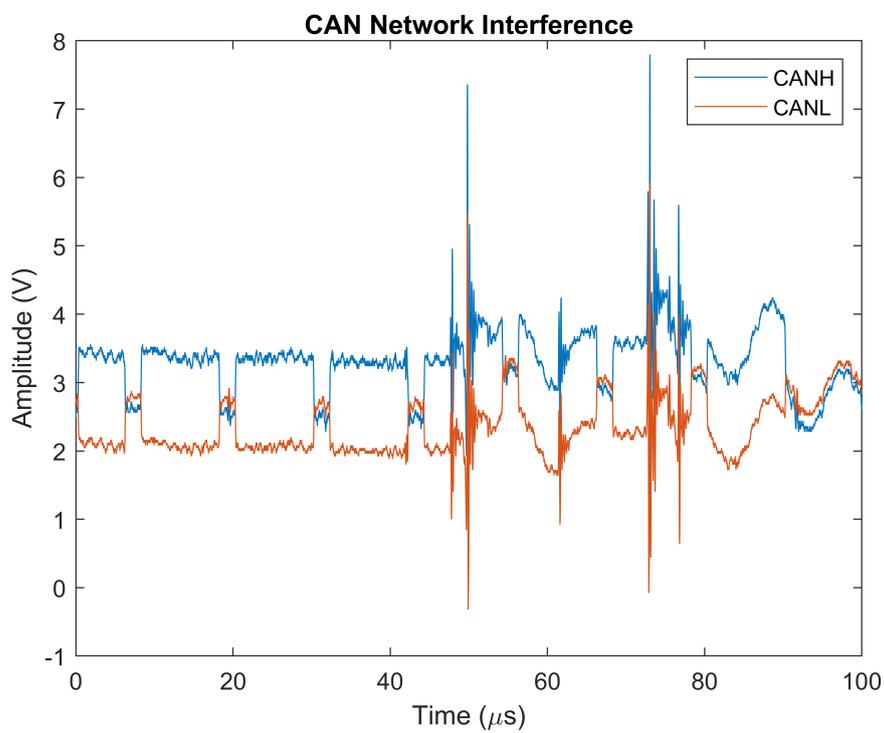


Figure 4.16: CAN bus voltages during initialisation of autonomous equipment.

Figure 4.16 shows a number of CAN bits at the precise moment the autonomous equipment was turned on. Multiple periodic sources of interference with a Δt of around $20\ \mu\text{s}$ (50 kHz). This is likely interference from the fundamental component of the inverter used to power the equipment, with the subsequent higher frequency ringing with a period of $0.2\ \mu\text{s}$ (5 MHz) resulting from reflections caused by the 50 kHz component.

Lower frequency common-mode interference can also be observed, but this would not result in incorrect data arbitration by the CAN controller as the voltages are within the maximum permissible for the device, and are still of the common-mode form.

Another example of the serious induced voltages on the CAN bus can be seen in Figure 4.17. Here a large transient is induced on the bus - in excess of the 8 V limit of the instrument. By extrapolation, it can be deduced that the peak of this transient was around 12 V - 6 V. This is over the maximum electrical limit for the MCP2515 CAN controller. Over time, such transients would lead to unreliable operation of the controller, even when operating in quiet EM conditions, and other permanent damage. Furthermore, the interference is so severe that the differential-mode can be seen to break between $45 - 80\ \mu\text{s}$. If any essential control commands are sent during this period then, in a best-case scenario, the car will not respond. This suggests that the intra-vehicular interference is the cause of the unanticipated loss of control of the vehicle.

In order to confirm this theory, the loom was probed with the CAN network off and the inverter switched on. This would measure only the induced voltages. Furthermore, the timebase of the measurement was opened up in order to check the periodicity of the interference source. Figure 4.17 clearly shows the interference is time-dependent and not from a constant source, and so a wide timebase would allow for multiple transient events to be captured. An example of such events can be seen in Figure 4.18.

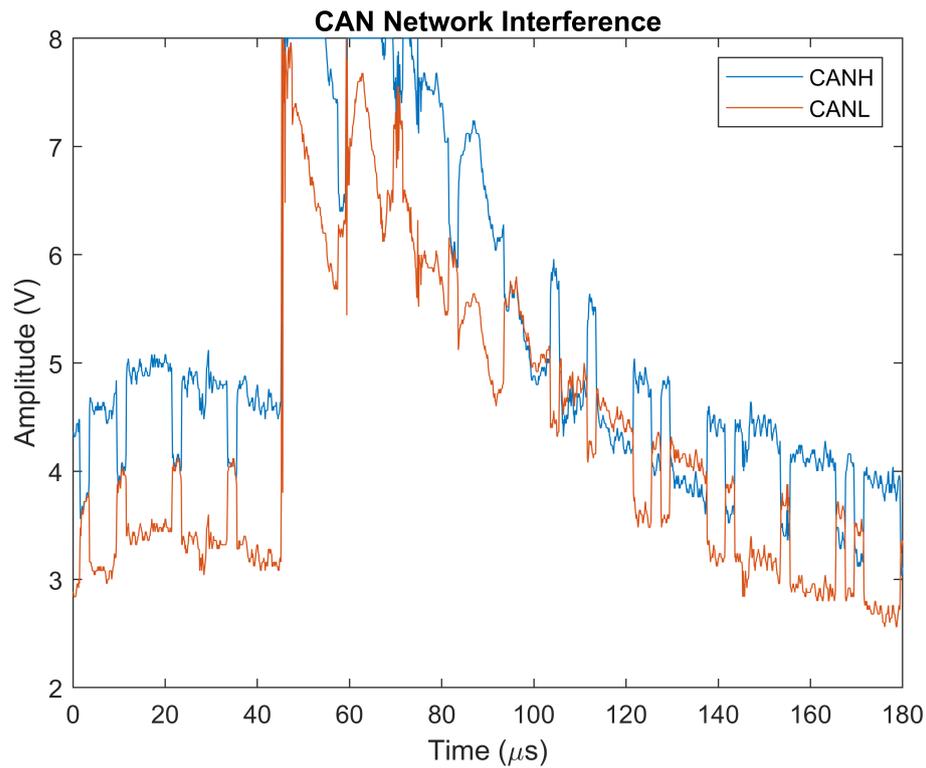


Figure 4.17: Severe induced common and differential-mode noise.

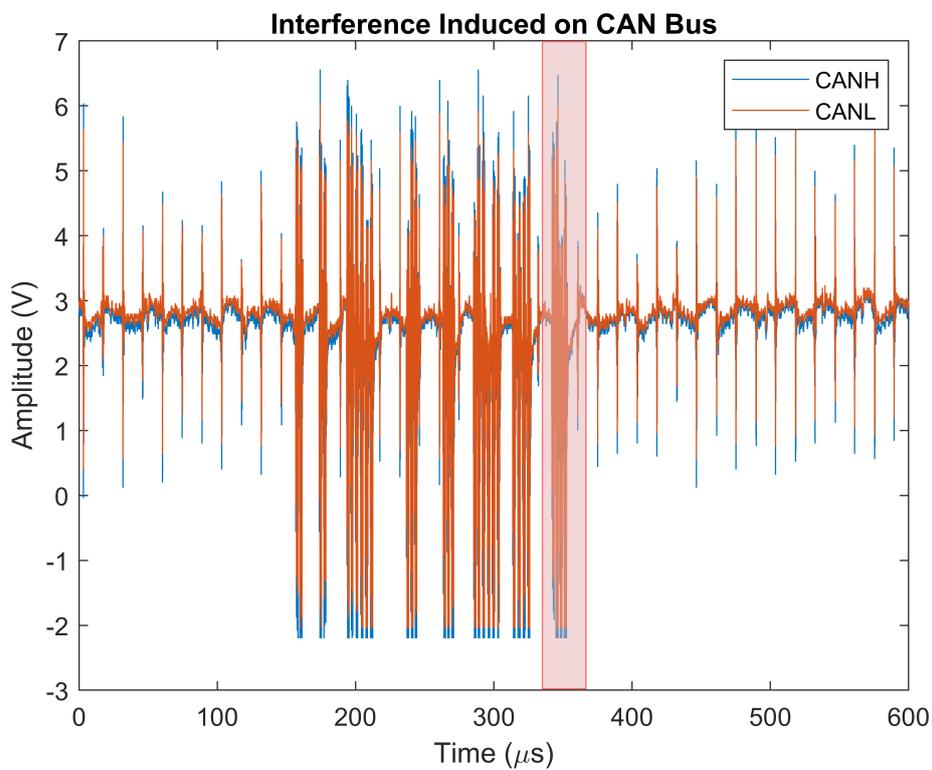


Figure 4.18: Induced voltage on an idle CAN bus.

Two main emission modes can be observed from this plot. The first, occurring every $15\ \mu\text{s}$ (67 kHz), appears to be highly repetitive in the temporal dimension, but varies considerably between 0 V and 6 V in amplitude.

The second mode appears to be much more destructive to the transmission of data on the CAN bus. Whilst timing of this second mode is highly irregular, they appear to occur in bursts of 5 - 10 transients, each separated by around $25\ \mu\text{s}$. With voltage swings between -2 V and 6 V, these bursts are far more likely to cause complete failure of the data arbitration at the CAN receiver. As the frequency content of each transient burst is, relative to the wide timebase, very high, Figure 4.19 provides closer detail of one of the transient bursts as indicated by the region highlighted in red in Figure 4.18.

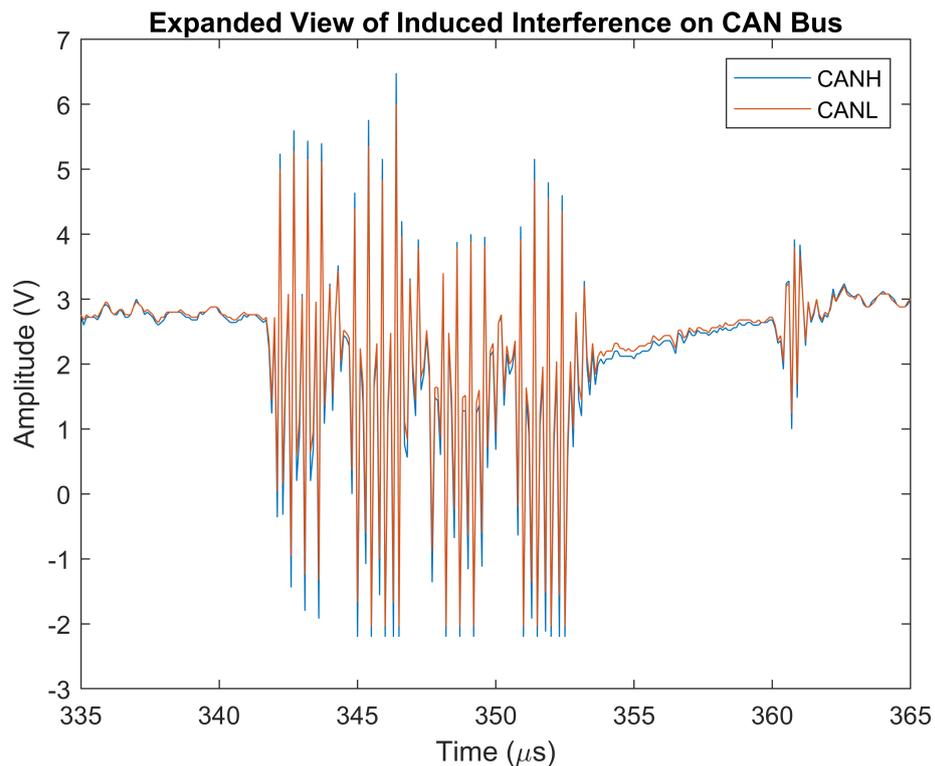


Figure 4.19: Expanded view of the region highlighted in red from Figure 4.18.

In this case, the larger transient can be seen to consist of four smaller bursts with a period of $0.8\ \mu\text{s}$ (1.25 MHz) - far faster than any other transient measured thus far. This closely matches a peak at 1.25 MHz in the frequency

domain measurement, which is of the highest amplitude in the boot and roof measurements (Figure 4.5). These two measurements in both the time and frequency domain corroborate exceedingly well to highlight the potential for damage to occur when EMI is not accounted for when designing electrical power and communication networks. In this case, the emissions from the additional autonomous driving sensors, computing and power equipment proved to have a great enough effect on the CAN bus that data was severely compromised.

It is also interesting to note the form of the interference which was induced on stubs of the CAN bus loom. Figure 4.17 shows a very suddenly transient edge, followed by a trailing curve and some ripple in the common-mode. This is highly characteristic of a waveform with high spectral content which reduces through time, and the resultant ringing voltages are indications of voltage reflections at impedance boundaries. Such a response is impossible to deduce through examination of the frequency-domain spectra alone, as all phase information is destroyed in a frequency-domain measurement. The measurements presented in this chapter highlight the importance of considering EMI in the time domain, just as much (if not more so) than in the frequency domain. However, as contemporary standards can trace their routes to the dawn of EMC engineering, when the primary concern was ensuring broadcasting radio stations would not interfere with one another, the consideration of time domain interference in many standards is generally limited in scope. Therefore, a greater examination of the effect of time domain interference on wired communication links would prove useful to better understand its potential impact upon networks.

4.4 EMC Recommendations

These measured examples of potentially dangerous EMI on the CAN bus provide a clear indication of the importance of ensuring that high bandwidth and high power systems are properly selected and designed around at a systems level. Whilst the frequency domain spectra indicate that very high levels of attenuation of EMI can be achieved within very short distances in the automotive environment, the same measurements also indicate that many frequency components are able to propagate throughout the car. Before this experiment was concluded the following suggestions for ensuring EMC were presented to the MLV group for their consideration.

1. Shielded, twisted-pair conductors should be used for all CAN stubs. This would offer an increased shielding effectiveness between emissions sources and the victim conductors. Furthermore, the use of twisted-pair conductors would ensure that, on average, each conductor is the same distance from an emission source, thereby increasing the common-mode rejection.
2. A common theme throughout measurements in both the time and frequency domains was the EM emission contributions from the inverter. This inverter was of a high power design, but suspicions arose regarding the quality of the device to provide electrically clean power. By investing in a supply from well-regarded manufactures of power supplies, the emissions from the noisy inverter could be easily reduced.
3. System components which do not require wireless transmission of data (such as power supplies, network cables, etc) are housed within metallic enclosures to provide better shielding. This would reduce emissions from noisy sources, as well as protect potential victims from the effects of EMI.

In this specific case, more fundamental EMC techniques were not considered appropriate for suggestion, due to the lack of access to many of the more fundamental operating characteristics of the system devices (such as slope control, data rate, modulation schemes, etc). Personal correspondence with members of the MLV group indicate that, since these techniques were implemented, the operational reliability of the vehicle has improved.

4.5 Chapter Conclusion

This chapter has demonstrated the high variability of the EM environment, even within small differences in distances from sources, as found in a typical car. Specifically, the measurements presented offer unique examination of the emission sources and variability of the autonomous cars of the future. Whilst previous works, such as [19] have described the EM environment using similar methods, these were performed on vehicles with internal combustion engines and limited digital components. This information is used later in the thesis to allow for realistic induced voltage levels to be used in a laboratory environment, permitting easy testing for EMC which translates effectively into real-world applications.

These measurements also highlight the importance of considering EMI in the time-domain. The voltages induced on the CAN bus were not from steady-state sources, though they were observed in frequency-domain measurements. This raises the question of whether or not the phase components of source components are critical to inducing errors on a system, which is covered in the next chapter.

Chapter 5

Radiated Interference on Wired Networks

The understanding of the EM, intra-vehicular environment afforded by the measurements in the previous chapter demonstrates the potential for the existing wired communication links to be grossly affected by EMI. The distinct impact of the time-domain interference on the CAN bus was particularly visible. Time-domain interference may be considered as interference which varies through time, such as EMI induced from power supply switching or crosstalk from neighbouring data network conductors. Traditionally, however, standards for EMI compliance verification have concentrated efforts on frequency-domain testing of systems - where the spectral content of an interfering source does not vary through time. This is evident from an examination of standards (which, with particular consideration given to the context of this thesis, have been chosen due to their utilisation in verifying the susceptibility of automotive systems to EMI) such as ISO11898 [62], CISPR 25 [76] and CISPR 12 [84], wherein power limits are prescribed for frequency bands under a given test regime.

This focus on frequency-domain-driven standards is a result of the more primitive communication methods used during the formative years of EM

communication. Once, the only electronic communication equipment encountered by the ordinary person was an AM or FM radio, transferring audio information. In such a case, standards which examine the susceptibility of such devices to the relatively steady-state spectral emissions of broadcast radio are entirely appropriate. However as a greater range of high-speed, time-varying systems became more ubiquitous, so does the need for standards which account for the time-domain variability of both the emissions themselves and the system's susceptibility. In other words, does the information pertaining to a signal component which is absent in the frequency-domain (such as phase) affect a system's response to EMI, or is it simply the signal's frequency-domain characteristics (frequency and power) which determine susceptibility? If the time-domain information is important, then to what degree?

Methods for measuring EMI in the time-domain have long been proposed, such as in [85] - published in 1972. More contemporary work, such as [86], have described time-domain measurement techniques, which have historically been limited in practical implementation by the speed of analogue-to-digital converters and dynamic range limitations. Since the advent of lower-cost, high-speed devices, time-domain sources have started to become accounted for in some EMC standards, such as [87], but these remain the minority and account for highly simplified time-domain waveforms such as pulses. The complex interference patterns observed in Chapter 4 highlight the need for a greater understanding of time-domain interference phenomena, especially as cars contain an increasing number of sensitive devices.

As a result of the time-domain interference observed in Chapter 4, several wired communication standards were examined under a carefully chosen time-domain-EMI source. This chapter details observations made about the magnitude of the impact that time-domain phenomena have on contemporary wired communication systems.

5.1 Inducing Bit Errors on Wired Networks under Time Domain EMI

Three of the wired communication protocols, described in Chapter 3 and found inside the Kia Soul EV from Chapter 4, were chosen to be subjected to time-domain interference: TTL, RS232 and CAN communication protocols. Purpose-built networks of all three protocols were built, which were simplified as far as possible to facilitate measurements of bit errors at a receiving module. This was realised by designing networks with a logically simple master/slave control scheme, with data movement entirely from the master to the slave node devices, as per Figure 5.1.

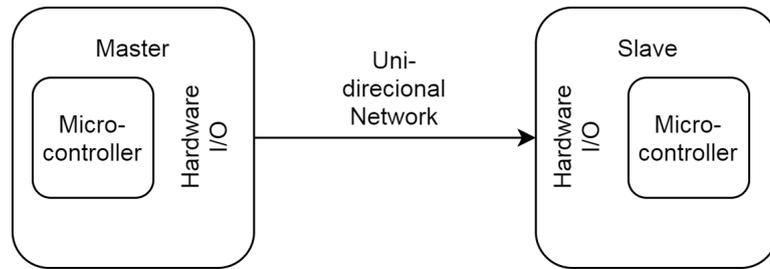


Figure 5.1: Hardware configuration of wired protocols under test.

Measuring the effect of EMI on such a simple network rules out any errors caused by the breakdown of higher-level OSI model systems, ranging from packet collisions (where multiple network nodes attempt to send packets simultaneously) to incorrect node contention resolution in the CAN system. Whilst important to consider these possible failure modes on a real-world network, it is beyond the scope of this thesis to consider events resulting from layer failures so far abstracted from the physical.

The experimental setup can be divided into two primary components: the RF and the communication network. The RF portion consisted of the equipment required to synthesise the desired interference waveform in a suitable physical environment for the network under test, whilst the communication network setup required a working communication link and

a means of measuring and processing the data bits transmitted through the network.

The components of the RF setup, as well as their interconnections, can be seen in Figure 5.2.

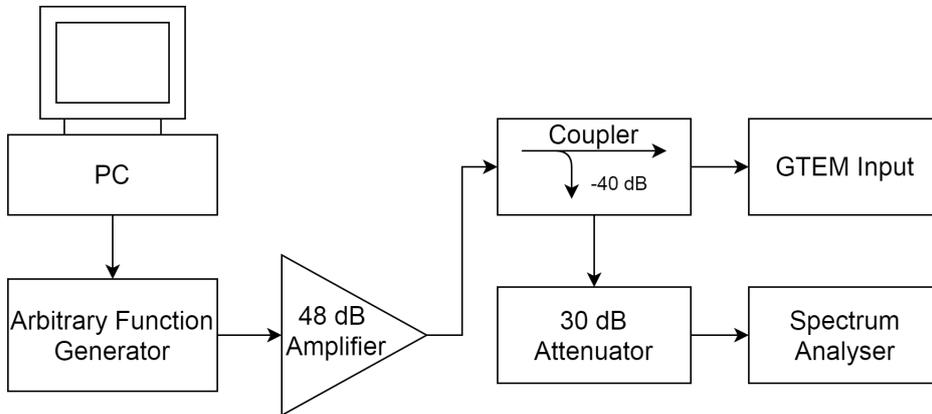


Figure 5.2: RF layout for synthesising interfering waveforms.

A PC was used to calculate the required the sample values for a given waveform using the BenchLink Waveform Builder Pro Software, created by Keysight Technologies [88]. These sample values were then sent through a General-Purpose Interface Bus (GPIB) to an Agilent 81150A arbitrary waveform generator, which synthesised the interfering waveform's shape. The output of the 81150A was then fed into a Prana DT70 amplifier. The amplified waveform was connected to an Amplifier Research DC3010A coupler, which allowed a -40 dB component of the amplified signal to be sampled by an Agilent N9344C spectrum analyser, allowing for examination of the amplified interference waveform in the frequency domain to ensure that no distortion of the signal occurred during the signal's synthesis or amplification. Whilst the use of a spectrum analyser to observe waveforms in the frequency domain would not permit analysis of signal in the time-domain this would potentially highlight inadequacies in the heavily frequency-domain approach of many contemporary standards. The N9344C was fitted with an Anritsu 42N50A-30 30 dB attenuator to reduce any

potentially harmful RF power from damaging the spectrum analyser's input. With a measured response flatness of less than ± 0.5 dB across the attenuator's response between 0-1 GHz, and a coupling flatness of ± 0.6 dB from the coupler's response, such sampling of the actual interference waveform would be sufficiently accurate for the purposes of this experiment as the level measurement uncertainty of the N9344C itself is specified as ± 1.3 dB in the 1 MHz - 7 GHz band [89].

A Gigahertz Transverse Electromagnetic (GTEM) cell was chosen to expose the network under test to an even, time-varying electric field. A GTEM cell, as shown in Figure 5.3, is a pyramidal chamber of metallic construction which is used for radiated EM emissions onto victim devices within a given working volume [90].

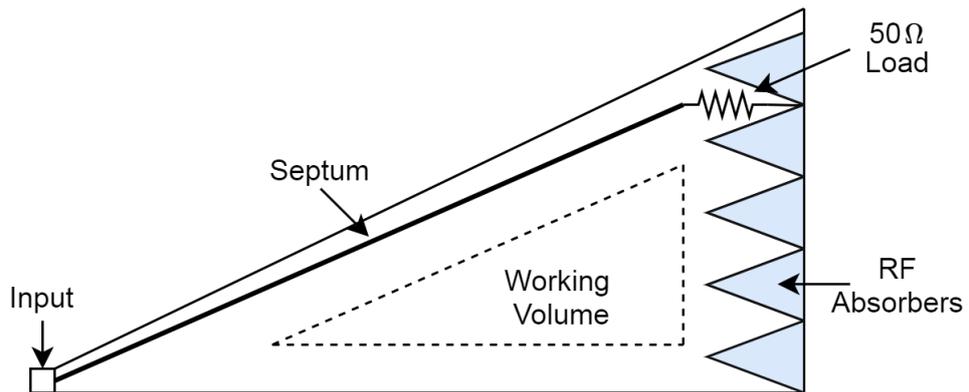


Figure 5.3: Diagram of GTEM construction.

The GTEM can be thought of as an expanded coaxial transmission line, with the offset septum acting as the inner conductor, surrounded by a grounded enclosure. Signals applied to the inner conductor will generate a vertical E-field above and below the septum. A working volume, measured and defined by the manufacturer of the GTEM, is created in the space below the septum where a statistically even E-field is created. The lower septum height in the GTEM results in a stronger E-field for the same given input power, but also decreases the working volume. The means of calculating the average E-field strength within this working volume of the GTEM is

given by (5.1), as per [90].

$$E = \frac{V}{h} \quad (5.1)$$

Where E is the electric field strength, V is the voltage on the septum and h is the height above the GTEM floor. The voltage may be calculated from (5.2).

$$V = \sqrt{Z_0 P} \quad (5.2)$$

Where Z_0 is the characteristic impedance of the cell (50Ω) and P_{in} is the input power to the GTEM cell. (5.1) and (5.2) may be combined to create a convenient form as per (5.3).

$$E = \frac{\sqrt{Z_0 P}}{h} \quad (5.3)$$

A hybrid termination, consisting of a 50Ω load and pyramidal RF absorbers is used to absorb all RF currents and EM fields at the end of the cell. Pyramidal absorbers are particularly well suited to absorbing EM fields at frequencies above 100 MHz, as reflections from incident waves cancel out, with higher frequency waves cancelling out further along the pyramidal structure - in effect creating a continuously-varying impedance through the structure's primary axis. For the pyramidal absorber to be effective, each pyramid must be at least one half wavelength in length, though as the wavelength of a given field is shorter in the pyramidal material, as given by [91] it may be expressed as per (5.4).

$$\lambda_r = \frac{1}{\sqrt{\epsilon_r}} \quad (5.4)$$

Where λ_r is the wavelength inside the media, and ϵ_r is the relative permittivity.

For absorbing frequencies lower than 100 MHz, ferrite tiles are used. These tiles rely on tuning the thickness of the tile so that reflected waves destructively act with incident waves, forming resonances within the material which appear as tunable nulls in the return loss response [92].

The communication network measurement component of the experiment can be seen in Figure 5.4.

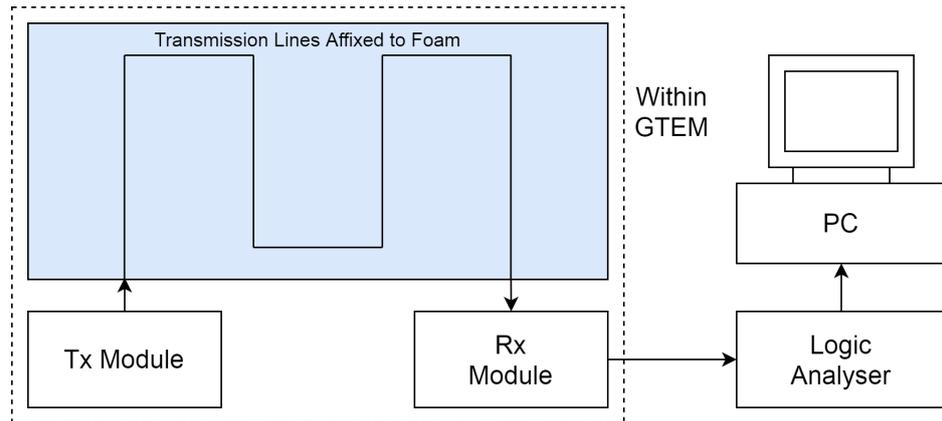


Figure 5.4: Communication network measurement setup.

Within the GTEM cell, two modules are used to act as a master transmitter and a slave receiver. Between each device, transmission lines consisting of multi-strand, copper wire was affixed to extruded polystyrene foam to allow the transmission line to be fully illuminated by the E-field, allowing the electric field to couple onto the network. A Saleae Logic 8 logic analyser was used to capture the physical data bits at the receiver. the use of a logic analyser permitted the gathering of large quantities of data in real time on a PC which, in turn, facilitated the processing of data with relative ease.

On the PC, a Python script was written to import the captured data from the software communicating with the logic analyser, and compare it against the known value of each transmitted bit. A ratio of the number of incorrectly received bits to the total number of transmitted bits could then be calculated. This ratio, known as the Bit Error Ratio (BER), is a unitless measure of the ability of a system to transmit bits correctly and can be

calculated using (5.5).

$$BER = \frac{N_{err}}{N_{bits}} \quad (5.5)$$

Where N_{err} is the number of incorrectly transmitted bits, and N_{bits} is the total number of transmitted bits.

The entire experimental setup can be seen in Figure 5.5, with both the RF and communication network components in place for a measurement. An inset image, bordered in red, shows the interior of the GTEM cell. The pink wiring mounted to the closed-cell foam can be seen, as well as a receiver module for RS232 measurements. Behind the closed-cell foam mount the pyramidal absorbers are visible, as well as a fraction of the septum above the working space.



Figure 5.5: Complete experimental setup for GTEM measurements.

The PC in view acts as both the source of sample data for the arbitrary function generator, and as storage for the data bit measurements from the logic analyser.

Initial measurements on all three networks were performed using a standard sine wave to record the level at which a conventional frequency-domain interference source would begin to induce errors on each network. With

each network operating at 1 Mbps, and the fundamental frequency of the interfering waveform at 1 MHz, it was found that for E-field strengths between 3.4 Vm^{-1} and 34.3 Vm^{-1} the sinusoidal interfering waveform was incapable of generating errors. A first-order differential Gaussian waveform, chosen as an easy to replicate form of time-domain interference, was then tested on the communication networks, but again no errors were recorded. Finally, a square wave was used as the interference source, which resulted in non-zero BERs for the TTL and CAN protocols, as recorded in Table 5.1. These error values were recorded for an E-field strength of 21.7 Vm^{-1} over 23 million bit samples.

Table 5.1: BERs for a square wave interference source

Protocol	BER
TTL	0.164
RS232	0
CAN	0.017

A time-domain measurement of the induced interference on the conductors allows for an examination of the real effect the interference sources each have on the networks, and indicate how such interference may be interpreted by receivers. Figure 5.6 indicates that the square-wave interference source can generate much higher levels of interference on the underlying data signal than the sinusoidal and first-order differential Gaussian waveforms, even when the sources are at the same power. The sinusoidal waveform is capable of inducing voltage swings of only a few millivolts on the data signal which is far from the voltage swing required to erroneously arbitrate the data bit, as discussed in Chapter 3. The first-order differential Gaussian is better at inducing voltage shifts, but a short term interference of 1 V is evidently not enough to mislead the receiver. However, the square-wave interference source induces huge voltage swings of $\pm 3 \text{ V}$, with subsequent ringing occurring for nearly the entire bit duration.

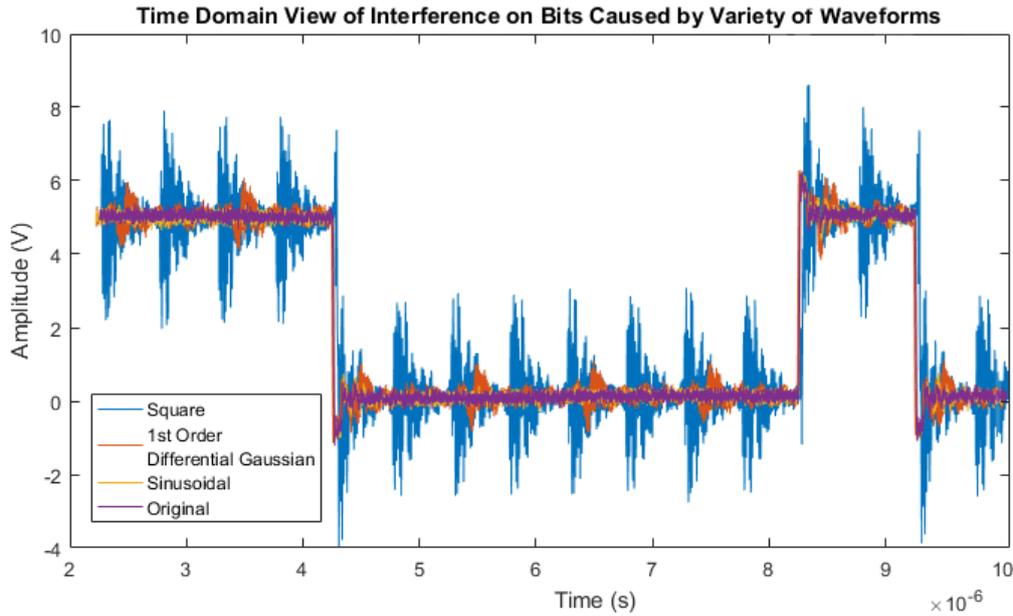


Figure 5.6: Time-domain measurement on the conductor of a 5 V TTL compatible communication network.

This is, of itself, not entirely surprising. Literature is full of design methodologies for reducing the potential for EMI from sources with high $\frac{dV}{dt}$ or $\frac{dI}{dt}$ characteristics [93, 94, 95, 96] by reducing the rise times of the voltages and currents in everything from digital communication networks to switching power converters. Typical design methodology in these circumstances is to decrease the $\frac{d}{dt}$ of the source as far as applicably possible for the device (whether it is a power converter or communication link) to still function, whilst causing as close to the required BER on the victim network as possible. This seems to imply that there is some boundary between a zero BER and a non-zero BER, but this is really only accounted for in standards by considering the induced interference levels by varying power. The square-wave interference case proved to be useful for examining the time-domain BER boundary - that is, where a source (for the same amplitude) will go between generating no errors and some errors. This was primarily due to the square wave's convenient form when expressed as a frequency domain representation via the Fourier transform.

Consider a square-wave source of time-domain interference, which is expressed as per (5.6).

$$V(t) = \begin{cases} 0, & -\pi \leq t < 0 \\ 1, & 0 \leq t < \pi \end{cases} \quad (5.6)$$

Where V is the voltage of the signal, t is time, and T is the period of the waveform. The Fourier series representation in (5.7) assumes that the series converges and is a continuous function over the interval 0 to T [97].

$$V(t) = a_0 + \sum_{n=1}^{\infty} (a_n \cos(nt) + b_n \sin(nt)), \quad -\pi \leq t < \pi \quad (5.7)$$

The coefficients a_0 , a_n and b_n must then be calculated for the given function. In this case of this square wave, these coefficients can be found as per (5.8)–(5.10)

$$a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} V(t) dt \quad (5.8)$$

$$a_n = \frac{1}{\pi} \int_{-\pi}^{\pi} V(t) \cos(nt) dt, \quad n = 1, 2, 3, \dots \quad (5.9)$$

$$b_n = \frac{1}{\pi} \int_{-\pi}^{\pi} V(t) \sin(nt) dt, \quad n = 1, 2, 3, \dots \quad (5.10)$$

In this instance, the function is odd, and so all a_n values are zero. Meanwhile, all b_n coefficients can be calculated using (5.11). If the waveform was shifted in time by $+\frac{T}{2}$ then the even function would simply swap coefficient values.

$$b_n = \frac{2}{n\pi}, \quad \text{for all odd } n \quad (5.11)$$

It is therefore relatively simple to create a time-domain form of interference that can be reduced to a single component (a sine wave as used for frequency-domain testing) by varying the number of Fourier terms, n , as

illustrated in Figure 5.7.

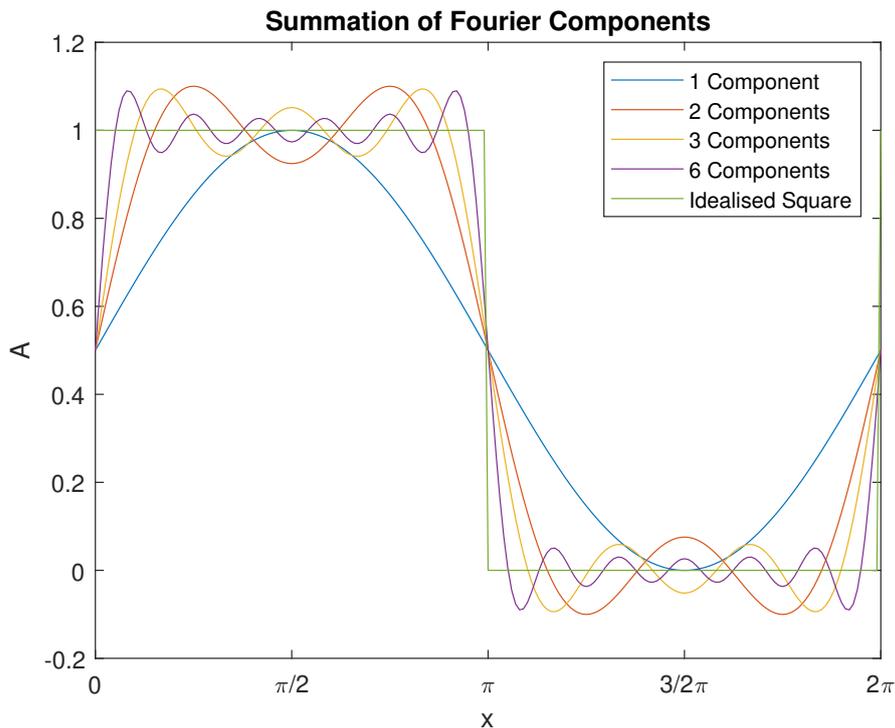


Figure 5.7: Effect of increases number of Fourier components of a square wave.

As the number of components is increased, so too does the $\frac{dV}{dt}$ of the edges of the square wave. This progressive series of waveforms can be expressed by (5.12).

$$x(t) = \frac{4}{\pi} \sum_{n=1,3,5\dots}^m \frac{\sin(2\pi nt)}{n} \quad (5.12)$$

Measurements of the BER of each system was taken using this series, up to $n = 59$. However measurements were performed only on every fifth value for n , due to the lengthy procedure for measuring and calculating the BER, which had to be controlled manually. Each measurement would take around 20 minutes to fully process.

Each interference waveform was examined in the frequency domain, and as the example in Figure 5.8 shows, each waveform was verified to have the correct harmonic components at the correct amplitude. This measurement

was performed using a linear sweep, and a resolution bandwidth of 1 kHz to ensure that spurious tones or even harmonic distortions could be observed down to -75 dBm.

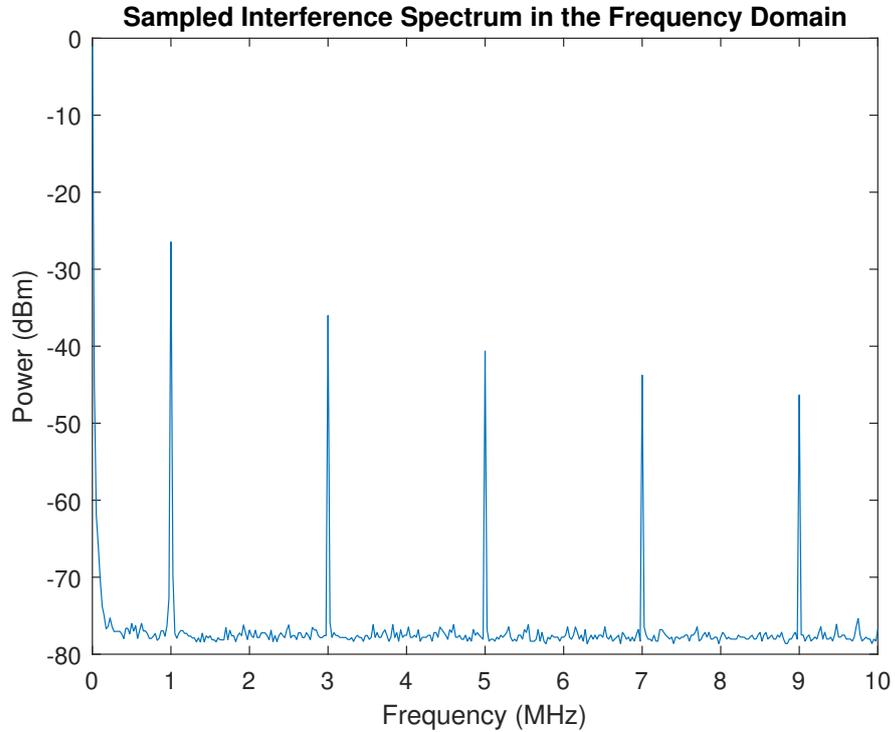


Figure 5.8: Frequency-domain sample of interference signal.

As shown in Table 5.2 the harmonic amplitudes matched the theoretical values with a high degree of accuracy, indicating low distortion of the time domain waveform in the RF signal conditioning setup.

Table 5.2: Harmonic verification of interfering waveform.

Harmonic (MHz)	Harmonic Amplitude (dBm)		
	Theoretical	Measured	Difference
1	-26.4	-26.5	-0.1
3	-35.9	-36.0	-0.1
5	-40.3	-40.6	-0.3
7	-43.3	-43.6	-0.3
9	-45.5	-45.8	-0.3

It was found that the relationship between the number of Fourier terms and the BER of both the CAN and TTL networks was highly non-linear, as shown by the plot in Figure 5.9.

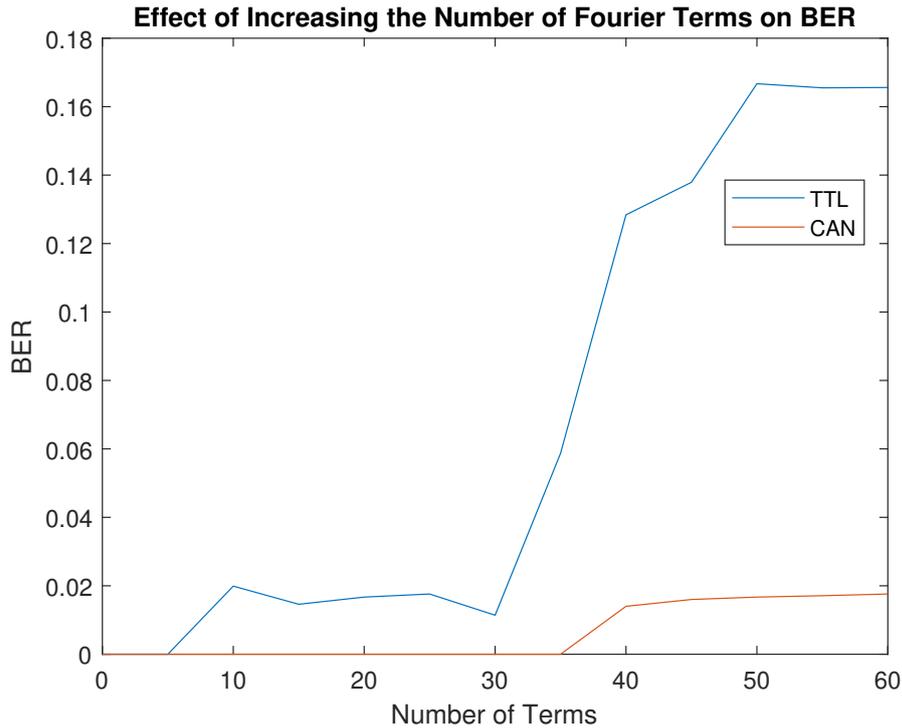


Figure 5.9: Effect of increasing the number of Fourier terms on TTL and CAN systems in a GTEM cell.

In the case of the TTL network, errors were first observed at 10 terms. The BER then appeared to reach a steady-state condition, fluctuating by no more than 3×10^{-3} . However, between the measurements in the 30-50 term region, the BER increases by more than 0.11 before settling to a value around 0.163. This BER value, found at above 50 terms, closely matches the BER of 0.164 measured using the square-wave output of the 81150A indicating that this value was the real steady-state BER as a function of the number of Fourier terms.

