

The Development of a System that Emulates Percussion to Detect the Borders of the Liver

by

Hanz Frederick Rauch

Thesis presented at the University of Stellenbosch

in partial fulfilment of the requirements

for the degree of

Masters of Science in Mechatronic Engineering



Department of Mechanical and Mechatronic Engineering

University of Stellenbosch

Private Bag X1, 7602 Matieland, South Africa

Supervisor: Prof C. Scheffer

Co-Supervisor: Dr G-J van Rooyen

March 2009

Declaration

I, the undersigned, hereby declare that the work contained in this thesis is my own original work and that I have not previously in its entirety or in part submitted it at any university for a degree.

Signature:

H.F. Rauch

Date:

Copyright © 2009 University of Stellenbosch
All rights reserved.

Abstract

Percussion is a centuries old bedside diagnostic technique that is used to diagnose various conditions of the thorax and abdomen, among these, abnormalities of the liver. The physician taps the patient's skin in the area of interest to determine the qualities or presence of the underlying tissue or organ, by listening to the generated sound.

The research contained in this thesis views percussion as a system identification method which uses an impulse response to identify the underlying system. A design employing an electromagnetic actuator as input pulse generator and accelerometer as impulse response recorder was motivated and built.

Tests were performed on volunteers and the recorded signals were analysed to find methods of identifying the presence of the liver from these signals. The analyses matched signals to models or simply extracted signal features and matched these model parameters or signal features to the presence of the liver. Matching was done using statistical pattern recognition methods and the true presence of the liver was established using MR images. Features extracted from test data could not be matched to the presence of the liver with sufficient confidence which led to the conclusion that either the test, apparatus or analysis was flawed.

The lack of success compelled a further test on a mock-up of the problem – a silicone model with an anomaly representing the organ under test. Results from these tests showed that signals should be measured further from the actuator and the approach followed during this test could lead to the successful location of the anomaly and discrimination between subtle differences in the consistency thereof.

It is concluded that further research should aim to first validate percussion as performed by the physician and increase complexity in a phased manner, validating results and apparatus at each step. The approach followed was perhaps too bold in light of the lack of fundamental understanding of percussion and the underlying mechanisms.

Uittreksel

Beklopping is 'n eeu-oue mediese tegniek wat gebruik word om toestande van die buik en borskas te identifiseer en diagnoseer. Onder andere word dit ook gebruik om abnormaalhede of siektes van die lewer te identifiseer. Die dokter gebruik sy vinger om die area van belang te klop en gebruik dan sy ontwikkelde tas en gehoor sensasies, gebind met sy ondervinding, om die kwaliteit of teenwoordigheid van onderliggende weefsel of organe te bepaal.

Die navorsing vervat in hierdie tesis beskou beklopping as 'n stelsel identifiserings metode wat 'n impuls respons gebruik om die onderliggende stelsel te identifiseer. 'n Meganisme was ontwerp om beklopping na te boots, waar die klop uitgevoer word met 'n elektromagnetiese vibrator en die reaksie met versnellingsmeters gemeet word. Die ontwerp is uit literatuur en praktiese oorwegings gemotiveer, ontwerp en gebou.

Toetse was op vrywilligers uitgevoer en die opgeneemde seine is geanaliseer om maniere te vind om die seine te manipuleer om die teenwoordigheid van die lewer te voorspel. Seine is gekoppel aan modelle of eenvoudig verwerk om parameters uit die seine te onttrek. Statistiese patroon herkenning metodes het gepoog om model of sein parameters te koppel aan die teenwoordigheid van lewer weefsel, waar laasgenoemde deur MR beelde vasgestel is. Die parameters kon egter nie met genoegsame vertroue gekoppel word aan die posisie van die lewer nie – dus die gevolgtrekking dat die metode, apparaat of verwerkingstegnieke inherent foutief is.

Die mislukking in die voorafgaande toets het gelei tot 'n volgende toets waar 'n model van die probleem gemaak is. Die model bestaan uit 'n silikon blok en 'n voorwerp van 'n ander materiaal binne-in om die lewer of ander orgaan voor te stel. Resultate wys dat tydens toetse op vrywilligers die sein te na aan die intree gemeet is. Verder word ook gewys dat die posisie van die voorwerp suksesvol geïdentifiseer kan word en dat daar gediskrimineer kan word tussen toestande van die voorwerp.

Die gevolgtrekking word gemaak dat verdere navorsing moet fokus daarop om eers beklopping, soos deur 'n dokter beoefen, te verifieer. Hierna kan die kompleksiteit van die metode op 'n gefaseerde manier verhoog word deur stelselmatig funksies van die dokter (hand bewegings, gehoor en laastens interpretasie) te vervang. Verhoogde kompleksiteit en apparaat moet ook in elke fase geverifieer word. Die aanslag was moontlik oorhaastig gewees terwyl die onderliggende meganismes van beklopping nie goed verstaan word nie.

Acknowledgements

Special thanks are given to:

Prof Cornie Scheffer

Dr Gert-Jan van Rooyen

Dr Dirk Koekemoer

Thanks to all the students who volunteered for tests during the study. Without them the project would not have been possible.

Thanks to fellow researchers who not only shed light on some dark corners, but also served as motivation and inspiration during the course of this study.

Dedications

I dedicate this thesis to my wife and love of my life, Diane-Maree.

Contents

Declaration	i
Abstract	ii
Uittreksel	iii
Acknowledgements	iv
Dedications.....	v
Contents.....	vi
List of Figures	x
List of Tables.....	xiv
Nomenclature	xv
Chapter 1 Introduction	1
1.1 Background.....	1
1.2 Motivation, Goal and Scope	2
1.3 Overview.....	3
Chapter 2 Background and Literature Review	5
2.1 Percussion	5
2.1.1 Technique	6
2.1.2 History and Origin.....	7
2.1.3 Criticism and Limitations	9
2.1.4 Emulation of Percussion.....	10
2.2 The Liver.....	11
2.2.1 Characteristics	12
2.2.2 Diseases of the Liver	13
2.2.3 Detection of Liver Disease	14
2.3 Conclusions.....	16
Chapter 3 System Development.....	17
3.1 Conceptual Design	17
3.2 Actuator	18

3.3 Sensors	21
3.4 Data Acquisition System.....	24
3.5 Signal Conditioning System	25
3.6 Computer System.....	27
3.7 Frame	29
3.8 Conclusions.....	30
Chapter 4 Human Subject Tests	32
4.1 Volunteers	32
4.2 Test Procedure	33
4.3 Signal Processing and Feature Extraction.....	37
4.3.1 Direct Feature Extraction.....	37
4.3.1.1 Pre-processing of Data	38
4.3.1.2 Identifying Features.....	39
4.3.2 Model Fitting Methods	42
4.3.2.1 Autoregressive Models.....	43
4.3.2.2 Mechanical Model.....	46
4.4 Classification.....	52
4.4.1 Classification Introduction	53
4.4.2 Extracting Ground Truth	58
4.4.3 Classification of Extracted Features	60
4.4.3.1 Directly Extracted Features	60
4.4.3.2 Autoregressive Models.....	61
4.4.3.3 Mechanical Model.....	64
4.5 Conclusions.....	66
Chapter 5 Model Tests	68
5.1 Experiment Design.....	68
5.2 Test Procedure	71
5.3 Evaluations.....	72
5.3.1 Actuation Signal for Impulse Estimation	73
5.3.2 Position of Actuation and Measurement.....	78
5.3.3 Finding the Phantom.....	81
5.3.4 Repeatability of Procedure	84
5.3.5 Effect of Actuation Parameters.....	85
5.3.6 Differentiating Anomaly Types.....	87
5.4 Conclusions.....	88

Chapter 6 Conclusions and Recommendations	90
List of References.....	93
Appendix A Electronic Designs	97
A.1 Design of Signal Conditioning and Routing PCB.....	97
A.1.1 Requirements	97
A.1.2 System Block Diagram	97
A.1.3 MUX system	98
A.1.4 Analogue Signal Conditioning System.....	101
A.1.5 Layout Considerations	103
A.1.6 Detailed Design Drawings.....	104
A.2 Design of Power Source PCB	106
A.2.1 Requirements	106
A.2.2 System Block Diagram	106
A.2.3 AC-DC Power Conditioning System.....	107
A.2.4 Fixed Voltage Regulation Systems.....	109
A.2.5 Layout Considerations	109
A.2.6 Detailed Design Drawings.....	110
A.3 Design of Accelerometer PCBs	111
A.3.1 Requirements	111
A.3.2 Detailed Design Drawings.....	112
A.4 FSR Calibration.....	113
Appendix B Software Designs	115
B.1 <i>get_aver_pulse.m</i>	115
B.2 General Feature Extraction Flow	116
B.2.1 Direct Feature Extraction.....	117
B.2.2 Extracting Autoregressive Features	117
B.2.3 Extracting Mechanical Model Features	118
Appendix C Mechanical Designs	120
C.1 Frame Design	120
C.2 Phantom test setup.....	127
Appendix D Additional Information	128
D.1 Body Planes and Orientations	128
Appendix E Datasheet Excerpts	129

E.1 Vibrational Shaker V201-PA25E from LDS Group.....129
E.2 Smooth-On Dragon Skin Silicone Rubber Compounds131

List of Figures

Figure 2.1: Diagrammatical representation of percussion ([6])	6
Figure 2.2: Percussion of the liver [7].....	6
Figure 2.3: Anatomical views and position of the liver [35].....	12
Figure 3.1: Vibrational shaker as implemented.....	19
Figure 3.2: FSR and Voltage Divider Configuration	20
Figure 3.3: Force Sensing Resistor signals	21
Figure 3.4: Image of Accelerometer Sensor Unit.....	23
Figure 3.5: Accelerometer saturation	24
Figure 3.6: Accelerometer discontinuities.....	24
Figure 3.7: PMD 1608FS USB Block Diagram and pin-out [53]	25
Figure 3.8: Signal Conditioning and Routing Block Diagram	26
Figure 3.9: Power Source PCB system block diagram.....	27
Figure 3.10: <i>DataloggerGUI.m</i> general software flow diagram	28
Figure 3.11: <i>DataloggerGUI.m</i> graphical user interface.....	29
Figure 3.12: Frame design and measurement setup	30
Figure 4.1: Experimental setup - Volunteer under frame.....	34
Figure 4.2: Measurement points on patient	34
Figure 4.3: Example of signals measured	35
Figure 4.4: Examples of signals measured - enlarged.....	36
Figure 4.5: Typical MRI Image and analysis thereof.....	36
Figure 4.6: FSR signals showing time jitter	38
Figure 4.7: Time aligned and averaged FSR signals.....	39
Figure 4.8: Series of time aligned accelerometer signals	39
Figure 4.9: Extracting rise time, settling time and frequency	40
Figure 4.10: Amplitude spectra of multiple signals	41
Figure 4.11: Features (percentage of energy in subsets of spectra) extracted from amplitude spectra.....	41
Figure 4.12: Example 1 of fit with ARX model.....	45
Figure 4.13: Example 2 of fit with ARX model.....	45
Figure 4.14: Distribution of fitness values for data fitting to an ARX model.....	46
Figure 4.15: A dual Mass-Spring-Damper system.....	46
Figure 4.16: Simulink model of dual Mass-Spring-Damper system.....	48
Figure 4.17: The basic cycle of the GA [56].....	49
Figure 4.18: Using mechanical model fit with simulated noise	50

Figure 4.19: Example 1 of good fit on mechanical model	50
Figure 4.20: Example 2 of good fit on mechanical model	51
Figure 4.21: Example 1 of bad fit on mechanical model	51
Figure 4.22: Example 2 of bad fit on mechanical model	51
Figure 4.23: Distribution of fitness values for data fitting to mechanical model	52
Figure 4.24: An example of mapping data using PCA.....	54
Figure 4.25: Difficult data set.....	54
Figure 4.26: Fisher Linear Discriminant example.....	55
Figure 4.27: Quadratic Bayes discriminant classifier example	56
Figure 4.28: The MOG classifier results	56
Figure 4.29: The ANN classifier results.....	57
Figure 4.30: Labelled labels on data capture locations for two subjects.....	59
Figure 4.31: Scatter plot of first three directly extracted features	60
Figure 4.32: PCA of a ARX data set.....	63
Figure 4.33: PCA of a mechanical model data set	64
Figure 5.1: Silicone model on base	69
Figure 5.2: Phantom tests measurement setup	70
Figure 5.3: Numbering of input/outputs on silicone model	70
Figure 5.4: Model Dimensions.....	71
Figure 5.5: Impulse response estimation methods	74
Figure 5.6: Comparison of impulse response estimation methods.....	75
Figure 5.7: Comparison of impulse estimation methods - bode plots.....	75
Figure 5.8: Time domain signals and estimations from model	76
Figure 5.9: Frequency domain signals and estimations from model.....	76
Figure 5.10: Coherence between input and output.....	77
Figure 5.11: Input-output combinations.....	78
Figure 5.12: Comparing signals to model (Air Phantom)	79
Figure 5.13: A signal trace from image depiction.....	79
Figure 5.14: Correlation between distance and peak time	80
Figure 5.15: Signals with jelly phantom with zero offset	81
Figure 5.16: Example of extracting peak time and amplitude.....	82
Figure 5.17: Example of extracting $f_{std\ dev}$ and A_{60Hz}	82
Figure 5.18: Comparing extracted features to silicone model parameters (offset of five positions).....	83
Figure 5.19: Predicting location of anomaly using extracted parameters	83
Figure 5.20: Scatter plot of features vs. disrupted distance.....	84
Figure 5.21: Repeatability of some signals and their fits	85

Figure 5.22: Effect of variation of input angle (using Chirp method).....	86
Figure 5.23: Effect of variation of input angle (using Pulse method).....	86
Figure 5.24: Effect of variation of offset force (using Chirp method).....	86
Figure 5.25: Effect of variation of offset force (using Pulse method).....	86
Figure 5.26: Identifying the anomaly - visual analysis.....	87
Figure 5.27: Identifying the anomaly - scattered data.....	88
Figure A-1: Signal Conditioning and Routing Block Diagram.....	98
Figure A-2: MAX4531 Block Diagram [64].....	98
Figure A-3: MUX subsystem circuit diagram.....	100
Figure A-4: Signal conditioning sub system simulation model.....	101
Figure A-5: Signal conditioning amplitude response.....	102
Figure A-6: Signal conditioning sub system circuit diagram.....	103
Figure A-7: Signal Conditioning and Router PCB Schematics (part 1).....	104
Figure A-8: Signal Conditioning and Router PCB Schematics (part 2).....	105
Figure A-9: Signal Conditioning and Router PCB Board Top.....	105
Figure A-10: Signal Conditioning and Router PCB Board Bottom.....	106
Figure A-11: Power Source PCB system block diagram.....	107
Figure A-12: AC-DC Power conditioning system.....	107
Figure A-13: Fixed voltage regulation systems.....	109
Figure A-14: Power PCB top view.....	110
Figure A-15: Power PCB bottom view.....	111
Figure A-16: ADXL330 MEMS accelerometer block diagram (extracted from datasheet).....	111
Figure A-17: Accelerometer PCB schematic.....	112
Figure A-18: Layout of Accelerometer PCB (single layer).....	113
Figure A-19: FSR resistance-force curve [51] and curve fit.....	114
Figure B-1: <i>getting_aver_pulse.m</i> flow diagram.....	115
Figure B-2: General Feature Extraction Flow.....	116
Figure B-3: Direct feature extraction.....	117
Figure B-4: Extracting autoregressive features.....	117
Figure B-5: Extracting mechanical model features, Genetic Algorithm.....	118
Figure B-6: Extracting mechanical model features, Genetic Algorithm.....	119
Figure C-1: Frame concept.....	120
Figure C-2: Mechanical drawing: frame back part.....	121
Figure C-3: Mechanical drawing: frame half circle.....	122
Figure C-4: Mechanical drawing: frame base.....	123
Figure C-5: Mechanical drawing: frame feet.....	124

Figure C-6: Mechanical drawing: frame slider	125
Figure C-7: Mechanical drawing: frame clamps	126
Figure C-8: Phantom test setup	127
Figure D-1: Summary of body planes and orientations [66].....	128

List of Tables

Table 2.1: Summary of percussion sounds [17].....	8
Table 4.1: Actuation point labels	34
Table 4.2: ARX Fits Results (sub set).....	44
Table 4.3: Confusion matrix from Fisher Linear Discriminant example	55
Table 4.4: Confusion matrix from Quadratic Bayes discriminant classifier example.....	55
Table 4.5: Confusion matrix from MOG classifier example.....	57
Table 4.6: Confusion matrix from ANN classifier example	57
Table 4.7: Summary of assigned labels (ground truth)	58
Table 4.8: Confusion matrix results on directly extracted features	61
Table 4.9: Percentage of ARX data points used.....	62
Table 4.10: Preliminary evaluations on ARX data.....	63
Table 4.11: Confusion matrix results on mechanical parameter data.....	64
Table 4.12: Confusion matrix results on mechanical parameter data.....	65
Table 5.1: Fitness measures for repeatability measurement.....	85
Table A-1: MAX4531 Truth Table	99
Table A-2: MUX system connection table.....	99
Table A-3: Power budget for positive voltage	108
Table A-4: Power budget for negative voltage	108

Nomenclature

A – Amplitude	LED – Light Emitting Diode
AC – Alternating Current	LFT – Liver Function Test
ACC – Accelerometer	LPF – Low Pass Filter
ANN – Artificial Neural Network	m – Meter, measurement of distance
ARX – Autoregressive model with exogenous inputs	MEMS – Micro-electromechanical Systems
CT – Computed Tomography	MOG – Mixture of Gaussian (Classifier)
DAQ – Data Acquisition (Device)	MRI – Magnetic Resonance Imaging
dB – Decibel, measure of power	MUX – Multiplexer system
DC – Direct Current (but used to indicate an unchanging voltage signal component)	N – Newton, measure of force
DS – Direct Search Algorithm	ODE – Ordinary Differential Equation
ECG – Electrocardiogram	Op-amp – Operational Amplifier
EMI – Electromagnetic Interference	PC – Personal Computer
F – Force	PCA – Principal Component Analysis
f – Frequency	PCB – Printed Circuit Board
FFT – Fast Fourier Transform	QDC – Quadratic Bayes Discriminant assuming Gaussian densities
G – Conductor or conductance (inverse of resistance)	R – Resistor or resistance
g – Gram, measurement of weight	RMS – Root-mean-squared (VAC rating).
GA – Genetic Algorithm	US – Ultrasonography
GND – Ground referring to potential	USB – Universal Serial Bus
GP – General Practitioner or Physician	V – Volts
GUI – Graphical User Interface	VA – Volts-amps indicate power rating in AC situations.
Hz – Hertz, measure of frequency	VAC – AC Voltage
I/O – Input/Output	Ω – Ohm, measure of electrical resistance
IC – Integrated Circuit	ω – Radians/second, measure of frequency
LDA – Linear Discriminant Analysis	
LDV – Laser Doppler Vibrometer	

Chapter 1 Introduction

The study summarised in this thesis involves the development of a device used for the emulation of *medical percussion* in order to determine the borders or condition of the human liver. As an introduction, this chapter will give background on percussion and how it is used to detect liver disease, discuss the goal and scope of this study, give some insight into the motivation of the study and give an overview of this thesis report.

1.1 Background

Medical percussion is a centuries old technique that is used by physicians to assess the condition of the thorax and abdomen. The technique forms part of the clinical or physical examination (which also includes palpation, auscultation and patient history), performed by a physician when investigating the patient for primary signs of disease. Percussion mainly involves the tapping of areas of the body by the fingers of the physician – either directly with two fingers; or indirectly, where one hand is used as a base on the skin and the other to tap the middle finger of the first hand. The physician listens to the resulting sound to determine the presence of underlying tissue, bone, air or abnormality.

Unfortunately, percussion is not well understood and the technique itself is not agreed upon. Even so, percussion is simplified for the purposes of this study: when percussion is applied to a patient's body, one can view the input stroke, underlying organs and output sound as a very complex system responding to an impulse. The underlying system is identified by a highly specialized detection system – the trained physician's ear (and tactile sense) and mind. This thesis views percussion as a system identification method, a method that could perhaps be emulated by a machine. Previous studies that involve the emulation of percussion were identified from literature and, although none of these studies were focused on actual emulation of percussion, they do provide valuable insight into possible approaches and solutions.

Percussion is applied to detect and diagnose various organs, but the scope of this thesis is limited to the liver, where percussion is used to detect an over- or undersize or abnormality thereof. The liver is a complex, but important organ in the human body. It serves various important bodily functions (among these are the metabolism of energy and synthesis of regulatory proteins without which the body cannot function) and is therefore vulnerable to numerous diseases – most of which would cause the size or elasticity of the liver to change and therefore might be detectable by percussion. Most diseases of the liver are either symptomatic of other diseases or lead to other diseases. If these conditions go unchecked it could have dire effects and even lead to death.

More insight into percussion and the liver is given in Chapter 2 Background and Literature Review.

1.2 Motivation, Goal and Scope

Identification of liver abnormality is difficult at best and often early diagnosis could save a life. Currently, a physical exam by a general practitioner (GP) is the first step in detection and diagnosis of various diseases of the liver. The doctor uses a combination of percussion, palpation and auscultation to determine the boundaries and hardness of the liver. If the doctor finds that the liver has an abnormal sound or that it is abnormally large or small, the doctor will refer the patient for a sonogram (ultrasound imaging), Elastogram (ultrasound elasticity imaging), X-ray, liver function test, liver biopsy, CT (Computerized Tomography) or MRI (Magnetic Resonance Image). These techniques are effective in further diagnosis of liver disease, but are expensive, require specialised facilities and rely on the doctor's experience and presence to initially detect abnormalities during the physical exam. Furthermore, some of these methods rely on potentially dangerous electromagnetic radiation or are invasive as with a liver biopsy.

The problems with high cost and lack of skills are quite apparent in rural areas where access to experienced doctors and services are limited and/or the patients are sometimes too poor to afford proper health care. Physical examinations are performed by nurses, who might not have the experience and expertise of a doctor and referral to a secondary institute for further investigation might be too expensive. Wrongful primary investigation can prove to be costly, resulting in either very late diagnosis or unnecessary expenses. This underlines the necessity for a better, but inexpensive, unbiased liver disease screening device in rural areas.

The University of Stellenbosch's Mechanical Engineering Department was approached by *GeoAxon Holdings (PTY) Ltd*, a medical research company, with the request to research, design and develop a device that uses the principles of percussion to non-invasively extract information about the condition of the liver, in an unbiased fashion and at a fraction of the cost of existing techniques. Further work might include application of the device or derivatives thereof in diagnostics of other medical problems, including other diseases of the abdomen, diseases of the lungs, bone fractures and diseases of the lower abdomen.

Unfortunately, the applications of percussion is very broad (lungs, spleen, heart, bones, cranium, etc.) and it would be an unrealistic expectation to apply this study to all of these applications. It was therefore decided to narrow the scope of the study to focus on one of the most important organs in the human abdomen, the liver. Percussion is mainly used to detect the boundaries of the liver, which would indicate whether it is over or under size, or sometimes to determine liver quality like hardness. The problem is therefore to develop a system to emulate percussion for the purposes of detecting abnormalities of, or just the liver (and therefore boundary). In order to solve this problem one would need to understand the technique of percussion and especially how it applies to the liver. If a deep understanding of percussion existed one could perhaps simulate a mathematical model of the problem with different actuators and sensors, but no such deep understanding exists. A dedicated system would therefore need to be developed with an actuator to generate an input similar to (or better than) that of the physician as well as a sensor system to record the reactionary signal(s). The design choices need to be motivated as far as possible and an applicable experimental setup needs to be designed to test the concept.

If one could exactly emulate the action of the physician and exactly mimic the sensor that a physician uses to assess the reaction (sound and tactile – ears and hands), one would expect to measure what the physician hears – however, this scenario is highly unlikely and one would most probably have to develop a signal processing and classification system that generates the outputs that a physician would have for a particular point on a patients body. Where percussion is applied to the liver, this output is basically one of the following: ‘Liver’ or ‘No liver’, and in a more advanced situation: ‘Abnormal liver’ or ‘Normal liver’. The topography of these outputs is also of relevance as they indicate where the liver ends and therefore indicates the size of the liver.

The scope of the project can be summarized by the following steps:

- Research and understand percussion as it applies to detection of the liver.
- Review literature that relates to this study in an attempt to better understand the problem and identify possible solutions.
- Develop and acquire the necessary tools.
- Develop an experimental setup for data collection and collect data.
- Develop signal processing techniques to identify the liver.
- Validate results.
- Document research and results.

Unfortunately, an initial system did not yield a solution that could successfully identify the liver and a secondary study followed – a mock-up of the problem was created and various uncertainties regarding the method were investigated. This secondary investigation was not part of the initial scope of the project, yet was seen as a necessary step to clarify why the initial approach failed. If information was revealed that could help in the analysis of the initial study’s data, this would have been used and applied, but this was not the case. The initial experiment was done on volunteers and as such is referred to as *Human Subject Tests*. The subsequent experiment was applied to a silicon model with a phantom anomaly inside and is thus referred to as *Model Tests*.

1.3 Overview

This thesis report documents the relevant information collected during the research process, the processes followed in the attempt to reach the research goals, the results of these processes, and conclusions based on the results, findings and experiences.

Chapter 2 contains in depth information about the technique and history of percussion as well as information about the liver. Percussion has been criticised over the centuries and various references are invoked to review criticisms and perceived limitations – these findings show a lack of understanding about the physics of percussion which contributed to of the lack of success of the project. There have been various studies that have, inadvertently, copied components of the technique of percussion, but for completely different reasons – these are also discussed with a view as to how they apply to this study.

The importance of the liver is emphasized with the main focus on liver abnormalities potentially detectable by percussion. Current techniques that are used to detect liver abnormalities are pointed out with special attention to their perceived shortcomings and how a suggested device may be useful.

System design is documented in Chapter 3 which contains information about the acquisition and development of the tools required to perform the two main actions of the physician – to generate an input pulse and to record and interpret reactionary signals. Design choices are motivated based on various literature sources and deductions. Critical evaluation of selected sensors' performance is also performed and reference is made to subsequent analyses of sensors in Chapter 4.

Chapter 4 contains details on the test procedure and subsequent analysis of the experiment on human subjects. The experiment involves percussing fixed points on a volunteer's body while measuring the surrounding skin's vibrational response. The signals produced by the skin surrounding the percussion point are analysed and the analyses are focused on extracting signal parameters and subsequently matching these parameters to underlying physical features.

The experiment on human subjects was unsuccessful in arriving at a method that could successfully predict the presence of liver tissue. Chapter 5 discusses a subsequent experiment which was performed on a silicon model with an anomaly. This experiment was aimed at evaluating other approaches (i) as far as the input signal for percussion is concerned – instead of using a pulse as input to the actuator, chirp and white noise signals are used to estimate an impulse response; and (ii) as far as the location of the actuator input and sensor output is concerned – it may be that better signals are acquired further from the input and with the input and output at different positions relevant to the organ under investigation.

The results of both these experiments are discussed at the end of each respective chapter and collectively summarised with conclusions and recommendations for further investigations in Chapter 6.

Chapter 2 Background and Literature Review

Percussion is a technique that has been developed over centuries. This chapter provides insight into the technique, origin and limitations of medical percussion and also looks at studies which involved the emulation of percussion or touch on points that should be considered when such an attempt is made. An overview of the liver and its functions is also included with a look at diseases that may be detectable by percussion. Modern techniques to detect liver abnormalities are evaluated with attention to shortcomings that may be solved by the suggested device or system..

This chapter deals at some length with information that isn't used directly for research purposes – the lack of consistency in how percussion should be performed and how it actually works. This information is presented to emphasize the lack of conclusive information, a factor that contributed to the lack of success in this project.

2.1 Percussion

In brief, percussion is the method where the surface of the skin is tapped to determine the underlying structures' quality by assessing the resulting sound [1]. It forms part of the clinical or physical examination performed by a physician when investigating the body of a patient for signs of disease. It is mainly used for assessment of the thorax (lungs heart) and abdomen (liver, spleen, kidneys, stomach), but can also be used to examine bones [2] or the head [3]. Other components of the clinical examination are inspection, palpation and auscultation, all of which would be performed after taking a medical history.

The technique of percussion, along with the other clinical examination techniques, has declined in popularity with the advent of modern diagnostic techniques [4]. Even so, the technique is still applied by physicians and is essential in the early detection of abnormalities in the chest and abdomen – without applying some form of bedside diagnostic, a physician could not even consider further, more advanced, diagnostic techniques.

2.1.1 Technique

Percussion is described in medical textbooks [5] and journals [4]. One hand is used as a base and placed against the body in the area to be percussed. The middle finger of the base hand (also called pleximeter) is placed on the area to be percussed, away from bony prominences. The dorsal side of the middle phalanx (finger bone) of the middle finger is struck by the tip of the middle finger of the other hand (also called the plexor) – see Figure 2.1. The resultant normal percussion sound can be placed in one of three categories – resonance, tympany and dullness (signal trace depictions in Table 2.1). Normally resonance is expected over lung tissue, tympany over portions of the abdomen and dullness over solid organs like the liver. Sounds between resonance and tympany are called hyper resonant and relative dullness describes the sound between tympany and dullness. Abnormalities are detected by comparing sounds on both sides of symmetric organs or by detecting an abnormal sound.




Figure 2.1: Diagrammatical representation of percussion ([6])

As this investigation is focused on percussion of the liver, only the analysis of liver percussion sounds is discussed. The liver is percussed on the right mid-clavicular line (the line vertically down from the middle of the right clavicle). The dull sound of the liver is usually audible from the upper edge of the sixth rib down to the lower edge of the tenth rib, just at the costal margin. These anatomical areas are depicted in Figure 2.2.

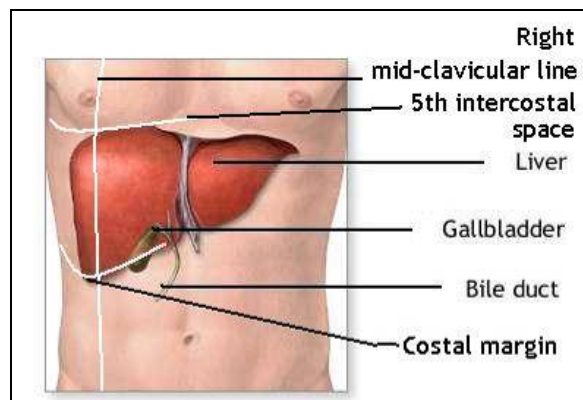


Figure 2.2: Percussion of the liver [7]

2.1.2 History and Origin

According to Fedorowski [4], McGee [8] and Yernault et al. [9], Leopold Auenbrugger published his *Inventum Novum ex Percussione Thoracis Humani* ("New Invention of Human Chest Percussion") in 1761 which contained the first detailed description of chest percussion. Percussion as such, was however not invented by Auenbrugger, and was often performed to diagnose the human abdomen at the Hippocratic School. Swiss veterinarians even percussed the heads of cattle to detect cysticerci. Auenbrugger was the first person to regard the technique of percussion as an essential part of the physical diagnosis. Auenbrugger suggested [8] that the chest should be percussed directly with the tips of the fingers of one hand, that a healthy sound resembled the sound of a stifled drum and a duller sound on the one side of the chest may indicate pleural effusions, pneumonia or asthma (among various other).

Unfortunately, Auenbrugger's work was not taken seriously by physicians of his time and it was only when Frenchman Jean Corvisart came upon his work that percussion became an integral part of the physical examination. According to Yernault and Bohadana [9], Corvisart believed that the physician's sense organs, education and exercise were essential practices in becoming good bedside physicians. Corvisart used percussion as his main method in discovering and localisation of many diseases. He added the use of the palmar surface of extended and approximated fingers to percuss regions of the chest. He introduced the technique to his students of whom René Laënnec was one. Laënnec felt that percussion on its own was insufficient for the diagnosis of diseases of the chest which led him to develop the technique of auscultation (as well as the stethoscope) which also forms part of the suite of techniques available to the medical practitioner in the physical examination.



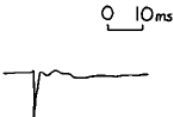
In 1828 Pierre A. Piorry introduced *Indirect* (or mediate) *Percussion* where an impulse is given by the right middle finger (called the plexor) to a solid resonant body (called pleximeter), like an ivory sheet, which is applied to the organ or area under investigation. Resulting sounds supposedly reveal the condition of the area or organ under investigation. Risse [10] describes three circumstances where Piorry felt direct percussion (as suggested by Auenbrugger) was inadequate: (i) it can cause pain where the patient suffers from inflammation of the pleura; (ii) thick tissue under the skin limits the transmission of *auditory sensations*; and (iii) the use of one's palm vibrates a large area which might cause the practitioner to miss smaller lesions. To Piorry's dismay, two of his students, William Stokes and James Hope, simplified the technique by replacing the ivory sheet with their left middle finger and this is the method that is still used today.

Auenbrugger, Corvisart and Laënnec's technique can be described by the term *Direct Percussion*, as the finger is directly in contact with the skin during the action. *Topographic* and *Comparative Percussion* describe how the technique is applied - Topographic Percussion involves identifying underlying tissue by assessing the percussion note in that particular location, and Comparative Percussion identifies abnormalities by comparing the result obtained by percussing at different areas on the surface of the body.

The third technique of percussion (pertaining to the way it is applied) is *Auscultatory Percussion* which, according to [8], was introduced by Camman and Clarke in 1840. Developed to determine the size and density of organs, most notably the heart and the liver, it was practiced with a solid stethoscope placed against the chest with the one hand and the other striking the wall of the chest whilst listening for the sound [11]. In a study by Sehrwald in 1904, however, it was found that the sounds had little to do with the underlying organs – this finding led to the failure of Auscultatory Percussion [8]. Recent studies by Guarino [12], [13], [14] and a very recent study in India [15] have resurrected Auscultatory Percussion as a diagnostic technique to detect pleural effusions, yet McGee remains sceptical [8]. As pleural effusions are not related to the liver, this method can be abandoned for the purposes of this study.

The descriptions of percussion sounds have varied greatly over the centuries starting with the descriptions by Auenbrugger of tympanic, dull and indistinct [9], to Joseph Skoda (who pioneered percussion throughout Europe after Piorry) who described percussion sounds with the terms *full-empty*, *clear-dull*, *tympanic-nontympanic* and *high-low* [16]. The terms introduced by Skoda lead to ambiguous definitions being given to some sounds and these terms were eventually discarded and replaced with the terms used to describe sounds today: *tympany-resonance-dullness* [8]. Sound pressure waveforms of these sounds, as extracted from Murray and Neilson [17], are summarized in Table 2.1. Also shown in the table are the topographical locations where the notes are expected to be found: Tympany over the abdomen, Resonance over the lung, and Dullness over the liver.

Table 2.1: Summary of percussion sounds [17]

	Tympany (over the abdomen)	Resonance† (over the lung)	Dullness‡ (over the liver)
Sound pressure waveform (y axis = N/M ²)			
Waveform description	Dampened single frequency (200 - 600 Hz)	Rarefaction pressure spike and two compression peaks	Rarefaction pressure spike only
Duration	>40 msec	Approximately 15 msec	<3 msec
Frequency content	Dominant single frequency	Assessed by the ear to have a low-frequency component, compared with dullness§	Assessed by the ear to be high in frequency vs. resonance

Unfortunately, graphical representations are not accessible to a physician when performing percussion and, according to McGee [8], the interpretations of these descriptions are varied. McGee cites various references describing and deliberating the qualities of these sounds as heard and interpreted by an individual. Modern evaluations of percussion sounds, like that by Murray and Neilson [17] & [18], have gone a long way in clarifying these definitions.

One element of percussion theory that is still in dispute is the factors that cause the percussion note. In the study by McGee [8] two theories originating from the 1800s are discussed: *Cage Resonance Theory* and *Topographical Percussion Theory*.

Cage Resonance Theory originated in 1852 from experiments performed by a Russian, J.F. Mazoon. He observed that any external pressure on the chest (like a stretcher, hand or pillow) dampened the percussion note by impeding the movement of the chest wall. He removed the abdominal organs from cadavers and found that the percussion note becomes more resonant at distant sites, even the upper part of the chest. Upon replacing pieces of the liver and holding it against the thoracic wall, the percussion note over that area was dull. The degree to which distant notes were dampened due to the presence of liver tissue on the inside wall of the thorax was also proportional to the force of applied to the piece of the liver. In a study by Murray and Neilson [19], attempts are made to account for the frequency of the percussion note by percussing single ribs and assessing the acceleration thereof – it is found that the measured oscillation is half the frequency of the note measured, which is accounted for by the attachment of the rib to the rest of the chest. In the studies by Sapira et al. [20] and Sullivan et al. [21] it is also shown that a lighter percussion stroke resulted in over estimation of the liver size, which is contrary to what is expected by Topographic Percussion Theory – as one moves to the sides of the liver, the liver tissue lies deeper and a harder percussion stroke should therefore penetrate deeper. This results is explained by McGee [8] as the Cage Resonance Theory – as one percuss the liver near where it touches the body wall, the heavy stroke is sufficiently strong to generate distant resonant notes and the light not note is not; as one moves away from this point, the light stroke generates a more resonant note – yet the author believes that McGee [8] forgot that resonance is not associated with the liver and that the liver span should then have been underestimated.

Topographic Percussion Theory has its origins with Pierre Piorry who believed that Topographic Percussion was possible because underlying organs absorbed sound waves – if the organ was in contact with the skin its absorption of the sound would cause the resulting sound to be dull; should the organ and surface be separated by lung tissue no absorption would occur and the sound would be resonant. Piorry's theory was evolved by Skoda, a mathematician and physicist, who believed that the percussion tone was dependant on the amount of air that separates the organ from the surface. Subsequent studies by Weil quantified the depth that an organ could be to be detectable as 6 cm – 2 cm of body wall and 4 cm of lung – which was subsequently adopted as the “several centimetre rule” that is still used in text books today.

2.1.3 Criticism and Limitations

Except for the disagreement about the factors that cause the percussion note, other factors that led to the criticism of percussion include the discrepancies about the technique of percussion and the performance of percussion.

The strength of the percussion stroke was believed to have to be as strong as possible which actually led to the development of indirect percussion – to alleviate the pain caused when percussing the abdomen of a critically ill patient. Subsequent studies came to the conclusion that the strength of the stroke made no difference in the results obtained. In two studies [20] & [21], however, it was found that the strength of the stroke did have an effect on estimated dimension of the liver span.

In a study by Burger et al. [22], the researchers discovered a great variability in various parameters concerning the technique - the velocity of the plexor (ranging between 0.45 and 2.2 m/s), the duration of contact between the plexor and pleximeter (ranging between 0.06 and 0.18 s) after the stroke and the contact pressure of the pleximeter (ranging between 0.9 and 3.2 kg). These variances could account for inter-observer discrepancies found in the various evaluations cited by McGee [8].

