



A Comparison of Binaural Ultrasonic Sensing Systems

A thesis submitted in fulfillment of the
requirements for the award of the degree

Doctor of Philosophy

from

UNIVERSITY OF WOLLONGONG

by

Benjamin David Stanley

B. Sc. (Hons), Australian National University

School of Information Technology and Computer Science

December 2003

© Copyright 2003
by
Benjamin David Stanley
All Rights Reserved

Declaration

I, Benjamin David Stanley, declare that this thesis, submitted in fulfillment of the requirements for the award of Doctor of Philosophy, in the School of Information Technology and Computer Science, University of Wollongong, is wholly my own work unless otherwise referenced or acknowledged. This document has not been submitted, in whole or in part, for qualifications at any other academic institution.

Benjamin David Stanley
December 13, 2003

When I first started . . . I wrote “the Asdic is a feminine, fickle and unpredictable instrument”, but on reconsideration I withdraw this libellous statement. The Asdic is very reliable as an instrument. What is unpredictable is the behaviour of the sound waves in the sea.

The account of a British naval officer using the Asdics system to locate a sunken British submarine in 1951, from Hackmann [\[48\]](#).

Abstract

Autonomous robotic systems have been under development for a number of years. This development is being driven partly by the need to explore the processes and logic that humans use in accomplishing repetitive tasks and partly by the need to provide probes and vehicles in environments that present threats to humans. It is this need to send robots into harsh environments which has provided the impetus for this thesis to explore the use of binaural ultrasonic sensors for *robotic agents* thereby enabling them to take the place of *living agents*. In all, five sensor systems were built and investigated. Two systems were based upon *matched filters* with Barker codes. The first used quadrature demodulation while the second used unsynchronised demodulation. The remaining three systems used *continuous tone frequency modulated* signals. The first was implemented with a *discrete Fourier transform* spectral estimator, the second with a *Yule-Walker* spectral estimator and the third using a *least squares modified Yule-Walker* spectral estimator.

The construction, implementation and testing of these five sonar systems reveals that there is no clear preferred system despite the significantly different methodologies adopted to determine the basic time of flight and received power information. The more classical *matched filter* systems have superior *distance of flight* precision compared to the newer *continuous tone frequency modulated* class. However it is clearly demonstrated that the *continuous tone frequency modulated* class is able to insonify a larger area and operate with a better *signal to noise ratio*. Once calibrated, this class of sonar is able to at least equal the performance of the *matched filter* class.

The thesis focuses on the rigorous determination of the range and bearing tuple from both the theoretical or model viewpoint and the experimental viewpoint. A rigorous theoretical development of both the *inter-aural distance difference* and the *inter-aural power difference* methods of computing bearing and the resolution of the correspondence problem are included. Experimental work was undertaken to validate the models and to provide data on which comparative study may be based. The conventional metric quantities were acquired with a large level of redundancy thereby allowing the determination of strong statistics associated with accuracy, precision and resolution.

This thesis shows that, while there are performance differences between the systems, the selection of the best system is application dependent. In this work the aim is to equip a *robotic agent* with a sensory system enabling it to function in the place of a *living agent*, especially in a harsh or dangerous environment. For this purpose the system suggested by this thesis is a *continuous tone frequency modulated* sonar system with a *least squares modified Yule-Walker* spectral estimator.

Contents

List of Figures	xvii
List of Tables	xxiii
List of Algorithms	xxv
Glossary	xxvii
Acknowledgements	xxxvii
1 Introduction	1
1.1 Aim	1
1.2 Goals	1
1.3 Continuous Tone Frequency Modulated Sonar	2
1.3.1 Sonar Origins	3
1.3.2 CTFM Air Sonar as a Sensory Aid for Humans	4
1.3.3 CTFM Air Sonar as a Robotic Sensory System	5
1.4 Outline	6
1.5 Contributions	8
I Foundations	11
2 Review of Robotic Location Sensing Technology	13
2.1 Criteria for Comparison	13
2.2 Infra-Red LED Sensors	14
2.3 Scanning Laser Range-Finders	15
2.4 Radars	17
2.4.1 Technology	17
2.4.2 Application to Mobile Robots	19
2.5 Ultrasonic Sensors	19
2.5.1 Polaroid Type	21
2.5.2 HiSonic Sensor	21
2.5.3 Plant and Target Recognition with CTFM Sonar	23
2.5.4 Measuring Bearing with a Single Receiver	23
2.5.5 Direction Finding Sonar Ring	25

2.5.6	Impulses and Template Matching	25
2.5.7	Barker Codes and Matched Filters	26
2.5.8	3D CTFM Ultrasonic Sensor	26
2.5.9	CTFM Sonars for <i>Robotic Agents</i> by Leslie Kay	28
2.5.10	CTFM Arrays	29
2.6	Combined Visual and Sonar Sensor	30
2.7	Properties of Ultrasound	30
2.8	Conclusion	33
3	Operation of CTFM	35
3.1	Overview	35
3.2	Transmit Signal	36
3.3	Demodulation	38
3.3.1	Single Demodulation	38
3.3.2	Dual Demodulation	42
3.4	Conclusion	47
4	The Signal Path	49
4.1	Domains	49
4.2	The Sonar Equation	50
4.2.1	Power Domain	53
4.2.2	Decibel Power Domain	54
4.2.3	Time Domain	55
4.2.4	Frequency Domain	58
4.2.5	Comparison of Domains	59
4.3	Sonar Equation Terms	59
4.3.1	Transducer Response	59
4.3.2	Spreading Loss and Propagation Delay	59
4.3.3	Air Absorption	60
4.3.4	Target Reflecting Area	61
4.3.5	Processing Gain	62
4.4	Multiple Targets	62
4.5	Variable Gain Compensation	62
4.6	Theoretical Range Resolution	63
4.6.1	The Time-Bandwidth Product	63
4.6.2	Processor Compression	65
4.6.3	Range Resolution Limit	65
4.6.4	Evaluating Time and Frequency Widths	67
4.7	Precision and Signal to Noise Ratio	69
4.8	Conclusion	69
4.8.1	Acknowledgements and Contributions	69

II	Supporting Equipment	71
5	Experimental System Overview	73
5.1	Chirp Input and Output Cards	73
5.1.1	Chirp Generator Card	75
5.1.2	Chirp Capture Card	78
5.1.3	Chirp IO Software	80
5.2	Precision Positioner	80
5.2.1	Hardware	83
5.2.2	Software	83
5.3	Environment Sensor	84
5.3.1	Electronics	85
5.3.2	Software	87
5.4	Host Computer	88
5.4.1	Transducer Response Measurement	88
5.4.2	Feature Extraction Dialogue	88
5.4.3	Hardware Abstraction	90
5.5	Auxiliary Equipment	91
5.5.1	High Voltage Power Supply	91
5.5.2	Ultrasonic Microphone	91
5.5.3	Calibrated Source	91
5.5.4	Vaisala HM-34	91
5.6	Conclusion	91
5.6.1	Acknowledgements and Contributions	93
6	Transducers	95
6.1	Construction of the Kay Transducer	95
6.1.1	Dimensions	96
6.2	Operation of the Kay Transducers	96
6.3	Models of Transducer Operation	99
6.3.1	Transduction Model	99
6.3.2	Radiating Plane Piston Model	99
6.3.3	Kuc-Siegel Impulse Model	103
6.4	Total Effect of a Transducer	103
6.5	Beam Directivity	104
6.6	Circuits	105
6.6.1	Transmit Circuit	105
6.6.2	Receive Circuit	106
6.6.3	Transmit Circuit with External Bias	107
6.7	Experimental Kay Transducer Response	108
6.8	Response of Polaroid Transducers	111
6.9	Conclusion	114
6.9.1	Acknowledgements and Contributions	116

7	Sensor Head Design	119
7.1	Requirements of the <i>Robotic Agent</i>	119
7.1.1	Precision and Resolution	119
7.1.2	Horizontal Beam Width	120
7.1.3	Vertical Beam Width	121
7.1.4	Update Time	122
7.2	Transducer Selection and Arrangement	123
7.2.1	Transducer Size	123
7.2.2	Receiver Orientation	126
7.2.3	Selection of Transducers and Splay Angles	129
7.2.4	Distance Between Receivers	130
7.3	Manufactured Sonar Head	131
7.3.1	Crosstalk	133
7.4	Conclusion	133
7.4.1	Acknowledgements and Contributions	135

III Sonar DOF Sub-System 137

8	Matched Filter System	139
8.1	Matched Filters	139
8.2	Barker Code	141
8.3	Barker Coded Transmit Signal Generator	141
8.4	ADC Output	145
8.5	Barker Code Matched Filter	145
8.6	The Envelope Extractor	148
8.6.1	The Need for Envelope Extraction	150
8.6.2	Quadrature Demodulation	151
8.6.3	Unsynchronised Demodulation	152
8.7	Target Detection	155
8.7.1	Peak Finder	155
8.7.2	Peak Validator	156
8.7.3	Peak Interpolator	158
8.8	Range Resolution	158
8.8.1	Theoretical System Resolution	159
8.8.2	Time-Bandwidth Measurements	159
8.8.3	Theoretical Resolution Using Time Width	161
8.8.4	Theoretical Resolution Using Bandwidth	161
8.8.5	Summary of Resolution Results	161
8.9	Software Implementation	162
8.10	Conclusion	162
8.10.1	Acknowledgements and Contributions	164

9	Implementation of CTFM	165
9.1	System Parameters	165
9.2	Sampling Rates	169
9.2.1	Receive Sampling Rate	169
9.2.2	Transmit Sampling Rate	169
9.2.3	Adjusted Sweep Time	170
9.3	Transmit Sweep Generator	171
9.4	Dual Demodulation Sweep Generator	174
9.5	Pre-Filtering	174
9.5.1	Cancelling Filter Delay	175
9.6	Multiplication	176
9.7	Band Pass Filter and Downsampler	178
9.7.1	Decimation	178
9.7.2	Bandpass Filtering	180
9.7.3	Decimation of CTFM Signals	180
9.7.4	Frequency and Distance of Flight	184
9.7.5	Design Considerations	184
9.7.6	Output Signals	185
9.8	Software Implementation	185
9.9	Range Resolution	188
9.9.1	Time-Bandwidth Measurements	188
9.9.2	Theoretical Resolution Using Frequency Width	190
9.9.3	Theoretical Resolution Using Time Width	190
9.9.4	Summary of Range Resolution Results	191
9.10	Conclusion	191
9.10.1	Acknowledgements and Contributions	192
10	DFT Spectrum Estimation	193
10.1	The Fourier Transform	193
10.1.1	Time and Frequency Sampling	194
10.2	Relationship Between the DFT and CTFM	195
10.2.1	Signal Time	196
10.2.2	Spectral Line to Distance of Flight	196
10.2.3	DOF Quantisation	196
10.2.4	Minimisation of DOF Quantisation	197
10.3	Echo Extraction	198
10.3.1	Echo Time Selection	198
10.3.2	Discrete Fourier Transform	199
10.3.3	Peak Finder	199
10.3.4	Peak Interpolator	201
10.3.5	Frequency to TOF Converter	204
10.3.6	User Interface	204
10.4	Conclusion	204

10.4.1	Acknowledgements and Contributions	206
11	Auto-Regressive Spectral Estimation	207
11.1	Theory	208
11.1.1	The Z Transform	208
11.1.2	Rational Model for Spectral Estimation	209
11.2	Yule-Walker Spectral Estimator	212
11.2.1	Linear Predictive Filter	212
11.2.2	Autocorrelation of a Sequence	215
11.2.3	Spectral Estimation	216
11.3	Least Squares Modified Yule-Walker Spectral Estimation	217
11.4	Implementation	219
11.4.1	Echo Time Selection	219
11.4.2	Autocorrelation Estimator	220
11.4.3	Yule-Walker Equation Solver	221
11.4.4	Spectral Peak Finder	221
11.4.5	Frequency to TOF Converter	224
11.4.6	User Interface	224
11.5	Conclusion	224
11.5.1	Acknowledgements and Contributions	226
12	The Speed of Sound	227
12.1	Literature	227
12.2	Formulae for Calculation	229
12.2.1	Poole's Formulae	229
12.2.2	Cramer's Formula	230
12.2.3	Dispersion	230
12.3	Error Analysis	231
12.3.1	Relative Importance of Error Terms	231
12.3.2	Combining Errors	231
12.4	Calibrating the Environment Sensor	232
12.4.1	The Temperature Scale	233
12.4.2	Humidity	235
12.4.3	Pressure	235
12.4.4	Experiments	236
12.4.5	Uncertainty	236
12.5	Conclusion	236
12.5.1	Acknowledgements and Contributions	236
13	DOF and Power Measurement Performance	239
13.1	The Experiments	240
13.1.1	Five Types of DOF Sub-system	240
13.1.2	Configurations	241

13.1.3	The Observations	241
13.2	The DOF Precision Experiment	243
13.2.1	DOF Precision Analysis	243
13.2.2	Power Precision Analysis	254
13.2.3	Cross-correlation Analysis	262
13.2.4	Noise Level Analysis	265
13.2.5	Summary of DOF Precision Results	276
13.3	The Range Resolution Experiment	278
13.3.1	Matched Filter QD Resolution Results	279
13.3.2	Matched Filter UD Resolution Results	281
13.3.3	CTFM-DFT Resolution Results	283
13.3.4	CTFM-YW Resolution Results	287
13.3.5	CTFM-LSMYW Resolution Results	290
13.3.6	Summary of Resolution Results	290
13.4	Conclusion	295
13.4.1	Acknowledgements and Contributions	296
IV	Range and Bearing Sub-System	297
14	Range and Bearing Calculation	299
14.1	Inter-Aural Distance Difference	301
14.1.1	Target Types	301
14.1.2	The Effect of Ignoring the Target Type	307
14.1.3	Range and Bearing Precision	309
14.1.4	IDD Bearing Precision Experiment	318
14.2	Inter-Aural Power Difference	323
14.2.1	Target Types	324
14.2.2	Beam Power Modelling	324
14.2.3	Log Power Difference	327
14.2.4	IPD Bearing Precision	330
14.2.5	IPD Calibration Experiment	331
14.2.6	Targets Lying Outside the Horizontal Plane	337
14.3	Correlation Between IDD and IPD	337
14.4	Conclusion	337
14.4.1	Acknowledgements and Contributions	340
15	Correspondence Problem	341
15.1	DOF Window	344
15.2	All Pairings	345
15.3	Correspondence Solvers	345
15.3.1	Simple Method	345
15.3.2	Statistical Method	347

15.4	Combining IDD and IPD Bearings	350
15.5	Simulation Experiments	353
15.5.1	Modelling	353
15.5.2	Noise-Free Echoes with Simple Solver	354
15.5.3	Noisy Echoes with Simple Solver	356
15.5.4	Noisy Echoes with Statistical Solver	356
15.5.5	Unresolved Noisy Echoes with Statistical Solver	358
15.5.6	Discussion	360
15.6	Bearing Resolution	362
15.7	Conclusion	362
15.7.1	Acknowledgements and Contributions	363
V	Conclusion	365
16	Conclusion	367
16.1	Summary of Aims and Results	367
16.2	Comparison of Results with Other Systems	371
16.2.1	Range	371
16.2.2	Bearing	374
16.2.3	Spectral Estimates	375
16.2.4	Summary	376
16.3	Limitations	376
16.4	Discussion of Contributions	377
16.5	Further Work	378
VI	Appendices	379
A	Transducer Models	381
A.1	Transduction Model	381
A.2	Radiating Plane Piston Model	384
A.3	Kuc-Siegel Impulse Model	388
A.3.1	Single Point Receiver	389
A.3.2	Rotated Two Point Receiver	390
A.3.3	On-Axis Transducer	390
A.3.4	Rotated Transducer	391
A.3.5	Comparison with the Radiating Plane Piston Model	393
B	Matched Filter Theory	397
B.1	Derivation of Matched Filter Criteria	397
B.2	Signal Selection	401

C	Optimisation of AR Spectral Estimators	405
C.1	Properties of Reversals	405
C.2	Durbin Algorithm	405
C.3	LSMYW Optimisation	410
C.4	Checking the Code	411
D	Geometrical Sonar Simulation	413
D.1	Review	413
D.2	This Work	413
D.3	Operation	413
D.4	Example	416
D.5	Effect of Non-90° Corners	416
D.5.1	Acute Corners	416
D.5.2	Obtuse Corners	416
D.6	Conclusion	420
E	Atmospheric Effects	421
E.1	Variables	421
E.2	Saturation Pressure	422
E.3	Molecular Vibrations	423
E.4	Absorption	423
E.5	Speed Of Sound	424
E.5.1	Uncertainties	426
E.6	Correction for Dispersion	426
F	Calibration Device	429
F.1	Theory	429
F.2	Design	429
F.2.1	Manufacture	429
F.3	Results	431
F.3.1	Measuring Target DOFs	431
F.3.2	Measuring Time Of Flight	431
F.3.3	Slope Results	433
F.4	Problems	433
G	Improved Head Design	435
G.1	Planes and Ninety Degree Corners	435
G.2	Edge Reflectors	436
G.3	Error Results	436
H	Modelling of Transducer Path Angles	439
H.1	Fitting for the Gaussian Parameters	439
H.2	Transducer Path Angles	440
H.3	Gaussian Products	443

H.4	Modelling the Combined Beam Pattern	445
H.5	Solving For Transducer Gaussian Parameters	446
H.6	Conclusion	447
I	Curve Fitting with Error Analysis	449
I.1	The Nature of Errors	449
I.2	Error Analysis of Multi-Parameter Functions	450
I.3	Curve Fitting	452
I.3.1	Weighted Measurements	453
I.3.2	Curve Fitting Solution	454
I.4	Meaning of the Variance Covariance Matrix	455
I.5	Forward Error Calculation	456
I.6	Reverse Error Calculation	457
I.7	Significance of the Regression	459
I.8	Weighted Averaging	460
I.9	Conclusion	460
	Bibliography	461

List of Figures

1.1	The KASPA navigational aid for the blind	5
1.2	Alternative ways of building a sonar DOF sub-system	7
2.1	Precision and accuracy	14
2.2	Resolution with sinc functions	15
2.3	Infra-red beacon sensor	16
2.4	Scanning laser range-finder	17
2.5	The Polaroid transducer	21
2.6	Geometry of a single transducer sonar system	23
2.7	Research CTFM probe system	24
2.8	A CTFM sonar head with a transmit beam-forming array	30
2.9	Sonar reflection	32
3.1	Overview of a simple CTFM sonar system.	35
3.2	Example of a frequency sweep	36
3.3	Transmit and receive signals for a single target	37
3.4	The operation of a single demodulation CTFM sonar.	39
3.5	Multiple echoes and demodulation output signals	41
3.6	The operation of a dual demodulation CTFM sonar	43
3.7	Parameters of dual demodulation sweeps	44
4.1	Relationships between time and frequency domains	50
4.2	The reference path of the sonar signal	51
4.3	Received echo power as a function of range	55
4.4	Graphs of the terms making up the sonar equation	56
4.5	Method of virtual sources for finding the path from a plane reflector	61
4.6	Gabor elementary signals	64
4.7	The effect of receive processors upon the time bandwidth product	66
5.1	The equipment used to support the sonar experiments.	74
5.2	The chirp generator card	76
5.3	Reconstruction filter circuit and response	77
5.4	Chirp capture card	79
5.5	Dialogue for generating sonar waveforms for experimentation	81
5.6	Chirp capture dialogue	81
5.7	The precision positioner	82

5.8	The dialogue used to control the positioner.	83
5.9	Graph displayed by the positioner control software	84
5.10	Environment sensor photograph	85
5.11	Environment sensor circuit	86
5.12	Environment sensor program	87
5.13	The feature extraction dialogue	89
5.14	The universal IO selector dialogue	90
5.15	Auxiliary equipment	92
6.1	Construction of old Kay transducer	96
6.2	The backing plate of the Kay transducer	97
6.3	Kay transducer insert	97
6.4	Movement of the transducer diaphragm	98
6.5	The radiating plane piston model	100
6.6	Sound pressure field calculated according to the radiating plane piston model	100
6.7	Beam pattern calculated according to the radiating plane piston model.	101
6.8	A beam spectrogram calculated according to the radiating plane piston model	102
6.9	Beam directivity	105
6.10	Original transmit circuit.	106
6.11	The receiver circuit	107
6.12	Modified transmit circuit.	108
6.13	Experimental on-axis response of the Kay transducer.	109
6.14	Three-dimensional plot of the Kay transmitter beam-forming response	110
6.15	Comparison of beam spectrograms for Kay transducer.	112
6.16	Experimental on-axis Polaroid frequency response	113
6.17	Experimental Polaroid beam spectrogram.	115
7.1	Horizontal dual-sensor geometry	121
7.2	Obstacles in the environment of a <i>robotic agent</i>	122
7.3	Vertical sensing area	123
7.4	Effect of transducer size on beam power	124
7.5	Effect of transducer size on beam width	124
7.6	Responses of transducers of various sizes to a CTFM signal	125
7.7	Beams of the three transducers	127
7.8	The extent of the field of binaural audition	128
7.9	IPD curves for various splay angles	130
7.10	Sensor head design	131
7.11	Sensor head photograph	132
7.12	Geometry of the sensor head	132
7.13	Crosstalk signals	134
7.14	Crosstalk paths	135
7.15	Sensor head internal photographs	136

8.1	The matched filter problem	140
8.2	Autocorrelation of a Barker code	142
8.3	Overview of Barker code matched filter sonar system	143
8.4	Barker coded transmit signal, with autocorrelation and spectrum	144
8.5	Experimental Barker coded signal, matched filter output and spectrum	146
8.6	Barker coded signal as a convolution	147
8.7	Operation of the three-step matched filter procedure	149
8.8	Unit sampled one-cycle sine-wave	150
8.9	Signals from quadrature demodulation	152
8.10	Block diagram of unsynchronised demodulation	153
8.11	Dual synchronised demodulation	154
8.12	Signal spectrum before and after the absolute value operation	154
8.13	Filter for removing the carrier	155
8.14	Signals from unsynchronised demodulation	156
8.15	Peak validation	157
8.16	Curve fitting	158
8.17	Settings for the matched filter quadrature demodulation sonar system	163
9.1	Overview of the CTFM sonar system	166
9.2	Digital signal processing for dual demodulation	168
9.3	Adjusting the sweep time T_{sw}	170
9.4	Mis-matched joining zone	172
9.5	Transmit signal spectrum	173
9.6	Dual demodulation signal spectrum	174
9.7	Received signal spectrum	175
9.8	Filter used to pre-filter received echoes	176
9.9	Received signal spectrum after pre-filtering	177
9.10	Signal frequencies after multiplication	179
9.11	Decimation by four	181
9.12	Decimation filter band design	182
9.13	Map of output frequency space	182
9.14	The downsampling filter	183
9.15	Demodulated echo spectrum	185
9.16	CTFM tonal echo signals	186
9.17	User interface module for CTFM dual demodulation	187
10.1	Relationship between time and frequency domains	194
10.2	Overview of DFT spectral estimation and target detection	198
10.3	Experimental DFT output	200
10.4	Closeup of DFT spectrum	203
10.5	Settings for the DFT and peak extraction modules	205
11.1	The Z domain	208

11.2	An auto-regressive system	211
11.3	Input and output of a linear predictive filter	212
11.4	Relationship between the LPF and the PEF.	214
11.5	The prediction error filter and the AR model	216
11.6	A poor least squares fit using the minimum number of data points	217
11.7	Steps performed in Yule-Walker spectral estimation and peak finding	219
11.8	Block diagram of the autocorrelation estimator	220
11.9	Computational complexity of autocorrelation	221
11.10	Steps performed by the spectral peak finder.	222
11.11	Relationship between the pole and the position of the spectral peak	223
11.12	Settings for the YW spectral estimator and peak extraction module	225
12.1	Errors in the speed of sound	232
12.2	Temperature transducer mounted behind fan	234
12.3	The LM335 calibration graph	234
12.4	The relative humidity sensor calibration graph	235
12.5	Variation in speed of sound, LM335 transducer	237
13.1	Two experimental configurations to characterise DOF measurement performance	240
13.2	The configuration used for range and bearing precision experiments	242
13.3	The configuration used for the range resolution experiment	242
13.4	Results of the DOF precision experiment, left channel	245
13.5	Results of the DOF precision experiment, right channel	246
13.6	Environmental parameters for the MF range precision experiment	247
13.7	Environmental parameters for the CTFM range precision experiment	248
13.8	The effect of the choice of speed of sound calculation upon the measured DOF .	250
13.9	Analysis of speed of sound compensation	255
13.10	Results of the left channel echo power precision measurements	257
13.11	Results of the left channel echo power precision measurements	258
13.12	Returned echo power estimated by the CTFM-YW system	260
13.13	Correlation between the left and right channels	263
13.14	MF-QD noise level	268
13.15	MF-UD noise level	269
13.16	CTFM-DFT noise level	270
13.17	CTFM-YW noise level	272
13.18	CTFM-LSMYW noise level	273
13.19	Secondary reflections from the positioner base	279
13.20	Legend for range resolution results	279
13.21	Resolution results for the MF-QD sonar system	280
13.22	Noise level in the MF-QD sonar system for the resolution experiment	282
13.23	Raw echo waveform for the matched filter sonar	283
13.24	Resolution results for the MF-UD sonar system	284
13.25	Output waveforms of the matched filter sonar system	285

13.26	Resolution results for the CTFM-DFT sonar system	286
13.27	Resolution results for the CTFM-YW sonar system	288
13.28	Spectra from the CTFM-DFT and YW sonars	289
13.29	Resolution results for the CTFM-LSMYW sonar system	291
13.30	Spectra from the CTFM-LSMYW sonar	292
14.1	The configuration used to characterise bearing measurement performance	300
14.2	Three types of sonar reflector	301
14.3	Geometry of a plane reflector	302
14.4	Geometry for an edge or point reflector	304
14.5	Corner and plane reflectors with virtual sources	306
14.6	Edge reflector with pseudo source	307
14.7	Error due to considering a point reflector as a plane reflector	308
14.8	Conversion of uncorrelated DOF measurement noise to polar coordinates	311
14.9	Conversion of DOF measurements with correlated errors to polar coordinates	312
14.10	Isopleth plot of bearing standard deviation σ_θ	314
14.11	Isopleth plot of the range standard deviation σ_r	315
14.12	Range and bearing precision of MF sonars as a function of range and bearing	316
14.13	Range and bearing precision of CTFM sonars as a function of range and bearing	317
14.14	Configuration used in the bearing precision experiment	318
14.15	Results for bearing calculation by IDD with $d_R = 15\text{ mm}$	319
14.16	Results for bearing calculation by IDD with $d_R = 13.5\text{ mm}$	321
14.17	Variation in echo amplitude with target angle	324
14.18	Variation in echo power with target angle, in decibels	326
14.19	Residual analysis of the IPD Gaussian fits	326
14.20	Log power difference between right and left channels	328
14.21	Residuals from the IPD fitting	329
14.22	Three possible root configurations encountered when solving for bearing by IPD	330
14.23	IPD Bearing precision as a function of power precision and correlation	332
14.24	Summary of results for IPD bearing calculation	333
14.25	Received echo power vs target bearing for CTFM-DFT sonar	335
14.26	Received power vs target bearing for CTFM-YW sonar	335
14.27	Correlation between θ_{IDD} and θ_{IPD}	338
15.1	A scene with three targets, and the echoes	342
15.2	Road-map of the solution of the correspondence problem	343
15.3	Difference between right and left channel DOF as a function of range	344
15.4	Difference between right and left channel DOF as a function of bearing	345
15.5	The DOF window of echo L1	346
15.6	Targets found using good and bad echo pairs	346
15.7	Isopleth plot of the 95% bearing difference rejection criterion	349
15.8	Isopleth plot of the standard deviation of the combined IDD and IPD bearing	352
15.9	Simple correspondence solver with noise-free measurements	355

15.10	Simple correspondence solver correct echo pairings, measurements with noise	357
15.11	Duplicate echoes output by the statistical correspondence solver	359
15.12	Duplicate echoes output by the statistical correspondence solver	361
A.1	Equivalent circuit of a transducer	382
A.2	Sound waves radiating from a baffled point source	385
A.3	Geometry for circular plane piston model	386
A.4	Point transmitter and single point receiver.	389
A.5	Point transmitter and two point receivers.	390
A.6	Disk divided into infinitesimal elements.	391
A.7	Angled disk divided into infinitesimal elements.	392
A.8	Geometry for change of variables.	394
B.1	A graphical dictionary of functions and their autocorrelations	402
D.1	Construction of intermediate virtual sources	415
D.2	2D model example output for a corner with trimming	417
D.3	Echo paths through an acute corner	418
D.4	Echo paths through an obtuse corner	419
E.1	Absorption of sound in air as a function of the environmental parameters	425
E.2	Plots of the speed of sound in air as a function of the environmental parameters	427
F.1	Design of the calibration device	430
F.2	Complete calibration device	430
F.3	Calibration device dimensions	431
F.4	Calibration device TOFs measured by the sonar	432
G.1	The improved head design with $d_T = 0$	435
G.2	Errors due to considering a point reflector as a plane reflector	436
H.1	Sonar paths from the transmitter to the receivers	441
H.2	Closeup of the sonar path showing the angles subtended to transducers	441
H.3	Behaviour of specular reflection angles β_1 and β_2	443
I.1	Straight line fit	453
I.2	Ellipsoid of standard deviation	455
I.3	Ellipsoids of standard deviation with varying correlation	456
I.4	Error bounds on the fitted line	457
I.5	Reverse estimated error bounds	458

List of Tables

2.1	Properties of infra-red LED sensors.	16
2.2	Properties of the SICK scanning laser range-finder.	18
2.3	Properties of MMW radar as reported by Langer	20
2.4	Properties of the Polaroid ultrasonic range-finder kit.	22
2.5	Properties of the single receiver direction measuring system built by Yata.	24
2.6	Properties of the ring sensor system built by Yata.	25
2.7	Properties of the Kleeman and Kuc impulse ultrasonic sensor system.	27
2.8	Properties of the matched filter sonar system built by Peremans.	28
2.9	Properties of the scanning array CTFM sonar systems built for robotic use by Kay.	29
2.10	Properties of the combined sonar and visual sensor system built by Akbarally.	31
6.1	Self-generated bias	106
7.1	Widths of the field of binaural audition	129
8.1	Time-Bandwidth Products for MF Signals	160
8.2	Results of theoretical MF resolution calculations.	161
9.1	Parameters used in the implementation of dual demodulation	167
9.2	Dependent variables used in this implementation of dual demodulation.	168
9.3	Corner frequencies for the decimation filter	184
9.4	Time-Bandwidth Products for CTFM Signals	189
9.5	Summary of range resolution results for the CTFM system	191
13.1	Experiments to quantify DOF and echo power measurement performance	241
13.2	The offset times t_{off} obtained with the different sonar types.	249
13.3	Measured differences from the speed of sound, c_ϕ	251
13.4	Results for the precision of the different sonar systems	254
13.5	Received echo power precision	259
13.6	Normalised measured source level, A_1 , at 1 m	261
13.7	Air absorption, α' , in $\text{dB} \cdot \text{m}^{-1}$	262
13.8	Summary of noise level results.	274
13.9	Predicted maximum range	275
13.10	Range resolution results	293
14.1	Estimated range precision	313

14.2	IDD bearing accuracy and precision results	323
14.3	IPD bearing precision results	336
15.1	Bearing calculated by IDD and IPD for each possible pairing of echoes	347
15.2	Computed minimum bearing error rejection criterion	348
15.3	Standard deviations of combined IDD and IPD bearing estimates	351
16.1	Results for the performance criteria applied to range measurements	369
16.2	Results for the performance criteria applied to bearing measurements	370
16.3	Results for the qualitative performance criteria applied to range measurements	371
16.4	Signal durations and length in air of the various sonar systems.	373
A.1	Analogies used in transduction analysis	382
A.2	Variables used in equivalent circuit	383
E.1	Cramer's coefficients for calculating the speed of sound	426
F.1	Distances used with the calibration device	432

List of Algorithms

8.1	Method of validating peaks and removing sidelobes.	157
15.1	Simple method of solving the correspondence problem.	347
C.1	The Durbin algorithm with order $O\left(\frac{3}{2}p^2\right)$ flops	409
C.2	The Durbin algorithm with order $O\left(p^2\right)$ flops	410

Glossary

Acronyms

ADC, Analogue to Digital Converter. Converts an analogue signal to digital samples, with a sampling frequency f_s , p. 78.