The CAN system proved to be much more resilient to the interference waveform, primarily due to the differential mode signalling. A BER of 0.014 was measured at 40 terms, before rapidly stabilising around the 0.17 BER of the 'ideal' square wave interference pattern, with a response unlike that of the common-mode TTL network.

These results indicate that, at least in cases with common-mode commu-

nication networks, the actual effect of an increasing $\frac{dV}{dt}$ or $\frac{dI}{dt}$ does not always increase the BER of a system. Whilst such an increase in the gradient of a rising or falling edge of a power converter or data signal may result in great analogue interference on a victim network, it is evident that there is not necessarily a linear translation between interference in the analogue and interference in the digital domains of such communication systems. Indeed, the results in Figure 5.9 indicate that as the number of terms increases, the BER may actually *decrease*.

Such findings indicate that it is not entirely accurate to make the claim that slopes must be lengthened to reduce interference on victim devices, providing that the interference is considered at the more abstracted, *digital* level instead of the analogue. If victim devices operate using some form of transform between analogue and digital representations of information, then this data suggests that in some cases, it may be more appropriate to increase the slope of a waveform to achieve an acceptable BER in neighbouring systems. This would have implications ranging from more efficient switching in power converters to increasing data transmission rates in systems where the period of the rising and falling edges of bits is a relatively high proportion of the total bit period.

Whilst the networks under test were deliberately designed to be as simple as possible (in terms of hardware) they still utilised basic level 2 operations, such as framing, to transport data through the network. In order to investigate this further, a greater level of simplification was required, to identify the causes of this BER response.

5.2 Induced Errors on a Level 1 System

By reducing the network under test to a simple level 1 system, the data flow on the network resembles a continuous stream of data with no framing. As

such, any interference on the network will manifest during a bit transmission. With no framing bits, any interference on bits on the network is also guaranteed to occur during a data bit. By using the BER as a measure of the effect of any EMI on the network, a high level of abstraction is achieved when compared to contemporary EMC standards that use power or voltage levels to measure and classify the extent of potential interference on a network. At the same time, this approach is less abstracted than measuring the effect on a system with level 2 or 3 components, which facilitates an examination of the effect of EMI on the receiver of a device.

Steps were taken to create an experimental setup that would better resemble a real-world scenario. Instead of using a GTEM cell to expose a single data network to even E-field interference, a wiring setup was created to facilitate radiated emissions from a source conductor onto a victim conductor, as shown in Figure 5.10 and Figure 5.11. Such a setup provides a more realistic network scenario with two close parallel conductors whilst also providing a well-defined case for replicating and simulating.

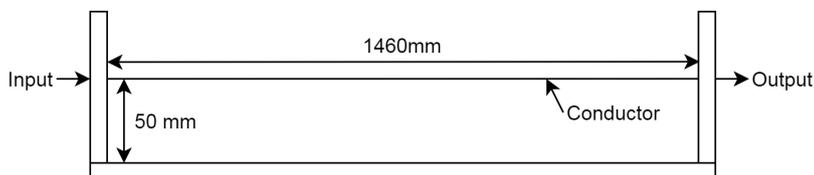


Figure 5.10: Side view of wiring loom (not to scale).

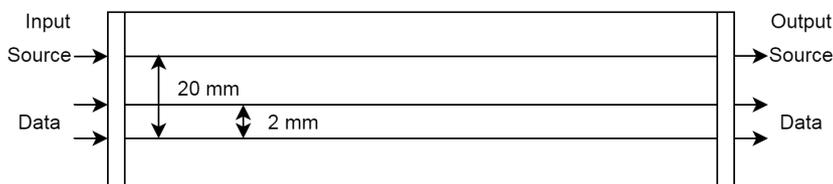


Figure 5.11: Above view of wiring loom (not to scale).

This loom was placed inside an anechoic chamber to minimise an emissions from the setup affecting external equipment, as well as to mitigate the effect of external noise sources (such as telecommunication systems,

lighting, and other experiments) on the experiment. An anechoic chamber is a fully enclosed room, lined with RF absorbent materials. Typically, the room is made of metal with a single, RF-gasket lined door. A combination of ferrite and pyramidal absorbers may be used inside the chamber to reduce reflections within the working volume. Where all surfaces of the room are covered in RF absorbent materials, the chamber is considered fully-anechoic. However, when test cases do not require attenuation of reflections in all directions a surface may be left uncovered, resulting in a semi-anechoic chamber. This is typically the case when large or heavy equipment must be used inside the chamber, which would otherwise crush or deform any materials on the floor [98].

A bit-error-ratio tester (BERT) was used as a means of measuring the BER of the victim network. Such a device acts as both a transmitter and receiver of a continuous flow of data bits, with a defined threshold value arbitrating between whether the receiver arbitrates a bit to be a high or low. By comparing the digital representation of each bit with the transmitted representation, the real effect of interference on the successful transport of data can be measured. Initially, problems with acquiring a suitable BERT forced work into developing an open-source instrument, which is discussed further in Chapter 6. An Anritsu MP8302A was later used to perform the majority of measurements found in this work, due to the new availability of the instrument for out-of-the-box measurement capability. With an output frequency between 1 kHz to 20 MHz, and adjustable threshold voltage selection, this instrument was well-suited for measuring the effect of radiated, time-domain EMI from neighbouring sources on 1 Mbps systems [99].

Both the anechoic chamber and BERT were used in conjunction with the source synthesis equipment from the experiment in Section 5.1 to create the experimental setup shown in Figure 5.12.

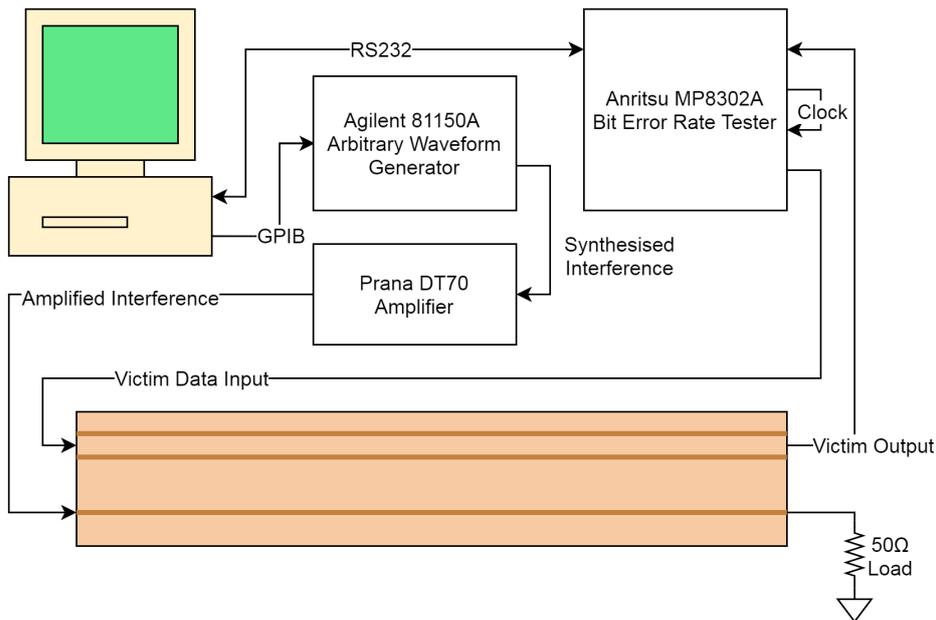


Figure 5.12: Experimental setup for measuring the BER of a network exposed to time-domain EMI.

This setup makes use of the same synthesis procedure for generating the same series of scaling, time-domain interference waveforms from (5.12). However instead of feeding into a GTEM cell, with a 50Ω load between septum and rear wall, the emission source conductor is simply terminated with a 50Ω load to local ground.

The BERT can be seen to have not only connections for transmitting and receiving data bits, but also for a clock signal. This clock is used to determine the sampling point of the BERT. As a point of reference in this experiment, the sampling point was chosen to be at half of the bit duration, and an appropriate length of cable was used to account for the electrical delay of the data path through the victim conductor.

A photo of this experimental setup can be seen in Figure 5.13. Note the PC is outside of the chamber, and that this particular photo includes a spectrum analyser for verification of the interfering waveform.

A comprehensive plot of the resulting data can be found in the heat map plot in Figure 5.14 and in the multi-line graph in Figure 5.15. In this



Figure 5.13: Photograph of the interior of the anechoic chamber for crosstalk measurements.

experimental run, a 1 MHz data signal was subjected to a sweep through the interference series from (5.12), as well as through multiple square-wave amplitudes. Note, in the case of Figure 5.15, the lines are in amplitude order from 0.63 V to 19.775 V. The threshold voltage for the TTL data was set to 1.4V and a termination impedance of $1\text{ M}\Omega$ was chosen for this initial test to facilitate the generation of errors on the network.

Each plot contains the same data but provides differing perspectives on the captured data. Both plots indicate that the phenomena measured in the results from Figure 5.9 is at least partially caused through physical layer mechanisms, and is not entirely derived from interference timing in the inter-frame period. In the case of Figure 5.14 this is evident from the colour oscillations along a given row. An initial high value BER in the 1-3 term region quickly decays to some lower value across both the amplitude and term axis, resulting in a low BER trench sweeping through the 6-10 term region.

The line plot in Figure 5.15 allows for a more detailed examination at specific BER values. At higher interference amplitudes, a clear peak of BER can be seen at a low number of Fourier terms. This provides a stronger

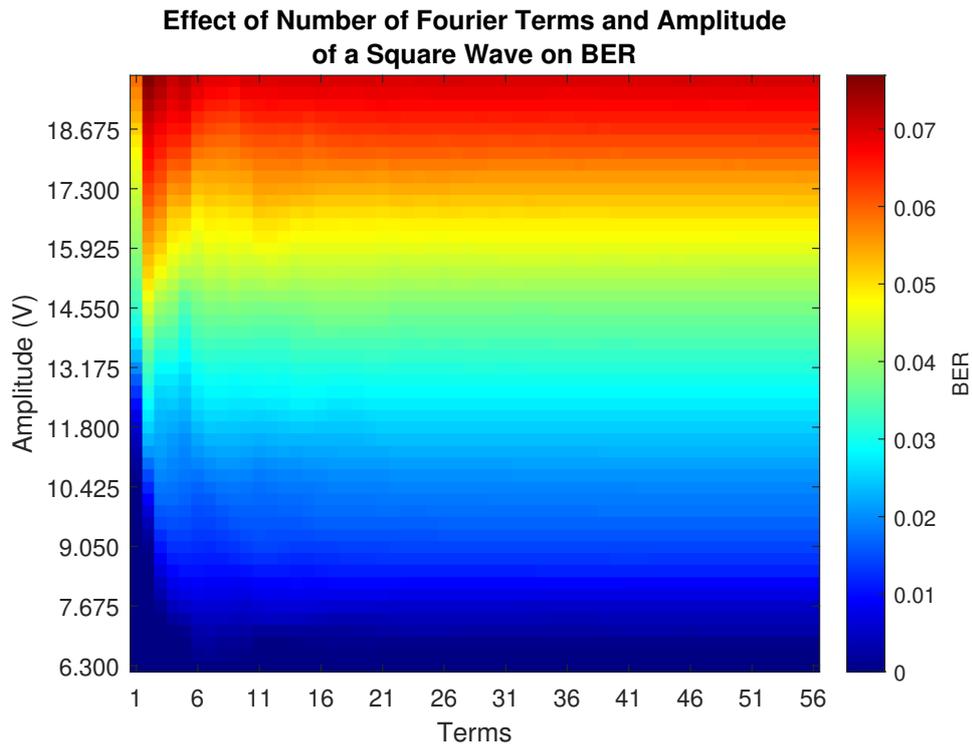


Figure 5.14: Heat map of BER for systems with $1\text{ M}\Omega$ impedance.

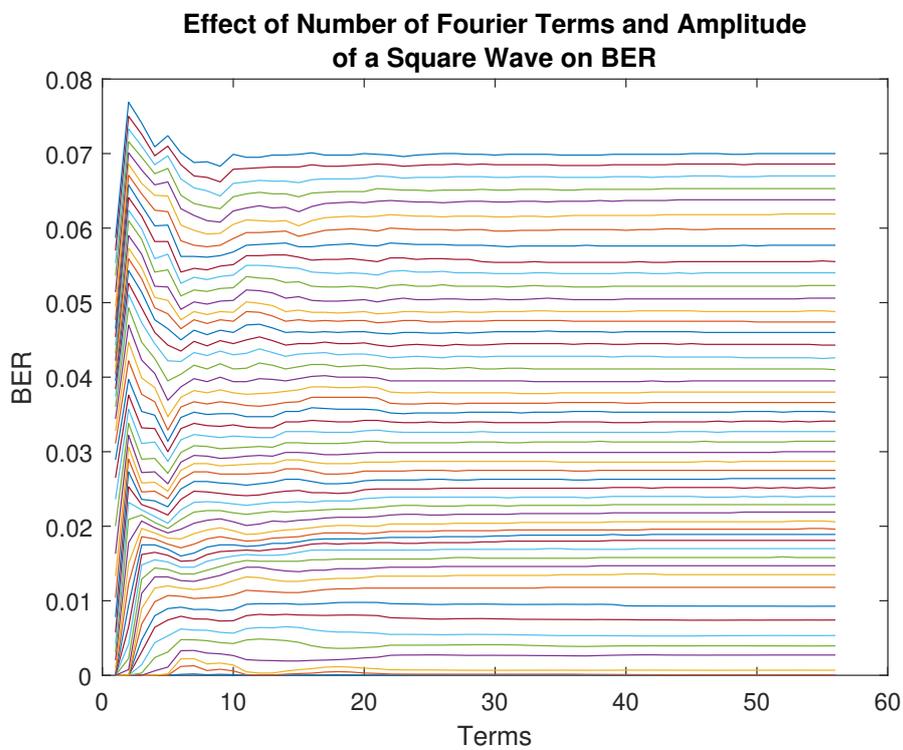


Figure 5.15: Linear plot of BER against number of Fourier terms. Note, lines are naturally sorted by interference amplitude.

impetus to investigate this phenomena further, as such a result indicates that it is possible to increase the slope of an interfering waveform and *increase* the number of errors on a victim network. As already discussed, in typical scenarios where such adjustment of slopes is performed, this increasing the slope (emulated by decreasing number of Fourier terms) is done on systems that would normally benefit from smaller slope periods. It is possible that doing so blindly may result in more errors on the victim network as well as lower performance from the source - whether that is lower switching efficiency or lower data throughput.

This is particularly evident at lower interference amplitudes, as shown from the logarithmic graph in Figure 5.16.

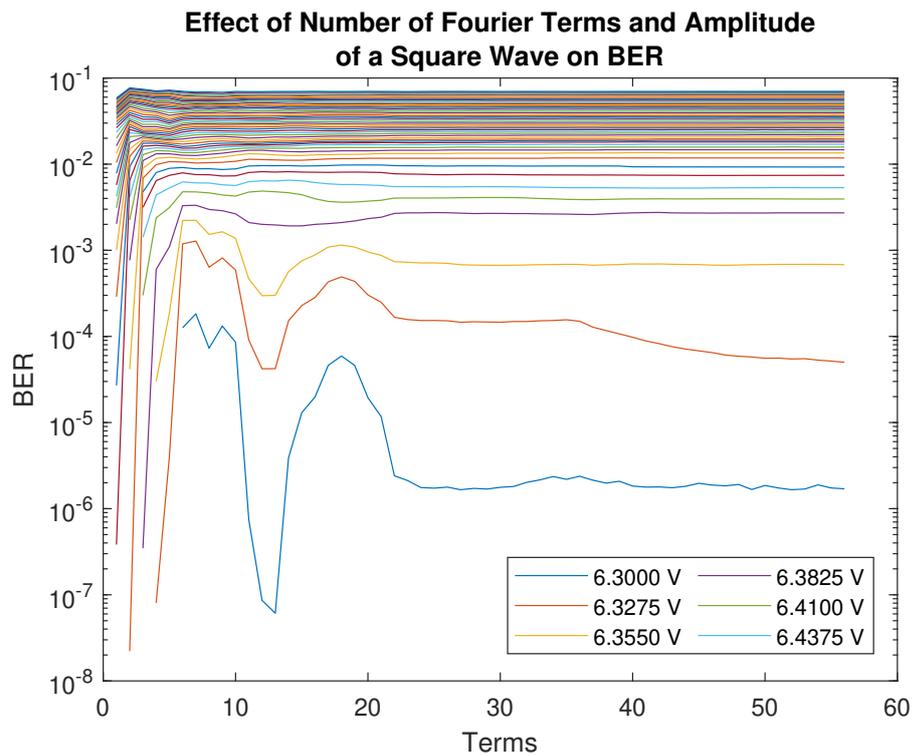


Figure 5.16: Logarithm plot of BER against number of Fourier terms, showing detail of BER for lower interference voltages.

This peak-oscillation behaviour can be seen to increase dramatically as the interference amplitude decreases, where peak to trough differences may be as great as 2×10^{-4} , with much greater variation in secondary and

tertiary peak BER values.

This new data has the potential to change how the real effects of an interfering source slope are considered and designed in systems. Distinct regions can be observed in the BER plots, which may be used in the EMC design process. These are:

- Where no errors are generated (for low interference amplitudes).
- Where the BER may vary drastically between a range of Fourier terms.
- Where the BER converges to some steady-state value after a given number of Fourier terms.

Further tests proved that this was not a special network configuration case, with an example in Appendix B showing another series of plots with data measured from a different network configuration with even more noticeable BER variations.

5.3 Modelling Distinct BER Regions

In order to begin to understand the reasoning for this BER variation phenomena, a computer-based model was developed to investigate the effect of adjusting circuit parameters on the BER. Whilst measurements on the wiring loom provided a real-world indication of the effect of time-domain interference on the victim network, it was limited by two primary characteristics. The first was the inability to adjust potentially important circuit elements such as receiver characteristics (sampling point, duration, etc). Secondly, the time period required for a single 3D plot of BER ranged from days to weeks depending on the number of required measurements. This was despite the use of complete automation in the control of instrumentation and in data recording, which allowed for continuous operation.

Work, inspired by that undertaken by C. Smartt at the George Green Institute for Electromagnetic Research (GGIEMR), began on developing a means of creating models of the network as well as simulating the coupled interference on the victim and generating a reasonable approximation of the BER. This was done through a four-step procedure:

1. Generate a model of the wiring loom.
2. Derive a transfer function between source and victim on the wiring loom.
3. Use the transfer function to calculate the interference waveform on the victim conductor.
4. Calculate the probability of a bit error using the interference waveform.

Models of the wiring loom were generated using the State-of-the-Art Cable Models for SPICE (SACAMOS), a software for creating sub-circuits for the Simulation Program with Integrated Circuit Emphasis (SPICE) analogue circuit simulator [100]. SACAMOS was chosen as it provides an intuitive user interface for quickly generating cable models through the use of a number of techniques, such as domain decomposition, modal decomposition and the method of characteristics. Further discussion of these techniques is beyond the scope of this work, though more information can be found in [101]. Furthermore, SACAMOS was familiar to the author as they were involved in generating test cases for the verification of the models, utilising CAN transceivers as termination circuits, as can be found in [102] and [103].

Figure 5.17 shows a typical SACAMOS use case, where a conductor's physical properties are being defined. Multiple cable models are then com-

binned into a bundle, which represents a 2D slice of the loom. This bundle is then converted into a model by specifying the length of the loom.

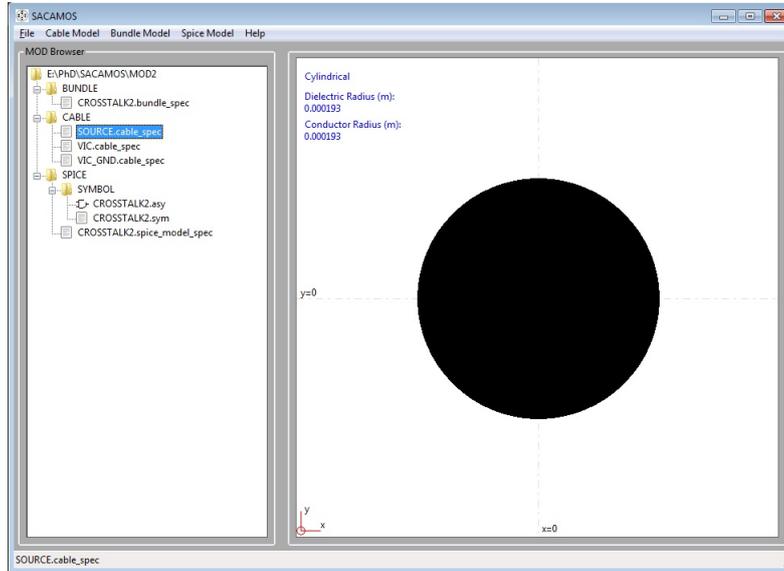


Figure 5.17: SACAMOS software for creating the loom model.

The output from SACAMOS can then be used to define the behaviour of a single sub-circuit in the SPICE simulation software. LTspice was chosen to perform the simulation due to it being a well-maintained package, as well as being available for no cost. The test circuit for this simulation can be seen in Figure 5.18.

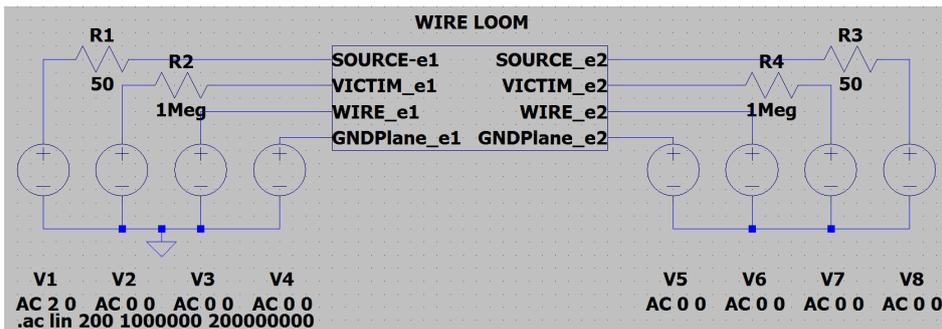


Figure 5.18: SACAMOS-generated model incorporated into LTspice.

In this instance a source conductor is driven by an AC source of $50\ \Omega$ impedance, representing the arbitrary function generator. The opposite end of the conductor is terminated with a $50\ \Omega$ load. The victim conductor has a similar configuration, but only has source and load impedances of

1 M Ω connected. The remaining AC sources (V2-V8) are connected for permitting the convenient adjustment of circuit behaviour, but in this test case are all set to zero volts and have no effect on the simulation.

When the simulation is started, the SPICE directive (denoted by the .ac command) performs a linear sweep of the output frequency of the V1 source from 1 MHz to 200 MHz of 200 points. The voltage at the VICTIM_e2 node is measured for each frequency and is used to formulate the complex transfer function of the circuit between source and victim. This is, in effect, a measurement of the complex S_{21} parameter.

With this in mind, the data from LTspice was exported to Matrix Laboratory (MATLAB) and a comparison was performed between the simulated and measured transfer functions which can be seen in Figure 5.19.

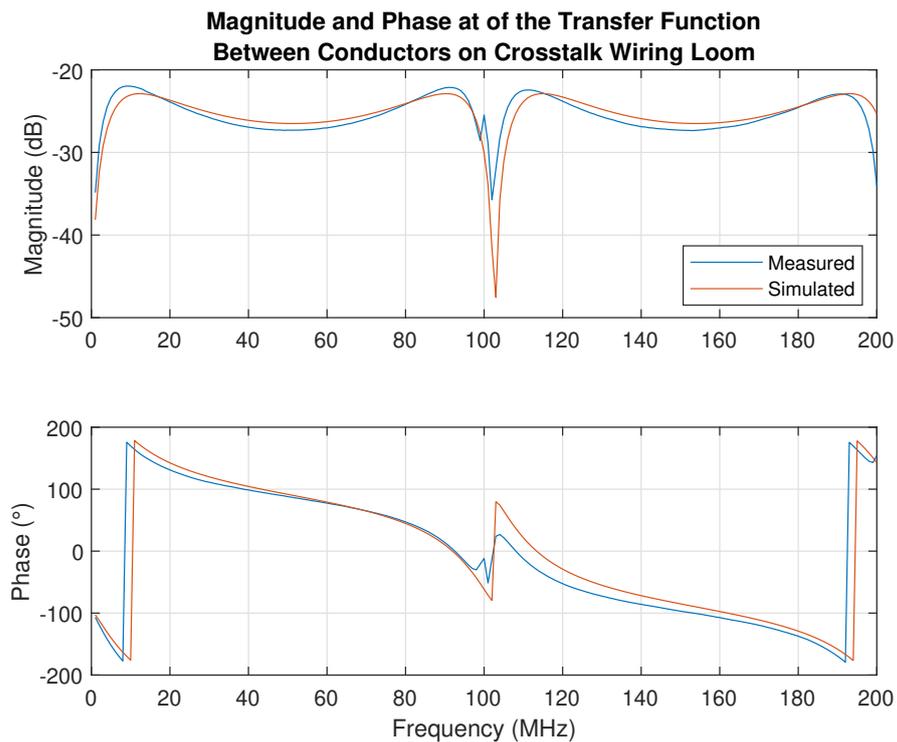


Figure 5.19: Comparison of the simulated and measured complex transfer functions.

Both the phase and magnitude response of the simulated output closely resemble measured values across the frequency range, though with some

small differences. The measured magnitude response can be seen to have higher peak values, though by only one or two dB. The simulated response also has a much deeper null of -48 dB at 103 MHz, when compared to the -36 dB at 102 MHz of the measured response. This is primarily due to subtle differences in the physical dimensions of the actual wiring loom, which the sub-circuit SPICE model does not account for. An example would be sag in the conductors - or a variation in height above the ground plane as a function of distance along the conductor. The wiring loom was fabricated to minimise such effects, but some differences inevitably remain.

This satisfactory simulation of the real-world complex transfer function was then processed further by MATLAB to generate a time-domain representation of the coupled waveform. A MATLAB script was written which first generated the required scaling coefficients for an odd square wave, as explained in Section 5.1. This Fourier term scaling component is therefore written as (5.13).

$$Noise_{input}(n) = -\frac{j^4}{\pi n} \quad (5.13)$$

Where n is the term number and j is the imaginary number. Each of these input terms were then scaled by the complex transfer function from the simulation in LTspice as per (5.14).

$$Noise_{output}(n) = H(n) \cdot Noise_{input}(n) \quad (5.14)$$

Where $H(n)$ is the complex transfer function at the n^{th} term. Note that it is important to ensure that the measurement points in the LTspice simulation align with the calculated component terms in the MATLAB script to ensure accuracy in the final output.

Through the use of an inverse discrete Fourier transform and some amplitude scaling factor, (5.15) shows how this complex components of the noise

output value was used to create the interfering waveform as a function of time.

$$Noise(t) = A \cdot \sum_{n=1}^{nterms} \Re(Noise_{output}) \cos(\omega_0 nt) - \Im(Noise_{output}) \sin(\omega_0 nt) \quad (5.15)$$

Where A is the amplitude scaling factor, $nterms$ is the total number of terms, ω_0 is the fundamental frequency of the waveform and t is time.

Figure 5.20 demonstrates the reasonable accuracy the simulated induced interference waveform has, when compared to a measured interference waveform.

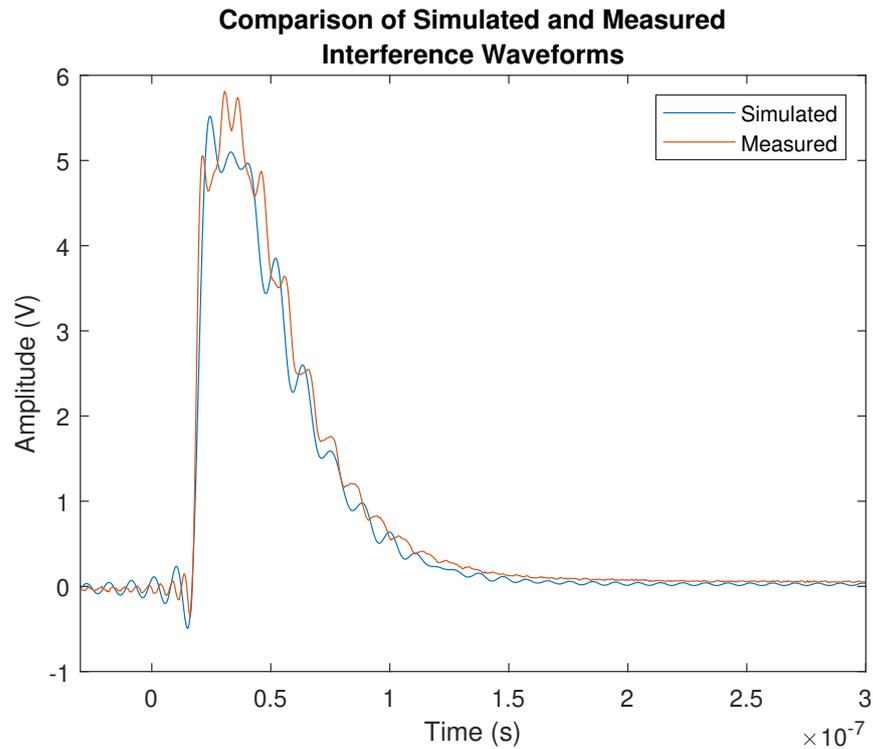


Figure 5.20: Comparison of the measured and simulated induced interference waveforms.

The simulated waveform closely matches the profile of the measured waveform, though with some noticeable differences in some feature details. For example, the simulation accurately models the steps in the decaying waveform (a result of higher-frequency components), but has a slightly offset

which accumulates over time. Furthermore, the peak of each of these steps is more pronounced than in the measured waveform. Additional detail is lost at the peak. The simulated waveform makes an attempt to properly replicate the two peaks, but fails to generate the required amplitude. By using the measured transfer function as the input to this method, it is possible to see that these discrepancies in the induced interference waveform are primarily due to the limitations of the simulated wiring loom. This is evident from the waveform comparison in Figure 5.21 which compares the measured induced interference waveform on the wiring loom with the output of the simulated waveform which uses a measured transfer function as input. A much higher amount of detail is simulated using this approach, suggesting better modelling of the loom would produce a more accurate simulation of the interference waveform.

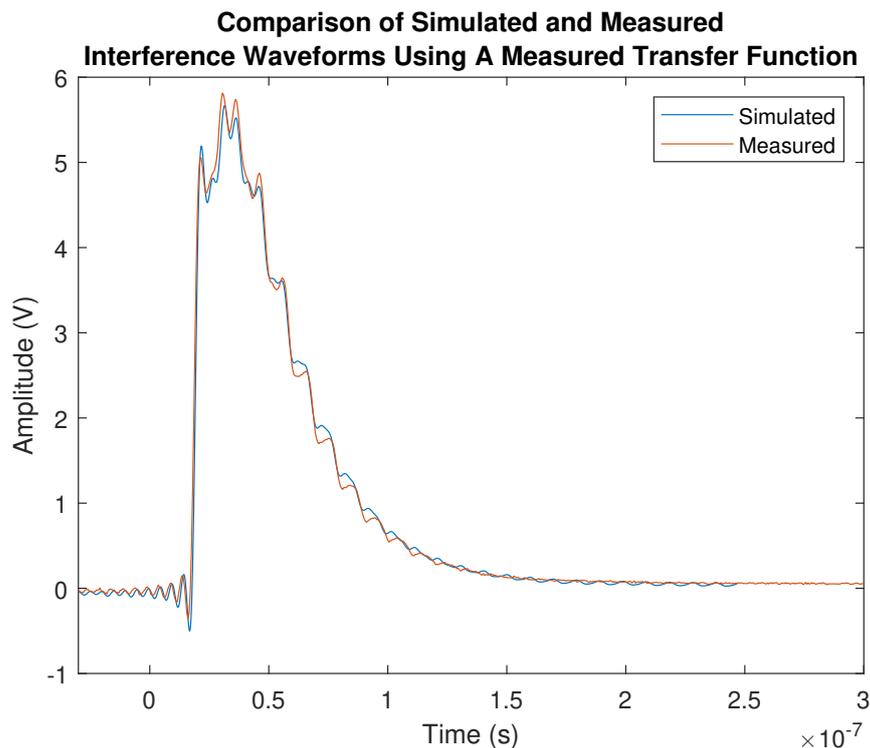


Figure 5.21: Comparison of measured and simulated induced interference waveforms; simulation input from a measured transfer function.

A moving average, taken across the induced waveform across a prede-

terminated number of samples, serves as a simplistic model of a receiver's front end circuit. An incrementing count of the number of averages which exceeded the threshold voltage was divided by the number of moving average samples, creating a probability of exceeding the threshold voltage. When fitted to an array with dimensions equal to the number of interference terms by the number of amplitudes, an array of probability of bit errors was created with noticeable similarities to the 3D BER plots from Section 5.2. An example, using the LTspice-generated complex transfer function and a moving average period of 1×10^{-8} s, can be seen in the line plot in Figure 5.22. Although the magnitude of the initial probability peak is much greater than that observed in the measurements, the simulation does clearly show the same oscillatory relationship between BER and number of Fourier terms, indicating that the effect is not due to some unaccounted experimental variable.

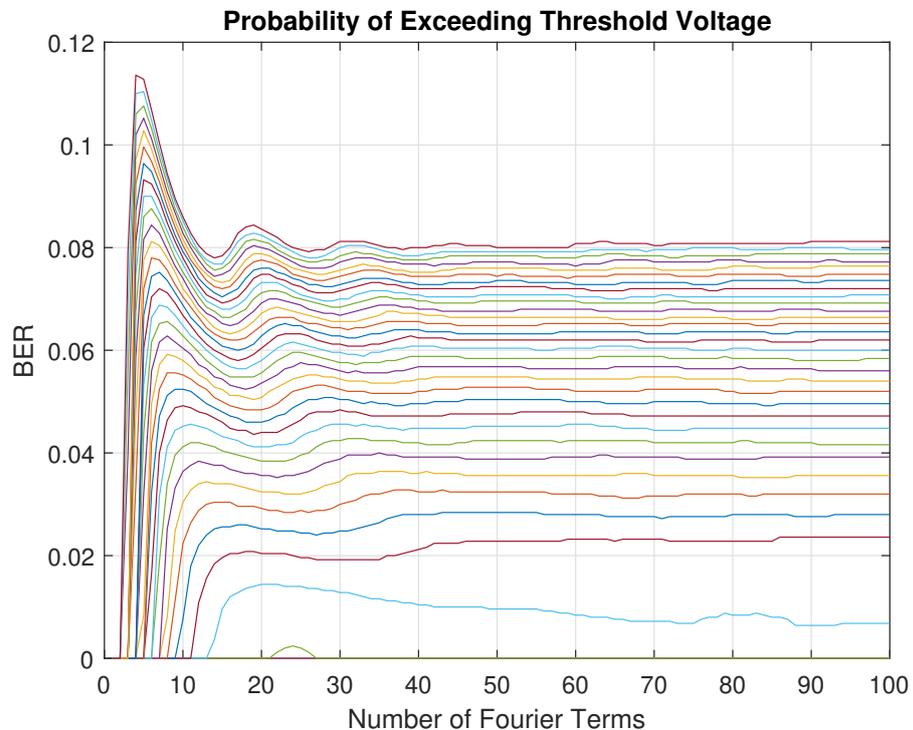


Figure 5.22: Simulated probability of bit errors. Interfering waveform amplitudes are represented with each individual line.

Adjustment of the moving average duration reveals that the peaks and troughs in the BER response can be explained, at least in part, through the sampling period during bit arbitration at the receiver. By increasing the sampling duration in the simulation to 6×10^{-8} s, the probability plot in Figure 5.23 can be generated.

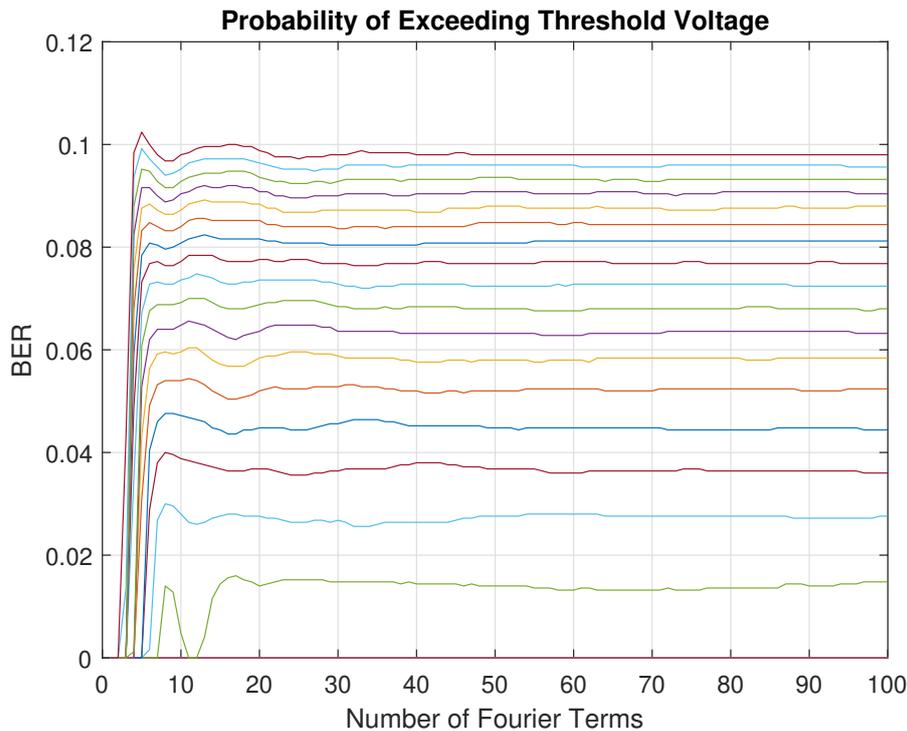


Figure 5.23: Simulated probability of bit errors with a 60 ns sampling period. Interfering waveforms amplitudes are represented with each individual line.

The steady-state probability value can be seen to increase for all amplitudes. For example, the topmost interference amplitude of 30 V in Figure 5.22 has a steady-state probability of error of around 0.08 by 30 terms. The same amplitude in Figure 5.23 converges on 0.1, but within only 10 terms. Furthermore, the peak deviation from the converged value can be seen to reduce as the sampling period increases. When the sampling period is 1×10^{-8} s, the peak has a probability of error approximately 0.034 above the steady-state value. This decreases to 0.004 when the sampling period is 6×10^{-8} s.

This response to varying the sampling period suggests that the BER is also greatly affected by the design of the receiver, even at the silicon level, and that the design of the receiver can have a large impact on the BER response of a system. By decreasing the sampling period, the window for averaging a voltage that is above the threshold for an error is also smaller. Therefore for the same interference waveform, caused by some slope, the number of sampling windows with an average voltage exceeding the threshold will also decrease, lowering the probability of a bit error. This can be seen from the varying steady-state probability values, when the sampling period is adjusted.

The reason for the rising and declining bit error probabilities can be determined by examination of multiple induced interference waveforms, an example of which is given in Figure 5.24.

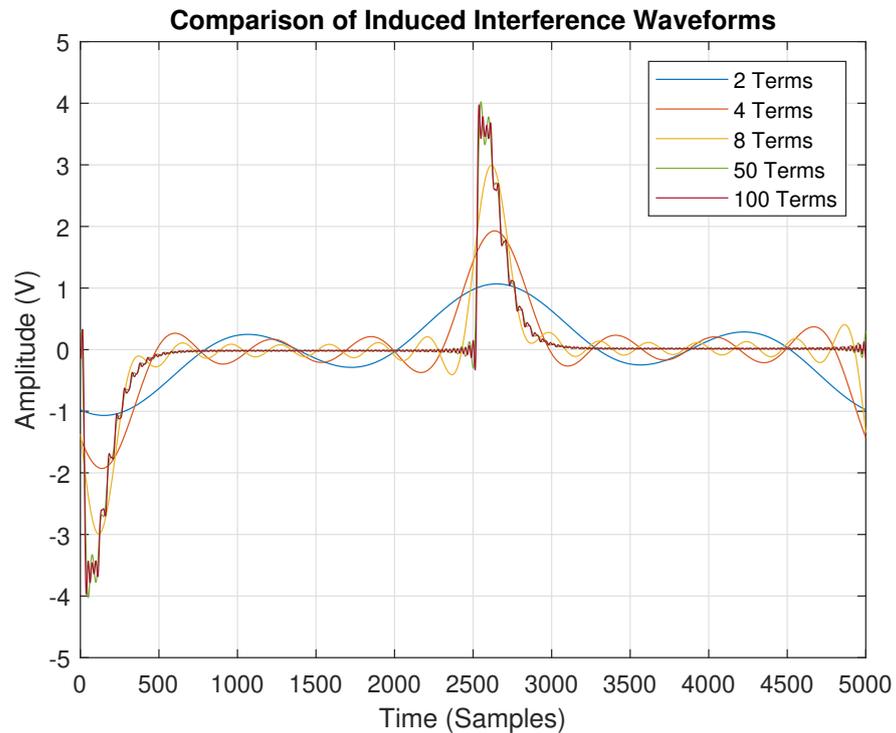


Figure 5.24: A comparison of induced interference waveforms.

At low term numbers, the interfering waveform may not have a sufficiently fast slope to induce a voltage that is capable of exceeding the threshold

voltage. As the number of terms increase, so too does the amplitude of the induced waveform, until the threshold voltage is exceeded and some number of sampling windows will have an average voltage above that threshold. This trend continues, though as the number of terms increases, the width of the interference waveform begins to decrease. Therefore, even though the induced waveform may be greater in amplitude, the reduced time period above the threshold voltage does not affect the mean of a sufficient number of sampling windows to increase or maintain the probability of error. As a result the probability of error decreases, as *two* fundamental criteria must be met for a bit error to occur:

1. The induced waveform must exceed the threshold voltage.
2. The mean threshold voltage must be exceeded across a sampling window.

Therefore the amplitude of the induced waveform is only responsible for bit errors up to a point. Induced waveforms from fast slopes are tall and narrow, and as a result fixed moving sample periods have a high probability of not sampling any of the induced waveform. The BER reaches a steady-state value when the source cannot couple higher-frequency terms onto the victim conductor. This is apparent by comparing the induced waveforms from a source of 50 terms and 100 terms in Figure 5.24. For a clearer numerical example, see Table 5.3.

Table 5.3: Samples above threshold from Figure 5.24.

Number of Terms	Number of Samples Above 1.4 VThreshold
2	0
4	285
8	237
50	198
100	196

In real-world receivers, the factors determining the sampling window period are numerous. Indeed, the approach taken by this simulation is to still utilise a very simplistic approach; only implementing a moving mean to determine sampling window voltage, and measuring the probability of a bit error from a zero-volt state on the victim conductor. That being said, this simulation method has proved faithful enough to begin to understand the mechanisms for this BER phenomena.

These previous test cases used the transfer function generated from LTspice, though examples of BER plots generated using a measured transfer function may be found in Appendix C.

5.4 Corroboration of Time-Domain Effects

Further evidence for the relationship between the time-domain EMI source and the effect of the induced waveform on a victim has also be investigated by other researchers in the GGIEMR using a different simulation method.