According to a modern bedside diagnostic textbook [5] the plexor finger should be withdrawn immediately after the percussion stroke, but in a study by Coleman [23] in 1939, clinicians could not distinguish notes created by rapid withdrawal of the plexor from notes by resting the plexor gently on the pleximeter after the stroke.

On the argument of whether direct or indirect percussion performs the best, Sapira et al. [20] found that the direct technique gave more accurate results for some cases, yet no other investigations that considered direct percussion could be found.

There have been many studies that compare the results of Topographic Percussion for the estimation of liver size to gold standards [20], [21], [23], [24]. The golden standard in these studies are either only scintiscan [25] (a radio nucleotide scan) or both ultrasound and scintiscan [21], [26] & [24] or ultrasound and x-ray [20]. Sullivan et al. [21] and Peternel et al. [25] found poor correlation between ground truth and the percussion finding of the span on the mid-clavicular line (vertical) with mean errors of >2 cm and 3.5 cm respectively, over ranges of 0 – 9.5 cm and 0.5 and 11.5 cm. On the other hand, Sapira et al. [20] found a reasonably high correlation coefficient of $r=0.53-0.61$ (the results are cited as a correlation coefficient). The study by Leevy and Naftalis [26] calculated a very high correlation of $r=9$, but these results have been disregarded by McGee [8] due to the fact that investigators (percussors) were allowed to repeat their evaluation after seeing the result on the scintiscan.

Despite the criticism, Skrainka et al. [24] presented a study aimed at evaluating the reliability of percussion (and other bedside diagnostics that form part of the clinical investigation). They found that clinical estimation (use of percussion and palpation) of the liver span should definitely be considered a valid technique.

2.1.4 Emulation of Percussion

During the previous century there have been various studies that attempted to bring technology and percussion together. This section discusses these studies and point out some ideas that are relevant to this study. These references are also cited in Chapter 3 and Chapter 4 where they influenced design choices.

In a study by Murray and Neilson [17] & [18], an attempt is made to characterize percussion sounds and subsequently classify them. This study was documented in 1975 and already showed promise with limited computing capability. By comparing the frequency spectrum of percussion notes they found that the energy of the percussion notes was below 600 Hz, that the frequencies at which notes peak differs between sounds (resonance, tympany and dullness) and that the total energy in each differs. To eliminate the influence of noise multiple acquisitions were taken and an average signal used for analysis. This information provides guidelines as to how data acquisition equipment should be used (acquisition bandwidth).

A subsequent study by Murray and Neilson [19] made use of accelerometers, tied to the pleximeter fingers, to compare the acceleration of the pleximeter finger to percussion sounds during the act of percussion. The reason for this investigation was that some physicians believe that the tactile sensations felt by the base hand contribute additional information to the percussion note itself. The researchers found strong correlation between the percussion note and the acceleration signal indicating that no additional information is presented. The acceleration signal and sound waveform discrepancy could be explained by the fact that the sound does not originate from a point source but a surface, therefore generating plane waves. This information is helpful as it shows that accelerometer could also be considered to acquire a waveform of the percussion sound.

Sandrin et al. [27] and Ophir et al. [28] & [29] introduced the concept of elastography – where a shear wave is introduced on the body wall and its propagation through the liver tissue is recorded by way of ultrasound. This technique uses an electromagnetic vibratory shaker to generate the shear wave and its propagation through the liver tissue is fundamental in the success of the technique.

Bohadan and Kraman [30] attempted to evaluate the consistency of the percussion stroke during Auscultatory percussion by using the strokes of a mechanical thumper as a standard. It found that the strokes by percussors were similar to that of the thumper. The thumper operated with a rubber tipped cylinder connected to a solenoid valve which, in turn, is connected to a source of compressed air.

Royston et al. [31] assessed the use of abdominal percussion for pneumoperitoneum (when free air is present in the abdominal cavity) detection. They performed their tests on dogs with the use of a modal hammer and a steel bar and recorded the signal using an air-coupled omni-directional microphone. In this study, and subsequent studies by Mansy et al. [17] & [18], an averaging technique was used to acquire a signal with reduced noise. They also mention specific features – *attack rate*, *decay rate* and *dominant frequency* – that were used to predict the volume of air present.

A short paper concerning the types of sensors that could be used to measure body surface vibration [32], evaluated the use of air coupled microphones, contact piezoelectric elements, accelerometers, an electronic stethoscope, implanted microphones and an acoustic stethoscope. Results were compared to that of a Laser Doppler vibrometer and it was found that all devices performed satisfactorily and the air-coupled microphone was the most sensitive to ambient noise. The electronic stethoscope performed the best as it is designed to detect a very specific signal.

2.2 The Liver

The goal of this project is to use the emulation of percussion to detect abnormalities of the liver. The liver as organ is of utmost importance to the normal functioning of the human body as it performs an array of functions; therefore it is susceptible to various diseases. Particular attention is given to diseases that can and should be detectable by percussion as they influence the physical properties of the liver. Further attention is given to modern methods for detecting abnormalities of the liver, and practical problems with these methods are pointed out as a motivation for a different solution.

2.2.1 Characteristics

According to Hawker [33] and Encyclopaedia Britannica [34], the adult human liver normally weighs between 1400 and 1600g and is a soft, pinkish-brown "boomerang shaped" organ. It is the second largest organ (the largest organ being the skin) and the largest gland within the human body. It is positioned immediately under the diaphragm on the right side of the upper abdomen extending from the right fifth intercostal space (space between ribs) in the mid-clavicular line down to the right costal margin. The liver lies on the right of the stomach and makes a kind of bed for the gallbladder (which stores bile). Figure 2.3 shows the normal anatomical position of the liver in the adult male (as extracted from Netter [35]).

The liver is supplied by two main blood vessels on its right lobe: the hepatic artery and the portal vein which enter the liver at the porta. The hepatic artery normally comes off the celiac artery, the first major branch of the abdominal aorta. The portal vein brings venous blood from the spleen, pancreas, and small intestines, so that the liver can process the nutrients and by-products of food digestion. The hepatic veins drain posteriorly directly into the inferior vena cava.

The functions of the liver [33] & [36] include bile metabolism of energy, lactate, protein, bile acids, bilirubin, lipids, drugs and toxins; and synthesis of proteins and bile; synthesis of Kupffer cells which are critical in defence of the human body against infection; and regulation of total blood volume (the liver stores 15 % of the total blood volume at any given time). The functions of bile include digestion, recirculation of bile acids, excretion of endogenous compounds like cholesterol, bilirubin and aged proteins, excretion of foreign compounds like drugs, environmental toxins and heavy metals and mucosal immunity by secretion of immunoglobulin A (an antibody).

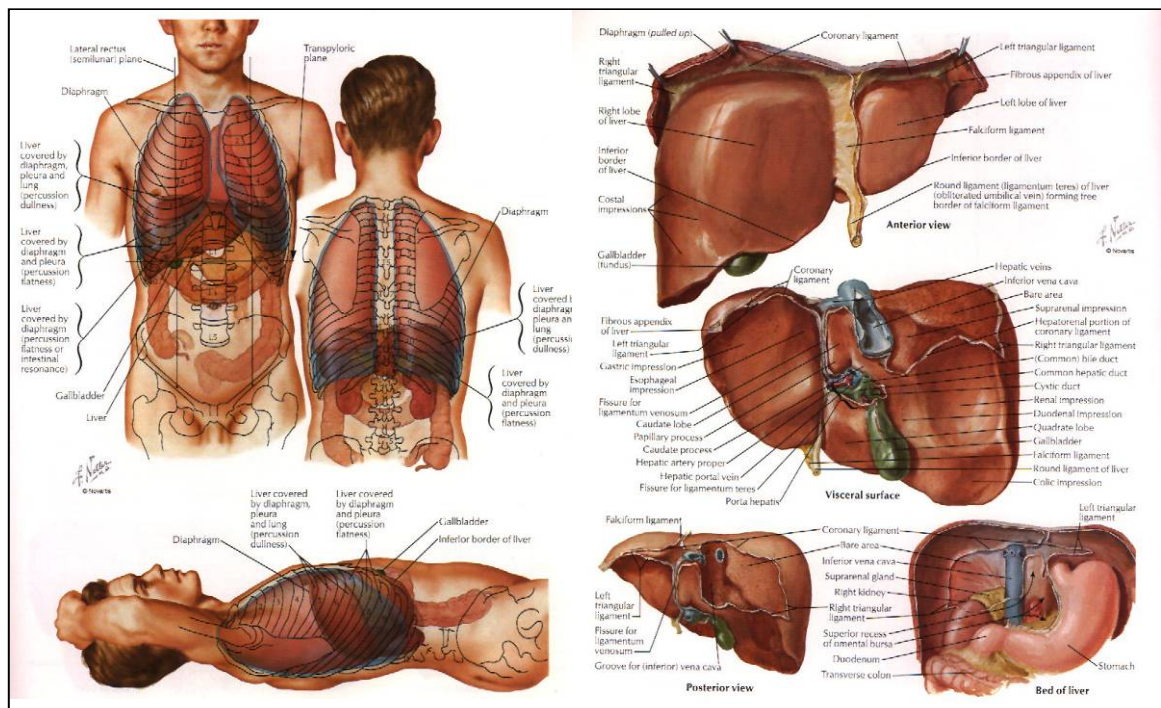


Figure 2.3: Anatomical views and position of the liver [35]

It is apparent that the liver is a critical organ in the human body. Any damage to the liver will negatively impact on the body as a whole and complete impairment of the liver's functions, without intervention, could lead to death.

2.2.2 Diseases of the Liver

There are various diseases or actions (like drug abuse) that impact on the liver and, depending on the liver function being influenced or liver cells being affected, these diseases will manifest itself as a symptomatic disease of the liver or a symptom directly caused by the liver's impairment (like jaundice).

Most of the diseases leading to symptomatic diseases of the liver are only detectable by liver function tests (LFT) or liver biopsies, yet the symptomatic diseases are differentiable/detectable by traditional screening techniques like percussion, ultrasonography and CT-scanning.

The most common symptomatic diseases are: (i) Hepatitis, (ii) Cirrhosis, (iii) Cholestasis, (iv) Hepatic Steatosis, (v) Cancer of the liver (hepatic tumours), and (vi) Jaundice. This section will briefly cover the description and symptoms of each.

(i) **Hepatitis** [37] & [38]

The term hepatitis generally refers to inflammation of the liver by inflammation of hepatocytes, but can also imply general cellular alterations including necrosis (death). Chronic hepatitis eventually leads to cirrhosis or necrosis of hepatocytes.

(ii) **Cirrhosis** [37] & [39]

Cirrhosis is a process characterised by fibrous alteration of normal liver architecture into structurally abnormal nodules of liver cells surrounded by fibrosis (the formation of excess fibrous connective tissue as a reparative process to injured liver cells). The symptoms of cirrhosis are very wide and unspecific and mostly associated with impaired liver function (jaundice, hypogonadism, muscle wasting). Untreated, cirrhosis will lead to eventual failure or cancer of the liver.

(iii) **Cholestasis** [37], [40] & [41]

Cholestasis is any condition in which bile flow is blocked. The blockage of bile flow could lead to inflammation of the liver and hepatocytes which in turn could lead to cirrhosis.

(iv) **Hepatic Steatosis** [37], [42] & [43]

Lipids (fat) represent approximately 5% of the normal liver. In the case of Hepatic Steatosis (also Fatty Liver or steatorrhoeic hepatitis) small vacuoles of lipid (liposomes) accumulate in hepatocytes. As the size of the vacuoles increase, the cell nuclei are pushed to the periphery and large fat vesicles are formed. Large vacuoles eventually coalesce to form fatty cysts which represent irreversible lesions.

(v) **Cancer of the liver (hepatic tumours)** [37], [42] [44] & [45]

Hepatic tumours can be benign or malignant (cancerous) and present as either hepatomegaly (enlarged liver), abdominal pain, jaundice or impaired liver function.

(vi) **Jaundice** [42], [44] & [46]

Jaundice in general refers to a yellowing of the skin and the mucous membranes caused by increased levels of bilirubin in the human body. Bilirubin is mainly formed during the breakdown of haeme (which forms part of the haemoglobin in red blood cells) in the spleen and Kupffer cells in the liver and is secreted with bile. Jaundice can be divided into three classes associated with the cause: (i) pre-hepatic (or haemolytic) caused increased rate of haemolysis (breakdown of red blood cells); (ii) hepatic caused by liver cell impairment; and (iii) post-hepatic caused by an interruption to the drainage of bile. Jaundice is one of the first symptoms of impaired liver function and is important in initial screening. It is not a disease of the liver as such, but a symptom primarily caused by liver diseases. A measurement of the level of jaundice in a patient can assist in the diagnosis of that patient.

These six symptomatic diseases are important to this study for two reasons:

- i. All except Jaundice represent as macroscopic changes to the liver which should be detectable as anomalies during percussion, palpation or auscultation. In the case of hepatitis, hepatic steatosis, cholestasis and possibly tumours, the liver will increase in size. In the case of cirrhosis the liver could either increase or decrease in size depending on the prognosis. The elasticity of the liver could change in all cases, and the consistency mostly changes in tumour cases. Jaundice is a symptom that can indicate any of the other conditions.
- ii. These conditions are, in various cases and as stated, caused by other diseases. By the time these conditions present themselves the prognosis is severe and it is imperative that the patient presenting these symptoms be presented with the appropriate medical care. As is often the case in rural areas, the appropriate screening tools are not available and the bedside diagnostic methods (history and physical exam) need to be well understood to detect these symptoms.

2.2.3 Detection of Liver Disease

Modern methods for the detection of liver diseases can be divided into four categories: (i) Clinical Evaluation (History and Physical Exam), (ii) Laboratory Evaluation, (iii) Radiological Imaging and (iv) Morphologic Evaluation (Liver Biopsy). This section will briefly mention the relevance of each of these methods.

(i) Clinical Evaluation [37]

The clinical evaluation is usually performed by the general practitioner (GP) and involves taking a complete history of the patient and a physical examination. The GP relies on his own experience and knowledge, and then makes an assessment. The general exam includes percussion, palpation, examination of the body and taking the patient history. The physical exam is rather inexpensive in comparison to any other screening technique, but is highly dependent on the GP's experience and knowledge.

(ii) Laboratory Evaluation [33] & [37]

Laboratory evaluation refers to Liver Function Tests (LFTs) and concerns a wide variety of blood tests. The standard LFT measures the levels of serum enzymes and other markers. These tests are generally inexpensive once the equipment is attained, but still require skilled operators and large initial investment.

(iii) Radiological Imaging [33], [37], [47] - [48]

Radiological imaging techniques include X-ray, Ultrasonography, Computed (Axial) Tomography (CT and CAT), Medical Resonance Imaging (MRI), Elastography and Radionuclide Imaging (Single Photon Emission Computed Tomography (SPECT) and Positron Emission Tomography (PET)). These techniques all involve some form of radiation, but only X-ray, SPECT, PET, CT and CAT involve dangerous ionizing radiation, however, at safe exposure levels (for single doses/assessments).

Of these methods, X-ray, SPECT, PET, CT, CAT and MRI are all relatively expensive and require special facilities to perform – these are therefore irrelevant to the study as they are ill-suited for rural areas. Even so, they are all adaptable to scan soft tissue including the liver.

Ultrasonography (US) can detect the size and shape of the liver, surface irregularities, alterations in hepatic parenchyma (fatty, necrotic, fibrous) and space occupying lesions within the liver such as tumors, cysts, abscesses, calcifications and stones. US is relatively inexpensive, mobile, produces images in real-time, causes minimal discomfort for the patient and has no long-term side effects. Although ultrasound might have a lot of commonality with percussion on face value, it must be pointed out that sound at the frequencies used by ultrasound are collimated (they do not diverge) whereas the low frequency waves of percussion are not. For this reason the principles of ultrasound cannot be extrapolated to percussion theory.

Elastography [27] - [29] is a relatively new technique that involves producing elastograms (images of the elasticity/stiffness) of the tissue under test with the aid of either MRI or ultrasound. The tissue is excited by a transducer that produces waves at 50 to 1000 Hz. The acquired data is processed to determine Young's modulus or Poisson ration of the underlying tissue. The technique of elastography is particularly valuable in the diagnosis of cirrhosis (because the liver becomes harder/stiffer), liver tumours, hepatic steatosis (liver becomes softer/more elastic) and cholestasis (blockage of tracts causing some softer regions). The cost of elastography is comparable to US. It still requires the skills of a technician to operate, is too costly to be used in rural areas and can only diagnose a few pathologies of the liver.

(iv) **Morphologic Evaluation** [37] & [49]

Morphologic evaluation mainly refers to taking a biopsy of the liver, but specimens can also be obtained surgically. A liver biopsy involves the removal of cells or tissue by means of percutaneous needle or transvenous catheter; which is then histologically (chemically or microscopically) analysed. The main advantage of morphologic evaluation is that it can give a very clear picture of the pathological state of the liver when the parenchymal cells of the liver are involved. It does, however, fail to diagnose many pathologies that are not directly caused by hepatocytes. The greatest disadvantages involved with a biopsy are that it can be very painful for the patient, require highly skilled and educated pathologists and, inherently, can be very expensive.

2.3 Conclusions

Although medical percussion is explained in medical textbooks, the underlying physics is not well understood and ambiguous. It mainly uses some variant of the hand as a striking tool and in most cases the other as a mediate and the resulting note is used to diagnose or identify underlying pathology or organs. Percussion is applied to assess various regions of the body, one area being the liver – this application of the technique is explained in Section 2.1.1.

Percussion has evolved over the past four centuries and has been adopted by most physicians and medical textbooks. There are, however, various versions of the technique with most of them still in use today and even studies focused on establishing which version is superior could not identify a clear winner.

Over the previous century there have been various studies that have both discredited and affirmed the technique in all areas where it is applied. These studies evaluate inter-observer agreement on percussion notes, the correct technique and the results obtained. The results between studies are conflicting, which makes it difficult to make clear conclusions.

There have been numerous modern studies not aimed at emulating percussion, but inadvertently copying the actions of the hand in order to obtain a more repeatable stroke/result. Other studies were aimed at evaluating the percussion note by computer – these used crude signal analysis techniques and little computing power to successfully distinguish between percussion notes.

In Section 2.2 the importance of the liver as an organ in the human body is underlined. Early detection of liver problems is of utmost importance and the *correct application* of percussion could succeed in identifying these.

Techniques other than percussion, that evaluate the liver's condition are superior, but more expensive and less accessible (depending on where a patient is) than percussion. Unfortunately, skill and experience is of utmost importance to the successful application of percussion. Lack of these attributes can lead to misdiagnosis.

It is possible to copy the technique of percussion by some mechanical device, but there is no scientific basis to draw a parallel between percussion and artificial percussion (whatever that may be) in order to validate results. Although signal traces have been depicted for percussion sounds recorded with a microphone, other sensors might perform better, but record different signal traces (acceleration vs. sound pressure) – whether the same separability would be attainable is questionable. It may also be possible that such an investigation may present with far superior results.

Chapter 3 System Development

Hardware and software needs to be acquired, or designed and built, and finally integrated to perform the primary tasks performed by the physician when performing percussion – to percuss the patient’s body wall (actuate) and to listen to the result (sense). The identification of signals would require experimentation and evaluation (like the training of the physician’s ear) with the systems that perform these primary tasks – the experimentation and analyses are discussed in Chapter 4 and Chapter 5. This chapter deals with the development of a system that can actuate the body wall and record reactionary signals.

3.1 Conceptual Design

Except for the physical devices needed to perform the primary tasks, one would need to control these devices, condition the signals and store the acquired signals for later analysis. The following components were identified as necessary and their individual designs and/or selections are discussed in subsequent sections:

- Actuator (a mechanical device that could deliver the percussion stroke)
- Sensors to measure changes generated after the percussion stroke.
- Signal Conditioning System to remove noise and unnecessary components as well as provide power to sensors.
- Data Acquisition (DAQ) System to digitize sensor data for computer analysis.
- Computer to collect data and control the actuator
- Rigid frame (after some initial investigation it became apparent the patient’s movement manipulates the recorded data as well as the actuator input) to attach the actuator to the bed.

Some of these subsystem requirements could be met by calibrated laboratory equipment, but laboratory equipment is often ill-suited for testing on human subjects and extensive development may have been required if such equipment was chosen. Furthermore, unless equipment was not being used for other laboratory work, it would need to be acquired which would be expensive especially given budgetary constraints. The custom designed sensors that were employed in the end, may have been better adapted to human testing and far cheaper, but their performance were found to be lacking, unpredictable and could not be validated. The approach of developing custom sensors should perhaps have been preceded by a proof of concept with standard, tested and calibrated equipment and only once the concept was proven. Unfortunately, the author continued on this path of customised sensors based on the assumption that they would work and are working and also because of a lack of knowledge of standardised laboratory equipment. More attention will be given to these problems in subsequent sections.

3.2 Actuator

The actuator should ideally be a robotic hand that can move like that of a physician during the physical exam, only more repeatable, but this approach would be impractical as the control of such a robotic hand would be a project in itself. Realistically the actuator should be some device that would allow an investigator to provide a repeatable mechanical thump to the body of a patient without hurting the patient. Additionally, having the ability to be actuated with different signals would be useful.

Two options were investigated – a modal hammer and an electromagnetic vibrational shaker.

The modal hammer is a standard modal analysis tool used in Mechanical Engineering for Vibrational Analysis. Hammer models mostly contain a load cell in the tip of the hammer and the input force signal is combined with output sensor signals to estimate a frequency response function (FRF) – this effectively normalises the output to the input. The use of a modal hammer for human testing was not researched, but tested and the tests revealed that very light thumps against the ribs caused severe discomfort to volunteers. The main researcher had minimal exposure to modal analysis and was not aware of FRF estimation techniques at this time and rejected the use of the modal hammer based on the lack of repeatability and discomfort to the volunteer. The lack of bandwidth in the input force signal used for the FRF estimate may have proven to be an additional dismissing factor.

The electromagnetic vibrational shaker is another standard Vibrational Analysis tool. The devices work on different principles, but in essence have some piston that can be moved by modulating the current through electromagnetic coils. The shaker is therefore versatile in that it can be driven with various analogue signals – a pulse to simulate the percussion stroke, a frequency ramp that could help determine a frequency response, and white noise to estimate an impulse response (time domain version of a complex FRF). A vibrational shaker is used in 1-D transient elastography, a technique used to determine the elasticity of, amongst others, the liver [27]. The technique uses ultrasound to image the wave as it passes through the liver, proving the signal penetrates at least as deep as the liver – this result further motivates the decision to use the shaker as actuator. The problems that one would have with the vibrational shaker is that it has a short offset position (<25 mm) which means that one would have to place the head of the shaker against the skin of the patient; the shaker does not have a standard load cell, which means that an additional force sensor would need to be developed; and the shaker models that were accessible had rubber diaphragms which generated a sound at the same frequency as the excited frequency – this would render sound sensors useless. The advantages of the shaker would be the ability to modulate the piston movement (repeatable); the shakers are mountable (if one should need to connect it to a frame); a head could be designed to make the input to the patient more comfortable; and the modulation capabilities could make the FRF estimation more dependable.

These considerations led to the decision to make use of the electromagnetic vibrational shaker. The mechanical shaker that was used in this project is the V201-PA25E from LDS Group [50]. An excerpt of the datasheet can be found in Appendix E.1. A photo and drawing of the shaker is shown in Figure 3.1.



Figure 3.1: Vibrational shaker as implemented

The use of the vibrational shaker necessitated additional custom design work which included a customised tip and stem and a custom force sensing element.

The coupling element that transfers the energy from the electromagnetic shaker is designed in such a way that the weight and inertia of the shaker is decoupled from the target (this is the function of the narrow stem) and to distribute the energy evenly over the contact area. In addition to distributing the energy, the spherical shape of the tip limits the effect of non-perpendicular input (offset angle) which is a factor when actuating a human body which is covered by non-regular surfaces. Other exciter head shapes were not considered.

The sensing of input force requires some force sensor. Although a load cell was not considered, it would have required some custom design to mount it between the shaker and its tip and may not have recorded the exact input force. Mounting the load cell to measure the load delivered to the human subject may also have caused discomfort for the patient. A force sensing resistor (FSR), the Model 400 FSR from Interlink Electronics [51], was therefore placed in the centre of the tip's contact area. Although the FSR may not be the ideal sensor for measuring input force, it was easily acquired at the time of design and very easy to implement as it is made from a thin film and does not interfere with the contact between the shaker tip and the subject's skin. An FSR changes resistance inverse proportionally to the force applied and is connected in a voltage divider configuration with 10 k Ω resistor. The FSR and voltage divider configuration are shown in Figure 3.2.

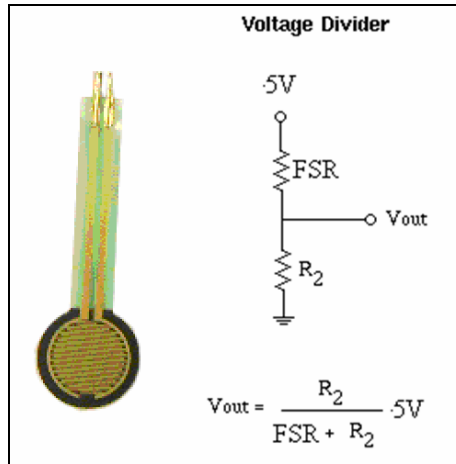


Figure 3.2: FSR and Voltage Divider Configuration

The input force vs. output voltage relationship is given by equation 1 below. This relationship is derived in Appendix A.4.

$$F = 0.584 \cdot \left(\frac{5}{V_{OUT}} - 1 \right)^3 \text{ N} \quad (1)$$

The exact relationship between input force and measured voltage was not verified by testing, but equation 1 was used to calculate set points so that two tests could be performed using the same pre-tension (assuming zero drift). An example of an FSR signal recorded during experimentation on human subjects is shown in Figure 3.3. The top left plot of the figure shows the signal where a sequence of pulses were delivered to the human subject and the top right and bottom plots are depiction of the same signal, but zoomed in sections. The FSR voltage signals (shown in blue) in general seemed unnatural with components that look like discontinuities. The FSR force signals (shown in red) were obtained by converting the voltage signal to force (N) using equation 1, and seem more natural, but still do have some saturation at 0 N which can be attributed to the actuator tip not making any contact with the skin for a short period of time. The slowly changing offset can be attributed to the subtle movement of the volunteer under test. Although it is not clear here, some signals still showed discontinuities or sudden jumps in the offset voltage which could not be attributed to environmental factors.

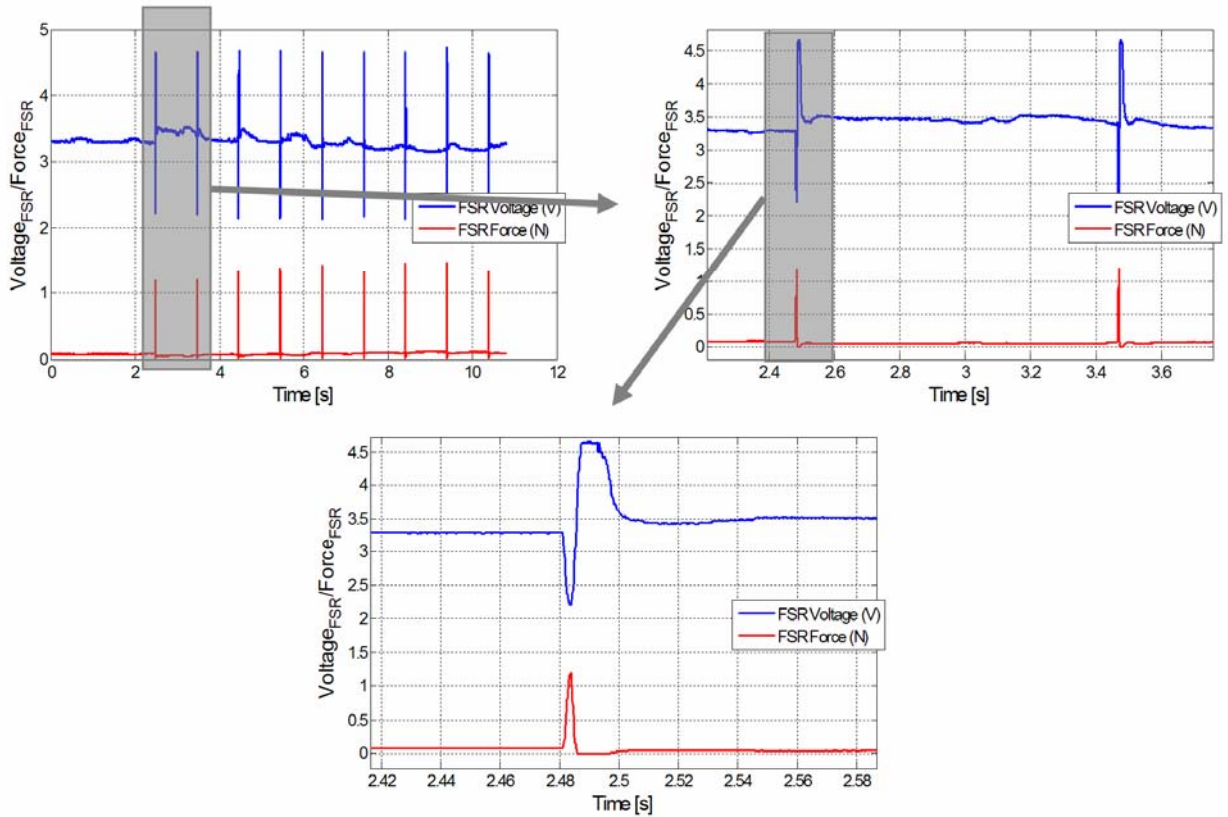


Figure 3.3: Force Sensing Resistor signals

The electromagnetic shaker is also driven from a single channel sound amplifier which amplifies the sound output from a PC (or any other sound generating device). The sound amplifier was borrowed and is developed primarily for this function and is not an off-the-shelf device (therefore, there is no datasheet). The amplifier is AC coupled.

3.3 Sensors

The function of the sensor or sensors is to measure the reactionary signal of the human body wall. The physician relies on the sound generated by the body wall during percussion to decide whether the sound is liver dullness, but, according to McGee [8], it is unclear whether a physician practising percussion is making deductions based on what he hears or whether he includes what he feels (with his hands) in making deductions. In the discussion about the actuator it was pointed out that the electromagnetic shaker has a diaphragm which generates a sound representation of the input force signal and this sound is louder than any reactionary signal one could hear. This places an unfortunate constraint on the selection of the sensor as it cannot be an air coupled microphone and needs to be some sensor measuring the vibration of the skin.

The spectral content of the sound recorded by the physician's ear will not exactly match the spectral content of the vibrational signal measured on the skin as the transfer of vibration from skin to air is not ideal and other acoustic effects in a room may influence the sound heard. Be that as it may, the signal bandwidth should be similar and there should be distinguishing features in the surface vibration signal just as there are in the sound, although these may not be the same features. According to Murray and Neilson [17] & [18], the energy spectrum of the percussion note detectable above the abdomen is contained between 200 and 600 Hz, which means that whichever sensor is implemented needs to have a bandwidth of 600 Hz.

A study by Mansy et al. [32] gives some guidance as to how to select sensors for surface vibration measurements in humans. They compared the use of an electronic stethoscope, accelerometers, contact piezoelectric elements, an air coupled microphone, an acoustic stethoscope with built in microphone and a Laser Doppler vibrometer (LDV) by measuring vibrations on a silicon phantom with an internal speaker. They concluded that all of these sensors performed adequately as far as signal-to-noise is considered, yet the stethoscopes performed better for measurement of the desired signal (this evaluation assessed typical auscultation sounds which are related to the heart beat and lung ventilation) and the air-coupled microphone was very sensitive to ambient noise. The LDV should be immune to ambient noise and was therefore considered as the golden standard to which the other devices were compared. In this selection of sensors only two were considered for this thesis: the electronic stethoscope and accelerometer. The LDV was not considered as none was accessible at the time.

An electronic stethoscope from GeoAxon was evaluated for the purposes of recording percussion sounds. It was found that it was too sensitive and saturated, therefore eliminating the option.

Both industrial standard piezoelectric accelerometers and smaller, more modern MEMS accelerometers were considered. Industrial piezoelectric accelerometers are heavy and expensive, but MEMS accelerometers are small in size and weight and relatively inexpensive, which led to the selection of the latter for this project.

MEMS accelerometers operate in various ways, but usually have some microstructure with a known weight loaded on some micro spring mechanism and some method to detect deflection of the structure. The measured signal is related to the earth's gravitational field and the actual acceleration of the device. The devices are usually manufactured in flat square or rectangular packages with acceleration axes being parallel to the flat surface unless it has a third axis, in which case it would be perpendicular to the package's surface. The third axis, labelled the z-axis, is of interest as the device can be mounted to a surface with this axis being perpendicular to this surface and therefore most sensitive to vibration if only one vertical movement is considered.

A formal literature investigation into the actual required sensitivity revealed no exact quantity as to what the typical acceleration of the skin under percussion is and it was decided to build the sensor and evaluate its performance. The ADXL330 from Analog Devices [52] was selected. The ADXL330 has a sensitivity of 300 mV/g (where g denotes gravitational acceleration), dynamic range of 3 V or 10 g and bandwidth of 1600 Hz in the x and y directions and 550 Hz in the z direction – the bandwidth of the z-axis is a bit below the 600 Hz requirement, yet it represents a roll-off of 20 dB per order of magnitude (5.5 kHz), therefore a dominant frequency below 600Hz should still be measurable.

An accelerometer PCB was designed as a substrate for the ADXL330. The design is detailed in Appendix A.3. The PCB with the ADXL330 is engulfed in an epoxy to protect the circuitry. A small clip is glued to the base of the unit and clips to an ECG sticker. The ECG sticker sticks to the patient's skin in the target position and the accelerometer unit clips into the sticker. The combination of ECG sticker and clip may have some energy transfer effects, but the exact transfer function between a vibrating source and MEMS accelerometer was not established. However, a frequency sweep with the ECG sticker attached to the actuator head revealed that the signal bandwidth remained in tact. It was assumed that the combined effect of the sticker and clip would be consistent throughout measurements, if there was any effect at all.




Figure 3.4: Image of Accelerometer Sensor Unit

In a preliminary investigation into the applicability of the ADXL330 as vibration sensor, the ECG sticker and sensor unit were placed on various places on the thorax and abdomen of a volunteer and both the electromagnetic vibrational shaker and modal hammer were used as thumper. Accelerometer signals were measured with an oscilloscope and it was found that the sensor could detect signals generated from distant locations (>20 cm) as well as having the necessary dynamic range to record strong signals (without saturation) in close proximity (<5 cm) to actuation. From this experiment it was concluded that the sensor should perform adequately as a body wall vibration detector.

It was decided to measure at fixed distances from the point of actuation and at the same height, therefore only requiring two accelerometers. An additional accelerometer unit was built and could have been used if there were certain points of interest to be measured, but there were no such instances.

Although preliminary investigations showed that the accelerometer produced normal looking signals, some of the signals recorded during human subject tests showed seemingly unnatural artefacts. These artefacts included saturation which could be ascribed to maximum acceleration limits being reached (see Figure 3.5), and high frequency glitches or sudden changes in the voltage signal which could not be explained (see Figure 3.6). The signal processing techniques used in Chapter 4 assumes that these artefacts would be rejected by matching models to the input-output data or by obtaining averages of multiple signals.

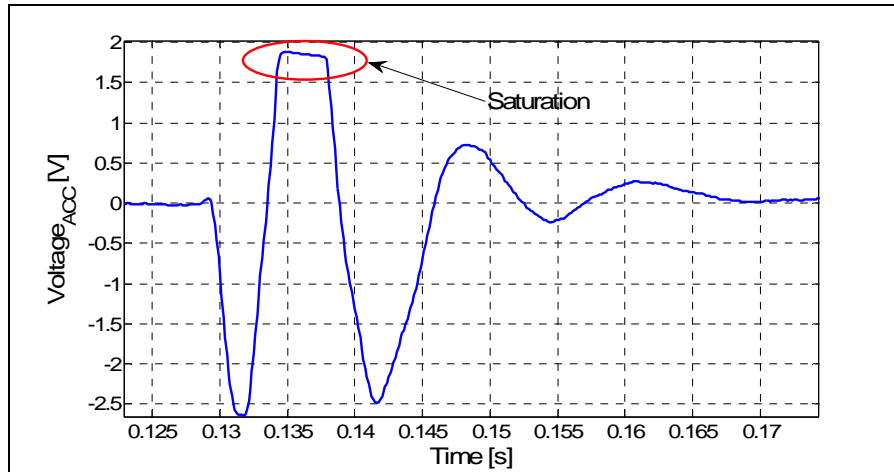


Figure 3.5: Accelerometer saturation

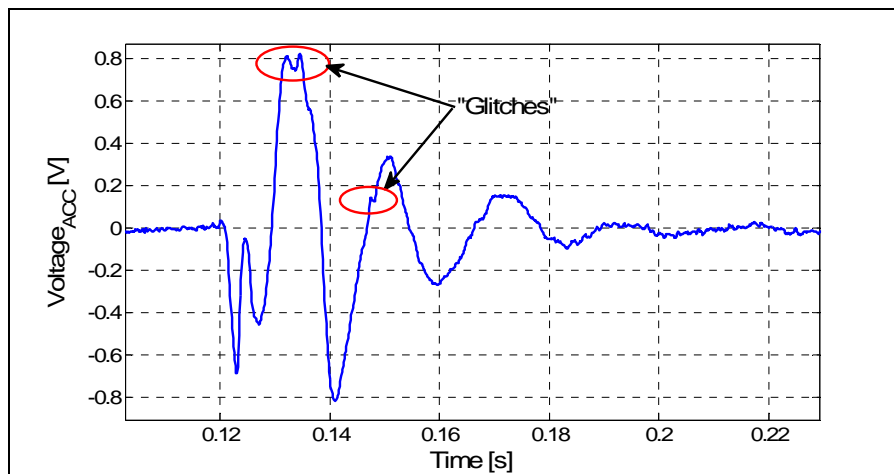


Figure 3.6: Accelerometer discontinuities

3.4 Data Acquisition System

The function of the data acquisition system is to digitize signals before sending it to a computer for storage and/or analysis. For this purpose an already available module, the Measurement Computing PMD-1608FS USB Analogue and Digital I/O Module, was used. The device has eight single-ended 16-bit analogue-to-digital conversion channels with software selectable input ranges of ± 10 V, ± 5 V, ± 2 V and ± 1 V and sample rates of 50k samples/s per channel, but 100k samples/s total. There are also eight digital input/output channels with built in pull-up resistors which are used to control the multiplexers of the Signal Conditioning System. The device uses an USB 2.0 port, comes with the necessary driver software and can be controlled through the Matlab Data Acquisition Toolbox (Matlab was used as the software development and evaluation environment). The user guide and specification is available from the manufacturer [53].

Figure 3.7 shows a picture of the PMD-1608FS USB and its system block diagram respectively, with the device pin-out.