AR, Auto-Regressive. A system which computes the next output point based upon the previous output points. The system function, in the z plane, is composed entirely of poles, p. 9.

ARMA, Auto-Regressive Moving Average. A system which computes the next output point based upon the input points and the previous output points. The system function, in the z plane, is composed of both poles and zeroes, p. 259.

CPU, Central Processing Unit. The primary point of algorithmic execution within a computer, p. 62.

CTFM, Continuous Tone Frequency Modulated. A continuously generated tone, whose frequency is modulated, usually by a linear sawtooth function, p. 35.

CWFM, Continuous Wave Frequency Modulated. A continuously generated tone, whose frequency is modulated, usually by a linear sawtooth function. Historical term for CTFM, p. 3.

DAC, Digital to Analogue Converter. Converts digitally recorded samples, with a sampling frequency f_s , into an analogue representation, p. 52.

DFT, Discrete Fourier Transform. Method of performing a Fourier transform upon discrete data, p. 193.

DOF, Distance Of Flight. Distance of sonar flight, usually from the transmitter to the receiver, p. 6.

DSP, Digital Signal Processing. Signal processing upon a signal defined by discrete samples, p. 6.

FFT, Fast Fourier Transform. An efficient implementation of the DFT which significantly reduces the number of multiplications required compared to the direct implementation, p. 4.

FFTW, Fastest Fourier Transform in the West. An implementation of the FFT [41], p. 199.

FIFO, First In First Out. A queue structure where the longest waiting entries are removed first, p. 75.

FIR, Finite Impulse Response. A digital filter using only zeroes in its rational model, which produces a response of finite duration in response to an impulse input, p. 87.

FM, Frequency Modulated. Referring to the changing of frequency with time, p. 3.

FMCW, Frequency Modulated Continuous Wave. A frequency modulated continuous wave signal, used in radar literature to refer to CTFM, p. 3.

GUI, Graphical User Interface. The current standard for building human computer interfaces utilising graphic iconic representations of objects on a graphics screen, p. 93.

IAD, Inter-aural Amplitude Difference. A method of calculating the direction of an incoming echo by using the difference between the log amplitude of the left and right echoes, p. 323.

IDD, Inter-aural Distance Difference. A method of calculating the direction of an incoming echo by using the distance difference between the left and right echoes, p. 301.

IDFT, Inverse Discrete Fourier Transform. Method of performing an inverse Fourier transform upon discrete data, p. 220.

IFD, Inter-aural Frequency Difference. A method of calculating the direction of an incoming echo by using the frequency difference between the left and right tonal echo signals, p. 301.

IID, Inter-aural Intensity Difference. A method of calculating the direction of an incoming echo by using the difference between the log intensity of the left and right echoes, p. 323.

IIR, Infinite Impulse Response. A digital filter which uses poles and zeroes in its rational model. Such a filter may produce a response of unbounded duration when excited by an impulse input, p. 175.

IO, Input and Output. A hardware device allowing a computer to input and output signals or information, p. 73.

IPD, Inter-aural Power Difference. A method of calculating the direction of an incoming echo by using the difference between the log power of the left and right echoes, p. 323.

ITD, Inter-aural Time Difference. A method of calculating the direction of an incoming echo by using the time difference between the left and right echoes, p. 301.

LED, Light Emitting Diode. Electrical component which emits monochromatic but not coherent light, p. 13.

LPF, Linear Predictive Filter. A linear filter which seeks to predict the next point in a sequence, given all of the previous points as input, p. 212.

LSMYW, Least Squares Modified Yule-Walker. A modern method of performing spectral analysis by directly utilising more than the minimum number of discrete values of the autocorrelation of the signal, p. 217.

LTI, Linear Time-Invariant. A system which only performs linear transformations, which are not a function of time, p. 52.

MA, Moving Average. A system which computes the next output point based upon the current and previous input points. The system function, in the z plane, is composed entirely of zeroes, p. 210.

MF, Matched Filter. Refers to a sonar system utilising a filter which has been matched to the expected received signal. In this thesis the signal is a Barker coded sinewave, p. 139.

MMW, MilliMetre Wave. A form of radar using frequencies around 30 gigahertz, p. 19.

PEF, Prediction Error Filter. A linear filter formed by taking the difference between the Linear Predictive Filter (LPF) output, and the actual value of the next point in the sequence, p. 213.

PFA, Probability of False Alarm. The probability of detecting a target when no target is physically present, p. 266.

PSD, Power Spectral Density. Power per unit frequency, p. 207.

QD, Quadrature Demodulation. A technique for demodulation which assumes that the carrier frequency is approximately four times the sampling rate, p. 151.

SNR, Signal to Noise Ratio. Power per hertz of signal divided by power per hertz of noise, p. 140.

SSR, Signal to Sidelobe Ratio. The ratio between the maximum power level in the main lobe and the maximum side-lobe power level, p. 142.

STP, Standard Temperature and Pressure. The standard temperature and pressure is defined to be $T_0 = 293.15$ K and $p_{s0} = 101.325$ kPa, p. 228.

TB, Time-Bandwidth product. The product of the duration and bandwidth of a signal, p. 63.

TOF, Time Of Flight. Time of sonar flight, τ , usually from the transmitter to the receiver, p. 25.

UD, Unsynchronised Demodulation. A technique for demodulation by rectification and low-pass filtering, p. 152.

VCV, Variance CoVariance. The variance covariance matrix contains variances down the diagonal, and covariances between different terms in the off-diagonal positions, p. 455.

YW, Yule-Walker. A modern method of performing spectral analysis by directly utilising the autocorrelation of the signal, p. 212.

Functions

$f(k)$	Frequency of DFT output sample number k , in hertz, p. 195.
$f_{d0}(t)$	Upper demodulation signal frequency as a function of time, p. 43.
$f_{d1}(t)$	Lower demodulation signal frequency as a function of time, p. 43.
$f_{gaus}(\theta, a, \theta_0, b)$	Gaussian function which is fitted to the transducer bearing response, p. 325.
$f_t(t)$	Transmit frequency sweep, p. 36.
$h(t)$	Impulse response as a function of time and perhaps other variables, p. 388.
$h_{opt}(t)$	Optimal matched filter, p. 397.
$J_1(x)$	Bessel function of the first kind, p. 101.
$N(\mu, \sigma^2)$	The normally distributed probability density function, p. 451.
$R_x[l]$	Autocorrelation of a sequence, evaluated at the l th lag, p. 214.
$s(t)$	Signal as a function of time and perhaps other variables, p. 388.
$s_d(t)$	Demodulation signal as a function of time, p. 39.
$s_{d0}(t)$	Upper demodulation signal as a function of time, p. 45.
$s_{d1}(t)$	Lower demodulation signal as a function of time, p. 45.
$s_{out}(t)$	Output signal as a function of time, p. 39.
$s_r(t)$	Received echo signal as a function of time, p. 39.
$s_t(t)$	Transmit signal as a function of time, p. 38.
$t(l)$	Time of sample number l in seconds, p. 195.
$\delta(t)$	Impulse function, p. 388.
$\phi_{d0}(t)$	Upper demodulation signal phase as a function of time, p. 43.
$\phi_{d1}(t)$	Lower demodulation signal phase as a function of time, p. 43.
$\phi_t(t)$	Transmit signal phase, p. 37.
$\tau_g(\omega)$	Group delay of a filter, in seconds, p. 175.
$\Theta(\omega)$	Filter phase delay, in radians, p. 176.
$\chi_m(\tau)$	Maximum signal to noise ratio, p. 398.
$\mathcal{P}(f)$	Power spectral density as a function of frequency, p. 209.

Greek Symbols

α	Absorption of sound by the air in nepers per metre, p. 421.
α'	Absorption of sound by the air in decibels per metre, p. 421.
α_r	Receiver skew angle in sonar head in degrees, p. 126.
β_i	Curve fit parameters, p. 452.
γ	The ratio of specific heats C_p/C_v at constant pressure and constant volume, p. 228.
Δd	Spacing between DFT spectral lines when converted to distance of flight, in metres, p. 196.
Δf	Spacing between samples in output from DFT in the frequency domain, in hertz, p. 195.
Δt	Spacing between samples in input to DFT in the time domain, in seconds, p. 195.
θ	Angle from x axis of coordinate system, usually used with r to form a coordinate (r, θ) .
θ'	The bearing from P_{HC} to T_V , p. 303.
λ	Wavelength in metres, p. 101.
μ	Sweep rate in hertz^2 , p. 37.
ρ_0	Volume density of material in $\text{kg} \cdot \text{m}^{-3}$, p. 385.
τ	Sonar time of flight in seconds, p. 39.
ω	Radial frequency in $\text{rad} \cdot \text{s}^{-1}$, $= 2\pi f$, p. 385.
ω_s	Scaled radial frequency in $\text{rad} \cdot \text{s}^{-1}$, $= 2\pi f/f_s$, p. 209.

P

PAL, Programmable Array Logic. A programmable component for performing logical operations upon digital signals, often with reference to a clock signal, p. 78.

R

RMS, Root Mean Square. The root mean square of a set of observations is equivalent to the standard deviation, p. 267.

Notations

\oplus	Anti-logarithmic addition, $A \oplus B = \log \{ \exp(A) + \exp(B) \}$, p. 50.
*	Complex conjugate, z^* , of z , p. 194.
*	Convolution, $f(t) * g(t)$, of two functions $f(t)$ and $g(t)$, p. 55.
*	Correlation, $f(t) \star g(t)$, of two functions, $f(t)$ and $g(t)$, which is defined in terms of the convolution as $f(-t) * g(t)$, p. 141.
$\hat{}$	Estimated value \hat{x} of x , p. 213.
$E\{ \}$	Expectation value, $E\{g(x)\}$, of a function, $g(x)$, p. 213.
\mathcal{F}	Fourier transform, $\mathcal{F}\{g(t)\}$, of a function, $g(t)$. The transformed function is capitalised, and the independent variable is ω , p. 193.
$\text{Im}\{ \}$	Imaginary part of a complex number, p. 387.
	Mean value, \bar{x} , of variable x , p. 67.
$\text{Re}\{ \}$	Real part of a complex number, p. 387.
\sim	Reversal of a matrix or vector, p. 405.
T	Transpose of a Matrix or vector, p. 218.
\mathcal{Z}	Z transform, $\mathcal{Z}\{g(t)\}$, of a function, $g(t)$. The transformed function is capitalised, and the new independent variable is z , p. 208.

Symbols

A	Transducer sensitivity (dimensionless), p. 390.
$\alpha_{v,N}$	Molecular absorption caused by the vibrational relaxation of N_2 , in nepers per metre, p. 423.
$\alpha_{v,O}$	Molecular absorption caused by the vibrational relaxation of O_2 , in nepers per metre, p. 423.
a	Transducer radius in metres, p. 101.
B	Signal bandwidth in hertz, p. 42.
B_d	Demodulated signal bandwidth in hertz, p. 42.
C_ϕ	Phase offset in radians, p. 45.
c	Speed of sound in air in $\text{m} \cdot \text{s}^{-1}$, p. 39.
c_0	Speed of sound in air in $\text{m} \cdot \text{s}^{-1}$, at zero frequency, p. 424.

c_1	Speed of sound in air in $\text{m} \cdot \text{s}^{-1}$, calculated using temperature $c_1(T_a)$, p. 229.
c_2	Speed of sound in air in $\text{m} \cdot \text{s}^{-1}$, calculated using temperature and pressure $c_2(T_a, h_r)$, p. 229.
c_{ab}	Correlation coefficient between two quantities a and b , p. 264.
c_{int}	Statistical confidence interval, where 0.66 corresponds to one standard deviation, p. 450.
c_ϕ	Speed of sound in air in $\text{m} \cdot \text{s}^{-1}$, calculated using temperature, humidity, pressure and frequency, p. 426.
d	Distance of flight in metres, referring to the total path length out and back, p. 39.
d_{max}	Maximum operating sonar travel distance (not range) in metres, p. 42.
d_R	Offset distance of receivers in sonar head from P_{HC} in metres (base length), p. 133.
d_T	Offset distance of transmitter in sonar head from P_{hc} in metres, p. 133.
f	Frequency in hertz, p. 101.
f_0	Transmit frequency sweep start frequency in hertz, p. 37.
f_{00}	Upper demodulation sweep start frequency in hertz, p. 43.
f_{01}	Lower demodulation sweep start frequency in hertz, p. 44.
f_1	Transmit frequency sweep end frequency in hertz, p. 37.
f_{10}	Upper demodulation sweep end frequency in hertz, p. 43.
f_{11}	Lower demodulation sweep end frequency in hertz, p. 44.
f_{bus}	Bus frequency common to both chirp generator and capture cards, p. 169.
f_c	Carrier frequency for the Barker coded signal, p. 142.
f_e	Enhancement factor, p. 426.
f_{maxd}	Un-demodulated frequency in hertz corresponding to d_{max} , p. 167.
f_{mind}	Un-demodulated frequency in hertz corresponding to 0 DOF, p. 167.
$f_{r,N}$	Vibrational relaxation frequency of N_2 , in Hz, p. 423.
$f_{r,O}$	Vibrational relaxation frequency of O_2 , in Hz, p. 423.
f_s	Sampling frequency of a digitally sampled signal, p. 209.
f_{sd}	Sampling frequency after decimation, p. 183.

f_{sr}	Sampling frequency of received signal, in hertz, p. 169.
f_{st}	Sampling frequency of transmitted signal, in hertz, p. 169.
h	Absolute humidity as a mole fraction, in %, p. 422.
h_r	Relative humidity, the ratio between the partial pressure of water vapor and the saturation pressure, expressed as a number between 0 and 100 (%), p. 229.
j	Imaginary number, $\sqrt{-1}$, p. 193.
k_{mf}	Matched filter scaling constant, p. 399.
k_w	Wavenumber or inverse wavelength in $\text{rad} \cdot \text{m}$, p. 101.
l	A positive integer, p. 171.
M	The molar mass, in $\text{kg} \cdot \text{mol}^{-1}$, p. 228.
m_{rr}	CTFM maximum range ratio (dimensionless), p. 42.
N_{DFT}	Number of samples in the time domain input to DFT, p. 194.
N_{spec}	Number of samples in the frequency domain output from DFT, p. 195.
n_{dr}	Divisor of bus clock to obtain receive sampling frequency, p. 169.
n_{dt}	Divisor of bus clock to obtain transmit sampling frequency, p. 169.
n_{sr}	Number of samples received in one sweep period T_{sw} , p. 169.
n_{st}	Number of samples transmitted in one sweep period T_{sw} , p. 169.
p	Number of poles in rational model, p. 209.
p_0	Root-mean-square air pressure of a radiating source in pascals, p. 422.
p_d	Dynamic air pressure in pascals, p. 384.
P_{HC}	Head centre point, halfway between the two receivers, p. 131.
p_s	Static air pressure in pascals, p. 228.
p_{s0}	Reference ambient air pressure, p. 422.
p_{sat}	Saturation vapor pressure in pascals, p. 230.
p_z	Complex pressure field in pascals, p. 384.
Q	Source strength in $\text{m} \cdot \text{s}$, p. 385.
q	Number of zeros in rational model, p. 209.
R	The universal gas constant, in $\text{J} \cdot \text{mol}^{-1} \cdot \text{K}^{-1}$, p. 228.

r	Radius from centre of coordinate system, usually used with θ to form a coordinate (r, θ) , p. 302.
r'	Used in the derivation of the radiating plane piston model (Section A.2) for the distance from the point (r, θ) to an infinitesimal element. Also represents approximately the distance from the head centre point, P_{HC} , to the target. The precise definition is that the length of $P_{HC}T_V$ is equal to $2r'$, p. 386.
r_{max}	Maximum CTFM range in metres, p. 42.
T_0	Reference ambient temperature, p. 422.
T_3	Temperature of the triple point of water in Kelvin, p. 422.
T_a	Absolute air temperature in Kelvin, p. 229.
T_{DFT}	Duration of samples in time domain over which the DFT is taken, p. 195.
T_{do}	CTFM output signal duration in seconds, p. 41.
T_{ice}	Definition of the melting point of ice, which is also used to convert from ° Celsius to Kelvin, p. 229.
T_{max}	The time at which the echo from a target at r_{max} will arrive, in seconds, p. 41.
T_{sig}	Time over which all echoes are present in the demodulated CTFM output, in seconds, p. 42.
T_{sw}	Sweep time in seconds, p. 37.
t	Signal time in seconds, p. 384.
t_p	Periodic time in seconds, used in the definition of periodic waveforms, p. 37.
U_0	Speed amplitude of vibrating surface in $\text{m} \cdot \text{s}^{-1}$, p. 386.
z	A complex number, composed of real and imaginary parts, p. 208.

Acknowledgements

The research and preparation involved in the production of a Ph.D. thesis are not possible without the support and encouragement of a great many people. I was fortunate to be surrounded by people who willingly offered their time and energies to assist and support me in a great many ways. The following list is by no means exhaustive, and I apologise to those I have inadvertently omitted.

I sincerely thank my supervisor, Ass. Prof. Philip McKerrow, for his patient guidance and advice throughout this endeavour. His encouragement and friendship have been invaluable.

I am grateful to my colleagues in the Intelligent Robotics Laboratory, and the School of Information Technology and Science in general, at the University of Wollongong, for their advice and their useful and challenging discussions during the course of my research. I have particularly enjoyed the offerings of Danny Ratner.

Numerous people have offered their expert advice and encouragement along the way. My colleagues from the Intelligent Robotics Laboratory, University of Tsukuba, Japan, started me on the path to ultrasonic sensing and an understanding of the inner workings of the Yamabico mobile robots. The technical staff at the School of Information Technology and Science, University of Wollongong, have been of great help. Janice Sendt, of Thales Underwater Systems, Sydney, reviewed my work at various stages and offered much advice and encouragement to her sorcerer's apprentice. Prof. John Moore, of the Australian National University, Canberra, advised me about different spectral analysis techniques I should consider. Ass. Prof. Peter Morgan and Dr Ian Lisle, of the University of Canberra, were of great assistance with their advice, Ass. Prof. Morgan concerning statistical techniques and Dr Lisle concerning mathematical derivations. Dr. Penny Probert-Smith, University of Oxford, and Dr. Akihisa Ohya, University of Tsukuba, provided much valuable feedback on the content of the thesis. Segher Boessenkool, from the Ogg Vorbis project, discussed signal processing and code optimisation with me, while John Levon and others of the LyX mailing list provided fantastic help and support with LyX and \LaTeX problems. Javier Bezos even went so far as to add features to the \LaTeX gloss package specifically to meet my requirements.

I acknowledge the financial assistance of the Department of Education, Employment, Training and Youth Affairs, in the form of a scholarship under the Targeted Institutional Links program. I also acknowledge the financial assistance of the University of Wollongong, which enabled me to attend an international conference and visit several institutions around the world to discuss my work with colleagues.

My friends in Canberra, Wollongong, Campbelltown and around the world, have offered a listening ear, a chance to unwind and to think out loud, and their never-ending encouragement,

love and prayers over the last few years. Their support has helped me to make it this far. Graham Kettlewell has been of considerable help, spending many hours listening to my confessional debugging of computer code. His advice on the design and assistance with the construction of the calibration device was particularly valuable. Peter and Wendy Lee, Natasha Hanney, and other members of the Rosemeadow Anglican Church, have always been a great source of encouragement. Their constant prayers and concern for my progress and well-being have been greatly appreciated.

Thanks are also due to my colleagues at Thales Underwater Systems for their encouragement and practical support over the last year. It is fantastic to work in a place where I can continue to develop and advance my working knowledge of sonar systems. The support of management during the last few months in allowing me to leave early to work on the thesis is greatly appreciated.

The unconditional love, support and prayers of my family have made these past years possible. They have done all that was in their power to assist me at every turn. In particular, I thankfully acknowledge the loving support of my parents, Graeme and Helen Stanley. My father assisted with the printing and binding of this thesis, while my mother offered a listening ear, her encouragement and prayers. I am grateful to my sister-in-law, Roanna Carn, for assistance with proof-reading. I am indebted to my parents-in-law, Peter and Carol Morgan, who welcomed me into their home for six months, travelled to Campbelltown each weekend for weeks on end, adjusted their routines around my work, and spent countless hours proof-reading the many drafts and making suggestions about style. Their efforts have undoubtedly improved my grammar, style, and even spelling, and have greatly improved the quality of the completed volume. They have also constantly encouraged me, doing far more than I could ever have asked or expected of them.

Finally, without the love and support, patience and proof-reading of my long-suffering and beloved wife, Tammy, this thesis would not even have been a possibility. Tammy has put up with not just one but two extended periods of intense thesis writing, during which she has put up with far less time and attention than she deserves. Furthermore, she has interrupted her own studies to support the completion of this thesis. Tammy, thank you for everything.

To all of you, my appreciation and gratitude. Thank you can never be enough for all the help and support I have had in researching and preparing this thesis, and the untold hours of assistance to remove the spelling mistakes, inconsistencies and errors in the drafts. Ultimately however, any mistakes that remain in the final version are my own.

*The LORD is my rock, my fortress and my deliverer;
my God is my rock, in whom I take refuge,
my shield and the horn of my salvation.*

2 Samuel 22:2-3 (NIV)

Chapter 1

Introduction

1.1 Aim

The English word *robot* is relatively new. It entered the language in 1923 and is derived from a Slavonic stem which is related to the Greek word for work [88]. A common meaning is

A machine devised to function in place of a living agent; one which acts automatically or with a minimum of external impulse.

The *living agent* referred to above is usually a human being who is equipped to sense the environment in which he is situated, to remember and recognise objects in this environment, to perceive geometrical relationships between these objects, and to make reasoned use of this information.

By way of example the *living agent* plays a sport, called orienteering, which exercises the skills of sensing and navigation. The *living agent* is given a map and a compass, placed in an unfamiliar environment, and expected to navigate his way to a series of points. To succeed, the *living agent* must utilise sensory information to localise his position based on the observed position of landmarks in his environment, plan a route to the goal, navigate along the planned route, and avoid any obstacles which may appear in his path.

Thus if a *robotic agent* is to sense and navigate in an environment in place of a *living agent*, it must be equipped with the appropriate sensors, especially those which emulate sight. In addition, the *robotic agent* is usually expected to construct and use its own map as it traverses and explores its environment.

This thesis aims to probe and investigate the use of two binaural sonar systems, the [Matched Filter \(MF\)](#) sonar system and the [Continuous Tone Frequency Modulated \(CTFM\)](#) sonar system, as a technique to sense the environment, in much the same way as sight is used by humans, thereby allowing the *robotic agent* to function in place of the *living agent*. Situations where this is desirable include performing dull and repetitive work, or functioning in an environment hazardous to humans such as is found in coal mines or nuclear reactors.

1.2 Goals

The *living agent* referred to above uses a complex process, usually described as vision, to navigate. This process incorporates the imaging and recognition of objects, the measurement of their distance from the agent, the measurement of a bearing relative to a fixed or defined direction, and

the identification of the texture and colour of the surface. This information is processed by a visual sensory system to provide a picture of the environment.

Visual sensors estimate the bearing of an object from the position of that object in the image. For nearby objects range is estimated by measuring the difference in position of the object between two images. When distant objects present identical projections in both images, range estimates are performed using geometric cues.

Sonar sensory systems are not unlike visual systems in that binaural range measurements are needed to obtain range and bearing information. Sonar systems compute bearing by utilising the difference between binaural range measurements, or by other novel means. This provides the two observational types which must be recorded for a two dimensional environment.

Matched filters have long been recognised in the field of radar as being a useful method for increasing the power of the transmitted signal, by lengthening the chirp, while maintaining the ability to resolve close targets. The matched filter sonar system utilised in this study replicates and extends the previous work of Peremans [108].

A measure of the ability of a sonar system to provide a suitable system for object recognition, localisation and navigation in a living agent is given by bats [126]. All members of the suborder microchiroptera use sonar to hunt and navigate. While not blind, their eyes are generally very small and their vision is poor. Bats utilise many different types of call, depending upon whether they are hunting, flying in a densely forested area, or flying in open space (Suga [128]). The calls have various features in common with CTFM sonars. Thus, CTFM sonar is a technology utilised effectively by a *living agent* to perform navigation tasks.