The interference waveforms were created ahead of time, for use in a transient analysis in LTspice. Samples for these source waveforms were synthesised in the R software environment and converted to Piece-Wise Linear (PWL) file format, and assigned to a voltage source in LTspice. A SACAMOS model, identical to that used in this work, was used as the basis for the simulation. However, instead of running an AC analysis, the work undertaken by A. E. Pena-Quintal and K. Niewiadomski performs a transient analysis entirely within the LTspice environment. Therefore the effect of the transfer function between source and victim conductors is accounted for entirely within this simulation step, and the induced interference waveform is realised entirely within LTspice.

However LTspice does not easily facilitate a means of arbitrating an analogue voltage to a digital state, nor does it natively support a means

of calculating a BER or probability of bit error. As a result, for each interference waveform, the voltage-time data of the victim conductor must be exported into R. An example of the output from LTspice, showing bit voltages with no interference and whilst under the influence of a source with 49 terms can be seen in Figure 5.25.

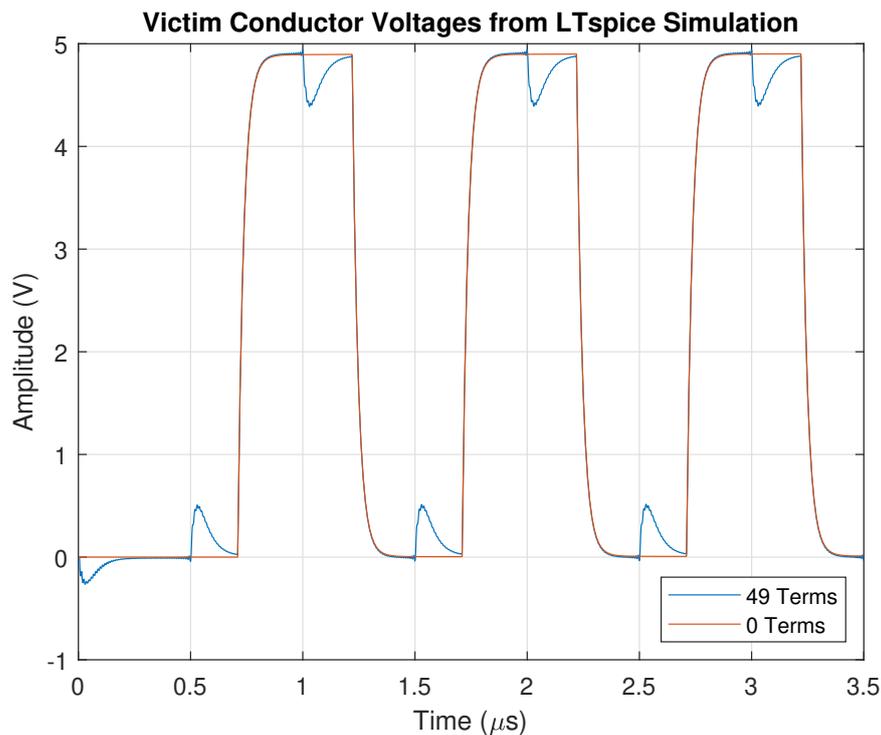


Figure 5.25: The effect of induced voltages on data, from transient LTspice simulations.

Each victim waveform under interference was then subtracted from the unaffected victim waveform in order to derive the induced interference waveform, thus allowing for only the effects of the interference source to be accounted for rather than, for example, any distortions on the victim waveform caused by the characteristics of its own transmission line. The duration of time above the threshold voltage was then taken. The process, from synthesising the interference waveform to measuring time above a threshold voltage, was then repeated for the complete summation of up to 60 Fourier terms.

The time above the threshold voltages were then plotted against the number of Fourier terms in the interfering waveform, as shown in Figure 5.26. The researchers claim that this duration above threshold voltage may be used to derive a probability of bit error.

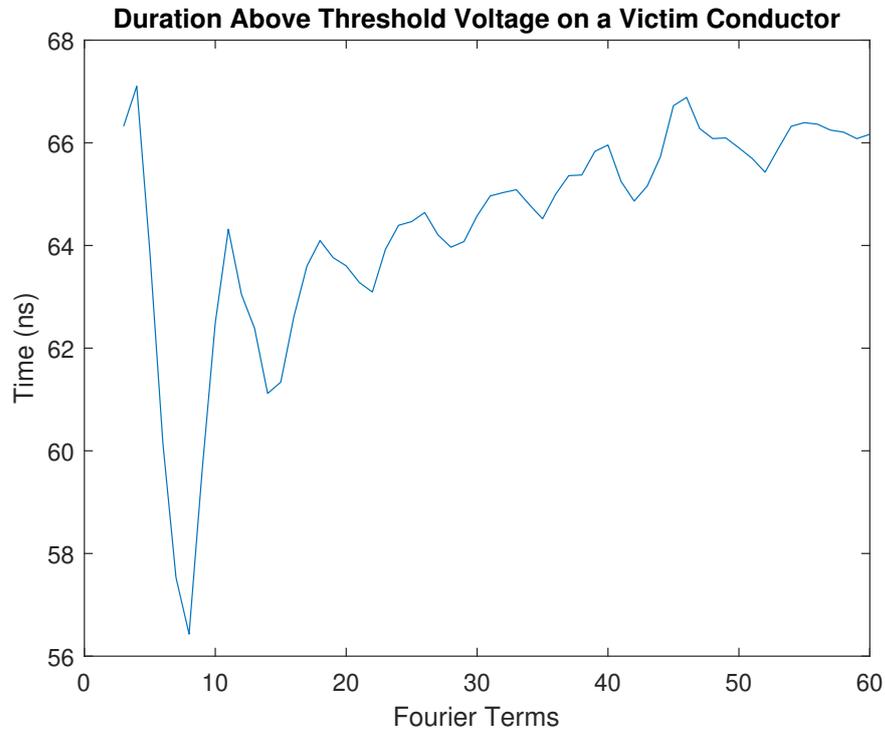


Figure 5.26: Relationship between number of Fourier terms and time above a threshold voltage.

As shown, the duration above the threshold voltage varies greatly depending on the number of Fourier terms, in much the same as the probability of error in the MATLAB simulation. An initial peak duration is observed at 4 terms, which immediately falls to a trough at 8 terms. The period above the threshold then increases to converge on some steady-state value, though with noticeable oscillations as it does so.

However there is also a clear difference between the results of the two approaches. The steady-state period reaches a point very close to the initial peak value, differing by around only 1 ns. In this thesis, it has been shown, through both experimental and simulated results, that this steady-state

value is often much lower than initial peaks, due to the quick reduction in the induced interference waveform's period above a threshold.

Such differences are attributable to the limitations of the cable model or due to the overly simplistic receiver model (in this case, assuming that the period above a threshold voltage is the only primary factor in bit arbitration). At this stage of development it is difficult to precisely indicate where improvements to the modelling can be made. For example, this approach is severely limited in its ability to perform real-world comparisons as a measured transfer function cannot be simply substituted for the simulated example. Furthermore, a more realistic means of arbitrating a probability of bit error is required - limitations were already apparently in the sliding window approach used in Section 5.3.

That being said this approach strongly indicates the findings throughout this chapter - that the relationship between the slope of a interference source and the real effect on the bit arbitration is not so simplistic.

5.5 Chapter Conclusion

This chapter has examined the effect of time-domain EMI on the BER of a variety of standard communication protocols, as well as various abstractions of these networks. It has been found that traditional rule-of-thumb estimations of interference caused by faster rising edges from time-domain sources are not accurate when using the BER of a network as a measure of the amount of interference experienced.

Whilst not necessarily suitable for every form of communication network, by considering the BER of a victim network as the measure of the severity of EMI it is possible to realise improvements in the efficiency and speed of sources of fast-edge switching interference sources. For example, in switching power supplies it is desirable to minimise the switching period in order

to reduce losses and maximise efficiency of the converter. However, utilising very fast switching periods results in very fast waveform edges which has been considered detrimental to the performance of victim communication networks. Another example is in the case of fast digital communication networks, where the transition period between high and low voltage levels must be minimised in order to maximise the throughput of data - in other words, the time period between bits must be minimised in order to reduce time spent not transmitting information. Again, this naturally results in fast switching periods with associated fast slopes. Where such networks are causing interference, the slope may be lengthened to mitigate the effect on victim networks, but at the cost of a reduction in data throughput. The consistent BER response across a variety of network configurations suggest that by taking a more holistic approach to determining the real effect of EMI on a system, such slopes may be properly engineered to reduce the bit corruption on a victim whilst maximising the efficiency of the source system - whether that be an efficiency of power conversion or data throughput.

This holistic approach to determining the effect of EMI would measure the effect a coupled waveform has on the information transmitted through a network, rather than strictly examining only the amplitude of induced or emitted interference at a specific frequency, as per the methodologies underlining many contemporary EMI standards. Negating the effects of a catastrophic breakdown in hardware, caused by extremely high interference amplitudes, the rationale behind the development of EMI standards is to ensure the safe transportation of data - not to ensure emission targets are met for their own sake. Whilst these traditional approaches may prove useful for quick, generalised testing, the results in this chapter indicate that significant benefits in system performance are achievable by understanding the coupling of time-domain waveforms on the information in a network.

This chapter has also described the development of a means of simulating

these BER responses, in order to better understand the reasoning behind the measured phenomena. The simulation methodology creates a reasonable representation of the measured behaviour, though noticeable differences still exist due to the differences between the real and simulated cable bundles, as well as the use of a very basic receiver model. However, significant portions of the simulation procedure have been shown to accurately replicate measured responses, and can be fed measured as well as simulated inputs to facilitate a more accurate output. A different simulation methodology, undertaken by additional researchers, has been able to replicate some aspects of the BER response, such as oscillations tending towards a steady-state, though some overly simplifications currently limits the ability to make direct comparisons to the simulation method in this thesis, as well as real-world measurements.

Future work in this area is needed to better understand the potential for this BER response to be exploited in real-world systems - rather than those under laboratory conditions. If shown to provide significant benefits to real-world systems, then a means of better simulating the receiver portion of a network would facilitate a more accurate representation of BER through the interference form and amplitude dimensions.

Chapter 6

Design of an Open-Source Bit Error Ratio Tester

As discussed in Chapter 5, the work on radiated emissions on wired communication links required the use of a BERT. However, the cost of procuring a commercial BERT for use in the work presented in this thesis was, for a period, far too high - preventing any bit error ratio measurements from being taken. Indeed, prices for off-the-shelf BERT systems can exceed £75,000 [104]. Furthermore, many commercially available BERTs have very high bandwidth capabilities, with even the most basic systems being capable of measuring data rates in excess of hundreds of Megabits per second (Mbps) [105]. Such high data rates were not required in the work performed by the author in this thesis and published in [106], which did not exceed 1 Mbps. Indeed such high data rates are not necessarily required by other academics, particularly if they are working on relatively low data rate systems, such as CAN. It was therefore decided to develop a simplified BERT which would allow for BER measurements to be taken, whilst keeping the costs of the instrument to a minimum.

In the spirit of facilitating ready access to such instrumentation for other academics and persons at large, it was also decided that all documentation

including circuit schematics, HDL, software and PCB files should all be made Open Source (OS), without any restrictive licenses. The design and operation of the BERT (known as OSBERT) as it is today is presented in this thesis demonstrating the flexibility of the system. All of this is achieved with a budget of less than £150. This chapter aims to describe the operation of the OSBERT system, and the novel design choices employed to realise a BERT with great flexibility for general and specialised testing of the effects of EMI on the successful transmission of data through an intra-vehicular network - indeed any data network.

6.1 Fundamental BERT Topology and Function

Fundamentally, a BERT is an instrument which must find the ratio of incorrectly transmitted data bits to the total number of transmitted data bits after they have travelled through a communications channel, whether that is wired or wireless, and have been resolved by a receiver. Some BERTs facilitate the transmission of a wide range of test patterns [105, 99], which can replicate the layer 2 and 3 framing and packeting methodologies found in various communication network protocols, such as CAN and RS232. Other BERTs may make measurements which utilise patterns of bits designed to stress a system in a particular manner. This might be as simple as a string of alternating 1s and 0s for testing tone detection systems in telecommunication systems [107] or pseudorandom binary sequences to ensure an statistically even distribution of 1s and 0s over a measurement duration as in the work performed by the authors in Chapter 5. To do this, the BERT requires five fundamental components; a pattern generator, an error detector, a clock generator, a transmitter, and a receiver [108]. A brief

description of the functionality of these components is given in Table 6.1, though it must be stressed that these components are by no means the only building blocks found within BERT systems, which may also include components such as data network analysers and packet header generators [109].

Table 6.1: Fundamental BERT Components

Component	Description
Pattern Generator	Creates the sequence of data bits to be transmitted through the communications channel.
Error Detector	Compares the transmitted and received data bit and counts the number of incorrect bits.
Clock Generator	Creates a variable frequency clock for synchronising the generation and comparison of data bits.
Transmitter	Prepares the outgoing bit stream for transmission through the channel.
Receiver	Accepts incoming data bits for comparison with their transmitted equivalent.

The basic interconnections between these components is presented in a graphical format in Figure 6.1.

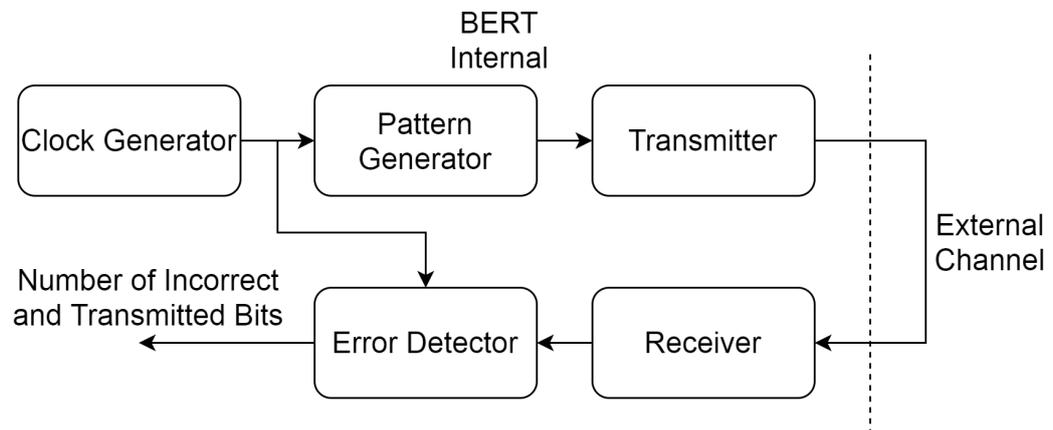


Figure 6.1: Interconnections between fundamental BERT components.

In theory, the operation of a BERT is quite simple. The clock generator creates a variable frequency clock (typically through the use of a PLL) which is used by a pattern generator to clock a data bit. This data bit forms part of a larger pattern, as defined in the setup of the pattern generator. The

data bit is then fed through a transmitter section, where analogue signal conditioning is performed on the data bit, so as to create a representation of the desired level 1 characteristics of the data under test, for example converting the single-ended data into differential mode. The data bit then travels through the channel where a receiver detects the analogue representation of the data bit, and decides the digital representation 1 or 0. This data bit is fed into the error detector module, which is synchronised to the pattern generator. The transmitted and received data bits can then be compared. If they are the same, then only a transmitted bit counter is incremented. If they differ, then an error counter is also incremented.

Whilst the theory of operation of a BERT may be succinctly described, the practice of implementation can quickly become complex. As data rates increase, the design of the transmitter and receiver becomes more critical to avoid poor representation of the desired bit form, and the demands on high speed, high depth memory increases for counting the increased number of bits. This design complexity, coupled with the arms race for faster data systems, has resulted in a market comprised of high-speed, high-cost BERTs, with numerous features ideal for high-level industrial testing, but results in systems which are over-specified for fundamental research work, such as that in this thesis, and which may exceed research budgets. Table 6.2 explores several commercially available systems.

It can be seen that whilst BERT systems do exist which would suitably measure the performance of existing intra-vehicular networks, the data rates for many devices far exceed the maximum specification for systems like CAN. Systems which are able to measure such data rates, such as the MP8302A, are only available second-hand, which is far from ideal when trying to acquire reliable instrumentation. Furthermore, these BERTs are limited in the number of data channels they can measure, and lack any sort of ability to generate synchronisation pulses for other test equipment, limiting

Table 6.2: Comparison of commercially available BERTs

Item	MP8302A	BSX125	N4962A
Cost	£3500 (Used)	£118500	\$37996
Output Signal Levels	TTL, ECL	LVPECL, LVDS, CML, ECL, SCFL	300 - 1800 mV _{pp}
Min. Data Rate	1 kbps	1.5 Gbps	0.5 Gbps
Max. Data Rate	20 Mbps	12.5 Gbps	12.5 Gbps
Precision	2 Decimal Places	Single Bit	Single Bit
Channels	1	1	1
Ext Clk. Inputs	1	1	1
Ext Sync. Outputs	None	None	None
Max. BER	0.094	1	1

their use for simultaneous measurements of individual lines of differential-mode communication networks. A serious limitation of the MP8302A, which was used in the work in Chapter 5, was the maximum BER limit of the instrument. During repairs of the MP8302A, it was found that this instrument was built upon a Maxim DS2174 EBERT integrated circuit (IC). This device has an upper limit on the BER measurement of 0.094, as it detects a BER of 6 in every 64 bits as a synchronisation loss [110]. Whilst such a feature may be useful in some applications, it was frustrating to design around during the work in this thesis.

Therefore, the need for a low-cost, easily re-configurable BERT was apparent. A brief description of the methods used in the design of this novel instrument are presented in the subsequent sections.

6.2 OSBERT - The Design

At the heart of OSBERT is an Intel field programmable gate array (FPGA) Cyclone IV. The use of an FPGA permits the parallelisation of processes within the same device, ensuring that absolute-time dependant operations may happen repeatably and without unexpected interruptions, as might be expected from the use of a microcontroller [111]. In essence, an FPGA is a combination of thousands of configurable logic blocks (CLB) which are linked via programmable interconnections to each other and input/output (IO) blocks, as determined by the developer. By utilising hardware-based logic, an FPGA is capable of performing many operations completely in parallel. This makes an FPGA a distinct class of device, completely separate from central processing unit (CPU) based devices, as is evident from the structure of a typical FPGA in Figure 6.2.

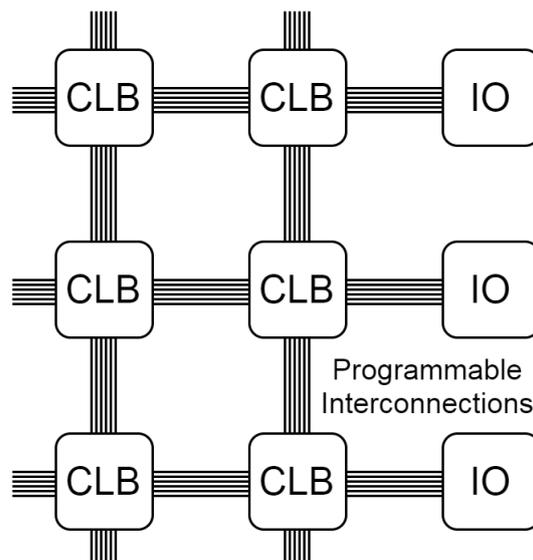


Figure 6.2: Internal layout of a typical FPGA [112].

The Quartus Prime Integrated Development Environment from Intel was selected to create the programming files for the FPGA, as the Lite edition of this software was freely available for download [113], thereby offering a cheap and convenient access for any OS development by other users. Quartus Prime also includes the ability to generate NIOS II soft-CPU's - CPU's which

are entirely synthesisable by the CLBs inside an FPGA. The ability to program some sequential functionality using the C programming language would provide a more familiar means of ensuring non-critical timing events and processes, such as BERT initialisation and communication with a PC, could be easily implemented without the use of complex, parallel state-machines written in a hardware description language (HDL).

As OSBERT would require custom hardware to function, it is possible to divide the design of OSBERT into two main categories: the hardware, and the Hardware Description Language (HDL) and firmware. The following sections describe the reasoning for some of the novel design choices incorporated into OSBERT. Full schematics and printed circuit board (PCB) layouts are available for reference in Appendix D.

6.2.1 Firmware and HDL Design

Figure 6.3 details the high-level modules and interconnections of OSBERT within the FPGA, as implemented by their respective constituent HDL modules.

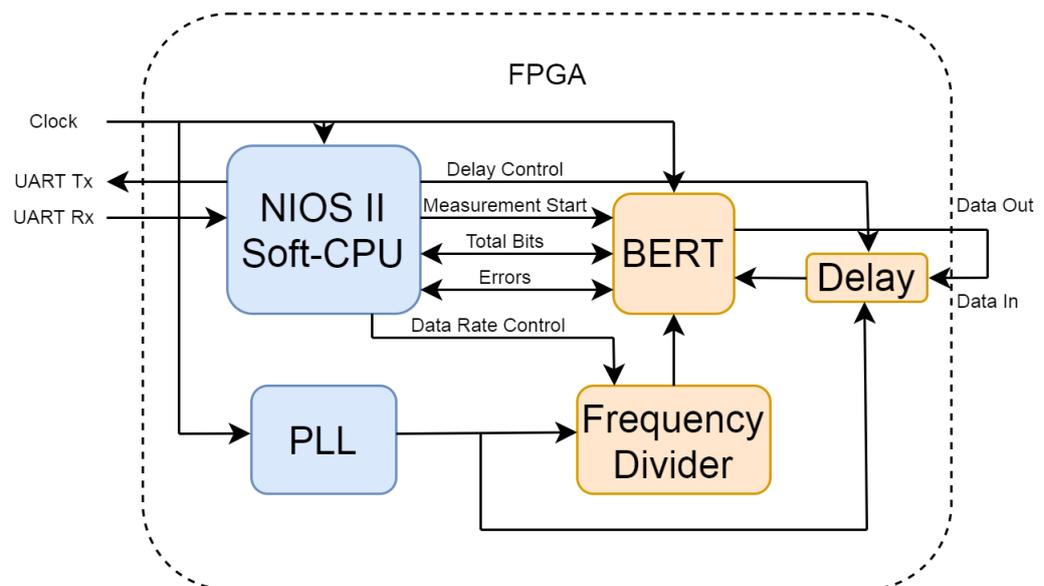


Figure 6.3: High-level layout of OSBERT HDL modules, with annotated interconnections. Modules synthesised by Quartus prime are shaded blue, whilst modules unique to OSBERT are shaded orange.

A NIOS II soft-CPU was utilised to handle the communication with a PC through a serial interface, as well as control the non-critical timing functions of the system such as resetting measurement buffers. The digital-domain BERT components of pattern generation, transmitter, receiver and comparator are handled within the module marked BERT. Analogue-domain functionality of the transmitter and receiver is presented in subsection 6.2.2. To permit the adjustment of data rate the BERT module derives a clock sourced from a frequency divider which controlled by the soft-CPU. This is itself fed by a PLL module which converts the 48 MHz input clock to 320 MHz allowing for the generation of data at rates far exceeding the capabilities of externally-sourced clock sources. A delay module is placed immediately after the data input to the FPGA. This module allows users to set the precise moment a sample is taken of a data bit, as well as accounting for sampling offsets caused by electrically-long transmission lines. It accomplishes this by utilising a string of D-type flip flops. The D-type flip-flop, as shown in Figure 6.4 is a component which consists of two inputs and a single output. An input D determines the state of the next output, Q , whilst the clock determines when that change of output occurs. This is summarised by the truth table in Table 6.3 [57].

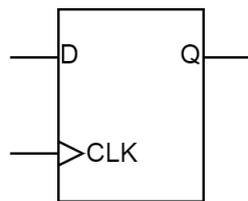


Figure 6.4: Schematic symbol for a D type flip flop.

Table 6.3: Truth table for D type flip flop.

Inputs		Outputs
CLK	D	Q
0	X	Q_{prev}
1	0	0
1	1	1

By arranging these D type flip flops in series, with outputs connected to inputs as per Figure 6.5, then a delayed measurement of the state on the initial input can be taken by sampling any given branch.

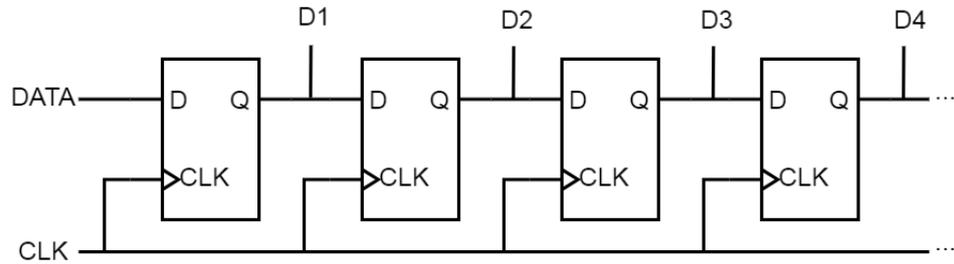


Figure 6.5: Delay module implementation.

If the clock is running at a frequency f_{clk} , in a delay module with n flip flops, then a delay of t_{delay} can be realised by (6.1).

$$t_{\text{delay}} = \frac{n}{f} \quad (6.1)$$

By varying the point at which the tap is taken, (given by D1, D2, ... in Figure 6.5), a variable delay line implemented entirely within the digital domain, is achieved.

It is helpful to consider a worked example to understand the operation of the divider in detail. Consider an input clock of 50 MHz, from which a 1 Mbps data rate is desired. This baud rate is first bit-wise shifted to the left by the number of bits in the accumulator. In the case of a 16-bit accumulator, this can be written as (6.2).

$$1000000_{10} \ll 16_{10} = 65536000000_{10} \quad (6.2)$$

This result is then divided by the width of the accumulator minus one, as shown in (6.3).

$$\frac{65536000000_{10}}{32767_{10}} = 25_{10} \quad (6.3)$$

By incrementing the accumulator by 25 on each clock cycle, then at 50 MHz the 16th bit will toggle at 2 MHz. If the pattern generator clocks only on a positive or negative edge the switching bit, then the desired baud of 1 Mbps will be achieved. The soft-CPU is capable of deriving the desired increment value for a given accumulator width, which allows for precise control over the data rate of OSBERT.

The construction of the soft-CPU, though synthesised using Quartus Prime tools, still requires a highly intricate design approach. This is highlighted by the topology, designed specifically for OSBERT, as shown in Figure 6.6.

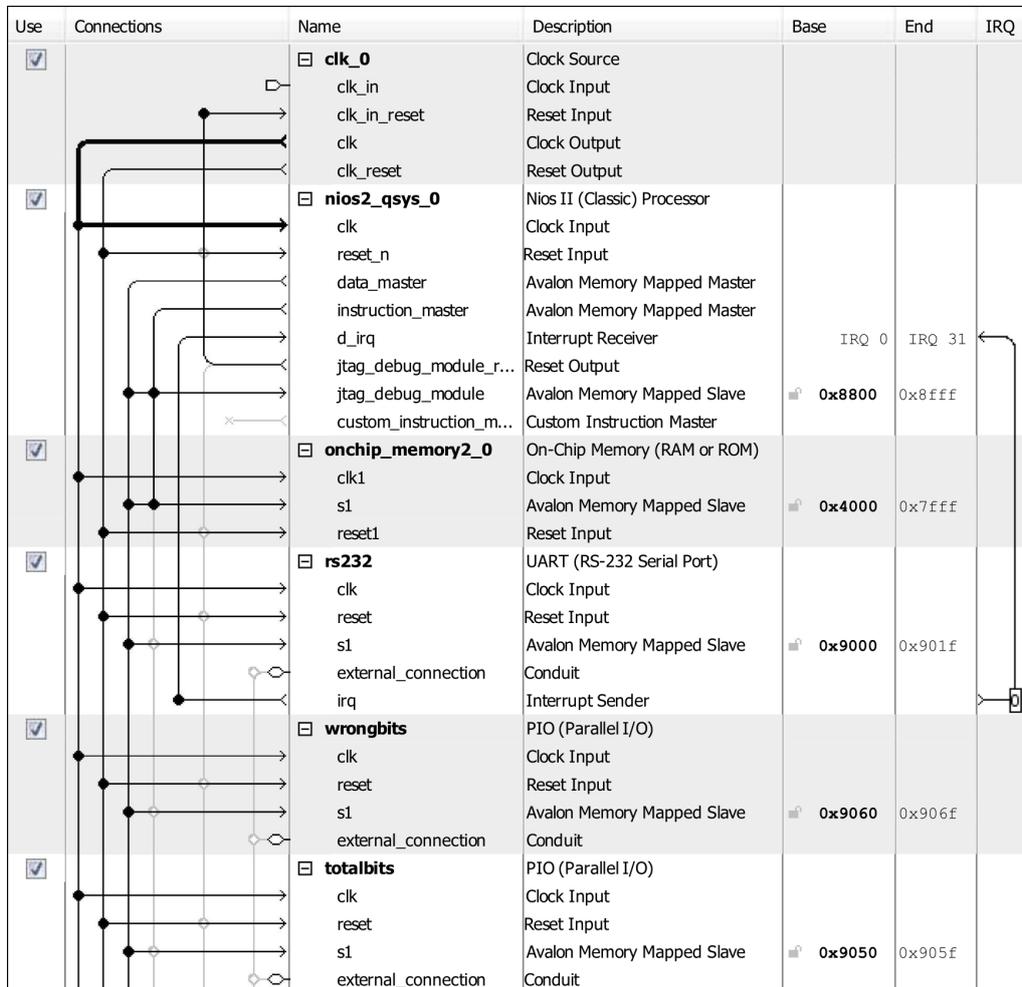


Figure 6.6: The design of the custom soft-CPU in OSBERT, based upon a NIOS II core. Note that several IO ports have been removed from the layout for clarity.

At the heart of the design is the NIOS II soft-CPU module. A UART module is used to communicate with an external PC, reducing the cost of expensive interfacing hardware. 16 kB of on-chip RAM is included in the design, which eliminates the need for external RAM and simplifies the construction processes of OSBERT.

The design of OSBERT deliberately refrains from using a general or multi-purpose data/control IO ports. This not only helps to simplify the implementation of the software within the CPU (as each important parameter can be directly accessed by the CPU with no abstraction), but it also permits the direct implementation of generalised HDL modules with little to no modification required as data is simply written or read at will from the module's dedicated bus.

Figure 6.7 presents a more detailed examination of the key components of the BERT HDL module from Figure 6.3.

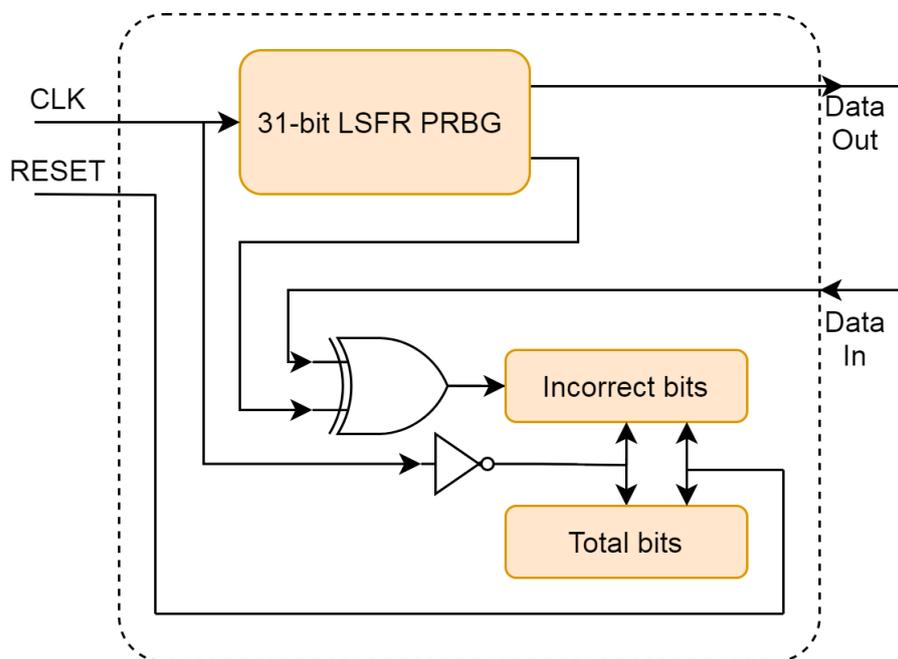


Figure 6.7: Detailed layout of BERT HDL module.

In short, a clock signal is used by a 31-bit linear-feedback shift register (LFSR) to generate a pseudo-random bit on each clock cycle. An LFSR consists of a serial string of D type flip flops, much like the digital delay

module. However, unlike the digital delay module, certain IO taps are connected through an XNOR gate to other IO taps, creating a linear feedback mechanism. The precise taps for these gates is important to ensure that the LFSR produces the maximum number of pseudorandom bit patterns, which is equal to $2^n - 1$, where n is the number of shift registers. Whilst there is no way to mathematically predict if a given LFSR configuration is of maximal-length [114], tables derived by [115] can be utilised to determine connections of the LFSR. These are conveniently available for designers in application notes, and can be readily sourced from [116]. By examination of the LFSR, as 31-bit long LFSR was deemed to provide a sufficient maximal-length random string length, whilst also requiring few connections (two XNOR connections at bits 31 and 28). This is represented graphically in Figure 6.8.

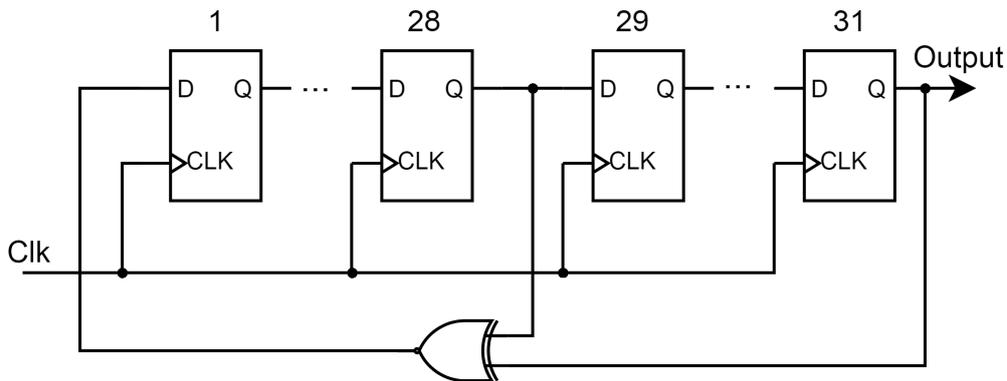


Figure 6.8: Connections of the LFSR in OSBERT.

The clocked, output data stream from the LFSR is then divided into two paths. One serves as a victim channel - in other words it forms the data transmitted through the channel subject to interference and is therefore routed out of the FPGA. The returning data from the channel under test is then fed into an XOR gate for comparison with the unaffected data stream. The exact point of digital state determination is governed by clocking the incrementing incorrect bits and total bits registers. By inverting the clock signal (thereby offsetting the rising edges by 180° , which is shared with

the LFSR, the sample will be taken at the exact midpoint of each data bit. This was suitable for the purposes of the author's experiments, but for experiments which require alternative positioning of the sampling point, the digital delay module can offset the sampling point as required.

Before a measurement, the soft-CPU will reset the values in the both the incorrect bits and total bits registers before enabling the frequency divider output. As each bit is received, the total bits register is incremented and an XOR operation applied with the received data bit and the known-good bit. The output of the operation is then summed with the contents of the incorrect bits. When the CPU is ready to stop the counting, the clock signal is disabled and the two registers are read from their own dedicated IO registers. A maximum measurement duration is not necessarily set by any time limitation on OS BERT, but rather by the size of the total bits register. As a 32-bit register, this limits the maximum number of bits to just over 2.14×10^9 . Depending on the data rate, this may facilitate measurement periods of up to 43 seconds at 50 Mbps, or 2147 seconds at 1 Mbps.

In order to minimise complexity and cost, OSBERT interfaces only through a serial interface to a host PC. Through the use of terminal software such as Termit [117], or through a user's own software or test routines, full control of the OSBERT feature set can be achieved. At the time of writing, the full list of commands is listed in Table 6.4.

Table 6.4: OSBERT commands

Command	Description
start	Starts the measurement and returns BER data.
freq	Change the data rate for a BER measurement.
duration	Sets the duration of a BER measurement.
outon	Turns an output channel on.
outoff	Turns an output channel off.
delay	Sets the delay module length.
h	Displays the help menu.
version	Returns the current version of OSBERT.

An example of the operation of OSBERT, using the Termite software, is shown in Figure 6.9.

```
UoN BERT v1_1
Startup...
Ready. For help send 'h'.

outon
Output On.
Ready. For help send 'h'.

start
Resetting values - done!
Counting... Please wait....

Finished count.
Total bits:      13858035
Incorrect bits: 0
BER:
Please wait... DO NOT SEND COMMANDS

Ready. For help send 'h'.
```

Figure 6.9: Operating example of OSBERT through Termite terminal software. Commands sent to OSBERT are coloured blue, the responses from OSBERT in green.

6.2.2 Hardware Design

At the time of writing, the latest version of OSBERT does not utilise any particularly novel hardware features or requirements, outside of the required support hardware for operation of the FPGA. The precise voltage levels OSBERT transmits data through (and receives data from) the channel under test is determined entirely through the front end hardware, which may be altered by adjusting component values to suit the desired application. Whilst this does present some limitation in terms of ease of reconfigurability, the experiments presented in this thesis used the same LVTTTL transmission protocol voltage levels, operating between 0 and 3.3 V for most of the work. As such, the ability to reconfigure such analogue transmission parameters

via software was not considered in this design, and was omitted in order to simplify the design and reduce costs. However, the front ends have been designed with reconfigurability for other protocols in mind.

OS BERT is designed with three fully independent data channels, which may be reconfigured for use as data clocks for synchronous communication protocol testing, as external clock inputs for synchronising from other instruments, or as external sync outputs for the synchronisation of equipment from the BERT. Each channel consists of a transmit and receive channel, which are again fully independent of one another. Whilst such an approach may be detrimental to certain experimental regimes, where data synchronisation must be measured, this approach allows for the measurement of BER values higher than some commercially available BERTs (see Table 6.2).

Each transmitter utilises a Potato Semiconductor PO49FCT32805 clock buffer driver IC which boasts a maximum 600 MHz TTL operating frequency, as shown in the schematic excerpt in Figure 6.10.

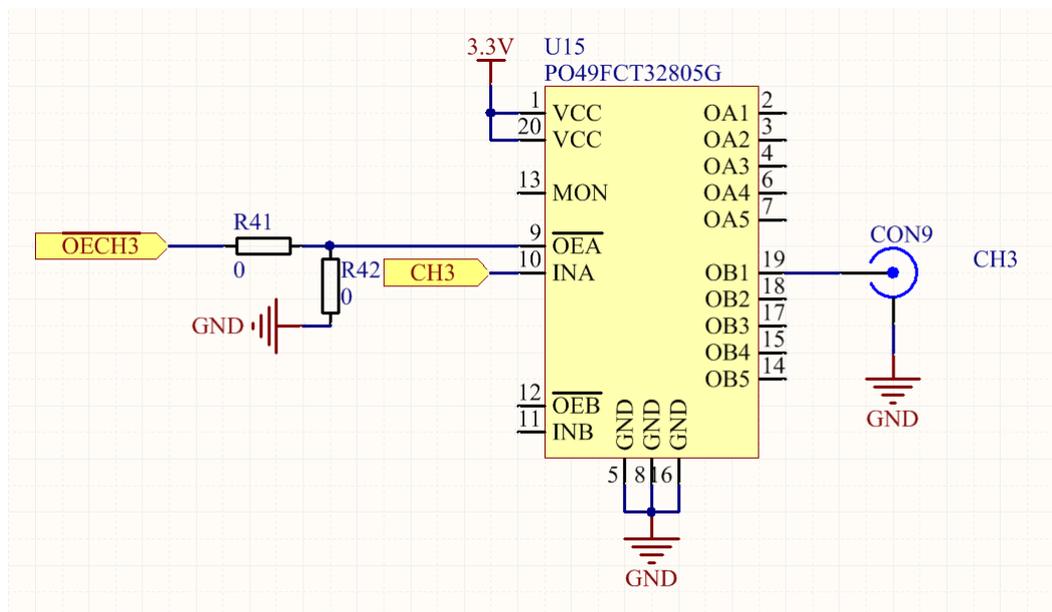


Figure 6.10: Single transmit channel front end of OSBERT.

The ample bandwidth overhead for the target data rates of OSBERT (10s of MHz), along with output drivers stiff enough to drive $50\ \Omega$ transmission

line impedances, ensuring sharp rise/fall times in the order of 1 ns can be achieved [118]. Furthermore, with a 2.6ns propagation delay, the effect of the series PO49FCT32805 on the timing of a general setup is negligible when compared to cable propagation delays. Note the two 0Ω resistors in Figure 6.10 are used to allow for re-configuring of the output enable control by FPGA or hardware.

The data receivers are each based around a ADCMP601 comparator. With rail-to-rail operation and single-supply capability [119] it is a simple and low-cost option to convert the analogue representation of a data bit into a digital state for the FPGA to process. Without the use of a comparator to arbitrate any given analogue state, it is possible for the FPGA to sample an incoming bit at an undefined voltage - in a worst case scenario this metastability may result in damage to the FPGA. This can be thought of by using a ball on a hill as a metaphor, as per Figure 6.11, as per [120].

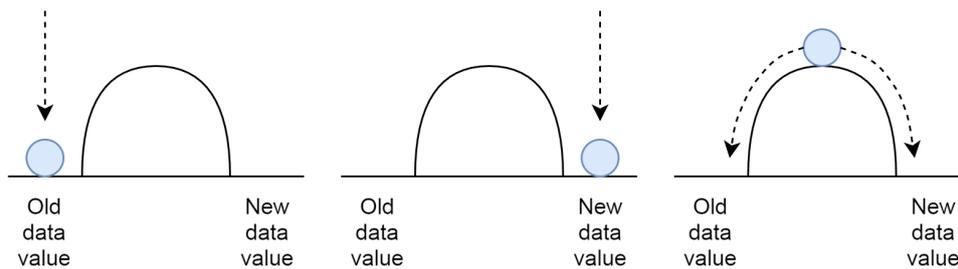


Figure 6.11: Example of metastability of a ball.

The sides of the hills represent stable digital states within the FPGA - that is, states in which the FPGA will behave in a deterministic fashion. The hill itself represents a metastable state. Dropping the ball (metaphor for a voltage transition between the levels for a logical 1 or 0) directly onto either of the side regions results in the signal transition being counted as an old data value or a new data value, depending on when the transition occurs relative to timing clocks within the FPGA. However, if the ball is dropped onto the hill, it may either balance perfectly at the very peak of

the hill but in practice it will tend to roll down towards either side of the hill. In this case, the ball (or rather the voltage representation of a bit) may not reach a stable state within the time period required by the FPGA to allocated register values. In this crude model, the closer the voltage is to the midpoint of states, the longer it will require to settle into a stable state, though in practice such models are more complex (as in [120]) and are beyond the scope of this work.

However, by using a comparator to arbitrate the incoming voltage to one of the two digital states, such issues are avoided. The comparator front-end, along with supporting components, is shown in Figure 6.12.