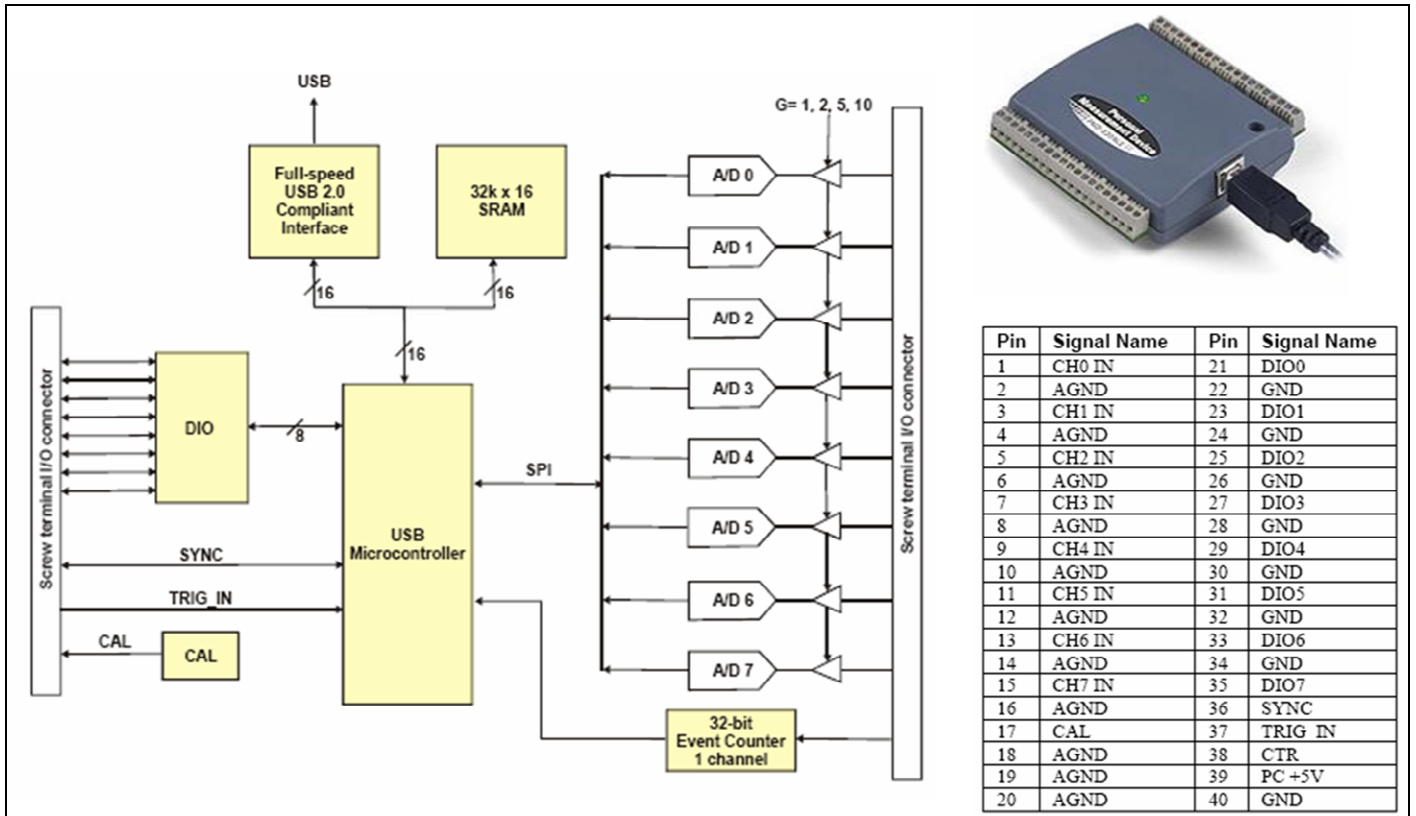


Figure 3.7: PMD 1608FS USB Block Diagram and pin-out [53]

The device was used for data acquisition and controlled from one Matlab m-file called *DataloggerGUI.m* discussed in Section 3.6. Analogue (analogue-to-digital conversion) channels 1 to 3 (CH1 in, CH2 in and CH3 in) were used for the three accelerometer signals from the Signal Conditioning and Routing PCB. The FSR signal was connected to analogue channel 0 (CH0 in) and the modulation input to the electromagnetic vibrational shaker was sampled through analogue channel 4 (CH4 in) in some experiments. Two digital outputs (DIO0 and DIO1) are used to control the multiplexing function on the Signal Conditioning and Router PCB.

3.5 Signal Conditioning System

The signal conditioning system's requirements were driven by the hardware choices in the previous sections. The accelerometer signals needed to be manipulated slightly to make better use of the data acquisition device's dynamic range. The data acquisition device also has limited acquisition channels which necessitates the ability to switch between signal sources. Different components also have different power requirements.

The ADXL330 MEMS Accelerometers require ground (GND) and 3.3 V power to operate and the sensor outputs contain a DC component – this is due to the offset acceleration caused by the earth’s gravitational field and the zero acceleration output being 1.65 V (halfway between the power rails). The sensors also have three outputs each; therefore, with three sensors one would already require nine analogue-to-digital conversion channels if each channel was of interest. In addition to this, the FSR sensor discussed in Section 3.2 and other possible sensors would also require analogue-to-digital channels.

The signal conditioning system has the following functions:

- Provide power to three ADXL330 MEMS Accelerometers.
- Multiplex the nine input signals to three outputs by either mapping the three z-axes or the x-, y- and z-axes of a single accelerometer.
- Remove the DC component from the accelerometer signal as it is not of interest (this will only give information as to the sensors orientation relative to the centre of the earth).
- Amplify the signals to make better use of the dynamic range in the analogue-to-digital conversion of the Data Acquisition System therefore improving signal to digitization noise ratios.

Two PCBs were developed to perform these functions. The first board is called the Signal Conditioning and Routing PCB and the second the Power Source PCB. The Signal Conditioning and Routing PCB uses multiplexing techniques to map input signals to output signals (as outlined above) determined by two control input lines. Mapped signals are high pass filtered to remove the DC component in the signal and then amplified by 10 dB. Connectors on the board provide connection for the accelerometer PCBs to transfer power and signals, power connectors and output connectors to transfer signals to the Data Acquisition System. The system block diagram is shown in Figure 3.8 and the design is detailed in Appendix A.1.

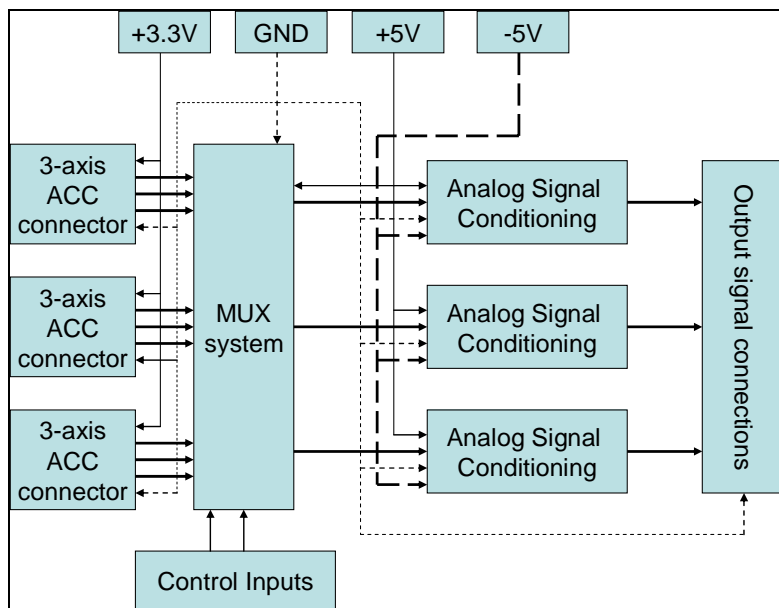


Figure 3.8: Signal Conditioning and Routing Block Diagram

The Power Source PCB converts 230 VAC_{RMS} to 3.3 V, +5 V and -5 V. The output voltages are used to power the first board and the accelerometers. The design of the Power Source PCB is detailed in Appendix A.2. The block diagram is shown in Figure 3.9.

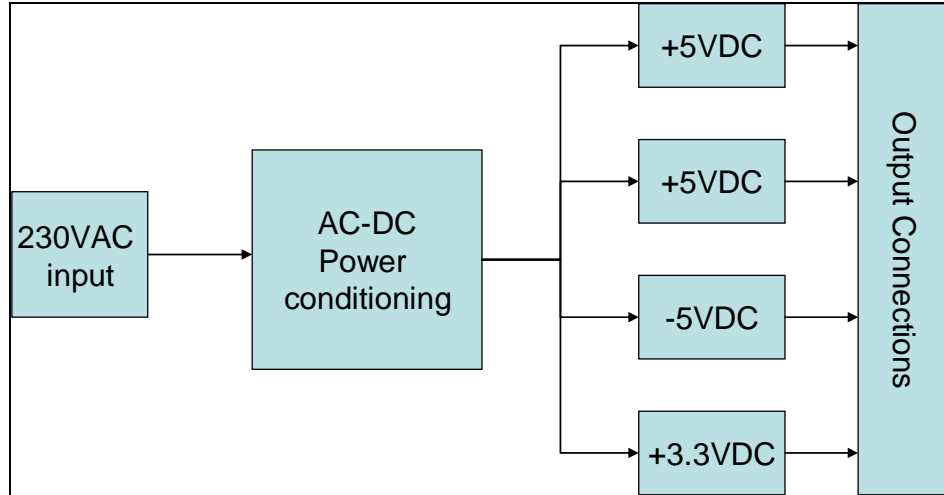


Figure 3.9: Power Source PCB system block diagram

3.6 Computer System

The functions of the ‘computer system’ are the following:

- Generate the sound signal for the actuator.
- Control the Signal Conditioning and Router PCB to return the correct signals.
- Record and store measured signals.
- Analyse recorded signals.

These functions are all controlled through software that was developed in and for the Matlab environment. The software can run on any computer running the Windows XP operating system with a USB 2.0 port and sound output. The software was not developed with distribution in mind; it was simply developed for use by the author.

The first three listed functions have to do with the data acquisition phase of the project and are performed by one program or m-file called *DataloggerGUI.m*, which has a graphical user interface (GUI) built from a base generated with the Matlab development tool called GUIDE (Graphical User Interface Development Environment). An overview of the software flow is shown in Figure 3.10. When the application starts it initialises values, updates logs and dates, initialises the DAQ device (if available) and generates the user interface. Hereafter the program is in a wait state where it waits for the user to change values on the panel, start an acquisition, save data or terminate the program. The click of a button, changing of a value, or entering of text calls functions or *callbacks* in the program that will perform an action related to that button. Underlying functions or *callbacks* can be called from others and not all of them are accessible by the user (e.g. plotting data).

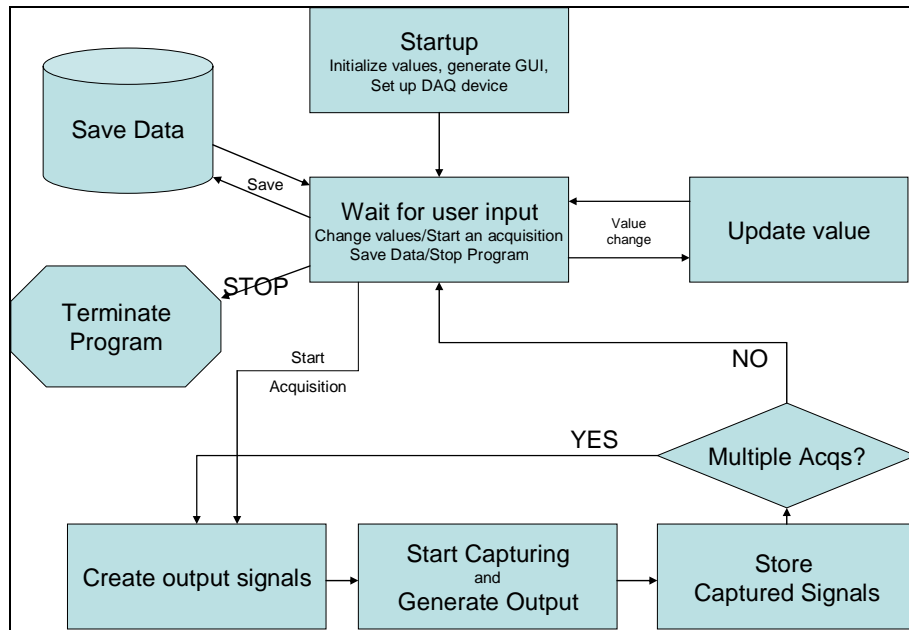


Figure 3.10: *DataloggerGUI.m* general software flow diagram

Should the user change a value, the underlying *callback* will update the value and return control to the interface immediately. If more than one control is clicked while an update is performed, that action is stored in a buffer and executed as soon as possible.

The software was developed to allow the user to perform tests with one or all of three actuation signals (white noise, chirp or pulse). When an acquisition sequence is started, the software generates the output signal internally and then starts capturing data and outputting the signal simultaneously. Hereafter the data is stored in internal variables and presented on a graph.

Once data has been collected, the user can choose to overwrite the data with a new single acquisition or sequence thereof; the user can save the stored variables in a .mat standard Matlab data file or the user can choose to end the program in which case data will be lost.

Additional functionality includes loading previously recorded data for display in the window. This is not included in the software flow as it is a redundant function as far as the project is concerned as the data analysis functions generate various other representations.

An image of the user interface is shown in Figure 3.11

Data analysis software or programs were also developed as m-files in the Matlab environment. These functions do not include GUIs as they were not developed as tools, but simply used for evaluation. These programs are discussed within the sections where they were employed.

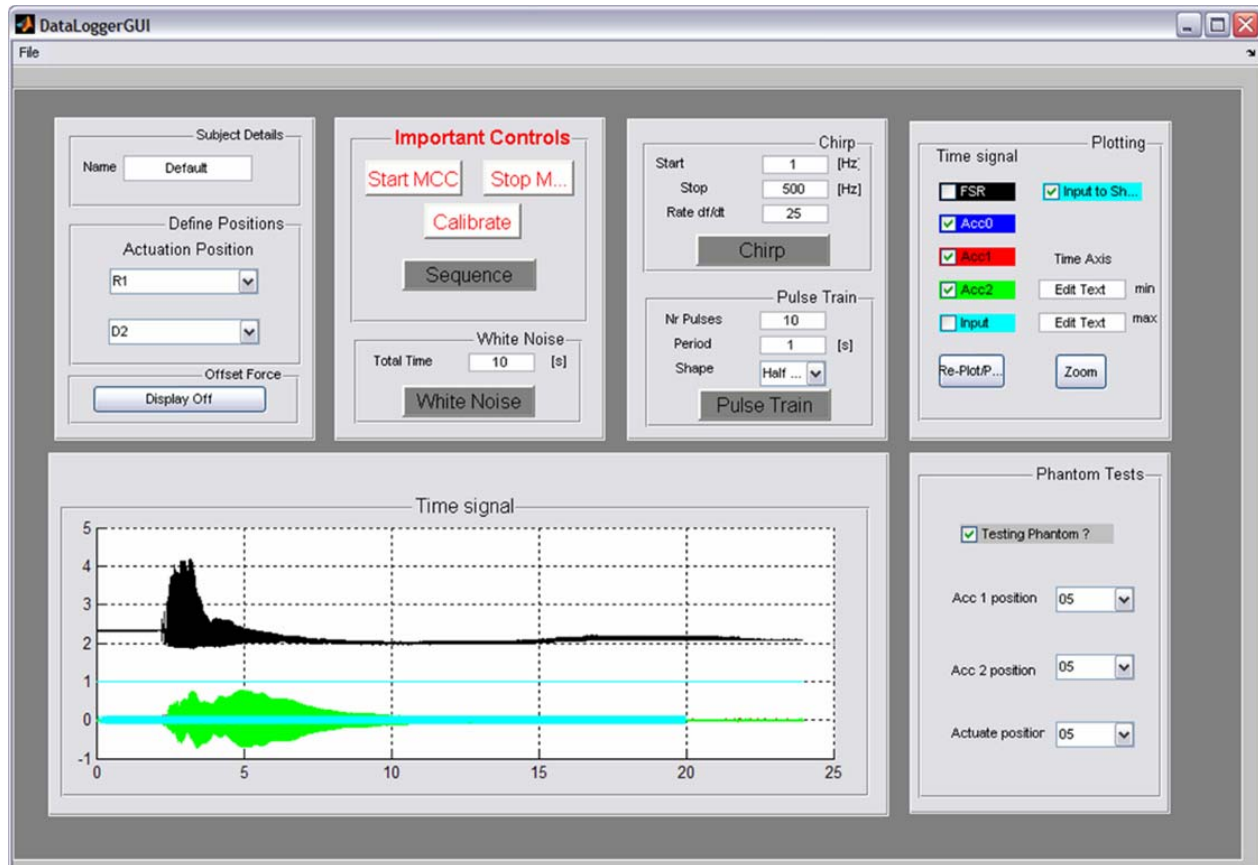


Figure 3.11: *DataloggerGUI.m* graphical user interface

3.7 Frame

If a patient is percussed with a mechanical device while he/she is standing up, the patient will tend to sway either simply periodically or in reaction to the percussion stroke (this was confirmed with primary investigations where the vibrational shaker was simply clamped to a bench and the patient asked to stand against it). If a patient is lying on his/her back, swaying is not possible, but if an operator is holding the device (especially if it is heavy) he/she will also tend to sway. To prevent patient or operator sway it was decided to design and construct a frame from rigid steel.

The designed frame must adhere to the following requirements:

- It must be attachable to a bed.
- The position of the actuator should be lockable for two reasons:
 - So that the actuator stays in place for repeat sequences.
 - If the actuator or parts of the frame slides out of place it could injure a patient.
- It must allow an operator to move the actuator to any position on the patient's body.

Of the concepts that were considered the one represented in Figure 3.12 was the one that was finally decided upon. The electromagnetic shaker (“shaker”) is attached to a base with a pipe extrusion (“shaker base”). The pipe extrusion fits into another curved structure (“curve slider”) with a locking screw to lock the shaker base into position. The curved base is curved so that it can slide over curved pipes (“curved track”) which fit into side supports structures (“straight slider”). These side supports are partially made from pipes so that they can slide over the straight track, therefore the name “straight slider”, and the straight track end in bed clamps. The straight and curved tracks have lock screws to keep the frame locked in position whilst a measurement is being done.




Figure 3.12: Frame design and measurement setup

Detailed design documentation can be found in Appendix C.1.

3.8 Conclusions

To perform emulation of percussion various devices and frameworks are required. A concept was developed by referencing literature and taking some practical considerations into account and this concept led to individual component requirements and finally component selection.

In the light of the outcome of this project, critical evaluation of the performances of sensors is required and this was done in Sections 3.2 and 3.3. It is shown that the FSR and accelerometer did perhaps not perform ideally, and should have been compared against standard sensors. The artefacts could perhaps be due to weak performance or the signal conditioning system, but that they truly reflect what is happening can also not be refuted. The evidence is inconclusive, but selection of trusted, calibrated laboratory equipment would have eliminated this problem. In Chapter 4 model matching techniques are used to define systems that best describe input-output data and it was hoped that these fitting techniques would reject signal artefacts – this theory is also verified using artificial data.

The choice of the electromagnetic vibratory shaker allows for some versatility over the modal hammer, but impacts on the sensor selection in that an air coupled microphone cannot be used to measure reactionary signals. The rejection of the modal hammer was perhaps premature and could have allowed the use of an air coupled microphone which is closer to percussion.

In Chapter 4, an experiment employing this system is developed and performed on the assumption that the method and devices are sound and would allow some method of determining the presence of the liver.

Chapter 4 Human Subject Tests

The desired system to be used for liver border detection does not only consist of hardware, software and data acquisition tools; it requires a method or algorithm to recognise which signals belong to liver presence and which do not. A method can be developed by simulation, but as it was pointed out, a model of the problem cannot be constructed easily and thus experimental data is required to develop and test an algorithm. This chapter is therefore involved with experimentation to collect data, a mass of signal analysis and feature extraction, and feature matching to underlying tissue (classification).

The experiment uses the tools developed in Chapter 3 to perform tests on volunteers (discussed in Section 4.1) in accordance to the test procedure set out in Section 4.2. In brief, the test involves volunteers lying on a bed underneath the frame being actuated by a series of pulses at different locations in a grid on the surface of the volunteer's body. Data is recorded using the tools developed in Sections 3.3 to 3.6.

Subsequent to these tests, volunteers are scanned using MRI and the actual location of the liver extracted from these images – this is called ground truth extraction which is discussed in Section 4.4.2.

Signal processing and feature extraction (Section 4.3) involves simplifying the mass of information contained in each time signal by representing it by features. Two feature extraction approaches were used – direct feature extraction, where components of the measured signals are used as features; and model fitting techniques where data is fitted to simple models and the parameters of the model that best describe the data is used as features.

Once features are extracted for each measurement point for each volunteer, classification techniques are evaluated to map parameters to underlying ground truth. The ability of a classifier to identify the underlying tissue (ground truth) is used as measure of the success of the method as a whole. Each feature extraction method combined with a type of classifier is considered a method.

The evaluations contained in this chapter ultimately showed that the approach followed was flawed and possible reasons are discussed. A subsequent investigation using a model of the problem (in Chapter 5) showed more promising results, but the chapter is far shorter than this one – the reason is that this chapter represents almost 90% of the effort of this thesis and the author needed to emphasize this and also indicate that no stone was left unturned in an attempt to solve the problem using the particular setup.

4.1 Volunteers

Volunteers consisted of eight male students, ages 22 to 25, of average build (1.75-1.95 m tall, 70-85 kg in weight and medium frame).

Three of the volunteers assisted in tests before the frame was built. The actuator was simply attached to a bench and these volunteers were asked to stand against the actuator whilst a series of pulses were performed. These tests showed significant problems in the repeatability of results and difficulty in pin pointing specific points of entry. These issues, among others, led to the subsequent development of the frame.

Another two volunteers assisted in the development of the experiment after the frame was built. During these tests, usage cases for the development of the control software were developed. Various issues in the movement of the frame and the placement of accelerometers were addressed.

Three other volunteers were used in actual data acquisition tests where data was useful, and this data is used in later analyses.

Each test took approximately an hour to complete in which time the volunteer had to be lying on his back and keep his body still. Although the percussion blows were weak, they were still uncomfortable. It was thus decided to limit the amount of test work and to focus on extracting features that could indicate the location of the liver.

4.2 Test Procedure

Volunteers were asked to lie on the bed underneath the frame described in Section 3.7 as shown in Figure 4.1.

ECG stickers were placed in fixed positions on the surface of the volunteer's skin. The bottom of the sternum was used as a reference point after finding it by palpation. From this point, three marks were made in each horizontal direction, 6 cm from the reference point and each other resulting in seven points. Another two rows of marks were made as the first, but 5 and 10 cm below the first row resulting in a total of 21 points. These points represent the points of percussion and ECG stickers were placed horizontally in between and outside of these points, 3 cm from the input, for attachment of accelerometers.

The actuation points are shown in Figure 4.2, as performed with one volunteer. The red, green and blue rows represent the first, second and third row respectively. The actuation points are labels as outlined in Table 4.1 where R3 represents the furthest position right (relative to the volunteer), LR represent the middle, UD represents the top (cranial) and D2 represents the lowest point (caudal). Some points are also indicated in Figure 4.2.



Figure 4.1: Experimental setup - Volunteer under frame

Table 4.1: Actuation point labels

R3/UD	R2/UD	R1/UD	LR/UD	L1/UD	L2/UD	L3/UD
R3/D1	R2/D1	R1/D1	LR/D1	L1/D1	L2?D1	L3/D1
R3/D2	R2/D2	R1/D2	LRD2	L1/D2	L2/D2	L3/D2




Figure 4.2: Measurement points on patient

The actuator was moved to each point with the head of the shaker placed firmly against the patient's skin. Accelerometers were placed on the ECG stickers directly left and right of the percussion point.

The *DataloggerGUI.m* Matlab program performed a calibration test to ascertain the orientation of the accelerometer and correct the z-axis (as the sensor might be lying skew). During calibration the thumper was excited with a 2 second 10 Hz sinusoidal signal for each accelerometer. The three axes of each accelerometer are recorded and the absolute value of acceleration, relative to the absolute value of the z-axis is calculated. This value indicates the factor the z-axis should be multiplied to get the actual perpendicular acceleration of the skin. (It must be stated that this calibration constant was not applied as signals were normalised in post processing.)

After calibration, a series of nine pulses at one second intervals were performed whilst recording the two accelerometer and FSR signals at each point of entry – this is also controlled from *DataloggerGUI.m*. During this actuation the volunteer was asked to hold his breath as to not influence the measurement. An example of the measured signals is shown in Figure 4.3. As can be seen, it is difficult to see what the actual signals look like unless one looks at less samples at a time and one removes the trend in the FSR signal to remove the large DC offset – this is shown in Figure 4.4. Accelerometer 1 was placed left and accelerometer 2 right of the actuator input, relative to the volunteer. The FSR voltage signal indicates when the force was applied. The FSR voltage may look very discontinuous and unnatural, but the voltage signal is very sensitive to slight force changes which could have been caused by slight movements or settling of the sensor. These effects were evaluated critically in Section 3.2, and calculated FSR force signals are shown in Sections 3.2 and 4.3.1.1.




Figure 4.3: Example of signals measured




Figure 4.4: Examples of signals measured - enlarged




Figure 4.5: Typical MRI Image and analysis thereof

After the patient had been subjected to this battery of tests, an MRI was performed of the patient's lower thorax and abdomen to determine the actual depth of the liver and other sub dermal presences (like ribs). An MRI slice is shown in Figure 4.5 with the relevant measurements extracted (with the help of SYNGO fastview software from Siemens AG). The extraction of ground truth is discussed in Section 4.4.2.

4.3 Signal Processing and Feature Extraction

The goal of signal processing and feature extraction is to process signals to remove noise and time align them, and then find features within signals which correlate well with *ground truth*. In the ideal case, features would correlate well with ground truth and a linear relationship could be established. In practice, however, the relationship may be far more complex (multiple groups or clusters of data points indicating underlying sub classes, or odd probability density functions, or simply too many features to linearise a relationship) and require a statistical classifier to be trained to map features to the correct class.

Feature extraction was done in two ways:

- Direct extraction from measured responses – by looking at the response signal one can extract features like rise time, decay time, resonant frequency, or amplitude. Note that these features ignore the input signal or assumes that it is consistent.
- Fitting of a model – one could fit a model that either matches the physics of the underlying problem (if one knew that the body wall behaved exactly like the membrane of a drum one would try to fit the wave equation, however this is not the case); or one could fit a model that seems to fit the response seen (this will be discussed later, but the measured responses seemed to behave like a mass-spring-damper system) – the model parameters that fit the measured signals best are compared to ground truth.

Direct extraction of features is the simpler approach, where the FSR signal is used as a time reference and features like maximum acceleration, settling time (decay rate) and peak time are extracted from the accelerometer signal. One could also look at the frequency content of a typical signal by looking at the amplitude spectrum, selecting specific frequency bands and calculating the energy contained therein; or one could simply look for the peak of the frequency content.

To fit a model to data might be a considered an easy option if an expected model for the situation existed – this is not the case. Two model types were considered: (i) a general autoregressive model; and (ii) a more specific mass-spring-damper system. An autoregressive model is general in that it can describe various model types (mechanical or thermal) with the same parameters called “time series coefficients”. The more specific mass-spring-damper system was considered after data collection and subsequent analysis of the data where it became apparent that the accelerometer signal behaved like the response of a single or dual mass-spring-damper system to the FSR signal.

4.3.1 Direct Feature Extraction

Direct feature extraction involves extracting features from time or frequency spectrum signals. In order to extract features from the captured data, the raw data needs to be pre-processed which involves time aligning multiple signals, allowing the possibility to visually compare time signals and phase information.

Section 4.3.1.2 is involved with the identification of features that separate signals from one another. The extracted features are used in Section 4.4.3.1 to train and evaluate classifiers.

4.3.1.1 Pre-processing of Data

The data that was captured during the test procedure is contained in time signals representing the input and output at each measurement point. If one wanted to compare signals to each other one would want to do this visually first, which means signals must be in phase with each other – i.e. the time axis is the same relevant to the percussion stroke in each measurement. Although signals are captured using the same sequence for each measurement point, the result is not aligned in time – an example of some FSR signals are shown in Figure 4.6.

Another consideration is contracting signal information which consists of a series of nine pulses that were performed at each measurement point. By taking the average of these signals one should be able to get a clear signal and reject the noise. The use of an average signal rather than absolute single signals to determine the signal trace of percussion sounds (recorded by microphone) was also performed by Mansy et al. [31].

A Matlab function called *get_aver_pulse.m* was created to time align signals and get the average. The algorithm finds the first pulse in the FSR signal measured at a particular measurement point (e.g. L1UD), then uses cross-correlation to find the time offset of the subsequent pulses. The other signals measured with the same DAQ sequence (Accelerometer 1, Accelerometer 2) should have the same time offsets and the result obtained is applied to each signal to obtain an average of that signal. More information on the algorithm used can be found in Appendix B.1.

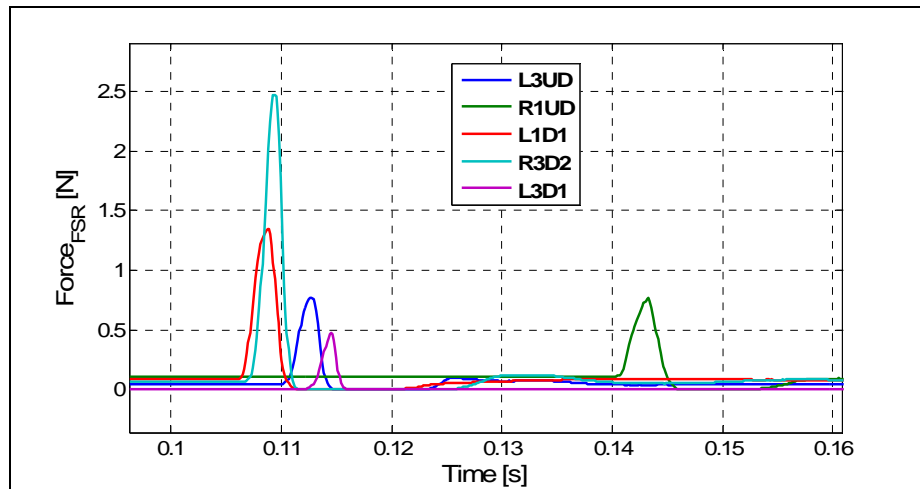


Figure 4.6: FSR signals showing time jitter

Processing the signals in Figure 4.6 yields the results shown in Figure 4.7. Using these averaged and time aligned signals allow the bulk comparison of numerous signals. The accelerometer signals from the top row of one data set are shown in a waterfall plot in Figure 4.8. Plots like these are used in the next section to identify features for extraction.




Figure 4.7: Time aligned and averaged FSR signals

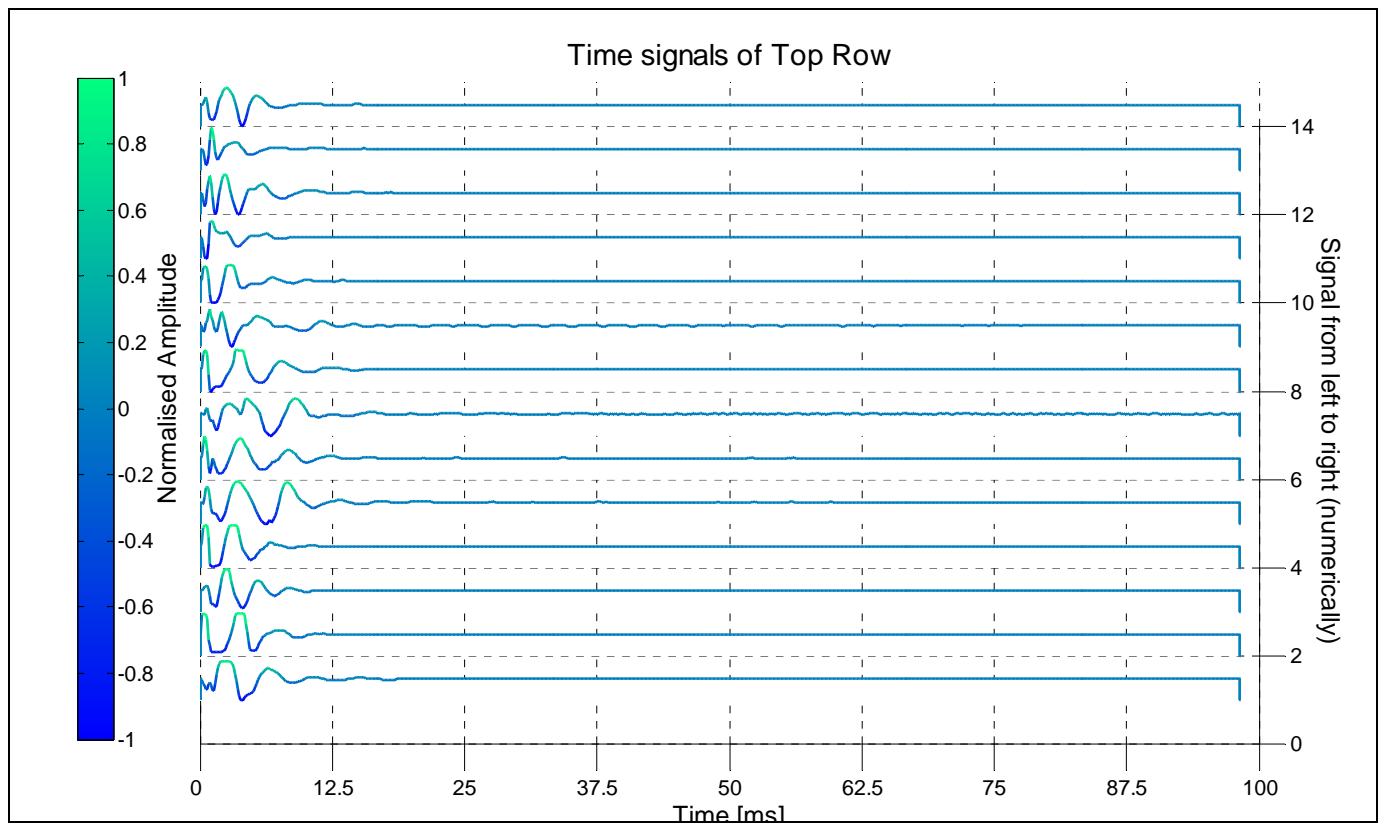


Figure 4.8: Series of time aligned accelerometer signals

4.3.1.2 Identifying Features

The averaging and time alignment discussed in the previous section allows one to look at multiple signals at the same time and identify possible distinguishing features. The goal of this section is to provide a list of features to be used for further evaluation (in Section 4.4.3.1) as well as the methods to extract them.

From the time signals depicted in Figure 4.8 it seems that most signals have a natural frequency multiplied with an exponential rise and decay. One can replace rise and decay rates with rise time and settling time as these are inter-related and the maximum component in the amplitude spectrum should be the dominant frequency. An example of the extraction of time features is shown in Figure 4.9 **Error! Reference source not found.** Note that some signal traces are noisy which leads to false detection of these features.

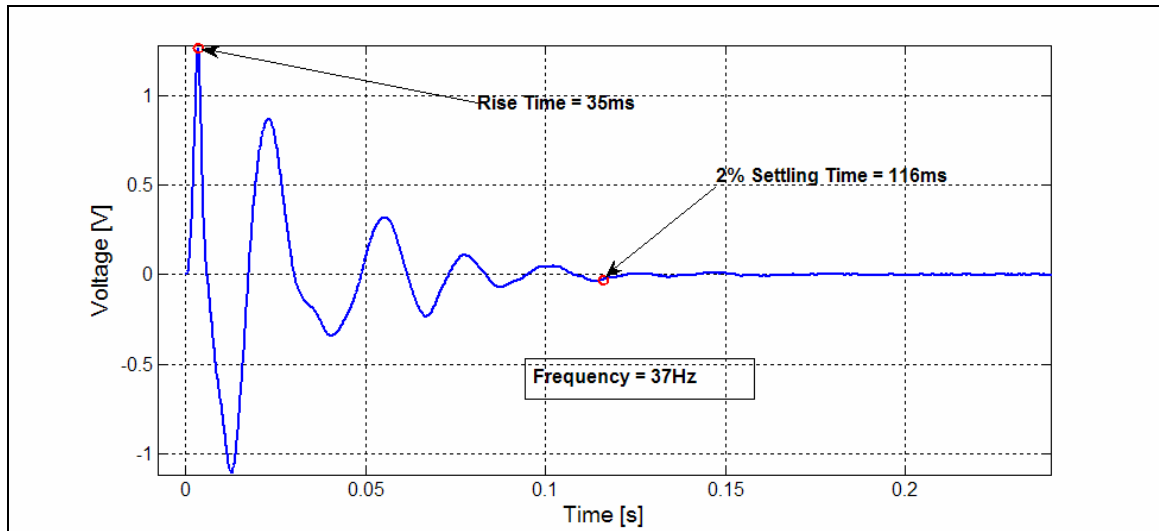


Figure 4.9: Extracting rise time, settling time and frequency

If one looks at the amplitude spectrum for a row of accelerometer signals one obtains the spectrum depicted in Figure 4.10. Most signals have a clear central lobe, but also side lobes. One way to generalise around this problem is to calculate the percentage of power over frequencies with subsets of frequencies overlapping. Eleven features are selected by calculating percentages of total energy in intervals of 50 Hz every 25 Hz from 0 Hz to 300 Hz. The amplitude spectra of the signals in Figure 4.10 have the features shown in Figure 4.11.