The goals of this thesis are therefore:

1. to build a MF binaural ultrasonic sensing system,
2. to build a CTFM binaural ultrasonic sensing system,
3. to investigate the ability of the MF and CTFM sonar systems to measure range and bearing, and to quantify their performance, and
4. to determine the best system or systems to emulate sight as a sensor for a *robotic agent*.

These goals will be demonstrated using untextured targets such as plane surfaces, corners and poles, which produce a well-defined echo replicating the transmitted signal. The consideration of textured surfaces, which have a filtering effect upon the transmitted signal, is beyond the scope of this thesis. This topic is the work of other laboratory members.

Achieving these goals will provide the basis for a sensor which should provide sufficient information to facilitate mapping and navigation, in a precise, robust, and reliable manner. In addition, it should be able to operate in both benign and hostile environments.

1.3 Continuous Tone Frequency Modulated Sonar

Continuous tone frequency modulated sonar has two major components. The first, continuous tone, refers to the uninterrupted transmission of a sinusoidal signal. The second, frequency modulation, refers to the changing of the frequency of the sinusoid with time. Air sonars typically use a

frequency range of 40-100 kHz, which falls within the ultrasonic range [146]. For this reason air sonars are commonly called ultrasonic sensors.

The **CTFM** sonar operates in three stages. The first stage is to play the **CTFM** signal through an ultrasonic transducer which projects the sound. In the second stage the sound is reflected by objects in the environment. In the third stage the reflected sound is heard by the ultrasonic receivers. The frequency modulation of the echo signal allows the time of flight to be measured. As sound propagates at a constant velocity in air, the path distance can be calculated. The theory of operation of **CTFM** sonar is described in detail in Chapter 3.

The following three sections briefly trace the development of air sonars.

1.3.1 Sonar Origins

When the Titanic sank with great loss of life in 1912, due to a collision with an iceberg, several people began to work seriously on a navigational aid using underwater sound. The first working system, constructed by Fessenden, detected the echo from an iceberg at a range of two miles in 1914.

At the onset of the First World War, in 1914, U-boats were used with great effectiveness by the Germans, because they could not be located. This prompted the British to develop the first underwater listening devices, which listened passively for the sound of the U-boat's propellers and machinery. As these were ineffective when the submarine machinery was halted, the British set about developing an active sonar echo location device. Utilising some technology from the other allies, the Asdics system was developed. The first successful test was in March 1918, when a submarine was detected at 500m. The Asdics system utilised a quartz piezoelectric transducer, and could indicate the range. However, the war ended before Asdics was deployed. (Hackmann [49] describes Asdics as a codeword. The first three letters refer to the Anti-Submarine Detection committee, while the ics was added to make it into a noun. This forms a word unrelated to the operation of the device, so as to keep its operation a secret during wartime.)

The Second World War saw the further development of the Asdics system, which gained the ability to determine the bearing of the submarine. These systems used a pulse as the transmit signal, which was shortened as development progressed so as to improve the accuracy of the echo features. At the end of the Second World War in 1945, ship based sonars could detect submarines at 1.2km, while the submerged submarine could detect a ship at 24km.

During the Second World War, there was intensive research into new sonar technologies. In the 1940s the benefits of continuously transmitting **Frequency Modulated (FM)** sonars were recognised. Kurie [86] details eleven benefits over the conventional pulse systems of the time, including being in constant contact with the target. However, the desired **FM** sonar could not be built due to the limitations of the electronics technology available. Development of the **FM** sonars flagged during the 1950s due to the introduction of chirped pulses, from radar, together with pulse compression techniques in the receiver. This led to the old **FM** sonars being called **Continuous Wave Frequency Modulated (CWFM)** sonars to differentiate them from the new chirped **FM** sonars. The **CWFM** sonars are referred to as **Frequency Modulated Continuous Wave (FMCW)** by the radar literature, with the term **CTFM** being adopted more recently.

The construction of a **CTFM** sonar requires an accurate and linear frequency modulator, and

a spectral analyser. The only spectral analyser available at the time was a bank of filters tuned to a range of different frequencies. These requirements hampered the development of CTFM sonars until digital techniques, such as the high speed digital frequency synthesiser and the Fast Fourier Transform (FFT), became available.

(Early sonar history taken from Hackmann [48]. CTFM information from Hayes [58] and de Roos [30]. History of the development of radar may be found in references [16, 129, 134].)

1.3.2 CTFM Air Sonar as a Sensory Aid for Humans

Kay [68, 69] began investigating CTFM sonars in 1959 by comparing them with pulse systems. Inspired by the bat's ability to navigate effectively in the dark, Kay sought to provide this ability to visually impaired persons. The first such system was a hand-held probe acting as a sonic torch [70]. This device gave promising results, and Kay began an extensive research program in 1966 to produce a wearable system. The first step was to build transducers which work effectively in air [90]. The second step was to display the sonar signal in a form suitable for interpretation by a visually impaired person. Rowell [120] determined that the most suitable form was to supply an auditory signal to the user through headphones. The sonar signal is demodulated to provide the auditory signal, which contains one tone for each reflecting target. Close targets produce low frequency tones while distant targets produce high frequency tones. The auditory signal is encoded in this way because the human ear is well suited for discriminating different tones, but is unsuited to detecting the duration of a tone or gap corresponding to the time of flight. Bearing information is presented to the hearer by the difference in the volume between each ear. The sonar system is worn by the user as a special pair of glasses. This couples the user's head movement with movement of the sonar beam, forming a tightly coupled perceptory system.

The product of this research, the ultrasonic glasses, was first built commercially in 1974 [71, 67]. More recently the KASPA model, which is worn as a headband, was introduced (Figure 1.1). Many blind people have successfully used these systems as a substitute for their loss of vision (Kay [74]). Users of the ultrasonic glasses have been able to perform independent navigation in their neighbourhoods by sensing and recognising hedges, moving cars, poles, and many other objects. Competent users are able to recognise composite objects which have previously only been observed separately, for example a vine growing over a trellis, or a table tennis table, described by one individual [73] as

... a table with a fence on top of it.

Furthermore Pivac (from Rowell [120]), a blind university student who learnt to use ultrasonic glasses, reports

Perhaps for the first time, since the need for vision, I have become aware of the relationship between the environment and myself. A relationship characterized by the "Self-Concept" in which I have more certainty in realising the distance, speed and grandeur of objects, moving or otherwise. Due to this factor, I find myself concentrating more on my own physical plane when moving through crowds. Previously, it was necessary to concentrate on the position of people and what they might do. But




Figure 1.1: The KASPA navigational aid for the blind, manufactured by SonicVision. This system was developed by Leslie Kay.

with the use of usg [ultrasonic glasses], I can obtain this information well in advance and therefore, am able to consider which for me would be the most suitable mode of bypassing as it were, gracefully.

Thus Pivac was able to perform path planning in a complex environment, demonstrating that **CTFM** sonars supply a sufficient amount of information to a human *living agent* to perform navigation tasks.

1.3.3 CTFM Air Sonar as a Robotic Sensory System

Given the success of **CTFM** systems as sensory organs for *living agents*, various researchers have worked on **CTFM** sensors for use by *robotic agents*. While these appear at first to be the same problem, the *robotic agent* faces extra difficulties. The sonar systems used by *living agents* present the sonar signal, after some pre-processing, to a neural processor in order to utilise the information present in the signal to recognise objects and assess their position. *Robotic agents* must replicate this object recognition and position estimation in a robust and reliable way. This task is assigned to the sensory system.

Kay [72] has attempted to build a 2D sensor suitable for use by a *robotic agent* utilising Fourier transform analysis. Davies [28] worked on building a 3D **CTFM** sonar for use by a *robotic agent* also utilising Fourier transform analysis. Politis [111] used **CTFM** sonar to recognise surface texture. Harper [55, 56, 93] used a **CTFM** sonar interfaced to a computer to recognise different species of plants based on many individual features exhibited by the echo. In these works, Politis and Harper demonstrated the ability to recognise landmarks, which are unique locations within

the environment. Landmark recognition facilitates localisation, or the estimation of the current position of the *robotic agent* on a previously recorded map. Ratner [117] worked on long distance navigation by following the border between a pathway and a lawn using a beam-forming CTFM sonar, allowing navigation between recognised map-based plant landmarks. Politis [112] has worked on recognising planes, edges and corners utilising their frequency response to a CTFM sonar.

Thus previous work has demonstrated that a variety of tasks may be accomplished by a *robotic agent* using a CTFM sonar.

1.4 Outline

This thesis aims to develop and quantify the performance of several sonar systems. The thesis is divided into five parts.

Part I reviews the foundation material from the literature upon which this research is built. In particular Chapter 2 reviews previous research on precision sonar systems for use by a *robotic agent*. This serves to inform us in our choice of sensor. Chapter 3 collates, discusses and analyses previous work on CTFM sonars for use by a *robotic agent*. Chapter 4 describes the path of the sonar signal from the signal generator in the transmitter, through the sensor head, environment, and receive processor, to the detector in the receiver. It shows how each element within the path may be modelled as a filter, and describes the effect that each element has upon the sonar signal. It also shows how the limiting precision and resolution characteristics of the sonar may be related to the system bandwidth. Thus Part I draws together the foundations from the literature for the work developed in Part III.

Part II of this thesis deals with the supporting equipment adopted to reach the goals listed in Section 1.2. Chapter 5 describes the experimental setup in which the experiments were made. The invariant nature of this setup was essential as experiments became more complex and demanding. This chapter also discusses the software architecture which was adopted to support the construction of several different kinds of sonar system by connecting together individual building blocks. Chapter 6 reviews the properties of ultrasonic transducers and derives the transducer filter function. It also contrasts the properties of two available types of transducers. Chapter 7 explains how the transducers were combined to form an integrated sensor which is capable of measuring both the range and bearing of targets, and describes the selection of the transducers.

Part III of this thesis concentrates on the sonar Distance Of Flight (DOF) sub-system, comprising the operation of the sonar signal and the extraction of individual echoes from the signal. This sub-system produces distance of flight and echo power measurements as its output. This part of the thesis considers several different solutions to this problem, as described by Figure 1.2. It begins with Chapter 8, which details the operation and construction of two matched filter sonar systems. The first system, which replicates previous work, utilises a Quadrature Demodulation system, and is designated MF-QD. The second system utilises an unsynchronised demodulator, and is designated MF-UD. Subsequent to this matched filter work, a CTFM sonar system was implemented using Digital Signal Processing (DSP) techniques, and is described in Chapter 9. The CTFM sonar system produces a signal containing one tone per target. Therefore spectral analysis is required

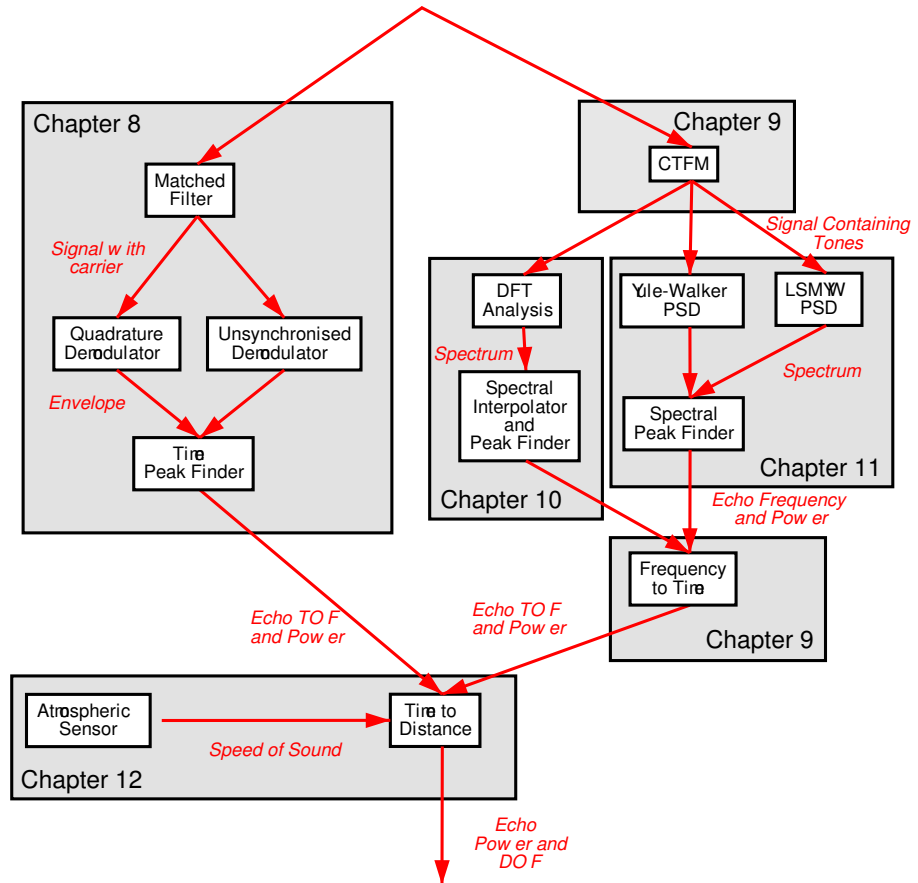


Figure 1.2: Several alternative ways of building a sonar **DOF** sub-system are evaluated in Part III of this thesis. Begin at the top and choose a complete path to reach the bottom. The stages traversed may be assembled to build a complete sonar **DOF** sub-system. This diagram illustrates only one channel. A multi-channel system would perform the same processing for each channel.

to extract information about each target. This part of the thesis investigates three different methods of performing spectral analysis of **CTFM** echoes. The first, which replicates previous work, uses a **Discrete Fourier Transform (DFT)**. The complete sonar system is designated CTFM-DFT, and is described in Chapter 10. The second, the **Yule-Walker (YW)** parametric spectral analysis method, is described in Chapter 11, and is designated CTFM-YW. A closely related third method, the **Least Squares Modified Yule-Walker (LSMYW)** parametric spectral analysis method designated CTFM-LSMYW, is also discussed in Chapter 11. Chapter 12 focuses on obtaining a reliable value for the speed of sound, and demonstrates that it is necessary to continuously compensate the sonar for variations in the speed of sound if accuracy is to be maintained. Chapter 13 quantifies and compares the performance of the five different **DOF** sub-systems, in terms of their precision, cross-correlation, and resolution.

Part IV of this thesis deals with the bearing calculation sub-system. This sub-system takes the distance of flight measurements from separate sonar channels and fuses them into a single estimate of the position of the reflecting object. The position is calculated in the form of range and bearing, thereby allowing a *robotic agent* to navigate as required in Section 1.2. In particular Chapter 14 details two independent methods of calculating the bearing to an object. The first method utilises the **Inter-aural Distance Difference (IDD)**. The second utilises the log power difference between the echoes detected by each receiver, or the **Inter-aural Power Difference (IPD)**. These pieces of information are independent. This chapter also details the results of experiments which quantify the bearing measurement performance of each of the five sonar systems. Chapter 15, the second chapter in this part, describes the method of forming a robust and reliable relationship between echoes from the left and right channels. This utilises the agreement between the two different methods of bearing calculation for a given echo relationship. It is this agreement which allows the construction of a robust and reliable sensor as required by Section 1.2. When agreement is established, the independent estimates of bearing are fused to obtain a more precise bearing estimate. The results of calculations in Chapter 15 indicate that the MF-UD sonar system provides the most robust echo relationships.

Part V of this thesis, Conclusion, discusses the results in terms of the aims and goals set in Chapter 1. The best sonar system overall, out of the five systems evaluated in this thesis, is found to be the CTFM-LSMYW sonar system.

1.5 Contributions

The aim of this thesis is to build an air sonar sensory system to enable a *robotic agent* to function in place of a *living agent*. In meeting the aim, this thesis makes several contributions to knowledge, which can be summarised as follows:

- A robust statistical quantification and comparison of the accuracy, precision and resolution of five sonar systems in range and bearing is performed.
- A robust and reliable method of solving the correspondence problem, utilising two different methods of bearing calculation, is demonstrated. Robust statistical criteria are developed for accepting or rejecting an echo pairing between the two channels.

-
- This thesis performs an in depth study of both classical and [auto-regressive](#) spectral analysis techniques as applied to [CTFM](#) sonar systems for the extraction of [distance of flight](#) and echo power data.
 - A method of self-calibrating the speed of sound continuously from temperature observations is developed.

These contributions combine with previous work to form an air sonar system which meets the goals listed in [Section 1.2](#).

Part I

Foundations

Chapter 2

Review of Robotic Location Sensing Technology

This chapter collects and summarises, from the literature, properties of various sensors which have been investigated for use by *robotic agents*. These sensors produce measurements of objects in the environment of the *robotic agent*. Each of these measurements can be appraised by several criteria: accuracy, precision, and resolution. These criteria are defined in Section 2.1.

The criteria are used to compare some of the many different types of sensors which have been studied and investigated for use by a *robotic agent*. These include:

- infra-red [Light Emitting Diode \(LED\)](#) sensors,
- scanning laser range-finders,
- microwave and millimetre wave radars, and
- ultrasonic sensors.

These sensors possess widely differing attributes and capabilities. The remainder of this chapter characterises each sensor and summarises the state-of-the-art performance for each type. Particular attention has been paid to ultrasonic sensors, providing an introduction for this research.

2.1 Criteria for Comparison

The sensory comparison focuses on range and bearing as the measured quantities. The three major criteria used to assess the measurements are defined as follows:

Accuracy: How close the measurement is to the true value. An accurate measurement is free from bias (Eisenhart [36]) (Figure 2.1).

Precision: The precision of a measurement is the standard deviation of a set measurements. Precise measurements are usually repeatable with a high probability (Eisenhart [36]) (Figure 2.1).

Resolution: The minimum separation which may be measured between two objects before the two separate objects appear to fuse into one. The Rayleigh criterion (Giancoli [44]) arbitrarily defines this to be when the maximum of one pattern coincides with the first minimum of the other, producing a dip in the sum of the two patterns (Figure 2.2).




Figure 2.1: The difference between precision and accuracy, shown with normal probability density functions. Precise measurements have a small standard deviation σ , and are therefore more likely to be close to the mean. Accurate measurements have the mean close to the true value, and are therefore unbiased.

Three additional minor criteria are assessed for completeness. These are:

Minimum: The minimum value which may be measured by the sensor.

Maximum: The maximum value which may be measured by the sensor.

Quantisation: Some sensors can only produce measurements in multiples of some quantity. This may be because of a design decision to return measurements as an integer, or due to a more fundamental property of the sensor.

According to Eisenhart [36], the accuracy of any sensor is unknown, as it is impossible to know the true value. It is possible to use statistical methods to estimate how close a measurement is to the true value (standard error of the mean, Appendix I), but most authors do not do this. Therefore the accuracy of sensors is not compared. Authors often quote the precision and call it accuracy.

The minimum and maximum bearings measurable by a sensor are always the same, so the half beam width is quoted instead. With these fundamentals in place, the precision, resolution, minimum and maximum will be listed for the range, and the precision, resolution and beam width will be listed for the bearing of each sensor. The quantisation will be noted where relevant.

2.2 Infra-Red LED Sensors

Flynn [40] describes an infra-red LED obstacle sensor for mobile robots. It operates by continuously emitting infra-red light, and sensing the brightness of the returned signal. The power of the radiation is governed by an r^{-2} spreading loss. Furthermore, the lustre of the reflecting surface




Figure 2.2: Resolution with sinc functions, which commonly occur when the [Discrete Fourier Transform \(DFT\)](#) is used for spectral analysis. Two different sinc functions (coloured), corresponding to two frequencies, are added together to obtain the observed spectral amplitude. The distance between the centres of the two sinc functions is marked. The two frequencies are resolvable when the sum exhibits a dip, which occurs when the maxima of one sinc function is superimposed over the minima of the other.

affects the amplitude of the reflected light. These concerns are held in common with both the sonar and radar problems. As the sensor cannot measure distance, it is unable to tell the difference between a close, poorly reflecting object, and a distant, highly reflective object. As the sensor can only detect light or dark, its operation is limited. Bearing precision can be improved by using optical elements to reject light returning from outside the region of interest. Flynn used parabolic reflectors to achieve this.

Ward [135] describes an alternative mode of operation where the light is not emitted from the robotic sensor, but from beacons in the environment. A ring of IR receptors can then coarsely report the bearing of the beacons (Figure 2.3).

The specifications, shown in Table 2.1, are adapted from Flynn [40].

2.3 Scanning Laser Range-Finders

Many laboratories have recently acquired scanning laser range-finders manufactured by SICK [121] (Figure 2.4). The 2D version returns an array of range measurements, precision ± 10 mm, at 0.25° increments in a plane. While this sensor provides accurate information, it is incomplete. In particular the scan is taken in a plane, so it fails to return information on anything which does not intersect that plane. For example, scanning at chest height does not reveal the presence of a table or a chair. Furthermore, scanning fences with vertical bars may fail because the beams from individual scanning positions may fail to illuminate the bars - thus there is a spatial sampling problem. It is difficult, at present, to extract features from the point cloud which is returned from the scanner. The 3D version addresses the plane sampling problem, but introduces a new problem akin to vision: how should all the range information be processed to recognise objects?

The propagating light wave is diffusely reflected by most objects in our environment. However, some objects reflect light specularly or allow it to pass straight through, such as windows, mirrors, and still water. The difficulty of detecting such objects using a light-based sensor is not limited to sensors for *robotic agents*. This author has known birds to fly into a window and knock themselves



Figure 2.3: Infra-red beacon sensor. Each small box contains an infra-red detector which is covered by a small slit. The sensors detect light from beacons placed in the environment. This sensor was used on a Yamabico mobile robot.

Table 2.1: Properties of infra-red LED sensors.

Range		
precision	none	The sensor only provides a thresholded light level output.
resolution	-	
min	-	
max	-	
Bearing		
precision	-	Coarse if no optical elements are used, but good with optical elements.
resolution	-	
beam width	-	
Benefits		
<i>Extremely low cost:</i> Requires common LEDs and a small amplifier and thresholding circuit, as well as a few digital computer inputs.		
<i>Beacons:</i> Can detect simple LED beacons placed in the environment (Ward [135]).		
<i>Wall Detection:</i> Effective at detecting the ends of walls when scanned with the parabolic reflector.		
<i>Speed:</i> Fast.		
Problems		
<i>Coarse measurements:</i> very coarse range and bearing measurments.		
<i>Ambiguity:</i> Ambiguity between range and reflectivity of objects.		
<i>Visibility:</i> Mirror-like surfaces in some orientations will reflect the light away from the sensor and evade detection. Similarly, clean glass surfaces will allow the light to pass through without reflecting.		



Figure 2.4: The scanning laser range-finder manufactured by SICK Corporation. This is the 2D version.

out. Presumably the birds are fooled by the window reflecting the bushland scene behind. Thus the deficiency is not attributed to a problem with the sensor, but to the properties of the propagating wave and its interaction with the target.

The properties of the SICK scanning laser range-finder are summarised in Table 2.2.

2.4 Radars

Radar technology has much in common with sonar technology, and historically many technological advances were first implemented in radar and then migrated to sonar. This section first gives a brief review of radar technology, and then turns to the application of radar as a mobile robot sensor.

2.4.1 Technology

Radar as a technology is a well researched and mature field. The MF techniques used within this thesis are standard radar techniques. Also well established are FM sweep chirps, which are related to the CTFM technique evaluated in this thesis. The difference lies in the signal length and the receive processor.

Radar systems usually allow the simultaneous measurement of target range, velocity, bearing and azimuth. The range and velocity information are separated by using specially crafted signals, receive processors and detectors. While these capabilities are interesting, the measurement of target velocity was not included in any of the systems evaluated here as it does not seem important

Table 2.2: Properties of the SICK scanning laser range-finder.

Range		
precision	$\pm 10\text{mm}$	Data sheet quotes $\pm 10\text{mm}$, however Bailey [7] states that quantisation errors of $\pm 30\text{mm}$ were observed.
resolution	-	Only one range is returned at each position.
min	-	
max	-	
Bearing		
precision	$0.25^\circ - 1^\circ$	The (quantised) bearing sampling grid is user selectable.
resolution	-	Twice the selected bearing quantisation.
beam width	$\pm 90^\circ$	
Benefits		
<i>Reliable:</i> Reliably senses most objects.		
<i>Correspondence problem:</i> The measurement is performed using one laser beam, so there is no need to associate returns from separate channels.		
<i>Speed:</i> Scans may be repeated tens of times per second. The bottleneck is the serial line out of the device, so finer and larger scans take longer to read.		
Problems		
<i>Interpretation:</i> It is difficult to extract features from the point cloud returned by the sensor.		
<i>Water and Glass:</i> Puddles or lakes can reflect the laser light away if the water is flat, introducing a blind spot. Clean glass is not reliably detected, but dirty glass is.		
<i>2D planar:</i> 2D version only senses objects in a plane. Objects just outside the plane escape notice.		
<i>Artifacts:</i> False readings along steep edges.		
<i>Dust:</i> Possible false readings from rain, steam and dust particles.		
<i>Sampling:</i> Spatial sampling problem means that picket fences and similar objects can remain undetected. The sensor beams are so tightly focussed that the positions of the beam do not combine to cover the sensing volume, in either the 2D or 3D version.		

in our application and would have dramatically increased the system complexity.

The techniques utilised in radar also have broad applicability to robust, high precision air sonar systems. The interested reader is referred to Berkowitz [11], Skolnik [122] and Rihaczec [118] for discussion of the techniques used and the benefits they provide.

2.4.2 Application to Mobile Robots

Microwave radar is useful for a long or medium range outdoor sensor. Clark and Durrant-Whyte [23] used a microwave radar to localise a container-moving vehicle for use in shipyards. Langer [87] built a [MilliMetre Wave \(MMW\)](#) radar as a sensor for a self-driving highway vehicle. This radar unit uses the synthetic aperture technique (Elachi [37]) in the receiver to measure the bearing of targets. Delphi [31] developed a commercial car radar system, which uses a mechanically scanned beam for detecting the bearing of targets.

Many radars use a [Frequency Modulated Continuous Wave \(FMCW\)](#) signal, which is similar to the [CTFM](#) signal used in this thesis.

The results summarised in Table 2.3 are reported by Langer [87] from an experiment conducted in an open area test site using a single target with a radar cross section of approximately 7 m^2 .

2.5 Ultrasonic Sensors

Ultrasonic sensors, in general, operate by insonifying (defined to be similar to illuminating, but with sound, Kay [74], also ensonify in Kinsler [78]) an area, and measuring the time taken for an echo to return. The signal may be of many different types, including impulses and continuous signals. Ultrasound has a wavelength between 17 mm and 3.5 mm for the frequencies of 20 kHz to 100 kHz commonly used in air.

Ultrasonic sensing systems for use in air may be classified according to how many elements are used in their construction and the type of encoding used for the signal. There are several main configurations of ultrasonic sensors. The main configurations, in increasing order of complexity, are:

1. A sensor with one transmitter and one receiver, possibly combined into single transducer sensor, and possibly mechanically scanned.
2. Ring arrays with transducers used separately or fired in unison. This configuration eliminates scanning.
3. A sensor with a small number of transducers pointing in the same general direction and used in unison.
4. A transducer array, using either transmit or receive beam-forming (synthetic aperture), possibly with electronic scanning control.

The properties of the ultrasonic transducer used in the system are often crucial to the results. Therefore the type of ultrasonic transducer is recorded where the information is known.

The following sections discuss several different ultrasonic sensing systems.

Table 2.3: Properties of MMW radar as reported by Langer [87].