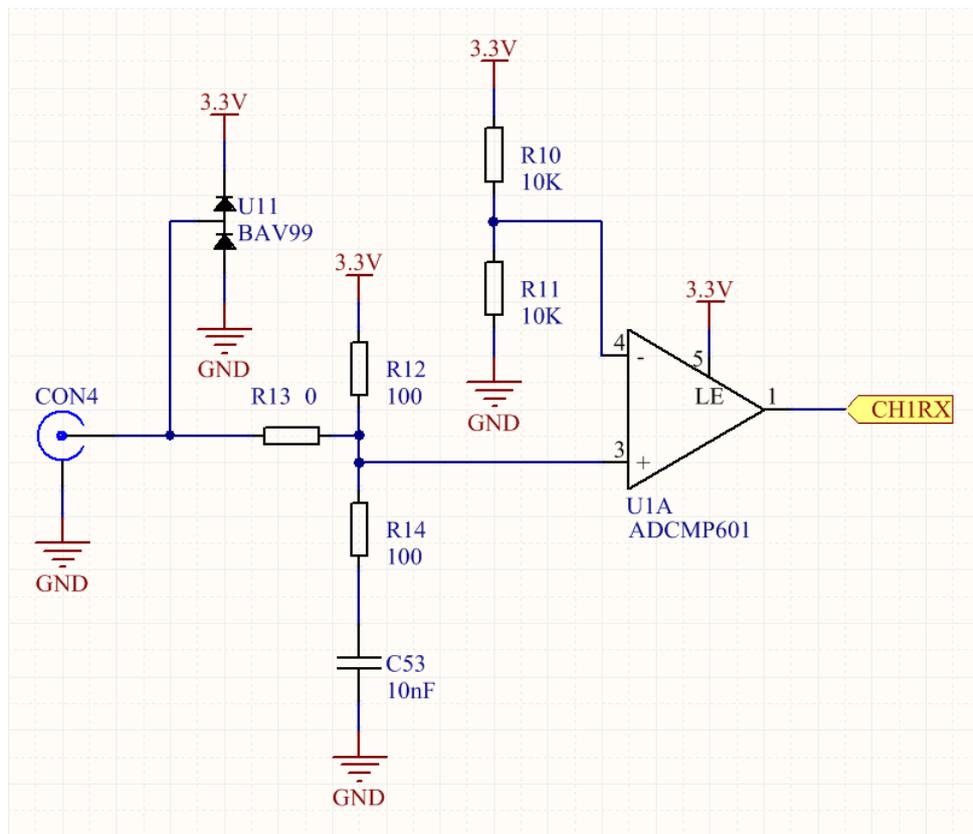


Figure 6.12: Excerpt from the schematics of OSBERT, showing the input comparator.

The output voltage, V_o , of the comparator can be adjusted by the voltage created by the ratio of R_{10} and R_{11} (which form a simple potential divider), as per (6.4).

$$V_o = \begin{cases} 3.3, & \text{if } V_+ > V_-, \text{ where } V_- = 3.3 \cdot \frac{R_{11}}{R_{10}}. \\ 0, & V_+ < V_-, \text{ where } V_- = 3.3 \cdot \frac{R_{11}}{R_{10}}. \end{cases} \quad (6.4)$$

Where the comparator inverting and non-inverting input voltages are denoted by V_- and V_+ , respectively. Other components, also visible in Figure 6.12, act as input protection against excessively high voltage transients (U_{11}) and allow for multiple impedance termination values and topologies to be implemented by a user.

Other hardware features were included in the design of OSBERT, such as an RS232 interface for communication with a host PC, and can be analysed in the full OSBERT schematics, along with PCB designs, in Appendix D.

6.3 OSBERT - A Qualitative Analysis

A fabricated OSBERT is shown, for reference, in Figure 6.13.

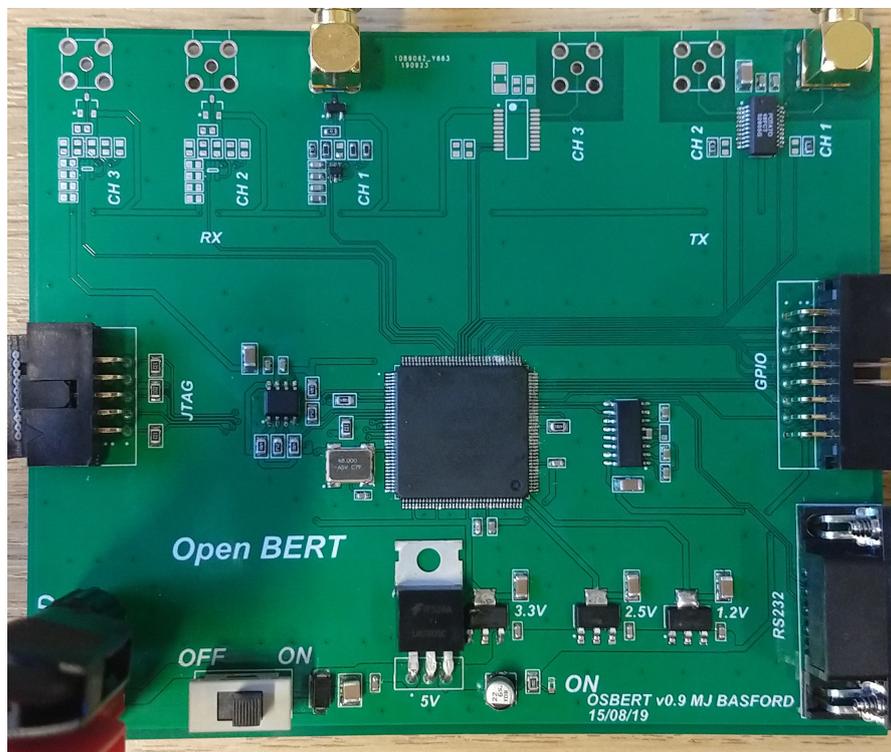


Figure 6.13: An operational OSBERT unit.

Note that this particular OSBERT only has one transmit/receiver channel populated to reduce costs, though functionality of all three channels is easily accomplished for other experiments.

It is difficult to perform a quantitative analysis on the functionality of a BERT. Indeed it may be considered meaningless to attempt to perform anything other than the most basic of quantitative testing on a BERT such as OSBERT, such as whether a 1 or 0 decoded properly when no stress is applied to a channel. This is because, unlike an instrument such as an oscilloscope or multimeter, which measures quantities which may be considered absolute (like voltage or current) with some degree of quantifiable accuracy, a BERT is simply a transmitter and receiver of bits in the digital domain and the ability to resolve an analogue representation of a data bit into a digital interpretation is, at least partially, a function of the front end hardware's properties. Therefore, with any difference in transmitter or receiver design or topology, one may expect a difference in the absolute value of a BER measurement. It may be possible, on the highest-end, instruments to take measurements of the analogue waveform, which can be post-processed to account for differing receiver properties, but this is far beyond the scope of OSBERT and other instruments in its bracket. In this sense, the BER measurement can be considered a relative measurement of a channel's performance against stresses.

The use of a FPGA for all of OSBERT's functionality permits an enormous amount of flexibility for experimental requirements. Other BERTs, such as the Anritsu MP8302A, often use dedicated ICs for at least some BERT functionality, and are therefore limited by any limitations of the IC - a good example of this is the maximum BER imposed by the DS2174 BERT IC. However, the reconfigurable nature of an FPGA allows users to adjust the fundamental inner workings of OSBERT to suit highly specialised experiments. An example of this, which is being realised at the time of writing by

members of the GGIEMR research group at the University of Nottingham, is a BERT with multi-channel BER measurements on individual differential-mode conductors for use on CAN networks. No BERT in Table 6.2 is capable of multi-channel measurements. To perform measurements on a differential network, such as CAN, a transceiver would have to be connected to the data output and again on the input. Whilst such an approach may be useful in some test scenarios, it only permits a measurement of the BER of channel *inclusive* of the transceivers - such a BER measurement would not reflect the behaviour of the channel itself. The flexible, modular nature of the HDL module's of OSBERT allows for easy reconfiguration of the BERT HDL module to the design shown in Figure 6.14.

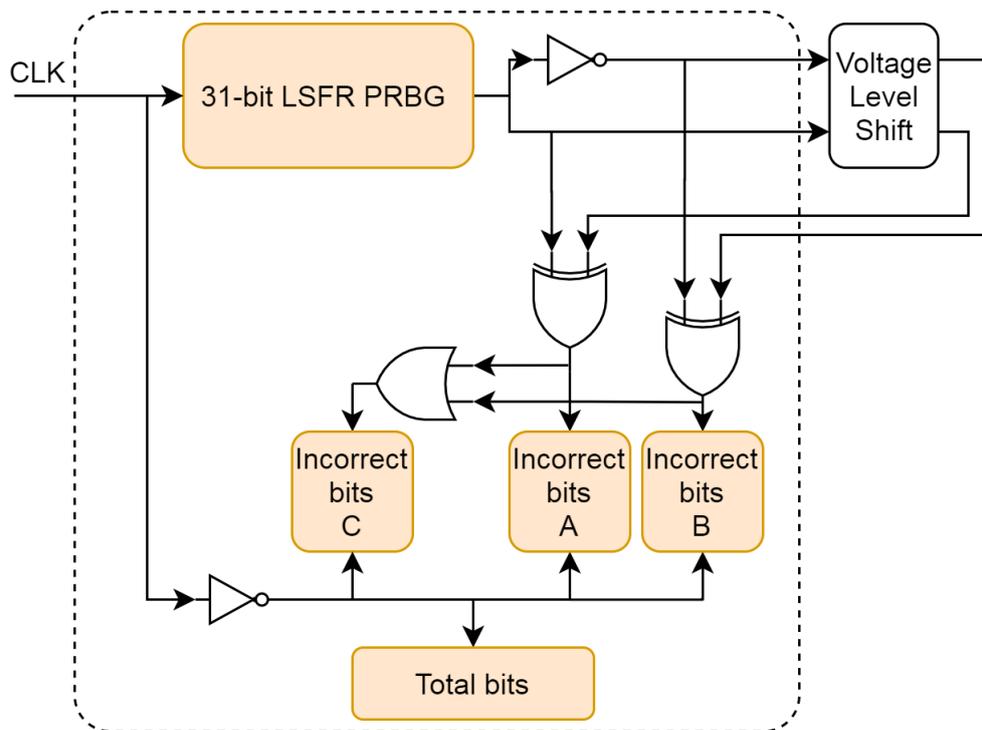


Figure 6.14: A differential BERT module for OSBERT.

The same LFSR from the BERT module in Figure 6.7 now divides into two channels - one for each differential mode. A single channel is then inverted before these two channels are split in two once again. One division from each channel leaves the BERT and uses a voltage shifter to create a true differential-mode data stream, compliant with relevant protocol speci-

cations. This continues through the channel under test, with no additional transceivers. Upon re-entering the BERT, each channel is independently measured, with incorrect bit transmissions on each mode being recorded. These outputs are then combined to form a single BER for the entire system. This demonstrates that OSBERT is capable of making highly specific measurements on just data channels, regardless of being single-ended or differential, without needing to also incorporate the effect of transceivers into a BER measurement.

A qualitative comparison can be made with the other BERTs from Table 6.2. This is presented in Table 6.5 where comparisons are made against the various performance capabilities deemed, in earlier sections, important to the work in this thesis.

Table 6.5: Comparison of OSBERT capabilities

Item	OSBERT	MP8302A	BSX125	N4962A
Cost	£100	£3500 (Used)	£118500	\$37996
Output Signal Levels	TTL, with provision for easy hardware adjustment	TTL, ECL	LVPECL, LVDS, CML, ECL, SCFL	300 - 1800 mVpp
Min. Data Rate	10 kbps	1 kbps	1.5 Gbps	0.5 Gbps
Max. Data Rate	50 Mbps	20 Mbps	12.5 Gbps	12.5 Gbps
Precision phantom	Single Bit	2 D.P. *	Single Bit	Single Bit
Channels phantom	3†	1	1	1
Ext. Clk Inputs	1‡	1	1	1
Ext. Sync. Outputs	Up to 2‡	None	None	None
Max. BER phantom	1	0.094	1	1

* D.P. = Decimal Places

† Fully reconfigurable

‡ Configurable

Each of the devices in Table 6.5 has their respective strengths. Compared to OSBERT, every other device is capable of recreating multiple communication protocols without hardware modifications, and are capable of longer measurement durations. Also, the commercially available systems all have the advantage for inexperienced users in that they are able to be used without construction and a technical understanding of the processes inside the device.

However the OSBERT shows a clear advantage in many other aspects. The cost is significantly lower than even the cheapest of pre-used units and it is capable of data rates applicable to the intra-vehicular communication protocols examined in Chapter 3. Whilst the maximum length of a measurement is limited by the size of the 32 bit registers, OSBERT offers a vastly superior precision compared to instruments in its class, such as the MP8302A. The former can report bit errors to a single bit, whilst the latter can only return figures to two decimal places, even when operated via a PC. The OSBERT also features three, fully reconfigurable data channels.

6.4 Chapter Conclusion

This chapter demonstrates a new, low cost, open-source BERT design which has been made fully open-source for the benefit of wider academic community. Whilst the design of a BERT is not entirely novel in of itself the simplified, modular HDL design methodology of OSBERT allows for unique and novel measurements on a wide variety of communication protocols. Indeed, the channel design is easily scalable, meaning dozens of channels can be measured simultaneously without *any* degradation in performance, due to the easy parallelisation of a hardware design in an FPGA.

The design of OSBERT has been accepted for publication, and was well-received by members of the academic community. The design has been

described as “robust and usable for many applications” by independent reviewers, and OSBERT had been appreciated by members of the academic community. With such feedback, the design of OSBERT clearly meets one of the thesis objectives in providing an open-source means of testing errors on communication networks under EMI. Its adoption by members of the academic community is a positive indicator for the continued development of OSBERT, with a dedicated website being planned at www.openbertuon.co.uk.

Chapter 7

Networks in a Wireless

Domain

The final technical chapter of this thesis examines the feasibility of replacing wired communication networks with wireless alternatives. As stressed in previous chapters, this requires robust communication with low BER, whilst simultaneously providing a relatively high data rate. However, both the financial and time costs of developing an entirely new communication standard are beyond the scope of this thesis. As an example the FlexRay consortium, consisting of BMW, Freescale Semiconductor, NXP Semiconductors and Robert Bosch GmbH [121] took four years of development before publishing the first FlexRay specification in 2004 [122]. Therefore a constrained goal is set for investigating the feasibility of implementing a wireless data network for intra-vehicular communication.

This goal is to create a wireless data link between two nodes in a vehicular environment which can reliably transmit data at 1 Mbps or higher. The data rate of 1 Mbps was chosen as a experimental goal, as this would facilitate the replacement of existing CAN infrastructure with no restrictions on data throughput. It should be noted that in this case, reliable data transmission is determined by the system having less than 10^{-7} failures per hour under

normal operation, as defined by other work on physical layer security for next-generation wireless networks such as 5G [123]. Finally, a measurement of the transmitted and received powers for communication would also be taken. Ideally, power requirements would be low to reduce interference between systems in both intra and inter-vehicular cases. It should be highlighted that the following experiments are, once again, concerned only with examining the effects on the physical layer of the OSI model. Higher level data processing and management techniques are beyond the scope of this work.

7.1 Creating Flexible Wireless Data Links

In order to attain this goal, a suitable means of creating a wireless link for undergoing physical-layer testing was required. In order to create a wireless connection, a transmitter and receiver would be need to be sourced, with carefully chosen features to perform wireless BER measurements. Ideally this transmitter and receiver (or pair of transceivers) would consist of a commercially available transmitter and receiver, thereby reducing the cost and complexity of the system.

In order to perform BER measurements in any way comparable to those in Chapter 5, the chosen wireless system would need to permit the use of wire mode data transfer, rather than packet mode, or at least include embedded BER measurement capability. In this case, wire mode describes a transmitter or receiver which transmits the state of a data input pin in real time, without any additional data processing or packet filling techniques. This would allow for a BERT's data output to be directly attached to the transmitter, and the wireless receiver to be attached to the BERT's data input. Any data processing or packet management would interrupt the flow of data, as expected by the BERT, and would make BER measurements

impossible. The alternative option was to find systems which were capable of making their own BER measurements internally, without any external BERT.

Whilst a real-world wireless data link would require many communications nodes placed at strategic points around the vehicle, as determined by important computing and sensor equipment in Chapter 4, a single point to point communication link was chosen for the work in this chapter. This simple wireless topology would allow for easy data gathering on a single data channel which, as will be described in more detail in Section 7.2, was the only means of performing measurements in a reliable, laboratory setting. A commercially available system would also permit a quick means of scaling the size and complexity of the network, therefore allowing for measurements more representative of a real-world intra-vehicular network.

Finally, the 900 MHz licence free band was chosen as the region of operation of the wireless network. This was due to the low amount of interference measured around the real-world autonomous vehicle in Chapter 4, which would easily facilitate a new communication channel for intra-vehicular communication. Furthermore, the licence-free nature of devices in this band would allow for easy integration of devices by manufacturers, without expensive and lengthy licensing procedures. However, as the work done in this chapter was performed in controlled laboratory conditions, the need to strictly adhere to any licence-free limitations was not generally required.

A number of potential wireless systems were considered for use in this experiment - three of these are listed with their important features in Table 7.1 [124, 125, 126], note that the specific value for the maximum data rate of the Universal Software Radio Peripheral (USRP) X310 is determined by the specific hardware configuration - in this case the limit was around 6 Mbps due to the interface's theoretical upper data transfer rate of 12.5 Mbps.

Table 7.1: Selection of Wireless Networking Equipment

Module	RM191-SM	Z9-T	USRP X310
Price	£14.20	£345.29	£4,750.00
Tx Frequency (MHz)	902-928	902-928	10-6000
BERT	None	None	Embedded
Interface	I2C, SPI	UART	Ethernet
Max Tx Power (dBm)	3	30	10
Max Data Rate (Mbps)	1	2.2	12.5
Channels	1	1	2
Ecosystem	Bluetooth	None	None

Ultimately, it was found that no low-cost, commercially available systems were suitable for use in this experiment. Whilst many systems such as the RM191-SM and Z9-T were optimised for use in the 900 MHz ISM band, systems with wire interfaces were generally capable of data transfer rates in the order of several kbps, whilst the cheapest modules capable of Mbps data transfer required interfaces such as UART, I2C and SPI which introduced unwanted packet bits. Furthermore, these modules did not include any internal BER measurement facility, making fundamental, physical-layer measurements of BER difficult. However the USRP X310 software defined radio (SDR), though more than one order of magnitude more expensive than any other device, proved to have the feature set required for this experiment. This SDR is designed for general-purpose wireless experimentation in laboratory settings, with an extremely large transmitter frequency range, Ethernet interface, and relatively high, maximum transmit power. As such, it should not be considered a drop-in device for any real-world application - it is simply too expensive. However the X310 does allow for easy adjustment of relevant operational parameters. Whilst, at present, there may be no suitable commercial system for plug-and-play operation and validation, the flexibility of the X310 allows for the determination of the necessary operating parameters for future commercial systems, which may be optimised for low cost production.

An example of the X310 unit is shown in Figure 7.1, showing the dual channel RF IO.



Figure 7.1: X310 software defined radio unit.

However, as the use of software in SDR indicates, the hardware of the device is only part of the radio system. SDR uses software, rather than hardware, to implement components such as filters, mixers and modulators. This may permit the faster implementation of new operating modes as no new physical components need to be altered in order to change the radio's characteristics, though it may take some time to learn how to effectively program each target system. In order to minimise the time required to operate the X310, the USRP QPSK transmitter and receiver examples provided by MathWorks were used as the basis for the software control of the radio [127, 128].

The upper level of the transmitter software, controlled through Simulink, is shown in Figure 7.2.

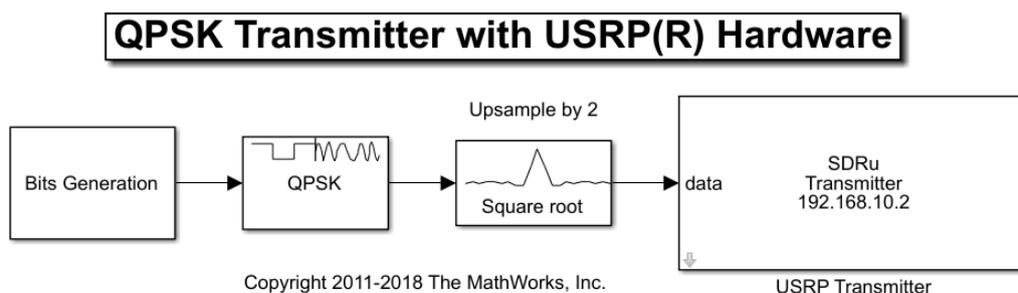


Figure 7.2: QPSK transmitter in Simulink.

The transmitter begins by generating a bit stream, which is shown in greater detail in Figure 7.3. It is inside this bit generation module that the data rate and contents of the transmitted message is generated. Figure 7.3 also shows the use of a Barker code as a header to the message, which allows the receiver to correlate its timing against a known pattern.

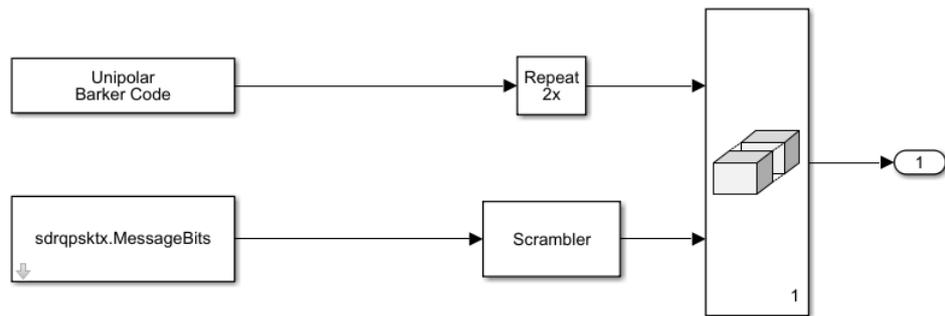


Figure 7.3: Bit generator module detail.

Before this header is combined with the message payload, the message bits are scrambled in a known fashion to redistribute the 1s and 0s in a more even fashion, which aids the timing recovery of the receiver. Once this sequence of bits has been created, it is fed through to the QPSK modulator, which behaves in the typical fashion as described in Section 2.3. This is then filtered through a raised cosine filter before being sent to the SDRu transmitter module. It is in this module that the various radio transmission parameters can be set, such as centre frequency, transmitter gain and channel number. It is worth noting that transmitter gain cannot be used to precisely set a determined output power from the SDR. Indeed, there is no published information on the output power of the SDR as a function of the gain parameter. Work performed by [129] indicates that output power can vary as much as 10dB across the bandwidth of a similar USRP device to the X310. Furthermore, the relationship between gain and output power at a given frequency is not always strictly linear. It is therefore important to understand that the gain parameter in the SDR transmitter can only

provide an approximate guide as to the behaviour of the output power. For quantitative measurements, the output must be measured through the use of a power meter or spectrum analyser.

The top level model for the SDR receiver, again implemented in Simulink, is shown in Figure 7.4.

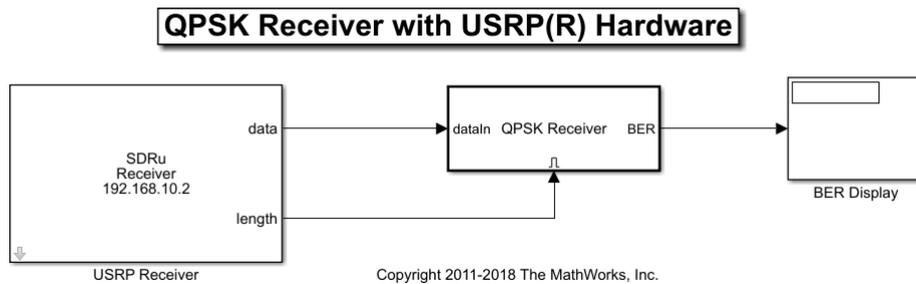


Figure 7.4: QPSK receiver in Simulink.

The SDRu receiver module samples the incoming signal and creates a datastream which is processed by the contents of the QPSK receiver block. This receiver block contains more functionality than just pure QPSK demodulation, which is described in greater detail in Section 2.3. As Figure 7.5 shows, the contents of the QPSK receiver block performs a number of tasks in order to adaptively adjust signal parameters to ensure a more reliable transmission of data. In this instance, the output from the receiver block is a simple BER measurement, though multiple measurements can be taken along the signal path to determine the system's behaviour. However, not all measurements were enabled during every experimental test run, due to the impact these had on the performance of the system. Ideally, the hardware of both the USRP and the host PC would not affect the radio's receive and transmit capabilities, but the various overheads and data bottlenecks which arise through the use of intensive routines such as data storage and display can affect the apparent characteristics of the USRP system due to limitations such as the bandwidth of Ethernet controllers, amount of Random Access Memory (RAM) and number of background applications.

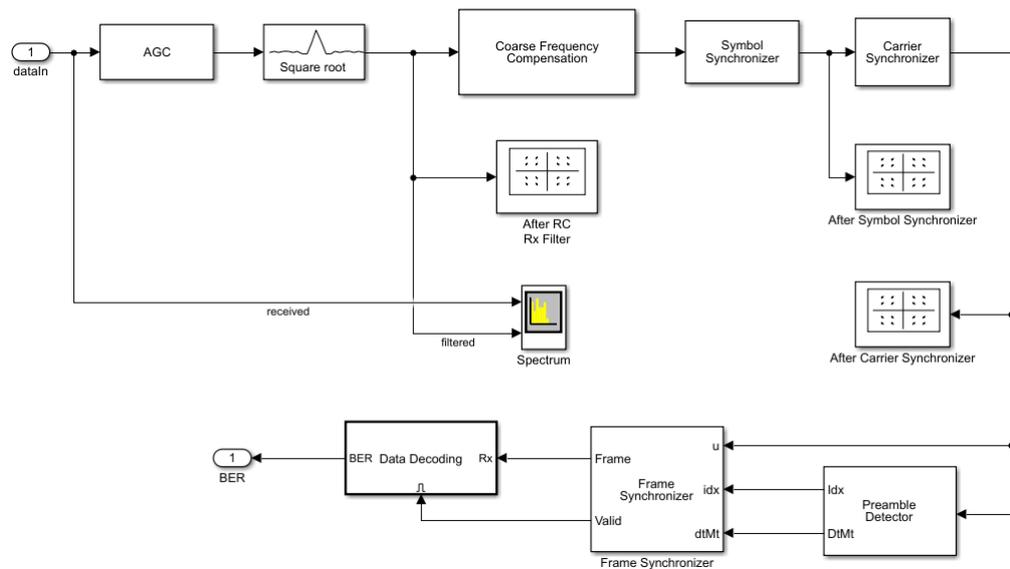


Figure 7.5: QPSK receiver module in Simulink.

First an automatic gain control (AGC) module is used to ensure that the incoming data signal is as close as possible to a desired power, in order for subsequent processing modules to function as optimally as possible. In this case, the desired output power is 2 W, with a maximum permissible gain of 60 dB which was determined experimentally. This is followed by another raised cosine filter, which feeds into a coarse frequency compensator. This module detects differences between transmit and receive carrier centre frequencies of up to 12.5 kHz and removes this offset. In the context of the constellation diagrams from Section 2.3, a carrier frequency offset would result in a spinning constellation in the time domain with no fixed symbol reference points. The coarse frequency compensation module uses a correlation-based estimation algorithm to match the received carrier frequency as closely as possible, such that symbol coordinates in the complex plane may move as little as possible. This refined output is then processed by the symbol and carrier synchronisation blocks, where symbol timing information and small frequency and phase corrections are implemented on the data. The payload is then processed by the frame synchroniser, after the preamble (consisting of the Barker code) is recovered by the preamble

detector module. When a frame is reconstructed, a BER measurement is made against the known transmitted data.

Many of the blocks in Figure 7.5 also contain smaller subsections, but these are generally beyond the scope of this work. As discussed in the following sections a great amount of low level fine tuning was not required in order to facilitate reliable data transfer, and so low level adjustments were not utilised in this work. For completeness, an example of a matching set of the MATLAB structures for both transmitter and receiver are given in Appendix E.

Through a trial and error based approach, the USRP was quickly configured to perform BER measurements on data through a given channel. Initially this channel was a length of coaxial cable - as close to an ideal channel as was available to the author. During this period it became apparent that hardware limitations of the host PC limited the values of several key target parameters of the experiment. Namely, the measurement period for high data rates (1Mbps or more) was severely limited to around 30 seconds before data overflows were experienced between the USRP and the host PC. As a result, uncharacteristically large BER values were measured after 30 seconds. Therefore, measurement durations were typically limited to 20 seconds, which at 1Mbps still resulted in a confidence levels greater than 95% for less than 1 significant figure variation. Furthermore, the trends in results in later sections indicate that the contributing bit errors were limited to only the initial startup of the system (more specifically the AGC convergence) and that during steady state operation no errors were measured, thereby increasing the effective confidence level of a steady state measurement of the system under test for a given BER.

7.2 Replication of an Automotive Environment

Due to the relatively large physical dimensions of the test equipment and associated PC systems required to create a functional wireless data link for BER testing, it was deemed highly impractical to perform tests inside a real vehicle. Instead, a reverberation chamber (RC) was utilised to replicate an example of a wireless communications channel in a vehicle.

An RC is an enclosure with boundaries which absorb the minimum amount of EM energy - in effect an opposite counterpart to the anechoic chamber described in Section 5.2. Instead of covering the interior of the enclosure with RF-absorbent materials such as pyramidal absorbers, the RC interior is made of highly reflective metal surfaces which both reflect interior test signals and shield the experiment from external RF sources. As a result, the RC may be considered as a cavity resonator with a high Q-factor. The Q-factor in this instance describes a component's ability to resonate when excited by some external energy [130], or rather the ratio of the stored energy in the system to the energy dissipated by the system during each resonant cycle which may be expressed as per (7.1).

$$Q = \frac{E_{stored}}{E_{dissipatedpercycle}} \quad (7.1)$$

This high Q-factor results in an environment with standing waves throughout. In order to achieve a more homogeneous spacial distribution of electric and magnetic fields in the chamber a large stirrer may be rotated to redistribute fields making them, on average, spatial isotropic and homogeneous. This makes RCs ideal for testing the EM compliance of physical devices as there is no need to rotate the device itself with respect to radiating antennas - a technique similar to that utilised by microwave ovens to evenly

heat food. Furthermore, the high Q-factor of RCs means that it is capable of very efficient conversion of an RF-source to high-intensity fields for testing of large systems [131]. In this instance, the need to replicate a wireless link between two separate points within a car would require spatial separation in the order of a couple of metres - accommodated easily by the RC owned by the GGIEMR at the University of Nottingham, UK.

Many important experimental parameters, such as lowest usable frequency, working volume and Q are dependent on the physical dimensions of the RC. A simplified form of the work in the National Bureau of Standards Technical Note 1066 [132] shows that number of modes above a cut-off frequency can be approximated by (7.2).

$$N \approx \frac{8\pi}{3} \cdot a_x a_y a_z \cdot \frac{f^3}{c^3} - (a_x + a_y + a_z) \cdot \frac{f}{c} + \frac{1}{2} \quad (7.2)$$

Where N is the unitless number of modes above the cut-off frequency f in Hertz, a is a physical dimension of the RC in metres, and c is the speed of light in free space in m/s. The product of the physical dimensions can be seen to have the most significant effect on N . Whilst the number of required modes can vary depending on the required statistical homogeneity, generally at least 100 modes above the cut-off frequency are required [133]. The GGIEMR's RC, built by Siepel, with the dimensions of 3.11 m by 3.72 m by 4.84 m has a specified lowest usable frequency of 200 MHz [134] - resulting in a comfortable 130 modes. The chosen centre frequency for transmission of 915 MHz has over 13000 modes, as shown by the plot of theoretical N values for the dimensions of the GGIEMR RC chamber in Figure 7.6. This figure also highlights the approximately cubic relationship between frequency and number of modes for a given volume. As a result, high frequencies easily facilitate an extremely high statistical homogeneity, but the number of modes for lower frequencies changes very little. This low

frequency response, primarily affected by the volume in (7.2) indicates that low frequency operation is not particularly well-suited for RCs and care must be taken to ensure spectral content results in statistically relevant homogeneity.

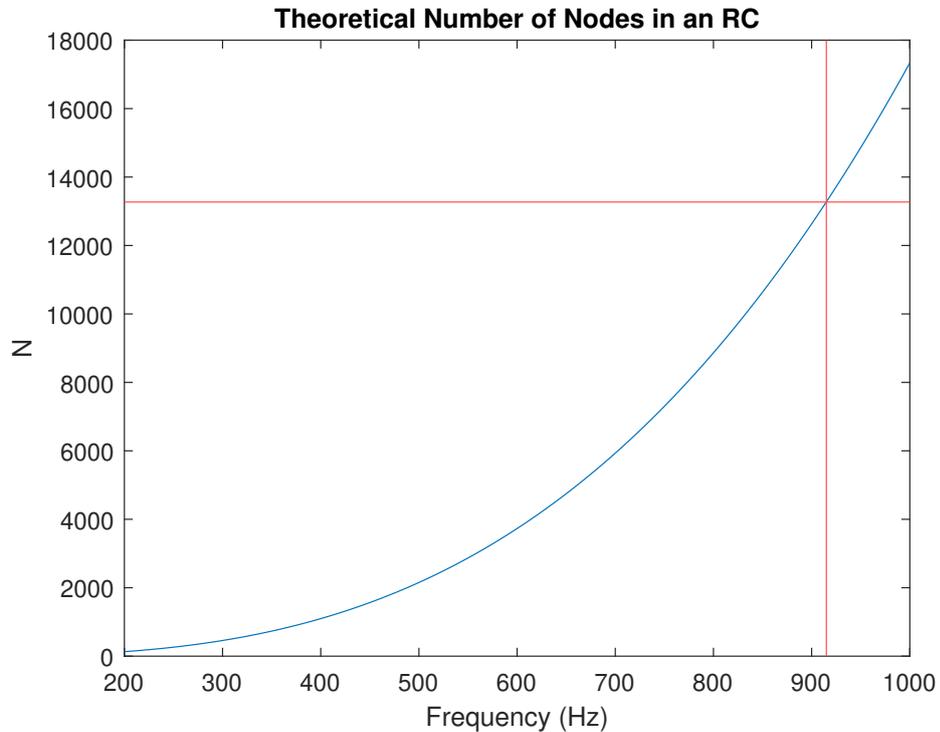


Figure 7.6: Theoretical number of modes for the RC. Red lines indicate the value at 915 MHz.

As well as a lowest usable frequency, the RC also has a defined working volume, to ensure boundary effects are negligible. This is primarily dependant on the wavelength of the lowest usable frequency, with the working volume inset between $\lambda/4$ to $\lambda/2$ from the interior surfaces[135]. Dimensions of and location of the stirrer also affect the final size of the working volume which, in the case of the RC owned by the GGIEMR, was approximately 2.7 m by 1.7 m by 2.1m, as per Figure 7.7. As long as any transmitter or receiver antennas are placed within this volume, then the near-field effect of the chamber boundaries and stirrer plates may be considered negligible in this work.

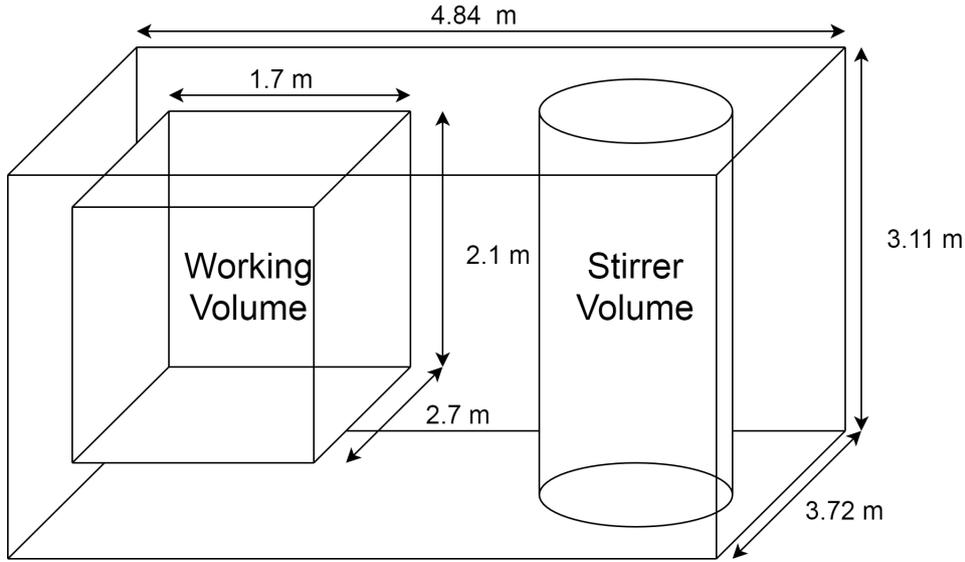


Figure 7.7: Working volume in the GGIEMR's RC.

Finally, the Q factor of the RC, for a narrow band Δf at frequency f , can be estimated using the equation provided by [136] as per (7.3).

$$Q(f) \approx \frac{3}{2} \cdot \frac{V}{\mu_r S \delta_s} \cdot \frac{1}{1 + \frac{3\pi}{8k} \left(\frac{1}{a_x} + \frac{1}{a_x} + \frac{1}{a_x} \right)} \quad (7.3)$$

Where, a is a chamber dimension in m , V is the chamber volume in m^3 , S is the total surface area of the chamber's interior in m^2 , δ_s is the skin depth in m (as per (2.11)), μ_r is the dimensionless relative permeability of the RC's material and k is the wave number in m^{-1} , as per (7.4).

$$k = \frac{2\pi f}{c} \quad (7.4)$$

However it should be noted that both [136] and [134] indicate that this method generally over estimates the Q factor of an RC, especially at sub-GHz frequencies. Furthermore, the work performed in [134] indicates a high variability in Q as a function of frequency when the chamber is 'loaded' - that is to describe the addition of RF absorbers to adjust the reflective environment of the RC. From examining data produced by [134], it is possible to estimate the Q of this particular chamber to between 30000

- 60000 under approximate loading with the same RF absorber type. In this work, a precise value of Q was not required.

The key to the replication of the automotive environment in this section is the reflective properties of the RC. Considering the automotive environment inside a car, the wireless channel may reasonably expect to operate in an environment with many reflected, or multi-path propagation, modes. The high number and close proximity of surfaces reflective to RF would cause a large amount of reflected signals from a transmitter which would be received at different time intervals by the receiver. This would naturally have a significant impact on the reliability of a wireless system if different symbols were being resolved by the receiver antenna at the same time due to this multi-path propagation. A measure of a channel's usable bandwidth is given by the coherence bandwidth, B_c , which may be defined as per (7.5).

$$B_c = \frac{1}{\tau} \quad (7.5)$$

Where τ is the temporal spread between the first and last significant spectral components - typically assumed to be -3 dB - of a transmission from an instantaneous moment, measured in s. In order to test the functionality of the proposed wireless link in such a complex environment, a measure of the multi-path (or scattered) components was required. This was done through the measurement and replication of a number of K-factors.

The K-factor is the ratio of a line of sight signal component and scattered components of a signal from an instantaneous point in time, received by an antenna, as per (7.6) in a multi-path environment.

$$K = \frac{\text{DirectComponent}}{\text{ScatteredComponents}} \quad (7.6)$$

Consider the E-field components inside a well-distributed RC. The total field, as a component of θ in a spherical co-ordinate system centred on the

transmitting antenna, may be expressed as in (7.7).

$$E_{\theta} = E_{s\theta} + E_{d\theta} \quad (7.7)$$

Where $E_{d\theta}$ is the direct E-field component and $E_{s\theta}$, the scattered component, may be written as a function of time, t in (7.8).

$$E_{s\theta}(t) = E_{0\theta} \cos(\omega t + \varphi_{s\theta}) \quad (7.8)$$

Where $E_{0\theta}$ is a constant vector, ω is the angular frequency, and φ is the phase. This can be written as (7.9), in complex notation with the base of the natural logarithm, e .

$$E_{s\theta}(t) = \Re\{E_{s\theta}e^{j\omega t}\} \quad (7.9)$$

Where j is the imaginary unit. Through the use of Euler's Formula in (7.10):

$$e^{jx} = \cos(x) + j \sin(x) \quad (7.10)$$

This equation may be written as per (7.11).

$$E_{s\theta} = E_{s\theta j} + jE_{s\theta i} \quad (7.11)$$

Supposing the E-field in the RC is a circular Gaussian field, as in [137], then for an E-field with time dependence $e^{j\omega t}$ in (7.9), the mean of each field component in (7.11) will be zero as a result of the central limit theorem which states that for a suitably large number of samples, the distribution will approximate a Gaussian, with variances in (7.12).

$$\langle E_{s\theta r}^2 \rangle = \langle E_{s\theta i}^2 \rangle = \sigma^2 \quad (7.12)$$

The work in [138] shows that through the use of Rayleigh and Rician distributions, which are beyond the scope of this work, the K-factor may be defined as in (7.13), which is a convenient form for experimentally calculating the K-factor.

$$K = \frac{|E_{d\theta}|^2}{2\sigma^2} \quad (7.13)$$

This convenience stems from the natural correlation between the equation for total E-field in the RC and the total scattering (S) parameters of the RC, as given by (7.14).

$$S_{21} = S_{21d} + S_{21s} \quad (7.14)$$

Where S_{21} is the ratio of the received waveform at the receiver antenna to the incident waveform on the transmit antenna. This complex quantity is, in effect, a transfer function of the system. S_{21d} is the contribution from direct components and S_s is the contribution from scattered. Note the similarity between (7.14) and (7.7). By defining values for $E_{d\theta}$ and σ from (7.13) through measured S-parameters, the K-factor for an RC, or other environments, may be determined.

[139] demonstrates that the complex variable S_{21s} are normally distributed with zero mean as in shown in (7.15).

$$\langle S_{21s} \rangle = 0 \quad (7.15)$$

Whereas the direct S_{21} component can be shown to have a non-zero mean, defined as (7.16).

$$\langle S_{21d} \rangle = d_R \quad (7.16)$$

Where d_R is some non-zero, unitless value.

Therefore, the mean for the total S_{21} is given by (7.17).

$$\langle S_{21} \rangle = d_R \quad (7.17)$$

Whilst the mean of the scattered S_{21} component was zero, the variances are all identical and non-zero. Furthermore, the variances of the direct S_{21} component are all zero. This results in (7.18):

$$\mathfrak{Var}(\Re(S_{21s})) = \mathfrak{Var}(\Im(S_{21s})) = \langle \Re(S_{21s}) \rangle = \langle \Im(S_{21s}) \rangle = \sigma_R^2 \quad (7.18)$$

Which can be rewritten as per (7.19).

$$\langle |S_{21} - \langle S_{21} \rangle|^2 \rangle = 2\sigma_R^2 \quad (7.19)$$

These value for variance and mean, (7.19) and (7.17) respectively, can then be combined with (7.13), resulting in (7.20) which allows for a K-factor to be calculated from a complex S_{21} value.

$$K = \frac{|E_{d\theta}|^2}{2\sigma^2} = \frac{d_R^2}{2\sigma^2} = \frac{(|\langle S_{21} \rangle|)^2}{\langle |S_{21} - \langle S_{21} \rangle|^2 \rangle} \quad (7.20)$$

As such, the K-factor for an environment may be calculated from measurements from a full 2-port Vector Network Analyser (VNA). However, the RF propagation environment inside a vehicle has been shown in Chapter 4 to vary greatly with small spatial variations. As such, a number of potential wireless node locations were specified around a 2009 Fiat Panda, chosen for its approximate size comparison to the Kia Soul EV, used in Chapter 4, and its experimental availability. Three test locations were determined to correspond with some of the important locations from the Kia Soul EV; the driver's console, the boot and the engine compartment, with testing rationale provided in Table 7.2.