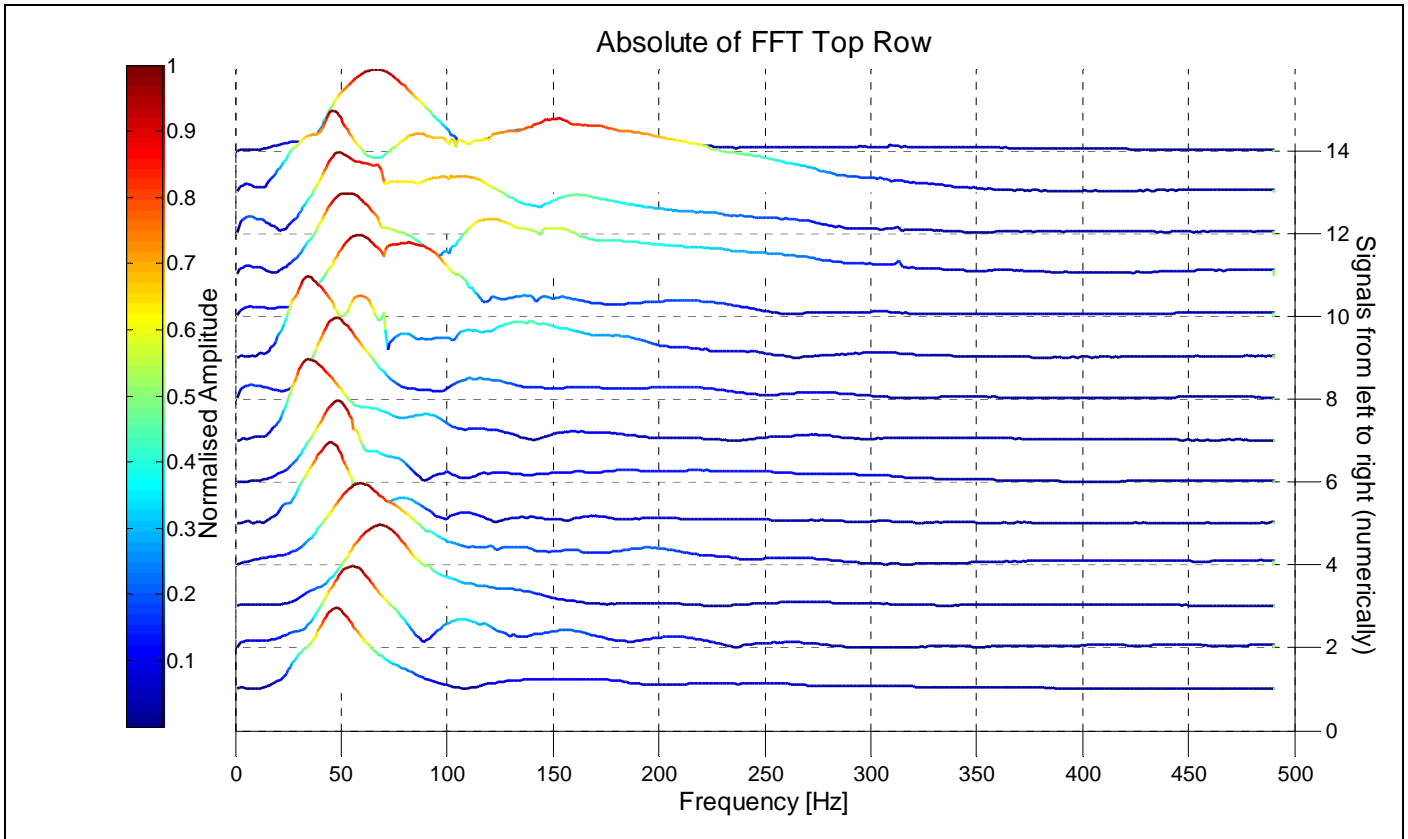


Figure 4.10: Amplitude spectra of multiple signals

An attempt was also made to generalise the response as that of a low pass filter – the point where the phase spectrum changes slope would be the point of the dominant pole. This method and feature was discarded as it could not be extracted with sufficient confidence.

Fourteen features were therefore extracted directly. These features are used for classifier training and testing in Section 4.4.3.1. Note that there were only three full data sets, with twenty-one points containing two accelerometer signals each – in total 126 data points’ features were extracted.

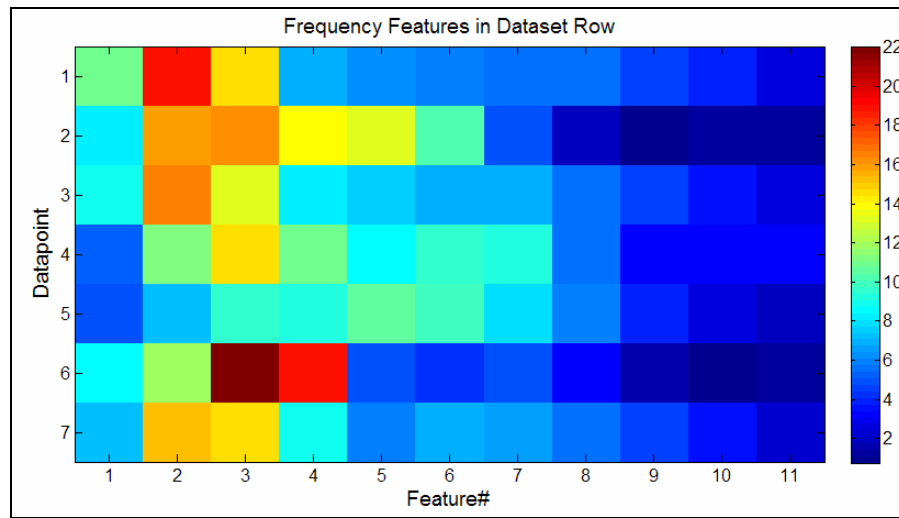


Figure 4.11: Features (percentage of energy in subsets of spectra) extracted from amplitude spectra

4.3.2 Model Fitting Methods

Fitting models to input-output data is an attractive approach as one reduces the data to the parameters that describe the model that best fit the data. As stated earlier, the underlying mechanisms that create a percussion sound is not agreed upon nor described technically. Therefore, no known model with unknown parameters exists. One solution is to fit data to a more generic model – like a higher degree filter, or simple mechanical system. The problem with these approaches, however, is that choosing a model structure that does not describe the problem well, could cause erroneous results even though the model can be made to fit. Increasing the complexity of a model means more describing which would lead to the case where one simply describes a signal by its N points and therefore require no model.

The various models that were considered can be categorized according to the tools that were used to develop/evaluate them:

- Matlab's *System Identification Toolbox* provides tools where one can determine the parameters of linear polynomial models (black box, system unknown) like ARX (Autoregressive model with exogenous inputs), ARMAX (Autoregressive moving average model with exogenous inputs); or linear and non-linear grey box (system known but parameters unknown) models described by unknown parameters of ordinary differential equations (ODE), that best relate input to output data. The parameters are estimated using linear regression and optimization using a gradient direct search method.
- Matlab was also used to fit input-output data to a single and dual mass-spring-damper system using stochastic optimization methods, more precisely, the *particle swarm* and *genetic* optimization algorithms (this could also be done using the linear and non-linear grey box models in the *System Identification Toolbox*, but the direct search method was a drawback). These search methods search over a whole search space continuously and use the best results to generate new points of evaluation. The methods are based on swarm and genetic theory, respectively, and work well for problems where multiple local optima and discontinuities in the search space might exist. Although the particle swarm optimization algorithm seemed promising it did not lead to good model fits and this idea was subsequently cast aside.

In Section 3.8 it was pointed out that sensor signals (FSR and accelerometer) possess unnatural artefacts some of which may have been caused by sensor error. One problem that the accelerometer signal experienced was saturation as the offset was too close to one of the rails. It was hoped that model fitting techniques would reject these saturations and noise and fit more natural looking responses (where appropriate). A short evaluation is done using the mechanical model and this evaluation can be found Section 4.3.2.2. The evaluation shows that data could be fitted well, but that the fitted parameters mismatched the original slightly.

4.3.2.1 Autoregressive Models

The *System Identification Toolbox* provides a multitude of possible models. Initially it was unclear which simple mechanical system could possibly closely resemble the data that was captured and it was decided to attempt to fit more general models to the data – therefore the ARX and ARMAX models (these belong to the more general class of time series models).

An autoregressive model with exogenous inputs (ARMAX) is simply a model where the next output value is a linear combination of the previous input and output values where the moving average adds a term for the moving average of the output. For ease of reference only the ARMAX model structure is shown in equation 2 [54].

$$y(t) + a_1 y(t-1) + \dots + a_{n_a} y(t-n_a) = b_1 u(t-n_k) + \dots + b_{n_b} u(t-n_k-n_b+1) + c_1 u(t-1) + \dots + c_{n_c} u(t-n_c) + e(t) \quad (2)$$

Here $y(t)$ and $u(t)$ are the output and input at time t , n_a is the number of poles, n_b is the number of zeros plus 1, n_c is the number of C-coefficients, n_k is the dead time and $e(t)$ is the white noise term. The ARX model structure does not have c-terms.

The ARMAX and ARX model structures therefore have $[n_a \ n_b \ n_c \ n_k]$ and $[n_a \ n_b \ n_k]$ degrees of freedom, respectively.

The toolbox uses the measure in equation 3 to determine the goodness-of-fit of the model output (y_h) to the actual output (y). This is the objective function that the search algorithm maximizes to find an optimal fit.

$$Fit(\%) = 100 \times \left(1 - \sqrt{\frac{\sum (y - y_h)^2}{\sum (y - \text{mean}(y))^2}}\right) \quad (3)$$

Initially the raw sensor output data was used with the FSR signal as the input and accelerometer 1 signal as the output and an average fit of only 40% was attained. The problem was attributed to noise in the signal. Subsequent filtering caused “smear” in the sharp rising edge of the FSR signal which led to the input lagging the output – an ARX model would not fit this scenario unless first input sample is multiplied infinitely to attain an appropriate output.

A vast improvement was attained when the input was shifted forward in time by 75 samples and the input and output signals zeroed before a non-noise change in the FSR signal was registered. A Matlab script ran through all the signals in multiple experiments and found the ARX fits by varying the number of poles (n_a) and number of zeros (n_b-1) from six to 15 (five or less did not yield adequate fits) and only using the dead time (n_k) (zero to 150 samples) that resulted in the best fit.

Additionally, signal filters were experimented with to ascertain which level of filtering might be the best out of those selected – the following Butterworth filters were experimented with: (1) none; (2) 300 Hz 4th-order; (3) 360 Hz 4th-order; (4) 400 Hz 4th-order; (5) 440 Hz 4th-order; (6) 480Hz 4th-order; (7) 600Hz 4th-order; (8) 600Hz 6th-order; and (9) 600Hz 8th-order.

The effective discrete search space therefore has two dimensions $-n_a, n_b$ and 9 possible filters; and at each point in the search space, three experimental datasets were evaluated each with 21 input-output combinations resulting in 10 (*possible poles and zeros*) x 9 (*possible filters*) x 42 (*data points*) x 3 (*subject*) = 11340 fits.

The average, variance, maximum and minimum of the fits for each possible combination (in the two dimensional search space) is listed in Table 4.2, sorted from best to worst performing and only a subset of the results.

Table 4.2: ARX Fits Results (sub set)

Mean of fits (%)	Variance of fits (%)	Max of fits (%)	Min of fits (%)	Order of Na	Order of Nb	Filter number
84.57	9.37	97.24	48.16	14	14	9
84.40	9.30	97.37	45.17	14	14	2
83.79	9.38	97.18	47.18	12	12	2
83.65	9.79	97.06	47.04	12	12	9
83.62	9.98	97.43	42.84	15	15	2
83.36	10.07	96.55	45.48	15	15	9
82.98	9.71	96.88	46.19	10	10	2
82.96	9.90	96.55	44.27	14	14	3
82.85	10.09	96.84	43.34	13	13	2
82.60	10.29	96.17	42.69	15	15	3
82.33	10.20	96.18	44.49	14	14	4
82.27	10.08	96.26	44.30	12	12	3
82.17	10.25	96.71	44.22	11	11	2
82.15	10.37	95.85	41.14	15	15	4
82.02	9.97	96.78	43.14	8	8	2
81.98	10.27	95.96	40.54	14	14	8

A typical example of an input-output signal and its fit is shown in Figure 4.12. In this particular case the model can be represented by the following 13 parameters: $n_k = 127$ samples; $a_{1-6} = [-1.9094, 0.8999, -0.2022, 0.2454, 0.093, -0.1249]$; and $b_{1-6} = [0.0349, -0.0449, 0.0113, -0.0026, 0.0114, -0.0078]$; there was also no filtering except for averaging. This particular fit seems to fit the initial part of the signal well, but not as well in the decaying stage.

A second example is shown in Figure 4.13. This model can be represented by the parameters: $n_k = 143$ samples; $a_{1-6} = [-2.5382, 1.9761, -0.3608, 0.3205, -0.7357, 0.3403]$; and $b_{1-6} = [0.041, -0.1139, 0.094, -0.0292, 0.0306, -0.0238]$; there was also no filtering except for averaging. Again, the initial part of the signal fits very well, but the rest of the output signal is not reflected by the model.

The signal traces in Figure 4.12 and Figure 4.13 show a significant delay between the input and output, but this delay is the one that was artificially induced (as stated in an earlier paragraph, the output was shifted 75 samples along the time axis) so that the model would find the time lag that best fits the input-output data. The artificially induced time lag is equal to 9.3 ms, to which the anomalies in signal traces can be attributed to.




Figure 4.12: Example 1 of fit with ARX model

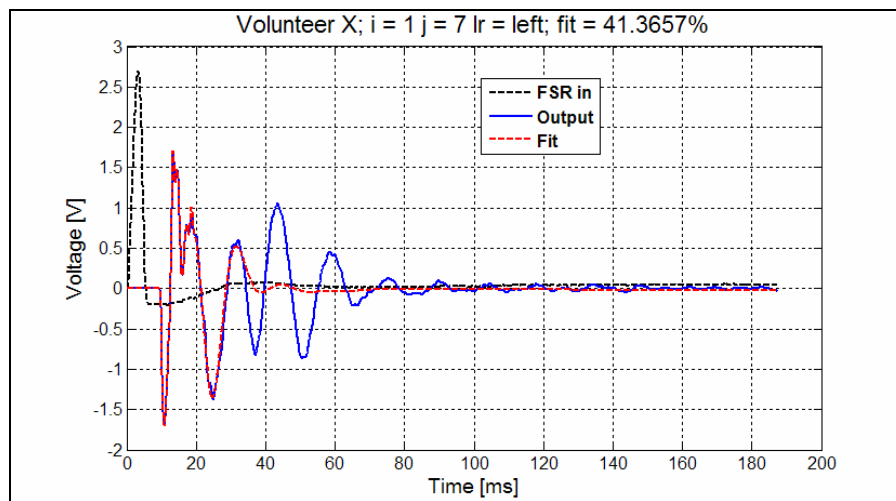


Figure 4.13: Example 2 of fit with ARX model

Once these parameters are extracted for each point in each dataset they can be used as features to train and evaluate classifiers that will map them to underlying ground truth – liver or none. In order to limit the error that is induced by noisy data (data that does not fit a model), thresholds for the required fit are set. Figure 4.14 shows a histogram of measured fitness of model to data for the case where no filtering was applied with orders of six for both “a” and “b”. In this case the threshold could be set at 50 %.

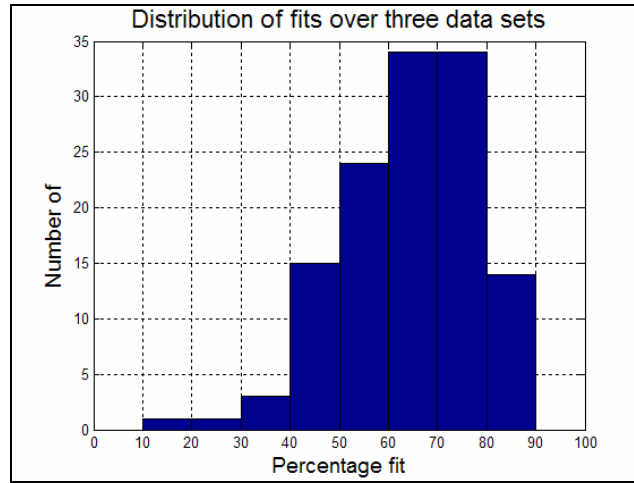


Figure 4.14: Distribution of fitness values for data fitting to an ARX model

4.3.2.2 Mechanical Model

Although the parameters describing the ARX systems produce seemingly good fits, the number of parameters describing each data point is large (30 in the most extreme case). If one looks at a typical input-output example (like that shown in Figure 4.12) the behaviour could belong to that of a mass-spring-damper or similar system. Various systems were considered (and evaluated) and it was found that a single mass-spring-damper system does not describe the system well; a non-linear spring did not contribute to the description; but a dual mass-spring-damper system does describe the system with reasonable correlation.

When considering a dual mass-spring-damper system, described by seven parameters ($m_1, m_2 = \text{mass [kg]}$, $k_1, k_2 = \text{spring constant [N/m]}$, $B_1, B_2 = \text{damper constant [N/m}^2]$ and $A = \text{linear gain [no unit, to compensate for overall signal gains/attenuations]}$) and as shown in Figure 4.15, one would have only seven features describing each measurement point.




Figure 4.15: A dual Mass-Spring-Damper system

The system in Figure 4.15 is described by the following ODEs (Ordinary Differential Equations):

$$m_1 \ddot{x}_1 = F - k_1(x_1 - x_2) - b_1(\dot{x}_1 - \dot{x}_2) \quad (4)$$

$$m_2 \ddot{x}_2 = k_1(x_1 - x_2) + b_1(\dot{x}_1 - \dot{x}_2) - k_2 x_2 - b_2 \dot{x}_2 \quad (5)$$

For each particular data set one would evaluate input and output data and attempt to find the model parameters that best fit the particular data set. Unfortunately, finding these parameters is not trivial and one would need a search algorithm to find the best fit within some search constraints (number of iterations, upper and lower bounds of parameters).

Equations 4 and 5 can be converted to their Laplace representations and then drawn in a block diagram for use with Matlab's Simulink. The Simulink model is shown in Figure 4.16. The seven variables are adjusted by an optimization algorithm to find a model that has a similar output to that measured. Additionally the model allows for some transport delay and saturation in the signal.

Finding the parameters that describe a model that fits the data is an optimisation problem, thus an objective function needs to be defined and it needs to be representative of the goal that is to be reached. Various objective functions were considered and it was found that typical fitness functions (like that used in equation 3 in Section 4.3.2.1) found solutions that are over damped to be a better fit than those that are under damped even though the true output was in fact under damped. Problems were also encountered due to the noisiness of the input (FSR signal) at low levels. An ad hoc measure was defined by identifying the start of the input and the possible end of the measured output and forcing the rest of the signal to be zero. This adapted output signal was then used as comparison to the model output with the same fitness measure as used in equation 3.

The most common optimisation algorithms are Direct Search (DS) based. These are good at finding a local optimum (which is sufficient in a search space where the local optimum is likely to be the global optimum), but do not perform well in situations where there might be many local optima that are far less optimal than a possible global optima. Stochastic optimization algorithms use behavioural patterns (like genetics or insect behaviour) to find global optima and in many difficult problems tend to outperform direct search methods. These algorithms search spaces by starting off with random positions in the search space and change positions based on some or all evaluations.

Two stochastic optimization algorithms were initially considered: the genetic algorithm (GA) and the particle swarm optimization (PSO) algorithm. Evaluations were done on small examples and it was found that the genetic algorithm found much better solutions than the PSO algorithm, which compelled the researcher to continue using only the genetic algorithm to find approximate solutions. The GA results were subsequently fed into a DS algorithm which found the local optimum around this particular solution.

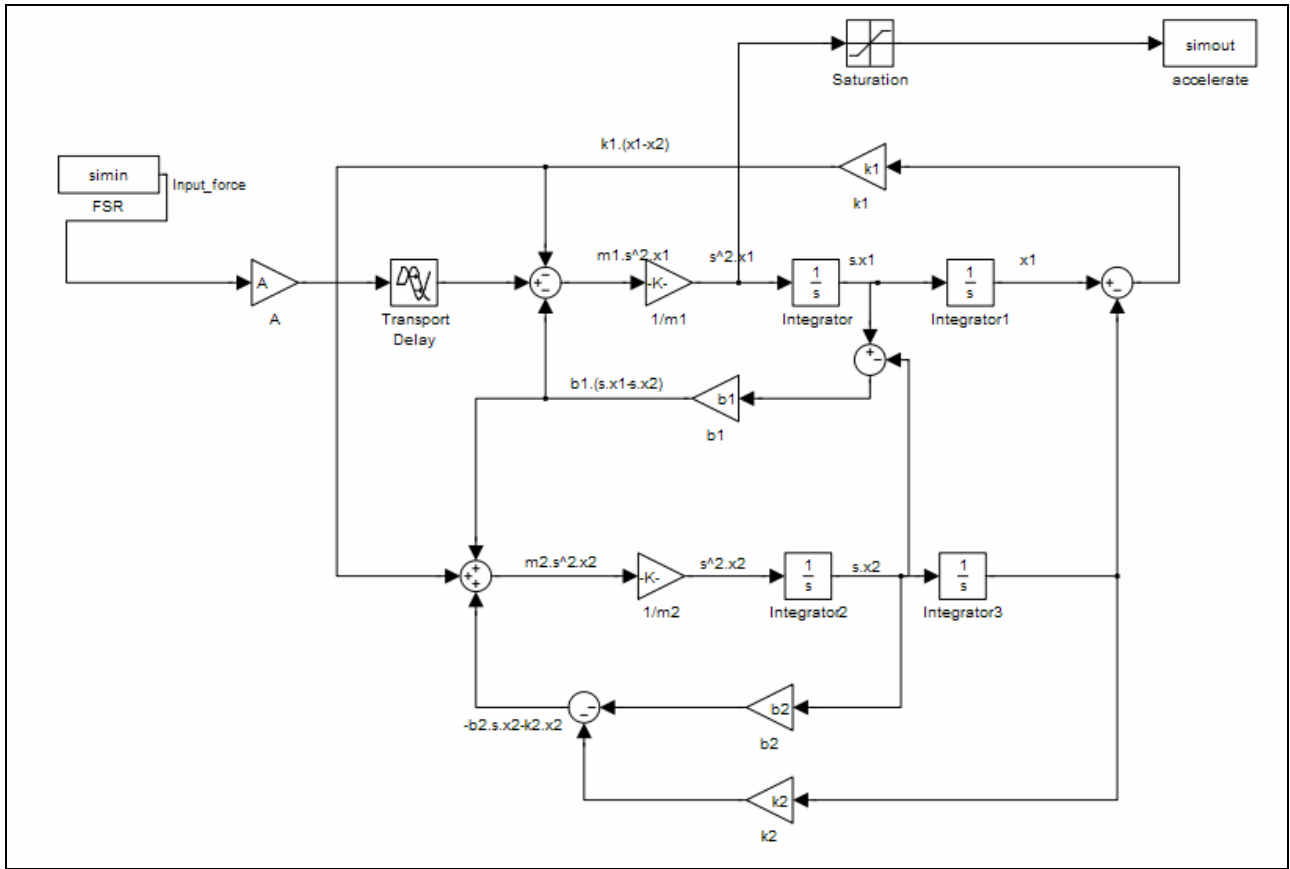


Figure 4.16: Simulink model of dual Mass-Spring-Damper system

The GA is based on natural selection in nature [55] & [56] – the process that (arguably) drives biological evolution. The algorithm repeatedly modifies a population of possible solutions to generate a new generation. Modifications are based on selection rules (how the parents that will reproduce are selected), crossover rules (how the solutions of the parents will be combined to form a new solution) and mutation rules (random changes to parent’s solutions). The GA flow is depicted in Figure 4.17.

A Matlab script was created to find optimal solutions using Matlab’s Genetic Algorithm and Direct Search Toolbox. Bounds were placed on the values for the spring constant, damping constant, gain and masses as follow: $\text{minimum}(m_1, b_1, A, k_1, m_2, b_2, k_2) = [2 \ 10 \ 10 \ 1000 \ 2 \ 10 \ 1000]$ and $\text{maximum}(m_1, b_1, A, k_1, m_2, b_2, k_2) = [50 \ 10000 \ 800 \ 10000000 \ 50 \ 10000 \ 10000000]$. These values were empirically tested and found to be the most appropriate. The algorithm population was limited to 20 individuals and the algorithm to 100 generations (these values were also empirically tested). Other parameters pertaining to the GA were left at default. Values (parameters) generated by the GA was passed to the objective script, which passes values to the Simulink model and compared the measured model output to actual output data using the objective function defined earlier. The optimal result obtained by the GA was passed to a direct search algorithm (from the same toolbox) to find the local optimum using derivative methods. The direct search algorithm was employed using its default parameters and the same boundaries as that of the GA. The software flow diagram can be found in Appendix B.2.3.

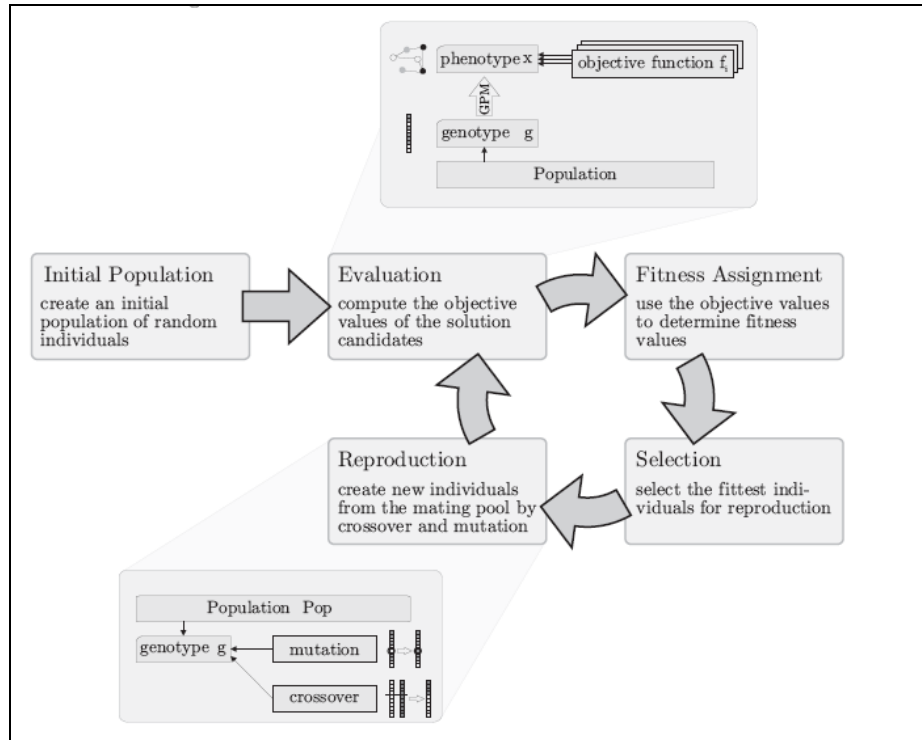


Figure 4.17: The basic cycle of the GA [56]

Before delving into actual results, a concern that is raised in Section 3.8 is that the sensors signals are non-optimal as it contains artefacts. The FSR signal “jumps” can be ascribed to small movements of the volunteer and saturation of the voltage signal is shown to be due to lack of force when the signal is converted to its force equivalent. Some accelerometer signals show sudden sharp changes, which seem unnatural but this cannot be proved. The saturation of the accelerometer signal, however, is due to actual saturation of the signal. The saturation effect is a clear unwanted artefact and this effect is simulated in the following example, with additive noise in the output.

A model was created with the following parameters: $A = 250$; $m_1 = 20$; $m_2 = 45$; $k_1 = 2.8M$; $k_2 = 7.9M$; $b_1 = 8000$; and $b_2 = 1300$, and simulated with a simple FSR input. The output acceleration was recorded with and without additive noise and saturation. The noisy signal was used to fit a model to and a final fit of 93% was obtained. The signals are shown together in, the first plot showing all signals together and therefore some not visible as they are shadowed; the second plot shows the saturation effect in the signal to be fitted and the fitted signal matching the noiseless signal; and the third plot shows the noise and clearly how close the fit is to the actual model output without noise. The fitted model has the following parameters (rounded off): $A = 225$; $m_1 = 15$; $m_2 = 36$; $k_1 = 2.4M$; $k_2 = 5.6M$; $b_1 = 9694$; $b_2 = 1710$.




Figure 4.18: Using mechanical model fit with simulated noise

This result shows that these artefacts could be ignored by fitting a mechanical model to the data, but that the fit is not perfect and the slight difference represents a system with different parameters. The difference between parameters represents up to 23% of the range of k_1 as an example. Although this is discouraging, a classifier may be able to find a linear combination of features that remain constant between the original parameters and the fitted ones. A critical evaluation of this suggestion was not performed, but the simulation does indicate that differences in a few parameters still result in a very similar system.

Figure 4.19 to Figure 4.22 show examples of both good and bad fits on real data. Bad fits were obtained for data that could not be represented by a model – this may be due to very noisy FSR signals and it was also found that the smaller the amplitude of the output signal the weaker the typical fit.

The data and model represented in Figure 4.19 has the following parameters (as an example):

$$A = 248; m_1 = 20.7; m_2 = 33.0; k_1 = 2.8\text{M}; k_2 = 7.9\text{M}; b_1 = 8838; b_2 = 31.$$

The data and model represented in Figure 4.20 has the following parameters (as a comparative example):

$$A = 30; m_1 = 8.2; m_2 = 9.6; k_1 = 2.6\text{M}; k_2 = 7.1\text{M}; b_1 = 940; b_2 = 5483.$$




Figure 4.19: Example 1 of good fit on mechanical model




Figure 4.20: Example 2 of good fit on mechanical model

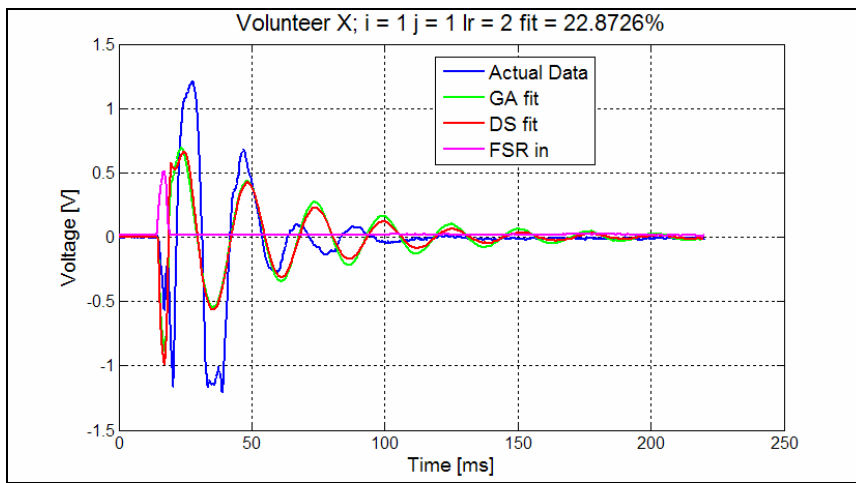


Figure 4.21: Example 1 of bad fit on mechanical model




Figure 4.22: Example 2 of bad fit on mechanical model

A distribution of fitness values obtained for the three datasets evaluated is shown in Figure 4.23. Only 40 of the 126 measurement points could fit a model with a fitness of more than 50 %. Data used in the next section for parameter mapping was limited by the requirement that the fitness of the model parameters has to exceed 40 % - there is a total of 59 data points that fit this requirement.




Figure 4.23: Distribution of fitness values for data fitting to mechanical model

4.4 Classification

Features were extracted from the data captured during the test procedure using the various methods discussed in Section 4.3. The extracted features form so called data sets – one data set for each feature extraction method. These data sets are all relatively small which limits possible classifier complexity (the evaluations will not be shown, but the use of combinations of classifiers did not improve results for classification of the extracted features).

A classification system is developed using training data in a process of training. Once a classifier is employed, it returns the likelihoods of a test feature belonging to any of the classes that the classifier is trained for. These likelihoods are so called *soft outputs* and are followed by *hard decisions* by applying a weighting or threshold before assigning the class with the highest likelihood to a feature (or data point) The practice of the classifier calculating likelihoods is also referred to as a mapping (the terminology used in this thesis matches that used by Van der Heijden et al. [57] and *PRTtools* a Matlab toolbox used for pattern recognition). Some mappings are linear or affine – matrix rotation projection and shift; and others are non-linear – these include various, e.g. Mixture of Gaussians (MOG), or Quadratic Bayes (QDC).

For each feature extraction method, the features are separated into training and test data. Training data is used to train a classifier to best separate classes in the training data and the subsequent evaluation with test data determines the validity (or performance) of a solution. Mappings can also be combined to create more complex mappings, but, as stated earlier, this is ill advised when such small data sets are considered.

Because the data sets are small, an extreme case of a cross validation technique is also used – for each point in the data set an evaluation is done using all the other points for training and the particular point for testing – this method is known as “leave-one-out” cross validation. The number of errors divided by the size of the data set yields the expected error rate for a particular classification system.

The next section will give an example of a classification problem to gain more clarity on some of the topics. The example is visual in terms of graphs and tabulated results, but the subsequent evaluations on experimental data are less visual and aimed at only giving performance results – this because numerous combinations were considered for each feature extraction method and each yields several plots and graphs, but can be summarised by a number indicative of its performance, the correct classification rate.

Before delving into evaluations on extracted features, the *ground truth* labels need to be extracted using a golden standard for this study, the MRI. MR images are used to determine which input points are close enough to the patient’s liver to be labelled as ‘liver’, or otherwise as ‘no liver’. Without these labels, a classification system will have no means of training a mapping, or evaluating its performance. Ground truth extraction is contained in Section 4.4.2

PRTools was used in all of this section’s evaluations. It is a Matlab toolbox for pattern recognition, developed by the Delft Pattern Recognition Group at the Faculty of Applied Physics at Delft University of Technology [58] and provides most of the tools one would need for basic and some advanced pattern recognition problems.

4.4.1 Classification Introduction

This section will attempt to introduce and clarify the topics of the subsequent sections. The topics are involved with statistical pattern recognition. The explanations are contracted for the sake of brevity and the reader can consult referenced sources for more information.

The feature extraction methods used in Section 4.3 mostly resulted in more than three features being extracted for each method – this makes it difficult to represent features visually. Principle Component Analysis (PCA) involves reducing the number of features (dimensionality) to a smaller number of features by taking linear combinations of the primary features. According to [59], PCA can be defined mathematically as an “orthogonal linear transformation” that transforms data to a new coordinate system such that the greatest variance by any projection of the data will lie on the resulting first axis, second highest variance on the second axis, and so forth. PCA is often performed prior to training a classifier or simply used to better visualise data. Unfortunately, PCA does not separate classes. To simulate PCA an example is shown in Figure 4.24. Here an artificial two class dataset with three features were generated and PCA maps it to two (three dimensions were simply selected as it can demonstrate the method).




Figure 4.24: An example of mapping data using PCA

PCA was used to gain insight into the scatter of some data sets and also as a pre-process to the training of some more complex classification systems.

A *PRTools* data example is depicted in Figure 4.25. The two classes are shaped in the form of opposing bananas, which is typically difficult to classify with a linear threshold (straight line). There are 2000 data points.




Figure 4.25: Difficult data set

Linear Discriminant Analysis (LDA) differs from PCA in that it attempts to find a linear combination of the features that best separate two classes [60]. *PRTools* implements LDA using Fisher's Linear Discriminant Classifier which does not assume uncorrelated features (LDA normally assumes this as a simplification). The algorithm maps data to a maximum dimension of the number of features minus 1 and then measures the distance from points on this axis to determine the class likelihoods – the results are shown in Figure 4.26. Performing a 100-fold cross validation yields an error rate of 14.4 % (or correct classification rate of 85.6 %). For one particular combination of 50 % training and 50 % test data, the confusion matrix result shown in Table 4.3 results from evaluation on the test data.

Table 4.3: Confusion matrix from Fisher Linear Discriminant example

True Labels	Estimated Labels		Totals
	1	2	
1	410	77	487
2	68	444	512
Totals	494	505	




Figure 4.26: Fisher Linear Discriminant example

A Quadratic Bayesian Classifier (QDC) assuming normal (Gaussian) densities estimates the Gaussian parameters (mean and covariance) of the training data and applies the threshold where class likelihoods are equal [60]. The outputs of the mapping are the likelihoods of any data point belonging to either class. The largest likelihood determines the class (in default cases), but this can be adapted should the cost of misclassification of a particular class be too high. The results are shown in Figure 4.27. Performing a 100-fold cross validation yields an error rate of 15.25 %. For one particular combination of 50 % training and 50 % test data, the confusion matrix shown in Table 4.4 results from evaluation on the test data.

Table 4.4: Confusion matrix from Quadratic Bayes discriminant classifier example

True Labels	Estimated Labels		Totals
	1	2	
1	427	77	504
2	70	425	495
Totals	497	502	999




Figure 4.27: Quadratic Bayes discriminant classifier example

A Mixture of Gaussians classifier clusters data in the number of clusters specified and determines the Gaussian parameters of each cluster (the underlying algorithm is far more involved, but not part of this discussion). The likelihood of the combined Gaussians determines the class likelihood for each class and the mapping maps features to these likelihoods. The results of training a mapping with a mixture of three Gaussian densities for each class are shown in Figure 4.28. Performing a 20-fold cross validation (the training of a MOG classifier takes much longer than training the previous two classifiers – therefore less folds are used) yields an error rate of 1.6 %. For one particular combination of 50 % training and 50 % test data, the confusion matrix shown in Table 4.5 results from evaluation on the test data.




Figure 4.28: The MOG classifier results

Table 4.5: Confusion matrix from MOG classifier example

True Labels	Estimated Labels		Totals
	1	2	
1	497	8	505
2	10	485	495
Totals	507	493	1000

An Artificial Neural Network (ANN) classifier involves the use of a network of processing nodes or neurons each connected with a weight (multiplication factor). The weights and nodes are trained to produce the desired output – a separation between classes. The ANN has recently lost popularity to the methods of Statistical Pattern Recognition (to which Quadratic Bayes, Fisher Linear Discriminant and Mixture of Gaussians belong, among various others). A back-propagation trained feed-forward neural net classifier with one hidden layer and three nodes in the hidden layer was trained on the example data, the results of which are shown in Figure 4.29. Performing a 20-fold cross validation (the training of an ANN classifier also takes much longer in training) yields an error rate of 5.7 %. For one particular combination of 50 % training and 50 % test data, the confusion matrix shown in Table 4.6 results from evaluation on the test data.

Table 4.6: Confusion matrix from ANN classifier example

True Labels	Estimated Labels		Totals
	1	2	
1	480	16	496
2	9	495	504
Totals	489	511	1000




Figure 4.29: The ANN classifier results

The classifiers shown here are used in further analysis on extracted features, but it must be stated that a multitude of other classifiers were considered (Simple LDA (uncorrelated features), Nearest Mean Classifier, Nearest Mean Scaled Classifier, K-Nearest Neighbour Classifier) all with similar or weaker results. Possibilities that weren't considered include Support Vector Machines and Hidden Markov Models, as well as combinations of used and unused classification structures..

4.4.2 Extracting Ground Truth

As explained previously, ground truth is the actual underlying class that a data point belongs to. Depending on the size of the data set, one could use very specific class labels like “ribs and liver more than 2 cm below surface”, “fat layer and liver more than 2 cm below surface”, etc.. The extreme case would be where measured data is fitted to measurable parameters, but the amount of data required would be enormous. Unfortunately very small data sets are available and it was hoped that simple class labels (like “liver”, “no liver”) could be distinguished as a proof of concept.

Three full data sets were acquired from three volunteers. Only two volunteers (Volunteers 1 and 2 in Table 4.7) were submitted to MRIs, but assessment of the images obtained revealed that these volunteers had very similar liver sizes and loci, leading to the assumption that the third volunteer (Volunteer 3) had a similar ground truth – discrepancies where assigned the label “none” (this is probably not a safe assumption at all, but the data is so limited that it was paramount to the study to at least process all the available data – the additional volunteer should also have a liver on the right side of his body (in very rare cases people have mirrored organs)).