Range		
precision	± 0.12 m	Theoretical precision ± 0.112 m.
resolution	0.93 m	
min	-	
max	180 m	Range depends upon the radar cross section of the target. Quoted results were 80m for people, 180m for cars, and 200m for trucks.
Bearing		
precision	0.07°	Only applies to targets which cannot be resolved in range.
resolution	3°	
beam width	$\pm 6^\circ$	Would have been wider but suffered from budget constraints.
Benefits		
<p><i>Reliable:</i> Returns data for most target types which are placed within the sensor beam.</p> <p><i>Beacons:</i> Road markings can be replaced with radar reflective tapes, and beacons can be placed into the environment, similar to the current optical retro reflectors, to aid reliable sensing for navigation.</p> <p><i>Harsh Environments:</i> Works in the dark and is immune to wind, rain, dust and fog, making it well suited for use outdoors.</p> <p><i>No moving parts:</i> Uses electronic scanning to measure bearing, so there are no moving parts to break down.</p> <p><i>Volume sensor:</i> The radiation fills the volume, and there is no spatial sampling as with laser range-finders.</p> <p><i>Speed:</i> Provides readings 10 times per second according to stated requirements.</p>		
Problems		
<p><i>Ghost targets:</i> Poor sidelobe suppression in bearing, which causes ghost targets from outside the main beam.</p> <p><i>Narrow vertical beam:</i> The road can rise or dip outside of the 3° vertical extent of the beam used by this sensor, but increasing the beam width leads to sensing undesired targets such as overhead bridges. No facility is provided for resolving objects vertically.</p> <p><i>Weak targets suppressed:</i> If two targets have the same range, and have different radar cross sections (i.e. a motor cycle next to a truck), then the weaker target will be lost in the sidelobes of the stronger target, and will not be observed by the sensor. This effect is amplified by the very wide bearing sidelobes of this sensor.</p>		



Figure 2.5: The Polaroid transducer. The black grille covers a mylar film with a gold conductor deposited on one side. The radius of the active area of the transducer is 16.2 mm. See Section 6.8 for details.

2.5.1 Polaroid Type

The Polaroid Corporation [12] developed an ultrasonic range-finder for auto-focusing their cameras. The range-finder uses a single electro-static transducer (Figure 2.5) to transmit and receive a sine wave packet of about 10 cycles (100 mm). The transducer requires a bias voltage to be applied before it can be used as a transmitter or receiver.

When the sonar is fired, a timer is started. This timer is stopped when a received echo exceeds a pre-set threshold and the range, r , is computed from the time of flight, τ , and the speed of sound, c , by $r = c\tau/2$. The object which caused the echo lies upon a spherical surface (Figure 2.6) whose extent is defined by the sonar beam width.

Polaroid supply a standard kit containing the signal conditioning circuits on a small board. The circuit includes a time dependent gain, empirically calibrated to compensate for spreading loss and air absorption. The circuit detects the received echo by comparison with a fixed threshold level. The timer circuit is usually supplied by the user. The properties of the sensor are summarised in Table 2.4.

2.5.2 HiSonic Sensor

The HiSonic sensor built for the Yamabico mobile robot by Ohno [103] is similar to the Polaroid sensor. It uses a Murata [98] piezo-electric transmitter-receiver pair, part numbers MA40A5S and MA40A5R. These transducers have only a small active area, but do not require a bias voltage. Besides using different transducers from the Polaroid system, the circuit also uses an exponentially decaying threshold. Changing the detection threshold from a constant to a modelled variable has the joint outcomes of reducing the effect of crosstalk from the transmitter to the receiver, and increasing the sensitivity to more distant targets. The sensor could detect small reflectors at a maximum range of 1.4m and larger targets, 0.1m wide, at a maximum range of 3.5m. The experience of this author with this sonar system indicates that careful adjustment of the decay threshold and offset time are required to achieve this range.

Insufficient data is provided on the HiSonic sensor to form a table of properties.

Table 2.4: Properties of the Polaroid ultrasonic range-finder kit.

Range		
precision	$\pm 1\text{ mm}$	Only measures the range to the closest target. The receiver circuit must be turned off while the transmitter is active. Typical
resolution	-	
min	100mm	
max	5 m	
Bearing		
precision	$\pm 15^\circ$	The object position must be estimated by combining measurements from different known positions.
resolution	-	
beam width	$\pm 15^\circ$	
Benefits		
<i>Reliable:</i> Reliably senses most objects.		
<i>Speed:</i> Measurements may be repeated at 30Hz for a maximum range of 5m.		
<i>Low cost:</i> Circuit is small and simple to build.		
Problems		
<i>No angle estimate:</i> The data from the sonar sensor is usually interpreted to have come from the same horizontal plane as the transducer, so echoes coming from small steps in the floor or door jambs cause confusion to the mapping system. The lack of bearing measurement requires additional complexity in the mapping system, as objects must be observed from multiple positions to determine their location.		
<i>Corner Detection:</i> According to Kleeman [82], measurements taken by a single transducer sonar sensor cannot tell the difference between a wall and a corner.		
<i>Objects obscured:</i> As the sensor can only measure the range to the closest object, more distant objects are effectively obscured.		

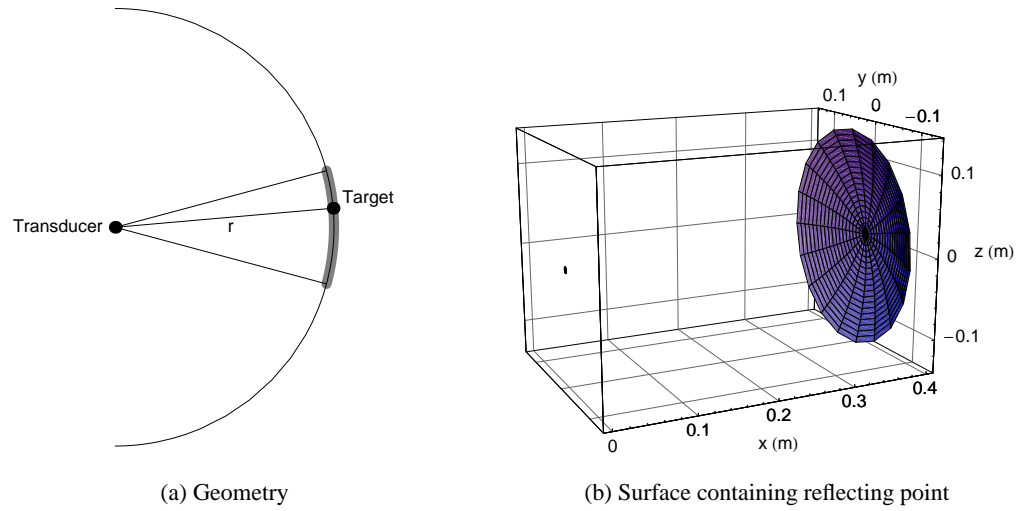


Figure 2.6: The geometry of a single transducer sonar system. (a) Shows a target, and the range which will be measured by the sensor. As the single transducer provides no information about the target bearing, the target may lie anywhere on a circular annulus (black). However, the sonar energy is constrained in a beam, so the object is most likely to be in the region marked in grey. (b) Shows the possible position of the object in 3D, with the transducer visible at $\{x, y, z\} = \{0\text{m}, 0\text{m}, 0\text{m}\}$.

2.5.3 Plant and Target Recognition with CTFM Sonar

Harper [56] demonstrated that it is possible to recognise plant species using a CTFM sonar (Figure 2.7). That work was intended to be used for landmark recognition for long range navigation. Other work was done by Politis [111, 112], which recognised corners, planes and edges from their echo signatures. Both of these systems use a single transmitter and single receiver sensor configuration, and CTFM electronics from Leslie Kay [73].

2.5.4 Measuring Bearing with a Single Receiver

Yata constructed and evaluated a direction-finding sonar system using a transmitter and a single receiver [142, 143, 144]. It operated by transmitting a broad-band impulse response signal, and exploited the frequency dependent beam pattern (Sections 6.3.2 and 6.3.3) to estimate the bearing of the target. The estimation of frequency is performed using the zero-crossing period of the second oscillation in the received echo. This method simplifies the signal processing considerably compared to traditional methods. The zero-crossing period was used to find the target bearing in a look-up table. The properties of this very simple sonar system are summarised in Table 2.5.

This author is saddened to learn of Yata's untimely death.



Figure 2.7: The head, driver electronics, and headphones comprising the research **CTFM** probe system used by Harper for recognising plants. This unit was supplied by Leslie Kay. The central element on the head is a transmitter, flanked by two angle-adjustable receivers. While this unit is binaural, and capable of measuring bearing, only one receiver channel was necessary for plant recognition work.

Table 2.5: Properties of the single receiver direction measuring system built by Yata.

Range		
precision	-	Not stated.
resolution	-	Not stated.
min	-	Not stated.
max	-	Not stated.
Bearing		
precision	$\pm 0.95^\circ$	For target at 0.8m. Depends upon target range. Bearing precision is $\pm 0.66^\circ$ for target at 0.5m.
resolution	-	Not stated.
beam width	8°	Results only presented for $\pm 4^\circ$.
Benefits		
<i>Simple:</i> Uses a simple hardware design.		
<i>Multiple targets:</i> Can detect multiple targets in one observation.		
Problems		
<i>Overlap:</i> Can not resolve targets with overlapping or interfering echoes.		
<i>Noise:</i> Not robust against noise.		
<i>Direction ambiguity:</i> The sensor can only measure the angle of the target from the normal of the sensor. If the target is assumed to lie in the plane, then a sign ambiguity remains. Yata overcomes this by using a bias angle, but this reduces the operating beam width of the sensor.		

Table 2.6: Properties of the ring sensor system built by Yata.

Range			
precision	-		Not stated.
resolution	170mm		
min	-		The increased range over the HiSonic sensor is attributed to the use of a horn.
max	5 m		
Bearing			
precision	$\pm 0.8^\circ$		Sensor system operates in a complete circle around the robot.
resolution	-		
beam width	360°		
Benefits			
<i>Simple:</i> Uses a simple hardware design.			
<i>Speed:</i> Update rate of 10Hz.			
<i>Multiple targets:</i> Can detect multiple targets in one observation.			
Problems			
<i>Complexity:</i> The main difficulty is the construction, calibration and maintenance of 30 sets of sonar circuitry.			

2.5.5 Direction Finding Sonar Ring

Yata also built a 30 element sonar ring which can measure the bearing to several targets simultaneously [142, 145]. Yata utilised piezo-electric transducers from Murata, model numbers MA40S4T and MA40S4R. The ring is comprised of 30 transmitter and receiver pairs. All of the transmitters are fired simultaneously, providing a 360° circular wavefront. This eliminates the need for sequential transmission found on other ring sensors. The echoes on all of the receivers are analysed together to detect and estimate the bearing of targets by [time of flight](#). Her work had its roots in the HiSonic system mentioned in Section 2.5.2, and utilises an exponentially decaying threshold level in its receiver. Unlike the HiSonic system, the receiver records samples of the output of the thresholding comparator, producing a 1-bit representation of the signal. The properties of Yata's ring sensor are reviewed in Table 2.6.

2.5.6 Impulses and Template Matching

Kleeman and Kuc [82] have developed several impulse ultrasonic sensors capable of measuring the range and bearing tuple of multiple targets. This sonar type has been further refined by Chong [21] and Heale [59]. The system utilises the Polaroid 7000 series electrostatic transducer. Transmission is performed by applying a 300 V, $10\mu\text{s}$ impulse to the transmitter. Thus the sonar transmits the impulse response, having a packet duration of $50\mu\text{s}$, equivalent to a length of 17.3 mm. A template matching system is used in the receiver. There are several configurations used, depending upon whether two or three dimensional sensing is required. The two dimensional configuration uses two transmitters and two receivers, while the three dimensional configuration uses three transmitters and three receivers. As the signal is quite short, one transducer in each configuration is used as

both a transmitter and a receiver. The receiver uses a 1 MHz sampling rate.

The sonar system measures the target bearing by triangulation. The templates are used to compensate for the effects of the transducer's response at different bearings. Additionally, this system can reliably classify targets as being a corner, plane or edge in a two transmission observation (2D version). The usual poor [Signal to Noise Ratio \(SNR\)](#) experienced with impulses is overcome by boosting the transmitter power well above the normal limits, and using large transducers.

Kleeman and Kuc [82, Section 9] provide experimental results for absolute accuracy. Their experiment consists of measuring the acoustic [time of flight](#) and geometrical [distance of flight](#). A line is then fitted to an x-y plot. The slope of this line is claimed to be $c/2$. This procedure will remove any bias in the fitting process. As the experiment was specifically designed to detect any bias in the sonar system, this is considered to be a serious weakness. Therefore their accuracy figure of ± 0.8 mm is brought into question.

The precision of ± 0.1 mm to ± 0.2 mm given for this sonar system appears to have been measured in a room with still air. A separate figure of ± 0.4 mm is given for turbid air driven by fans.

The properties of this sonar system are summarised in Table 2.7.

Hong [62] reports the development of a cheaper, simpler version of this sensor with a range accuracy of 10 mm and bearing precision of $\pm 2^\circ$. Hong found that the precision, identification range and resolution of closely spaced targets were compromised due to the very low 59 kHz sampling rate.

2.5.7 Barker Codes and Matched Filters

Barker codes were devised as a way to increase the energy of the transmitted signal, so as to improve the signal to noise ratio in the receiver. They originally came from radar. The extra energy is inserted by increasing the length of the transmitted signal. Usually, this would result in poor range resolution. However, the matched filter in the receiver compresses the signal into an impulse at the point where the echo starts. This allows echo signals which overlap to be resolved as separate targets.

This idea was applied to robotic sensing by Peremans [108, 109]. He used a single transmitter and three receivers, all using Polaroid transducers. The sensor is able to detect multiple targets, and measure their range and bearing tuple. Since multiple receivers are used, the echoes from each channel which correspond to a single target must be identified before the direction angle to the target can be estimated. Additionally, this sensor attempts to estimate whether the target is a plane, corner or edge reflector. The properties of this sensor are described in Table 2.8.

2.5.8 3D CTFM Ultrasonic Sensor

Davies [28] built and evaluated several 3D [CTFM](#) ultrasonic sensor systems using two different types of transducer. The first type of transducer was reported as the Warwick transducers, but no performance information is available. The second type of transducer was manufactured by Kay. Custom circuitry was used for transmit, receive, demodulation and sampling. A dedicated [Digital Signal Processing \(DSP\)](#) chip was used to perform the [Fast Fourier Transform \(FFT\)](#) analysis

Table 2.7: Properties of the Kleeman and Kuc impulse ultrasonic sensor system.

Range		
precision	$\pm 0.1\text{ mm}$ to $\pm 0.2\text{ mm}$	Precision is range dependent. Figures given up to 5m range for still air. Accuracy of $\pm 0.8\text{ mm}$ is claimed for ranges less than 4m. See notes in text. Not stated, but will have a minimum range due to turning off the receive circuit during transmission. For optimally aligned plane target.
resolution	10mm	
min	-	
max	8m	
Bearing		
precision	$\pm 0.1^\circ$	Accuracy claimed to be 0.2° over $\pm 10^\circ$. Targets separated in range can be resolved in bearing, even if the bearing is the same (unless obscured).
resolution	-	
beam width	$\pm 11^\circ$	
Benefits		
<i>Reliable:</i> Reliably senses most objects.		
<i>Speed:</i> Provides measurements in real time.		
<i>Multiple targets:</i> Can detect multiple targets in one observation.		
<i>Classification:</i> classifies planes, corners and edges in a double-transmission observation.		
Problems		
<i>Noise:</i> The energy level in the impulse signal is quite low, so the system will be vulnerable to disruption by noise. Leslie Kay [72] suggests that this can be quite significant in industrial environments.		
<i>Parameters:</i> The template matching space is multi-dimensional, including bearing, range, and temperature as factors.		
<i>Calibration:</i> The template matching is dependent upon the impulse response of the individual transducers used.		

Table 2.8: Properties of the [matched filter](#) sonar system built by Peremans.

Range		
precision	$\pm 1\text{ mm}$	Precision degradation of up to $\pm 10\text{ mm}$ for overlapping echoes. The sonar packet is 415 mm long.
resolution	-	
min	-	Not stated, but results given up to 4 m .
max	4 m	
Bearing		
precision	$\pm 0.6^\circ$	Resolution of targets in bearing is impossible when echoes are not resolvable in range.
resolution	-	
beam width	$\pm 12^\circ$	
Benefits		
<i>Reliable:</i> reliably detects most targets.		
<i>Multiple targets:</i> Can detect multiple targets in one observation.		
<i>Classification:</i> classifies planes, corners and edges in a double-transmission observation.		
Problems		
<i>Classification:</i> The classification of targets from a single observational location proved to be unreliable. Sensor motion was found to overcome this (Peremans [110]).		

of the demodulated signal. Davies describes some interpolation techniques which were used to improve the precision of the sensor. The software used is derived from the BAT system of Leslie Kay.

Davies’ systems measure range, bearing and elevation angle using [Inter-aural Distance Difference \(IDD\)](#) and [Inter-aural Power Difference \(IPD\)](#). However, the evaluation does not provide range or bearing precision. Instead, the evaluation is carried out in Cartesian coordinates in terms of Euclidean distances. Precisions of between 7.1 mm and 42 mm are reported, depending upon the bearing measurement technique. The errors appear to be largely in the bearing estimate. The [IDD](#) technique provides smaller errors than the [IPD](#) technique in Davies’ results.

Davies also describes degradation in system performance when a fan is used to add wind noise, and describes the adverse heat effect of a soldering iron placed underneath the sonar path. Some textured surface work is also included.

No table of results is included for Davies’ work due to the form in which the results are presented.

2.5.9 CTFM Sonars for *Robotic Agents* by Leslie Kay

Although Leslie Kay is best known for his pioneering work in the use of [CTFM](#) air sonar systems for use by *living agents*, he has also adapted his sonar technology for use by *robotic agents*. His early work [72] utilised a synthetic aperture technique to focus the transmitted beam. The beam was then scanned from left to right. Echoes detected at each scanning position could be uniquely attributed to a reflector at that bearing. Each 30° scan was completed in 15 ms. The properties of such a system are summarised in Table 2.9.

Table 2.9: Properties of the scanning array **CTFM** sonar systems built for robotic use by Kay.

Range		
precision	-	At 100kHz.
resolution	6mm	
min	-	
max	-	
Bearing		
precision	±0.4°	Determined by beam width, which is four wavelengths or 14 mm. Bearing precision calculated at 1 m range. Calculated at 100kHz.
resolution	0.8°	Width of focal point is four wavelengths or 14 mm. Bearing resolution calculated at 1 m range.
beam width	±15°	Area over which the beam is scanned.
Benefits		
<i>Speed:</i> Completes one scan in 15 ms. The sonar demodulation is implemented in analogue circuits, and the demodulated output requires only FFT processing.		
<i>Image:</i> Provides a range image of the insonified area.		
<i>Object classification:</i> provides information which has been used by Politis [112] to recognise reflector types and by Harper [55] to recognise plant species.		
<i>SNR:</i> CTFM provides a high signal to noise ratio as it uses a continuously transmitted signal.		
<i>Continuous Contact:</i> The CTFM sonar system can maintain continuous contact with the target, as it transmits continuously.		
Problems		
<i>Construction:</i> The construction is complicated by the number of elements in the beam-forming array. The large number of circuits driving the array elements must be carefully arranged to reduce crosstalk.		

More recent work [73] demonstrates a BAT computer program which processes the demodulated **CTFM** echoes from a binaural sonar system such as that shown in Figure 2.7, and computes the range and bearing of targets. The receivers are splayed to support the **IPD** method of calculating bearing. In this system a wide beam is transmitted, and the binaural echo information must be used to compute the bearing of targets. No performance evaluation is described for this system.

2.5.10 CTFM Arrays

Ratner [117] used a **CTFM** sonar array built by Leslie Kay (Figure 2.8) to recognise and localise plants from a mobile robot. This extends the work of Harper described in Section 2.5.3. The array is used as a synthetic aperture array transmitter with a 3.5° beam width. The vertical dimension is unconstrained, giving a fan shaped beam which is electronically scanned. The sensor was used to perform long range navigation by detecting and following the grass-concrete interface at the edge of a path.




Figure 2.8: A CTFM sonar head with a transmit beam-forming array (top), which is electronically steerable. The four transducers mounted below are the receivers. They are angled to support bearing computation from the IPD.

2.6 Combined Visual and Sonar Sensor

Akbarally [3] describes an unusual combination of a monocular visual sensor and a sonar sensor. The sonar sensor is similar to that previously described in Section 2.5.6. In Akbarally's work, the weaknesses of each sensor are complemented by the strengths of the other.

Most of Akbarally's thesis describes the operation of the sensor. Only one experiment was conducted to test the accuracy [3, Section 5.10.3]. The sensor produces a wire model of the observed scene. The accuracy experiment compares the lengths of these line segments with lengths measured by a ruler. Thus the range and bearing precision have not been independently assessed.

The properties of Akbarally's combined visual and sonar sensor are summarised in Table 2.10.

2.7 Properties of Ultrasound

The physical properties of ultrasound determine the benefits and limitations common to many ultrasonic sensor systems. These properties affect all of the ultrasonic sensors in Section 2.5.

Non-Normal Surfaces are Invisible

Ultrasound is reflected specularly from smooth surfaces. If a smooth plane surface reflects the ultrasound away from the receiver, then it will not be detected (Figure 2.9). Note that the corner produces a small echo [85]. The reflection from the edge is subjected to two spreading losses - first the spreading loss due to the ultrasound radiating from the transducer, and then the spreading

Table 2.10: Properties of the combined sonar and visual sensor system built by Akbarally. These results are slightly worse than those mentioned for the sonar sensor in Table 2.7, but the results quoted in that table are for more recent work.

Range		
precision	$\pm 0.15\text{ mm}$	Accuracy claimed to be $\pm 0.8\text{ mm}$.
resolution	-	
min	-	
max	6 m	
Bearing		
precision	$\pm 0.05^\circ$	Bearing precision at 1 m range.
resolution	-	
beam width	-	
Benefits		
<i>Complementary sensors:</i> The range measuring deficiency of the monocular vision is complemented by the range measuring ability of the sonar.		
<i>Detects multi-path:</i> Some sonar paths strike multiple objects, leading to the detection of phantom targets by the sonar sensing system. The vision processing system can be used to check that a target is present in the image at the coordinates reported by the sonar system, thereby eliminating multi-path reflections.		
<i>Atmospheric compensation:</i> The sonar system uses impulses and template matching, but the pulse shape changes with atmospheric conditions. Akbarally described a method of compensating the pulse shape for the atmospheric variation.		
Problems		
<i>Slow processing:</i> The sensor takes 90s to process a scene, with most of the time being consumed in processing the visual sensor data.		




Figure 2.9: An object with plane reflectors which directs most of the sonar energy away from the transducer. The transmitter produces radiating spherical waves (red) which strike the target. The plane reflections (black) are reflected away from the receiver. The edge diffraction (grey) radiates in all directions, and is the only signal to reach the receiver. The detected echo power, for plane and corner reflectors, is governed by an r^{-2} spreading loss and air absorption. The detected echo power from edge diffraction is governed by an r^{-4} spreading loss in addition to the air absorption. For many sonar systems, this will put the echo below the noise floor and the detection threshold, making such an object invisible to the sonar sensor except at close range. This topic is discussed further in Chapter 4. For clarity, diffraction from the other edges is not shown.

loss due to the ultrasound re-radiating from the corner. This additional spreading loss makes the reflections much weaker than those from corners and suitably oriented planes, and thus may not be detected by the receiver.

Ultrasonic Echo Features are Localisation Beacons

The physical features which return echoes are usually room features. Thus the echoes produce a higher ratio of relevant information to clutter than other sensors returning information on a dense grid such as vision or 3D laser range-finders. The information present in the ultrasonic signal may be extracted precisely without the use of fine sampling grids.

Visibility of 3D Corner Reflectors

Three dimensional corner reflectors have been shown by Hong [62] to be a useful room feature for robot localisation, because they have the same echo power as planes but without the problem of specular reflection. Echoes from corner reflectors may be seen over a wide range of positions relative to the orientation of the corner, allowing them to be used as beacons for robot localisation. An environment with few ultrasonic reflectors may have artificial 3D corner reflectors installed as

beacons to allow robotic navigation.

The use of 3D corners in radar and optics is well known. Their use in sonar, to provide similar control and calibration, should be encouraged.

Volume Sensor

The area of insonification is broad, and echoes will be returned from objects in a wide field of view. This allows more area to be sensed in a single reading than with narrowly defined sensors such as a 2D laser range-finder or a phased array beam-forming ultrasonic sensors.

In previous work, this last point has been considered a disadvantage. However, modern techniques allow the bearing of individual targets to be measured directly in a single or double transmission, eliminating the need to use a narrow beam to reduce the position uncertainty of sensed objects. Indeed, having a larger area of insonification allows more information to be extracted in a single measurement.

2.8 Conclusion

This review shows that the following desirable properties should be incorporated into the sensor:

- The sensor fills or insonifies the volume to be interrogated rather than digitally scanning the volume.
- The sensor should be self-insonifying, like radar, and not dependent, like most optical systems, on an external source.
- The sensor should be able to probe a significant object space with accuracy, precision and reliability.
- The scan data can be quickly and reliably processed.

Of the sensors reviewed only the impulse sonar, Barker codes with a matched filter, and the [CTFM](#) ultrasonic sensors are able to fulfil the above criteria. The impulse sonar system has been well documented elsewhere [82, 21, 59]. The matched filter and [CTFM](#) systems appear to be able to transmit more power and hence may be able to probe a large volume. Investigating these types of sonar system appears important. Additionally, the [CTFM](#) systems have been shown to support target and plant identification, making them suitable for localisation tasks.

Precisions of between ± 0.1 mm and ± 1 mm in range are reported. This sets a precision goal for the systems in this work. Furthermore, resolution between objects separated by 6 mm to 10 mm is reported for state of the art systems, setting the resolution goal for this work.

The bearing precisions and resolutions are much more poorly understood than the range component and thus a major contribution can be made by providing a rigorous study of range and bearing process and their interaction.

It is desired that the sonar system operate robustly in both benign and hostile environments. For an ultrasonic sensor system, hostile elements are ultrasonic noise, wind, and uneven temperature. Ultrasonic noise may be overcome by increasing the amount of power in the transmitted

signal, such as is used by the [matched filter](#) sonar system and in the [CTFM](#) systems discussed previously. The hostile elements of wind and uneven temperature affect the propagation of the ultrasound itself, so there is little that can be done to improve system performance where these are concerned. However, these problems are reduced considerably by operating in indoor environments. In addition to the requirements of operating in a hostile environment, the system chosen should be able to support the recognition of landmarks and other reflector types. Other work described in this chapter indicates that [CTFM](#) sonar systems are able to fulfil these requirements. With all these considerations, the [matched filter](#) and [CTFM](#) sonar systems were selected for further study and comparison in this thesis.

While cheap and simple sensors may be preferred by robotics researchers for application, the comparison between the capabilities of the various sensors shown here demonstrates that there is a cross-correlation between sensor complexity and the quality and quantity of useful information returned. In selecting sonar systems for study in this thesis we have chosen to emphasise capability over cost and complexity, so as to perform fundamental research into the capabilities of the ultrasonic sensor.

Chapter 3

Operation of CTFM

This chapter draws together the fundamental theory of [Continuous Tone Frequency Modulated \(CTFM\)](#) from the literature into a coherent mathematical framework to build an understanding of [Continuous Tone Frequency Modulated \(CTFM\)](#) sonars. This information is used to support design decisions which will be made in Chapters 6 and 7, and also provides the theory necessary to understand the digital [CTFM](#) sonar implementation described in Chapter 9.

The chapter begins with an overview of the [CTFM](#) sonar system. It then reviews the transmit signal and the mathematics used to generate it. Finally, it reviews two techniques for demodulating the received echoes: single and dual demodulation. Thus the chapter assembles the theory of [CTFM](#) into a coherent mathematical framework.

3.1 Overview

A single channel [CTFM](#) sonar system is shown in Figure 3.1. The sweep generator provides the ultrasonic [CTFM](#) signal for use by both the transmitter and demodulator. The signal is transmitted, reflected by objects in the environment, and returns to the receiver. The demodulator produces an output frequency in the audible range, with frequency being proportional to the distance of flight. Each target produces a unique frequency. Thus the range measurement problem becomes a spectral analysis problem when a [CTFM](#) sonar is used.

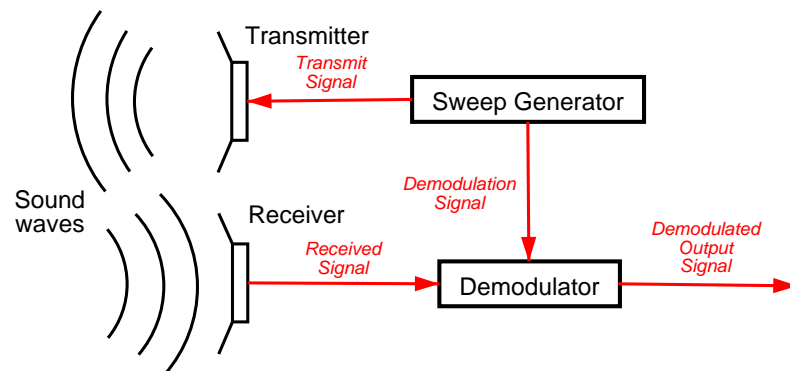




Figure 3.1: Overview of a simple [CTFM](#) sonar system.