Table 7.2: Fiat Panda 2009 test locations.

Number	Location	Rationale
1	Under Bonnet	Any engine unit, whether combustion or electric, requires control signals.
2	Central Driver Console	Important location for relaying car performance metrics to a driver.
3	Rear Left Boot	Space reserved for computing equipment in the modified Kia Soul EV.

The precise locations of these test points are shown in graphic form, below, in Figure 7.8.

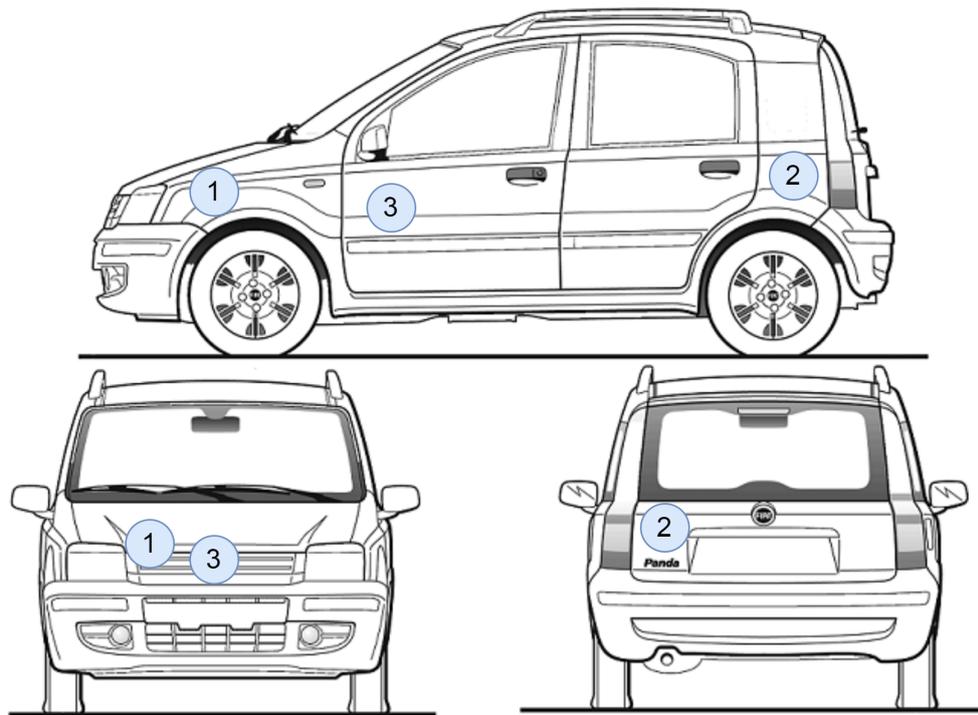


Figure 7.8: Location of test points in the Fiat Panda 2009. Edited from [140].

As the K-factor required only complex S_{21} -parameters the experimental equipment required for the setup was limited to only two antennas, each connected to a two port VNA as per Figure 7.9. The SAS-571, as used in Chapter 5 was chosen for use in this experiment, along with Pycom LoRa & Sigfox monopole antennas [141] - chosen for a better representation of a low-cost, compact antenna for use on a wireless node. Through experimental comparisons, an outdoor test area was utilised which provided ready access

to power for the VNA, whilst also minimising any affect on the measured K-parameter inside the vehicle. Furthermore, all measurements were taken in the evening to ensure no passing vehicles or persons would affect readings.

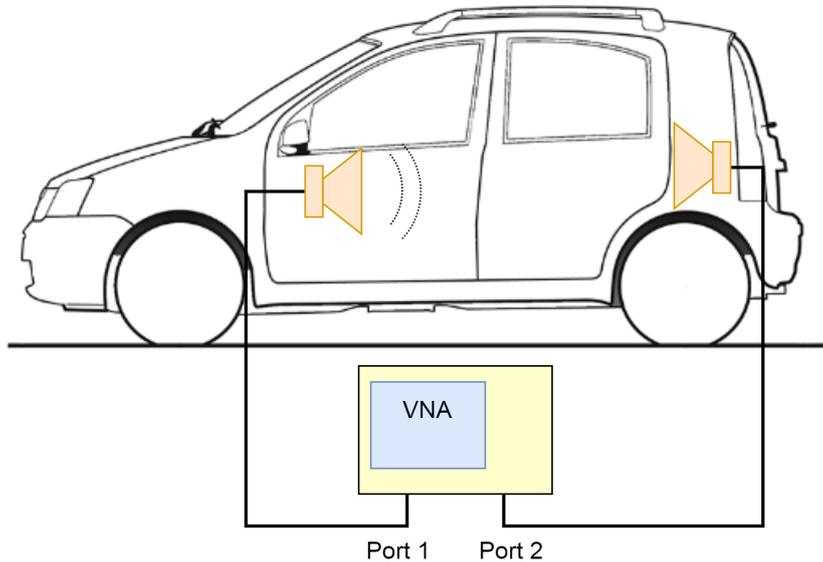


Figure 7.9: Experiment layout for measuring S-parameters around the car.

A photo of the test setup is provided in Figure 7.10. Note the VNA position, normally hidden from view, was adjusted for the photo. All coaxial connections between the VNA and antennas entered the car through slightly-opened windows to, again, minimise differences in the environment.



Figure 7.10: Photograph of the experiment around the vehicle.

Multiple measurements were taken with both antenna setups, altering between the test locations presented in Table 7.2. An example of a complex S_{21} measurement, taken between the driver’s console and the boot is shown in Figure 7.11. Additional plots are provided in Appendix F. All subsequent plots use a naming convention whereby the first two letters indicate the location of transmit and receive antennas (by D = driver’s console, B = boot, E = engine), whilst the third letter indicates the type of antenna used (H = SAS-571 horn, M = Pycom monopole). An additional set of measurements were taken, with the rear seats lowered (indicated by S in measurement plots). This was intended to provide a measurement for a strong wireless channel, without any of the metal bulkheads otherwise in the path of the boot or engine nodes.

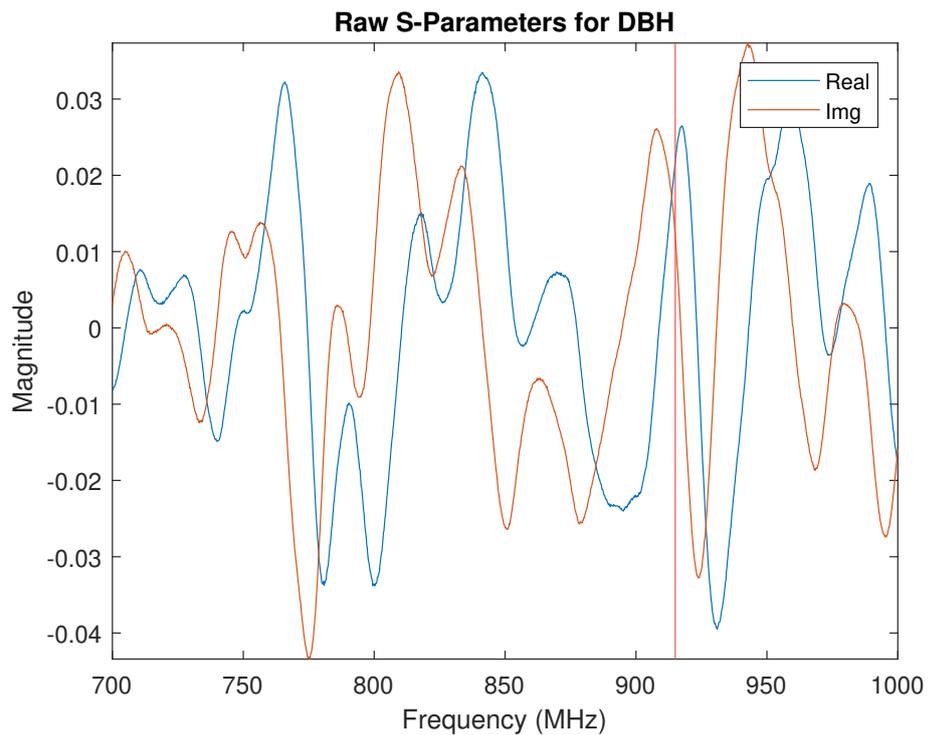


Figure 7.11: Complex S_{21} between the driver’s console and the boot with a horn antenna. Red line indicates 915 MHz.

The results from Figure 7.11 for each channel combination were then processed by the equation in (7.20). As the environment within the vehicle

is static, and therefore little variation in the multi-path propagation takes place, only one measurement for each node configuration was taken. The raw K-factor values, calculated from the values measured in the DBH case, are shown in Figure 7.12.

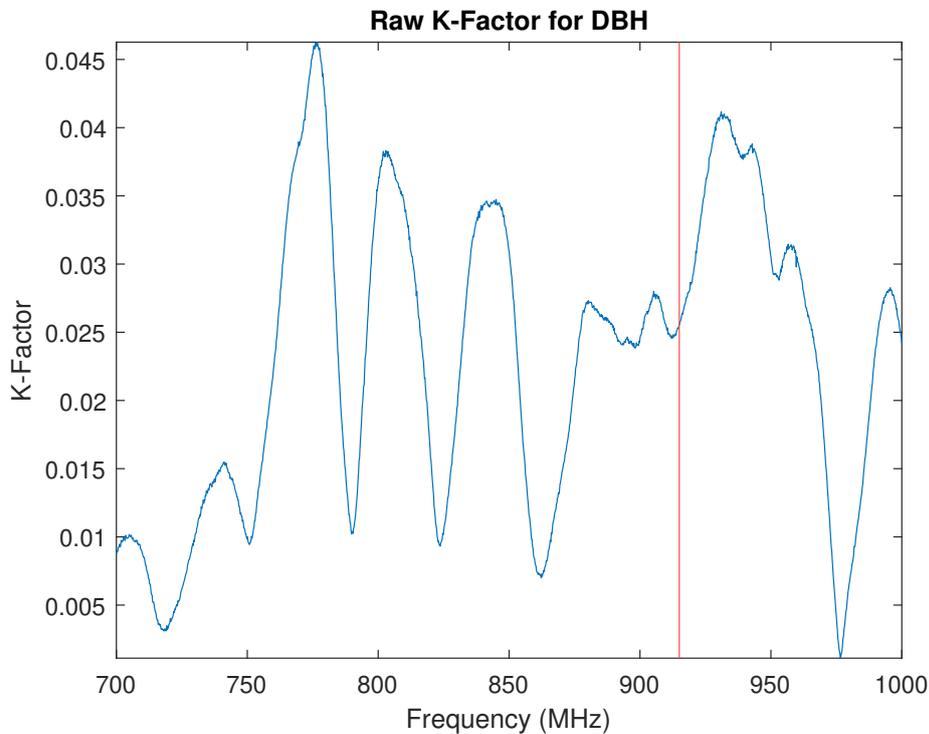


Figure 7.12: Raw K-factor values for DBH. Red line indicates 915 MHz.

Figure 7.12 shows a highly varying K-factor across the frequency spectrum, which is not unlike other K-factor measurements taken to measure wireless channels such as those found in offices or trains [134]. The relative low value of the K-factor in this instance indicates that the direct line of sight (LOS) component was very low when compared to the reflected, multi-path components. In turn, this may suggest that the coherence bandwidth of the channel may limit the upper data transmission rate. Careful testing and validation of data transmission rates and BER inside the RC would determine if this were the case, should the wireless system not reach required performance levels. This is compounded by the raw K-factor values for the several selected test cases presented in Figure 7.13.

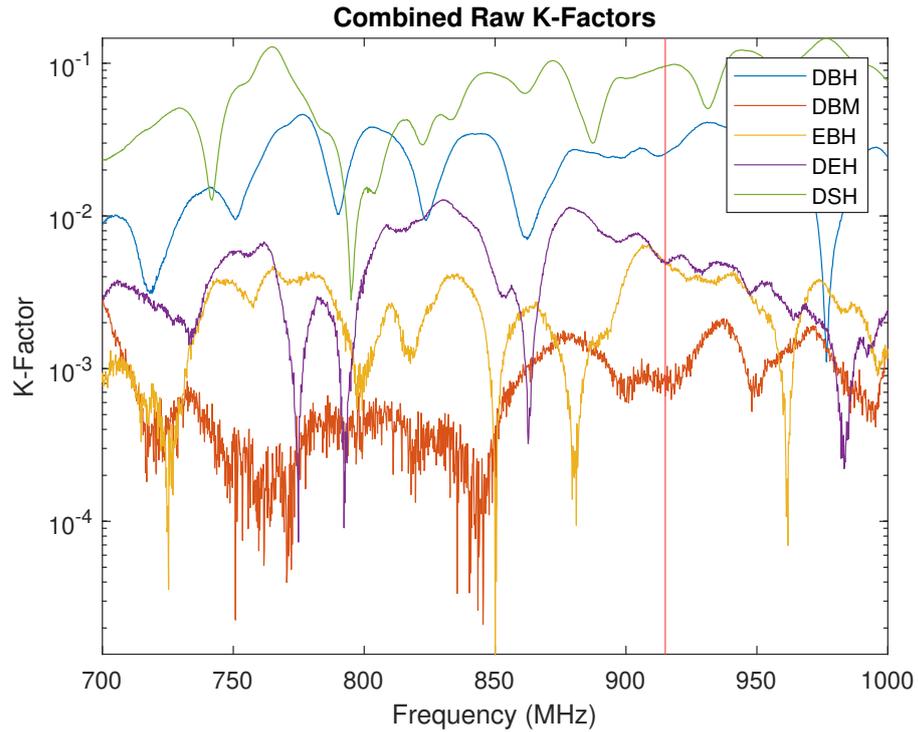


Figure 7.13: Selected K-factor values from around the vehicle.

These K-factors plots were chosen to represent a useful fraction of the intra-vehicular environment. The two driver-boot measurements were chosen as a wired channel of this nature was heavily utilised by the MLV in GIST, particular for monitoring vehicle and self-driving algorithm functionality. When using the horn antenna a relatively high K-factor is measured across the band, primarily due to the high directionality afforded by its construction. The monopole, in orange, shows an extremely low K-factor indicating a very large multi-path component in the received signal. Intuitively, this is due to the isotropic nature of a monopole antenna's radiation pattern, resulting in a smaller direct signal component. The effect of lowering the seats can be seen in general increase of K-factor, observed on the green line.

It is interesting to note that, whilst still very low in value, due to the large bulkheads between sections of vehicle, the K-factor measurements closely match in certain frequency bands of the driver-engine and engine-

boot channels. This is particularly true around the frequency of interest at 915 MHz. Therefore, it is possible to conclude that if one such test case were able to function, then the other would too, despite having to propagate through differing portions of the vehicle.

As the proposed test system only uses a single carrier frequency (as opposed to using multiple, wide-band carrier frequencies such as those often used in techniques such as FHSS) it is only required to replicate the K-factor at the same frequency in the RC. This serves to save a great deal of time, as whilst matching band-wide portions is entirely possible, it is extremely time consuming due to the high sensitivity of the RC to loading and positioning of loading elements and antennas. The precise values of the K-factor at 915 MHz, are therefore presented in Table 7.3, and were used to correctly load the RC to replicate the measured intra-vehicular environments.

Table 7.3: K-factor values at 915 MHz for selected channels.

Channel	K-Factor at 915 MHz
DSH	0.025
DBH	0.095
EBH	0.0049
DEH	0.00066
DSM	0.0049

The measurement of S_{21} was repeated using the same antennas in the RC. The K-factor for the particular antenna configuration was again calculated using (7.20). In order to adjust the K-factor inside the RC, multiple different antenna positions and orientations were tested, along with varying the number and position of the pyramidal RF absorber material inside the RC. Furthermore, altering the position of the RF stirrer would significantly alter the K-factor. It was found that even small differences in any of these parameters would result in very different K-factor values and responses across the frequency band. Multiple different configurations were tested and the position of all equipment was carefully noted for reference.

An example of the sensitivity of the K-factor to loading is shown in Figure 7.14, where the only difference between the experimental configuration in position 1 and position 2 was the removal of one RF absorber from the RC.

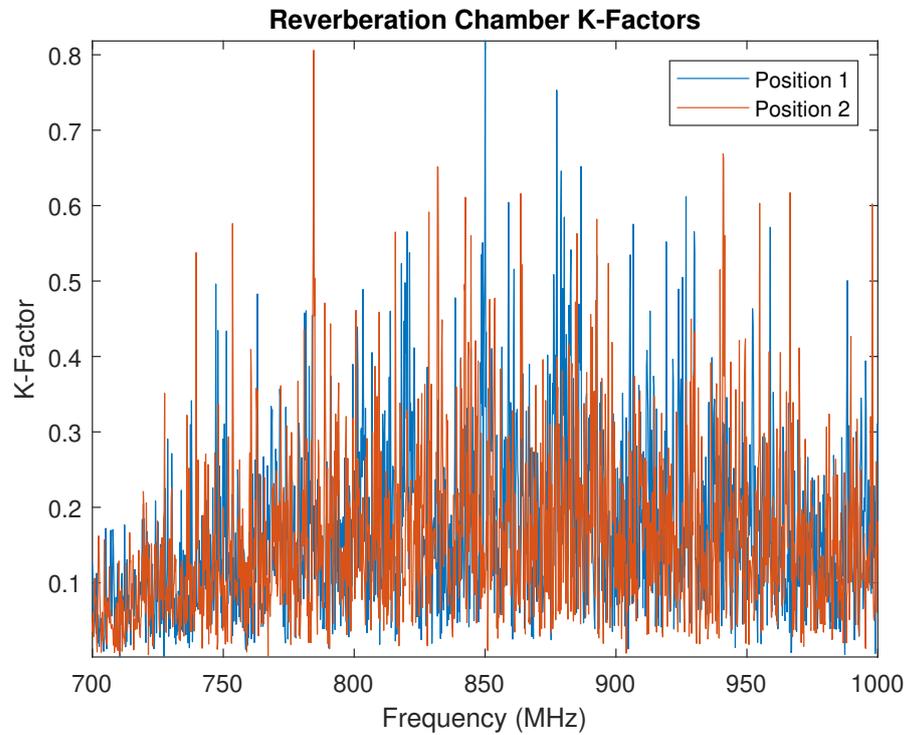


Figure 7.14: Difference in K-factor in the RC for two test configurations.

It is also worth noting the high amount of noise in the calculation - itself a result of high noise in the S-parameter measurement. This response is a result of the assumption that the only non-stirred component in the RC is the direct coupling term as in [138], fundamentally impossible to achieve in a real RC. The work in [138] suggests the use of curve smoothing to better visualise the data, an example of which can be seen in Figure 7.15. In this work, the smoothing of the data was accomplished using a moving mean method over a window of 20 samples. The window size was determined to provide an optimum balance between smooth out the high frequency noise components, whilst still allowing overall trends in the data to become apparent.

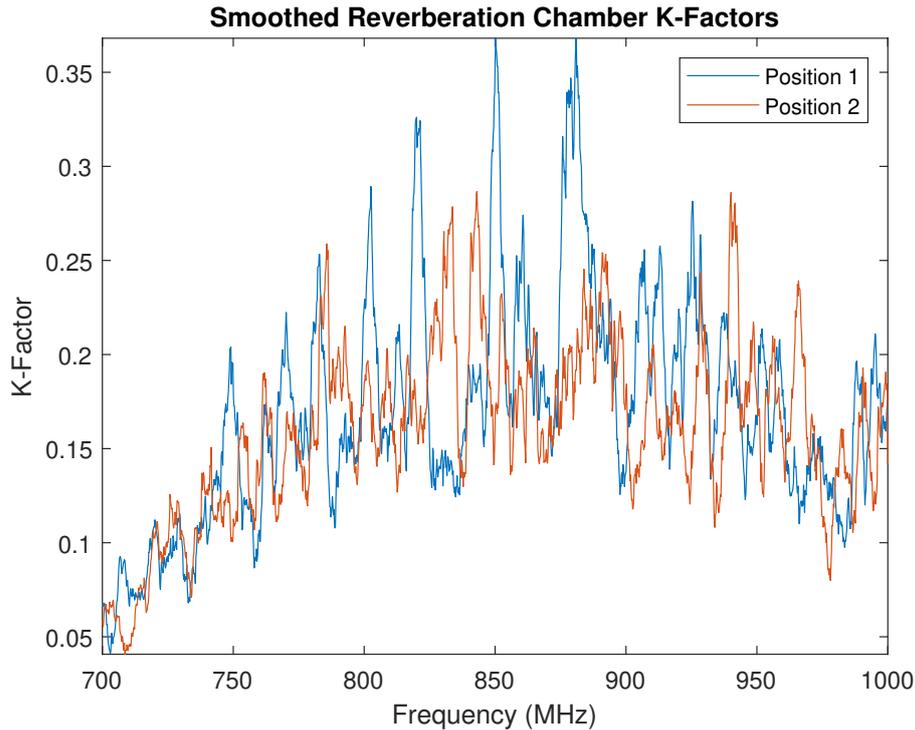


Figure 7.15: Smoothed K-factors in the RC for two test configurations presented in Figure 7.14.

There were substantial difficulties in recreating the K-factor values from Figure 7.13 across the measurement band for each wireless channel, due to the sensitive nature of the RC. As a result, the decision was made to recreate the K-factor only at the frequency of interest. Furthermore, it was found that reliably recreating a K-factor below a value of around 0.008 was extremely difficult in the large RC, owned by the GGIEMR. Therefore, only two test cases were able to be measured with a high degree of reliability, as presented in Table 7.4. In both the DBH and DSH test cases, the difference between the K-factor calculated from measurements in the car and the RC was below 0.002.

Table 7.4: Measured and RC K-factor values

Channel	Vehicle K-Factor	RC K-Factor	Δ K-Factor
DBH	0.025	0.026	0.001
DSH	0.095	0.093	0.002

However, whilst the experimental configurations were very close to the calculated K-factors from the intra-vehicular environment at the specified 915 MHz, Figure 7.16 shows the differences between these values across the measurement band.

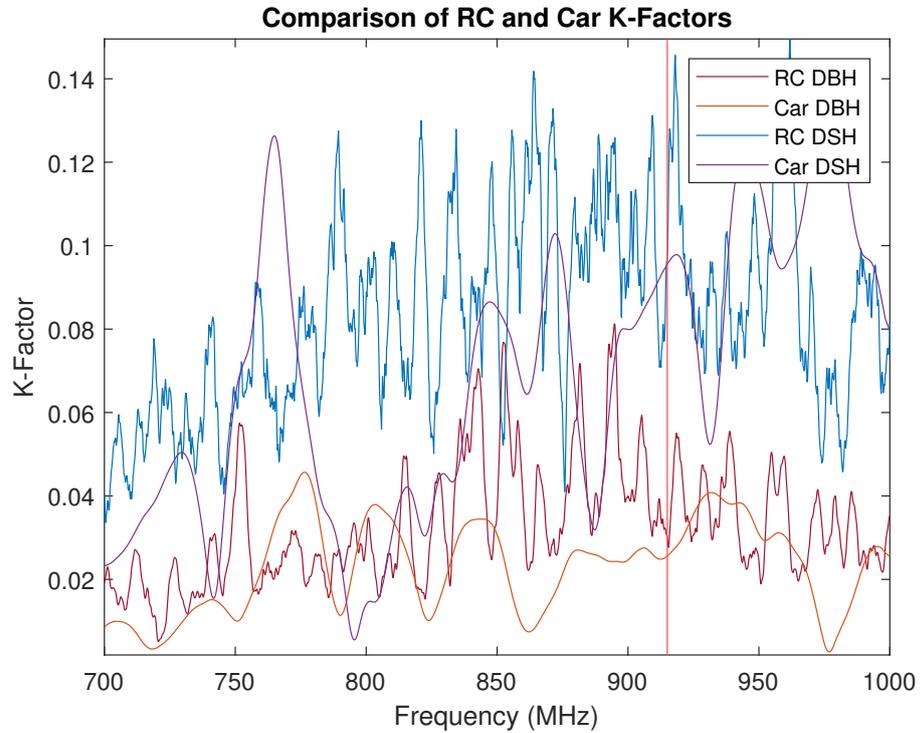


Figure 7.16: Comparison of RC and car K-factors. Red line indicates 915 MHz.

From this graph it is possible to see there are multiple points across the measurement band at which the corresponding channel pairs match and do not match. For example, the entirety of the RC K-factor for the DBH channel reasonably approximates throughout the spectrum, aside from being a little higher in value between 850 - 900 MHz. On the other hand, the RC replication of the DSH channel is consistently higher in value throughout most of the band, with much less low frequency variation. Despite this, it does regularly approximate the car's environment well at multiple frequencies and should therefore be suitable for experimental use, providing the limitations are considered.

Whilst this indicates that the experimental setups were not suitable for wide-band replication of the intra-vehicular environment, the narrow band K-factors were more than suitable for the fixed carrier frequency operation of the proposed QPSK system in the RC.

7.3 Wireless Transmission Performance

With the RC properly configured for the DSH channel case ($K = 0.093$), the USRP was configured to take a 10 second long measurement at 1 Mbps to verify the stability of the system, and its resilience to bit errors. This was repeated for varying gain levels of the USRP Tx module, which was easily reconfigured in the Simulink model (see Figure 7.2). Initial measurements of the spectra, shown in Figure 7.17 around the 915 MHz carrier frequency showed several clear spectra across all Tx gains. This measurement was made with an Agilent E4403B spectrum analyser.

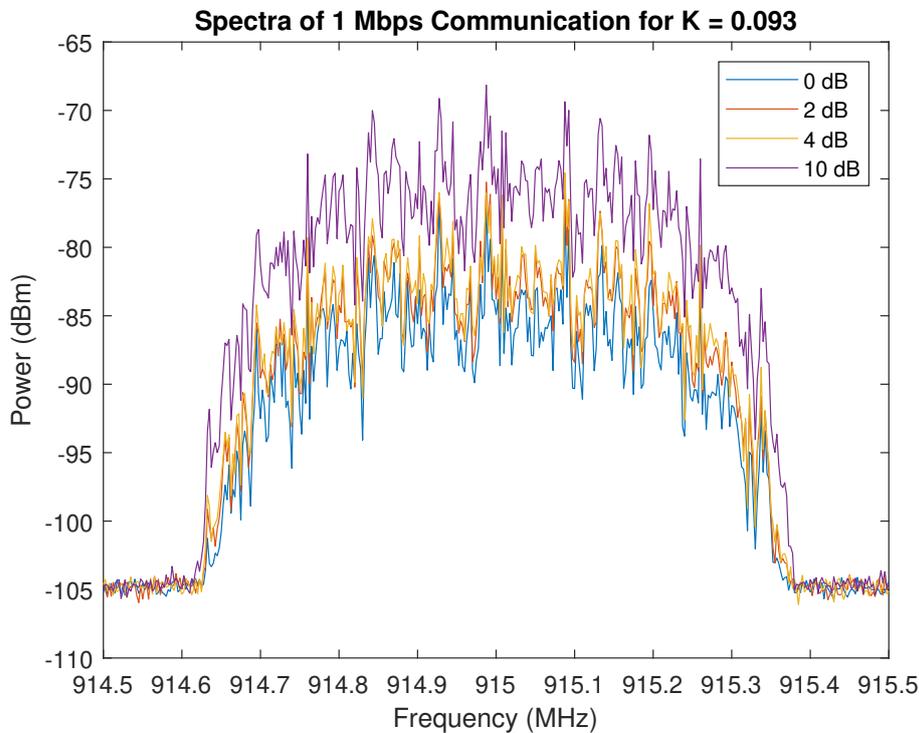


Figure 7.17: Transmission spectra for varying Tx gain for $K = 0.093$.

Each spectra is distinct and without obvious harmonics or distortion. The general trend of the measured spectra's power indicates proper functionality of the USRP's Tx gain control. It is possible to see that, whilst the full 1 MHz bandwidth of each spectrum is, at least, just distinct from the background, the roll-off of the 0 - 4 dB spectra is far less steep than the 10 dB, suggesting these gain levels are on the boundary of what might be suitable for reliable data transmission.

BER measurements were also taken at approximately 0.25ms intervals, allowing BER to be viewed as a function of time as per Figure 7.18.

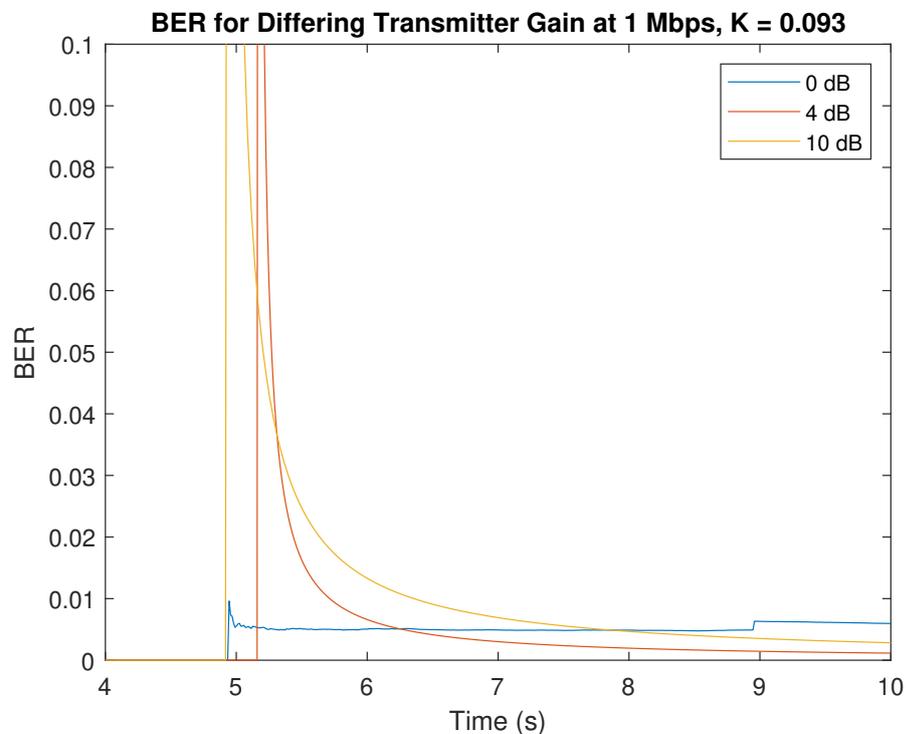


Figure 7.18: BER as a function of time for 1 Mbps transmission through the DSH channel.

It should be noted that due to variations in the Simulink model's interactions with the USRP, the actual start of data transmission can vary subtly between test runs. Therefore the start of the data transmission (indicated by the high jump in BER) in Figure 7.18 should be considered relative throughout the temporal axis. From Figure 7.18 it is clear to see

the effect that too low a gain will have on the performance of the system - the 0 dB case appears to quickly converge on a steady-state BER of around 5×10^{-3} , which would be far too high for most electronic systems to function properly - especially safety-critical systems. However, both the 10 dB and 4 dB plots indicate a high BER during the initial turn-on of the system, which quickly reduces through time. Indeed, aside from the initial errors (during the first 0.5ms, with the vast majority occurring in the first sampling period of 0.025ms), no other errors are observed during the experimental run. This indicates that both the 4 dB and 10 dB cases would be suitable for the reliable wireless transmission of data through the DSH channel in the vehicle.

Also visible in the 0 dB test case is an increase in BER around nine seconds into the experiment. This is an unfortunate result of bandwidth limitations in the Ethernet connection between the USRP and host PC, resulting in overflow conditions and subsequent errors in the transferred data. This BER is not a result of errors being induced in the wireless channel under test (which may otherwise rule against the reliability of the channel configuration), but is a result of the bottleneck between the USRP and host PC. This limitation exists due to the high throughput of data between USRP and host PC, especially when the wireless system is operating at high data rates. As a result, measurements with data rates exceeding 1 Mbps were not achievable with the current host PC hardware, but the somewhat artificial 1 Mbps limit does indicate that the proposed wireless system is capable of *at least* matching the *highest* specified data rate of CAN, as per the specification in [61]. It should be highlighted that this limitation is, with the current hardware, indicative of limitations of the experimental setup (more specifically, the host PC's hardware), and does not necessarily indicate upper data rate of the same wireless network with more capable rear-end control.

The experimental run was repeated, this time replication the DBH channel ($K = 0.026$) at 1 Mbps. The spectra, as shown in Figure 7.19 shows reasonable spectral form, though higher Tx gains were used when compared to the DSH channel.

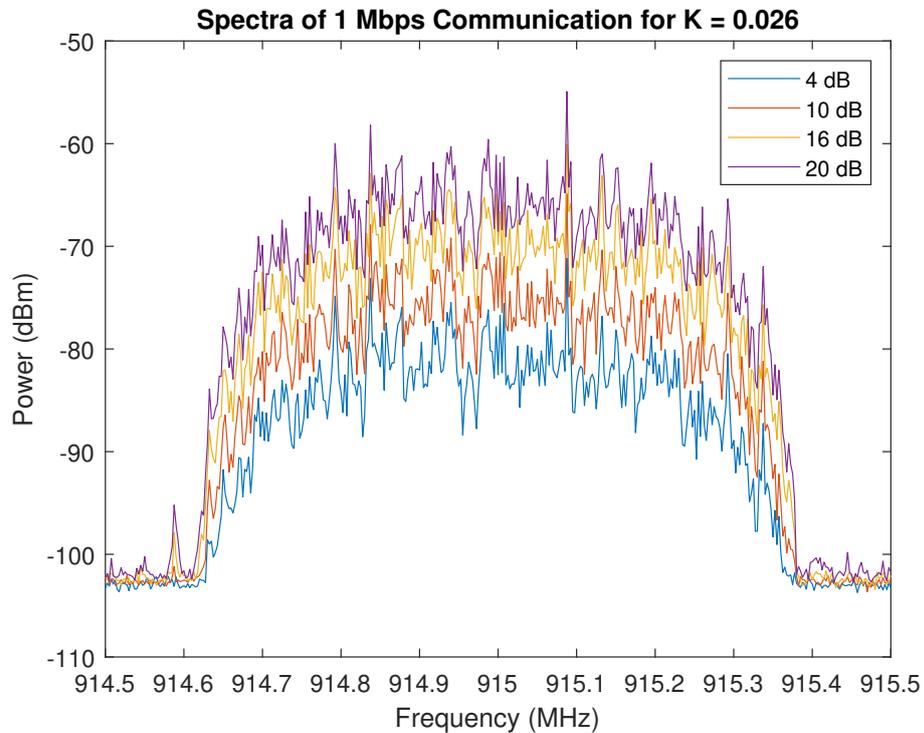


Figure 7.19: Transmission spectra for varying Tx gain for $K = 0.026$.

As per Figure 7.17, the separation of spectra is apparent from the Tx gain levels, but a spectral spur is visible at 914.6 MHz. This is likely a result of the imperfect nature of the raised cosine filter used to shape the QPSK spectra, which is more apparent in measurements as the gain is increased. In the tested cases, it is a relatively small component and therefore is of little concern, though care would need to be exercised for higher gains to ensure the additional component would not interfere with neighbouring systems.

The BER measurements also indicate a similar response versus time between the DSH and DBH channels, as shown in Figure 7.20. One apparent difference between individual plots is the BER steps at 4 dB. Unlike the single step in for the 0 dB case in Figure 7.18, these steps are the result of

the frequency acquisition stage losing lock due to a high signal to noise ratio (SNR). This distinction between BER steps caused by failure of the system through the channel versus BER steps due to USRP to host PC limitations was apparent during the experiment as no overflow warnings were logged during the tests.

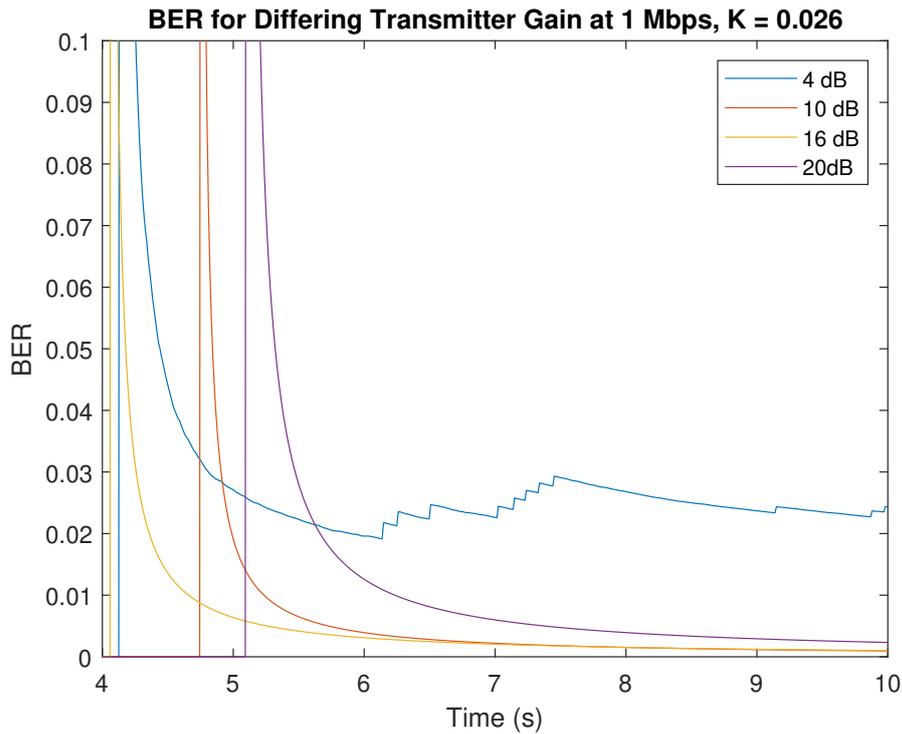


Figure 7.20: BER as a function of time for 1 Mbps transmission through the DBH channel.

Aside from this, a similar BER response is observed as in the DSH test channel with a large BER during initial startup and no subsequent bit errors after 5×10^{-3} seconds. This suggests that the higher amount of reflected signal components, arising from the lower K-factor, results in a higher SNR in all test cases. The higher gain values of more than 10 dB in the DBH test case mitigates this response. However, it must be considered that whilst using such a brute-force method is beneficial to providing reliable communication in laboratory conditions, using such a technique may prove complicated in the real-world application of such systems as the use of higher

transmission power is less efficient from both an energy and spectral-usage perspective.

Another set of measurements were taken using the same DBH test configuration whilst running at a reduced data rate of 400 kbps. Whilst contemporary intra-vehicular networks, such as CAN, are specified to run at up to 1 Mbps, they are capable of operating at lower data rates to improve the reliability of data transmission. The spectra, as shown in Figure 7.21, again shows excellent spectral form with minimal artifacts.

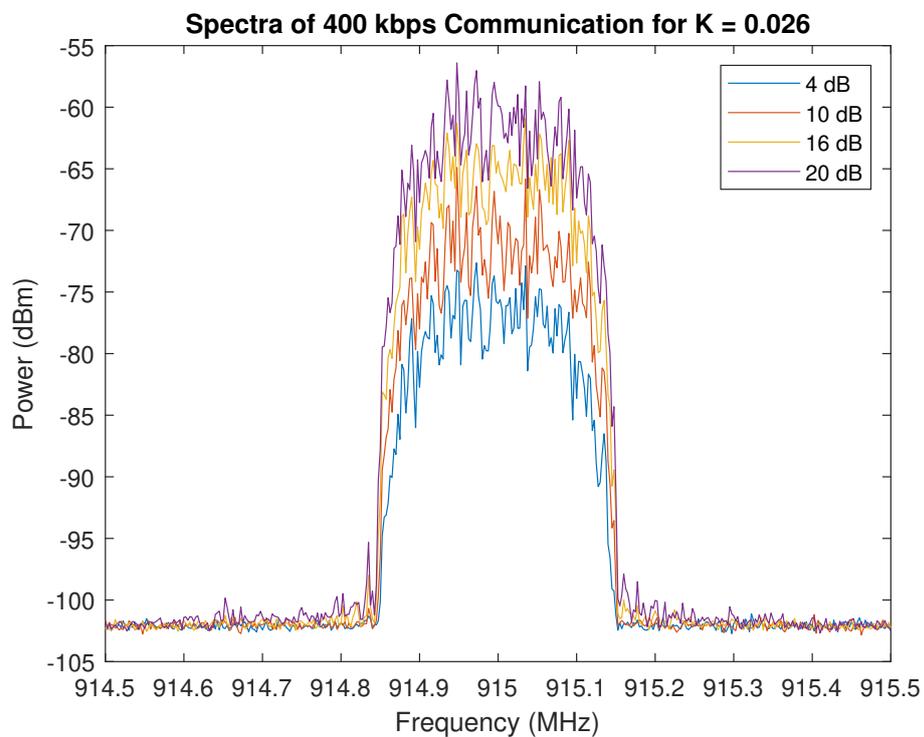


Figure 7.21: Transmission spectra for varying Tx gain for $K = 0.026$ at a reduced data rate.

Note the greatly reduced occupied bandwidth of around 400 kHz and steeper band roll-offs. The peak power levels are generally similar to the 1 Mbps test cases, but each spectrum has an improved SNR as a result of the narrower bandwidth - itself a result of the lower data rate. In terms of spectral management and maintaining a suitable SNR at low powers, it is clear from these measurements that wireless intra-vehicular networks would

greatly benefit from the ability to vary data rate during operation. Whether this is to provide greater resilience from noise at equal or lower transmission power, or facilitate efficient band usage in congested environments, simply lowering the data rate on a given wireless QPSK system would offering many practical advantages.

The plots of BER against time in Figure 7.22 highlight the importance of thorough testing and validation of the system's performance. From casual inspection of the spectra in Figure 7.21, one may infer that the SNR was sufficient to allow low-error system operation. However, it can be clearly seen that the 4 dB test case loses carrier synchronisation after only 1 second of operation and completely fails to resynchronise throughout the remainder of the experiment. However, both the 16 dB and 20 dB experiments show the same, reliable operation as in the previous measurements, at BERs around an order of magnitude lower.