Table 4.7: Summary of assigned labels (ground truth)

L=Liver; N=None.		R3	R2	R1	LR	L1	L2	L3
Volunteer 1	UD	N	N	L	L	L	N	N
	D1	L	L	L	L	L	L	N
	D2	L	L	L	N	N	N	N
Volunteer 2	UD	N	L	L	L	N	N	N
	D1	L	L	L	L	N	N	N
	D2	L	L	N	N	N	N	N
Volunteer 3	UD	N	N	L	L	N	N	N
	D1	L	L	L	L	N	N	N
	D2	L	L	N	N	N	N	N

The labels summarized in Table 4.7 were used for evaluation purposes.

Images of the MRI slices and the associated labels are shown in Figure 4.30 for these two volunteers with Volunteer 1 on the left and Volunteer 2 on the right.




Figure 4.30: Labelled labels on data capture locations for two subjects

4.4.3 Classification of Extracted Features

4.4.3.1 Directly Extracted Features

In Section 4.3.1.2 fourteen features were extracted directly from the accelerometer signals recorded during tests. These features are the rise and 2% settling times, the dominant frequency, and the percentages of energy contained in 11 subsets of the amplitude spectrum. Evaluations were performed on the first three and last eleven features separately and then together. A scatter plot of the first three parameters is shown in Figure 4.31.

As stated in the introductory section, four classification systems are used in the evaluation of mappings, but with the addition of a PCA combined with an MOG mapping. The five mappings are: (i) Fisher's linear discriminant analysis (FISHER); (ii) Quadratic Bayes assuming Gaussian densities (QDC); (iii) a Mixture of two Gaussians per class (MOG2) using all features; (iv) a MOG2 after PCA mapping features to three dimensions; and (v) an Artificial Neural Network (ANN) with one hidden layer of three nodes.




Figure 4.31: Scatter plot of first three directly extracted features

For each of the feature sets combined with each classification system a cross-validation process was followed with four folds for (i) and (ii) and three folds for (iii-v), as well as single 50 % training, 50 % test data confusion matrix results.

The confusion matrix results are summarised in the tables of Table 4.12. The cross validation results can be summarised as follow (values indicate correct classification rates):

- (i) FISHER – three features: 56 %; eleven features: 54 %; fourteen features: 52 %.
- (ii) QDC – three features: 53 %; eleven features: 61 %; fourteen features: 56 %.
- (iii) MOG2 – three features: 57 %; eleven features: 54 %; fourteen features: 52 %.
- (iv) MOG2+PCA – three features: 60 %; eleven features: 45 %; fourteen features: 64 %.
- (v) ANN 1(3) – three features: 50 %; eleven features: 53 %; fourteen features: 59 %.

Table 4.8: Confusion matrix results on directly extracted features

Three Features				Eleven Features				Fourteen Features			
Fisher				Fisher				Fisher			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None			Liver	None	
Liver	10	19	29	Liver	12	17	29	Liver	17	12	29
None	16	18	34	None	13	21	34	None	17	17	34
Totals	26	37	63	Totals	25	38	63	Totals	34	29	63
QDC				QDC				QDC			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None			Liver	None	
Liver	12	17	29	Liver	15	14	29	Liver	15	14	29
None	16	18	34	None	14	20	34	None	14	20	34
Totals	28	35	63	Totals	29	34	63	Totals	29	34	63
MOG2				MOG2				MOG2			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None			Liver	None	
Liver	15	14	29	Liver	9	20	29	Liver	29	0	29
None	15	19	34	None	15	19	34	None	34	0	34
Totals	30	33	63	Totals	24	39	63	Totals	63	0	63
MOG2+PCA				MOG2+PCA				MOG2+PCA			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None			Liver	None	
Liver	15	14	29	Liver	14	15	29	Liver	12	17	29
None	15	19	34	None	12	22	34	None	8	26	34
Totals	30	33	63	Totals	26	37	63	Totals	20	43	63
ANN 1(3)				ANN 1(3)				ANN 1(3)			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None			Liver	None	
Liver	25	4	29	Liver	14	15	29	Liver	14	15	29
None	32	2	34	None	18	16	34	None	19	15	34
Totals	57	6	63	Totals	32	31	63	Totals	33	30	63

The results are very weak and akin to a random classifier. From these results one can conclude that either the directly extracted features do not describe the signals well enough, or that the signals themselves are not representative of the underlying tissue/organs.

4.4.3.2 Autoregressive Models

As stated in Section 4.3.2.1 the following permutations were considered for models: (i) varying the number of possible poles and zeros from 6 to 15; (ii) considering various signal pre-processing filters or none; and (iii) varying the time delay until an optimum fit for the possible setup was found. The resulting fitness percentages were placed in histograms and appropriate thresholds for the required fitness percentages were set for each combination (pre-filter and number of poles and zeros) – see Figure 4.14. It was decided to ignore the data of very bad fits as these would taint the classification results.

The chosen thresholds and the percentage of data points that qualify for each combination of filter and number of poles and zeros are summarized in Table 4.9.

As it is impossible to discuss and review each permutation with each of the evaluated classifiers, it was decided to perform a single evaluation on each combination with three of the classifiers – (i) Fisher’s linear discriminant analysis (FISHER); (ii) Quadratic Bayes assuming Gaussian densities (QDC); (iii) a Mixture of three Gaussians per class (MOG2). The point that yields the best results is used in further classifier evaluation – the results are summarised in Table 4.10. The percentages in these results are the correct classification results and the ones with a greater rate than 60 % have been highlighted. There is no overlapping between the evaluations, but the pre-filter #8 (300Hz 6th order LPF) and 15 poles/zeros combination yields almost 60 % correct classification in each evaluation. This combination is used in subsequent evaluations.

Table 4.9: Percentage of ARX data points used

		Percentage of data points used									
		#Poles/Zeros									
		6	7	8	9	10	11	12	13	14	15
Pre-filter	1	84.13%	96.83%	96.83%	87.30%	89.68%	90.48%	78.57%	80.95%	82.54%	81.75%
	2	83.33%	85.71%	89.68%	88.10%	92.06%	89.68%	92.86%	91.27%	93.65%	92.86%
	3	81.75%	84.13%	88.10%	88.89%	88.89%	89.68%	91.27%	90.48%	92.06%	90.48%
	4	79.37%	78.57%	84.13%	86.51%	89.68%	88.10%	89.68%	90.48%	91.27%	90.48%
	5	77.78%	75.40%	83.33%	82.54%	88.10%	84.92%	89.68%	88.89%	89.68%	89.68%
	6	74.60%	72.22%	82.54%	77.78%	86.51%	84.13%	88.10%	86.51%	88.89%	88.10%
	7	73.02%	66.67%	81.75%	74.60%	84.92%	80.95%	87.30%	84.92%	88.10%	87.30%
	8	69.84%	62.70%	80.16%	80.95%	84.92%	84.92%	89.68%	89.68%	91.27%	89.68%
	9	71.43%	60.32%	79.37%	76.98%	88.89%	86.51%	92.06%	88.89%	92.86%	91.27%
		Threshold = 40%			Threshold = 50%			Threshold = 60%			

To visualise the data in at least some fashion, a PCA mapping from thirty to two features were trained on the selected data set. The result is shown in Figure 4.32. From this depiction one can see that some points are clustered together, which might make the problem slightly separable, however, a mass of different classes are collected at the origin – these will be difficult to separate.

Only three classifier topologies were considered (the Mixture of Gaussian training ran into singularities, even after PCA) – (i) Fisher’s linear discriminant analysis (FISHER); (ii) Quadratic Bayes assuming Gaussian densities (QDC); and (iii) an Artificial Neural Network (ANN) with one hidden layer of three nodes. In each case a cross-validation process was followed with four folds for (i) and (ii) and three folds for (iii), as well as single 50 % training, 50 % test data confusion matrix results.

The confusion matrix results are summarised in the tables of Table 4.12. The cross validation results can be summarised as follow:

- (i) FISHER – 59 %
- (ii) QDC – 55 %
- (iii) ANN 1(3) – 55 %




Figure 4.32: PCA of a ARX data set

Table 4.10: Preliminary evaluations on ARX data

FISHER results											
		#Poles/Zeros									
		6	7	8	9	10	11	12	13	14	15
Pre-filter	1	45.28%	47.54%	59.02%	58.18%	62.83%	52.63%	55.56%	47.06%	32.69%	51.46%
	2	56.19%	50.00%	52.21%	47.75%	60.34%	48.67%	52.14%	42.61%	44.07%	48.72%
	3	59.22%	50.94%	49.55%	60.71%	48.21%	55.75%	42.61%	54.39%	60.34%	45.61%
	4	56.00%	57.58%	56.60%	54.13%	46.90%	60.36%	57.52%	61.40%	44.35%	47.37%
	5	48.98%	45.26%	60.00%	51.92%	56.76%	58.88%	50.44%	60.71%	53.98%	53.98%
	6	51.06%	45.05%	55.77%	55.10%	65.14%	62.26%	54.95%	44.95%	55.36%	53.15%
	7	56.52%	45.24%	47.57%	51.06%	45.79%	52.94%	40.00%	53.27%	53.15%	52.73%
	8	52.27%	46.84%	58.42%	45.10%	43.93%	60.75%	52.21%	57.52%	54.78%	62.83%
	9	46.67%	42.11%	46.00%	48.45%	53.57%	57.80%	41.38%	44.64%	57.26%	49.57%
>60% correct											
QDC results											
		#Poles/Zeros									
		6	7	8	9	10	11	12	13	14	15
Pre-filter	1	45.28%	55.74%	55.74%	50.91%	59.29%	52.63%	47.47%	41.18%	50.00%	45.63%
	2	73.33%	64.81%	66.37%	58.56%	36.21%	46.90%	45.30%	54.78%	44.07%	60.68%
	3	53.40%	58.49%	60.36%	53.57%	50.00%	53.98%	58.26%	49.12%	53.45%	49.12%
	4	58.00%	47.47%	54.72%	63.30%	50.44%	53.15%	61.06%	40.35%	49.57%	54.39%
	5	53.06%	53.68%	56.19%	48.08%	44.14%	40.19%	52.21%	53.57%	61.06%	57.52%
	6	51.06%	36.26%	67.31%	61.22%	54.13%	58.49%	54.95%	50.46%	57.14%	53.15%
	7	52.17%	64.29%	51.46%	53.19%	51.40%	58.82%	52.73%	55.14%	53.15%	52.73%
	8	50.00%	54.43%	50.50%	45.10%	57.01%	45.79%	55.75%	53.98%	47.83%	59.29%
	9	62.22%	42.11%	54.00%	54.64%	55.36%	44.95%	34.48%	55.36%	50.43%	51.30%
>60% correct											
MOG3 results											
		#Poles/Zeros									
		6	7	8	9	10	11	12	13	14	15
Pre-filter	1	47.17%	54.10%	0.00%	63.64%	45.13%	47.37%	51.52%	54.90%	57.69%	53.40%
	2	0.00%	0.00%	50.44%	0.00%	50.00%	0.00%	65.81%	49.57%	61.02%	48.72%
	3	0.00%	56.60%	54.95%	51.79%	46.43%	62.83%	53.04%	59.65%	48.28%	52.63%
	4	0.00%	0.00%	0.00%	0.00%	57.52%	51.35%	55.75%	0.00%	42.61%	54.39%
	5	0.00%	53.68%	44.76%	0.00%	0.00%	51.40%	59.29%	42.86%	46.90%	52.21%
	6	0.00%	0.00%	0.00%	51.02%	50.46%	62.26%	47.75%	0.00%	44.64%	0.00%
	7	56.52%	0.00%	57.28%	51.06%	45.79%	43.14%	49.09%	55.14%	47.75%	49.09%
	8	0.00%	0.00%	0.00%	0.00%	57.01%	0.00%	0.00%	46.90%	63.48%	62.83%
	9	0.00%	0.00%	0.00%	0.00%	60.71%	0.00%	56.90%	55.36%	52.14%	0.00%
>60% correct											
Singularity problem											

Table 4.11: Confusion matrix results on mechanical parameter data

True Labels	Fisher		Totals	True Labels	QDC		Totals
	Estimated Labels				Estimated Labels		
	Liver	None			Liver	None	
Liver	13	13	29	Liver	6	20	26
None	10	20	34	None	13	17	30
Totals	23	33	63	Totals	19	37	56

True Labels	ANN 1(3)		Totals
	Estimated Labels		
	Liver	None	
Liver	8	18	26
None	6	24	30
Totals	14	42	56

The ARX classifier evaluation results are very weak for all three considered classifier structures. From these results one can conclude that either the ARX parameters extracted do not describe the signals well enough, or that the signals themselves are not representative of the underlying tissue/organs.

4.4.3.3 Mechanical Model

In the process of establishing the parameters of mechanical models as discussed in Section 4.3.2.2 only one combination of features were determined for each data point. As such, classifier evaluation is simpler than in the previous two sections.

As was done in Section 4.4.3.2 only data where the fit was greater than a specific threshold (40 %) were used in the evaluation and data was also split into 50 % for training and 50 % for testing. 59 out of 126 points qualified as far as fitness is concerned resulting in a 29/30 split between training and test data.

A PCA mapping from 7 to 2 dimensions result in the scatter plot shown in Figure 4.33. There may be some clustering of the data, but there is no clear trend.




Figure 4.33: PCA of a mechanical model data set

Five classifier topologies were considered (as was done in the previous sections) – (i) Fisher’s linear discriminant analysis (FISHER); (ii) Quadratic Bayes assuming Gaussian densities (QDC); (iii) a Mixture of two Gaussians per class (MOG2) using all features; (iv) a MOG2 after PCA mapping features to three dimensions; and (v) an Artificial Neural Network (ANN) with one hidden layer of three nodes. MOG3 was initially considered, but the training algorithm ran into matrix singularities which caused errors. In each case a cross-validation process was followed with four folds for (i) and (ii) and three folds for (iii-v), as well as single 50 % training, 50 % test data confusion matrix results.

The confusion matrix results are summarised in the tables of Table 4.12. The cross validation results can be summarised as follow (values indicate error rates):

- (i) FISHER – 56 %
- (ii) QDC – 58 %
- (iii) MOG3 – 73 %
- (iv) MOG3+PCA – 59 %
- (v) ANN 1(3) – 54 %

Table 4.12: Confusion matrix results on mechanical parameter data

Fisher				QDC			
True Labels	Estimated Labels		Totals	True Labels	Estimated Labels		Totals
	Liver	None			Liver	None	
Liver	9	6	15	Liver	7	8	15
None	8	6	14	None	8	6	14
Totals	17	12	29	Totals	15	14	29

MOG2			
True Labels	Estimated Labels		Totals
	Liver	None	
Liver	8	7	15
None	5	9	14
Totals	13	16	29

MOG2+PCA				ANN 1(3)			
True Labels	Estimated Labels		Totals	Labels	Estimated Labels		Totals
	Liver	None			Liver	None	
Liver	4	11	15	Liver	10	5	15
None	8	6	14	None	6	8	14
Totals	12	17	29	Totals	16	13	29

From the three feature extraction methods considered, the mechanical model seems to yield the best result. The MOG2 yielded the best result which indicates that there may be some clustering in the data. Using an MOG2 classifier should yield 70 % correct classification, should the data be matched to a two stage mass-damper-system and the parameters that best fit the data used for classification. The practicality of this solution is questionable as the Genetic Algorithm takes long to find an optimum fit.

4.5 Conclusions

A test procedure was developed to mechanically perform percussion on human volunteers and the tools developed in Chapter 3 were used to capture the acceleration of the skin next to the percussion site. Three human volunteers' data was captured and these were used for further analysis.

The presence of the liver was identified using MRI to find the percussion points that were close to the liver or not. These results were stored as ground truth and subsequent evaluations were targeted at identifying points that belong to the presence of the liver or not.

The collected data was processed in various ways to extract parameters that describe the signals. The processing methods included direct extraction from signals and model fitting techniques. The latter method used predefined models (autoregressive and mass-spring-damper models) and tuned the model parameters to fit the data – the model parameters that best fit the data were used for further evaluation. The signals from some data points could not be matched to models with great confidence (the source of the error was not evaluated), and these signals were ignored in further evaluations.

Various classifier topologies were used to attempt to match the data to underlying ground truth. Three of the classifier topologies (Mixture of Gaussians, Quadratic Bayesian Discriminant, and Fisher Linear Discriminant) assume some underlying probability density and trains itself accordingly, where an Artificial Neural Network simply tunes internal parameters to best separate classes. Classification results were calculated using cross-validation techniques and varied between 40 % and 73 % correct classification rates, with most of the results under 60 %. These results can be compared to a system that simply assigns classes at random – attaining a correct classification result as the prior probability of any point being a liver. The best results were attained using the mass-spring-damper parameter description and a Mixture of Gaussians classifier with two clusters (73 %).

Even though these tests were performed using a small data set, some correlation between features and the presence of the liver should have been apparent with most of the methods. The final method using a mechanical model to describe the system shows some promise, but the author is sceptic. It is most probable that the test procedure is flawed, that the selected hardware is suboptimal, or even that other parameter extraction and classification methods should be considered. It may also be possible that the problem is simply too complex and other factors, like the presence of bone or fat, are throwing signals out. The main concerns as far as the test procedure is concerned are the following:

- What is the effect of the actuator's offset force on the free vibration of the body wall – it may be possible that the actuator is simply damping the signal so that no *reaction* is measured, but only the input.
- The input angle and offset forces could not be controlled accurately during tests – it may be that these factors/parameters need to be controlled.

It was decided to perform an evaluation using a much simpler and controlled setup. This is the discussion of the next chapter. A mock-up of the human abdomen was created using silicone and an anomaly (hole) inside the silicone would simulate an underlying organ. The following investigations were to be performed:

- The possibility of using other actuator signals to identify the underlying system.
- What is the effect of distance between the actuator and accelerometer, and can one place the accelerometer close to the actuator?
- Can one find an underlying anomaly in such a simple setup?
- What is the effect of the offset force of the actuator?
- Does the angle of the actuator head influence the resulting signals?

Chapter 5 Model Tests

The measurement setup used during human subject tests was very specific and the various data processing techniques could not extract information that could predict underlying conditions with sufficient confidence. Some of the decisions that were made in development of the measurement setup for human tests were based on very vague information (e.g. the actuation signal was based on the only stroke possible for a medical practitioner, yet a vibrational instrument can be actuated using various signals) and, in hindsight, it would have been a better idea to create a mock-up of the problem and evaluate possible configurations before proceeding with such a specific measurement setup.

The function and goal of this experiment is therefore to use a mock-up of the problem (a simple model of the human body with some anomaly (or phantom) within the model to represent the liver) to gain more information on how different approaches to specific setup parameters (input signal type, input offset force and output measurement position) influence the ability to detect the liver or abnormalities thereof. This approach is similar to that used in elastography (measurement of the elasticity of liver tissue by ultrasound) research [27] where the use of a silicone mass with an anomaly (called a phantom) contained within were used to evaluate the concept.

Unfortunately, this experiment was performed very late in the project and very little time was available. Due the time limit, only the evaluations mentioned here were performed, but other evaluations that could have been done include assessment of the effect of ribs, fat and other organs.

5.1 Experiment Design

Except for the human subjects and the mechanical frame, the same tools were used as in the human subject tests, but some additional development was required:

- Model – a mass of material that would represent the human body and a phantom contained within it to represent the liver. For this purpose silicone rubber compound that closely resembles human tissue (boneless) was acquired and moulded into a flat cylinder. Whilst moulding, a rubber balloon was left in the mixture to leave a hole for simulants. A balloon filled with either (i) water; (ii) air; and (iii) jelly solution were used to simulate different anomaly conditions. The silicon rubber compound called Dragon Skin™ is produced by a company called Smooth-On and an excerpt from the datasheet can be found in Appendix E.2. The model was placed on a wooden base (which was glued to the work bench) with the plastic base of the model pushing against two guiding screws to make sure that it stayed in position when rotating.




Figure 5.1: Silicone model on base

- Clamp to hold the vibrational shaker in place – as the model was not quite the same shape as a human subject, the frame that was developed for human subject tests was not sufficient. The part of the frame that is connected to the vibrational shaker has quite a long pipe protruding from the back and this was simply attached to a clamp that was connected to the table on which experiments were performed.
- A DC power supply was used in series with the shaker amplifier to control the offset force of the vibrational shaker.

The measurement setup is shown in Figure 5.2. An image of the 3D model of the setup can also be found in Appendix C.2 - this was created to assist in the recreation of the experiment should it have been disassembled.

The silicone model has a cross-section and height 261 mm and 100 mm respectively. The circumference of the silicone model was divided into 13 approximately 60 mm intervals (at a specific point of actuation the accelerometers could not be placed closer than 30 mm from the centre of the actuator – this limitation was the motivation for this choice) and at each interval an ECG sticker was placed for connection of an accelerometer. Halfway between each pair of accelerometers is a designated point of actuation. Each input and output was numbered on the silicone model either with a circle around the number (output) or not (input) and this is shown in Figure 5.3.




Figure 5.2: Phantom tests measurement setup




Figure 5.3: Numbering of input/outputs on silicone model

The dimensions of the model along with the phantom hole's dimensions were used to generate a plot of the various elements of the models as seen from above. This representation is shown in Figure 5.4.




Figure 5.4: Model Dimensions

5.2 Test Procedure

Three different phantoms were used: (i) water filled balloons; (ii) air; and (iii) balloons filled with a “jelly” solution.

The silicone model with the phantom was rotated so that each of the 13 input positions was actuated. In each case the offset voltage was adjusted such that the offset force measured by the FSR (using the datasheet’s calibration curve – see Appendix A.4), without any actuation signal, read 1 N (Newton) – this value was never verified, but it was kept consistent throughout measurements as to negate the influence of differing offset force.

Only two accelerometers were available for these measurements, so the accelerometers were placed at two of the thirteen outputs. A sequence consisted of actuation by nine pulses; 10 s of white noise; and a chirp signal from 1 Hz to 1 kHz ramping at 25 Hz/s. The accelerometer signals, FSR signals and actuator input signals were stored automatically for each actuation type and the storage filenames contained the phantom type, the actuation type, the input position and the positions of the two accelerometers (e.g. Water_Chirp_Actuate_01_ACC1_10_ACC2_11.mat would be the file that contains the signals recorded when the phantom was a water filled balloon, the actuator was in position 1 and accelerometers 1 and 2 were placed at positions 10 and 11 respectively). After each sequence, the accelerometers were moved to the next set of outputs and the sequence repeated. Once all outputs for a particular input was covered, the silicone model was moved to the next actuation position.

All data collection was done with the *DataLoggerGUI.m* Matlab program that was written to perform all the required functions (as stated in Section 3.6, the computer generates output signals, sets up the DAQ device, records signals, and stores results with appropriate filenames). The GUI provides the user with controls to record where the actuator head is placed, where the accelerometers are placed, to display the offset force and control which type of actuation should be performed (or that a sequence should be performed). During a sequence, all data is saved with the appropriate filenames automatically, whereas manual selection of actuation would require manual saving, but with the default filename generated unless otherwise specified.

For each particular phantom type, two full datasets were recorded at different times to verify the repeatability of tests.

In addition to the aforementioned tests, multiple tests were performed at specific input and output combinations, but with the variation of input force and input angle. These tests should verify how critical these conditions are in obtaining repeatable results.

5.3 Evaluations

The goal of this experiment is to gain insight into the effect of the following factors/parameters whilst attempting to identify the location of and material in the anomaly:

- Actuation signal – in the human subject tests only pulses were used (to determine an impulse response), but the use of a chirp or white noise signal might yield better results in estimating an impulse response.
- Position of actuation and measurement – in human subject tests, measurement was done in very close proximity to the input, but it may be that there are no differentiating features as the actuator dampens the natural response.
- Finding the phantom or anomaly – in human subject tests features were extracted using various intricate methods, but there may be clear features in signals that indicate the position of the anomaly.
- Actuation factors – the influence of offset force and input angle could not be controlled accurately during human subject tests, and the sensitivity to these parameters may be of interest.
- Repeatability – during the human subject tests, each subject was only subjected to a single test, thus the repeatability of results could not be established. Multiple tests were performed on the silicone model and the repeatability can therefore be established for this experiment.
- Differentiability between underlying consistencies – can one identify the slight differences between water, air and jelly, if hidden in the silicone model? This would give clues as to how easy it would be to differentiate between healthy and diseased liver, should the elasticity of the liver be influenced by the disease.

Each of these evaluations is discussed in the following sections.

5.3.1 Actuation Signal for Impulse Estimation

If an ideal impulse could be delivered to a system, the impulse response should fully describe the system's transfer function – this theory resides in the theory of the Fourier Transform which can be found in many texts, but is cited from [61] for the purposes of this report.

Unfortunately the pulses that are delivered during pulse actuation are far from ideal and only represent an approximation of an impulse. Chirp and white noise actuation methods, on the other hand, can estimate an impulse response by use of cross correlation of the input and output signal according to Peebles [62]. The theory behind impulse responses and the use of noise and chirp signals to estimate them is quite involved, but an impulse can be estimated using equation 6 as extracted from Peebles [62], where N_0 is the noise power, R_{XY} is the cross-correlation between input x and output y and $h(\tau)$ is the impulse response estimate. A simple theoretical example that uses equation 6 is discussed next.

$$\hat{h}(\tau) \approx \left(\frac{2}{N_0} \right) \hat{R}_{XY}(\tau) \approx h(\tau) \quad (6)$$

Consider a system with the Laplace transfer function defined by $H(s)$, where

$$H(s) = \frac{1000}{s^2 + 20s + 1000} = A \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (7)$$

$$f_c = \frac{1}{2\pi} \omega_c = \frac{1}{2\pi} \sqrt{1000} \approx 5\text{Hz}$$

$$A = 1$$

$$\zeta = \frac{10}{\omega_n} = \frac{10}{\sqrt{1000}} = 0.31625$$

$H(s)$ therefore represents a 2 pole, 0dB low-pass filter with cut-off frequency at approximately 5 Hz and damping coefficient of 0.3, therefore being over-damped and should present with some oscillation.

This system is simulated with three different inputs: (i) a chirp signal that ramps from 0.1 to 500 Hz in 20 s; (ii) 20 s white noise; and (iii) a 20 s pulse train with 20 ms pulses spaced 1s apart. The simulation is done using Matlab's Simulink at a fixed sampling rate of 8 kHz. The block diagram and resulting signals are shown in Figure 5.5. To complicate matters, measurement noise with amplitude of approximately $1/10^{\text{th}}$ of the measured signals was added at the output of each signal.

The inputs and outputs of the systems are cross-correlated. The resulting waveforms for the noise method, chirp method and the measured pulse method is compared to what is returned when the *impulse* function in Matlab is used (the function estimates an ideal impulse response), in Figure 5.6. It shows that the chirp and noise methods should produce similar impulse responses to the pulse train and also that these methods are far more resilient to measurement noise in obtaining an impulse response. Additionally, it seems that the pulse train method obtains a phase lagged response which is not present in the other two methods – this is due to the high frequency harmonics of the sharp edge of the pulse that cannot be processed by the slow system.

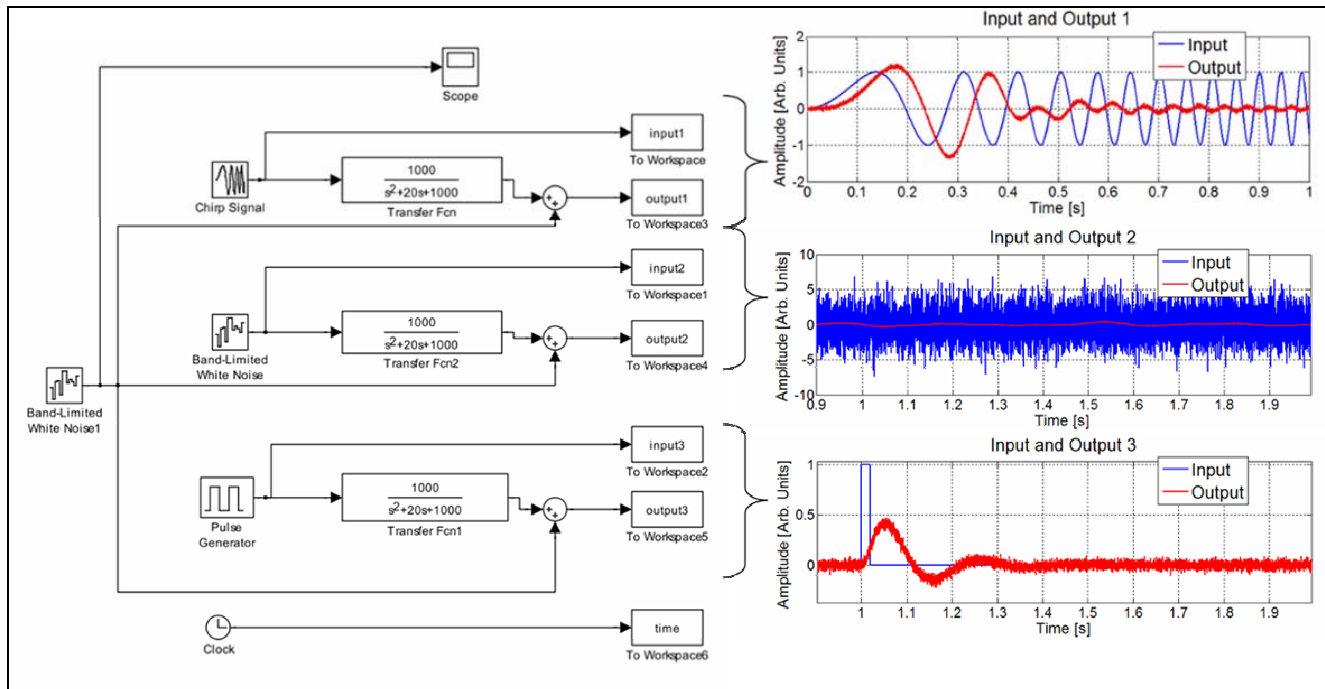


Figure 5.5: Impulse response estimation methods

The estimated bode plots of these three methods are compared to that obtained by Matlab’s *bode* function in Figure 5.7 – the bode plot of the pulse method result does not match the theoretical nearly as well as those of the other methods. The higher frequency phase responses of both the noise and chirp methods seem to deviate from the theoretical, more so for the white noise estimate. The white noise method also seems to generate results with a weaker signal-to-noise ratio (SNR). Although it is difficult to determine SNR for a decaying signal a Root-Mean-Squared (RMS) amplitude was determined for the expected signal and the difference signals for each case (noise) and used to calculate $SNR_{dB} = 20\log(A_{signal}/A_{noise})$: Chirp 33dB; Noise 30dB; Pulse 19dB. The SNR is therefore best for the Chirp method as far as noise is concerned, but only slightly better than the white noise method.




Figure 5.6: Comparison of impulse response estimation methods

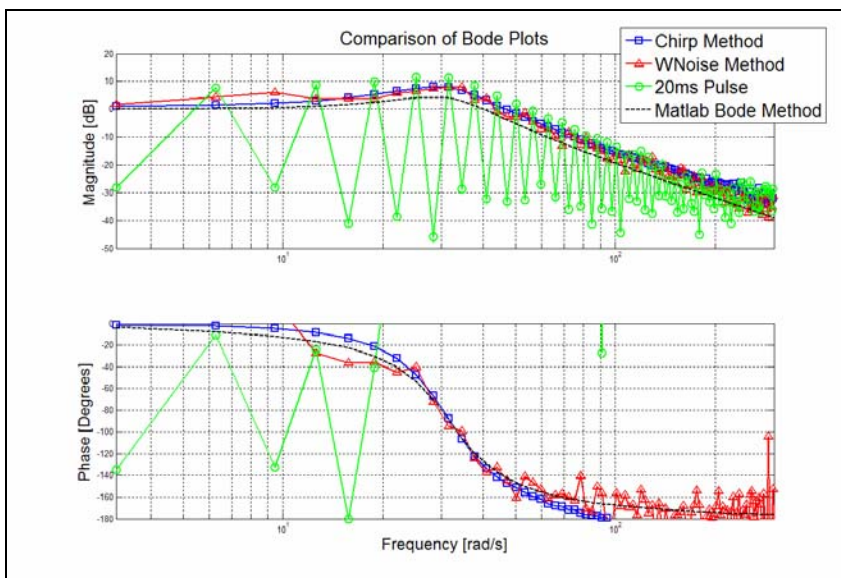


Figure 5.7: Comparison of impulse estimation methods - bode plots

All three of these actuation methods were used for each input output combination during data collection as it was still unclear which method would present the best results.

Figure 5.8 and Figure 5.9 show the time domain and frequency domain signals of the three methods for input at point 10 and output at point 5 with air as the phantom. The time domain signals look quite similar, but the frequency domain (amplitude spectrum) signals show that the pulse method has a reduced amount of high frequency information – this is due to the non-idealism and therefore reduced bandwidth of the input pulse.

The noise and chirp methods inherently have the correct time information and all evaluations are synchronised. The pulse method requires the matching of time alignment of input signals, which might induce an error and subsequently result in erroneous phase information. No comment can be made about the comparison of the chirp and noise amplitude spectra as the true response is not known, but it is possible that the same effect seen in the theoretical example is present – the noise estimate contains some induced noise.




Figure 5.8: Time domain signals and estimations from model




Figure 5.9: Frequency domain signals and estimations from model

The theoretical and practical evaluations show that the chirp and noise method provides better impulse response estimation than the impulse response generated by a non-ideal input pulse. The chirp estimate contains less noise and is used in most of the further evaluations.

Another evaluation that needs to be performed is testing the coherence between the input and output. Coherence is a statistical measure of the causality between input and output signals and is based on the assumption that signals are ergodic (statistical properties can be deduced from the sample length provided) and the system is linear [63]. It uses cross and auto spectral densities to determine the fractional part of the output signal power that is produced by the input at that frequency. Coherence (γ^2) is calculated by equation 8, where S_{XY} is the cross power spectral density between input X and output Y and S_{XX} and S_{YY} represent auto power spectral densities.

$$\gamma^2 = \frac{S_{XY}S_{YX}}{S_{XX}S_{YY}} \quad (8)$$

The coherence spectrum should be equal to 1 unless there is noise in the input or output signals, the system has non-linearities or the output is due to additional inputs besides the measured input. The same signals that were used to produce the model impulse estimates were used to determine their coherence and the result is shown in Figure 5.10.




Figure 5.10: Coherence between input and output

The coherence is less than ideal, but at least higher where the signal content is located. The setup was mostly immune to noise or other sounds which means that the output is not likely being caused by other input signals. The system is likely non-linear and/or the input and output signals contain noise. Further investigation of this kind was not performed, but should be considered in further studies or investigations that stem from this thesis.

5.3.2 Position of Actuation and Measurement

The question that is addressed in this section is ‘where should one actuate relative to a measurement point (or vice versa)?’ This is answered by looking at signals for different combinations and seeing whether there are differentiating features – if none, the combination is a bad choice; if some, the differences may be indicative of the phantom location (but this is the topic of the next section).




Figure 5.11: Input-output combinations

As stated earlier, the model was evaluated with thirteen possible inputs and thirteen possible outputs which make it difficult to gain insight into how signals compare when looking at them individually.

To compare multiple signals, they were depicted next to each other in an image, with the x-axis being the time/frequency axis and the y-axis being the relevant signals. Additionally, looking at each input and output combination, for a particular offset (e.g. input and output five positions to the left), one can compare the signal to parameters inherent to that combination – like “distance from input to anomaly”, “distance of signal path through anomaly”, etc. – these parameters can be extracted from the depiction in Figure 5.11. This can give clues as to which parameters in the signal are indicative of the location of the phantom.

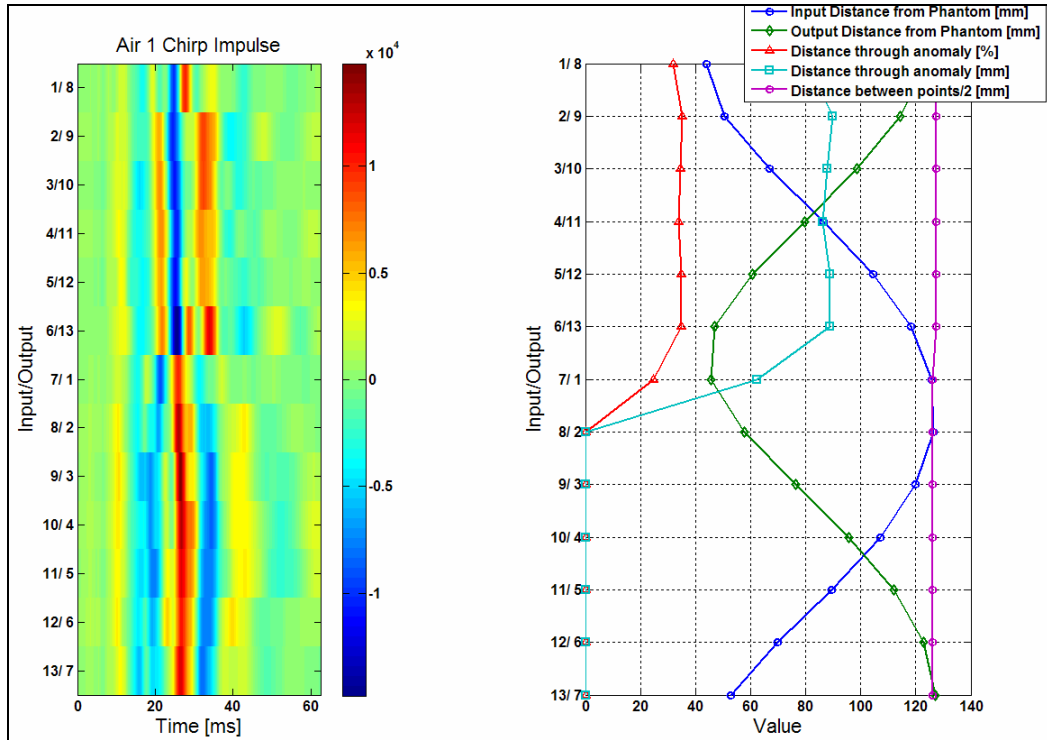


Figure 5.12: Comparing signals to model (Air Phantom)

An example of this depiction is shown in Figure 5.12 – with air as the phantom and employing the chirp method, the time domain impulse estimates are shown for the input-output combinations stated on the y-axis. The inputs combined with outputs seven positions to the left are used in this depiction. The signal at y equal to “1/8” is shown in Figure 5.13.




Figure 5.13: A signal trace from image depiction

From Figure 5.12, one can clearly see that there are signal variations as one moves inputs and outputs around the silicone model. One can even go as far as pointing these out – e.g. where the path crosses the anomaly, the peak around 27 ms is negative, otherwise it is positive. What one can conclude is that this particular combination shows differences for signals at different points relative to the anomaly.