(a) Time-frequency plot



(b) Time-amplitude plot

Figure 3.2: Example of a frequency sweep as (a) time-frequency plot, and (b) time-amplitude plot. The time-frequency plot shows the change in frequency of the signal over time. False sweep times and frequencies have been used for clarity in these plots.

The following two sections describe the operation of the sweep generator and the demodulator in detail. In particular, two different methods of demodulation are discussed. Discussion of spectral analysis is postponed until Chapters 10 and 11.

3.2 Transmit Signal

The sweep generator in Figure 3.1 produces a transmit signal and a demodulation signal. This section focuses on the transmit signal.

The transmit signal is a CTFM signal which is transmitted continuously while the sonar is operating. The frequency is modulated according to a sawtooth function (see Figure 3.2). The transmit sweep frequency $f_t(t)$ is determined by the following three equations:

$$f_t(t) = f_0 + \mu t_p \quad (3.1)$$




Figure 3.3: The transmit signal and received echo for a single target, showing the important features of the signal. The echo delay, or time of flight, is labelled as τ . The signals used in this thesis nominally use $f_0 = 100\text{kHz}$ and $f_1 = 50\text{kHz}$ with a sweep period of $T_{sw} = 0.16\text{s}$, as discussed in Chapter 9.

$$\mu = \frac{f_1 - f_0}{T_{sw}} \quad (3.2)$$

$$t_p = \text{remainder} \left(\frac{t}{T_{sw}} \right), \quad (3.3)$$

where the subscript t has been used to denote transmit signals; f denotes frequency, and the subscripts 0 and 1 denote the start and end of the sweep. Thus, the sweep start and stop frequencies are denoted by f_0 and f_1 respectively. The sweep time is denoted by T_{sw} , while μ is the sweep rate, in Hertz per second. The variable t_p denotes the periodic time in seconds, and is constrained by $0 \leq t_p < T_{sw}$ due to its definition. This is what causes the sawtooth to repeat with period T_{sw} . The transmit signal frequency $f_t(t)$ has units of Hertz and is shown in Figure 3.3. Note that μ is negative, indicating a downward sweep, which is the standard configuration for CTFM systems (Kay [71]).

The transmit signal $s_t(t)$ is written

$$s_t(t) = \sin(\phi_t(t)), \quad (3.4)$$

where $\phi(t)$ denotes a phase, with the subscript t denoting the transmit signal. The relationship between the phase $\phi(t)$ and frequency $f(t)$ is developed in the remainder of this section.

For a signal with constant frequency, the phase $\phi(t)$ is known to be

$$\phi(t) = 2\pi ft. \quad (3.5)$$

This does not hold for a variable frequency such as $f_t(t)$. Oppenheim et. al. [104, Section 7.6] describe the principle of instantaneous frequency for such cases. They suggest that phase $\phi(t)$ should be regarded as the integral of instantaneous frequency $f(t)$, i.e.

$$\phi(t) = 2\pi \int_0^t f(t) dt. \quad (3.6)$$

In the case of constant frequency, such as $f(t) = f$, this reduces to (3.5). In the case of the sweeping transmit frequency $f_t(t)$, substituting (3.1) into (3.6) yields

$$\phi_t(t) = 2\pi \int_0^t (f_0 + \mu t_p) dt.$$

Applying a change of variables $t = t_p$, over the range $0 \leq t < T_{sw}$, and integrating yields

$$\phi_t(t) = 2\pi f_0 t_p + \pi \mu t_p^2. \quad (3.7)$$

Allowing t to increase over T_{sw} would give the swept frequency a different phase offset at the beginning of every sweep. The omission of the phase offset is of no concern as the sweep is stored in a digital memory and simply repeated exactly for every sweep, and is computed so that there is perfect matching between sweeps (Section 9.3). Substituting (3.7) back into (3.4) gives the final transmit signal

$$s_t(t) = \sin(2\pi f_0 t_p + \pi \mu t_p^2), \quad (3.8)$$

with μ and t_p being calculated by (3.2) and (3.3) respectively. An example plot of this signal is shown in Figure 3.2b.

Formula (3.8) is used wherever a frequency sweep is required, substituting appropriate values for f_0 and f_1 . The sweeps usually occur over a bandwidth of one octave, such that the higher frequency is twice the lower one. All sweeps in the CTFM system use the same value for T_{sw} . Interestingly, (3.8) may be regarded as a constant frequency signal with a time varying phase modulation. Thus phase modulation and frequency modulation are equivalent in this context.

3.3 Demodulation

The demodulation stage of the CTFM sonar system transforms the ultrasonic frequencies of the sonar signal by lowering the frequency. In Kay's aid for the visually impaired [71], the demodulated output signal is in the audible frequency range. When the demodulation output is to be interpreted by the sensory system of a *robotic agent*, the demodulation output may cover any convenient frequency range, unless a *living agent* must listen to the signal during testing.

3.3.1 Single Demodulation

In single demodulation the demodulation sweep, which is identical to the transmitted signal, is multiplied with the received echo signal in the time domain. This produces sum and difference frequencies, which may be clearly seen in Figure 3.4. The distance of flight information is contained in the difference frequency. If the demodulation sweep and the echo sweep have the same

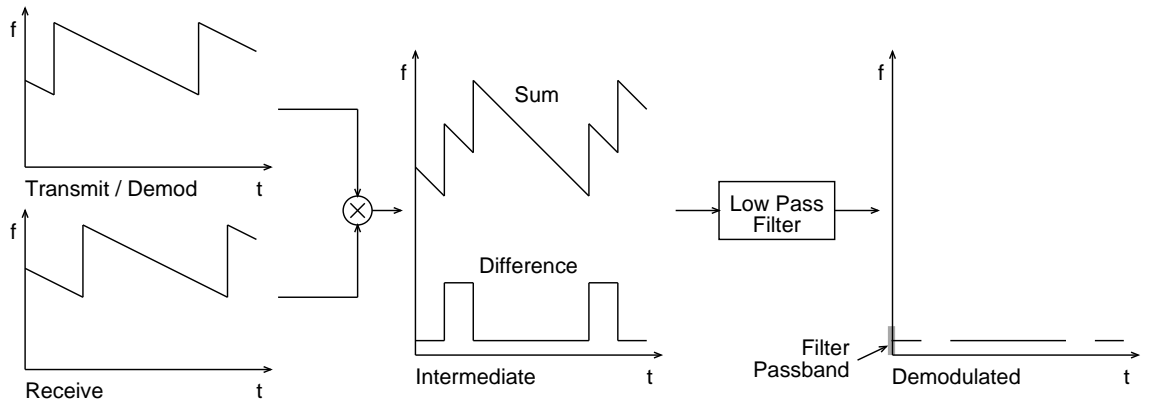


Figure 3.4: The operation of a single demodulation CTFM sonar.

slope, μ , constant for the duration of the sweep T_{sw} , then the difference frequency will be constant. For this to happen, the transmit signal must have a perfectly linear frequency sweep, and the reflecting object must be stationary relative to the sonar system.

The time domain multiplication also produces a sum frequency, but this signal is of no use so it is filtered out using a low pass filter. The demodulated intermediate signal contains steps, which correspond to the travel time from the transmitter to the target and back. The low pass filter mentioned previously also removes the top of these steps, leaving a signal with periodic gaps in its spectrum, as shown at the right of Figure 3.4.

It is important to obtain a qualitative understanding of the relationship between the **Distance Of Flight (DOF)** and the output signal. Consider first an echo with **DOF** $d = 0$. Such a signal is identical to the transmit signal, and thus also the demodulation signal. The difference frequency is thus zero. More distant targets may be considered by sliding the receive sawtooth to the right. This increases the difference between the frequencies of the received and demodulation signals, so the demodulated difference frequency increases for more distant targets.

DOF Calculation

This section explains how to compute the **DOF** from the demodulated output frequency. First, let d denote the distance of flight and let c denote the speed of sound. Then the time of flight τ may be calculated by

$$\tau = \frac{d}{c}. \quad (3.9)$$

The received signal $s_r(t)$ is demodulated by multiplying it with the demodulation signal $s_d(t)$ thereby giving the demodulator output signal $s_{out}(t)$:

$$s_{out}(t) = s_d(t) \cdot s_r(t). \quad (3.10)$$

In single demodulation, the demodulation signal $s_d(t)$ is identical to the transmit signal $s_t(t)$. The received signal is regarded simply as a time-delayed version of the transmit signal for this simple analysis, so

$$s_r(t) = s_t(t - \tau). \quad (3.11)$$

These two modifications bring us to

$$s_{out}(t) = s_t(t) \cdot s_t\left(t - \frac{d}{c}\right). \quad (3.12)$$

Substituting (3.4) into (3.10) and applying a trigonometric identity yields

$$s_{out}(t) = \sin \phi_t(t) \cdot \sin \phi_t\left(t - \frac{d}{c}\right) \quad (3.13)$$

$$= \frac{1}{2} \left[\cos \left\{ \phi_t(t) - \phi_t\left(t - \frac{d}{c}\right) \right\} - \cos \left\{ \phi_t(t) + \phi_t\left(t - \frac{d}{c}\right) \right\} \right]. \quad (3.14)$$

The frequency sum term on the right is of no interest, and is eliminated by the low pass filter. Rewriting (3.14) to reflect this yields

$$s_{out}(t) = \frac{1}{2} \cos \left\{ \phi_t(t) - \phi_t\left(t - \frac{d}{c}\right) \right\} \quad (3.15)$$

It has been established that the frequency of $s_o(t)$ is related to the **DOF** d . The frequency of $s_o(t)$ may be extracted by using the definition of instantaneous frequency (3.6) in the reverse form:

$$f(t) = \frac{1}{2\pi} \frac{\partial}{\partial t} \phi(t). \quad (3.16)$$

Applying this to $s_{out}(t)$ gives:

$$f_{out}(t) = \frac{1}{2\pi} \frac{\partial}{\partial t} \left(\phi_t(t) - \phi_t\left(t - \frac{d}{c}\right) \right) \quad (3.17)$$

$$= f_t(t) - f_t\left(t - \frac{d}{c}\right) \quad (3.18)$$

$$= f_0 + \mu \cdot \text{remainder}\left(\frac{t}{T_{sw}}\right) - f_0 - \mu \cdot \text{remainder}\left(\frac{t - \frac{d}{c}}{T_{sw}}\right) \quad (3.19)$$

This formula corresponds to the stepped difference frequency in the middle graph of Figure 3.4. The pulse comes about because of the difference in the time that the step of the remainder functions affects the signal. The longer interval may be written as

$$f_{out}(t) = \frac{\mu d}{c}, \tau \leq t < T_{sw}. \quad (3.20)$$

Remembering that μ is negative for a downward sweep, this formula yields a negative frequency for positive distance d . Noting that negative frequencies cannot be distinguished from positive frequencies without using a complex signal representation, the negative frequency may be safely ignored and the absolute value of the frequency is taken. Thus the final relationship between the demodulation output frequency and the **DOF** d is

$$f_{out}(t) = \begin{cases} \frac{|\mu|d}{c}, & \frac{d}{c} \leq t \leq T_{sw} \\ 0, & \text{otherwise.} \end{cases} \quad (3.21)$$

The frequency is constant but is discontinuous with time, as depicted in Figure 3.4. The frequency $f_{out}(t)$ is zero when $d = 0$, and increases with d , as required.




Figure 3.5: Multiple echoes and associated demodulation output signals. Note that echo times have been increased and the maximum range ratio, m_{rr} , has been set to 0.3 for clarity.

Signal Available for Spectral Analysis

The demodulation output frequency $f_{out}(t)$ only appears after the time of flight τ . It subsequently disappears at the end of the sweep time T_{sw} . This section explains how to calculate the length of the interval where the frequency $f_o(t)$ will always exist, even for distant targets.

The demodulation output frequency $f_{out}(t)$ exists for the time of the demodulation output, $T_{do} = T_{sw} - d/c$. The duration of this frequency is a critical factor in measuring it accurately. The time T_{do} must be as long as possible to overcome the limitations of the [Discrete Fourier Transform \(DFT\)](#), which will be discussed along with other relevant considerations in Chapter 10.

The sonar has a maximum working range, which is a factor in the design of the frequency sweep. All echoes with distances of flight

$$d = ncT_{sw} + \Delta, 0 \leq \Delta < cT_{sw}$$

where n is an integer, will be ambiguous. The spreading losses and air absorption limit the practical maximum range of the sonar system. The sweep time T_{sw} is normally set to be much longer than this practical maximum range, so the ambiguity is never an issue. The ambiguity distance in the sonar system described by this thesis is 55 m.

The spectral analysis of the signal is improved if analysis is begun after all of the echoes have arrived, as this avoids processing partial signals. The maximum range of the sonar is set by a design parameter, the maximum range time T_{max} (Figure 3.5). This is selected with regard to the maximum distance from which echoes can be picked up by the sonar. The maximum range time

T_{max} is specified as a fraction of the sweep time T_{sw} by the maximum range ratio, m_{rr} :

$$T_{max} = m_{rr} T_{sw}. \quad (3.22)$$

The sonar described in this thesis uses $m_{rr} = 0.068$. The signal time T_{sig} over which all echoes of a signal will be simultaneously available for processing is the time from the maximum range echo T_{max} to the end of the demodulation sweep at T_{sw} :

$$T_{sig} = T_{sw} - T_{max} = T_{sw} \cdot (1 - m_{rr}). \quad (3.23)$$

This is the amount of signal available for spectral analysis.

Maximum Range

The maximum distance of flight d_{max} which the sonar can measure is

$$d_{max} = c T_{max} = c m_{rr} T_{sw}. \quad (3.24)$$

The maximum range r_{max} is

$$r_{max} \simeq \frac{d_{max}}{2} = \frac{c m_{rr} T_{sw}}{2}, \quad (3.25)$$

where the approximation is required due to the head geometry. All distance computations refer to distance of flight along the complete sonar travel path to avoid the geometrical error incurred when using range.

The maximum range of the CTFM sonar discussed in this thesis is 1.88m, as discussed in Chapter 9.

Signal Bandwidth

For a given choice of f_0 and f_1 , the transmit signal $s_{out}(t)$ has a bandwidth

$$B = |f_0 - f_1| = |\mu T_{sw}|. \quad (3.26)$$

The sonar discussed in this thesis uses a bandwidth of 50 kHz, so broadband transducers are required. This is a requirement for any working CTFM sonar. The properties and selection of transducers is discussed in Chapter 6.

The bandwidth of the demodulated signal, B_d , is related to the signal bandwidth B by

$$B_d = m_{rr} B. \quad (3.27)$$

This is also the bandwidth available to the spectral analysis stage.

3.3.2 Dual Demodulation

The single demodulation scheme previously described has an interval where no echo signal is present, commonly called blind time. This presents a problem to systems which need to continuously track a target, and also to humans listening to the signal who become bothered by the gaps




Figure 3.6: The operation of a dual demodulation CTFM sonar. The separate demodulation sweeps, along with the signals they produce through the demodulation process, are shown in black and grey for clarity.

(de Roos [30]). It also presents a problem to the spectral analysis stage, which will be discussed in Chapters 10 and 11.

The scheme of dual demodulation, first described by Gough and de Roos [46], uses two demodulation sweeps to fill the gaps present in the single demodulation output signal (Figure 3.6). The two demodulation sweeps are arranged to provide the illusion of a single demodulation sweep over two sweep periods. When the received echo is demodulated using this signal, two sets of sums and differences are produced. In particular, the two difference signals join together, producing a continuous constant frequency where they overlap. It is this continuous frequency which is selected as the output of the dual demodulation CTFM sonar. It retains the property of increasing frequency in proportion to increasing distance of flight. However, zero DOF no longer corresponds to zero frequency, and the output frequency range is higher than for a single demodulation system.

The next few sub-sections provide a theory of operation for dual demodulation CTFM sonars, based on information from Hayes [58].

Dual Demodulation Sweeps

The demodulation sweeps (Figure 3.7) must both have the same slope μ as the transmit signal, and must be separated by the sweep bandwidth B . The upper and lower demodulation sweeps will be denoted by subscripts $d0$ and $d1$ respectively. Thus $\phi_{d0}(t)$ and $\phi_{d1}(t)$ represent the phases of the upper and lower demodulation sweeps, and $f_{d0}(t)$ and $f_{d1}(t)$ represent their frequencies. These are defined similarly to those of the transmit signal, but with an extension of the notation for the sweep frequencies. The first subscript continues to denote start (0) and end (1), while the second subscript denotes the first (upper) sweep with 0 and the second (lower) sweep is denoted by 1.

The frequencies of the sweeps are now determined using the notation. The starting frequency f_{00} of the top demodulation sweep may be freely chosen. As the choice of f_{00} determines the frequency band of the demodulated signal, it is discussed in detail in Section 9.6. The end frequency of the upper sweep f_{10} may be found by subtracting the sweep bandwidth B from the sweep start frequency f_{00} . Using (3.26) for B and removing the absolute value function (μ is negative) gives

$$f_{10} = f_{00} + \mu T_{sw}. \quad (3.28)$$




Figure 3.7: Definition of parameters of the dual demodulation sweeps $f_{d0}(t)$ and $f_{d1}(t)$.

The start frequency of the second sweep f_{01} must be identical to the end frequency of the first sweep, thus

$$f_{01} = f_{10} = f_{00} + \mu T_{sw}. \quad (3.29)$$

And finally the end frequency of the second sweep f_{11} , calculated similarly to (3.28), is

$$f_{11} = f_{01} + \mu T_{sw} = f_{00} + 2\mu T_{sw}. \quad (3.30)$$

The sweep period T_{sw} is the same as that used for the transmit signal, which remains the same as for single demodulation.

The upper demodulation sweep has a phase defined similarly to (3.7):

$$\phi_{d0}(t) = 2\pi f_{00}t_p + \pi\mu t_p^2. \quad (3.31)$$

To maintain the illusion that the dual demodulation sweep consists of a single sweep from f_{00} to f_{11} with period $2T_{sw}$, the time domain signal join between the upper and lower sweeps at time T_{sw} must be continuous. This is achieved by matching the phase of the two signals at that point. The phase of the lower sweep $\phi_{d1}(t)$ is matched to the phase of the upper sweep $\phi_{d0}(t)$ at the joining point:

$$\phi_{d1}(T_{sw}) = \phi_{d0}(T_{sw}). \quad (3.32)$$

As (3.3) requires that the signals are periodic in T_{sw} the signal at $t = T_{sw}$ is identical to that at $t = 0$, so

$$\phi_{d1}(0) = \phi_{d0}(T_{sw}). \quad (3.33)$$

The phase of the lower demodulation sweep $\phi_{d1}(t)$ is defined similarly to (3.7). Noting from (3.16) that frequency is the derivative of phase, it is possible to add an arbitrary constant C_ϕ to the lower demodulation signal phase $\phi_{d1}(t)$ without affecting the frequency $f_{d1}(t)$. Thus the phase of the lower demodulation signal $\phi_{d1}(t)$ becomes

$$\phi_{d1}(t) = 2\pi f_{01}t_p + \pi\mu t_p^2 + C_\phi. \quad (3.34)$$

Using (3.33) and (3.31) to solve for C_ϕ , and temporarily ignoring the periodicity requirement (3.3) gives

$$\begin{aligned} \phi_{d1}(0) = C_\phi &= \phi_{d0}(T_{sw}) \\ &= 2\pi f_{00}T_{sw} + \pi\mu T_{sw}^2 \\ &= 2\pi T_{sw} \cdot \left(f_{00} + \frac{\mu T_{sw}}{2} \right). \end{aligned}$$

Thus the final definition of the lower demodulation sweep phase $\phi_{d1}(t)$ is

$$\phi_{d1}(t) = 2\pi f_{01}t_p + \pi\mu t_p^2 + 2\pi T_{sw} \cdot \left(f_{00} + \frac{\mu T_{sw}}{2} \right), \quad (3.35)$$

which is matched exactly to the end of the upper demodulation sweep. Thus the upper demodulation sweep $s_{d0}(t)$, derived from (3.31) and a relationship similar to (3.4), is

$$s_{d0}(t) = \sin(2\pi f_{00}t_p + \pi\mu t_p^2) \quad (3.36)$$

and the lower demodulation sweep $s_{d1}(t)$, derived in a similar manner from (3.35), is

$$s_{d1}(t) = \sin\left(2\pi f_{01}t_p + \pi\mu t_p^2 + 2\pi T_{sw} \cdot \left(f_{00} + \frac{\mu T_{sw}}{2} \right)\right). \quad (3.37)$$

Demodulation

The demodulation proceeds similarly to that for single demodulation, except that the demodulation signal is now composed of two sweeps:

$$s_d(t) = s_{d0}(t) + s_{d1}(t).$$

The demodulation signal is multiplied with the received signal:

$$s_{out}(t) = s_d(t) \cdot s_r(t) \quad (3.38)$$

$$= (s_{d0}(t) + s_{d1}(t)) \cdot s_r(t) \quad (3.39)$$

$$= s_{d0}(t) \cdot s_r(t) + s_{d1}(t) \cdot s_r(t) \quad (3.40)$$

The derivation proceeds similarly to that for single demodulation - the received signal is substituted for one which is a time delayed replica of the transmit signal. Substituting (3.11) into (3.40) yields

$$s_{out}(t) = s_{d0}(t) \cdot s_t\left(t - \frac{d}{c}\right) + s_{d1}(t) \cdot s_t\left(t - \frac{d}{c}\right). \quad (3.41)$$

Substituting the phase representation of these signals into (3.41) gives

$$\begin{aligned} s_{out}(t) = & \sin \phi_{d0}(t) \cdot \sin \phi_t \left(t - \frac{d}{c} \right) \\ & + \sin \phi_{d1}(t) \cdot \sin \phi_t \left(t - \frac{d}{c} \right). \end{aligned} \quad (3.42)$$

Expanding this using a trigonometric identity and ignoring terms with phase sums (which are filtered out) yields

$$\begin{aligned} s_{out}(t) = & \frac{1}{2} \cos \left(\phi_{d0}(t) - \phi_t \left(t - \frac{d}{c} \right) \right) \\ & + \frac{1}{2} \cos \left(\phi_{d1}(t) - \phi_t \left(t - \frac{d}{c} \right) \right). \end{aligned} \quad (3.43)$$

The frequency of the first of these signals may be computed using instantaneous frequency (3.16):

$$\begin{aligned} f_{out}(t) &= \frac{1}{2\pi} \frac{\partial}{\partial t} \left(\phi_{d0}(t) - \phi_t \left(t - \frac{d}{c} \right) \right) \\ &= f_{d0}(t) - f_t \left(t - \frac{d}{c} \right) \\ &= f_{00} + \mu \cdot \text{remainder} \left(\frac{t}{T_{sw}} \right) - f_0 - \mu \cdot \text{remainder} \left(\frac{t - \frac{d}{c}}{T_{sw}} \right). \end{aligned}$$

This corresponds to the upper difference frequency (grey signal) on Figure 3.6. The frequency corresponding to the signal available for spectral analysis is

$$f_{out}(t) = f_{00} - f_0 + \frac{\mu d}{c}, \tau \leq t < T_{sw}. \quad (3.44)$$

The lower demodulation output frequency may be calculated by repeating the process with the part of (3.42) corresponding to the lower demodulation signal.

Equation (3.44) relates the target distance to the output frequency from the demodulator. Note that the final result for $f_{out}(t)$ is constant during the signal interval, being independent of t . Since μ is negative, an increase in range brings about a reduced frequency. However, there is now a frequency offset, and zero distance now corresponds to a constant frequency which can be adjusted at design time by choosing f_{00} . The negative frequency ambiguity which occurred for single demodulation does not occur here as long as the condition

$$f_{00} - f_0 \geq \frac{\mu d_{max}}{c} = \mu m_{rr} T_{sw} \quad (3.45)$$

is met.

Some Benefits of Dual Demodulation

The dual demodulation technique has filled in the gap in the demodulated signal at $T_{sw} \leq t \leq T_{sw} + d/c$ and periods thereafter which was present in simple demodulation. There is also phase continuity in the demodulated output signal at the sweep period, T_{sw} . However, there will be a

phase discontinuity at the travel time $\tau = d/c$ and integral periods of the sweep time T_{sw} thereafter. This affects the spectral analysis of the signal by [DFT](#) (Chapter 10) and also by [YW](#) (Chapter 11).

The biggest advantage of dual demodulation is that it provides a continuous signal for spectral analysis. There is a possibility of continuously tracking this frequency in order to continuously monitor the position of a moving target. However, there is a phase discontinuity in the frequency at the end of the echo interval, which causes high precision spectrum estimators to lose precision in their output.

3.4 Conclusion

This chapter has collected from the literature the theory behind the operation of two [CTFM](#) sonar types: single demodulation and dual demodulation systems. The single demodulation technique is simpler to implement, but has a 'blind time' in its output which may be a problem to the spectral analyser or to the human listener. The dual demodulation sonar eliminates this problem, but does not eliminate a phase discontinuity in the demodulation output. The dual demodulation technique was selected for use in the remainder of this thesis.

While the techniques outlined in this chapter are well documented in the literature, drawing it all together into the coherent mathematical framework presented here is the work of this author.

Chapter 4

The Signal Path

This chapter collects together various results from the literature to build a thorough understanding of the elements through which the sonar signal passes, and the various domains that are available in which to describe and analyse those signals. It describes the effect that each element, from the transmit signal generator to the detector in the receiver, has upon the signal. This in turn provides an understanding of the operation of the sonar system and the information about the objects in the environment that can be extracted from the echo signal. The elements of the physical sonar path are modelled by the sonar equation. The processing stage within the receiver is also examined. Thus this chapter describes all of the signal transformations performed between the signal generator and the detector.

Typical sonar echo returns vary over several orders of magnitude. Analysis of such signals is impractical in the linear domain, so several alternative domains which aid the analysis of the signals and the discussion of the sonar path are discussed. Furthermore, the alternative domains reduce time domain convolution to simple arithmetical operations, which further simplifies system analysis.

The last two sections analyse the precision and resolution of the sonar in terms of the [signal to noise ratio](#) and [time-bandwidth product](#). Some implications for the design of the transmit signal and the receive processor are also described.

4.1 Domains

The sonar signal, and the elements of the path which determine what will be received, can be analysed in any of several domains. The principal domains are the time domain, the complex frequency domain, the decibel total power domain and the decibel power frequency domain. The relationships between these domains are summarised in Figure 4.1. The domains may be grouped in two different ways. The first grouping is by time domain, frequency domain and total signal power. These groupings are separated horizontally in Figure 4.1. The second, orthogonal, grouping is by amplitude, linear power and decibel power. These groupings are separated vertically in Figure 4.1. As an example of how to use the diagram, the power spectral density, which is usually expressed in decibels, may be found by first taking the Fourier transform of the time domain amplitude signal, obtaining a complex amplitude frequency domain signal, and then performing the $20\log_{10}|z|$ operation to convert to power in decibels.