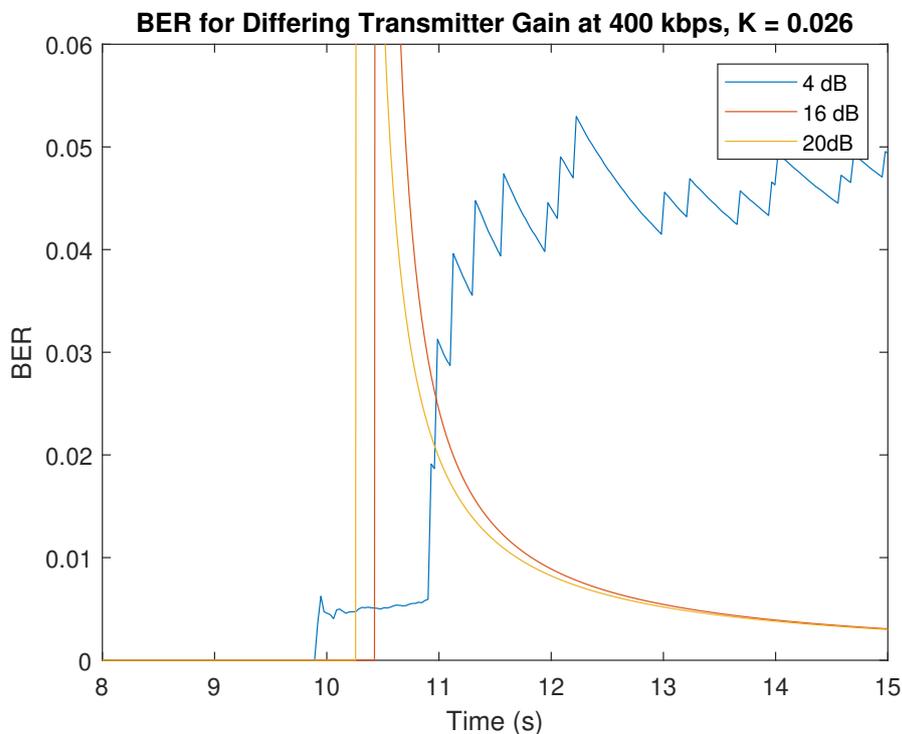


Figure 7.22: BER as a function of time for 400 kbps transmission through the DBH channel.

Characteristically, the plots of BER for the fully operational systems starts with a large BER step during startup. This large initial spike in BER in all operation test cases begs the question - if the number of errors is confined only to startup, what is happening during startup that causes these errors? This can be answered by examining the constellation diagrams.

The constellation diagram for the 16 dB gain at 400 kbps for the DBH test is presented in Figure 7.23. This specific constellation diagram also has a logarithmic colour filter applied to each individual sample, with early samples being a deep red and later samples being yellow, permitting observations on sample integrity through time.

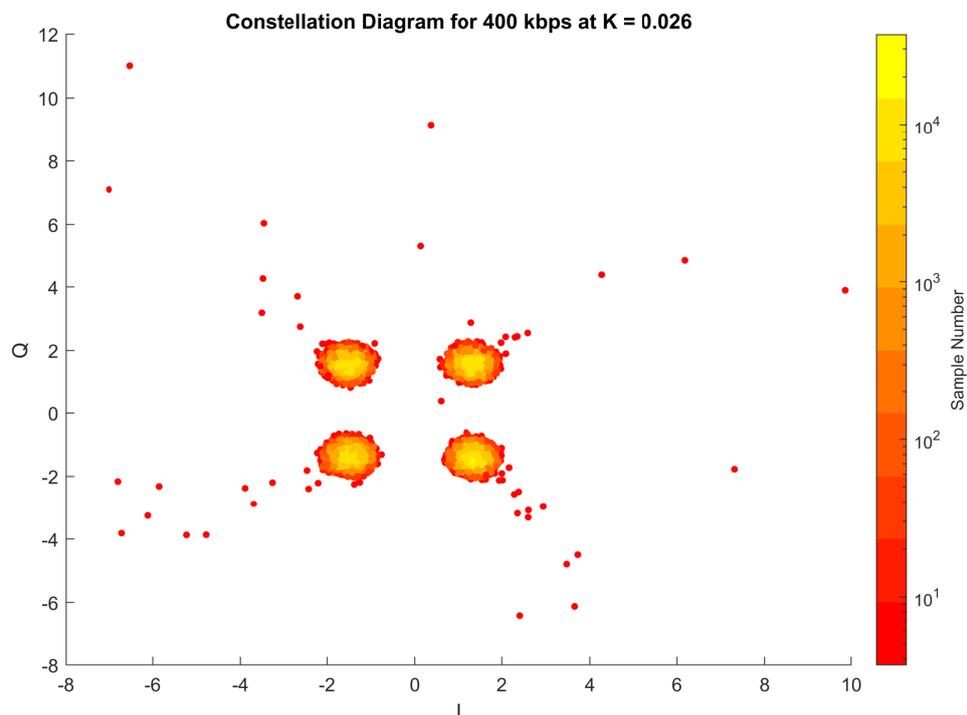


Figure 7.23: Constellation diagram with sample-ordered colour grading.

The four symbol locations are readily visible about the origin of the graph and each symbol location has sweeping trails of outlying samples. The USRP receiver would fail to properly decode these symbols, resulting in an increase of the BER. Figure 7.23 also indicates that *all* of these symbol outliers occur extremely early in the experimental run, with all outliers

occurring within the first 1000 samples. It can therefore be inferred that these outliers are the contributing errors in the BER plots against time.

The spiralling nature of the early samples indicates how the system is responding to these error-inducing samples. The size of the distances (called the error vector magnitude (EVM) between the idealised symbol coordinates and the outlier samples indicates the functionality of the AGC through time. Initially, samples have a very large EVM but through time this decreases as the AGC converges to the predefined output levels. No AGC is capable of an instantaneous response, though through careful optimisation the convergence period may be reduced further. The phase between of the samples reveals the operation of the frequency correction module. As small offsets in the carrier frequency are detected and cancelled by the USRP, the phase between the ideal and measured constellation point reduces. This is visible in Figure 7.23 by the rotation of samples about the ideal constellation points through time. This plot highlights that all bit errors are the result of an unoptimised receiver configuration for the start up period and that during normal, constant use, the system is highly reliable even at relatively low power.

7.4 Chapter Conclusion

This work indicates that reliable wireless networks are indeed feasible in the ordinary automotive environment. Within a short time frame, a working wireless system was created which was found to operate successfully through channels which were replicated in a reverberation chamber using real-world measurements around a car. However, not all channels were able to be replicated in the GGIEMR's reverberation chamber due to the low K-factors associated with transmission through particularly opaque channels such as the engine to driver. Future work in better-suited reverberation

chambers would allow for a more comprehensive survey of channel suitability throughout the car, though operating power levels would likely have to increase to reliably transport data - as indicated by the comparison of the two test cases in this chapter.

The work in this chapter demonstrates that, whilst there may be few commercially available wireless transceivers which would be suitable for at least matching the data rates of CAN in an intra-vehicular case, the actual design and specification of reliably working systems is not terribly difficult. Though this work used expensive USRP radios for design and evaluation of the wireless transmission parameters, it is becoming ever easier to design and manufacture specific hardware which would be more economically feasible to implement in consumer cars.

However there is more than just technical challenges to overcome before the widespread adoption of any wireless system can be realised, and that is primarily a result of the extremely tight legislation and safety standards imposed on the automotive industry. Whilst it may appear to be relatively trivial to implement a working wireless intra-vehicular network, it is another challenge to ensure it *always* works in a variety of environmental conditions. The multitude of testing and redesign of wireless systems would require at least some degree of model-specific testing and validation to ensure reliable operation, as even small variations in the configuration of the vehicle can result in noticeably different transmission channel behaviour.

Overall, this work shows that a new wireless alternative to wired communication is entirely possible using existing technology and is already capable of at least matching the data throughput of contemporary wired solutions. Limitations imposed by PC hardware and reverberation chamber parameters limited the breadth of the survey, but through careful measurement and proper interpretation of data from multiple domains, it is evident that the system presented in this work is capable of exceeding the

data throughput of some contemporary networks. System parameters are available in Appendix E and the measurement data for calculating K-factors in intra-vehicular cases can be found in Appendix F, which can be used by the wider academic community for research purposes.

Chapter 8

Summary and Reflections

This work has described the benefits of implementing wireless communication networks in the cars of tomorrow. Whether the benefits of a reduction in metal mass are felt by the manufacturer in reduced construction costs, the end user through reduced operating costs, or the environment through a reduction in greenhouse gas emissions, there is a clear motivation behind reducing the physical weight of components inside vehicles. Even if incremental steps are taken by reducing the mass of shielding, a partial switch of non-safety-critical networks to wireless, or improving the efficiency of common EMI emission sources, positive benefits would be felt by all throughout the car's lifetime.

The research, undertaken to construct this thesis, has not only directly facilitated the ability to make multiple quantitative and qualitative analyses regarding the feasibility of introducing wireless networks into the cars of tomorrow, but has also resulted in multiple data-sets describing the EM environment in the intra-vehicular case. Furthermore, insights into taking a more holistic approach to quantifying the effect of time-domain interference on a communications network have been made - potentially benefiting digital communications not only in the intra-vehicular case, but in a more general sense. This work has itself lead to the development of highly capable,

open source test equipment for the wider research community to utilise in BER measurements.

As described in Chapter 1, the objectives of this thesis can be expressed as follows:

1. To develop an understanding of the impact an increase in the electronic equipment, necessary for electric and autonomous cars, will have on the EM environment inside and around future cars.
2. To examine the effects of time-domain interference on wired communication links, looking at response of contemporary communication standards to interference, as well as develop an understanding of the influence of EMI on a system as a whole.
3. To develop an open-source means of testing the errors on communication networks under the influence of EMI.
4. To discover the feasibility of implementing a robust wireless communication link for intra-vehicular communications.

This thesis conclusion will, in turn, examine each of these objectives and reflect on the contributions of this work towards meeting the objectives.

8.1 Contributions and reflections

In order to develop an understanding of the EM inside and around future cars, an expansive EM survey was performed around a cutting-edge, autonomous, electrically-driven car at the MLV at GIST. After performing a qualitative analysis of the rationale behind the historical inclusion of electronic subsystems to cars, a re-emergent trend towards the inclusion of electronic systems for the purpose of safety and basic functionality was

identified, allowing for an academic reasoning for the importance of performing new EM environment testing around such a complex vehicle. The measurements themselves indicated the potential for susceptibility to EM emissions in cars of the near-future through identification of both congested and quiet frequency bands, and highlighted some surprising examples of the variability of EM emissions across small spatial variations across the vehicle. For example the motor compartment was one of the quietest test sites throughout the vehicle, whilst the boot (with its large collection of power converters, serves, wireless equipment, etc.) was the noisiest.

Furthermore, time-domain measurements were able to indicate the susceptibility of the CAN bus to damaging levels of EMI. Indeed this damage was not limited to just the information being transported through the network, but also to the hardware itself. These sources of time-domain interference highlighted the importance of better understanding the effect such interference has on victim networks.

As a result, multiple laboratory experiments were conducted to examine the effects of time-domain interference on wired communication links. It was found that the effects of faster rising-slope edge on a victim network were not directly proportional to the gradient of the source slope, when the effects are considered *after* the arbitration of the digital representation of a bit from the analogue representation. By considering the effect that this time-domain interference has on a system as a whole, it was shown that the conventional approach to minimise the gradient of the source slope could increase the number of bit errors of a system. Measurements demonstrating this were faithfully replicated using simulations, and the complexity of the relationship between bit error probability and the gradient of a interference waveform were further corroborated by other researchers using a separate, simulated approach. Although simulation methods at present are fairly primitive in their modelling of some system components such as the receiver,

it is anticipated that further development may result in a useful tool for modelling the BER of systems under different network configurations and time-domain emission sources.

It was described how such a relationship may be exploited in real-world systems to improve the efficiency of traditionally problematic emission sources. No longer must inefficiencies in power converters or data transport links must be tolerated, if a certain BER is acceptable - though such systems must still be carefully engineered.

In order to perform BER measurements at the physical layer specifications required for this work, an open-source BERT was developed. The instrument was developed to facilitate an easy means of performing gathering initial BER measurements, which in this thesis, allowed for the examination of the BER relationship with time-domain waveforms. OSBERT itself was also designed to facilitate easy expansion of the core device for more complex BER measurements in future work. The final BERT was not only more capable than some instruments in its class, but was also significantly cheaper and readily available to all persons through the complete release of all design files and documents.

Finally, an experimental approach was used to examine the feasibility of implementing wireless communication networks in the intra-vehicular case. Furthermore, new measurements were made around a car to replicate the EM environment of the inside of the car inside an RC. Whilst the findings of this experimental regime were somewhat limited by the properties of some of the equipment (such as the RC and the PC), it was shown that a highly reliable, wireless communication network ought to be practically feasible - at least in some communication channels. However, as discussed in the next section, more work is needed before such networks are likely to be implemented in commercial cars.

8.2 Future Work

There are a number of avenues for continuing the work in this thesis. From a wired communication perspective, developing simulation tools for more accurately determining the magnitude of the BER of systems exposed to sources of time-domain interference would prove beneficial to the wired intra-vehicular communication networks of tomorrow. In order to do this a more refined receiver model should be developed. This need is evident through the comparison of measured to simulated outputs from the cable model, transfer function, and induced interference waveform which have been shown to be highly capable of accurate modelling. As a result, discrepancies in the magnitude of the final BER output must be due to a simplistic receiver model. Work by A. E. Penal and K. Niewiadomski has started this process, though it is still in its infancy.

Further work will also be needed to examine the practicalities of exploiting these effects, observed under somewhat contrived conditions in a laboratory, on real-world systems. Whilst peaks and, perhaps more importantly, troughs can be clearly distinguished on the BER plots in this thesis, the magnitude of the BERs is still considerably higher than that encountered in real-world networks. If troughs may be observed more frequently at commercially-acceptable BERs, such as those in Appendix B then the case for the additional design complexity can be made stronger.

From a wireless perspective, the results in this thesis demonstrate that it is not necessarily a difficult technical challenge to implement a working wireless network in a intra-vehicular case. However, non-technical factors are likely to be the greatest inhibitor of immediate change in the design methodologies of commercial car manufacturers. As discussed in this work, the automotive industry uses a variety of strict EMC and reliability standards to ensure the highest degree of safety in commercial cars. Whilst

this is a generally positive design philosophy to implement, it comes at the cost of reducing the rate at which new technologies may be realised. Therefore, a great deal of additional work is required to not only increase the breadth of experimental cases (thereby addressing the limitations incurred in this work), but to also convey the reliability of such networks to car manufacturers in order to begin the process of implementing new styles of intra-vehicular networks.

Bibliography

- [1] P. Wadhvani and S. Yadav, “Embedded Systems Market Size By Component (Hardware, [ASIC & ASSP, Microcontroller, Microprocessor, Power Management Integrated Circuit (PMIC), Field Programmable Gate Array (FPGA), Digital Signal Processor (DSP), Memory], Software (OS, Middleware)], By Function (Standalone System, Real-Time System, Network System, Mobile System), By Application (Automotive, Consumer Electronics, Manufacturing, Retail, Media & Entertainment, Military & Defense, Telecom), Industry Analysis Report, Regional Outlook, Application Potential, Competitive Market Share & Forecast, 2020 – 2026,” *Global Market Insights*, 2020.
- [2] “Connected and Autonomous Vehicles 2019 Report/Winning the Global Race to Market,” *The Society of Motor Manufacturers and Traders*, 2019.
- [3] SAE, “J3016.201401 Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems,” 2014.
- [4] R. W. Cox, “Local Area Network Technology Applied to Automotive Electronic Communications,” *IEEE Transactions on Industrial Electronics*, vol. IE-32, no. 4, pp. 327–333, 1985.
- [5] O. Kiyotsugu, “Wiring Harnesses for Next Generation Automobiles,” *Fujikura Technical Review*, 2013, pp. 77–80, 2013.

- [6] macro trends, “Copper Prices - 45 Year Historical Chart,” 2020 (Last accessed 09 March 2020). <https://www.macrotrends.net/1476/copper-prices-historical-chart-data>.
- [7] X. Zhiran, *Cost Reduction in the Automobile Industry — Case Studies of the Chinese market*. PhD thesis, University of Halmstad Halmstad Halmstad, 2012.
- [8] Capgemini, “Improving Cost Management in the Automotive Supplier Industry,” 2010 (Last accessed 22 July 2020). <https://www.capgemini.com/wp-content/uploads/2017/07/Improving-Cost-Management-in-the-Automotive-Supplier-Industry.pdf>.
- [9] International Organization of Motor Vehicle Manufacturers, “World Motor Vehicle Production by Country and Type,” 2019 (Last accessed 09 March 2020). <http://www.oica.net/category/production-statistics/2018-statistics/>.
- [10] Ricardo, “Impact of Vehicle Weight Reduction of Fuel Economy for Various Vehicle Architectures,” *The Aluminium Association Project FB769*, 2008.
- [11] D. Ford, “As Cars are Kept Longer, 200,000 is New 100,000,” *New York Times*, 2012.
- [12] J. McDonald, “China to ban petrol and diesel cars, state media reports,” *Independent*, 2017.
- [13] NA, “France Set to Ban Sale of Petrol and Diesel Vehicles by 2040,” *BBC News*, 2017.

- [14] D. Sheridan, “Boris Johnson to Speed Up Ban on Petrol and Diesel Cars,” *The Telegraph*, 2020.
- [15] J. C. Maxwell, “A dynamical theory of the electromagnetic field,” *Philosophical Transactions of the Royal Society of London*, vol. 155, pp. 459–512, 1865.
- [16] M. S. Rogalski and S. B. Palmer, *Advanced University Physics*. 6000 Broken Sound Parkway NW, Suite 300, Boca Raton, FL 33487-2742: CRC Press, 2 ed., 2006.
- [17] C. Möhring, *Eine Geschichte des Blitzableiters. Die Ableitung des Blitzes und die Neuordnung des Wissens um 1800 [A History of the Lightning Rod. The Invention of the Lightning Rod and the Restructuring of Knowledge in the Period around 1800]*. PhD thesis, Bauhaus-Universität Weimar, Fakultät Medien, 7 2005.
- [18] J. Goedbloed, *Electromagnetic Compatibility*. Campus 400 Maylands Avenue, Hemel Hempstead, Hertfordshire, HP2 7EZ: Prentice Hall International (UK) Ltd, 1 ed., 1992.
- [19] F. Silva, M. Quilez, J. Llop, X. Torres, and P. Riu, “EMI From and Automotive CAN Bus,” *1999 IEEE International Symposium on Electromagnetic Compatability*, pp. 512–516, 1999.
- [20] J.-M. Liu, Y.-C. Huang, Y.-C. Ying, and T.-H. Kuo, “Slew-rate controlled output stages for switching DC-DC converters,” *2011 IEEE International Conference on IC Design & Technology*, pp. 1–4, 2011.
- [21] M. Kareppagoudr, E. Caceres, Y. Kuo, J. Shakya, Y. Wang, and G. C. Temes, “Passive slew rate enhancement technique for Switched-Capacitor Circuits,” in *2019 IEEE 62nd International Midwest*

- Symposium on Circuits and Systems (MWSCAS)*, pp. 913–916, Aug 2019.
- [22] M. Blank, T. Glück, A. Kugi, and H. Kreuter, “Digital Slew Rate and S-Shape Control for Smart Power Switches to Reduce EMI Generation,” *IEEE Transactions on Power Electronics*, vol. 30, pp. 5170–5180, Sep. 2015.
- [23] K. Malarić, *EMI Protection for Communication Systems*. Artech House, 1 ed., 2010.
- [24] W. Hart and J. A. Buck, *Engineering Electromagnetics*. McGraw-Hill Education, 9 ed., 2019.
- [25] Gu Chunhong and Li Shufang, “Shielding effectiveness of an enclosure with apertures,” in *2005 IEEE International Symposium on Microwave, Antenna, Propagation and EMC Technologies for Wireless Communications*, vol. 1, pp. 614–618 Vol. 1, Aug 2005.
- [26] “GORE® CAN Bus Cables,” Last accessed 31 March 2020. <https://www.gore.com/resources/canbus-cables-aircraft-datasheet>.
- [27] P. Horowitz, W. Hill, and T. C. Hayes, *The Art of Electronics*, vol. 2. Cambridge University Press, 1989.
- [28] K. Wesolowski, *Introduction to Digital Communication Systems*. Wiley, 2009.
- [29] L. Kirasamuthranon, J. Koseeyaporn, and P. Wardkein, “QPSK Modulator with Continuous Phase and Fast Response Based on Phase-Locked Loop,” *Radioengineering*, vol. 26, pp. 504–514, 06 2017.
- [30] M. Pursley, *Reference Data for Engineers*. Newnes, 2002.

- [31] “IEEE Standard for Information technology—Telecommunications and information exchange between systems Local and metropolitan area networks—Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” *IEEE Std 802.11-2016 (Revision of IEEE Std 802.11-2012)*, pp. 1–3534, 2016.
- [32] “Bluetooth Core Specification,” Standard, Bluetooth SIG, Dec. 2019.
- [33] “ZigBee Specification,” Standard, ZigBee Alliance, Aug. 2015.
- [34] R. Gottapu, *Cellular Mobile Communication*. Dorling Kindersley (India), 2013.
- [35] “Computer Chips Inside Cars,” (Last accessed 03 April 2020). <https://www.chipsetc.com/computer-chips-inside-the-car.html>.
- [36] R. W. Cox, “Local Area Network Technology Applied to Automotive Electronic Communications,” *IEEE Transactions on Industrial Electronics*, vol. IE-32, no. 4, pp. 327–333, 1985.
- [37] Feng Luo, Zhiqi Chen, Juexiao Chen, and Zechang Sun in *2008 IEEE Vehicle Power and Propulsion Conference, title=Research on FlexRay communication system*, pp. 1–5, 2008.
- [38] T. Herpel, B. Kloiber, R. German, and S. Fey, “Assessing the CAN communication startup behavior of automotive ECUs by prototype measurements,” in *2009 IEEE Instrumentation and Measurement Technology Conference*, pp. 928–932, 2009.
- [39] T. M. P. Simonik and J. Takac, “Principles and techniques for analysis of automotive communication lines and buses,” in *2014 ELEKTRO*, pp. 500–503, 2014.

- [40] *Automotive Repair Industry. hearings before the United States Senate Committee on the Judiciary, Subcommittee on Antitrust and Monopoly, Ninety-First Congress, first session, on Apr. 22-24, 29, 30, 1969.* U.S. G.P.O., 1969.
- [41] “Cadillac’s New Trip Computer,” *Texas Monthly*, p. 105, May 1978.
- [42] “Mercedes Benz and the Invention of the Anti-Lock Braking System (ABS) Ready for Production in 1978,” Jul 2008. <https://media.daimler.com/marsMediaSite/en/instance/ko/Mercedes-Benz-and-the-invention-of-the-anti-lock-braking-system-ABS-ready-for-production-in-1978.xhtml?oid=9913502>.
- [43] C. E. Dole, “Buick Goes for Performance. T-Type Cars Geared to Attract a Younger Buyer,” Jun 1986. <https://www.csmonitor.com/1986/0610/hriv-f.html>.
- [44] J. A. Zverina, “Ford, Motorola Form Major Technological Partnership,” Jun 1991. <https://www.upi.com/Archives/1991/06/11/Ford-Motorola-form-major-technological-partnership/5158676612800/>.
- [45] Jesda, “How In-Dash Navigation Worked In 1992 - Olds Was First,” Jan 2012. <http://jesda.com/2012/01/11/how-in-dash-navigation-worked-in-1992-olds-was-first/>.
- [46] B. Watson, “Chrysler Technological Innovations: The Big List of Chrysler Firsts.” <https://www.allpar.com/corporate/technology.html>.
- [47] Edmunds, “The Evolution of Front Airbags,” May 2009. <https://www.edmunds.com/car-safety/the-evolution-of-front-airbags.html>.

- [48] T. B. Lee and Utc, “Lidar Sensors are About to Become a Mainstream Car Feature,” Jan 2020. <https://arstechnica.com/cars/2020/01/lidar-sensors-are-about-to-become-a-mainstream-car-feature/>.
- [49] Microsoft, “Windows Network Architecture and the OSI Model - Windows Drivers,” Apr 2017.
- [50] Texas Instruments, *Automotive Logic Devices*. Texas Instruments, 2014. <https://www.ti.com/lit/ml/ssab009/ssab009.pdf>.
- [51] Texas Instruments, *TRS3223 Multichannel RS-232 Line Driver/Receiver*. Texas Instruments, 2008. <http://www.ti.com/lit/ds/symlink/trs3223-q1.pdf>.
- [52] L. Jiang and C. Yu, “Design and Implementation of Car Black Box Based on Embedded System,” in *2010 International Conference on Electrical and Control Engineering*, pp. 3537–3539, 2010.
- [53] T. Denton, *Automobile Electrical and Electronic Systems*. Routledge/Taylor & Francis Group, 2018.
- [54] R. De Andrade, K. N. Hodel, J. F. Justo, A. M. Laganá, M. M. Santos, and Z. Gu, “Analytical and Experimental Performance Evaluations of CAN-FD Bus,” *IEEE Access*, vol. 6, pp. 21287–21295, 2018.
- [55] T. Jiang, Z. Liu, and H. Zhou, “Design of FlexRay Network for In-wheel Independent Driving Electric Vehicle Distributed Control System,” in *Proceedings of the 32nd Chinese Control Conference*, pp. 7744–7749, 2013.
- [56] J. B. Buie, “Coupling Transistor Logic and Other Circuits,” Nov. 1 1966. US3283170A.
- [57] D. M. Harris and S. L. Harris, *Digital Design and Computer Architecture: From Gates to Processors*. Elsevier, 2007.

- [58] Texas Instruments, *KeyStone Architecture Universal Asynchronous Receiver/Transmitter*. Texas Instruments, 2010. <http://www.ti.com/lit/ug/sprugp1/sprugp1.pdf>.
- [59] M. Subaraman, “Communicating Two FPGA’s using UART,” *International Journal of Scientific & Engineering Research*, vol. 4, p. 633–636, Jun 2013.
- [60] Dallas Semiconductor, *Application Note 83 Fundamentals of RS-232 Serial Communications*. Dallas Semiconductor, 1998. <https://www.lammertbies.nl/download/dallas-appl-83.pdf>.
- [61] Bosch, *CAN Specification*. Bosch, 1991.
- [62] International Organisation for Standardisation, “Road vehicles — Controller area network (CAN) — Part 1: Data link layer and physical signalling,” 2015.
- [63] International Organisation for Standardisation, “Road vehicles — Controller area network (CAN) — Part 2: High-speed medium access unit,” 2016.
- [64] International Organisation for Standardisation, “Road vehicles — Controller area network (CAN) — Part 3: Low-speed, fault-tolerant, medium-dependent interface,” 2006.
- [65] Altera, “Simultaneous Multi-Mastering with the Avalon Bus,” 2002 (Last accessed 04 Jul 2020). <http://extras.springer.com/2001/978-0-306-47635-8/an/an184.pdf>.
- [66] J. Kime, “Can vs can fd: Efficiency,” 2016 (Last accessed 04 Jul 2020). <https://www.ixxat.com/technologies/all4can/can-news-blog/can-news-blog/2016/07/19/can-vs-can-fd-efficiency>.

- [67] F. Consortium, “FlexRay Communications System Protocol Specification Version 3.0.1,” 2010. <https://svn.ipd.kit.edu/nlrp/public/FlexRay/FlexRay%E2%84%A2%20Protocol%20Specification%20Version%203.0.1.pdf>.
- [68] I. O. for Standardisation, “Road vehicles — FlexRay communications system — Part 1: General information and use case definition,” 2013.
- [69] I. O. for Standardisation, “Road vehicles — FlexRay communications system — Part 5: Electrical physical layer conformance test specification,” 2013.
- [70] N. Instruments, “FlexRay Automotive Communication Bus Overview,” 2019. <https://www.ni.com/en-gb/innovations/white-papers/06/flexray-automotive-communication-bus-overview.html>.
- [71] R. Diamond, “FlexRay Protocol Overview,” 2018. <https://www.enigmatos.com/2018/03/06/flexray-protocol-overview-blog/>.
- [72] W. Zeng, M. Khalid, and S. Chowdhury, “A Qualitative Comparison of FlexRay and Ethernet in Vehicle Networks,” in *2015 IEEE 28th Canadian Conference on Electrical and Computer Engineering (CCECE)*, pp. 571–576, 2015.
- [73] “IEEE Standard for Information technology– Local and metropolitan area networks– Part 3: CSMA/CD Access Method and Physical Layer Specifications Amendment: Physical Layer and Management Parameters for 10 Gb/s Operation, Type 10GBASE-CX4,” *IEEE Std 802.3ak-2004 (Amendment to IEEE Std 802.3-2002 as amended by IEEE Stds 802.3ae-2002, 802.3af-2003 and 802.3aj-2003)*, pp. 1–56, 2004.

- [74] T. Instruments, *CAN bus, Ethernet, or FPD-Link: Which is best for automotive communications?* Texas Instruments, 2014. <https://www.ti.com/lit/an/slyt560/slyt560.pdf?ts=1597403343817>.
- [75] F. Silva, M. Quilez, J. Llop, X. Torres, and P. Riu, “EMI from an automotive CAN bus,” in *1999 IEEE International Symposium on Electromagnetic Compatibility. Symposium Record (Cat. No.99CH36261)*, vol. 1, pp. 512–516 vol.1, 1999.
- [76] C. I. S. des Perturbations Radio, “Limits and methods of measurement of radio disturbance characteristics for the protection of receivers used on board vehicles,” 1995.
- [77] C. I. S. des Perturbations Radio, “Industrial, scientific and medical equipment - Radio-frequency disturbance characteristics - Limits and methods of measurement,” 2016.
- [78] S. International, “Performance Levels and Methods of Measurement of Electromagnetic Compatibility of Vehicles, Boats (up to 15 m), and Machines (16.6 Hz to 18 GHz),” 2010.
- [79] D. Saranga, *Kia Soul EV (2018)*.
- [80] Kia, “Soul EV Specification,” 2017 (Last accessed 04 Jul 2020). <https://www.kia.com/content/dam/kwcms/kme/uk/en/assets/vehicles/soul-ev/specification/kia-soul-ev-specification.pdf>.
- [81] E. LINDGREN, *Model 7405 Near-Field Probe Set User Manual*. ETS LINDGREN, 2009 (Last accessed 04 Jul 2020). https://www.naic.edu/phil/rfi/antennas/NearFieldProbeSet7405_UserManual.pdf.
- [82] ITU, “Radio Regulations Articles, Edition of 2016,” 2016 (Last accessed 04 Jul 2020).

<http://search.itu.int/history/HistoryDigitalCollectionDocLibrary/1.43.48.en.101.pdf>.

- [83] Microchip, *MCP2515 Stand-Alone CAN Controller with SPI Interface*. Microchip, 2019 (Last accessed 07 Jul 2020). <http://ww1.microchip.com/downloads/en/DeviceDoc/MCP2515-Stand-Alone-CAN-Controller-with-SPI-20001801J.pdf>.
- [84] C. I. S. des Perturbations Radio, “Vehicles, boats and internal combustion engines - radio disturbance characteristics - limits and methods of measurement for the protection of off-board receivers,” 2007.
- [85] D. Stoll, “The advantage of time domain analysis for emc broadband measurements,” in *1972 IEEE International Electromagnetic Compatibility Symposium Record*, pp. 1–4, 1972.
- [86] M. A. C. Ltd, *AY4489 Time Domain EMC Emissions Measurement System*. Multiple Access Communications Ltd, 2004.
- [87] “IEC 61000-4-9:2016 ,” standard, IEC, July 2016.
- [88] Keysight Technologies, “Benchlink waveform builder pro.”
- [89] K. Technologies, *Keysight N9344C Handheld Spectrum Analyzer (HSA) 20 GHz Data Sheet*. Keysight Technologies, 2017.
- [90] E. LINDGREN, *Model 5400 Series Gigahertz Transverse Electromagnetic (GTEM!™) Cell Operation Manual*. ETS LINDGREN, 2018.
- [91] G. Dash, *How RF Anechoic Chambers Work*. Ampyx LLC, 1999.
- [92] F.-R. P. Corp., *Ferrite Tiles Absorbers for EMC Test Chamber Applications*. Fair-Rite Products Corp., NA.

- [93] M. K. Kazimierczuk, *Pulse-width Modulated DC-DC Power Converters*. The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ: Wiley, 1 ed., 2016.
- [94] N. Kularatna, *Power Electronics Design Handbook*. USA: Newnes, 1 ed., 1998.
- [95] C. Christopoulos, *Principles and Techniques of Electromagnetic Compatibility*. 6000 Broken Sound Parkway NW, Suite 300, Boca Raton, FL: CRC Press, 1 ed., 2007.
- [96] P. Chatterton and M. Houlden, *EMC Electromagnetic Theory to Practical Design*. Baffins Lane, Chichester, West Sussex, PO19 1UD: Wiley, 1 ed., 1992.
- [97] J. Stewart, *Calculus*. 10 Davis Drive, Belmont, CA 94002: Thomson, 6 ed., 2008.
- [98] C. L. Holloway, R. R. DeLyser, R. F. German, P. McKenna, and M. Kanda, "Comparison of Electromagnetic Absorber used in Anechoic and Semi-Anechoic Chambers for Emissions and Immunity Testing of Digital Devices," *IEEE Transactions on Electromagnetic Compatibility*, vol. 39, no. 1, pp. 33–47, 1997.
- [99] A. Corporation, *MP8302A Bit Error Rate Tester Operation Manual*. Anritsu Corporation, 2015.
- [100] C. Smartt *et al.*, *User Guide SACAMOS: State of the Art Cable Models for Spice Open Source Cable Models for EMI Simulations*. University of Nottingham, 2018. http://128.243.70.77/UoN/SACAMOS/raw/master/DOCUMENTATION/SACAMOS_UserGuide.pdf.

- [101] C. Smartt *et al.*, *Theory Manual SACAMOS: State of the Art Cable Models for Spice Open Souce Cable Models for EMI Simulations*. University of Nottingham, 2018. http://128.243.70.77/UoN/SACAMOS/raw/master/DOCUMENTATION/SACAMOS_TheoryManual.pdf.
- [102] C. Smartt, M. J. Basford, S. Greedy, D. W. P. Thomas, and M. Sumner, “EMC-Oriented Multi-Conductor Equivalent Circuit Cable Models for SPICE, Including Transfer Impedance Coupling and Incident Field Excitation,” in *2017 International Conference on Electromagnetics in Advanced Applications (ICEAA)*, pp. 1283–1287, 2017.
- [103] S. Greedy, C. Smartt, M. J. Basford, and D. W. P. Thomas, “Open source cable models for emi simulations,” *IEEE Electromagnetic Compatibility Magazine*, vol. 7, no. 3, pp. 69–81, 2018.
- [104] J. Langston. Private Communication, 2018.
- [105] K. Technologies, “N4962A Serial BERT 12.5 Gb/s User Guide,” 2014. <http://literature.cdn.keysight.com/litweb/pdf/N4962-91021.pdf>.
- [106] M. J. Basford, C. Smartt, D. W. P. Thomas, and S. Greedy, “On the Disruption of Wired Serial Communication Links by Time Domain Interference,” in *2018 IEEE International Symposium on Electromagnetic Compatibility and 2018 IEEE Asia-Pacific Symposium on Electromagnetic Compatibility (EMC/APEMC)*, pp. 183–186, 2018.
- [107] R. Grigonis, *Computer Telephony Encyclopedia*. CMP Books, 2000.
- [108] L. Brugarolas, “Simple BER Meter is Easy to Build,” Mar 2000. <https://www.edn.com/simple-ber-meter-is-easy-to-build/>.
- [109] G. C. Inc., *T1/E1 ATM BERT*. GL Communications Inc., NA.

- [110] Maxim, *EBERT DS2174*, 9 2002. Rev. 1.
- [111] P. Wilson, *Design Recipes for FPGAs: Using Verilog and VHDL*. Newnes, 2 ed., 2016.
- [112] M. Gschwind, V. Salapura, and D. Maurer, “FPGA Prototyping of a RISC Processor Core for Embedded Applications,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 9, no. 2, pp. 241–250, 2001.
- [113] “Quartus Prime Lite Edition Download,” Sep 2019. <https://fpgasoftware.intel.com/?edition=lite>.
- [114] J. Koeter, “What’s an LFSR?,” 1996. <http://www.ti.com/lit/an/scta036a/scta036a.pdf>.
- [115] W. W. PETERSON and E. J. WELDON, *Error-Correcting Codes*. M.I.T. Press, 2 ed., 1972.
- [116] Xilinx, “Efficient Shift Registers, LFSR Counters, and Long Pseudo-Random Sequence Generators,” 1996.
- [117] ITB CompuPhase, “Termite.”
- [118] P. Semiconductor, *PO49FCT32805 600MHz TTL/CMOS Potato Chip*. Potato Semiconductor, 2005. <https://pdf1.alldatasheet.com/datasheet-pdf/view/292094/POTATO/PO49FCT32805.html>.
- [119] A. Devices, *ADCMP600/ADCMP601/ADCMP602*. Analog Devices, 2011. https://www.analog.com/media/en/technical-documentation/data-sheets/ADCMP600_01_02.pdf.

- [120] Altera, “Understanding Metastability in FPGAs,” 2009. <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/wp/wp-01082-quartus-ii-metastability.pdf>.
- [121] O. Strobel, *Communication in Transportation Systems*. Information Science Reference (an imprint of IGI Global), 2013.
- [122] H. Li, H. Zhang, and D. Peng, “Research and Application of FlexRay High-Speed Bus on Transformer Substation Automation System,” *Web Information Systems and Mining Lecture Notes in Computer Science*, vol. 6318, p. 263–270, 2010.
- [123] R. Chen, C. Li, S. Yan, R. Malaney, and J. Yuan, “Physical Layer Security for Ultra-Reliable and Low-Latency Communications,” *IEEE Wireless Communications*, vol. 26, no. 5, pp. 6–11, 2019.
- [124] L. Connectivity, *LoRa/Bluetooth Low Energy (BLE) Module Part Numbers: RM186 and RM191*. Laird Connectivity, 2019. https://connectivity-staging.s3.us-east-2.amazonaws.com/2019-02/CS-DS-RM1xx%20v1_9.pdf.
- [125] Freewave, *ZumLink Covers Model:Z9-C or Z9-T*. Freewave, 2018. <https://www.freewave.com/wp-content/uploads/2018/12/Z9-C-Z9-T-User-Manual.pdf>.
- [126] E. Research, *USRP X300 and X310 X Series*. Ettus Research, 2010. https://www.ettus.com/wp-content/uploads/2019/01/X300_X310_Spec_Sheet.pdf.
- [127] “QPSK Transmitter with USRP® Hardware.” <https://uk.mathworks.com/help/supportpkg/usrpradio/examples/qpsk-transmitter-with-usrp-r-hardware-1.html>.

- [128] “QPSK Receiver with USRP® Hardware.”
<https://uk.mathworks.com/help/supportpkg/usrpradio/examples/qpsk-receiver-with-usrp-r-hardware-1.html>.
- [129] M. W. O’Brien, J. S. Harris, O. Popescu, and D. C. Popescu, “An Experimental Study of the Transmit Power for a USRP Software-Defined Radio,” in *2018 International Conference on Communications (COMM)*, pp. 377–380, 2018.
- [130] N. Storey, *Electronics: a Systems Approach*. Pearson, 5th ed., 2013.
- [131] M. T. Ma, M. Kanda, M. L. Crawford, and E. B. Larsen, “A Review of Electromagnetic Compatibility/Interference Measurement Methodologies,” *Proceedings of the IEEE*, vol. 73, no. 3, pp. 388–411, 1985.
- [132] B.-H. Liu, D. C. Chang, and M. T. Ma, *Eigenmodes and the Composite Quality Factor of a Reverberating Chamber*. National Bureau of Standards, 1983.
- [133] C. Bruns, *Three-Dimensional Simulation and Experimental Verification of a Reverberation Chamber*. PhD thesis, 2005.
- [134] A. M. Villalon, *Rail Internet of Things (RIoT) for Trackside Signalling*. PhD thesis, 2018.
- [135] Q. Xu and Y. Huang, *Anechoic and Reverberation Chambers: Theory, Design and Measurements*. WILEY-BLACKWELL, 2019.
- [136] A. K. Fall, P. Besnier, C. Lemoine, M. Zhadobov, and R. Sauleau, “Determining the Lowest Usable Frequency of a Frequency-Stirred Reverberation Chamber Using Modal Density,” in *2014 International Symposium on Electromagnetic Compatibility*, pp. 263–268, 2014.

- [137] L. Arnaut, “Mode-Stirred Reverberation Chambers: A Paradigm for Spatio-Temporal Complexity in Dynamic Electromagnetic Environments,” *Wave Motion*, vol. 51, pp. 673–684, June 2014.
- [138] C. L. Holloway, D. A. Hill, J. M. Ladbury, P. F. Wilson, G. Koepke, and J. Coder, “On the Use of Reverberation Chambers to Simulate a Rician Radio Environment for the Testing of Wireless Devices,” *IEEE Transactions on Antennas and Propagation*, vol. 54, no. 11, pp. 3167–3177, 2006.
- [139] D. A. Hill, *Electromagnetic theory of reverberation chambers*. National Institute of Standards and Technology, 1998.
- [140] D. Saranga, *Fiat Panda (2009)*.
- [141] “Pycom lora (868mhz/915mhz) & sigfox antenna kit,” Jun 2020.

Appendix A

Electromagnetic Measurements

Around a Kia Soul EV

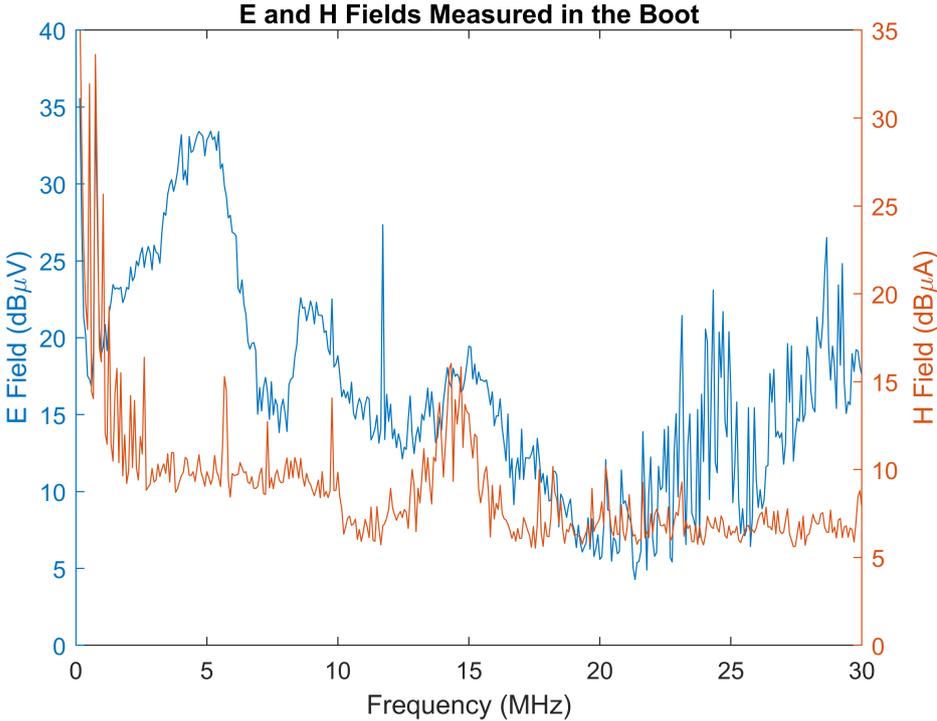


Figure A.1: Electric and magnetic field emissions in the boot.