A similar analysis was performed for each combination (offset) and the following conclusions can be made:

- With the output at the opposite side of the input (offset of six positions – remember that the positions are pre-numbered and output 1 is left of input 1) there exists some differentiation between signals, but it is unclear how anything correlates with the location of the anomaly.
- When the offset is between five positions to the left or right (offset smaller than four or larger than eight) a discrepancy is detectable as a peak time difference. Output 13 and input 1 has a slightly longer distance between them, subsequently some combinations of input/output points have a longer distance. The peak time correlates perfectly with the distance between input and output points and not with the presence of an anomaly close to points or in the path – an example is shown in Figure 5.14. Directly next to the input (offset of zero or 12), no differentiation is visible, which does not bode well for the method used in human subject tests where this was the case.

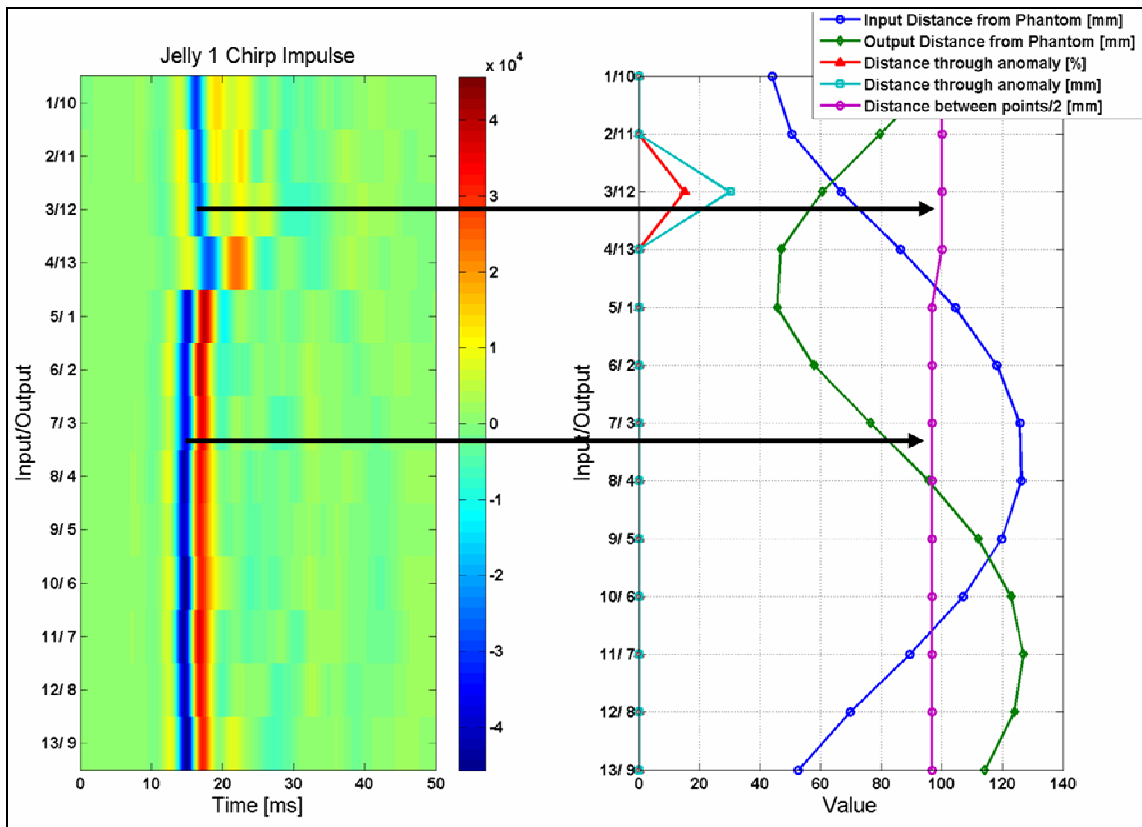


Figure 5.14: Correlation between distance and peak time

- The two remaining combinations (offset of five or seven positions) show promise. The signals all have a peak around 23-27 ms, and the peak is negative where the anomaly is present between points and positive otherwise.

Unfortunately, one cannot show each depiction, but the case with the output left of the input (zero offset) and jelly as the phantom is shown in Figure 5.15.

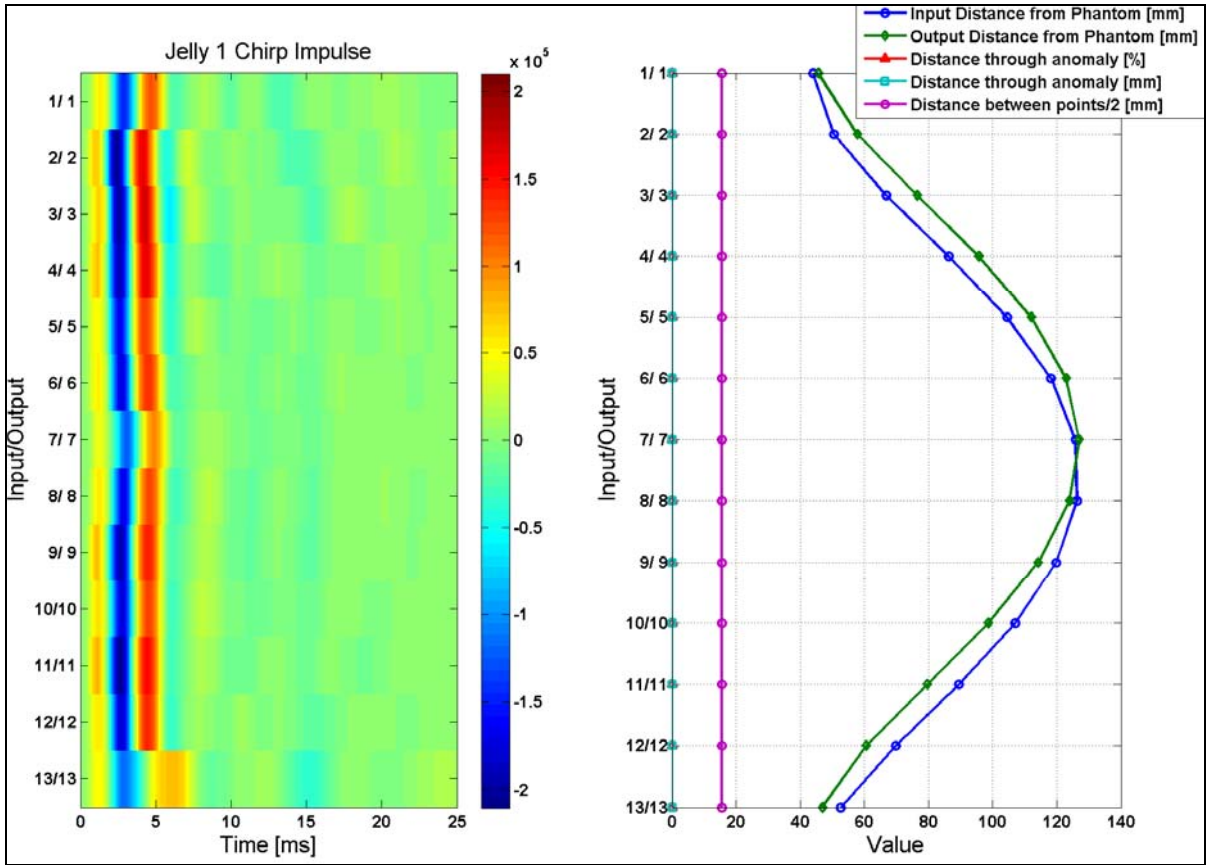


Figure 5.15: Signals with jelly phantom with zero offset

5.3.3 Finding the Phantom

The point of this evaluation is to identify where the anomaly is located. In the previous section it was shown that very little distinguishing features are visible when measuring close to the input; and some distinguishing features are visible when measuring directly opposite the input, but the features are not well correlated with the location of the anomaly. Signals at an offset of five or seven points showed one clear feature and some less clear features correlated with the presence of an anomaly between points. Other signal combinations show a strong correlation between the distance between points and the signal's peak time, as expected, and some unclear correlation between features in the signal and the presence of the anomaly.

The signal combinations at five and seven offset positions all seem to present with a peak close to 25 ms, the sign of which is correlated with the presence of the anomaly in the path. Furthermore, the amplitude without an anomaly in the path has a strong peak at 60 Hz, otherwise it is distorted with a wider standard deviation (should the amplitude spectrum be seen as a density spectrum).

The following features are therefore extracted from these signal combinations:

- Peak time (t_p)
- Peak Amplitude (A_p)
- Peak at 60 Hz in amplitude spectrum ($A_{60\text{Hz}}$)

- Standard deviation of amplitude spectrum ($f_{\text{std dev}}$)




Figure 5.16: Example of extracting peak time and amplitude




Figure 5.17: Example of extracting $f_{\text{std dev}}$ and $A_{60\text{Hz}}$

The extracted features can be plotted on an axis parallel to the underlying features to visually determine the correlation and the result for an offset of five positions is shown in Figure 5.18. A_p is clearly correlated with the presence of the anomaly and the $A_{60\text{Hz}}$ seems to have some correlation, but the other two features do not show any correlation. A similar result is obtained for an offset of seven positions.

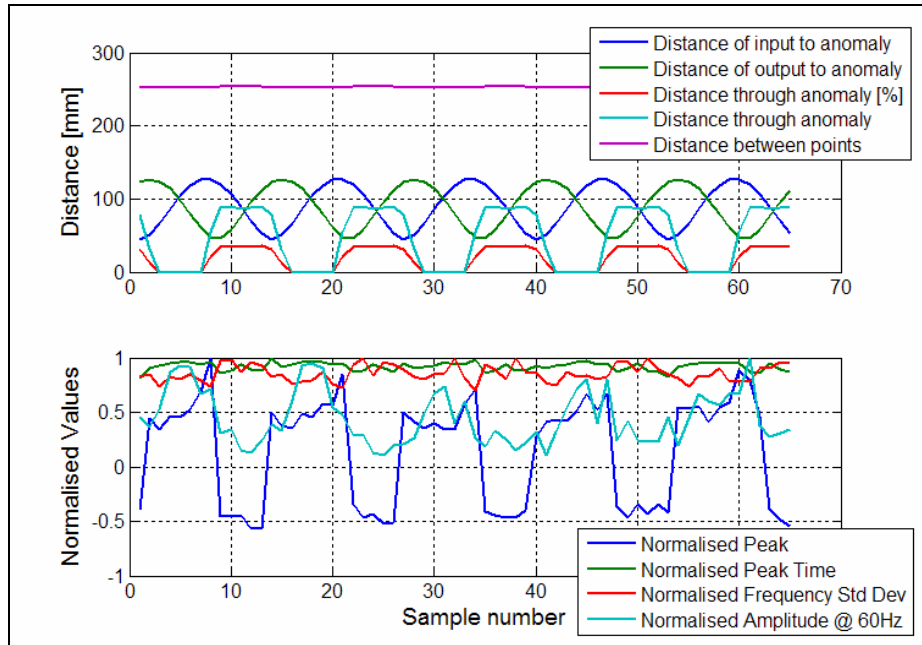


Figure 5.18: Comparing extracted features to silicone model parameters (offset of five positions)

These two features can be plotted against the distance through the anomaly to better emphasize the correlation – shown in Figure 5.20. It should therefore be possible to reconstruct an image showing the possible location of the underlying anomaly by looking at these two parameters – if the normalised amplitude at 60 Hz is smaller than 0.5 there may be an anomaly between points; or if the amplitude close to 25 ms is negative there may be an anomaly in the path. The use of these two methods is illustrated using the data obtained with a jelly anomaly and an offset of five positions between input and output, shown in Figure 5.19 – black lines indicate the presence of an anomaly between input and output and red lines not. Both features seem to predict the location of the anomaly with reasonable accuracy for the specific data set. A similar evaluation was performed on all available data and the use of $A_{60\text{Hz}}$ works for 60 % of the data, where A_p works for 90 % of the data.

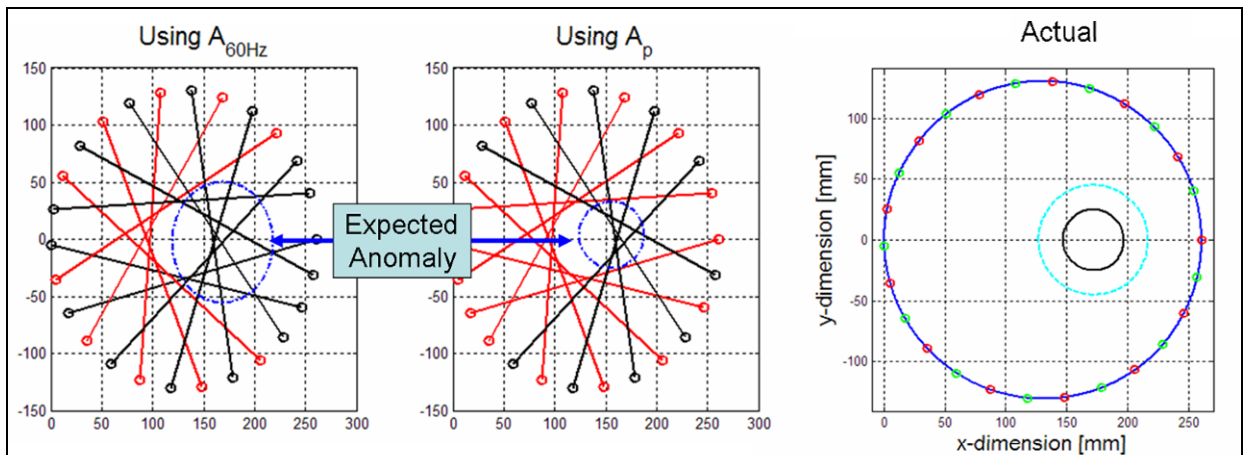


Figure 5.19: Predicting location of anomaly using extracted parameters




Figure 5.20: Scatter plot of features vs. disrupted distance

It is therefore concluded that, for this specific setup, one can predict the location of an anomaly with reasonable accuracy if one uses a chirp signal as an input at all of the 13 positions and measure the output for each input at an offset of five or seven positions. Measuring directly opposite the input does not yield usable results, nor does measuring next to the input.

5.3.4 Repeatability of Procedure

Repeatability is of importance to this evaluation as one would want to know whether data can be trusted. The previous section inadvertently evaluated repeatability by extracting the same parameters from signals that were recorded at different times and using the same setup. A different measure will be employed in this section – goodness of fit expressed as a percentage as expressed in equation 9, where y is the first signal and y_h the second.

$$Fit(\%) = 100 \times (1 - \sqrt{\sum (y - y_h)^2 / \sum (y - \text{mean}(y))^2}) \quad (9)$$

Air was used as a phantom in two separate experiments and these experiments can be compared for each signal combination, using the first as a base and the second as comparison. Table 5.1 summarises the results for each combination. The average fit is equal to 76.68 %, which is adequate. Some combinations show very bad fits, but looking at the actual comparisons show that these signals actually still do compare very well – see Figure 5.21.

Table 5.1: Fitness measures for repeatability measurement

		Output Position												
		1	2	3	4	5	6	7	8	9	10	11	12	13
Input Position	1	84.95	78.10	37.07	62.53	68.96	67.11	19.69	47.71	36.71	73.44	69.60	64.71	72.20
	2	72.21	72.09	68.81	65.00	71.34	51.39	64.61	54.54	70.36	65.39	79.55	73.61	78.55
	3	75.64	55.84	77.74	77.06	67.81	61.42	56.59	75.30	47.70	63.91	75.54	73.01	47.07
	4	63.22	67.62	70.14	79.70	75.94	72.79	69.67	73.06	77.64	35.04	71.71	75.47	76.51
	5	69.06	67.85	53.40	80.78	72.43	76.51	60.71	75.96	72.59	67.82	34.88	72.70	68.86
	6	46.93	61.35	77.81	73.36	77.88	78.09	73.34	60.76	70.39	70.96	73.17	53.09	63.57
	7	46.44	55.20	77.35	83.94	70.33	68.39	78.21	71.53	81.00	72.75	62.58	81.21	26.28
	8	59.97	54.50	77.84	81.59	75.25	72.50	69.86	73.25	81.59	52.85	59.86	65.28	81.58
	9	67.89	44.52	75.30	55.06	72.82	49.79	69.33	79.18	78.25	54.58	65.24	70.88	64.21
	10	39.29	79.81	49.99	54.48	70.29	58.13	63.81	80.58	83.07	75.19	71.83	69.59	74.51
	11	41.53	59.70	71.97	45.25	63.90	80.03	66.95	69.50	81.49	46.17	84.43	63.54	49.28
	12	54.56	64.67	50.67	69.99	22.03	79.82	65.87	81.95	66.85	59.98	55.16	78.75	48.93
	13	22.58	68.73	71.60	67.67	72.27	57.09	59.60	33.07	73.03	44.11	78.98	76.82	76.11




Figure 5.21: Repeatability of some signals and their fits

5.3.5 Effect of Actuation Parameters

There are two factors which could not be controlled during tests on human subjects – the angle of the actuator to the surface of the skin, and the offset force to the skin. It would be difficult to establish the exact effect of these parameters on the signals measured during human subject tests, but the effect on the silicone model tests should at least give some indication as to what the importance of control over these parameters are.

During human subject tests it was difficult to control the input angle of the actuator head as the human body has a fixed shape and the frame has a fixed shape, therefore for each volunteer a different relationship exists and it is almost impossible to have the input perpendicular to the body wall. Measurements were therefore performed with the actuator placed at angles of -45° , -10° , 0° , $+10^\circ$ and $+45^\circ$, of the perpendicular (minus sign refers the head facing output one and plus facing output thirteen). The input was at position two and output at position three (remember that this constitutes a two position offset as output one is after input one) – using the chirp method to estimate the impulse, as well as the pulse method and actual pulses.

The results of variation of the input angle are shown in Figure 5.22 and Figure 5.23. It seems that the amplitude is slightly higher where the actuator is directed toward the output and the phase is slightly delayed as the angle increases. The frequency content remains in tact.




Figure 5.22: Effect of variation of input angle (using Chirp method)




Figure 5.23: Effect of variation of input angle (using Pulse method)

During human subject tests, the offset force could have varied even though the head was moved as to exert as much force as was comfortable to the volunteer. Because this parameter can't be controlled it is important to determine its effect. The offset force was set at 0, 0.5, 1.5 and 2 N and the response calculated using the chirp method, and simply using the pulse response. The results are shown in Figure 5.24 and Figure 5.25.




Figure 5.24: Effect of variation of offset force (using Chirp method)




Figure 5.25: Effect of variation of offset force (using Pulse method)

The input force variation seems to have less influence on the pulse method's pulses, and more on the chirp method estimate. The input signal recorded during the chirp method gets distorted by the offset voltage used to generate the offset force – as such, the impulse estimates may be incorrect. The phase of the pulse method's pulses seems to be influenced slightly by the change in force, but the signal contents seem to remain in tact.

From these results one can conclude that the input force and angle does not influence signals extracted from the silicone model too much. Whether these results can be extrapolated to human subjects is debatable.

5.3.6 Differentiating Anomaly Types

Of interest to the research would be whether one could distinguish between healthy and diseased liver (where the disease influences the elasticity of the liver). The tests on the silicone model therefore involved the use of various phantoms to simulate different anomalies, just as the liver would present differently inside the human body.

Tests were performed with water, air and jelly as the anomaly, with two tests performed with water and air and a single test using jelly. Signals with the same input and output combination should ideally be the same for the same anomaly and differ between different anomalies. Should the signal differ between the same anomalies that signal is not dependable; otherwise, any differences seen between different anomalies are dependable.

Signals were compared by looking at both their time and frequency representations. All 13 by 13 combinations were looked at, and it was found that some signals from the second *water* set were very abnormal – it is believed that a problem crept in whilst the data was being captured, possibly the accelerometers were not connected properly. Fortunately some signal combination did present proper results (*air* signals matching each other and *water* signals matching each other) and a small subset of these showed some difference between the anomaly types. Using the combination of input at position 2 and output at position 10 (which means the anomaly was between the points) yields the results shown in Figure 5.26. The time signal (on the right) does not show clear discrepancies between signals, but the amplitude spectrum does – the *water* signals have their strongest components at 140 Hz, where the *jelly* signal has strong components at 80 Hz and 140 Hz. The *air* signals are relatively weak at these two frequencies, but have stronger components than the other two at 105 Hz.




Figure 5.26: Identifying the anomaly - visual analysis

The amplitudes at the three identified frequencies can be extracted and plotted pair wise. Only two combinations are shown in Figure 5.27. The discrepancies are clear for these data points and it should be possible to discriminate between anomalies for this simple problem. Unfortunately more data was not captured to verify these results.




Figure 5.27: Identifying the anomaly - scattered data

Evaluations on data points where the inputs and outputs were close to each other, showed no detectable discrepancies, again emphasising the need to measure further from the input when performing similar tests on human subjects.

5.4 Conclusions

The evaluations performed during silicone model tests led to the following conclusions:

- Method for identifying a system – during human subject tests only multiple pulses were used to obtain an impulse response estimate and subsequently derive unique parameters. The use of a chirp or white noise signal combined with cross-correlation generates almost ideal impulse responses, with the chirp signal generating superior results depending on the required bandwidth.
- The relationship between the input and output measurement positions – from the analyses on the silicone model it seems that almost no information can be deduced from impulse estimates calculated between inputs and outputs in close proximity. Measurements on the model showed that accelerometer readings at a distance of more than 300 mm from the input yield distinguishing results. Converting this measure to human testing has not been done, but it is clear that measuring next to the input will not reveal information about the underlying structure.
- The ability to identify the position of the phantom was demonstrated using only a few parameters extracted from subsets of the data. The parameters were extracted from data where the inputs and outputs were far removed further emphasising the need to have a larger distance between these points in human tests.
- Repeatability of the measurement procedure was demonstrated by comparing repeated experiments on the same anomaly and using a measure of fit between subsequent experiments. This shows that such a controlled environment can yield repeated results, but this was in fact also demonstrated by identifying the position using multiple data sets (separate experiments).

- The effect of the variation of offset force and input angle (not controllable in the human test setup) were evaluated on the silicone model and it was found that these variations only slightly influenced the signal obtained using pulsing. This could mean that their effects did not factor in the weak results obtained during human subject tests.
- The different anomalies could be differentiated by looking at the captured data. A point that is emphasised again is that information cannot be detracted from combinations where the input and output is close to each other.

This evaluation showed promise using the tools developed for human subject tests and identified different actuation methods to identify the system between input and output. Simple features are extracted from signals to identify the presence and type of the anomaly, which indicate that it should be possible to do the same on human subjects, depending on the effect of differences in fat presence and ribs. If more time was available, the effect of ribs and fat could also be simulated. One clear error that was made during human subject tests, however, is that the measurement was done too close to the actuator input.

Chapter 6 Conclusions and Recommendations

The research that has been reported shows that emulating percussion is difficult at best. Except for the lack of understanding about the underlying mechanics, there exists no guideline as to how to go about emulating the technique. It was concluded that the act of percussion uses some mechanical input pulse and the physician uses his/her hearing and experience to identify the underlying tissue.

The author investigated possible ways to emulate percussion. Subsequently, a design employing an electromagnetic actuator and accelerometer was motivated and built. The function of the accelerometer would be to listen to the body wall vibration and the actuator would deliver the input pulse. Additionally hardware was developed to collect data and perform experiments. The sensors that were used were not standard, calibrated laboratory equipment and their performances were shown to be lacking, but whether the lack in performance led to the lack of success could not be validated.

Upon completion of the system design, it was employed to collect data to develop signal processing methods to identify the borders of the liver from sensor data. The data collection method evolved with the assistance of human volunteers and only three volunteers were used in the final test setup – presenting the research with a very small data set. The method involved the use of successive pulses at various locations on a volunteer's body and measuring the acceleration of the skin next to the point of actuation.

To identify signals that indicate the presence of the liver, various signal processing and parameter extraction methods were considered (a small subset of all possible approaches). Each method involves some simplification of the data by extraction of parameters and a method for matching parameters to underlying truth (liver or no liver). Each method showed promise for some data, but not for all. Signal parameters could not be extracted with sufficient confidence, but even those parameters that were extracted with confidence could not successfully identify the presence of the liver. The method that showed the most promise involved matching signals to the parameters of a dual mass-spring-damper system and training a Mixture of Gaussians classifier with two clusters using these parameters – a 70% correct classification was shown. It is possible that such a result was an abnormality as most of the other techniques resulted in less than 60% correct classification.

These results compelled further investigation using a much simpler setup – a silicone model with an anomaly contained within, also using the same tools as were used in the human tests. The questions that needed to be answered were (i) are there perhaps better actuation methods; (ii) where should measurement take place relative to the position of actuation; (iii) could the position of the anomaly be identified using such a simple setup; (iv) is the procedure repeatable and how do factors like offset force and input angle influence results; and finally (v) could the physical properties of the anomaly be identified by such a process?

The tests on the silicone model revealed the following:

- i. A more involved method including actuation using a chirp or noise signal and cross-correlation could yield better results for defining a system.
- ii. There should be a large distance between the point of actuation and measurement. The lack of success in human tests can perhaps be attributed to this finding.
- iii. The position of the anomaly contained within the silicone model could be identified with some confidence – it may be possible to do the same in human testing and implementation, but the effect of other factors like ribs, other organs and fat cannot be predicted, nor could it be evaluated in the thesis' time frame.
- iv. The procedure is highly repeatable in such a controlled environment and factors like offset force and input angle influences signals only slightly.
- v. Signal features can be used to differentiate between anomaly types (water, jelly and air), but only when measuring at a distance from the input and with the anomaly between the input and output.

From this subsequent evaluation it does seem that the measurement location should at least be further from the input and that other actuation signals should be considered.

The outcome of this thesis is thus the following:

- A system and methods were developed to emulate the act of percussion as used to identify the border of the liver. This system and method was, however, shown to be invalid as it could not successfully identify the liver boundary.
- Further investigation in sensor performances may reveal that it also impacted on the negative outcome of this project.
- The methods that were used to process signals to perhaps find indicators of the liver are quite advanced, even though they could not be validated. They may still be useful in future evaluations.
- The source of error was investigated in subsequent tests and definite problems in the initial approach were pointed out.
- Although the thesis did not yield a positive result, the lessons that were learnt do contribute to the understanding of the problem and should not be ignored.

In the event of the investigation being taken further, the following should be considered:

- It is still unclear whether one would be able to extract information using the approach used and a more familiar starting point should perhaps be considered – percussion should be performed by physicians whilst the reactionary sounds are recorded using a microphone. The physician's remarks regarding the notes should be compared to the signal traces measured and it should be established whether it is possible to discern between percussion sounds.
- Once confidence has been established in how sounds differ from a signal processing point of view one can proceed to design a system that classifies resulting sounds automatically and establish such a system's performance or validity. A validated tool like this could contribute vastly to the understanding of percussion and could also be implemented as a learning tool for future physicians.

- Only after showing that there is promise in emulating percussion, one can systematically attempt to automate or mechanise the physician's functions, but these changes should happen systematically, validating each step. During these developments, the use of standard, calibrated laboratory equipment is strongly suggested. The use of non-standard untested sensors should only be implemented once validated against standard, calibrated laboratory equipment.

The human as a learning machine should not be underestimated – computer learning have proven superior in cases where a problem or process is well understood and the human is ill adept at taking various parameters into consideration. The technique of percussion, however, was developed by humans and for humans over centuries whilst never being well understood.

List of References

- [1] "Percussion" in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [2] S.L. Adams, P.R. Yarnold, and J.J. Mathews, "Clinical use of the olecranon-manubrium percussion sign in shoulder trauma," *Annals of Emergency Medicine*, vol. 17, May. 1988, pp. 484-7.
- [3] J.R. Guarino, "Auscultatory percussion of the head.," *British Medical Journal (Clinical research ed.)*, vol. 284, Apr. 1982, pp. 1075-1077.
- [4] J.J. Fedorowski, "Medical Percussion," *Hospital physician*, vol. 26, Sep. 1990, pp. 31-36.
- [5] E.L. DeGowin, *DeGowin & DeGowin's bedside diagnostic examination*, Macmillan, 1987.
- [6] "Clipart of Hands placed on a shoulder showing proper position and action to perform mediate percussion," *Fotosearch.com*, Jan 5, 1999. [Online]. Available: <http://www.fotosearch.com/LIF139/nu120011/>. [Accessed: Sep. 20, 2008].
- [7] "Liver," *howstuffworks.com*, Oct 3, 2005. [Online]. Available: <http://healthguide.howstuffworks.com/liver-picture.htm>. [Accessed: Sep. 20, 2008].
- [8] S.R. McGee, "Percussion and physical diagnosis: separating myth from science," *Disease-a-month : DM*, vol. 41, Oct. 1995, pp. 641-92.
- [9] J.C. Yernault and A.B. Bohadana, "Chest percussion," *European Respiratory Journal*, vol. 8, Oct. 1995, pp. 1756-1760.
- [10] G.B. Risse, "Pierre A. Piorry (1794-1879), the French "master of percussion"," *Chest*, vol. 60, Nov. 1971, pp. 484-8.
- [11] "A brief history of the Practice of Percussion.." *Medical Antiques Online*, Unknown Publish Date, [Online]., Available: <http://www.antiquemed.com/percus.html>. [Accessed: Sep. 17, 2008].
- [12] J.R. Guarino, "Auscultatory percussion. A new aid in the examination of the chest," *The Journal of the Kansas Medical Society*, vol. 75, Jun. 1974, pp. 193-4.
- [13] J.R. Guarino, "Auscultatory percussion to detect ascites," *The New England Journal of Medicine*, vol. 315, Dec. 1986, pp. 1555-6.
- [14] J.R. Guarino and J.C. Guarino, "Auscultatory percussion: a simple method to detect pleural effusion," *Journal of General Internal Medicine*, vol. 9, Feb. 1994, pp. 71-4.
- [15] S. Kalantri, R. Joshi, T. Lokhande, A. Singh, M. Morgan, J.M. Colford, and M. Pai, "Accuracy and reliability of physical signs in the diagnosis of pleural effusion," *Respiratory Medicine*, vol. 101, Mar. 2007, pp. 431-8.
- [16] A. Sakula, "Joseph Skoda 1805-81: a centenary tribute to a pioneer of thoracic medicine.," *Thorax*, vol. 36, Jun. 1981.

- [17] A. Murray and J. Neilson, "Diagnostic percussion sounds: 1. A qualitative analysis," *Medical and Biological Engineering and Computing*, vol. 13, Jan. 1975, pp. 19-28.
- [18] A. Murray and J. Neilson, "Diagnostic percussion sounds: 2. Computer-automated parameter measurement for quantitative analysis," *Medical and Biological Engineering and Computing*, vol. 13, Jan. 1975, pp. 29-39.
- [19] A. Murray and J. Neilson, "Diagnostic percussion: a study of chest-wall motion and the associated tactile sensation," *Medical and Biological Engineering and Computing*, vol. 16, May. 1978, pp. 269-273.
- [20] J.D. Sapira and D.L. Williamson, "How big is the normal liver?," *Archives of Internal Medicine*, vol. 139, Sep. 1979, pp. 971-973.
- [21] S. Sullivan, N. Krasner, and R. Williams, "The clinical estimation of liver size: a comparison of techniques and an analysis of the source of error.," *British Medical Journal*, vol. 2, Oct. 1976, pp. 1042-1043.
- [22] H.C. Burger, G. Castelyn, and F.L.J. Jordan, "How is percussion done?," *Acta Medica Scandinavica*, vol. 142, Feb. 1952, pp. 108-12.
- [23] W. Coleman, "The alleged dullness over the apex of the normal right lung," *American Journal of the Medical Sciences*, vol. 197, Feb. 1939, p. 145.
- [24] B. Skrainka, J. Stahlhut, C.L. Fulbeck, F. Knight, R.A. Holmes, and J.H. Butt, "Measuring liver span. Bedside examination versus ultrasound and scintiscan," *Journal of Clinical Gastroenterology*, vol. 8, Jun. 1986, pp. 267-70.
- [25] W. Peternel, J. Schaefer, and L. Schiff, "Clinical evaluation of liver size and hepatic scintiscan," *Digestive Diseases and Sciences*, vol. 11, May. 1966, pp. 346-350.
- [26] J. Naftalis and C. Leevy, "Clinical estimation of liver size," *Digestive Diseases and Sciences*, vol. 8, Mar. 1963, pp. 236-243.
- [27] L. Sandrin, M. Tanter, J. Gennisson, S. Catheline, and M. Fink, "Shear elasticity probe for soft tissues with 1-D transient elastography," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 49, 2002, pp. 436-446.
- [28] J. Ophir, B. Garra, F. Kallel, E. Konofagou, T. Krouskop, R. Righetti, and T. Varghese, "Elastographic imaging," *Ultrasound in Medicine & Biology*, vol. 26 Suppl 1, May. 2000, pp. S23-9.
- [29] J. Ophir, K. Alam, B. Garra, F. Kallel, E. Konofagou, and T. Varghese, "Elastography: imaging of tissue elastic properties in vivo," *Ultrasonics Symposium, 1998. Proceedings*. p. 1743 vol.2.
- [30] A.B. Bohadana and S.S. Kraman, "Consistency of sternal percussion performed manually and with mechanical thumper," *European Respiratory Journal*, vol. 5, Sep. 1992, pp. 1004-1008.
- [31] H. Mansy, T. Royston, and R. Sandler, "Use of abdominal percussion for pneumoperitoneum detection," *Medical and Biological Engineering and Computing*, vol. 40, Jul. 2002, pp. 439-446.
- [32] H. Mansy, R. Sandler, T. Royston, and D. Jones, "Testing sensors for body surface vibration measurements," *Proceedings of the First Joint BMES/EMBS Conference*, 1999, p. 800 vol.2.
- [33] F. Hawker, *The Liver*, Bailliere Tindall, 1993.

- [34] "Liver," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [35] F.H. Netter, *Atlas of Human Anatomy*, Icon Learning Systems, 1997.
- [36] "Reticuloendothelial System," in *The Free Dictionary*, Feb. 12, 2009. [Online]. Available: <http://www.thefreedictionary.com/reticuloendothelial+system>. [Accessed: Aug. 14, 2008].
- [37] C.M. Leevy, *Diseases of the Liver and Biliary Tract*, Year Book Medical Pub, 1977.
- [38] "Hepatitis," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [39] "Cirrhosis," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [40] "Cholestatic Jaundice," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [41] D.C. Dugdale III, "Cholestasis," in *MedlinePlus Medical Encyclopedia*, [U.S. National Library of Medicine](http://www.nlm.nih.gov/medlineplus/), Bethesda, US, 2008. [Online]. Available: <http://www.nlm.nih.gov/medlineplus/ency/article/000215.htm>. [Accessed: Jul. 20, 2008].
- [42] M.H. Beers, *The Merck Manual 18th Edition*, Merck, 2006.
- [43] "Fatty liver," in *Wikipedia – The Free Encyclopedia*, Wikipedia 2008. [Online], Available: http://www.wikipedia.org/wiki/Fatty_liver. [Accessed: Jul. 20, 2008].
- [44] S. Chopra, *Disorders of the Liver*, Lea & Febiger, 1987.
- [45] "Liver cancer," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [46] "Jaundice," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [47] "Ultrasound," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [48] "Single Photon Emission Computed Tomography," in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [49] "Biopsy" in *Encyclopaedia Britannica Online Academic Edition*, Encyclopaedia Britannica 2009. [Online]. Available: <http://www.search.eb.com.ez.sun.ac.za/>. [Accessed: Feb. 12, 2009].
- [50] Permanent Magnet Shakers Catalogue, *Permanent magnet vibration testing systems*, Ling Dynamic Systems LDS Ltd., Hertfordshire, England.
- [51] Interlink Electronics, "Force Sensing Resistor (FSR)", *Interlink Electronics*, 2006. [Online]. Available: http://www.interlinkelectronics.com/force_sensors/technologies/fsr.html. [Accessed: Nov. 27, 2008].
- [52] Analog Devices Technical Staff, *ADXL330*, Analog Devices Inc., Norwood, MA, US, Mar. 2006.

- [53] Measurement Computing Technical Staff, *PMD-1608FS User's Guide*, Measurement Computing Corp., Middleboro, MA, US, Mar. 2004.
- [54] L. Ljung, *System Identification Toolbox 7 - User's Guide*, The MathWorks Inc., Natick, MA, US, Mar. 2007.
- [55] M.A. Abramson, *Genetic Algorithm and Direct Search Toolbox 2 - User's Guide*, The MathWorks Inc., Natick, MA, US, Mar. 2007.
- [56] T. Weise, *Global Optimization Algorithms – Theory and Application*, University of Kassel, Kassel, Germany, 2008. [E-Book]. Available: <http://www.it-weise.de/projects/book.pdf>. [Accessed: Jul. 13, 2008].
- [57] F. van der Heijden, R. Duin, D. de Ridder, and D.M.J. Tax, *Classification, Parameter Estimation and State Estimation: An Engineering Approach Using MATLAB*, Wiley, 2004.
- [58] P. Paclik, "PRTools: The Matlab Toolbox for Pattern Recognition," Delft Pattern Recognition Group, Oct., 2008. [Online]. Available: <http://www.prtools.org/index.html>. [Accessed: Oct. 20, 2008].
- [59] I.T. Jolliffe, *Principal Component Analysis*, Springer, 2002.
- [60] R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*, Wiley-Interscience, 2000.
- [61] S. Riedel and J. Nillsson, *Electric Circuits*, New Jersey: Prentice Hall, 2000.
- [62] P.Z. Peebles Jr. , *Probability, random variables, and random signal principles*, McGraw-Hill, Inc., 2001.
- [63] P. S. Heyns, "Mechanical Vibrations", *Postgraduate Course Notes*, University of Pretoria 2007.
- [64] Maxim Technical Staff, *MAX4530-4532*, Maxim Integrated Products, Sunnyvale, CA, US, 1996.
- [65] D.A. Neamen, *Electronic Circuit Analysis and Design*, New York: McGraw-Hill, 2001.
- [66] NASA, *Man-Systems Integration Standards (MSIS)*, NASA, Houston, TX, US, Jul. 1995. [Online]. Available: <http://msis.jsc.nasa.gov/>. [Accessed: Aug. 27, 2008].

Appendix A Electronic Designs

A.1 Design of Signal Conditioning and Routing PCB

The purpose of the Signal Conditioning and Routing PCB is to route and pre-condition accelerometer signals to the DAQ device (PMB-1608FS USB [53]).

A.1.1 Requirements

The expected input signals have a DC component and are at levels between 0~3.3V, whereas the DAQ device has a 16-bit resolution over ± 5 V.

The accelerometers require power, which is also to be routed via this board so that they can connect only to one port.

Furthermore, this circuit is designed with the capability to provide power to up to three 3-axis accelerometers (ADXL330 [52]) which means nine channels of information. The DAQ device has eight analogue input channels of which some are allocated for other functions; therefore only three channels are demarcated for accelerometer signals.

Requirements are therefore as follow:

- Be able to switch between 4 configurations: 3 axes of one accelerometer (of which there are 3) or the 3 z-axes of the accelerometers.
- Remove DC component in signals.
- Amplify signals to better use dynamic range.
- Provide power routing to accelerometers.