The different domains are useful for solving different problems, or providing different views

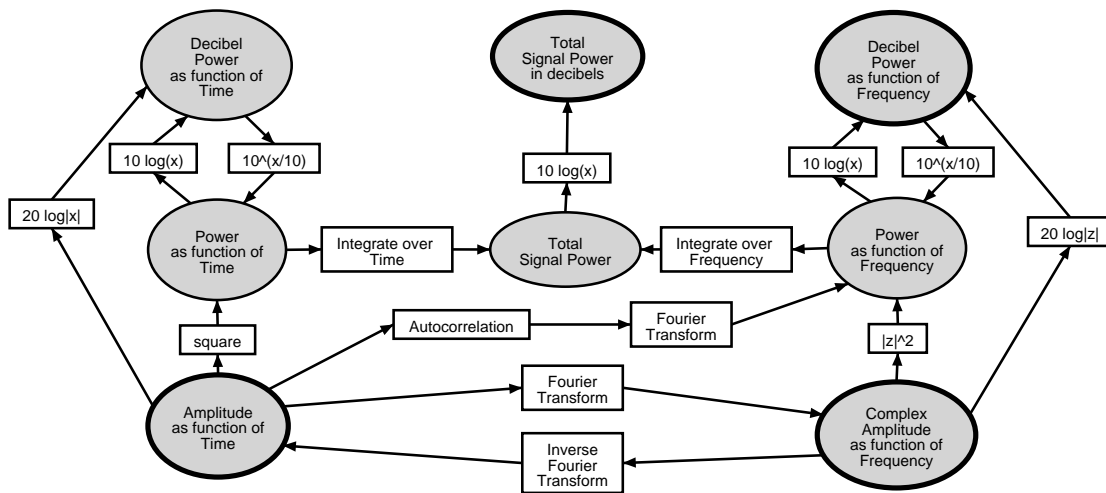


Figure 4.1: The relationships between the time and frequency domains, and the power and decibel power domains. Ovals represent signals, while boxes represent operations and transformations.

of the same problem. For example, the beam-forming operation of the transducer upon a particular signal may be found in the time domain by performing a convolution between the signal and the beam-forming impulse response as described by Kuc and Siegel [85]. Convolution is an integral transform. However, when the signal and the impulse response are transformed into the frequency domain using the Fourier transform, the effect of the filter upon the signal may be found by multiplying the two together (Bracewell [14]). The resulting time domain waveform may be found by taking the inverse Fourier transform. As efficient implementations of the Fourier transform exist, this procedure is often faster than direct time-domain convolution. Furthermore, the analysis and interpretation of the signals is made easier.

The decibel domain is useful as it transforms multiplication into addition, and division into subtraction, which further simplifies some problems. However, when non-decibel units are added, the corresponding operation upon decibel units is the anti-log addition function

$$A \oplus B = \log \{ \exp(A) + \exp(B) \},$$

which could not be considered to be a simplification. In most cases where the decibel domain is applied, the \oplus operator is not required.

Many quantities dealt with in decibels are transfer coefficients from the input to the output of a system and are therefore dimensionless, which simplifies this problem.

4.2 The Sonar Equation

The sonar equation describes the elements forming the path of the signal from transmitter input to the receive processor output. The sonar equation allows the maximum range of the sonar system to be determined. Furthermore, performance limitations can be identified, quantified and addressed. The sonar equation is derived in several texts under different contexts. For a radar derivation, see Berkowitz [11], while maritime sonar is treated by Nielsen [100] and Kinsler [78].




Figure 4.2: The reference path of the sonar signal, showing the principal effects upon the signal. Not all possible physical effects are included. The numbered elements are discussed in the text.

The reference sonar path is shown in Figure 4.2. The reference path is composed of ten elements which are listed in order from the model input to the model output.

1. The electrical signal is first transformed into acoustical energy by the process of transduction within the transmitter. This process is a function of frequency. It is independent of range and bearing.
2. The transducer vibrations are coupled into the air with a directional component. This is transmit beam-forming.
3. The radiating pressure waves become spherical in the far field, and suffer a spreading loss and a propagation delay.
4. The air absorbs some of the energy on the outward path.
5. The sound is reflected by the target. The reflected sound energy depends upon the area and sound reflectivity of the target. These two items collectively are also known as the sonar cross-section.
6. The sound suffers spherical spreading loss and propagation delay on the return path.
7. The air absorbs some of the energy on the return path.
8. The returning sound energy excites the diaphragm of the receiver in a manner which depends upon the incident angle. This is receive beam-forming.
9. The diaphragm vibration is transformed into electrical energy by the process of transduction in the receiver.
10. The receive processor amplifies the signal while rejecting noise, providing processing gain.

Those effects not explicitly considered above may be regarded either as noise or perturbations from the reference model of Figure 4.2. The potential sources of noise are as follows:

1. Environmental noise and interference, such as air-conditioners, music and industrial machinery.
2. Electrical noise in the receive amplifiers.
3. Quantisation noise in the [DAC](#).
4. Processing noise.

The noise amplitude combines additively with the signal. While the last three elements are under the control of the system designer, the environmental noise is not. Although these separate noise sources are recognised, they will not be incorporated into the model.

Potential sources of perturbation from the reference model of Figure [4.2](#) are:

1. Air currents.
2. Non-uniform air temperature, leading to refraction of the sound, causing bending in the sonar path.

Davies [[28](#), Section 7.2, p66] notes that hot objects such as 1) the surface level of a hot cup of coffee and 2) a hot soldering iron are invisible to sonar. It appears that the effect is due to refractory path bending, similar to that found in under-water sonar systems (Kinsler [[78](#), Chapter 15]).

Dealing with effects such as these requires knowledge of the air temperature field in the vicinity of the direct sonar path, so that the refraction may be computed. This problem is outside the scope of this thesis.

Having considered the physical part of the signal path, the electrical part is now considered. The received signal is processed by the receiver before being passed to the detector. The detector usually operates by finding the peak in the output of the receive processor. The target will only be detected if the peak exceeds the detection threshold, which is necessary to reject noise and prevent detection of false targets.

The modelling of the signal path outlined here assumes that each of the elements in the path behaves in a linear fashion, and that frequencies within the signal are not shifted by any element in the sonar path. That is, the elements are assumed to be [Linear Time-Invariant \(LTI\)](#) systems.

The transmitted signal has now been modified by both delays and shape changes in the time domain, by a change in the shape of the power spectra and of the phase response in the frequency domain, and by a change in echo signal power in the power domain. Each of these features may be exploited to extract information about the position and characteristics of the target from the processed echo signal.

Four different representations of the sonar equation are now reviewed. While these representations are equivalent in some respects, each form provides information which is useful in one situation or another.

4.2.1 Power Domain

Consider now the power contained within an echo from a single target. The elements of the signal path may be combined to form the sonar equation. The effect of each element is integrated over the bandwidth of the signal to obtain the total effect, and any travel or filter delays are ignored. This yields the power which can be expected at the detector. The method of combination may be performed in either the power domain or the decibel power domain. The power domain is treated in this section, while the decibel power domain will be treated in Section 4.2.2.

When the signal elements are combined in the power domain, i.e. total signal power in Figure 4.1, multiplicative composition is used to describe the effects of each element. The sonar equation in this domain is

$$p_R = p_T \cdot \left\{ b_T \cdot \left(e^{-\bar{\alpha}r} \right)^2 \cdot \left(\frac{1}{4\pi r^2} \right)^2 \cdot \sigma_c \cdot b_R \cdot p_G \right\}. \quad (4.1)$$

The curly braces distinguish the path elements from the model input and output, which are:

p_R , the received power (model output), and

p_T , the transmitted power (model input).

The terms inside the curly braces correspond to the numbered path elements shown in Figure 4.2.

b_T is the transmitter beam-forming coefficient and transduction (path elements ① and ②).

r is the target range.

$\frac{1}{4\pi r^2}$ is the spreading loss. Two terms are included, one for the outward path and another for the return path (path elements ③ and ⑥).

$e^{-\bar{\alpha}r}$ is the mean absorption of sound by the air over the signal spectrum, over the range to the target. One absorption term is required for each direction (path elements ④ and ⑦).

σ_c is the sonar cross-section of the target (path element ⑤).

b_R is the receiver beam-forming coefficient and transduction (path elements ⑧ and ⑨).

p_G is the processing gain in the receiver (path element ⑩).

These terms will be explained in more detail in Section 4.3. For terms which include a frequency component, the power contribution is determined by integrating the frequency component over the signal bandwidth.

4.2.2 Decibel Power Domain

Equation (4.1) can be converted into decibels of power, the total signal power in decibels block in Figure 4.1, using the relationship $10\log_{10}(x)$. In this domain, the multiplications are transformed to additions, and squares transformed to multiplication by two. The sonar equation in this form is

$$RP = TP + \{TB - 2AL - 2SL + TS + RB + PG\}. \quad (4.2)$$

The curly braces distinguish the path elements from the model input and output, which are:

RP , the received power (model output), and

TP , the transmitted power (model input).

The terms, including references to the numbered path elements shown in Figure 4.2, are:

TB , the transmitter beam-forming and transduction (path elements (1) and (2)),

AL , the absorption loss over range (path elements (4) and (7)),

SL , the spreading loss for range to target (path elements (3) and (6)),

TS , the target strength or cross section (path element (5)),

RB , the receiver beam-forming and transduction (path element (8) and (9)), and

PG , the processor gain (path element (10)).

The principal benefit of this form is that, for example, a 10dB increase in the transmitted power can be seen to directly provide a 10dB increase in the received power. This simplifies system analysis, once the correct parameters for the system at hand have been determined.

The echo excess, EE , being the amount of signal power in excess of the detection threshold DT , is

$$EE = RP - DT. \quad (4.3)$$

The detection threshold is a user specified parameter. It should be set by reference to the noise power at the input to the detector. Further guidance on setting the detection threshold may be found in Van Trees [132, Chapter 2]. The topic will be discussed further in Section 13.2.4.

Plots

The power in decibel form provides a convenient form for plotting the sonar equation, as several of the terms change magnitude by large amounts as a function of range. While features of a small order of magnitude become invisible on a linear plot, they are exposed in full detail in a decibel plot.

A plot of the received power, incorporating the terms of (4.2), for on-axis targets, is shown in Figure 4.3. The graph indicates the received signal power as a function of range for two types of




Figure 4.3: The dependency of the received echo power on range, for a target which is moved along the principal axis of the sensor. This diagram assumes that the transmitter and receiver are co-aligned, and that the path angles at the transducers are negligible. It also assumes unit transduction.

target, a point reflector and a plane reflector. The equations defining the curves will be discussed in Section 4.3. The available signal must exceed the detection threshold before a target can be detected. Thus the graph indicates that point reflectors will become invisible to such a sonar system at 1.4m, while plane reflectors will become invisible at 4.6m. The shape of the curves cannot be altered, but the range may be improved by shifting the vertical relationship between the received power curves and the detection threshold. For example, if the noise floor is dominated by electrical noise in the receiver, then reducing electrical noise and hence the detection threshold will allow the range to be extended. A similar effect may be obtained by increasing the transmit power or improving the processing gain, which shifts up the received power curves.

Plots of the individual terms of the sonar equation (4.2) are shown in Figure 4.4. The terms are discussed individually in Section 4.3.

4.2.3 Time Domain

A third representation of the sonar equation which we will examine is in the time domain. The effect of individual elements is represented by impulse response functions $h(t)$ which are convolved (denoted $*$) with the signal to obtain the output. This form explicitly represents the delay of the signal path, as the impulse responses contain Kronecker deltas, including offsets, to represent delays. The response of the system is

$$s_{rec}(t) = s_{trans}(t) * h_{path}(t), \quad (4.4)$$

where

$s_{rec}(t)$ is the received signal after processing (model output),

$s_{trans}(t)$ is the transmit signal (model input), and




Figure 4.4: Graphs of the terms found in the sonar equation. The top panel shows the only term which depends upon bearing, which is the transducer beam-forming term. This must be included twice, once for the transmitter and again for the receiver. The bottom panel shows the range dependent terms, being the two-way spreading loss and the two-way air absorption. The two-way spreading loss depends upon the cross-section, σ_c , of the target. The point reflector is the weakest possible reflector, while the plane reflector is the strongest possible reflector. Other reflectors fall somewhere in between these two, depending upon their cross-section. The air absorption depends upon the temperature, humidity and pressure. The curves shown are calculated for 20°C and 103.25kPa, with humidity as marked, and averaged over the frequency range 50-100 kHz. The graph shows that the spreading loss is the dominant term.

$h_{path}(t)$ is the impulse response of the path model, described by (4.5) below.

The path elements of the sonar equation are represented by

$$\begin{aligned} h_{path}(t) = & h_{trans}(\theta_T, t) * h_{rad}(r, t) * h_{air}(t) \\ & * h_{target}(t) * h_{rad}(r, t) * h_{air}(t) \\ & * h_{rec}(\theta_R, t) * h_{proc}(t). \end{aligned} \quad (4.5)$$

The path elements of (4.5) can be related to the numbered path elements shown in Figure 4.2, as listed below.

$h_{trans}(\theta_T, t)$ is the impulse response of the transmitter, incorporating transduction and beam-forming (path elements (1) and (2)).

$h_{rad}(r, t)$ is the impulse response corresponding to spherical spreading loss and propagation delay. The delay is modelled by $\delta(t - r/c)$, where c is the speed of sound (path elements (3) and (6)).

The air is assumed not to be dispersive. Figure E.2, shows the variation in the speed of sound to be $0.05 \text{ m} \cdot \text{s}^{-1}$ over the frequency range 50-100 kHz. If we use the mean value of the dispersive correction when calculating the speed of sound (this will be discussed in Section 12.2.3), then the non-dispersive assumption is valid in air.

$h_{air}(t)$ is the impulse response of the air (path elements (4) and (7)).

$h_{target}(t)$ is the impulse response of the target reflection, incorporating the sonar cross-section (path element (5)).

$h_{rec}(\theta_R, t)$ is the impulse response of the receiver, incorporating transduction and beam-forming (path elements (8) and (9)).

$h_{proc}(t)$ is the impulse response of the receive processor (path element (10)).

This form represents each element of the path as a filter. The filters are applied to the transmit signal by convolution, and the received signal is determined. The shape of each impulse response function determines how the signal is re-shaped as it passes through the element. Elements having non-delta impulse responses are the air absorption, $h_{air}(t)$, the transmitter, $h_{trans}(\theta_T, t)$, the receiver, $h_{rec}(\theta_R, t)$, the processor, $h_{proc}(t)$, and the response of the target, $h_{target}(t)$, in the case where the target surface has texture, which is outside the scope of this thesis. The other element, the spherical propagation and travel delay, $h_{rad}(r, t)$, is independent of frequency, and has a purely impulsive time dependency. The effects of the shaped impulse responses are, in general, more easily analysed in the frequency domain.

The impulse model, like the power model, is commutative, and the path elements may be re-ordered as convenient for computation.

Turning now to the benefits and drawbacks of the time domain representation of (4.5), it is noted that the convolution form is relatively inefficient to compute. If the model is only required

to predict the echo from a single target, then only the echo portion of the signal need be computed, and the representation of (4.5) is adequate for computation. However, when working with longer signals, improved efficiency may be obtained by moving to the frequency domain.

4.2.4 Frequency Domain

The computational inefficiency of the time domain convolution form (4.5) of the sonar equation may be avoided by transforming the sonar equation into the frequency domain. Utilising the convolution theorem (Bracewell [14]),

$$\mathcal{F}\{f(t) * g(t)\} = \mathcal{F}\{f(t)\} \cdot \mathcal{F}\{g(t)\} = F(\omega) \cdot G(\omega), \quad (4.6)$$

convolution becomes multiplication under the Fourier transform. The delays modelled using the Kronecker delta function in the time domain are transformed using the shift theorem

$$\mathcal{F}\{\delta(t - t_0)\} = e^{-j\omega t_0}, \quad (4.7)$$

which uses complex phase information to represent the delay. The relationship between the transmit and receive signals becomes

$$S_{rec}(\omega) = S_{trans}(\omega) \cdot H_{path}(\omega), \quad (4.8)$$

where

$S_{rec}(\omega)$ is the complex spectra of the received signal after processing (model output),

$S_{trans}(\omega)$ is the complex spectra of the transmit signal (model input), and

$H_{path}(\omega)$ is the complex amplitude, in the frequency domain, of the path model, described by (4.9) below.

The path elements of the sonar equation are represented by

$$\begin{aligned} H_{path}(\omega) = & H_{trans}(\theta_T, \omega) \cdot H_{rad}(r, \omega) \cdot H_{air}(\omega) \\ & \cdot H_{target}(t) \cdot H_{rad}(r, \omega) \cdot H_{air}(\omega) \\ & \cdot H_{rec}(\theta_R, \omega) \cdot H_{proc}(\omega) \end{aligned} \quad (4.9)$$

The path elements of (4.9) can be related to the numbered path elements of Figure 4.2, as listed below.

$H_{trans}(\theta_T, \omega)$ is the complex amplitude, in the frequency domain, of the transmitter, incorporating transduction and beam-forming (path elements (1) and (2)).

$H_{rad}(r, \omega)$ is the complex amplitude, in the frequency domain, of spherical spreading loss and propagation delay. The delay is modelled by a multiplicative constant phase term $\exp(-j2\pi r/c)$ (path elements (3) and (6)) which assumes that the air is non-dispersive (see Section 4.2.3).

$H_{air}(\omega)$ is the complex amplitude, in the frequency domain, of the air (path elements ④ and ⑦).

$H_{target}(\omega)$ is the complex amplitude, in the frequency domain, of the target reflection, incorporating the sonar cross-section (path element ⑤).

$H_{rec}(\theta_R, \omega)$ is the complex amplitude, in the frequency domain, of the receiver, incorporating transduction and beam-forming (path elements ⑧ and ⑨).

$H_{proc}(\omega)$ is the complex amplitude, in the frequency domain, of the receive processor (path element ⑩).

Equation (4.9) may be used to find the power form by integrating the frequency responses over the bandwidth of the signal used, as indicated in Figure 4.1.

4.2.5 Comparison of Domains

The four domains each have advantages and disadvantages, and can be used to obtain different types of information about the operation of the sonar. For example, the power domain representation is most useful when working out the physical mechanisms of propagation and loss terms. However, when working out the maximum range at which a sonar may operate, the decibel power domain is more convenient, as the multiplications of the power domain are reduced to additions and subtraction. When modelling signal waveforms, either the time or frequency domain representations are used. The most important benefit of the complex amplitude frequency domain representation is that convolution is transformed to multiplication, which may be computed more efficiently. The existence of the [Fast Fourier Transform \(FFT\)](#) algorithm provides an efficient method of transforming data between the time and frequency domains. Furthermore, it is more efficient to perform convolutions via the Fourier transform when long signals are involved. Thus the complex frequency domain is the preferred domain of operation in this thesis.

4.3 Sonar Equation Terms

Each of the elements of the sonar equation will now be discussed in turn.

4.3.1 Transducer Response

The transducer response, describing elements ①, ②, ⑧ and ⑨ in Figure 4.2, is composed of both a transduction term and a beam-forming term. These terms are the subject of Chapter 6.

4.3.2 Spreading Loss and Propagation Delay

The spreading loss and propagation delay correspond to elements ③ and ⑥ in Figure 4.2. The operation of the propagation delay is straightforward and has already been described in Sections 4.2.3 and 4.2.4.

A spherical wave radiating from a point source will cover an area of $4\pi r^2$. As energy must be conserved, the constant amount of energy present must be spread over a larger and larger area as

the wave travels and the sphere expands. Thus the energy density at a point on the sphere must be governed by $(4\pi r^2)^{-1}$. The wave amplitude is therefore governed by $(2\sqrt{\pi}r)^{-1}$. The power spreading loss term may also be converted to power decibels, giving

$$10\log_{10}\left\{\frac{1}{4\pi r^2}\right\} = -20\log_{10}(2\sqrt{\pi}r).$$

This term only includes one direction of travel.

The spreading loss is the dominant term in the sonar equation, and is primarily responsible for limiting the maximum range of the sensor.

Normalisation of Spreading Loss

The spherical spreading loss applies only in the far field, defined to be $r > a^2/\lambda$, where a is the transducer radius and λ is the wavelength. For points close to the transmitter, the power must be determined by other methods. The wavefront in the near field does not spread, so the signal power is conserved for a short distance. The plot of Figure 4.3 assumes that the power detected from a plane and the power detected from a point reflector are equal at the edge of the near field.

4.3.3 Air Absorption

The air absorption corresponds to elements (4) and (7) in Figure 4.2. The propagation of sound through air is not lossless. Some of the energy is absorbed by the air, by exciting various molecular rotational and vibrational modes of the constituent gasses of which the air is composed. The absorption is a function of frequency, temperature and humidity, and weakly depends upon static pressure and air composition, particularly the amount of CO_2 . The air absorption, $\bar{\alpha}$, used in (4.1), is the mean absorption over the bandwidth of the signal. Methods for calculating the absorption of sound in air are given in Appendix E. The information given in the appendix is sufficient to find the spectral magnitude, $|H_{air}(\omega)|$, as a function of frequency.

If the time domain impulse response, $h(t)$, of the air is required, then the spectral magnitude, $|H_{air}(\omega)|$, is insufficient. Simply taking the inverse Fourier transform of the spectral magnitude, $|H_{air}(\omega)|$, provides an impulse response $h(t)$ which is centred about $t = 0$. This may be viewed as a consequence of the symmetry relationship: the Fourier transform of an even (or symmetrical) function is real (has zero phase) (Oppenheim and Schaffer [105, Section 8.2.6]). This relationship implies that the system output signal would arrive before the input. This is clearly impossible. The remedy is to make the impulse response causal, i.e. $h_{air}(t) = 0$ for all $t < 0$, by supplying the missing phase information $\arg\{H_{air}(\omega)\}$. (This is the phase, $\arg\{H_{air}(\omega)\}$, as a function of frequency, ω , not of range, and thus is not a phase of α .) Kleeman and Kuc [82] describe how this may be done for the air absorption function. General methods are also given by Oppenheim and Schaffer [105, Chapter 11]. The solution, however, is not unique, as a system having the desired frequency response can be constructed with arbitrary phase, and hence delay. The Hilbert minimum-phase representation (Hahn [50, Section 7.2.3]) provides the smallest phase representation which makes the entire impulse response causal. Thus the minimum phase criteria must be used to construct unique phase information, providing sufficient information to obtain the complete complex amplitude frequency domain representation, $H_{air}(\omega)$, of Figure 4.1. From there, the inverse Fourier




Figure 4.5: The method of finding a sound path reflected from a plane, using a virtual source. A mirror image, T' , of the transmitter, T , is constructed. T' is the virtual transmitter. The plane reflector is thus a perpendicular bisector of the line TT' . A path may now be drawn from T' to the receiver. The portion of the line on the same side of the reflector as the receiver forms the real path. A line drawn from T to the reflection point completes the path, and can be seen to obey the law of angle of incidence = angle of reflection. The waves radiating from the transmitter have also been drawn. The wavefronts always lie perpendicular to the sonar path. When the waves radiating from T strike the surface, they are reflected, and have a new centre at T' .

transform is used to obtain the impulse response, $h_{air}(t)$.

4.3.4 Target Reflecting Area

The target reflecting area corresponds to path element (5) in Figure 4.2. The echo strength is in proportion to the cross sectional area of the reflector, σ_c . Furthermore, targets with surface texture may be modelled using this term, by incorporating a distributed impulse response rather than a Kronecker delta. Further discussion of texture modelling is outside the scope of this thesis.

Plane reflectors, having a large reflecting area, may be modelled using the principle of the virtual source. A mirror image of the transmitter is constructed behind the reflecting surface, as shown in Figure 4.5. When the reflecting surface is large, the power spreading loss, $(4\pi r^2)^{-2}$, is no longer incurred twice over the range, r , but only once over a range $2r$, becoming $(16\pi r^2)^{-1}$. Thus the spreading loss term should be modified. However, it is possible to create a range dependent source strength for plane reflectors which allows continued use of the double spreading loss terms which were previously defined. This is done by setting the known plane spreading loss, $(16\pi r^2)^{-1}$, equal to the product of the point spreading loss, $(4\pi r^2)^{-2}$, and the plane source strength, $\sigma_{c-plane}$, i.e.

$$\frac{1}{4} \cdot \frac{1}{4\pi r^2} = \left(\frac{1}{4\pi r^2} \right)^2 \sigma_{c-plane}.$$

Solving for $\sigma_{c-plane}$ yields

$$\sigma_{c-plane} = \pi r^2.$$

Such a source strength precisely cancels the double spreading loss and inserts a single spreading loss over a length $2r$ into the sonar equation.

The removal of the double spreading loss explains why plane reflectors provide such a powerful echo when compared with point reflectors, as seen in Figure 4.3.

4.3.5 Processing Gain

The processing gain corresponds to path element ⑩ in Figure 4.2. The receiver will typically provide processing which amplifies the signal while rejecting the noise. This thesis describes two such processors. These are the CTFM sonar system described in Chapter 3, and the MF receiver described in Chapter 8. The processor gain, p_G , is the amount of signal enhancement divided by the noise enhancement, after processing.

The processing described above may introduce delays into the signal path. However, as these filter delays are known absolutely, the signal time-stamp is offset by the processing delay to compensate. (This is not an adjustment for the CPU time taken to perform the processing.)

4.4 Multiple Targets

When an environment contains multiple reflectors which produce echo signals detectable by the receiver, the echoes are combined additively. The received signal is effectively

$$s_{total}(t) = \sum_i s_{rec}(t - \tau_i), \quad (4.10)$$

where $s_{rec}(t - \tau_i)$ refers to the signal received from each target, after passing through the elements described by the sonar equation. The addition physically occurs in the air, and the superimposed wave patterns act collectively upon the receiver. As the systems described are LTI, the addition may be moved to operate upon received and processed echo signals.

4.5 Variable Gain Compensation

In some situations where the sonar system is limited by the quantisation noise of the ADC, shown in Figure 5.4, the noise performance may be improved by using a time dependent gain as described by Bilber [12] and mentioned in Section 5.1.2. The time dependent gain is programmed to compensate for plane reflector spreading loss shown in Figure 4.4. However, the time dependent gain also amplifies the environmental noise, and thus does not help a sonar system which is limited by this noise source (Kinsler and Frey [78, Section 15.8]). Furthermore, it does not reduce internal system noise.

Time dependent gain compensation cannot be used with continuous transmission systems such as CTFM, as time is no longer proportional to range. The range information contained in such signals may only be compensated using a filter, after the demodulation stage (see Section 3.3), with a response calibrated to compensate for the spreading loss. When the sonar system is implemented digitally, there is no benefit to performing this operation, as the signal processing noise floor is very low (the digital signal processing is performed in floating point arithmetic). The principal

noise sources are the electrical amplification noise and ADC quantisation noise. Therefore CTFM systems must use a low noise DAC with a suitable dynamic range, have low system noise, and use increased output power and processing gain to improve long range performance.

4.6 Theoretical Range Resolution

The concept of resolving two targets was previously defined in Section 2.1. A complete practical understanding of resolution depends upon the ambiguity function (see Rihaczec [118, Chapter 1]), but some important theoretical limiting results can be derived from the system bandwidth. This section firstly shows how time and frequency are linked through the time-bandwidth product. Secondly, it discusses processor compression and the effect upon range resolution. Thirdly, the theoretical range resolution limit is discussed. Finally, it gives practical techniques for measuring the time and frequency widths of a signal so that the time-bandwidth product may be evaluated. This will allow a comparison between the theoretical limits and the values actually achieved by the sonar systems developed in this thesis.

4.6.1 The Time-Bandwidth Product

Resolution is related to the width of the signal (see Figure 2.2) as it is presented to the detector, after processing. In an MF sonar, the width is measured in time. In a CTFM sonar, the width is measured in frequency. These quantities are fundamentally linked through the time-bandwidth product relationship. The remainder of this section introduces the concept of the time-bandwidth product, relates it to the problem of resolution in the case of MF and CTFM sonars, and finally discusses implications for the design of the processor before the detector.

The time-bandwidth product relationship is derived using standard Fourier transform theory (Bracewell [14, Chapter 8]). A more fundamental derivation utilising the theory of quantum mechanics is provided by Gabor [43], leading to the relationship

$$\Delta t \cdot \Delta f \geq \frac{1}{2}. \quad (4.11)$$

The value of the constant depends upon the measure of width that is used.

The validity and the implications of relationship (4.11) are now investigated. Consider the Fourier transform of the eternal complex pure tone, $e^{j2\pi f_0 t}$, which is the delta or impulse function $\delta(f - f_0)$ (Bracewell [14, Chapter 6]). As the temporal signal is shortened to make a practical signal, its Fourier representation broadens in frequency, and the delta function begins to blur. The penultimate limit of this sequence is where the time domain signal has itself become a delta function, whose Fourier transform is an infinitely broad-band signal. The concepts of bandwidth and duration are fundamentally linked by the Fourier transform relationship. This is shown in Figure 4.6.