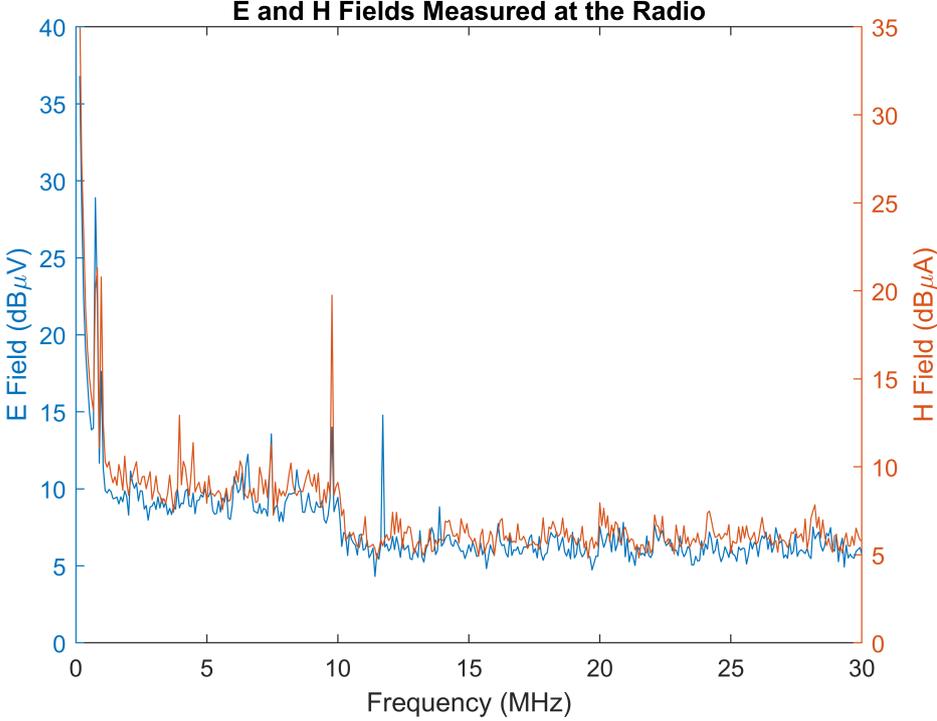


Figure A.2: Electric and magnetic field emissions at the radio antenna.

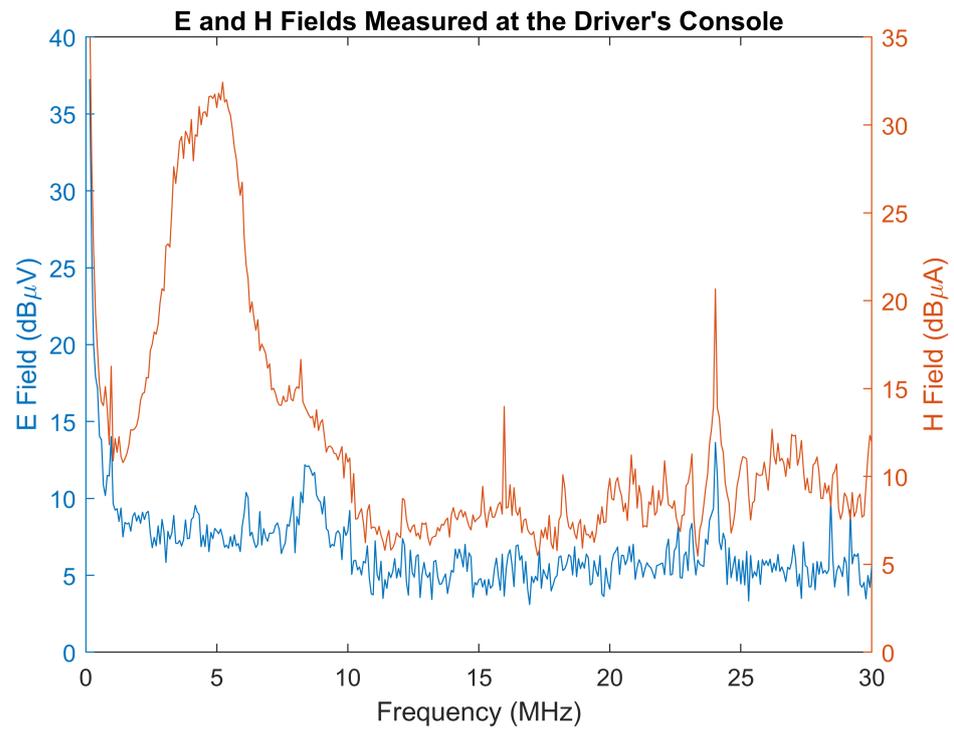


Figure A.3: Electric and magnetic field emissions at the driver's console.

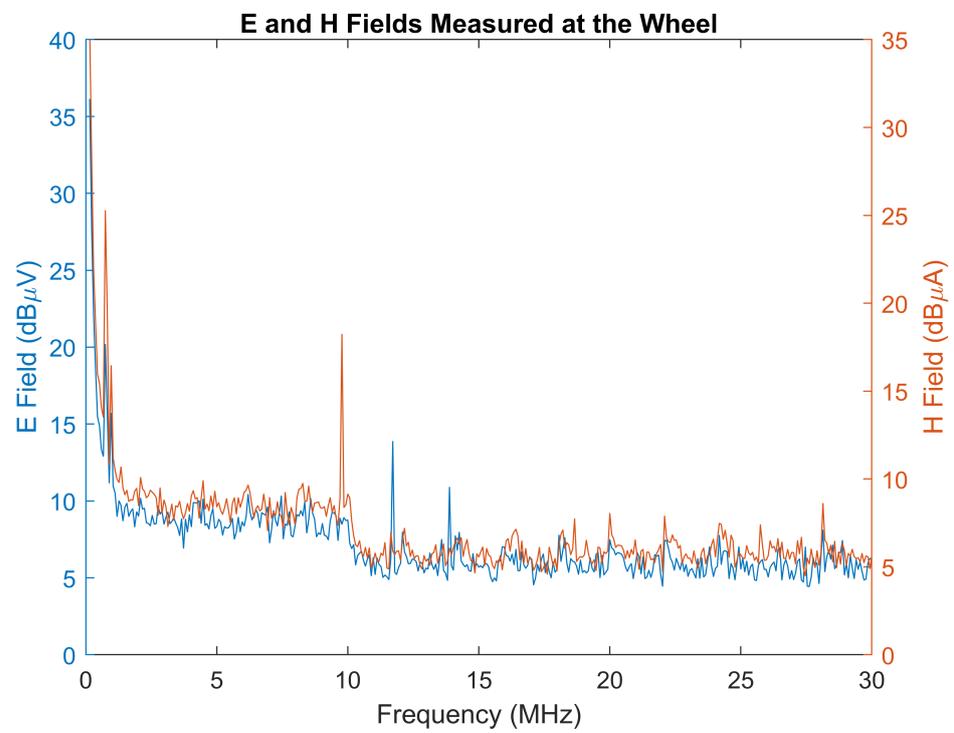


Figure A.4: Electric and magnetic field emissions in the wheel well.

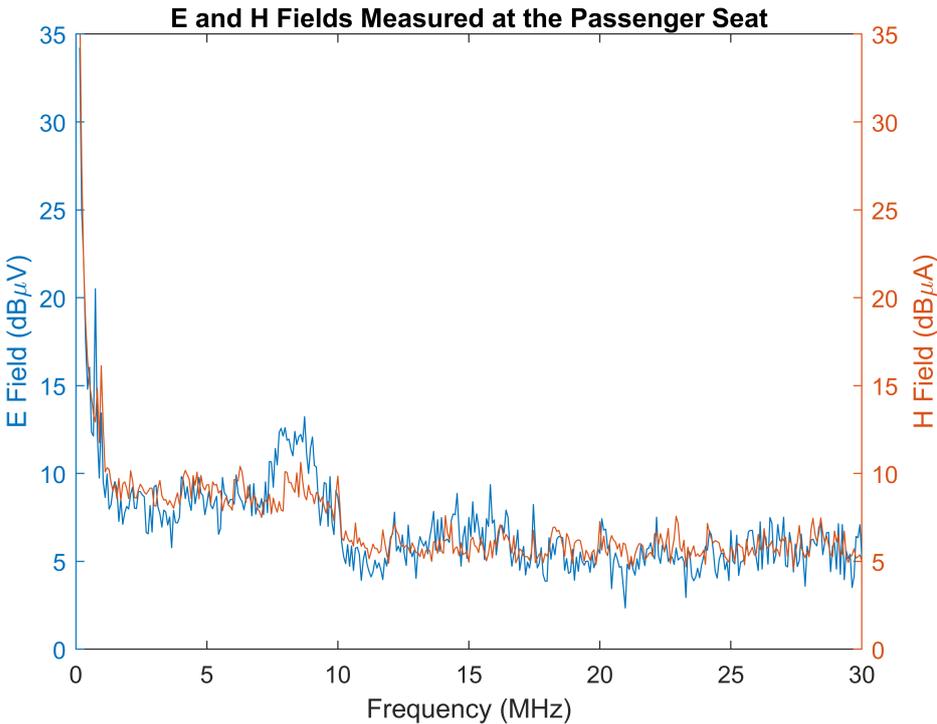


Figure A.5: Electric and magnetic field emissions at the rear passenger’s location.

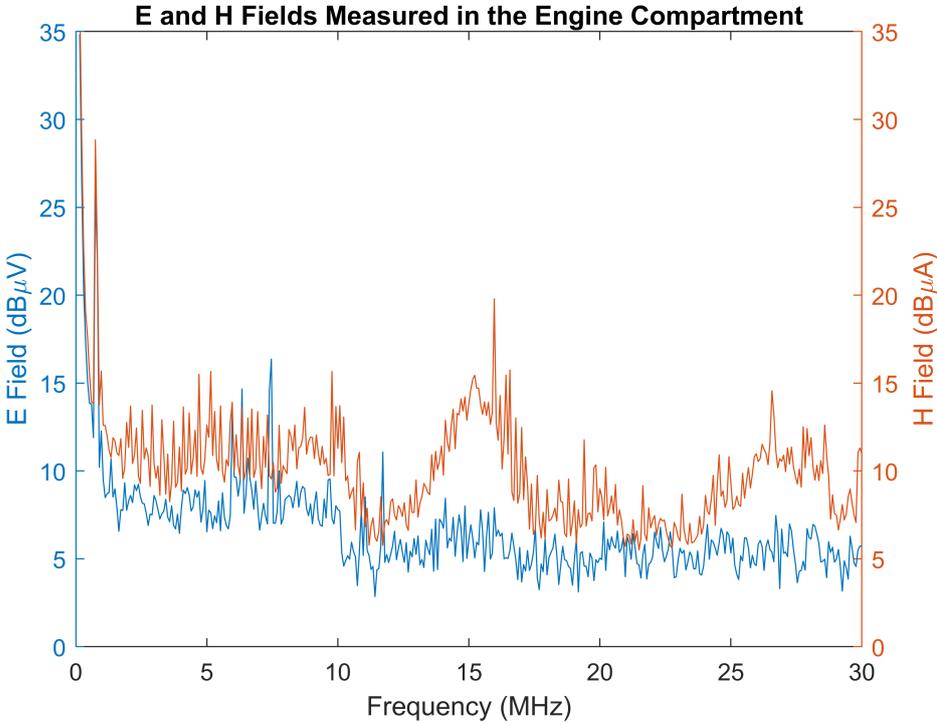


Figure A.6: Electric and magnetic field emissions in engine compartment.

Appendix B

Additional Study of BER Variations

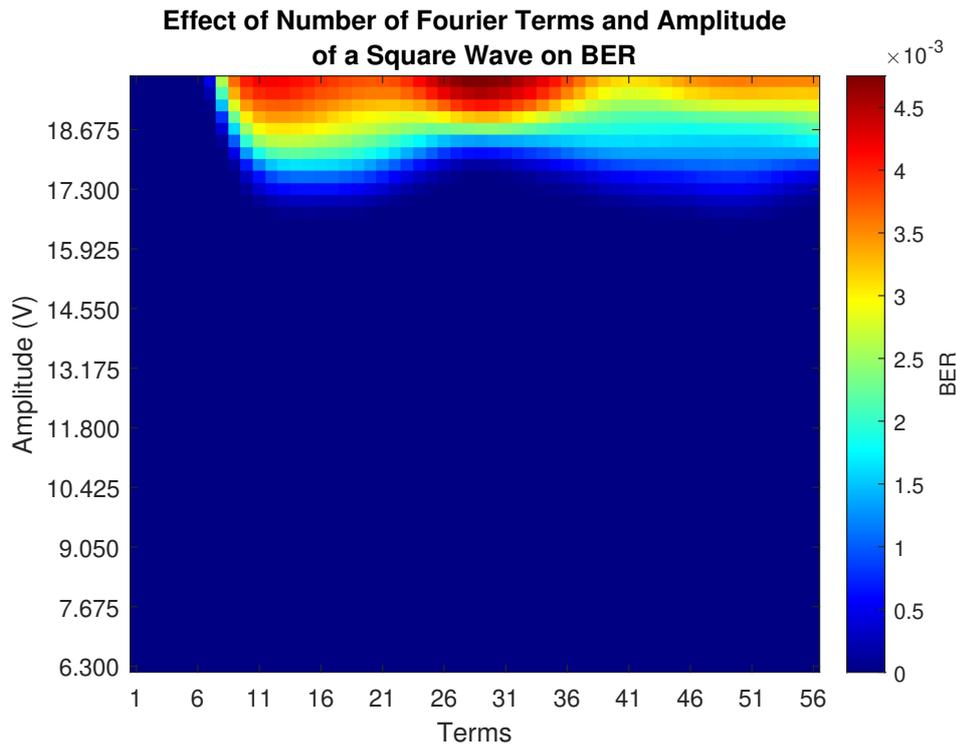


Figure B.1: Colour map of BER for systems with $50\ \Omega$ impedance.

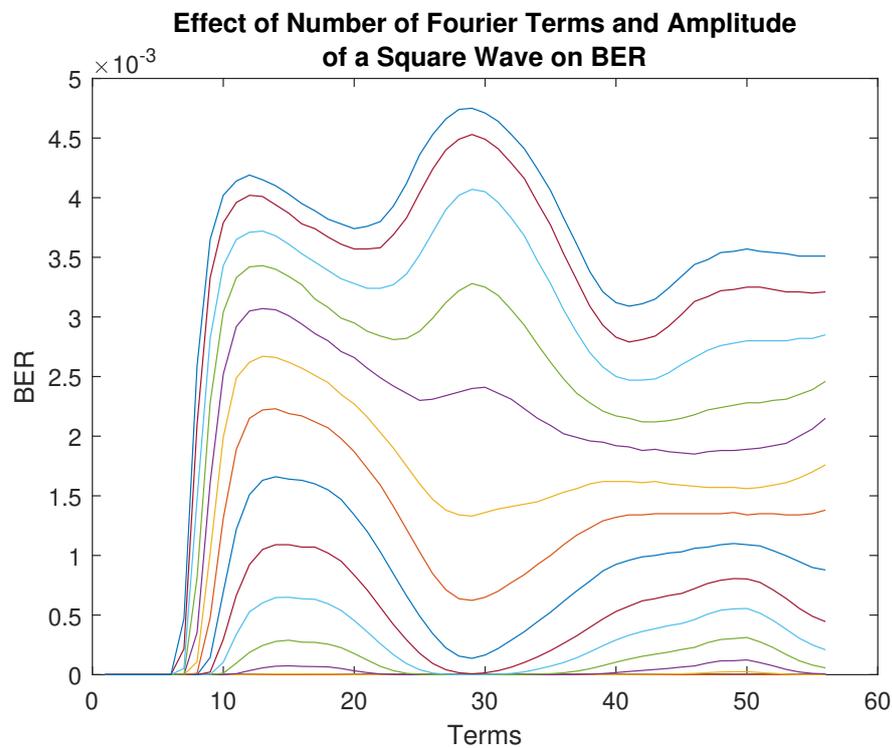


Figure B.2: Linear plot of BER against number of Fourier terms. Note, lines are naturally sorted by interference amplitude.

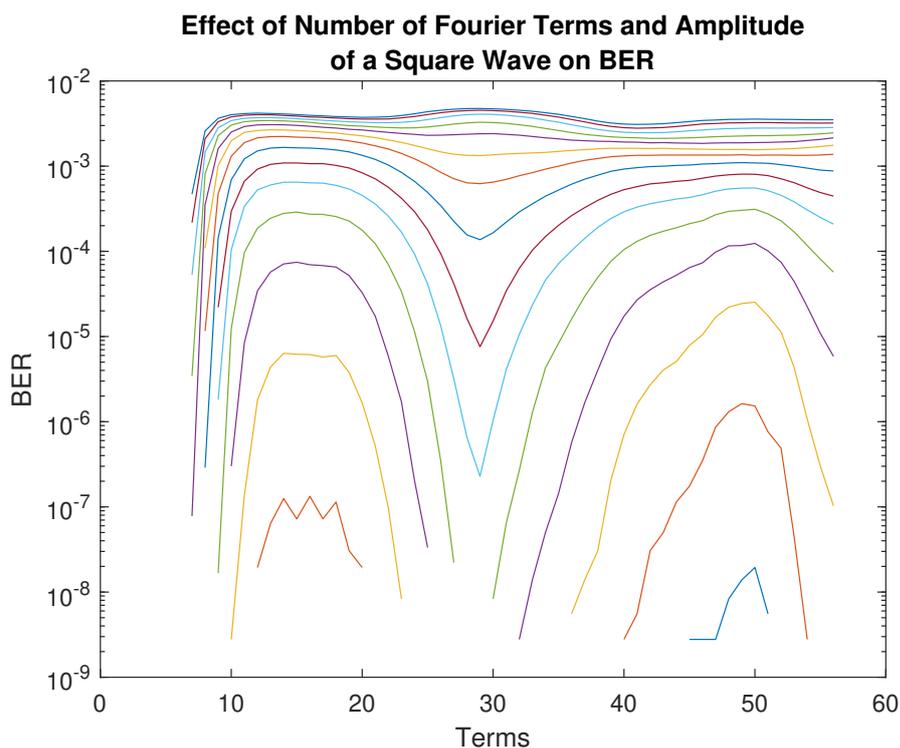


Figure B.3: Logarithm plot of BER against number of Fourier terms, showing detail of BER for lower interference voltages.

The experimental configuration in Section 5.2 is based on a data transmission network with a $1\text{ M}\Omega$ termination impedance. In this appendix, results are shown for a system of $50\ \Omega$ termination impedance.

The colour map in Figure B.1 immediately indicates the $50\ \Omega$ network's great resilience to EMI for the same waveform shape and amplitude - in this case BER values are non-zero at around 16V , when the $1\text{ M}\Omega$ has initial non-zero BER values at 6.4V . Furthermore, the magnitude of the BER is more than one magnitude lower at the maximum value, and extends through a much greater range of values.

Figure B.2 shows the highly variable BER response as a function of the number of Fourier terms. The variability in this case is much higher than that experienced by the $1\text{ M}\Omega$ system, with highly varying features across interference amplitude and shape. The greater range of BER values is best demonstrated in the logarithmic representation of this data in Figure B.3.

Discontinuities between similar colours in this plot indicate a zero BER value. It is evident that for lower interference amplitudes that the magnitude of the trough BER is much lower than for higher amplitudes at this $50\ \Omega$ - a trait similar to that experienced by the $1\ \text{M}\Omega$ system.

By designing an emission source waveform to have a rate of change similar to that achieved at around 25-30 terms, it would be possible to accommodate a nearby victim network, with minimal interference. For systems that may have been operating in the sub-25 term region, this would result in an improvement in both source *and* victim efficiency. For systems operating at equivalent rate of change that were greater than 30 terms, the information in Figure B.3 permits an optimal design solution to maximise efficiency and integrity of both source and victim.

Appendix C

Simulated BER Plots Using Measured Transfer Function

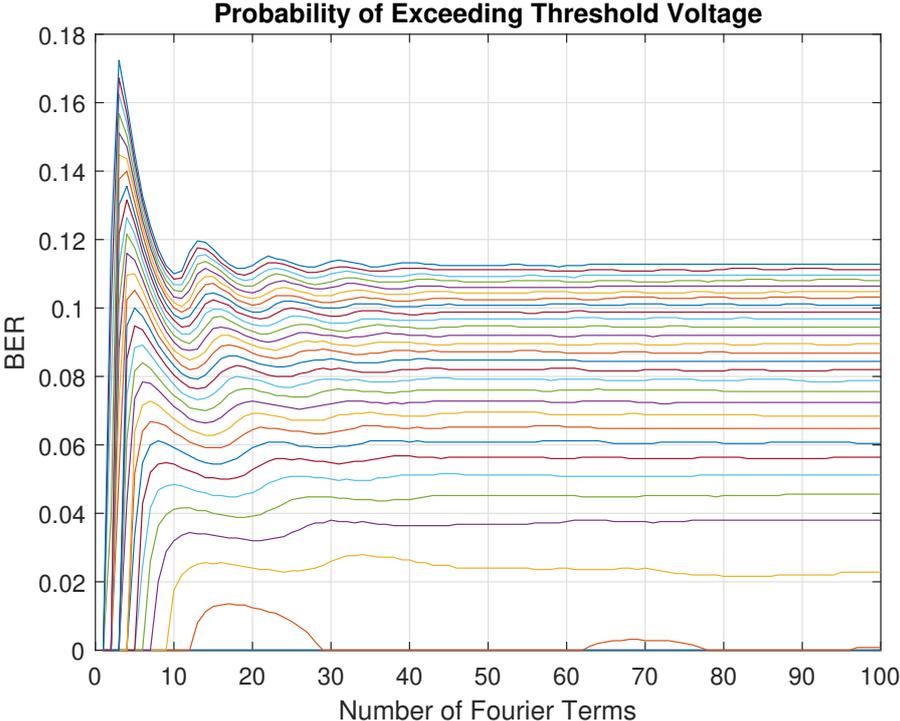


Figure C.1: Probability of interference with a 1 ns sample period. Uses a measured transfer function.

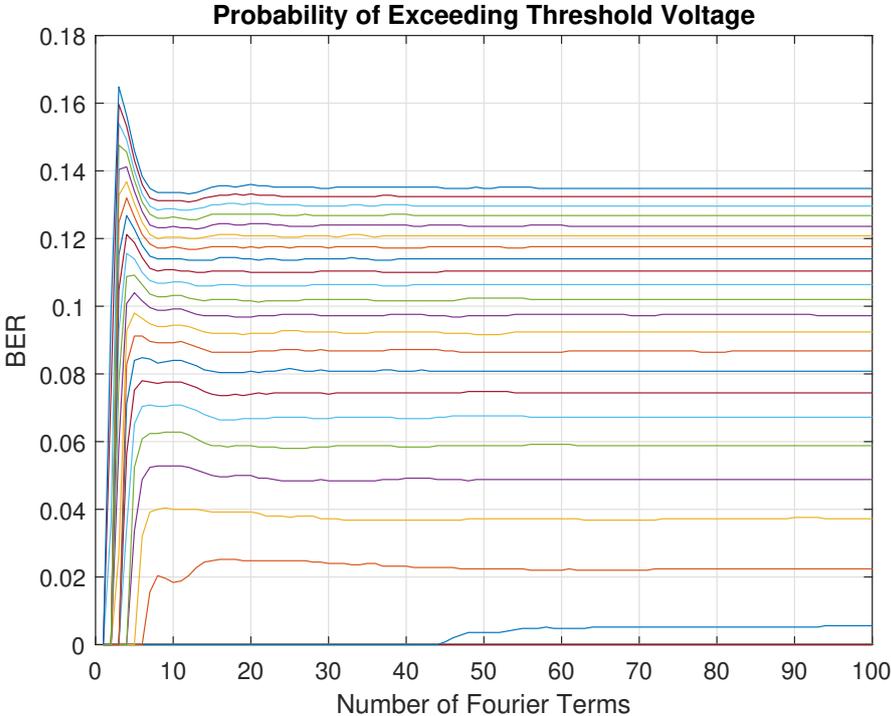


Figure C.2: Probability of interference with a 6 ns sample period. Uses a measured transfer function.

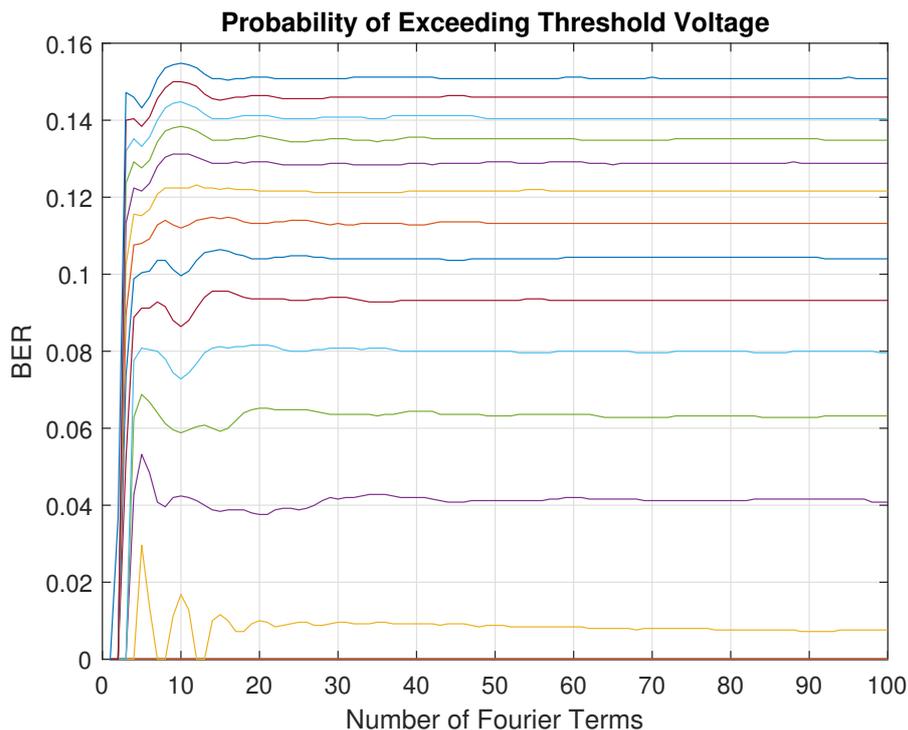


Figure C.3: Probability of interference with a 10 ns sample period. Uses a measured transfer function.

The simulation method described in Chapter 5 can also be used to examine the effect of the receiver design on the BER of a system, when exposed to identical forms and amplitudes of EMI. In this appendix, source amplitudes ranging from 10 - 35 V are simulated to generate BER curves.

Figure C.1 shows that for short sampling windows, the BER waveforms tend to have a greater difference between initial spike and converged BER values along with a much greater oscillatory response along each given amplitude. This is to be expected when considering the induced interference waveforms, examined in Chapter 5, as a smaller sampling window allows for a greater influence of the induced voltage spikes on the decaying edge of the interference waveform. Indeed, for the 1 ns case in Figure C.1, it can be seen that for low interference voltage amplitudes that the BER variation is significant across a much wider range of Fourier terms.

As the sampling window period is increased to 6 ns, the response can

be seen to homogenise somewhat, as per Figure C.2. The longer sampling window has reduced the influence the trailing peaks exercise on the bit arbitration, though at the expense of an increased converged BER.

This trend continues as the sampling window is increased to 10 ns, where the initial BER peak is completely reduced for higher voltage interference sources as shown in Figure C.3. For lower voltage sources, fine detail can be observed with very low BER regions between peaks.

By varying the sampling window duration of a receiver, it is possible to control the BER curves in the same manner as adjusting the interference source itself. By carefully designing receiver parameters, the BER of a system may be reduced significantly when the system is exposed to a given form of interference. Such an engineering methodology may be combined with altering the interference source waveform to facilitate even greater reduction of the BER.

Appendix D

OS BERT Schematics

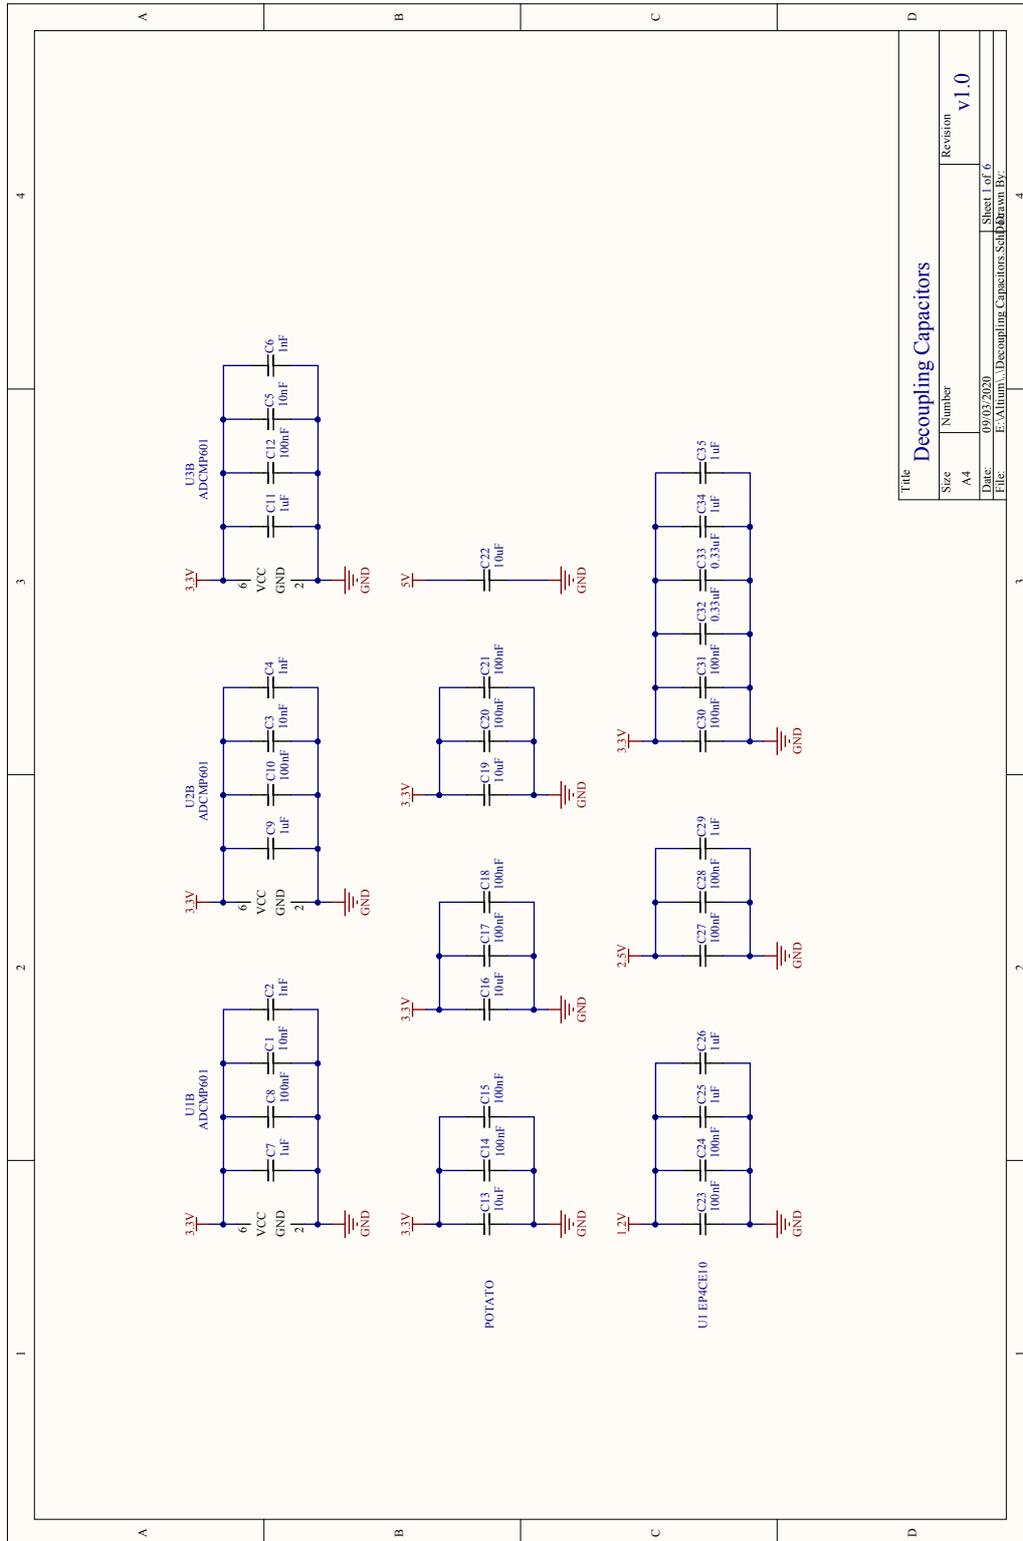


Figure D.1: OSBERT schematic page 1.

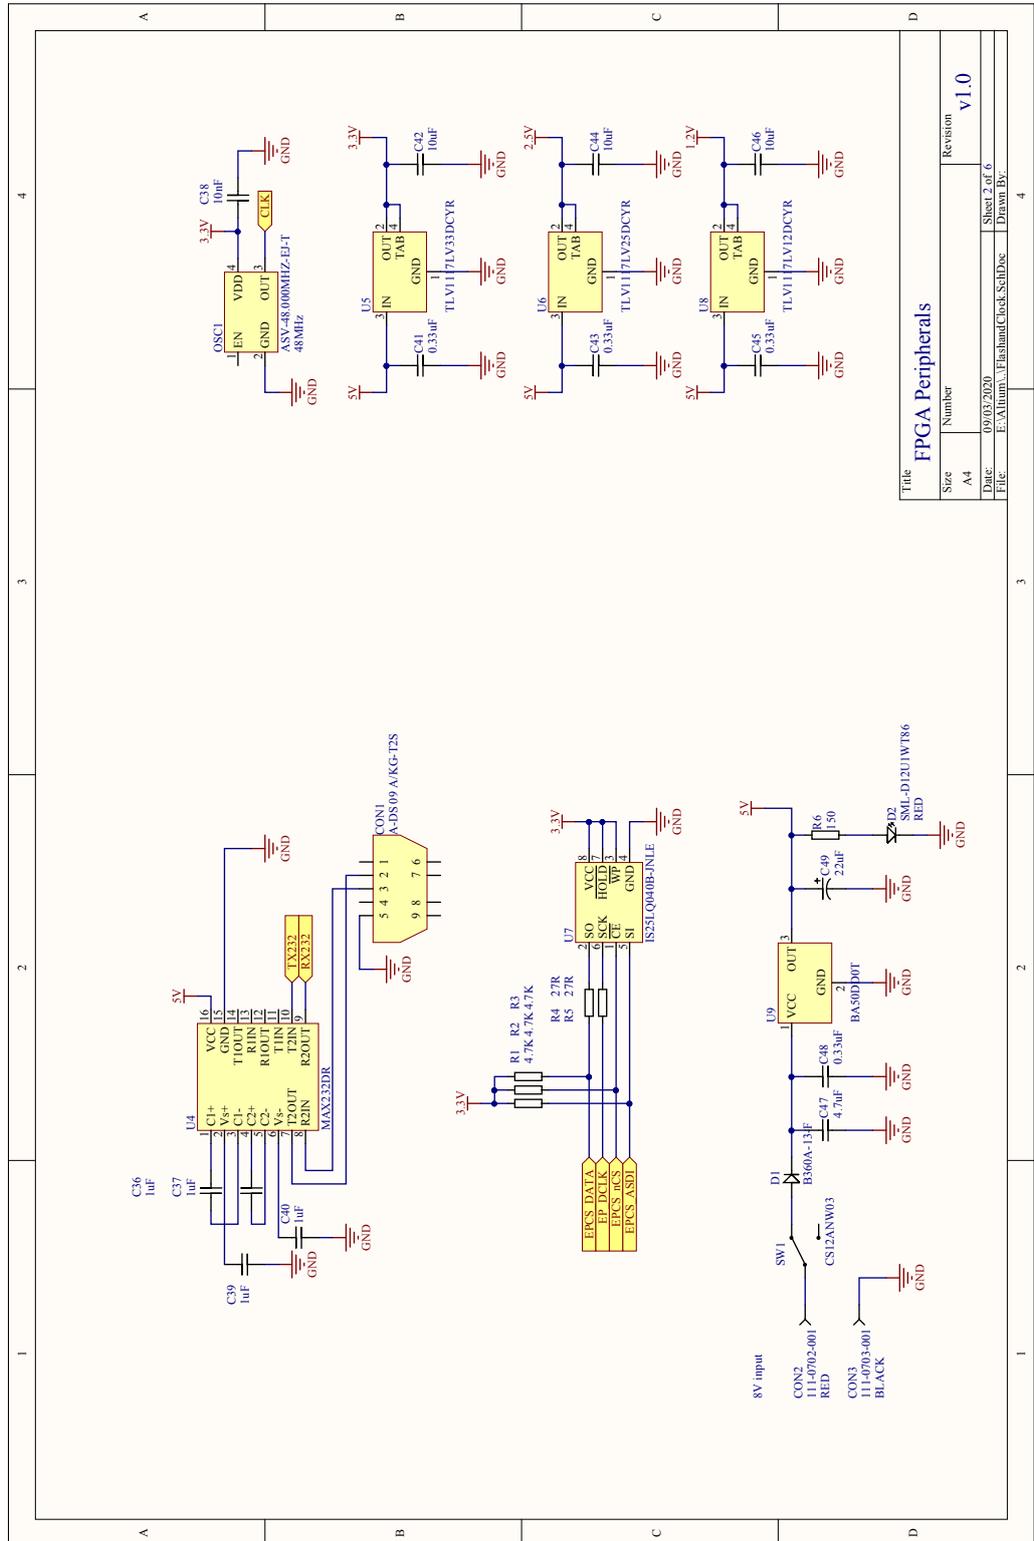
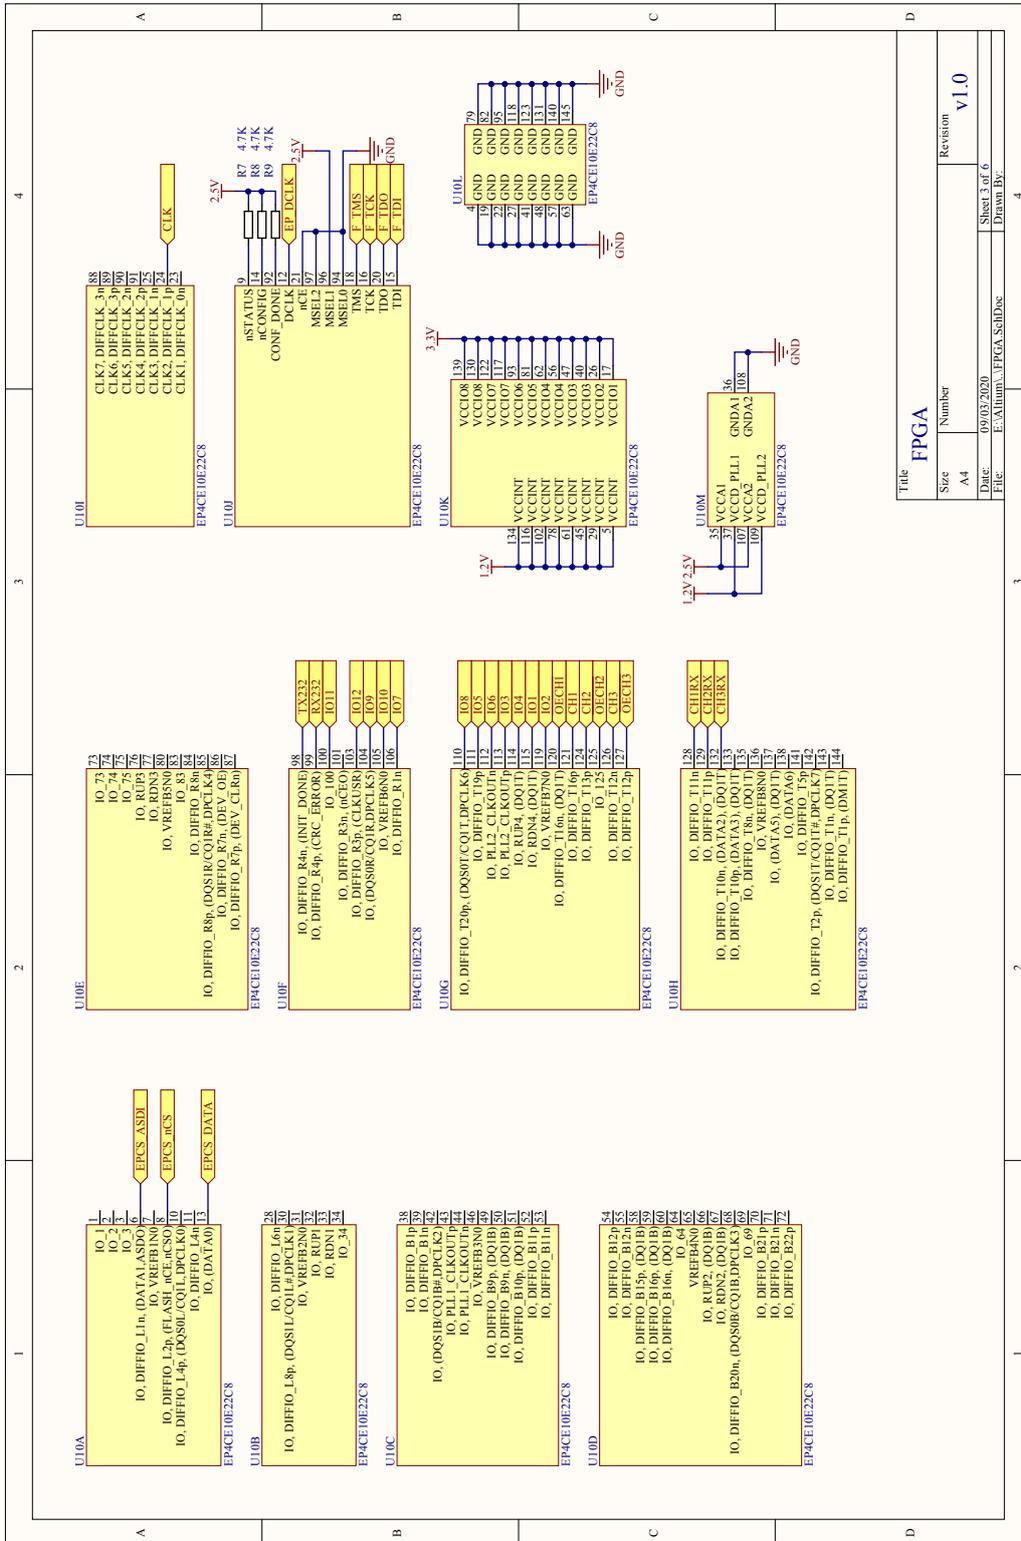


Figure D.2: OSBERT schematic page 2.