A.1.2 System Block Diagram

The system block diagram is shown in Figure A-1. There are power inputs for GND, +3.3 V, +5 V and -5 V which are connected to the blocks where these levels are required. Each accelerometer (ACC) connection provides three analogue signals as input which is switched by a multiplexer system (MUX) to three signals in total. Analogue Signal Conditioning blocks remove the DC components in the MUX output signals and amplify them, the result of which is provided through the output connections.

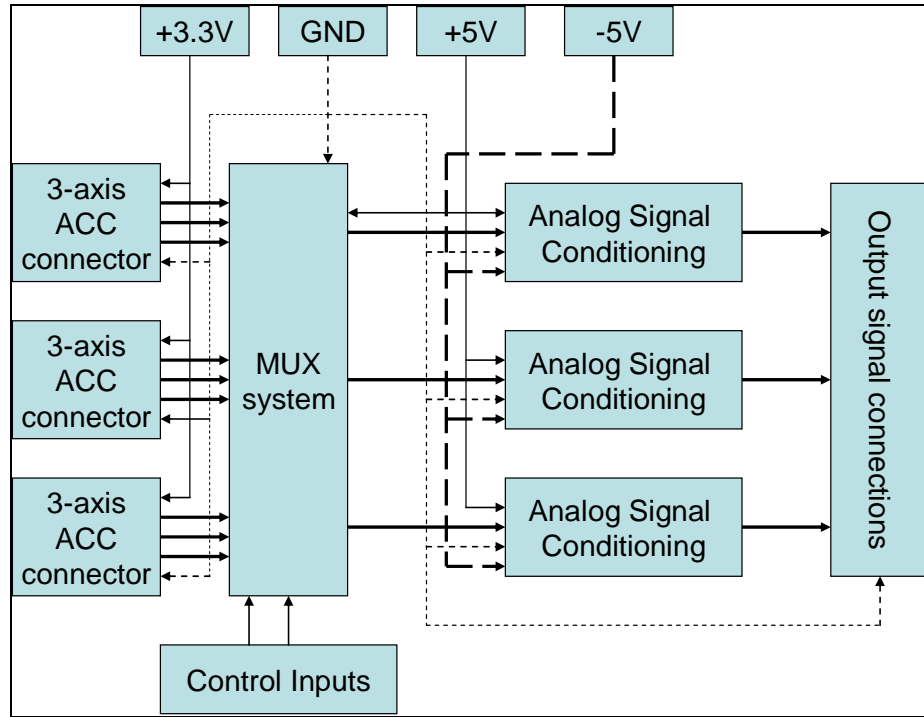


Figure A-1: Signal Conditioning and Routing Block Diagram

A.1.3 MUX system

The MUX system consists of two MAX4531 MUX ICs from Maxim Semiconductor [64]. Figure A-2 shows the device's block diagram. Each IC has two output channels (COMA and COMB), and each output is multiplexed from four inputs (NO0A, NO1A, NO2A, NO3A and NO0B, NO1B, NO2B, NO3B, respectively). Each IC has two address lines (ADDA and ADDB) to determine which input is switched to which output.

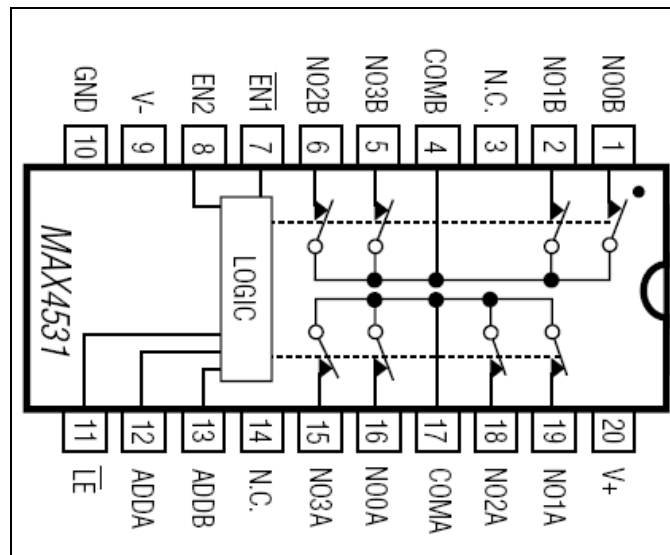


Figure A-2: MAX4531 Block Diagram [64]

Each IC has two output channels (COMA and COMB), and each output is multiplexed from four inputs (NO0A, NO1A, NO2A, NO3A and NO0B, NO1B, NO2B, NO3B, respectively). Each IC has two address lines (ADDA and ADDB) to determine which input is switched to which output. Table A-1 summarizes the truth table for the MAX 4531.

The function of this subsystem is to switch between four possible three-signal-sets:

1. Z0, Z1 and Z2
2. Z0, X0 and Y0
3. Z1, X1 and Y1
4. Z2, X2 and Y2

(e.g. Z0 refers to the Z-axis of accelerometer 0 signal).

To make these combinations possible two address lines are used for both MUXes and both ports of the one and one of the other are used (therefore 4 connections to each port mapping to one of the three output signals).

Table A-1: MAX4531 Truth Table

\overline{LE}	EN2	$\overline{EN1}$	ADDRESS BITS		ON SWITCHES
			ADDB	ADDA	MAX4531
0	1	0	X	X	Last address
X	0	X	X	X	All switches open
X	X	1	X	X	All switches open
1	1	0	0	0	COMA-NO0A, COMB-NO0B
1	1	0	0	1	COMA-NO1A, COMB-NO1B
1	1	0	1	0	COMA-NO2A, COMB-NO2B
1	1	0	1	1	COMA-NO3A, COMB-NO3B

Figure A-3 shows the circuit diagram of the MUX subsystem. As can be seen, the same address lines are used for both devices, labelled ADDA and ADDB. Pins /LE, EN2 and /EN1 are routed to GND, +5 V and GND respectively so that the circuit always operates in accordance to the address lines. Table A-2 summarizes how connections are made depending on address signals from a control device.

Table A-2: MUX system connection table

Address Lines		Connections
ADDB	ADDA	
0	0	Z0-IN1
		Z1-IN2
		Z2-IN3
0	1	Z0-IN1
		X0-IN2

		Y0-IN3
1	0	Z1-IN1
		X1-IN2
		Y1-IN3
1	1	Z2-IN1
		X2-IN2
		Y2-IN3

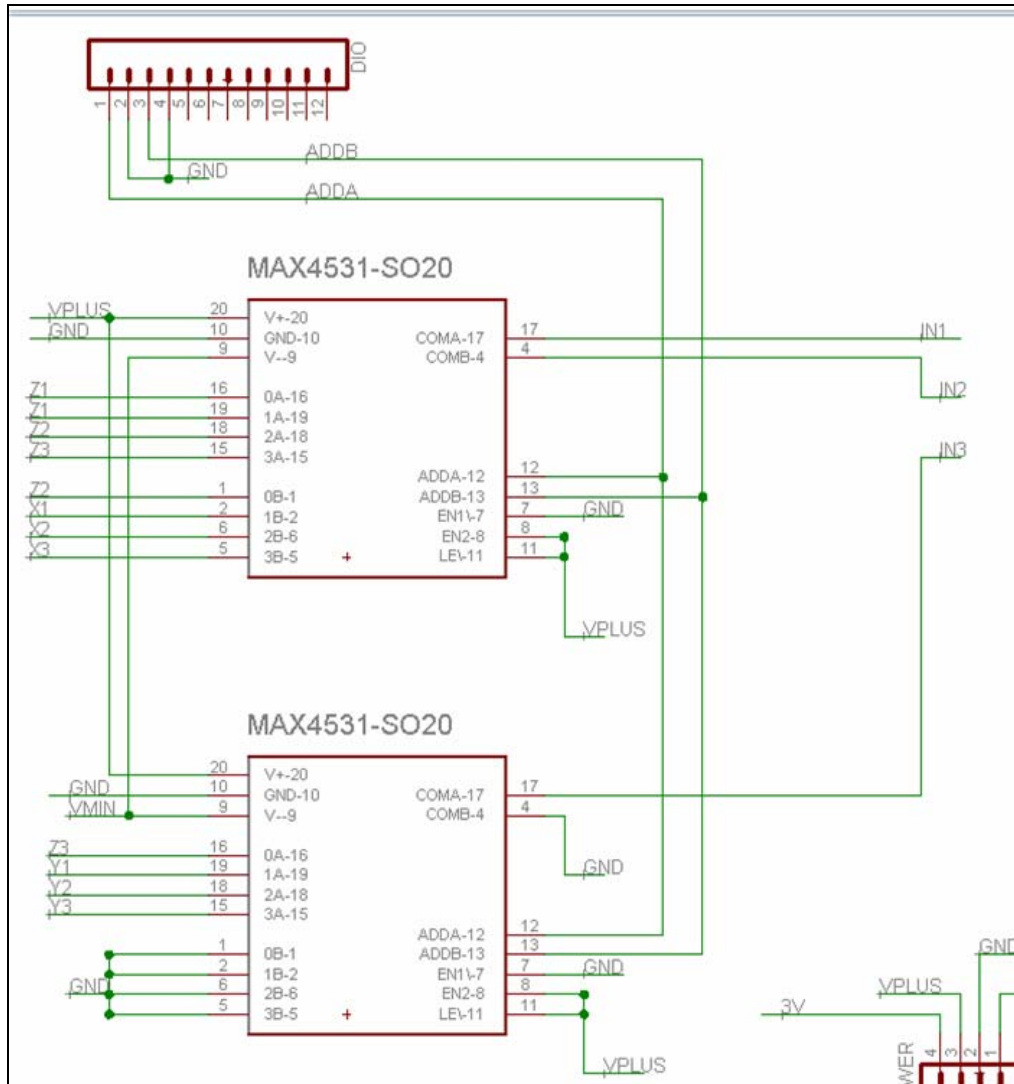


Figure A-3: MUX subsystem circuit diagram

Additionally the necessary power connections are made in accordance with the pin descriptions. The signal path has a resistance of 75 Ω , which is taken into account in the next section. The detailed design drawings in can be found in Appendix A.1.6.

A.1.4 Analogue Signal Conditioning System

The multiplexed signals originating from the accelerometer contain a DC offset voltage due to the orientation of the sensor and the zero acceleration point being offset due to the rails of the device being 0 and 3.3V. As stated earlier, the DC component is meaningless in this analysis, therefore this subsystem needs to remove this component and amplify the signal to make use of the $\pm 5V$ resolution of the DAQ device.

The accelerometer’s circuitry already low-pass filters the acceleration signals with one pole at 1600 Hz (X- and Y-axis) and 550 Hz (Z-axis), therefore the signal conditioning needs no low pass filtering.

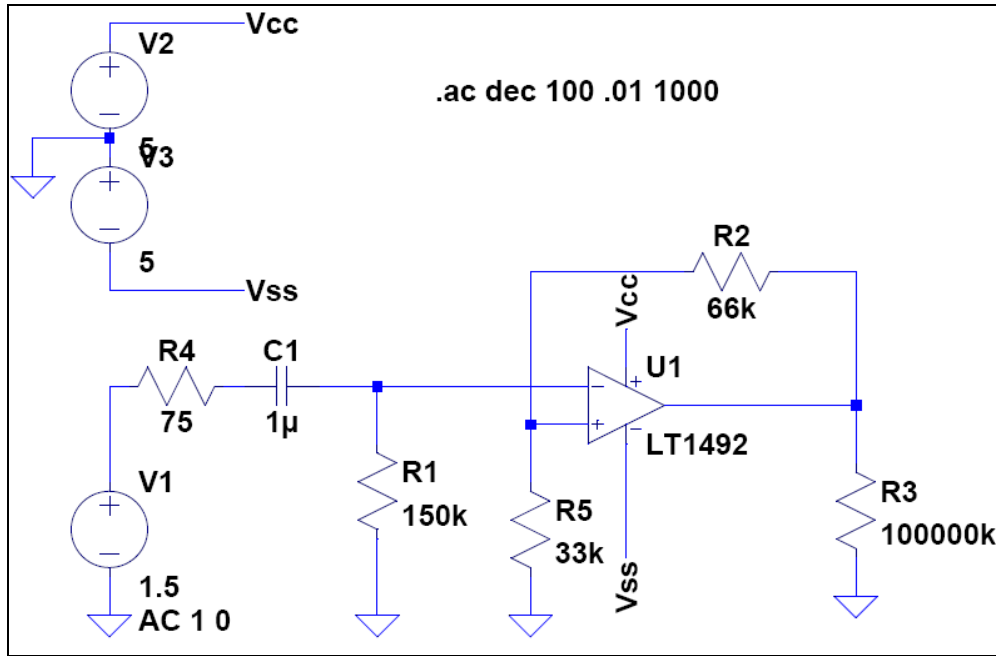


Figure A-4: Signal conditioning sub system simulation model

Figure A-4 shows a model of the suggested and implemented circuit for conditioning. A low voltage operational amplifier from Linear Technology (LT1492) is used in this model simply because of the use of the free simulation software, LTSpice, offered by Linear Technology. V1 simulates the accelerometer signal; R4 simulates the impedance of the MUX system. The combination of C1 and R1 acts as a passive 1-pole high pass filter; whereas R5, R2 and U1 (op-amp) acts as an active 3X voltage amplifier. R3 simulates the input impedance of the DAQ device (100 MΩ). V2 and V3 simply act as the $\pm 5 V$ power rails for the op-amp. The transfer function (in Laplace form) from V1 to V(R3) can be represented by equation 10 below.

$$\frac{V_{out}}{V_{in}}(s) = \left(1 + \frac{R_2}{R_5}\right) \left(\frac{sC_1R_1}{1 + sC_1(R_4 + R_1)}\right) \quad (10)$$

The first part of equation 10 represents the gain of the signal conditioning stage as the second part has unity gain where the signal frequency is higher than the second part's cut-off frequency. The second part of the equation forms

the high pass filter with unity gain in its pass band and -3dB cut-off frequency $f_c = \frac{1}{2\pi R_1 C_1}$ (Because $R_1 \gg R_4$ the effect of R_4 can be ignored). Equation 10 can also be expressed with gain (A) and cut-off frequency (f_c) parameters as is shown in equation 11 below.

$$Gain(f) = A \times \left(\frac{j \cdot f / f_c}{1 + j \cdot f / f_c} \right) \quad (11)$$

Where $A = \left(1 + \frac{R_2}{R_5} \right)$ and denotes the pass-band gain; and j denotes the complex number ($\sqrt{-1}$).

With resistor and capacitor values as shown in Figure A-4, a pass-band gain of 3 V/V or 9.54 dB and a cut-off frequency of approximately 1 Hz are calculated. SPICE simulation of the circuit results in the plot in Figure A-5 where very similar results are seen.




Figure A-5: Signal conditioning amplitude response

The actual circuit was designed and built with LF351 op-amps (these ICs are readily available at the University of Stellenbosch's Electronic Engineering Department in DIP form factor), space for an adjustable potentiometer (for adjustable gain) and fixed resistors (for fixed gain) to determine the pass-band gain, and power noise decoupling capacitors placed close to the op-amps. The final realisation of the circuit schematic is shown in Figure A-6. The detailed design drawings in can be found in Appendix A.1.6.

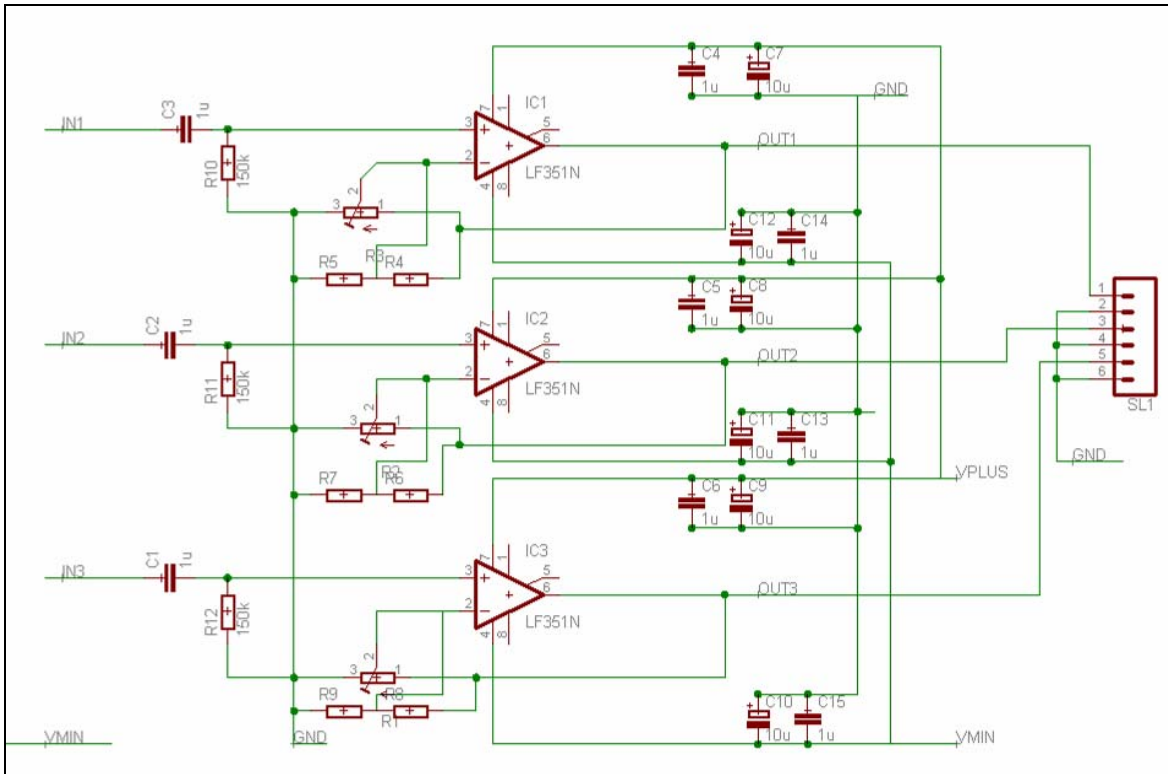


Figure A-6: Signal conditioning sub system circuit diagram

A.1.5 Layout Considerations

All components share the same ground plane which greatly simplifies the layout. In general the following approach was followed:

- Power supply decoupling capacitors were placed as close as possible to the IC they protect.
- Ground planes and connections were kept as large and wide as possible to limit EMI effects.
- Power supply tracks were kept as wide as was sensible.
- Tracks were kept as short as possible – connected devices were placed as close as possible to each other.
- The board size was kept to a minimum; however size was not a constraint.
- Layout was done on a dual (top and bottom) layer board.

A.1.6 Detailed Design Drawings

The Signal Conditioning and Routing PCB's system design was discussed in the previous section. The following figures should help the reader understand how the PCB was realised, but these cannot be used to directly recreate the PCB. The author should be contacted who will supply interested parties with the necessary design files. "EAGLE - Easily Applicable Graphical Layout Editor Version 4.16r1 for Windows Light Edition" was used as design software for the PCB. Figure A-7 and Figure A-8 show the schematics of the circuit. Figure A-9 and Figure A-10 show the top and bottom view of actual board. To give the reader an idea of the actual size of the board, the top edge is 93.2 mm and the left edge is 80 mm in length. The boards were populated and implemented as suggested by the components indicated on the board images.

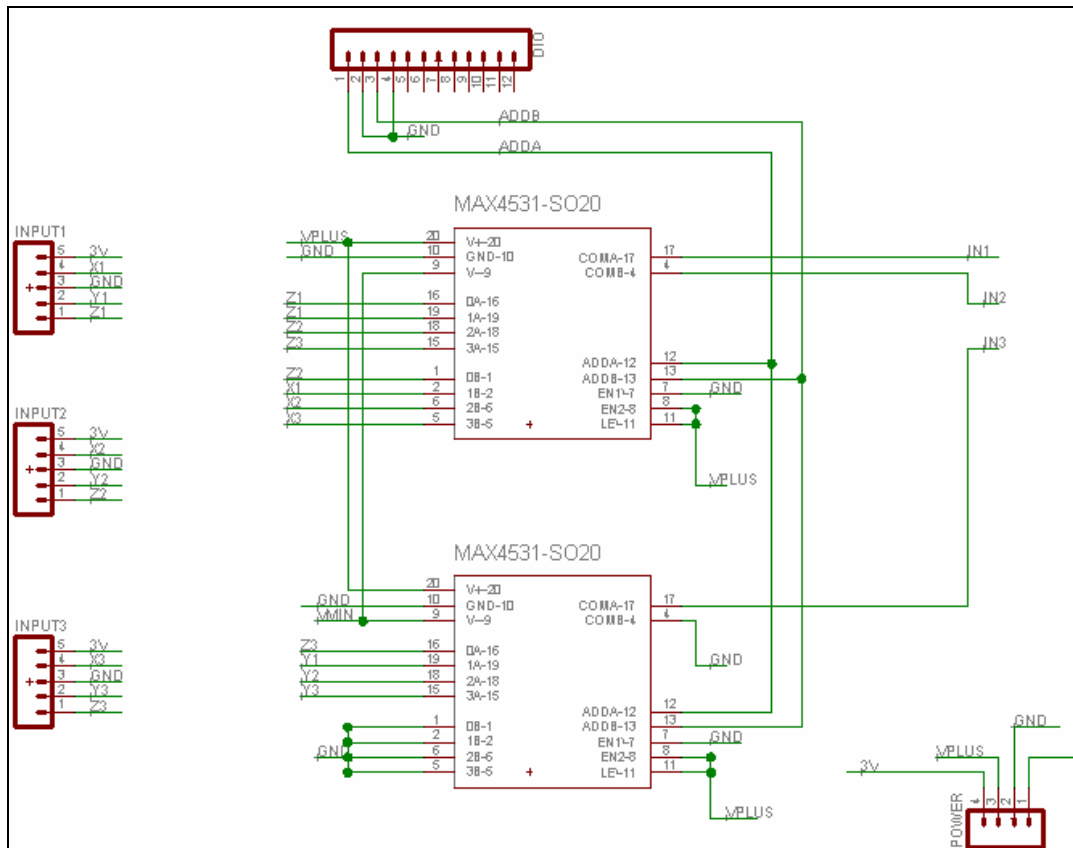


Figure A-7: Signal Conditioning and Router PCB Schematics (part 1)

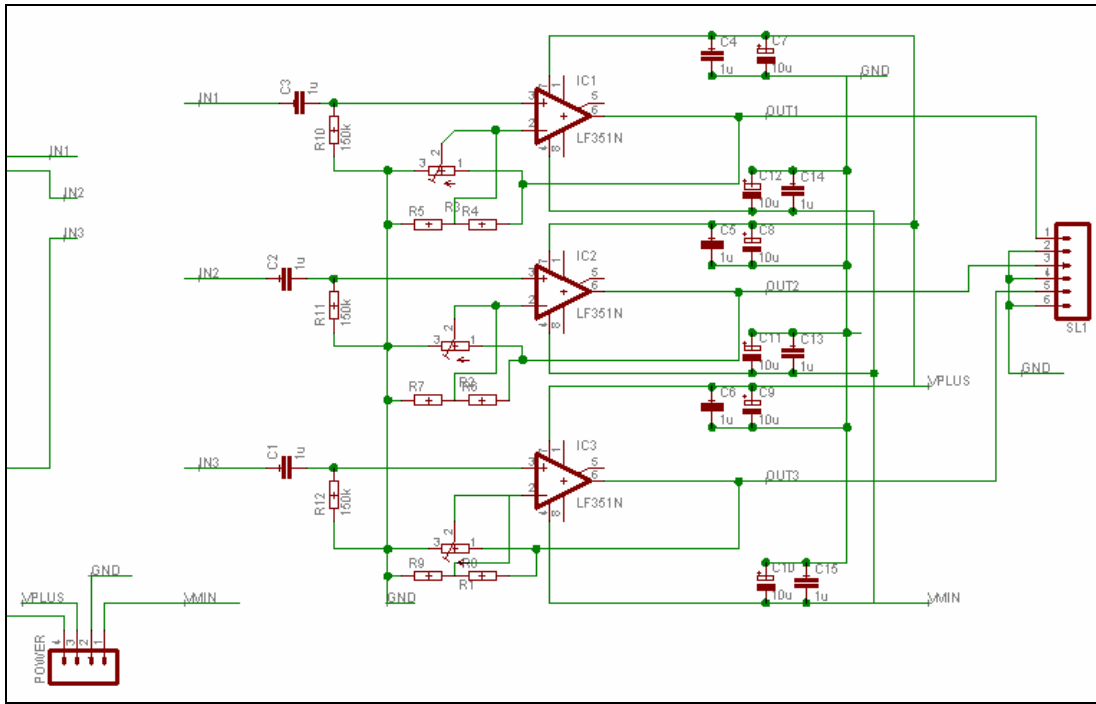


Figure A-8: Signal Conditioning and Router PCB Schematics (part 2)




Figure A-9: Signal Conditioning and Router PCB Board Top




Figure A-10: Signal Conditioning and Router PCB Board Bottom

A.2 Design of Power Source PCB

The function of the Power Source PCB is to provide the various power levels required by different components in the Signal Conditioning and Routing PCB (Appendix A.1) and the Accelerometer PCBs (Appendix A.3). The PCB receives 230 VAC input power and converts it into +3.3 V, -5V and +5 V.

A.2.1 Requirements

The PCB simply needs to have a 230 VAC input and DC +3.3 V, -5 V and +5 V outputs.

A.2.2 System Block Diagram

The system has a five power conditioning subsystems – an AC to DC power conditioning sub system; and the +3.3 V, -5 V and two +5 V regulator blocks/systems. The block diagram is shown in Figure A-11.

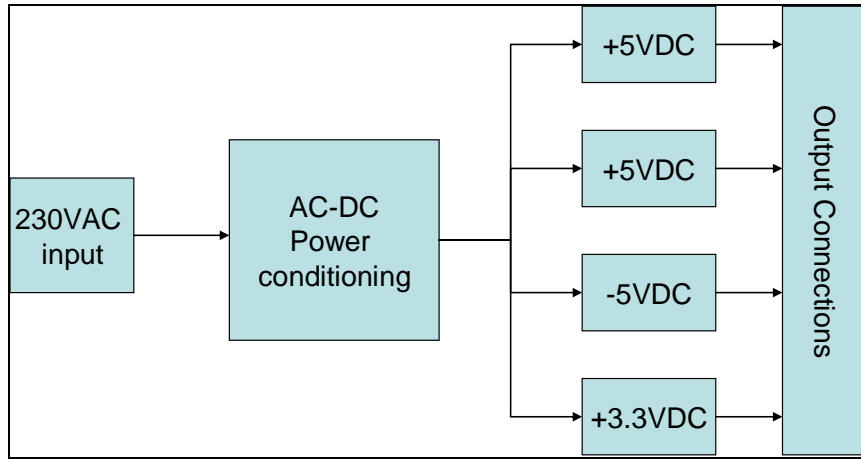


Figure A-11: Power Source PCB system block diagram

A.2.3 AC-DC Power Conditioning System

To obtain the +3.3 V, -5 V and +5 V voltages simple fixed voltage regulators are used with care taken to convert the 230 VAC input to a DC voltage within the specified input voltage ranges of the fixed regulators with a combination of transformer, diode rectifier bridge and capacitor output filter. These components are shown in the circuit diagram in Figure A-12.

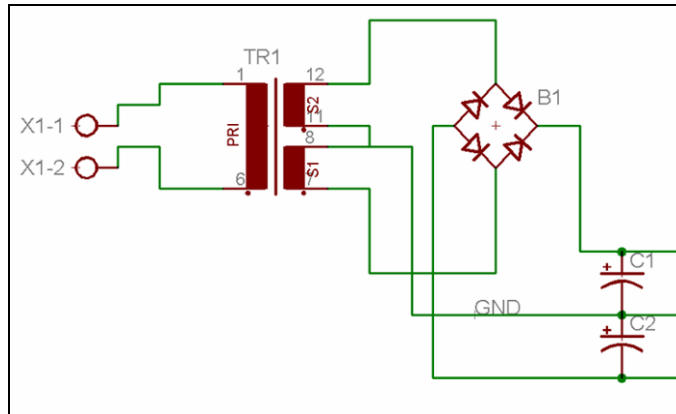


Figure A-12: AC-DC Power conditioning system

The transformer was selected to produce output voltages that are within the acceptable input voltages for the uA78M05 and uA79M05 regulators (to be discussed in the next section) which are 7~25 VDC and -7~-25 VDC. Taking the voltage drop over the diode rectifier bridge into account ($2 \times 0.7 \text{ V} = 1.4 \text{ V}$) and dividing by $\sqrt{2}$ to obtain the peak voltage results in a required output rated voltage of approximately 6~18 VAC_{RMS} and -6~-18 VAC_{RMS} . A 230/12 3.2 VA transformer with dual secondary coils were obtained to provide $2 \times 12 \text{ VAC}_{\text{RMS}}$ and a more than adequate power rating for the circuits that need power. A review of this design should decrease the size of the transformer and the capacitors (to be discussed later) to be closer to requirements.

The centre of the secondary coils are used as output ground and the other two outputs are rectified and capacitively filtered to produce ± 15.5 VDC (multiply peak with $\sqrt{2}$ and minus voltage drops across diode bridge to get to this value) with some ripple component depending on the load current.

The voltage ripple can be calculated by equation 12 shown below [65].

$$V_r = \frac{I}{100C} \quad (12)$$

Where I is the load current and C denotes the filter capacitor. This shows that the bigger the capacitance and smaller the load current, the smaller the voltage ripple. The reader is reminded, however, that voltage regulators reject voltage ripple by 60-90 dB. Table A-3 and Table A-4 summarise the power requirements for various components that get power from the +15.5 V and -15.5 V respectively. Standard practice would be to determine a required minimum voltage ripple and use the power requirements to calculate the required capacitance, however, 1mF capacitors were available at the time of the design and the actual size of the board is not a constraint and it was decided to implement these in the design. Using the capacitor values and the required current in equation 12, voltage ripples of 0.1796 V (1.2%) and 0.0576 V (0.4%) are calculated respectively. A voltage regulator with a mere 60 dB ripple rejection should bring this ripple down to a thousandth of the input at 0.18 mV and 0.06 mV respectively; and with additional output filter capacitors this noise source should be negligible and pose no threat to the operation of any other module.

Table A-3: Power budget for positive voltage

Device	Rated Voltage (V)	Rated Current (mA)	Number of	Total Power (mW)	Current@15V (mA)
Accelerometer	3.300	0.320	3	3.168	0.211
5V LEDs	5.000	3.000	2	30.000	2.000
3.3V LEDs	3.300	3.000	1	9.900	0.660
Op-Amps	5.000	0.250	3	3.750	0.250
MUXs	5.000	0.010	2	0.100	0.007
uA78M05s	15.000	4.500	3	202.500	13.500
REG1117-3.3	5.000	4.000	1	20.000	1.333
Total					17.961

Table A-4: Power budget for negative voltage

Device	Rated Voltage (-V)	Rated Current (mA)	Number of	Total Power (mW)	Current@-15V (mA)
-5V LED	5.000	3.000	1	15.000	1.000
Op-Amps	5.000	0.250	3	3.750	0.250
MUXs	5.000	0.010	2	0.100	0.007
uA79M05s	15.000	4.500	1	67.500	4.500
Total					5.757

A total current of 23.7mA is drawn from the transformer which equates to 0.568 VA which is well within the transformer's 3.2 VA rating.

A.2.4 Fixed Voltage Regulation Systems

The voltages produced by the AC-DC Power Conditioning System are ± 15.5 VDC where the devices that get power from the PCB require +3.3 V, -5 V and two +5 V power rails. This is accomplished by the circuit diagram shown in Figure A-13.

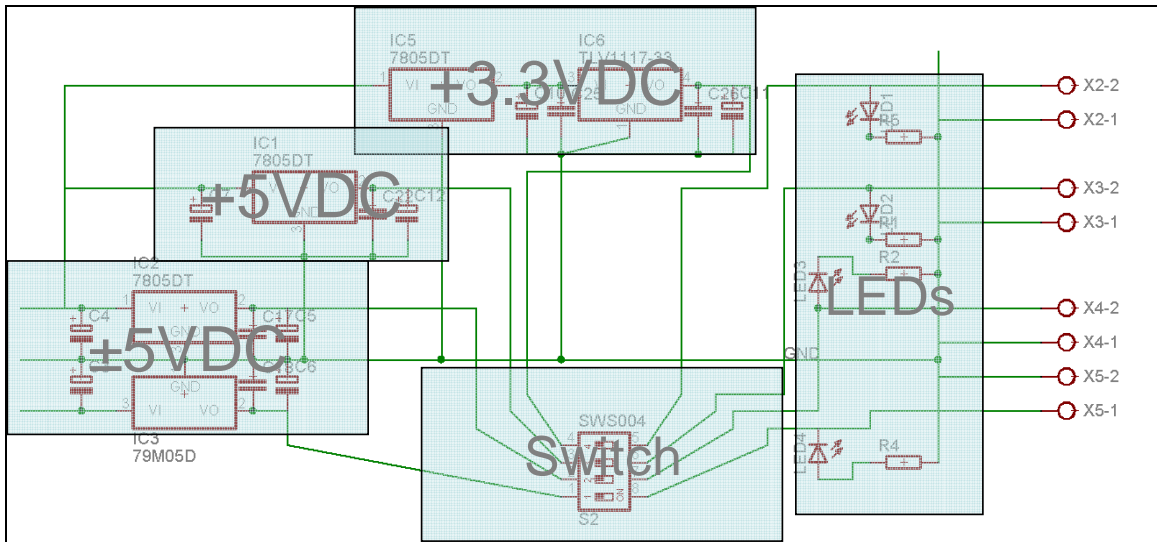


Figure A-13: Fixed voltage regulation systems

The ± 5 VDC voltages are generated by a uA78M05 and uA79M05 voltage regulator which generates +5 V and -5 V respectively. The input and output capacitors act as input and output filters placed close to the ICs.

The single +5 VDC is also generated by a uA78M05 voltage regulator with the necessary input and output capacitors as suggested by the datasheet.

The +3.3 VDC is generated by a combination of uA78M05 5 V voltage regulator and REG1117-3.3V 3.3 V voltage regulator. The reason for this combination is that the REG1117-3.3V has a low drop-out voltage (1V) and a maximum input voltage of 15 V (lower than anticipated input voltage). Input and output capacitors are placed as suggested by the datasheet of the REG1117-3.3V and not the uA78M05 (as the filtering effects of the former component's design should be sufficient).

A DIP switch is placed between the regulators and the output connectors to make it possible for the user to disable some voltages should it not be required.

Each output is connected to a through-hole LED and resistor combination which illuminates if the output is active. Each resistor is selected so that 4 mA flows through the LED (therefore all LEDs have the same luminescence) – therefore R1, R2, and R4 are selected as 1 k Ω , and R5 as 430 Ω .

A.2.5 Layout Considerations

All components share the same ground plane which greatly simplifies the layout. In general the following approach was followed:

- Power supply decoupling capacitors were placed as close as possible to the IC they protect.

- Ground planes and connections were kept as large and wide as possible to limit EMI effects.
- Power supply tracks were kept as wide as was sensible.
- Tracks were kept as short as possible – connected devices were placed as close as possible to each other.
- The board size was kept to a minimum; however size was not a constraint.
- Layout was done on a dual (top and bottom) layer board.

The input power supply (230 VAC) was kept separate from the DC circuitry as it may cause 50 Hz noise on output power rails.

A.2.6 Detailed Design Drawings

The detailed schematics can be found in Figure A-12 and Figure A-13 of the previous two sub sections. These and the following figures should help the reader understand how the PCB was realised, but these cannot be used to directly recreate the PCB. The author should be contacted who will supply interested parties with the necessary design files. “EAGLE - Easily Applicable Graphical Layout Editor Version 4.16r1 for Windows Light Edition” was used as design software for the PCB. Figure A-14 and Figure A-15 show the top and bottom view of actual board. To give the reader an idea of the actual size of the board, the top edge is 108 mm and the left edge is 97 mm in length. The boards were populated and implemented as suggested by the components indicated on the board images.

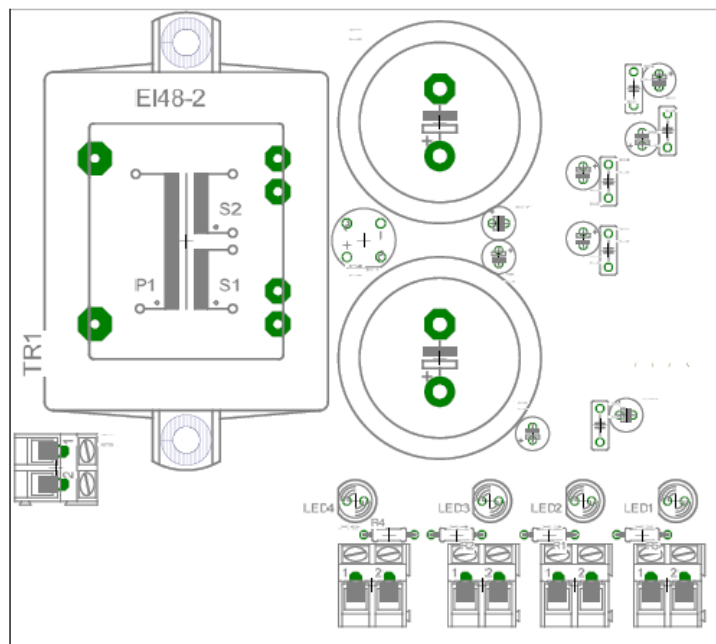


Figure A-14: Power PCB top view

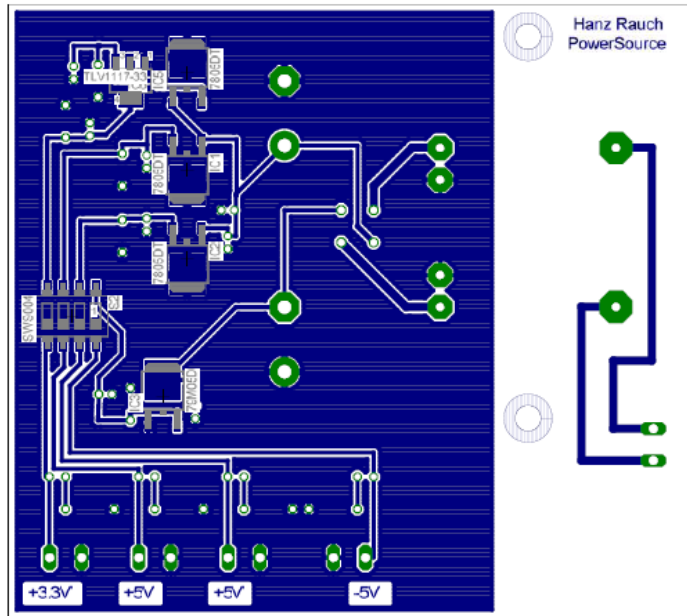


Figure A-15: Power PCB bottom view

A.3 Design of Accelerometer PCBs

The accelerometer PCBs acts as a surface for the ADXL330 3-axis MEMS accelerometers to attach to and contain the necessary power supply noise decoupling and output low pass filter capacitors. Due to the simplicity of the PCB, there is no system breakdown in this section.