The direct implication is that in order to measure the frequency of a pure tone precisely, it must have a long duration. If the tone is only present for a short time, the uncertainty in its measurement prevents precise establishment of its frequency. Conversely, if the position of a temporal signal is to be precisely measured, it must have a wide bandwidth. Narrowing the bandwidth will cause the

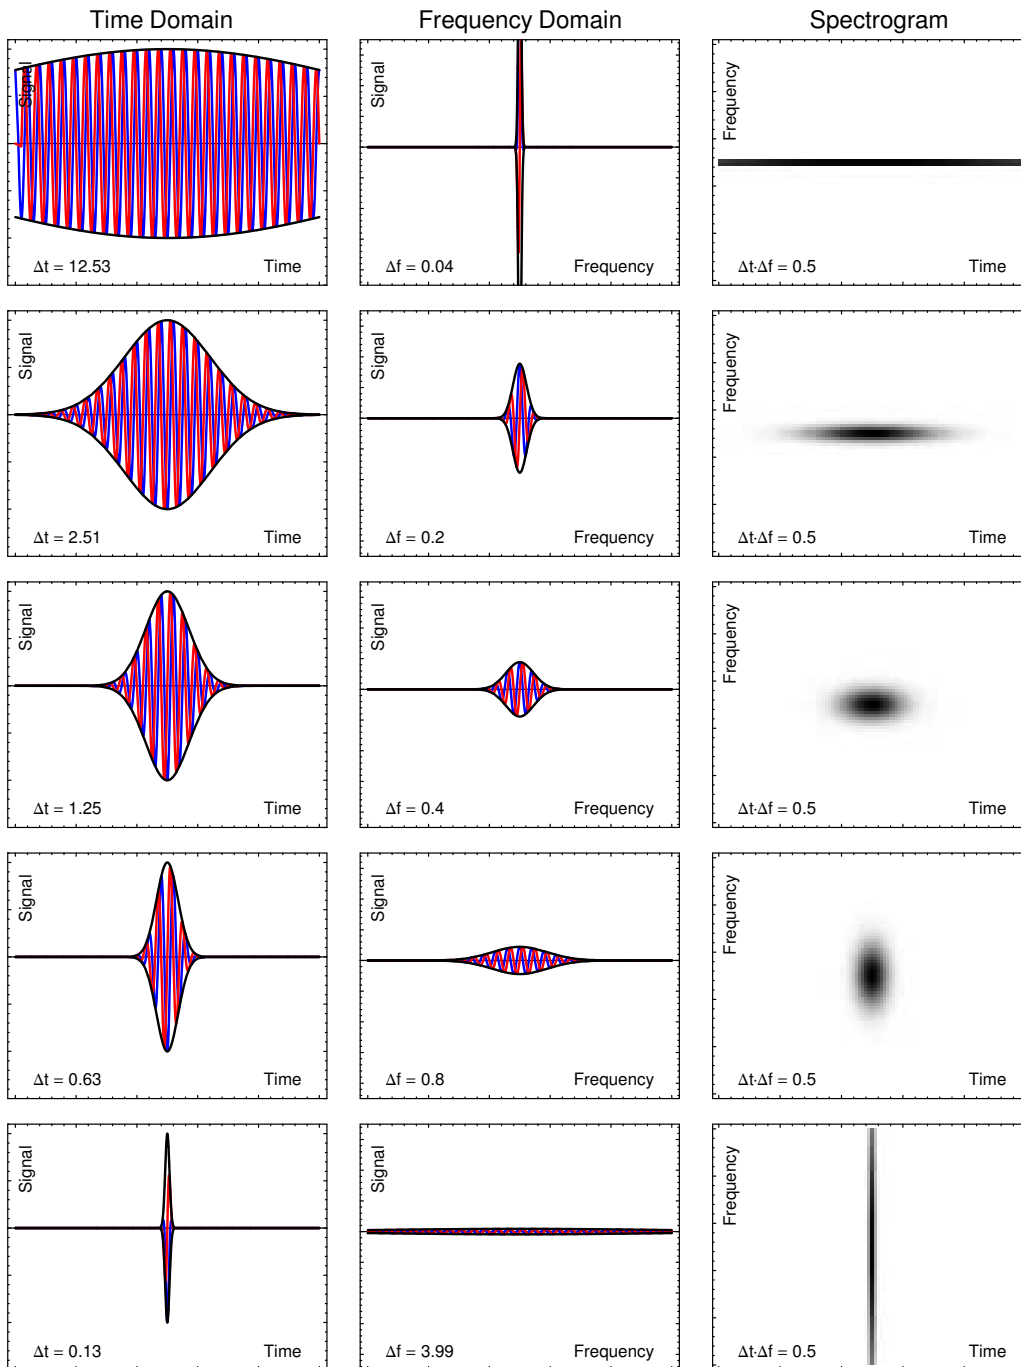


Figure 4.6: Gabor elementary signals, derived from Gabor [43]. The left panels show the time domain signal. The centre panels show the frequency domain signal generated by the Fourier transform (all on the same vertical scale). The right panels show the (normalised) spectrogram of the signal. The time width parameter is varied by row. These signals exhibit the minimum possible [time-bandwidth product](#) $\Delta t \cdot \Delta f$ in accordance with (4.11). In each of the plots blue is real, red is imaginary, and black is envelope. Note that the amplitude of the frequency domain signal changes to reflect the energy content of the time domain signal.

peak to broaden, which in turn broadens the probability distribution of the detector output, causing a precision degradation.

Finally, the implications of (4.11) upon the design of the processor will be considered. It is noted that in the case of a CTFM system, the detection of the signal is carried out by resolving signals in the frequency domain. Therefore the goal of the CTFM processing is to compress the signal into as small a frequency space as possible. On the other hand, in the case of a MF system, the detection of the signal is performed in the time domain. Therefore the goal of the MF processing is to compress the signal into the shortest possible time. The literature commonly refers to this as pulse compression.

4.6.2 Processor Compression

The signal to noise ratio of a sonar system plays an important part in determining its maximum range. The SNR, and hence the range, may be increased by increasing the energy of the transmit signal. The simplest way of doing this is by increasing the voltage used to drive the transmitter. However, there are practical limits upon the operating range of the transducer, which will be discussed in Section 6.2. A signal designed for an amplitude constrained system must be extended in time to obtain the required signal energy and hence SNR. This step will degrade resolution unless some form of receiver processing is used.

The receiver processing compresses the signal in either time or frequency before presenting it to the detector, as shown in Figure 4.7. The MF sonar compresses the signal in the time axis, while the CTFM sonar compresses the signal in frequency. The compression achieves two goals:

1. It improves the signal to noise ratio and facilitates precise detection, and
2. It allows the echoes from two closely spaced targets to be resolved.

Various methods of performing the compression processing are described by Rihaczec [118, Section 3.2] and by Benjamin [10, Section 5]. The most commonly used is the matched filter technique.

4.6.3 Range Resolution Limit

The system bandwidth determines the theoretical range resolution limit of both MF and CTFM sonars (de Roos [29]).

Matched Filter Sonar Resolution Limit

In the case of the MF system, the time width, Δt , of the signal presented to the receiver has a lower bound, which is determined by the Gabor time-bandwidth product relationship (4.11) in terms of the signal bandwidth, Δf . The processed signal bandwidth is limited by the bandwidth of the signal path, as given by equations (4.8) and (4.9). Thus a quantified relationship between range resolution and system bandwidth may be constructed by combining the relationship between

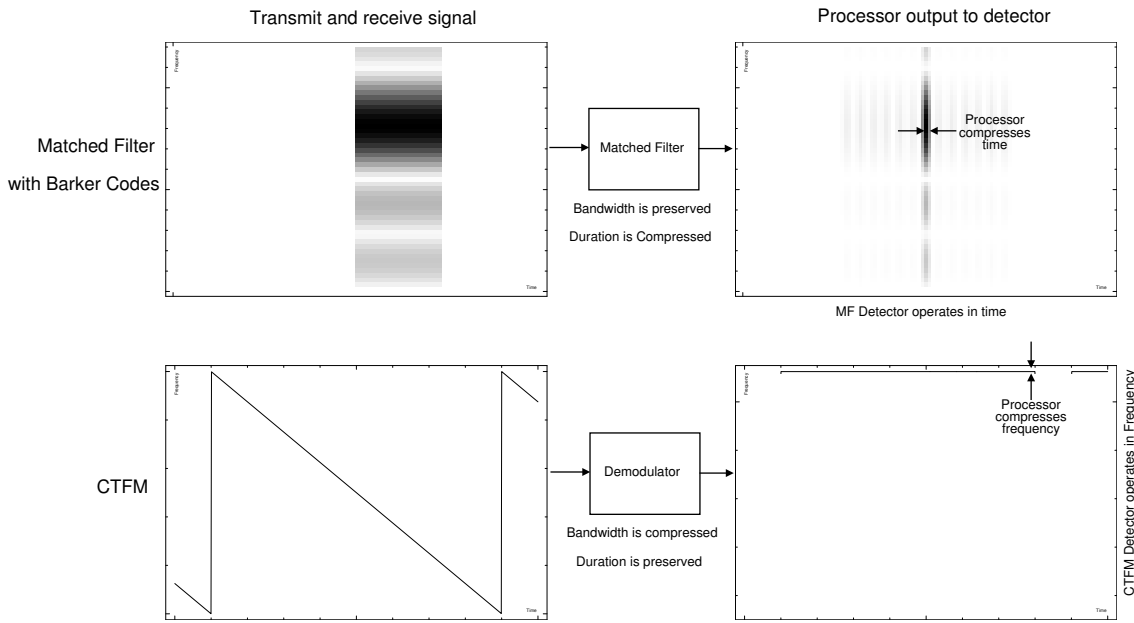


Figure 4.7: Spectrograms showing the effect of receive processors upon the time bandwidth product of two classes of sonar signals. The MF and CTFM classes are duals, with one performing detection in the time domain and the second in the frequency domain. Thus the MF processor performs time compression, while the CTFM processor performs frequency compression.

range width and time width, $\Delta r = c \cdot \Delta t / 2$, with the time-bandwidth product equation (4.11). This procedure yields

$$\Delta r \geq \frac{c}{4 \cdot \Delta f}. \quad (4.12)$$

This provides the theoretical resolution limit for a sonar system utilising a time based detector, such as the MF system. When applying (4.12) to obtain the theoretical resolution limit, the system and signal path bandwidth must be used rather than that of the transmit signal. Resolution degradation will be caused by the following items:

1. If the transmit signal fails to exploit the full bandwidth of the channel, then the receive signal bandwidth will be lower than the maximal amount.
2. Receive processing may reduce the signal bandwidth beyond the reduction imposed by the channel (for example, the MF processor squares the spectrum).
3. The off-axis transducer response, as shown in Figure 6.15, has a reduced bandwidth which will in turn reduce precision.
4. The air absorption, discussed in Section 4.3.3, attenuates the high frequency components of the signal, and thus reduces the bandwidth with increasing range.
5. The signal presented to the detector must have the minimum TB of $1/2$ to achieve the limiting value. In practice, the signal will usually have a larger TB product, contributing to a reduction in precision.

The first of these five causes of resolution degradation may be avoided by proper system design. Item 2 is a consequence of the design of the processor, while items 3 and 4 are unavoidable properties of the transducers and the air. Thus the range resolution predicted by (4.12) is a lower bound rather than a predicted value.

CTFM Sonar Resolution Limit

The relationship (4.12) may also be derived for the CTFM class. Starting with the relationship between the range and DOF widths Δr and Δd respectively,

$$\Delta r = \frac{\Delta d}{2}, \quad (4.13)$$

and substituting the relationship between DOF width and demodulated frequency width, Δf_d , (using (3.21)), $\Delta d = c \cdot \Delta f_d / |\mu|$, yields

$$\Delta r = \frac{c \cdot \Delta f_d}{2 \cdot |\mu|}. \quad (4.14)$$

The demodulated frequency width, Δf_d , is related to the time width, Δt , by the Gabor relationship (4.11). Substituting this into (4.14) yields

$$\Delta r \geq \frac{c}{2 \cdot |\mu|} \cdot \frac{1}{2 \cdot \Delta t}. \quad (4.15)$$

The time width, Δt , is in turn related to the system bandwidth by the slope of the frequency sweep, $\Delta t = \Delta f / |\mu|$ (see Figure 3.2 and equation (3.2)). Substituting this into (4.15) yields

$$\Delta r \geq \frac{c}{4 \cdot \Delta f}, \quad (4.16)$$

which is identical to the resolution limit given by (4.12) for time domain sonars such as the MF system. This demonstrates that the two types of sonar system have the same fundamental resolution limit, and that it is dependent upon the system bandwidth, Δf .

The CTFM sonar system has a reduction in resolution in addition to the list given for the MF system. The CTFM system has the additional step of frequency estimation before passing the signal to the detector. The resolving ability of the spectral estimator adds a sixth potential cause of resolution loss.

4.6.4 Evaluating Time and Frequency Widths

The evaluation of the time-bandwidth product for a particular signal is straightforward if the start and finish times and the lower and upper frequencies are well defined. However, for signals for which the end points are not well defined, measures for an equivalent duration, Δt , and an equivalent bandwidth, Δf , may be defined in the manner shown below. These measures of equivalent width are consistent with the form of the time-bandwidth product given by (4.11).

The measures of equivalent width are computed using the first and second moments of the time and frequency domain signals. These may be computed according to the formulae given by Gabor [43], which have been adapted and summarised here for convenience. The overbar, \bar{x} , will

be used to denote the mean or the first moment. The first moment or mean time position, \bar{t} , of the signal $s(t)$ is given by

$$\bar{t} = \frac{\int_{-\infty}^{\infty} t \cdot s^2(t) dt}{\int_{-\infty}^{\infty} s^2(t) dt}, \quad (4.17)$$

while the second moment, $\overline{t^2}$, is given by

$$\overline{t^2} = \frac{\int_{-\infty}^{\infty} t^2 \cdot s^2(t) dt}{\int_{-\infty}^{\infty} s^2(t) dt}. \quad (4.18)$$

The frequency moments are calculated similarly to the time moments, but the negative frequencies are omitted to avoid the symmetry of the Fourier transform, $S(f)$, of the real signal, $s(t)$. Gabor [43] avoids this by requiring the signals to be analytic, that is, having no negative frequency component. This requirement is avoided in this work by taking only the positive half of the frequency symmetrical spectrum, $S(f)$. Therefore the first moment or mean frequency, \bar{f} , is given by

$$\bar{f} = \frac{\int_0^{\infty} f \cdot |S(f)|^2 df}{\int_0^{\infty} |S(f)|^2 df}, \quad (4.19)$$

while the second moment, $\overline{f^2}$, is given by

$$\overline{f^2} = \frac{\int_0^{\infty} f^2 \cdot |S(f)|^2 df}{\int_0^{\infty} |S(f)|^2 df}. \quad (4.20)$$

The formulae for the moments, (4.17), (4.18), (4.19) and (4.20), may be evaluated upon sampled data by adapting the continuous integrals to equivalent techniques for sampled data.

The moments are then combined to form the measures of equivalent width. The equivalent duration, Δt , is given by

$$\Delta t = \sqrt{2\pi(\overline{t^2} - (\bar{t})^2)} = \sqrt{2\pi(\overline{t^2} - (\bar{t})^2)}, \quad (4.21)$$

while the equivalent bandwidth, Δf , is given by

$$\Delta f = \sqrt{2\pi(\overline{f^2} - (\bar{f})^2)} = \sqrt{2\pi(\overline{f^2} - (\bar{f})^2)}. \quad (4.22)$$

These measures of equivalent width are compatible with the form of the [time-bandwidth product](#) given in equation (4.11), and will be used when estimating the time and frequency widths of various signals.

Practical Note

The widths measured by the techniques described in the previous section may be erroneous when the formulae are applied to a function with a low level sidelobe or noise floor. The [MF](#) auto-correlation function has low level sidelobes which cause problems with the direct application of (4.21). The spectra produced by [CTFM](#) system has a background noise level which causes problems with the direct application of (4.22). These signals must be trimmed before reliable width estimation can be performed using these techniques. Rihaczec [118, Chapter 1] notes that the practical system range resolution is determined by the width of the central auto-correlation or ambiguity function peak. This practical technique is in contrast to equations (4.12) and (4.16) which provide a theoretical resolution limit.

4.7 Precision and Signal to Noise Ratio

The precision with which a target's position can be repeatably measured was previously defined in Section 2.1. Rather than being related to the width of the signal presented to the detector, it is determined by the width of the probability distribution governing the detection position. The detection position is usually defined by the peak of the signal presented to the detector.

Rihaczec [118, Section 3.1] shows for a specific example that precise range measurements may be obtained when there is no system noise, which is equivalent to having an infinite [signal to noise ratio](#). In the more realistic case where noise is present, range performance will be degraded. The effect may be quantified using a technique outlined by Van Trees [132, Section 4.2.3-4, eqn. (110)], but this requires specific knowledge of the sonar signals used. The calculations are quite involved, and will not be applied in this thesis. However, graphs indicating the noise floor will accompany the results where relevant so that the impact of a low [signal to noise ratio](#) upon the results may be identified.

4.8 Conclusion

The time, frequency, power, and decibel power domains have been described, and have been shown to be useful in various circumstances for both understanding and simplifying the sonar equation. Of the four domains presented, the simplest is seen to be the decibel power domain. The convolution of signals in the time domain is transformed to addition in the decibel power domain. However, while this domain is convenient for system evaluation, it cannot support signal modelling. In this case, the complex amplitude frequency domain is preferred, as convolution is reduced to multiplication. This domain is therefore the preferred domain of operation for this thesis.

The sonar equation describes the effects of the sonar path upon the signal, and can be used to predict some aspects of system performance, including the determination of maximum range. Furthermore, in the power domain, the sonar equation can be used to analyse the maximum range, and helps to identify the limiting elements. The time domain and frequency domain representations can be used to build a simulation of the sonar system if sufficient information is known about the transducers and the environmental parameters. Knowledge of the elements of the signal path is necessary to properly interpret the echoes which are returned by the sonar system.

The [signal to noise ratio](#) and the [time-bandwidth product](#) have been shown to be important considerations in the design of a sonar system, having important roles in determining the precision and resolution of the sonar system. The [MF](#) and [CTFM](#) sonar systems have been shown to be equivalent with respect to resolution capabilities, before considering the resolving power of the spectral estimator used. The theoretical range resolution limit will be calculated in Section 8.8 when the system bandwidth is known.

4.8.1 Acknowledgements and Contributions

While this chapter is a synthesis of pre-existing work, especially that of Gabor [43], Bracewell [14], Rihaczec [118], Kuc and Siegel [85], and Kinsler and Frey [78], the organisation of the

discussion into complementary domains which facilitate understanding and processing efficiency is a contribution of the author. The discussion of the variable gain compensation in Section 4.5 is also due to this author.

Part II

Supporting Equipment

Chapter 5

Experimental System Overview

The literature review of Chapter 2 indicates that our *robotic agent should be* equipped with either a matched filter or a [CTFM](#) sonar in order to maximise the sensory volume, and obtain the greatest precision and resolution. This chapter details the supporting sub-systems, including computer software necessary to investigate the performance of [MF](#) and [CTFM](#) sonars.

The pieces of equipment used in the experimental work, and the relationships between them, are shown in Figure 5.1. The central component is the host computer, which controls the experiments. Installed in this computer are two custom-designed [Input and Output \(IO\)](#) cards, the chirp generator card and the chirp capture card. These cards perform digital to analogue and analogue to digital conversion of the sonar signal. The [IO](#) cards, together with the high voltage power supply, are connected to the sensor connector box which serves as a patch panel collecting the signals required by the sonar head into a single cable for convenience. The design of the sonar head will be described in Chapter 7.

The host computer performs a number of tasks including:

- Controlling the three axis positioner which is used in various configurations by the different experiments to quantify the characteristics of the sonar systems.
- Generating and capturing sonar signals, the graphical analysis of those signals and the associated signal processing for the various sonar systems.
- Acting as the client for the environmental sensor system, an independent system responsible for observing temperature, pressure and relative humidity.

There are also some pieces of auxiliary equipment which are mentioned for completeness.

5.1 Chirp Input and Output Cards

At the outset of the project there was no commercial high speed [Digital to Analogue Converter \(DAC\)](#) which met our performance requirements. Hence Phillip M^cKerrow and Michael Milway designed versatile chirp generator and chirp capture cards for use in ultrasonic sensing research. The cards were constructed using wire-wrap technology, making it possible to add new features to the design as required. Both cards plug into a Macintosh computer with a NuBus socket.

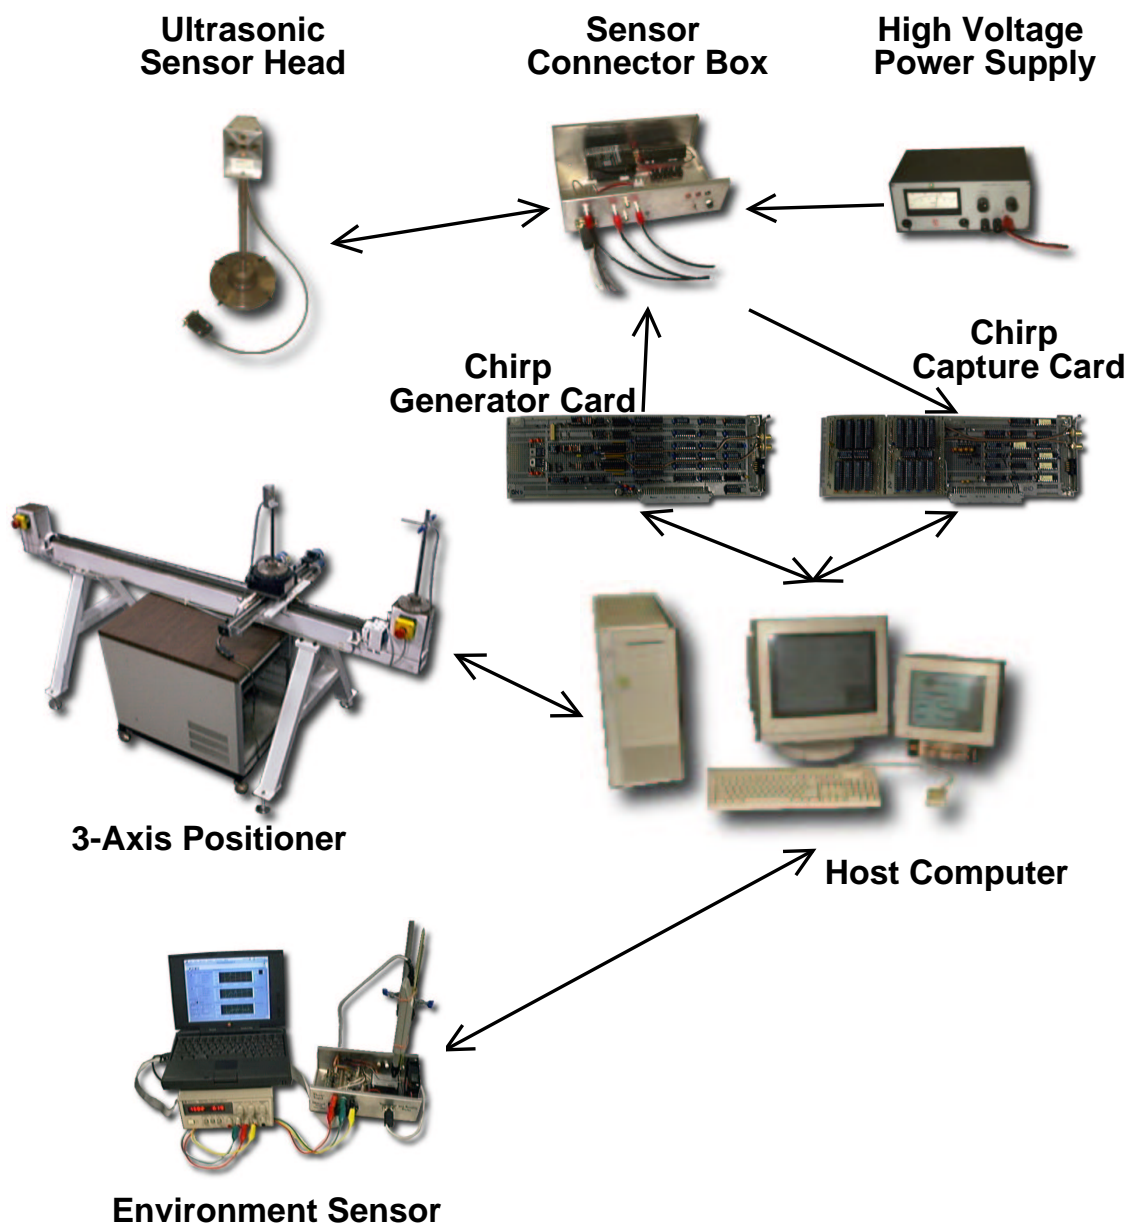


Figure 5.1: The equipment used to support the sonar experiments.

5.1.1 Chirp Generator Card

The chirp generator card is used to convert digital sonar signals into analogue waveforms. A block diagram is shown in Figure 5.2a, and a photograph in Figure 5.2b.

The card contains 2 channels which are functionally identical. Each channel consists of identical functional blocks as shown in Figure 5.2a.

A **First In First Out (FIFO)** memory is used to buffer the data. The **FIFO** memory consists of $4 \times$ Cypress CY7C466A chips, two per channel, structured as $64 \text{ k} \times 9$ bits per chip. The chips are used in parallel to obtain 18 bit words. Our design uses 12 bits to represent the digital signal and 4 bits to provide external trigger signals. Therefore 2 bits remain unused. The trigger bits are used to synchronise chirp capture so as to begin simultaneously with the commencement of chirp generation. The **FIFO** memory can store 163 msec of data at a 400 kHz sampling rate. Thus an arbitrary chirp waveform may be pre-loaded to the **FIFO** memory and subsequently synthesised without **CPU** overhead. As the **CTFM** signals, described in Chapter 3, repeat periodically, the design allows the data in the **FIFO** memory to be re-used and transmitted repeatedly.

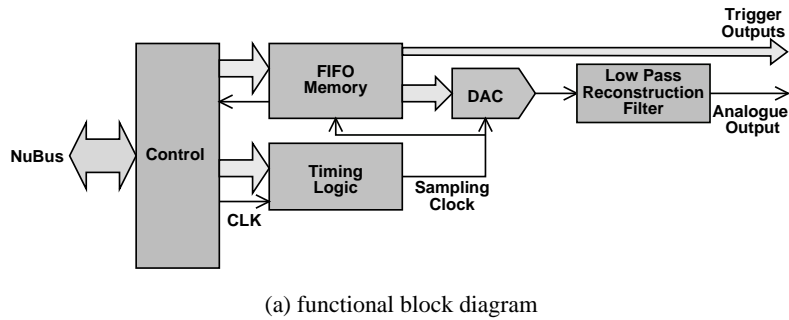
The **DAC** is a 12 bit Analog Devices AD7111. The sampling frequency is program selectable, dividing the 10 MHz NuBus clock frequency by an integer in the range [3...255]. The available clock rates are thus 39.2 kHz to 3.3 MHz.

The low pass reconstruction filter removes harmonic multiples of the synthesised signal which appear due to step reconstruction in the output of the **DAC**. The reconstruction filter serves to smoothly interpolate the signal. The filter is a 4th order Butterworth filter implemented using two Sallen and Key circuits [133, Chapter 6.5]. Its circuit is shown in Figure 5.3a, and the frequency response is shown in Figure 5.3b.


The resistive dividers on the output stage (Figure 5.3a) are to compensate for the filter gain. They are selected to have low values because the cable to the sensor head introduces capacitance to ground. If the output of this circuit has a high impedance, the capacitance shunts the high frequencies to ground.

Construction of an analogue card for use inside a digital computer is a difficult process. Digital circuits tend to radiate lots of high frequency harmonics due to their square waveforms. Keeping this noise out of the analogue circuits is a challenge. Some of the techniques adopted to reduce the noise were:

1. Separating the earth circuits used for the digital and analogue parts of the circuit.
2. Using star earthing throughout the circuit.
3. Using shielded cable for long runs of analogue signal.
4. Separating the digital power supply from analogue circuits. The wire-wrap board has a ground plane on one side and a +5V digital supply plane on the other. Where necessary, this plane was cut.
5. Separating the analogue and digital power supplies by installing separate regulators for the analogue power supplies.




(a) functional block diagram




(b) photograph

Figure 5.2: The chirp generator card. The functional block diagram is shown in (a). The control block contains the bus interface circuitry and the control and status registers. The **First In First Out (FIFO)** memory holds the sonar transmit signal. The timing block contains a programmable clock generator. The **Digital to Analogue Converter (DAC)** and the low pass reconstruction filter serve to remove any digital artifacts from the analogue signal output. Additional trigger outputs are available to indicate the beginning of the transmit phase. Photographs of the front and back of the card are shown in (b). The large grey socket is the NuBus connector, and the two gold plated RCA sockets are the analogue outputs. Two channels are provided. The circuit was constructed using wire-wrap technology.