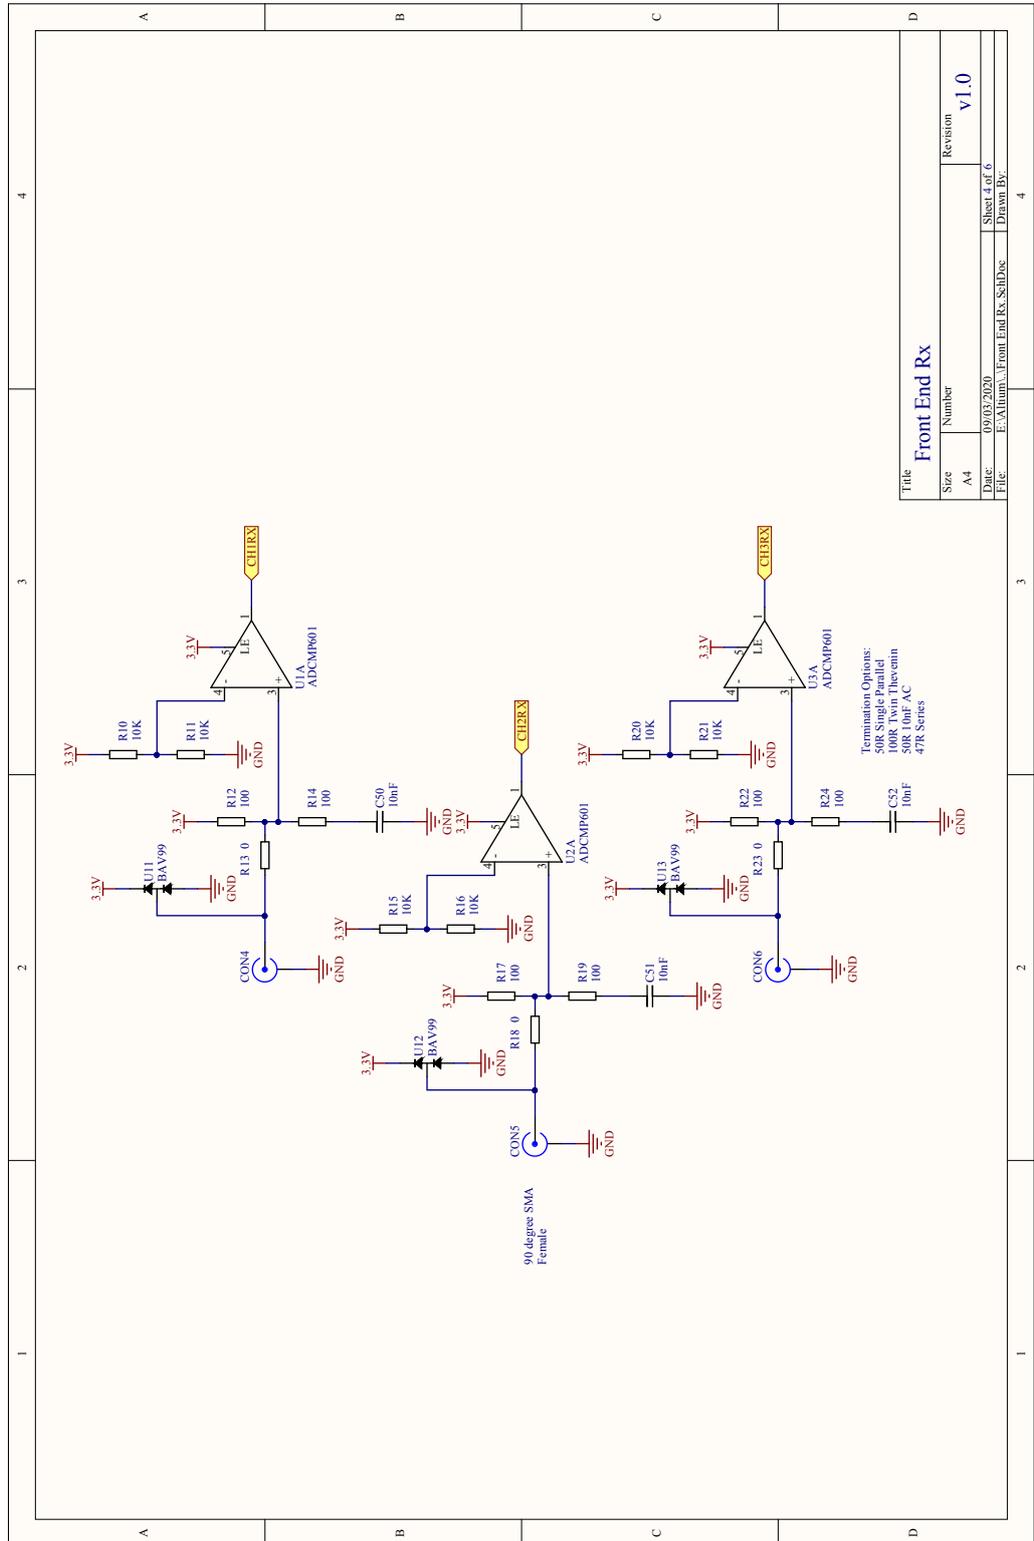


Figure D.4: OSBERT schematic page 4.

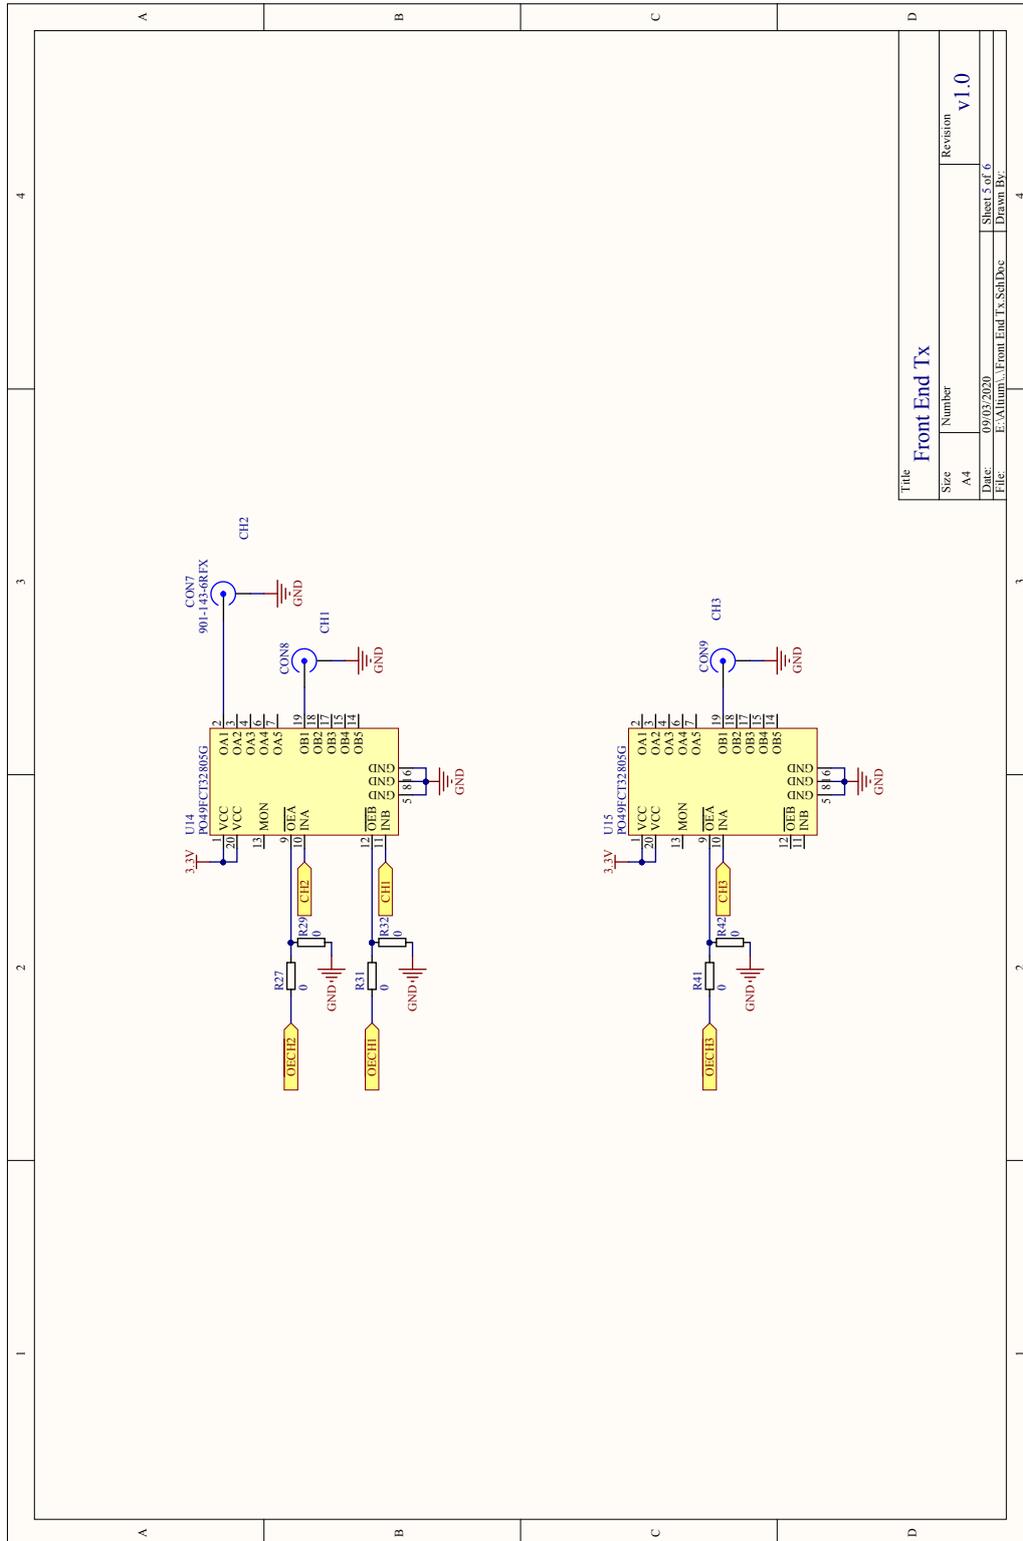


Figure D.5: OSBERT schematic page 5.

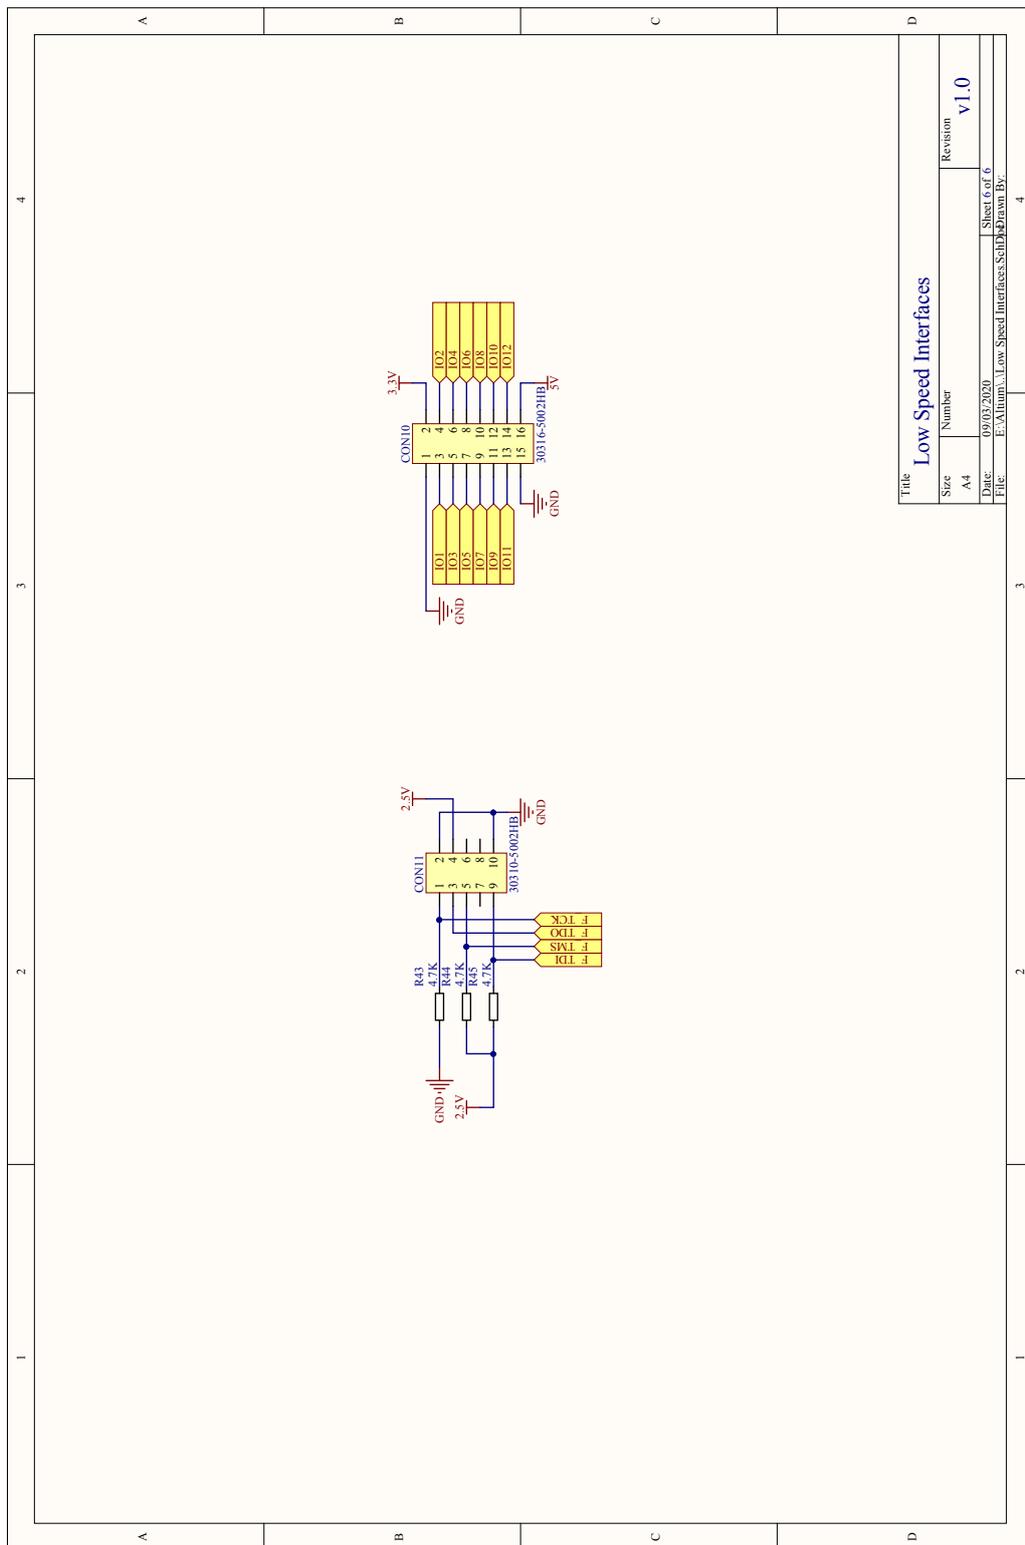


Figure D.6: OSBERT schematic page 6.

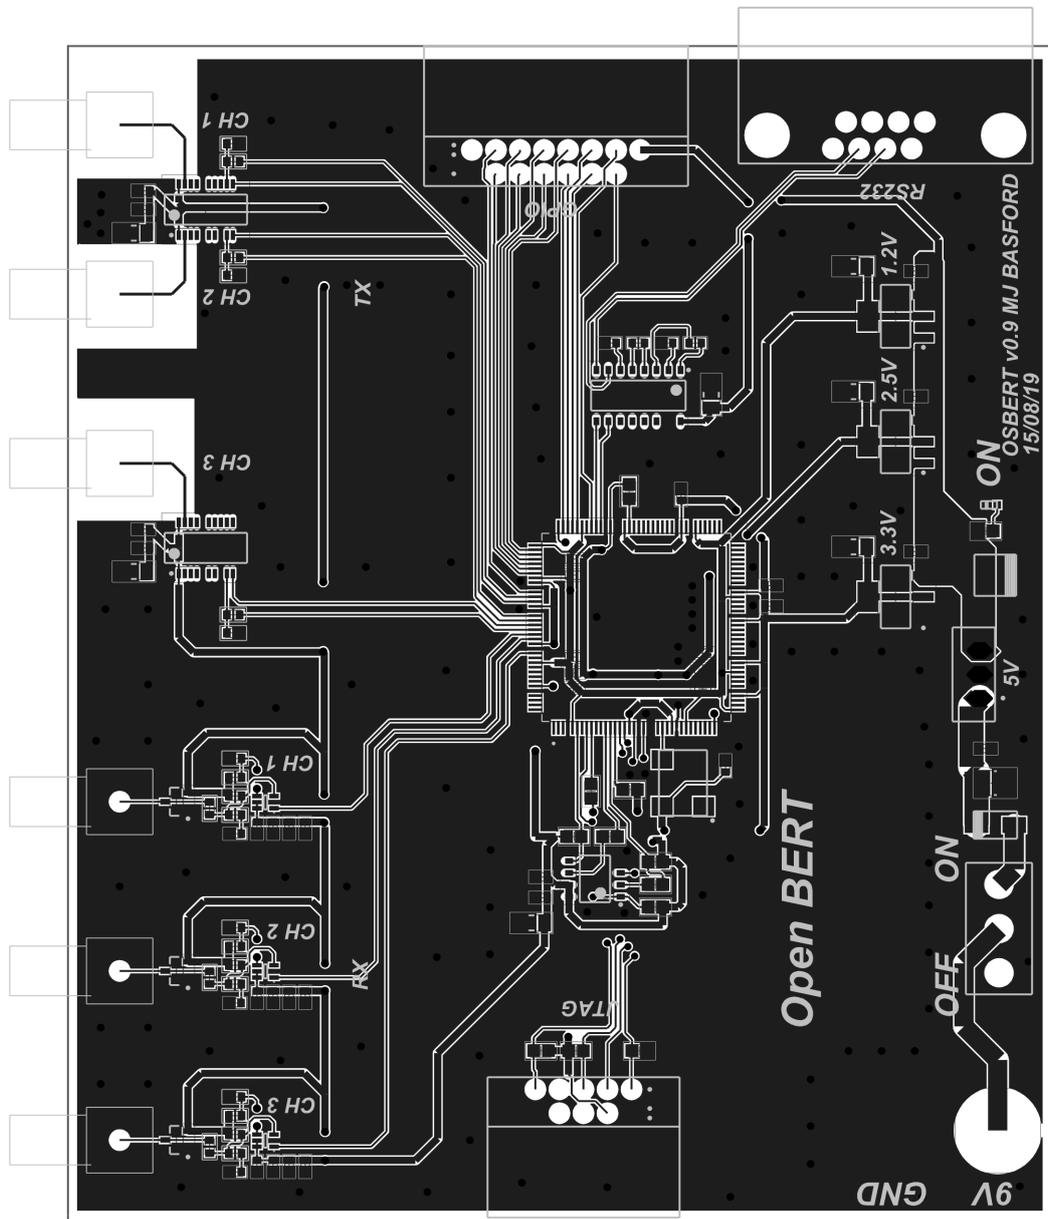


Figure D.7: PCB of OSBERT, top

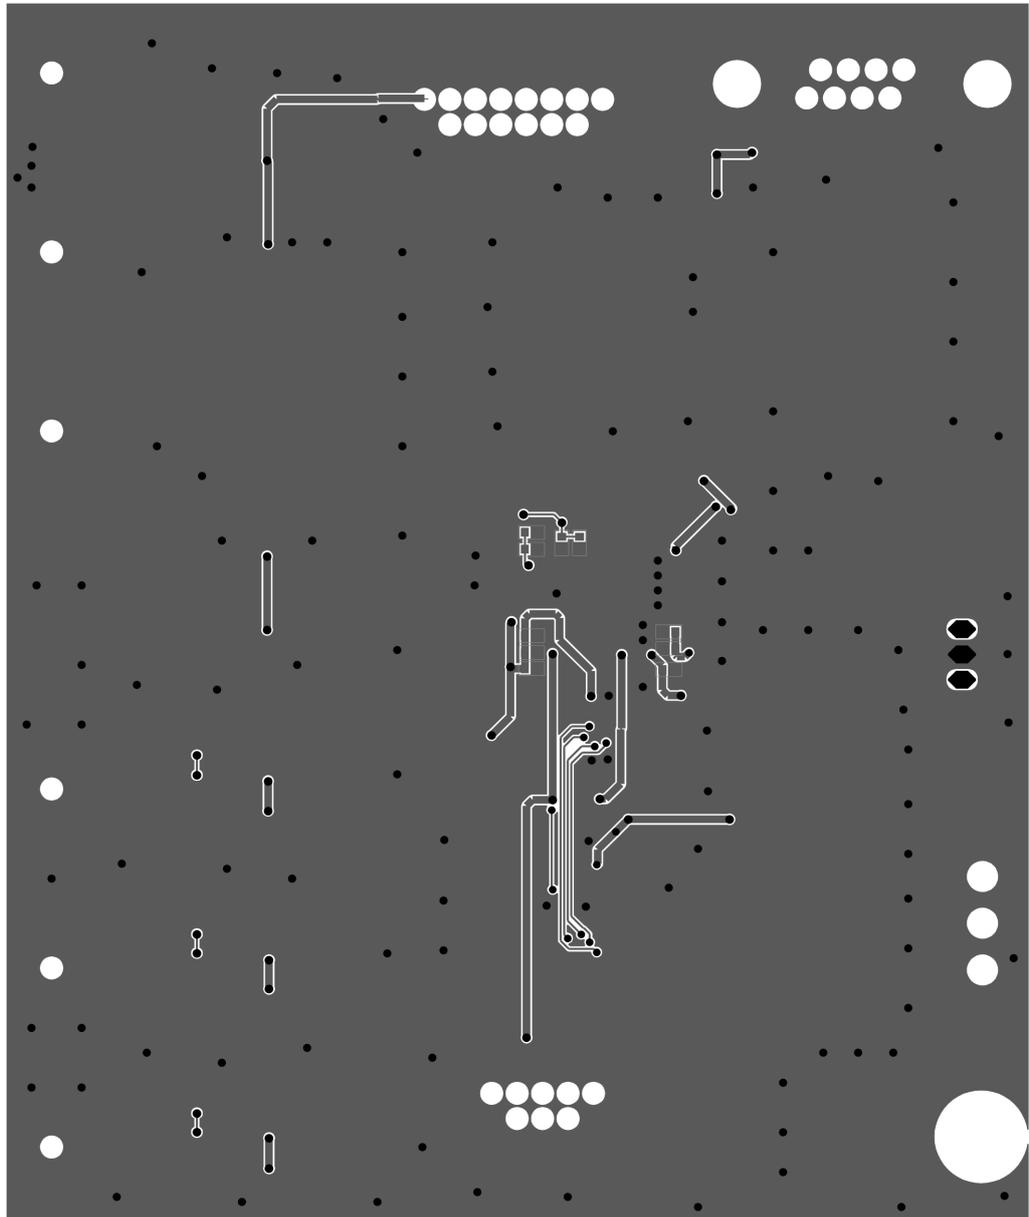


Figure D.8: PCB of OSBERT, bottom.

Appendix E

USRP QPSK Parameters

Table E.1: USRP QPSK transmitter structure

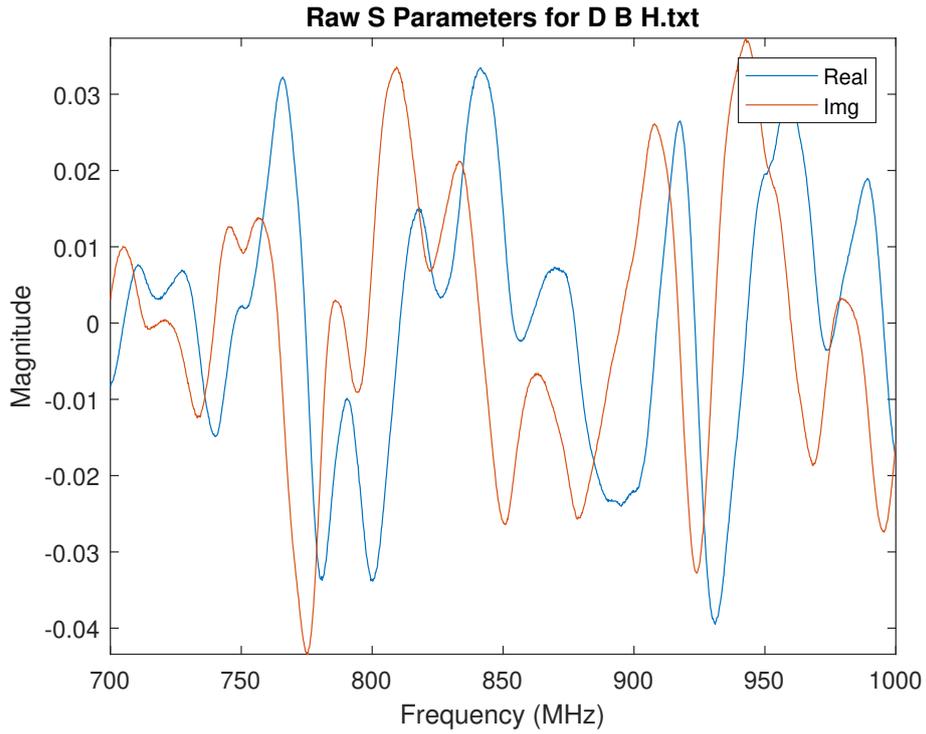
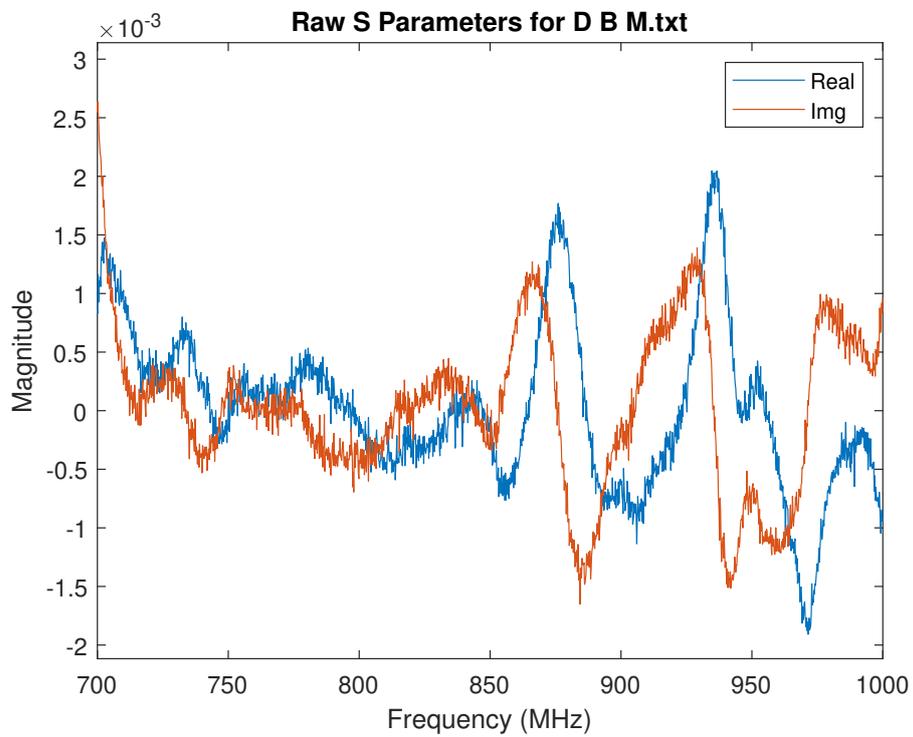
Field	Value
Rysm	200000
ModulationOrder	4
Interpolation	2
Decimation	1
Tsym	5.000e-06
Fs	400000
BarkerCode	[1 1 1 1 1 -1 -1 1 1 -1 1 -1 1]
BarkerLength	13
HeaderLength	26
Message	'Hello world'
MessageLength	16
NumberOfMessage	100
PayloadLength	11200
FrameSize	5613
FrameTime	0.0281
RolloffFactor	0.5000
ScramblerBase	2
ScramblerPolynomial	[1 1 1 0 1]
ScramblerInitialConditions	[0 0 0 0]
RaisedCosineFilterSpan	10
MessageBits	[11200*1 double]
MasterClockRate	100000000
USRPCenterFrequency	915000000
USRPGain	25
USRPFrontEndSampleRate	400000
USRPInterpolationFactor	250
USRPFrameLength	11226
USRPFrameTime	0.0281
StopTime	1000

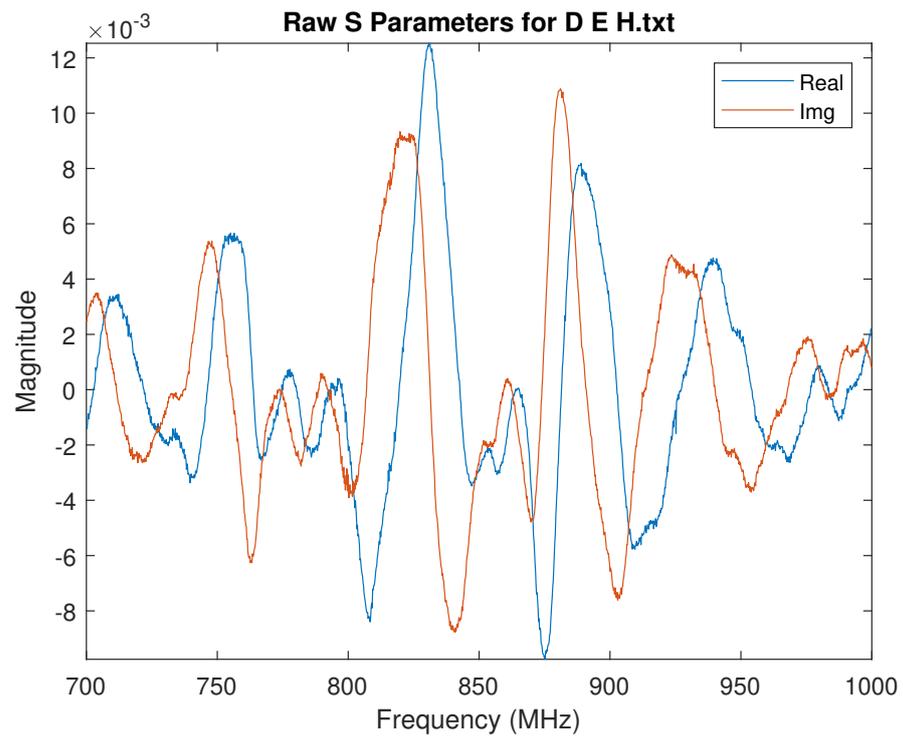
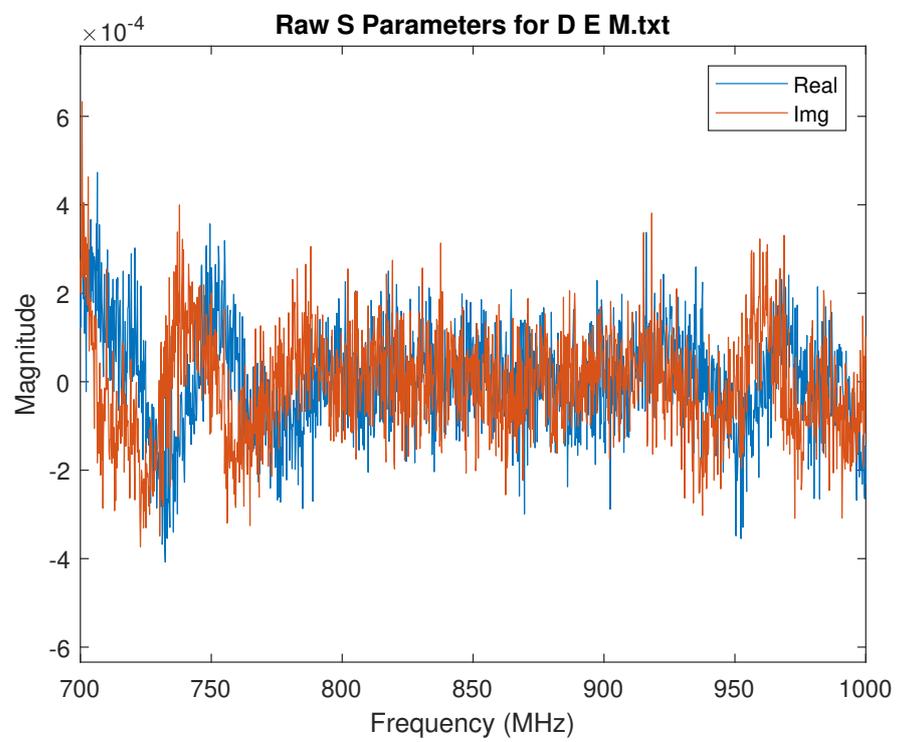
Table E.2: USRP QPSK receiver structure

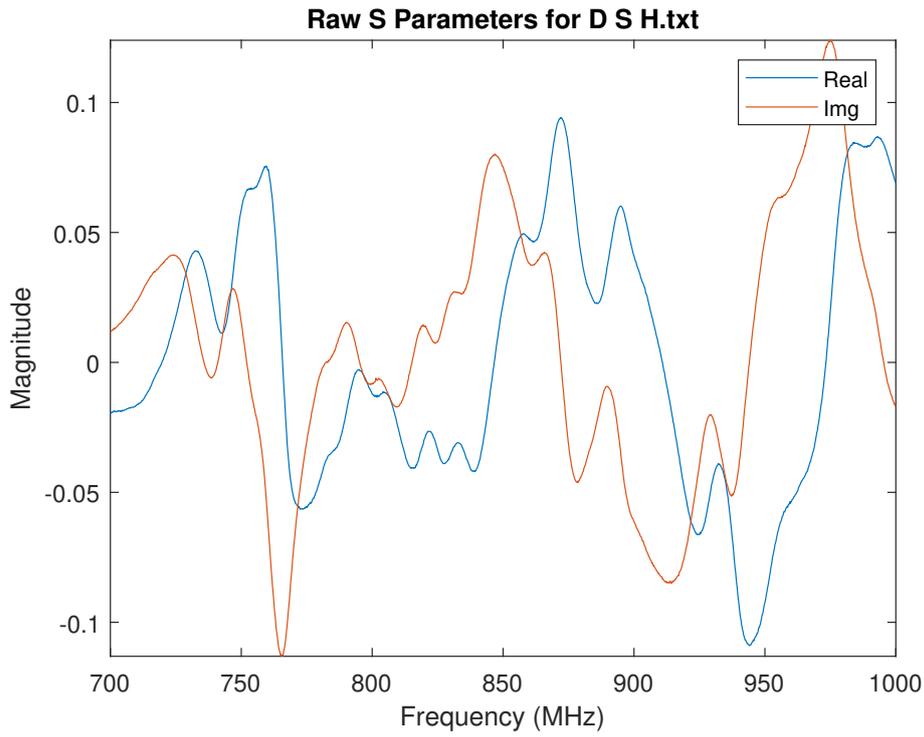
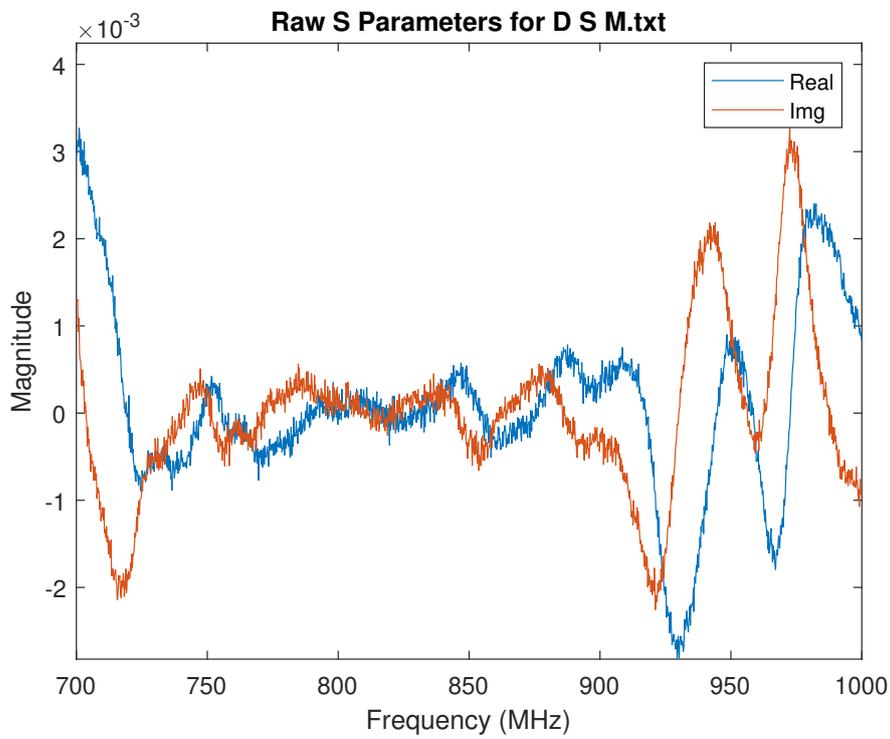
Field	Value
Rysm	200000
ModulationOrder	4
Interpolation	2
Decimation	1
Tsym	5.000e-06
Fs	400000
BarkerCode	[1 1 1 1 1 -1 -1 1 1 -1 1 -1 1]
BarkerLength	13
HeaderLength	26
Message	'Hello world'
MessageLength	16
NumberOfMessage	100
PayloadLength	11200
FrameSize	5613
FrameTime	0.0281
RolloffFactor	0.5000
ScramblerBase	2
ScramblerPolynomial	[1 1 1 0 1]
ScramblerInitialConditions	[0 0 0 0]
RaisedCosineFilterSpan	10
Desired Power	2
Averaging Length	50
MaxPowerGain	60
MaximumFrequencyOffset	6000
PhaseRecoveryLoopBandwidth	0.0100
PhaseRecoveryDampingFactor	1
TimingRecoveryLoopBandwidth	0.0100
TimingRecoveryDampingFactor	1
TimingErrorDetectorGain	5.400
PreambleDetectorThreshold	0.800
MessageBits	[11200*1 double]
BerMask	[7700*1 double]
MasterClockRate	100000000
USRPCenterFrequency	915000000
USRPGain	31
USRPFrontEndSampleRate	400000
USRPInterpolationFactor	250
USRPFrameLength	11226
USRPFrameTime	0.0281
StopTime	10

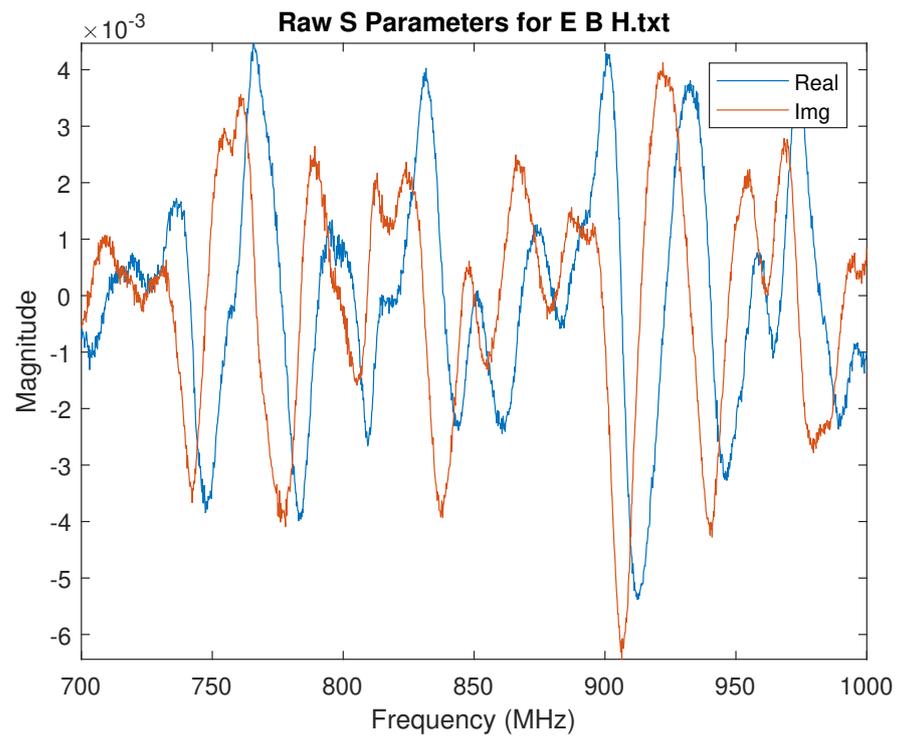
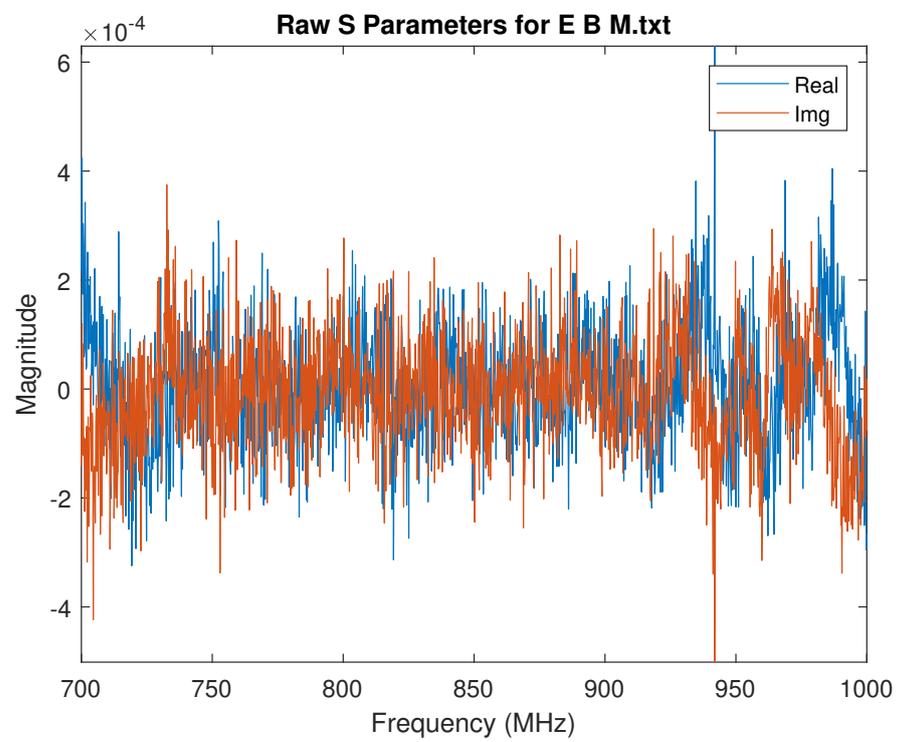
Appendix F

Measurements for K-Factor Replication

Figure F.1: Complex S_{21} for DBH channel.Figure F.2: Complex S_{21} for DBM channel.

Figure F.3: Complex S_{21} for DEH channel.Figure F.4: Complex S_{21} for DEM channel.

Figure F.5: Complex S_{21} for DSH channel.Figure F.6: Complex S_{21} for DSM channel.

Figure F.7: Complex S_{21} for EBH channel.Figure F.8: Complex S_{21} for EBM channel.