A.3.1 Requirements




Figure A-16: ADXL330 MEMS accelerometer block diagram (extracted from datasheet)

The ADXL330 3-axis MEMS accelerometer functional block diagram is shown Figure A-16. The internal operation is detailed in the datasheet [52].

For the accelerometer to be light enough to not significantly dampen the PCB needs to have a small footprint. Furthermore, the necessary power supply decoupling capacitors are to be supplied close to the IC (0.1 μ F) as well as the three output's output filtering capacitors. The selection of output filter capacitors depends on the desired signal bandwidth which is governed by equation 13.

$$F_{-3dB} = \frac{1}{2\pi(32k\Omega) * C_{X,Y,Z}} \quad (13)$$

The X and Y output capacitors are set to 4.7 nF to give a bandwidth of approximately 1 kHz (the actual signal bandwidth for these two channels is 1.6 kHz) and the Z input is set to 10 nF to give a bandwidth of 500 Hz (the Z signal bandwidth is 500 Hz). The Z-axis, as the device is mounted on the patient's body, is parallel to the expected motion of the skin – this could therefore be considered the actual signal bandwidth.

A.3.2 Detailed Design Drawings

The detailed schematics and actual layout can be seen in Figure A-17 and Figure A-18 below. These and the following figures should help the reader understand how the PCB was realised, but these cannot be used to directly recreate the PCB. The author should be contacted who will supply interested parties with the necessary design files. “EAGLE - Easily Applicable Graphical Layout Editor Version 4.16r1 for Windows Light Edition” was used as design software for the PCB. To give the reader an idea of the actual size of the board, the top edge is 20mm and the left edge is 18 mm in length. The boards were populated and implemented as suggested by the components indicated on the board images.

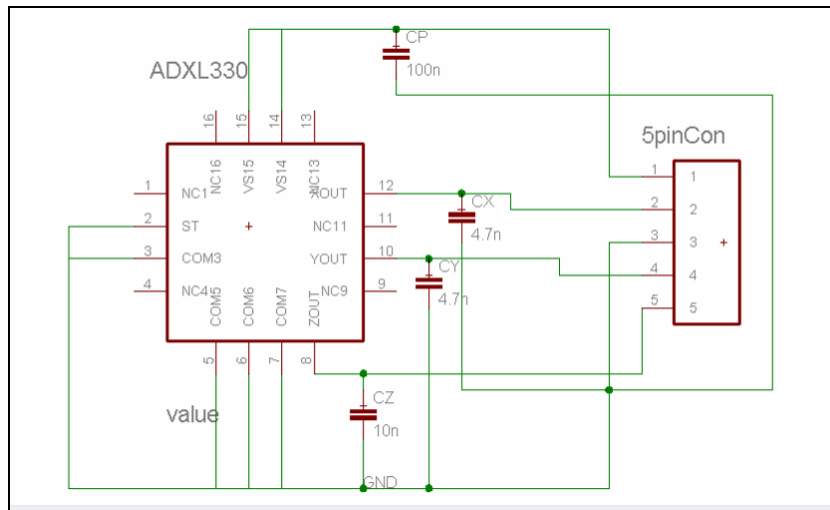


Figure A-17: Accelerometer PCB schematic




Figure A-18: Layout of Accelerometer PCB (single layer)

A.4 FSR Calibration

The Force Sensing Resistor (FSR) was *calibrated* using the conductance-force curve on the Interlink Model 400 datasheet [51]. The author acknowledges that this method is scientifically incorrect and that the supplier datasheet even says that the sensor should not be used for precision measurement, but the result should be repeatable and make individual measurements comparable. It assumed that the sensor does not drift.

The FSR is connected in a voltage divider topology with resistor (R_{REF}) and power source (V_{cc}) – depicted in Figure 3.2. The FSR resistance (R_{FSR}) relative to the voltage measured above the FSR (V_{FSR}) is given by equation 14 and the result is converted to conductance (G_{FSR}) in equation 15.

$$R_{FSR} = \frac{V_{FSR}}{V_{cc} - V_{FSR}} \cdot R_{REF} \quad (14)$$

$$G_{FSR} = \frac{V_{cc} - V_{FSR}}{V_{FSR} \cdot R_{REF}} \quad (15)$$

The conductance-force relationship for the FSR is shown in Figure A-19. The values are extracted at each point and a cubic curve is fitted through the points (notice that the curve fitting axes are swapped and converted).

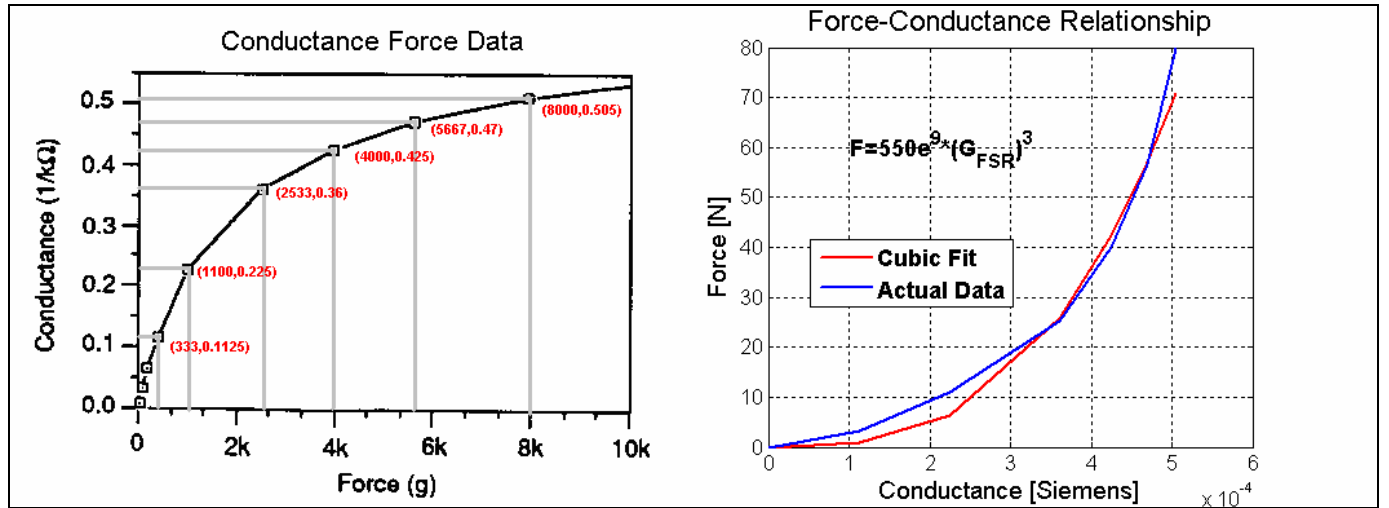


Figure A-19: FSR resistance-force curve [51] and curve fit

The cubic fit through the points is given by equation 16.

$$F = 550 \cdot 10^9 \cdot G_{FSR}^3 \quad (16)$$

The force-voltage (F - V_{FSR}) relationship can now be obtained by combining the cubic fit and G_{FSR} - V_{FSR} relationship – shown in equation 17.

$$F = 550 \cdot 10^9 \cdot \left(\frac{V_{cc} - V_{FSR}}{R_{REF} \cdot V_{FSR}} \right)^3 = 550 \cdot 10^9 \cdot \left(\frac{1}{R_{REF}} \left[\frac{V_{cc}}{V_{FSR}} - 1 \right] \right)^3 \quad (17)$$

In the implemented circuitry, a 9.8 kΩ resistor was used as R_{REF} and V_{cc} was equal to 5 V. Replacing the variables in equation 17 with these values, gives equation 18.

$$F = 0.584 \cdot \left(\frac{5}{V_{FSR}} - 1 \right)^3 \quad (18)$$

Appendix B Software Designs

The source code is available on the CD that accompanies this thesis report. Should the CD not be available the author can be contacted.

B.1 *get_aver_pulse.m*

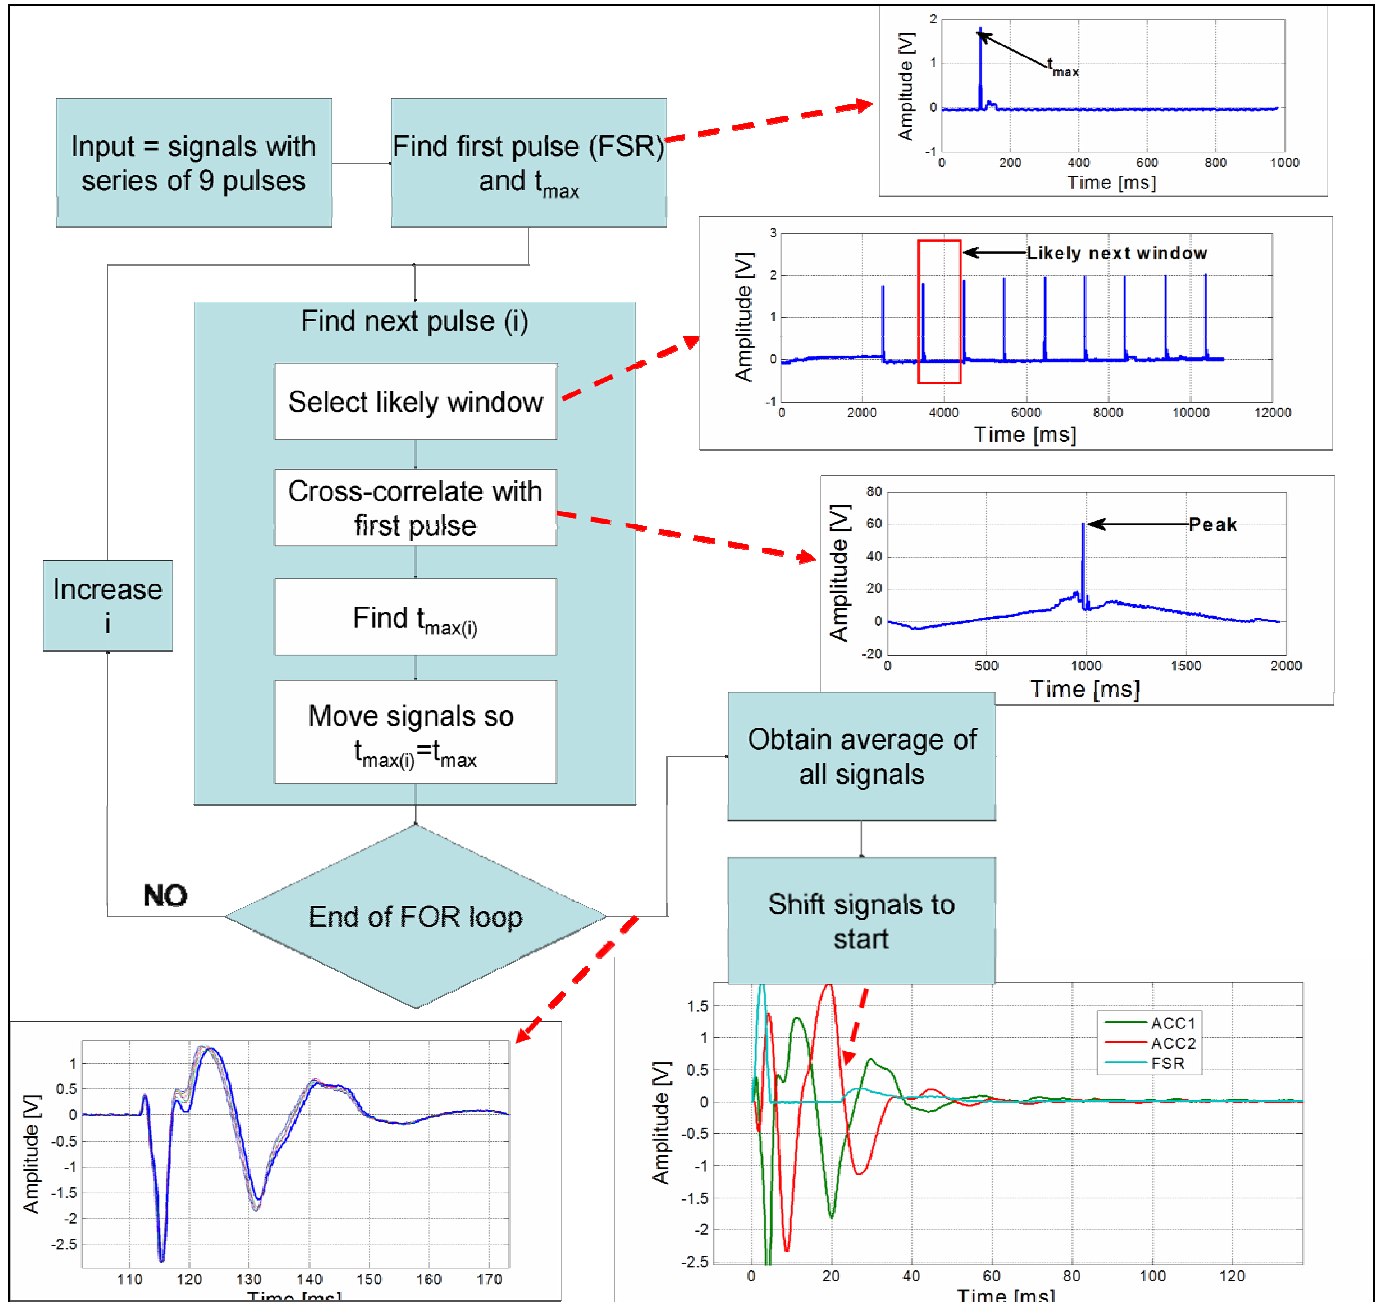


Figure B-1: *getting_aver_pulse.m* flow diagram

The task of this function is to take a set of signals from a particular acquisition – FSR, Accelerometer 1 and Accelerometer 2 – extract the pulses in each signal, obtain an average of that signal, and adjust the signals so that the FSR signal starts at 0 ms and the other two are shifted accordingly (temporal relationship between signals are maintained).

This task is accomplished by using some prior knowledge as to where signals should occur, predicting the time windows where each pulse should be located, using cross-correlation to determine the exact point at which two signals match the best and finally to move each pulse so that it matches the first. These signals are then averaged and the start of the FSR signal is identified. All signals are the shifted forward in time so that the FSR signal starts at 0 ms.

The algorithm flow is depicted in Figure B-1.

B.2 General Feature Extraction Flow

The human subject data was converted into features using three different methods, but the general flow was the same – generate a list of the data files to be processed, loop through each file to obtain an average set of signals and subsequently extract the features. Temporary results are stored in variables and log files whilst the algorithm loops, and when it finishes all data is stored in a data file. This general flow is depicted in

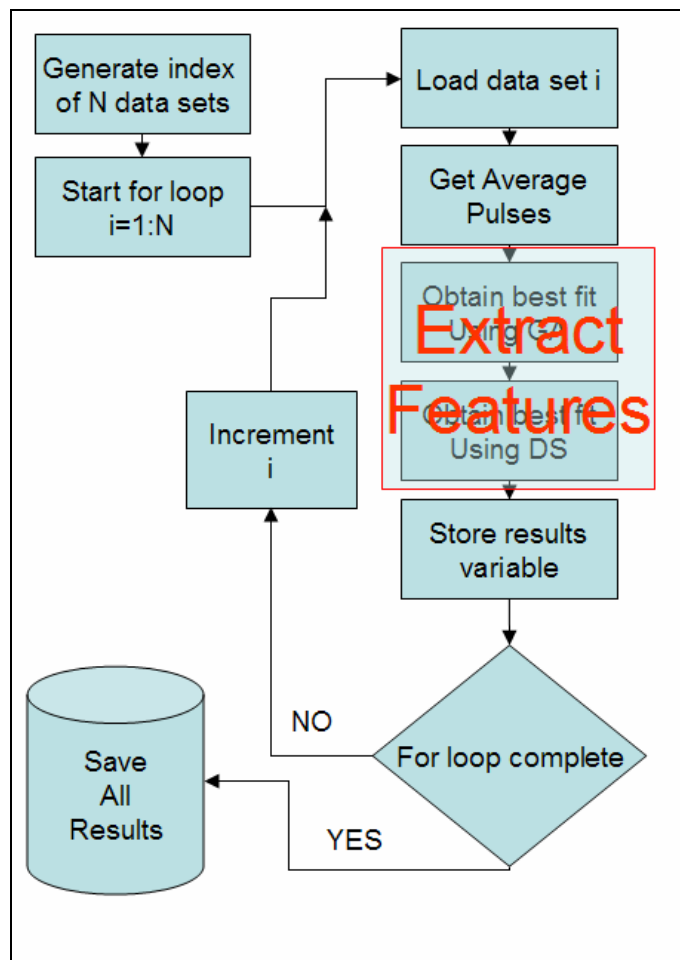


Figure B-2: General Feature Extraction Flow

The flow diagram for each extraction method is shown in broad overview in the following sections.

B.2.1 Direct Feature Extraction

Direct feature extraction is quite simple - values are simply calculated from the signals iteratively. This simple flow diagram is depicted in Figure B-3.

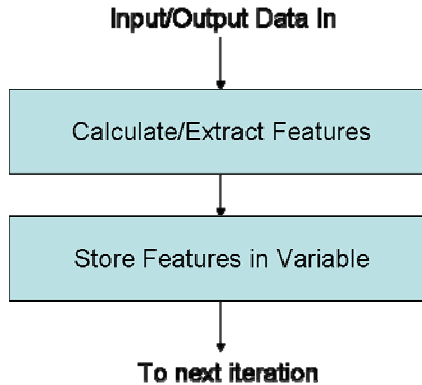


Figure B-3: Direct feature extraction

B.2.2 Extracting Autoregressive Features

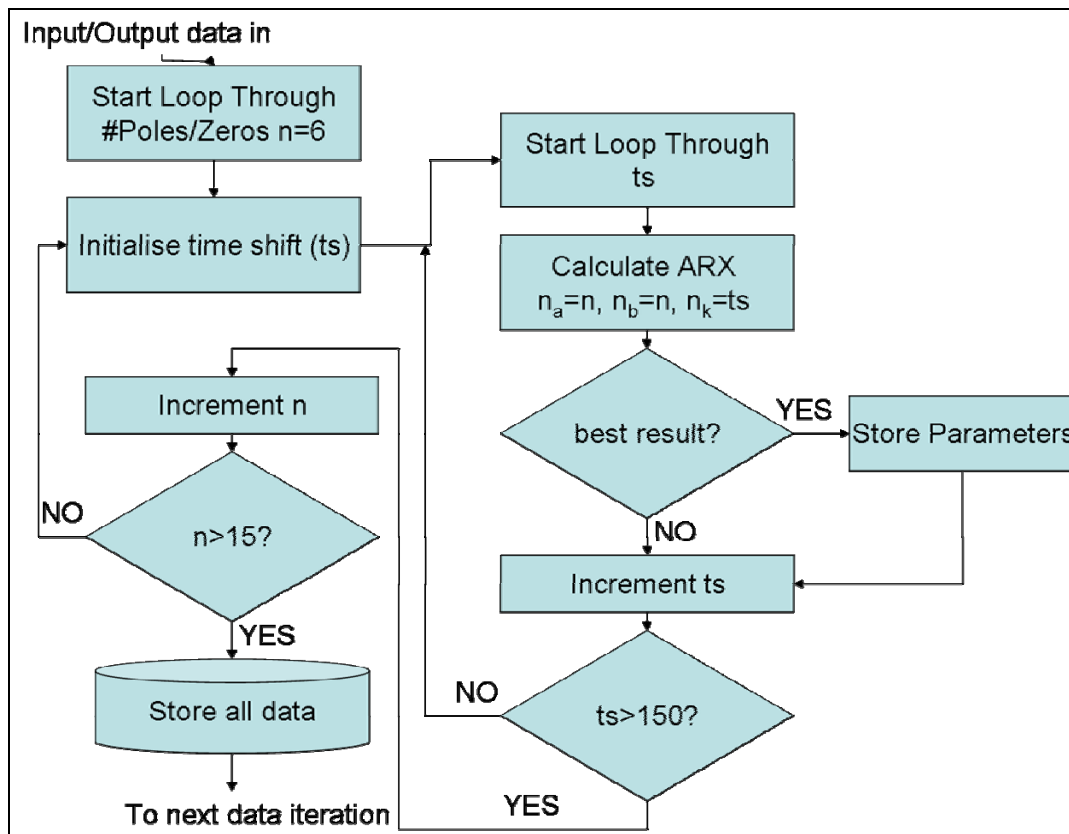


Figure B-4: Extracting autoregressive features

The autoregressive (with exogenous input) system identification extracts points for systems with six to 15 poles. The process that is followed with each input/output data combination is the following: use the toolbox to find the time series coefficients that best fit the data with a specific time shift; if a best result is obtained, update, otherwise don't; increase time shift; when all time shifts have been considered, increase number of poles/zeros. This flow is depicted in Figure B-4.

B.2.3 Extracting Mechanical Model Features

The parameters of the dual mass-spring-damper system are tuned to fit data using the Genetic Algorithm (GA). The GA passes values from the search space to an objective function which simulates the model using these parameters. A cost function evaluates the fit and returns the fitness value to the GA. Based on the number of generations and the resulting fit the GA will terminate or generate a new population. The best fitting population is sent to the Direct Search Optimisation (DSO) algorithm to find the maximum in the local space surrounding the GA result. The DSO algorithm uses a gradient approach to find the optimum solution and evaluates the solution by using the same objective function. The DSO algorithm terminates when the improvement becomes small enough or when the number of iterations exceed the upper limit (settings are left at the toolbox default). These two stages are depicted in Figure B-5 and Figure B-6 respectively.

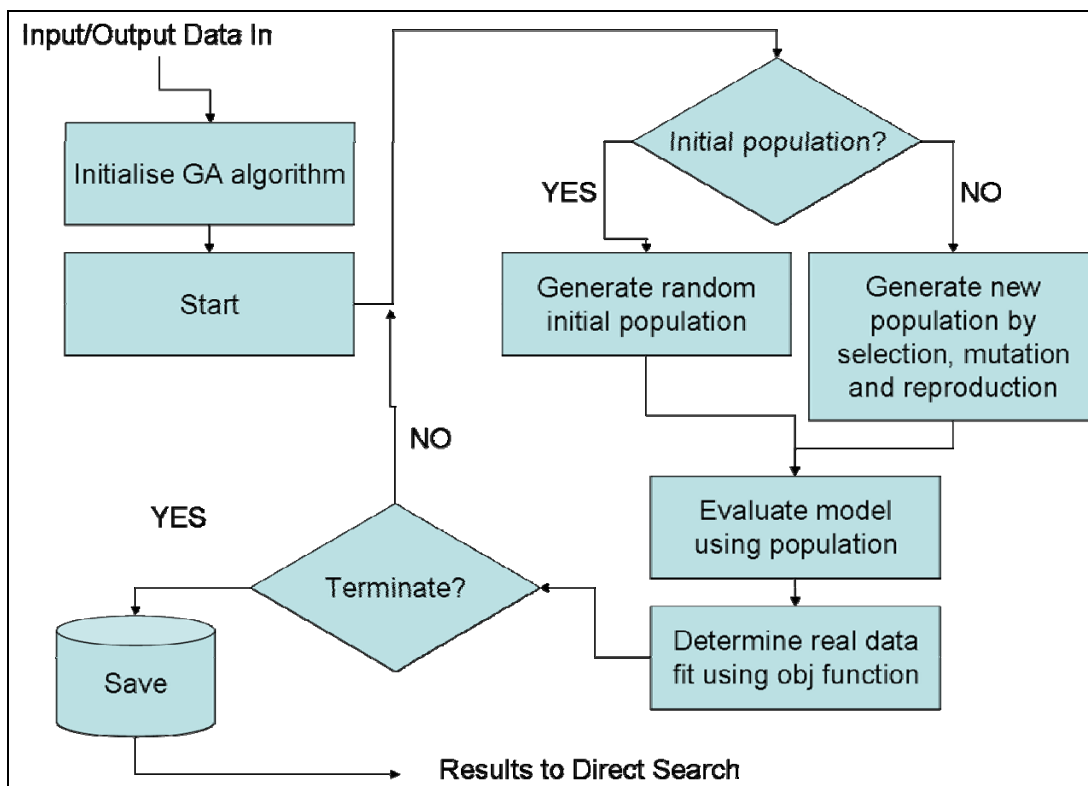


Figure B-5: Extracting mechanical model features, Genetic Algorithm

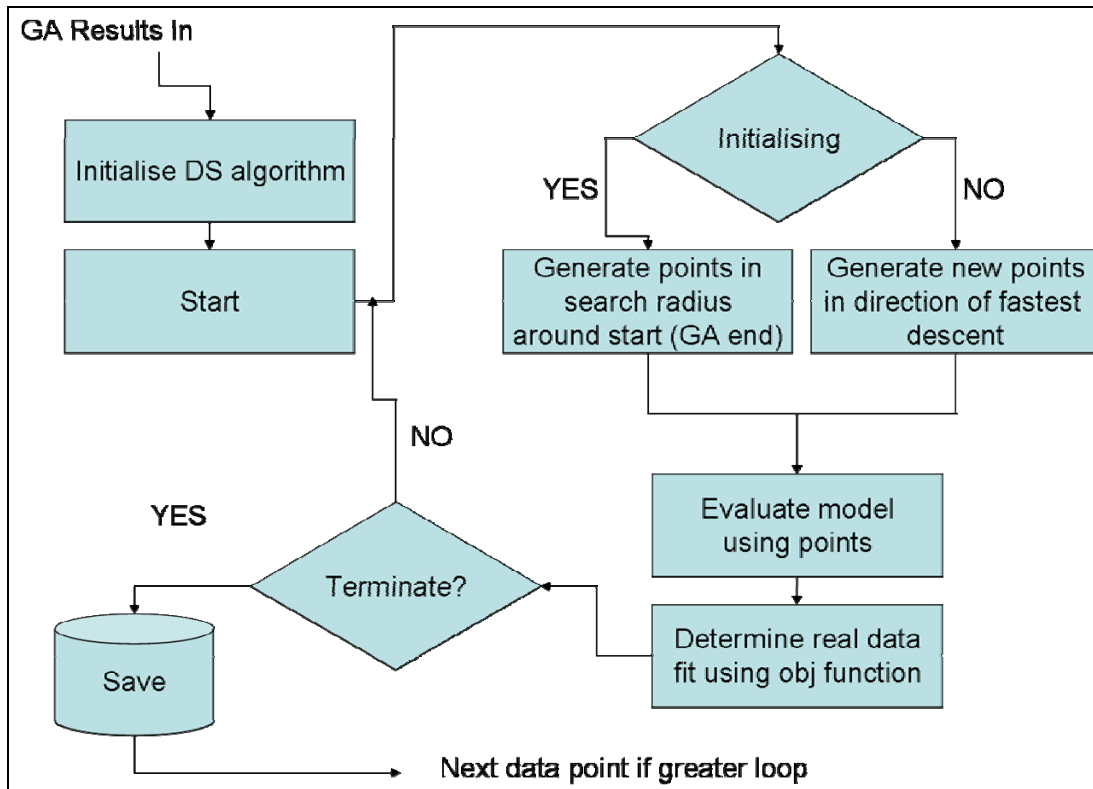


Figure B-6: Extracting mechanical model features, Genetic Algorithm

Appendix C Mechanical Designs

C.1 Frame Design

The experimental setup for tests on human subjects requires the design and manufacture of a rigid adjustable frame that can be attached to a bed for experimental purposes. The frame should allow an operator to move a vibrational shaker to various points for “percussion” of a patient or volunteer.




Figure C-1: Frame concept

The mechanical drawing for the parts that make up the frame are contained in Figure C-2 to Figure C-7.




Figure C-2: Mechanical drawing: frame back part





Figure C-3: Mechanical drawing: frame half circle



Particulars	Quantity	Unit	Remarks	Drawn	Checked
Base				1/1	

Figure C-4: Mechanical drawing: frame base




Figure C-5: Mechanical drawing: frame feet




Figure C-6: Mechanical drawing: frame slider




Figure C-7: Mechanical drawing: frame clamps

C.2 Phantom test setup

The setup used for tests on the silicone model was drawn up so that it could be recreated, if need be. Figure C-8 shows a depiction of the 3D model of the setup.




Figure C-8: Phantom test setup

Appendix D Additional Information

D.1 Body Planes and Orientations

Figure D-1 shows the body planes and orientations as is referenced in the text. This image was extracted from [66].


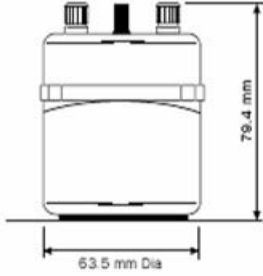
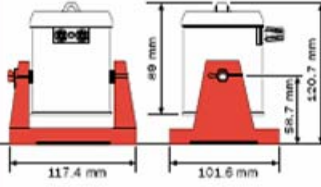


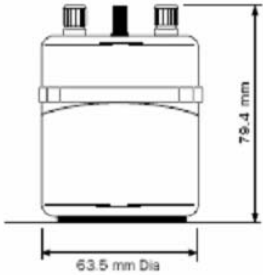
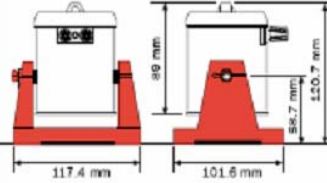
Figure D-1: Summary of body planes and orientations [66]

Appendix E Datasheet Excerpts

This section contains excerpts of datasheets which are not easily attainable off the internet, or may be removed in future.

E.1 Vibrational Shaker V201-PA25E from LDS Group

Model	V101/2 - PA 25E		V201/3 - PA 25E	
				
	Metric	American	Metric	American
Armature diameter	-	-	-	-
System sine force peak - naturally cooled	8.9 N	2 lbf	17.8 N	4 lbf
System sine force peak - forced air cooled	-	-	17.8 N	4 lbf
Shaker sine force peak* - forced air cooled	-	-	26.7 N	6 lbf
System random force rms (ISO5344)	-	-	-	-
Shaker max random force rms*	-	-	-	-
Armature resonance frequency	12000 Hz	12000 Hz	13000 Hz	13000 Hz
Useful frequency range	5-12000 Hz	5-12000 Hz	5-13000 Hz	5-13000 Hz
Effective mass of moving element	0.0065 kg	0.0143 lb	0.020 kg	0.044 lb
System velocity sine peak	1.31 m/s	51.6 in/s	1.49 m/s	58.7 in/s
Shaker velocity sine peak*	1.31 m/s	51.6 in/s	1.83 m/s	72 in/s
Max acceleration sine peak	1373 m/s ²	140 gn	1334 m/s ²	136 gn

Model	V101/2 – PA 25E		V201/3 – PA 25E	
				
	Metric	American	Metric	American
Amplifier rating	0.048 kVA	0.048 kVA	0.048 kVA	0.048 kVA
LDS amplifier	PA 25E	PA 25E	PA 25E	PA 25E
Suspension axial stiffness	3.15 N/mm	18 lbf/in	3.5 N/mm	20 lbf/in
Aux. suspension axial stiffness	-	-	8.76 N/mm	50 lbf/in
System displacement (continuous) pk-pk	2.5 mm	0.1 in	5 mm	0.2 in
Shaker displacement (continuous) pk-pk*	2.5 mm	0.1 in	5 mm	0.2 in
Vibrator mass-optional trunnion mounted	-	-	3.17 Kg	7 lb
Cooling air flow rate	-	-	0.0012 m ³ /s	2.5 ft ³ /m
Shaker mass base mounted	0.91 kg	2 lb	1.81 kg	4 lb
Max. working ambient temperature				
Shaker:	30°C	86°F	30°C	86°F
Amplifier:	35°C	95°F	35°C	95°F
Heat rejected to air				
Shaker cooling fan:	-	-	-	-
Amplifier:	0.067 kW	0.067 kW	0.067 kW	0.067 kW
Electrical requirement				
Amplifier:	0.09 kVA	0.09 kVA	0.09 kVA	0.09 kVA
Acoustic noise at 2m**				
Shaker:	-	-	-	-
Amplifier:	silent	silent	silent	silent

* Shaker ratings are those which can be achieved with a larger amplifier than that supplied as standard

** Maximum noise level during operation at full thrust, measured at 2m distance

† Includes cooling fan

E.2 Smooth-On Dragon Skin Silicone Rubber Compounds



Dragon Skin™

Addition Cure Silicone Rubber Compounds

PRODUCT OVERVIEW

Dragon Skin_™ and Dragon Skin_™ Q are high performance platinum cure silicone rubbers that can be mixed 1A:1B by weight or volume and cures at room temperature with negligible shrinkage to a Shore 10A. Cured Dragon Skin_™ is very soft, very strong and very “stretchy”. It will stretch many times its original size without tearing and will rebound to its original form without distortion.

Dragon Skin_™ is suitable for many special effects applications, especially animatronics where repetitive motion is required. It is water white translucent and will accept pigments for creating color effects or for creating a paint that can be painted on to an existing Dragon Skin_™ (or other) silicone surface. Because of the superior physical properties and softness of Dragon Skin_™, it is also used for orthopedic & cushioning applications. **Dragon Skin_™** has a pot life of 20 minutes and a demold time of 5 hours.

In A Hurry? Dragon Skin_™ Q (Q = Quick) has a pot life of 8 minutes and a demold time of 75 minutes.

TECHNICAL OVERVIEW

Key Values: ~Mixing Ratio: 1A to 1B by weight or volume		~Shore A Hardness: 10			
~Color: Translucent Clear					
Properties	Viscosity	G/CC	Cu. In./Lb.	Tensile Strength	Mix Ratio
Part A	-	-	-	-	100 pbw/pbv
Part B	-	-	-	-	100 pbw/pbv
Mixed	23,000 cps	1.07	25.8	475 psi	-
Dragon Skin		Dragon Skin Q		~Elongation At Break . . . 1000%	~100 % Modulus . . . 22 psi
~Pot Life*: 20 minutes		8 minutes		~Die B Tear Strength . . . 102 pli	~ShrinkageNegligible
~Demold*: 5 hours		75 minutes		* at room temperature: 75°F / 23° C	

Preparation

Safety – Use in a properly ventilated area (“room size” ventilation). Wear safety glasses, long sleeves and rubber gloves to minimize contamination risk. Wear vinyl gloves only. Latex gloves will inhibit the cure of the rubber. **Store & Use At Room Temperature** (72°F / 22°C). Warmer temperatures will drastically reduce working time and cure time. Storing material at warmer temperatures will also reduce the usable shelf life of unused material. These products have a limited shelf life and should be used as soon as possible.

Cure Inhibition – Addition-cure silicone rubber may be inhibited by certain contaminants in or on the pattern to be molded resulting in tackiness at the pattern interface or a total lack of cure throughout the mold. Latex, sulfur clays, certain wood surfaces, newly cast polyester, epoxy or urethane rubber may cause inhibition. If compatibility between the rubber and the surface is a concern, a small-scale test is recommended. Apply a small amount of rubber onto a non-critical area of the pattern. Inhibition has occurred if the rubber is gummy or uncured after the recommended cure time has passed. **To prevent inhibition, one or more coatings of a clear acrylic lacquer applied to the model surface is usually effective.** Allow any sealer to thoroughly dry before applying rubber. **Note: Even with a sealer, platinum silicones will not work with modeling clays containing heavy amounts of sulfur. Do a small scale test for compatibility before using on your project.**

Applying A Release Agent? Although not usually necessary, a release agent will make demolding easier when casting into or over most surfaces. Ease Release 200_™ is a proven release agent for making molds with silicone rubber. Mann Ease Release_™ products are available from Smooth-On or your Smooth-On distributor. **~IMPORTANT:** To ensure thorough coverage, lightly brush the release agent with a soft brush over all surfaces of the model. Follow with a light mist coating and let the release agent dry for 30 minutes.

If there is any question about the effectiveness of a sealer/release agent combination, a small-scale test should be made on an identical surface for trial. Also, you can call **Smooth-On for technical assistance at (800) 762-0744.**

Measuring & Mixing . . .

Materials should be stored and used in a warm environment (72° F / 22° C). This product has a limited shelf life and should be used as soon as possible. Wear safety glasses, long sleeves and rubber gloves to minimize contamination risk. Before you begin, pre-mix Part B (base) thoroughly. After dispensing required amounts of Parts A and B into mixing container (1A:1B by volume or weight), **mix thoroughly for 3 minutes** making sure that you **scrape the sides and bottom of the mixing container several times**. After mixing parts A and B, vacuum degassing is recommended to eliminate any entrapped air. Vacuum material for 2 -3 minutes (29 inches of mercury), making sure that you leave enough room in container for product volume expansion.

Pouring

Curing

Mold Performance

For best results, pour your mixture in a single spot at the lowest point of the containment field. Let the rubber seek its level up and over the model. **A uniform flow will help minimize entrapped air.** The liquid rubber should level off at least 1/2" (1.3 cm) over the highest point of the model surface.

Curing . . . Allow rubber to cure as prescribed at room temperature (75°F/23°C) before demolding. Full physical properties are attained after 24 hours cure at room temperature. Do not cure rubber where temperature is less than 65°F /18°C. Allow mold to cool to room temperature before using.

If Using As A Mold . . . When first cast, silicone rubber molds exhibit natural release characteristics. Depending on what is being cast into the mold, mold lubricity may be depleted over time and parts will begin to stick. No release agent is necessary when casting wax or gypsum. Applying a release agent such as Ease Release 200 (available from Smooth-On) prior to casting polyurethane, polyester and epoxy resins is recommended to prevent mold degradation.

Thickening With Thi-Vexsm Thickening Agent . . . A very small amount of Thi-vex will thicken the rubber brushable for vertical surface application. 0.25% - 0.5% Thi-vex (% of the total weight of the mixture, A+B) is all that is necessary.

Example: 100 Parts A + 100 Parts B (mix thoroughly) + 0.5 Parts by weight = brushable, some sag
100 Parts A + 100 Parts B (mix thoroughly) + 1.0 Part by weight = thick brushable, no sag

Thinning Dragon Skin – Smooth-On's 'Silicone Thinner' will lower the viscosity of Dragon Skin for easier pouring and vacuum degassing.

Safety First

The Material Safety Data Sheet (MSDS) for this or any Smooth-On product should be read prior to use and is available upon request from Smooth-On. All Smooth-On products are safe to use if directions are read and followed carefully.

Be careful. Use only with adequate ventilation. Contact with skin and eyes may cause irritation. Flush eyes with soap and water for 15 minutes and seek immediate medical attention. Remove from skin with waterless hand cleaner followed by soap and water.

Important: The information contained in this bulletin is considered accurate. However, no warranty is expressed or implied regarding the accuracy of the data, the results to be obtained from the use thereof, or that any such use will not infringe upon a patent. User shall determine the suitability of the product for the intended application and assume all risk and liability whatsoever in connection therewith.

**Smooth-On offers a complete line of Liquid Rubber, Liquid Plastic and Release Agent products for hundreds of industrial and art related applications.
Chances are there is a distributor in your area to offer local support.**

Toll-free: (800) 762-0744

Fax: (610) 252-6200

Website: www.smooth-on.com

040808