(a) circuit



(b) filter response

Figure 5.3: The reconstruction filter as used in the generator card. (a) Shows the two Sallen and Key circuits in series. The frequency response of the reconstruction filter is shown in (b). The corner frequency is about 150 kHz, providing a flat response over the region 20kHz – 120kHz.

5.1.2 Chirp Capture Card

The chirp capture card was designed under the assumption that the host computer would be unable to read data at the sampling rate required for ultrasonic data, so a [FIFO](#) memory was used for each channel to allow the computer to read the data after capture has taken place. A block diagram of the capture card is shown in Figure 5.4a, and a photograph is shown in Figure 5.4b.

Two cards were manufactured, each with two channels, to support the four receivers available on the sonar head. The original computer, a Quadra 950, had sufficient space to install all of the cards. However when this computer failed, the replacement PowerMac 8100/80AV had insufficient space to install two capture cards. (Each capture card occupies two slot spaces.) Therefore one card with two channels was used for the remainder of the work.

The incoming signal is first conditioned by a programmable gain amplifier. This gain can be programmed to increase with time, so as to amplify distant echoes more than closer echoes. The gain compensates for the spreading loss mentioned in Section 4.3.2. This feature is useful when working with short sonar signals such as pulse echo or matched filter waveforms, but cannot be used with continuous signals used by [CTFM](#) sonars. For these sonars the gain may be programmed to remain constant.

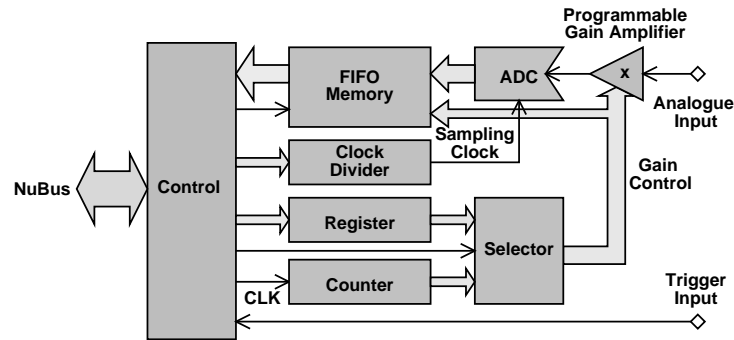
Normally an anti-aliasing filter would be required before the [Analogue to Digital Converter \(ADC\)](#). However the Nyquist frequency (>150 kHz), which is half the sampling frequency (>300 kHz), is always selected to be greater than the maximum frequency in the input signal, so the anti-aliasing filter may be neglected. Even if aliasing is present, it will not overlap the highest frequency of interest (100 kHz) as long as the input frequencies are no higher than 200 kHz. In practice, the receive transducer and amplifier provide the necessary low pass filter to guarantee that no significant aliasing occurs.

The [ADC](#) is a 12 bit Analog Devices AD1671, which has a maximum sampling rate of 1.25 MHz. The capture clock is generated by dividing the 10 MHz NuBus clock frequency by a programmable count, in the range [8...256]. The allowable range of capture frequencies is thus 39 kHz to 1.25 MHz.


The programmable gain bits and the data from the [ADC](#) are then stored in the [FIFO](#). This is constructed from eight Cypress CYC466A 64 k \times 9 bit memory chips. We use two chips side by side to obtain 18 bit words, and 4 chips deep to obtain a 256 kword memory. This is sufficient to record 655 ms of data at a sampling rate of 400 kHz. The [ADC](#) output samples are stored in 12 bits of the [FIFO](#), while a further 3 bits are used to store the programmable gain information and another 1 bit is used to indicate an overflow in the conversion. Therefore 2 bits remain unused.

The chirp capture may be triggered either by programming a register in the card, or through an external trigger input. This mode is usually used in conjunction with the trigger outputs from the chirp generator card. It requires a short ribbon cable connecting the two cards. The trigger signal supplied by the chirp generator card is synchronous with the 10MHz NuBus clock signal, and the first capture sample is initiated by the next NuBus clock cycle divisible by the receive clock divider. The [Programmable Array Logic \(PAL\)](#) components were specifically programmed to obtain triggering which was reproducible down to the underlying 10MHz timebase.

Initial testing of the card revealed a poor signal to noise ratio. This led to a partial re-design of the card, as noted in the previous section. In addition to the techniques noted there, the following



(a) functional block diagram



(b) photograph

Figure 5.4: The chirp capture card. (a) Block diagram of one channel of the capture card. The control block contains bus interface circuitry and control and status registers. The analogue input signal passes through a programmable gain amplifier, which may utilise a fixed gain or a time-dependent gain to compensate for spreading loss. The [Analogue to Digital Converter \(ADC\)](#) produces digital samples of the input signal at a sampling frequency controlled by the clock divider. The digital samples are stored along with the gain which was used at the time in a [FIFO](#) memory, for later retrieval by the host computer system. The circuit is replicated for the second channel. (b) Photograph of the card, front and back. The grey plug is the NuBus connector, while the two gold plated RCA sockets are the analogue inputs. The front view of the card has one of the [FIFO](#) modules removed which is then shown, front and back, underneath the front view. The white paper covers the top of a brass shield, which was found to be necessary to prevent noise from entering the [ADC](#) while writing to the [FIFO](#) memory immediately above.

were also applied to the chirp capture card:

1. Electrostatic shielding over the [ADC](#) chips. These parts are mounted directly underneath the [FIFO](#) memory chips. It was found that as the [FIFO](#) memory chips mounted directly above the [ADC](#) were being written to, the noise increased.
2. The analogue amplifiers have separate drive and sense outputs and inputs. These are designed to be wired separately and connected together at the input to the next circuit, so as to cancel noise in the transfer from one circuit to the other. In the initial construction, these pins had been connected directly together, thereby defeating their design.

With these modifications, the noise has been contained to within 2 bits (3-4 counts) of the [ADC](#).

5.1.3 Chirp IO Software

The selection of the Apple Macintosh platform, and the associated Mac OS 8-9 operating systems, allowed the chirp card driver to be developed as part of the chirp application program without any operating system drivers. This is due to the lack of a protected mode.

The chirp application program was written in the C++ programming language (Stroustrup [127]) using Metrowerks CodeWarrior <http://www.metrowerks.com/>. The graphical user interface is built upon the Power Plant C++ framework which comes with CodeWarrior.

The chirp generation software is divided into two main parts. The first is a low level driver which talks to the chirp generator card directly. The driver code is part of the chirp application program. The second part is a versatile dialogue (Figure 5.5) which allows the user to select the waveform and envelope. A repetition rate may also be set, which is useful for [CTFM](#) signals. The sampling rate of the signal is automatically chosen to be the highest rate possible subject to the constraint of running out of [FIFO](#) memory. Triggers may be placed at arbitrary locations in the signal. This dialogue is mostly used for quick investigations and for generating suitable signals for tracing hardware faults. There is also an additional layer of software, which allows a file or network stream to be substituted for the actual chirp card hardware. This feature will be described more fully in Section 5.4.3.

The chirp capture software comprises a driver and a dialogue for setting up the various options on the capture card, such as the sampling rate, the gain and the duration of the capture. A picture of the chirp capture dialogue is shown in Figure 5.6. Chirp capture may be initiated asynchronously from chirp generation, or may be synchronised through the use of triggers.

Once a signal has been captured, it is displayed in a zoomable window. Various operations are available, such as Fourier transform and low pass filtering, for performing rudimentary analysis. The data may also be saved to a file in various formats for further investigation in programmable environments such as Matlab [130] or Mathematica [139].

5.2 Precision Positioner

The precision positioner, shown in Figure 5.7, allows all five sonar systems to be tested with repeatable and computable geometry. This avoids the odometry errors, typically associated with

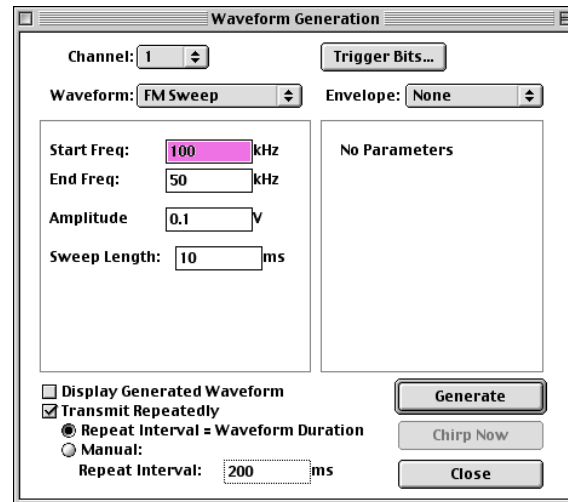


Figure 5.5: Dialogue for generating sonar waveforms for experimentation. Several types of waveforms (sine-wave, pulse, Barker code, FM Sweep) may be combined with any of several envelope functions (trapezoidal, triangular, Gaussian) and transmitted to the sonar. This dialogue also allows the repetition function of the chirp card to be controlled.

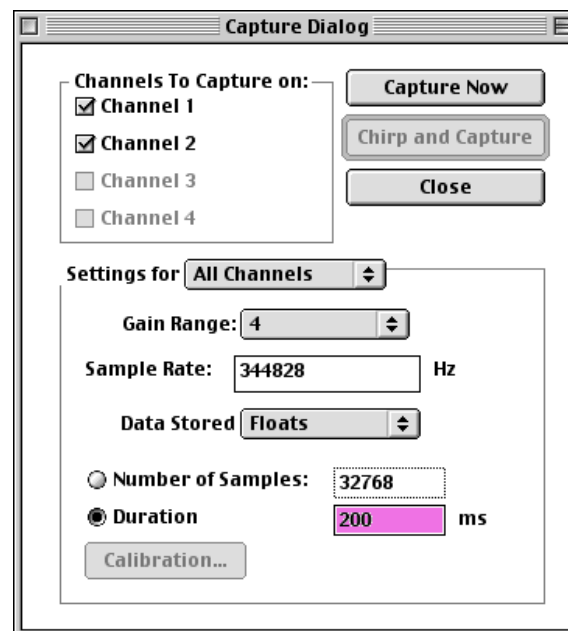


Figure 5.6: The chirp capture dialogue which allows specification of sampling frequency, the gain used for the pre-amplifier, selection of the channels to use, and capture duration. This dialogue may be used to capture an ultrasonic waveform or to transmit and capture simultaneously.




Figure 5.7: The precision positioner. Three axes of movement are defined as machine coordinates. These are longitudinal, x , along the base, transverse, y , along the cross slide, and rotational, θ , through a rotating turret mounted on the cross slide. Positive x movement is toward the upper left and positive y movement is toward the lower left. Positive rotation is defined by the software to be anti-clockwise, providing a right-handed coordinate system. The home position is at the rear right. The box underneath houses stepper motor control modules.

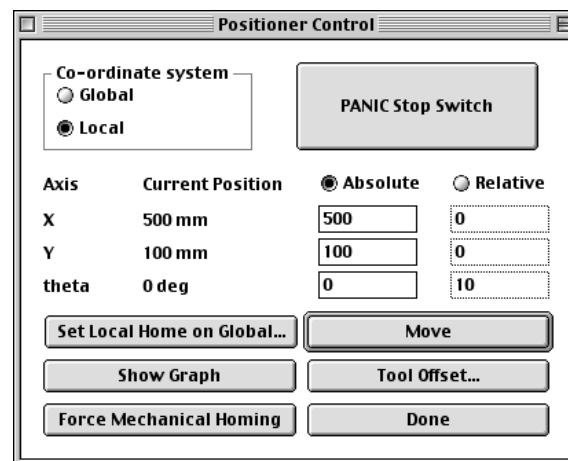


Figure 5.8: The dialogue used to control the positioner. Movement may be controlled in global (machine) coordinates or local coordinates, and may be to an absolute position or a relative position. The show graph button brings up the graph shown in Figure 5.9.

mobile robot platforms, which would otherwise require extra effort to avoid (for example Chong [22]).

The positioner provides three interchangeable mounting points for attaching sonar heads or targets, one on the movable platform, and one at each end. The sensory equipment shown in Figure 5.7 is mounted atop short poles. The poles are mounted in special bases with keys matching the key holes in the positioner. The bases are precisely manufactured to hold the poles vertical to avoid positioning errors as the turret is rotated.

The control of position is achieved through an integrated hardware and software system.

5.2.1 Hardware

The precision positioner was built to a design by Phillip M^cKerrow and Michael Milway, and allows movement under computer control on three axes, longitudinal, transverse and rotational, denoted by the right handed tuple $\{x, y, \theta\}$. Stepper motors produce indexed motion through ball-screw drives. The position step quanta are $\Delta x = 0.0254 \text{ mm}$ ($= 0.001 \text{ inch}$), $\Delta y = 0.05 \text{ mm}$ ($= 0.002 \text{ inch}$) and $\Delta \theta = 0.02^\circ$. The stepper motors are controlled by indexing modules, which are connected through a serial cable to the host computer.

5.2.2 Software

Harper [54] wrote the original positioner control software in Modula-2. The code was re-written in C++ so that it could be integrated with the chirp application program running on the host computer. A programming interface is provided, as well as a dialogue for interactive use (Figure 5.8).

The software utilises four different coordinate systems (Figure 5.9). The first is the machine coordinate system, which is defined by the extents of travel. The second is the global coordinate system which extends the machine coordinates outside the limits of the machine. The third is the local coordinate system, which is designed to be matched to the experiment at hand and is defined




Figure 5.9: The graph displayed by the positioner control software. The graph is in global coordinates, whose origin coincides with machine coordinates as defined by the limit switches. The black axes denote the machine coordinates and the movement extent of the centre point of the turret. The x -axis moves from one end of the machine to the other, while the y -axis corresponds to the cross-slide, as shown in Figure 5.7. The rotating turret is drawn as a black circle, with a black radial line indicating the direction which is aligned with the machine x -axis when the machine is homed to the limit switches. The software defines local, turret and sensor coordinate systems which are shown. This graph shows the sensor coordinates are rotated by 90° with respect to the turret coordinates. This graph updates to reflect the movement of the positioner.

to be parallel to global coordinates. The offset between them is set so that the origin of the local coordinate system coincides with the target or the sonar in the experimental setup. There is also a θ offset, allowing the mounting holes on the turret to be pointed parallel to the x -axis of the local coordinates.

The third coordinate system is the sensor coordinate system. There is a tool offset between the turret origin and the sensor coordinate origin. This may be used to compensate for an off-centre sensor or target. The $\{x, y, \theta\}$ coordinates entered by the user control the origin and direction of the sensor coordinate system. The rotation is about the origin of the sensor coordinates, and is measured from the x -axis of the local coordinates. The coordinate transformations were computed using homogeneous transformation matrices (McKerrow [95]).

The three coordinate systems can be shown interactively on a live graph. A sample graph is shown in Figure 5.9. The live graph was found to be useful feedback while setting up an experiment, as it allows the experimenter to verify that the physical positioner location is the same as the computer model of the positioner location.

After properly setting up the local coordinates and tool offsets, experiments can be performed by manipulating only the position of the sensor coordinate system in the local coordinate frame. This feature simplified the code which controls each experiment.

5.3 Environment Sensor

The formulae for calculating the speed of sound, which will be discussed in Chapter 12, require temperature, relative humidity, and pressure. The environment sensor measures these environmental parameters and makes them available to the chirp application program.




Figure 5.10: The environment sensor computer, with the connection box, power supply and probe.

5.3.1 Electronics

The environment sensing system consists of two main parts (see Figure 5.10). The first is the interface electronics box, along with the associated probes and transducers which measure the environmental parameters. The second part is a computer which contains an ADC card and runs custom software to take the measurements and make them available to the sonar program. It was necessary to have the environment sensor separate from the host computer due to a shortage of slot space in the host computer. This arrangement proved to be convenient as the environment sensor could be moved around to perform calibrations.

The connection box, shown in Figure 5.10, contains a Vaisala pressure transducer, model PTA427, and a Vaisala temperature and humidity probe, model HMP-35A. The HMP-35A does not contain any readout or signal conditioning electronics.

The computer system was built using a Macintosh PowerBook 5300c with a National Instruments DAQCard-1200 PCMCIA credit card adaptor (<http://www.natinst.com/>). The DAQCard provides 8 analogue inputs in single ended mode or 4 analogue inputs in differential mode. The differential input mode was selected as it gave better noise rejection characteristics. Unfortunately the four channels cannot be read simultaneously with anything other than a unity gain applied, so the full scale of 0-10V must be used to achieve the stated 12-bit (4096 count) resolution. The signal conditioning circuit was re-designed (Figure 5.11) so that the signals were presented to the ADC with their expected full range of variation covering the 0-10V range.

The signal conditioning for the LM335 zener diode required the subtraction of a voltage before scaling. Standard circuits for solving this problem require highly matched resistors for accurate operation. This requirement was avoided by using an instrumentation quality differential amplifier circuit (Hambley [52, Section 11.9.1]) which does not require precision resistor matching for accurate operation. Alternatively, one of the many available signal conditioning integrated circuits could have been used.

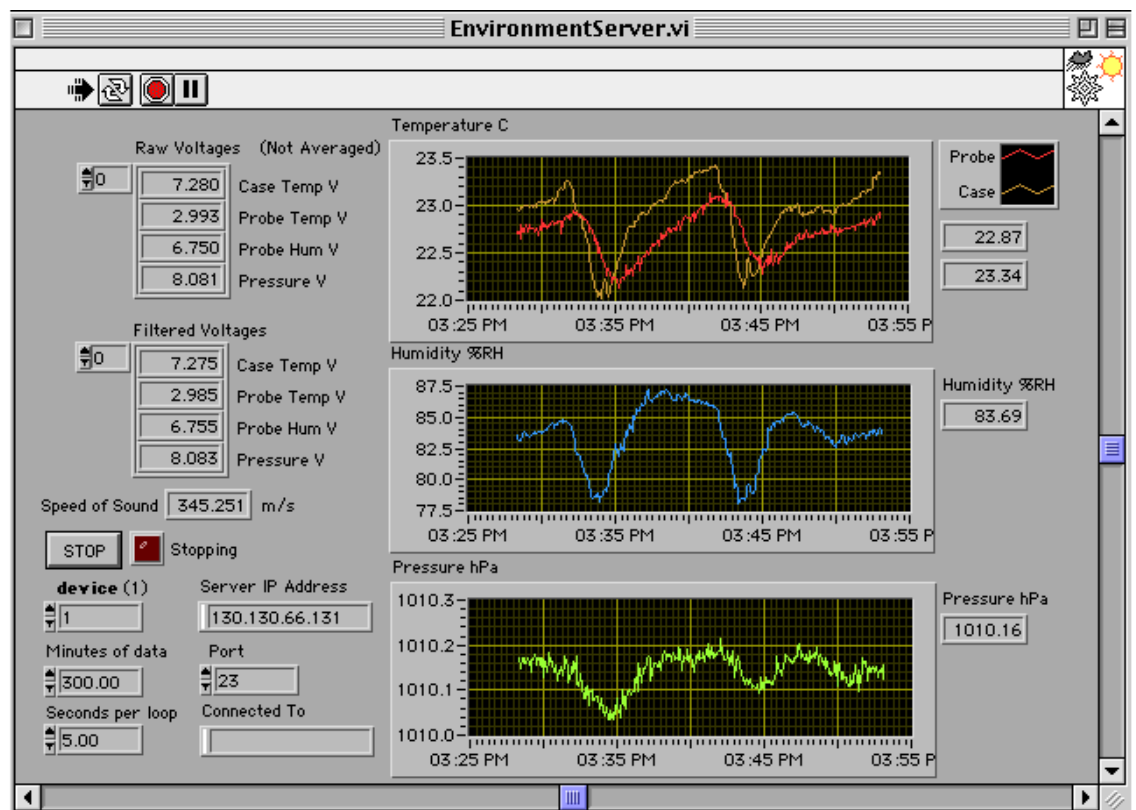


Figure 5.12: Screen-shot of the environment sensor program running in LabView. The three graphs show the temperature, humidity and pressure. The two temperature traces are from the HMP-35A probe (red) and the LM335 zener diode (brown).

5.3.2 Software

The environment sensor software was written using LabView 5.1 from National Instruments (<http://www.natinst.com/>). James Martin implemented the first LabView prototype. The final version was written, debugged and tested by this author in a matter of days, which compared favourably with the weeks it took to write and debug the initial C++ program.

Initial tests showed that there was some noise in the voltage readings. This was dealt with by taking readings from the ADC at a sampling rate of 100Hz and then low pass filtering each of the 4 channels using a Finite Impulse Response (FIR) filter to remove any frequencies above 5 Hz. The number of samples taken was equal to the filter length. The output of the filter (after the initial transient) was used as the sample point at that sample time.

The environment sensor program contains the calibration formulae which convert the raw voltages read by the ADC into Kelvin, % Relative Humidity and Pascals. The procedure used to obtain these calibration formulae will be described in Section 12.4.

The front panel of the LabView program is shown in Figure 5.12. The graphs were constructed by taking one (filtered) reading once every 5 seconds over a period of several hours. The speed of sound shown on the screen-shot is calculated using (12.7).

Sometimes the observed pressure has a downward spike, or performs an up or down step. At

first it was thought that there was noise in the system. However, our laboratory is pressurised by the air conditioner. When the door is open, a strong draught is felt. The downward spikes in pressure happen when the door is opened momentarily while the air conditioning is on. The steps in the pressure are the air conditioning being turned on and off.

Another phenomenon typically observed in our laboratory is the air conditioner cooling turning on and off, which can be seen in Figure 5.12. Note that the HMP-35A probe fails to respond adequately to the temperature variations. This will be discussed further in Section 12.4.1.

5.4 Host Computer

The chirp application program which runs on the host computer accumulated many features in addition to the ones described so far. The following list describes some of the more important features.

5.4.1 Transducer Response Measurement

A module was written to allow automated measurement of a transducer beam pattern. This subroutine utilised the chirp generation and capture code as well as rotation of the sonar head using the positioner. The signal to be used is chosen using the chirp generation dialogue, and is then loaded into the generator card. Similarly, the settings for the chirp capture are set up using the chirp capture dialogue. The subroutine accepts as arguments the starting and ending positions of the traverse, the number of stops in between, and which axes should be moved. When the experiment is started, it progresses automatically and saves the individual captured chirp data to a file. The current position and readings from the environment sensor are also stored for later analysis in environments such as Matlab or Mathematica.

This arrangement is capable of measuring the on-axis frequency response of a transducer, by setting the transmitted signal to contain frequencies of interest by adopting a frequency sweep. It can also be used to generate a complete picture of the lobe structure of the transducer, in both frequency and direction angle. An example of such a plot is shown in Figure 6.15. The plots are generated by Mathematica, using data captured by the chirp application program.

5.4.2 Feature Extraction Dialogue

The feature extraction dialogue, shown in Figure 5.13, allows the selection of all sonar systems described in Figure 1.2.

This dialogue allows the user to choose a complete pre-configured sonar system and shows its various processing stages. The panels in the dialogue allow the parameters of each stage of the sonar system to be adjusted during a sequence of trials. Problems may be investigated using debugging output, in the form of graphical or textual output, by checking a box on a per-stage basis.

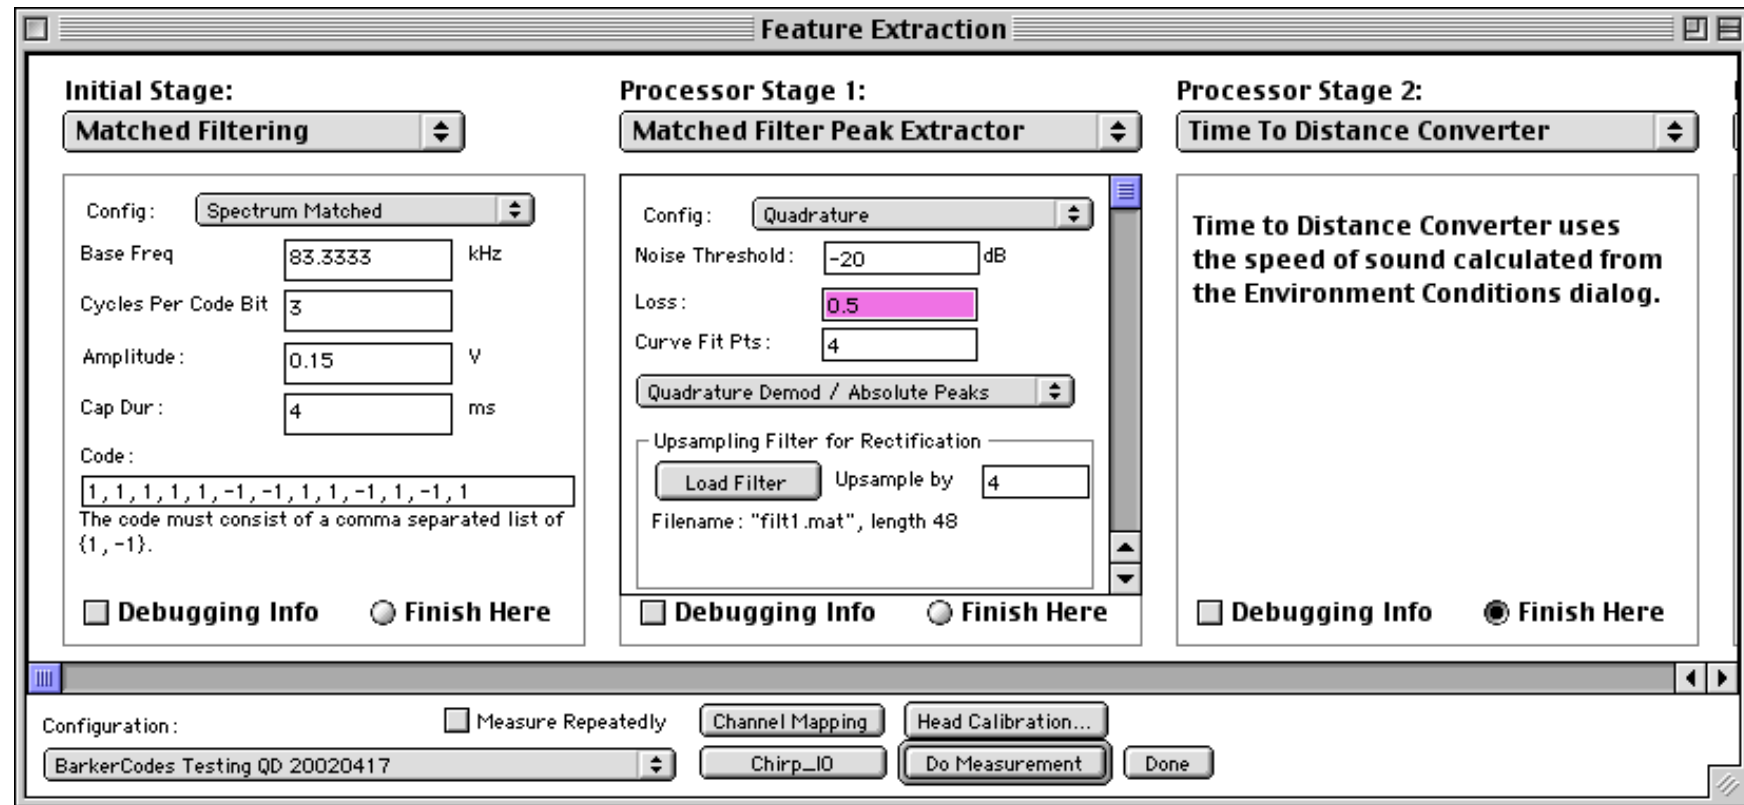


Figure 5.13: The feature extraction dialogue. This is the heart of the sonar system. Each panel represents a stage in the processing of a sonar system. Stages execute from left to right. The leftmost stage generates the transmit signal and performs some basic echo processing. Subsequent stages take input and produce output like a pipeline. Stages may be interchangeable as long as the exchanged data types (not shown) are compatible. Each stage is capable of producing debugging output, usually in the form of graphs. A configuration management system is implemented for each stage and for the dialogue as a whole. In practise, the different sonar systems are selected from a pre-configured arrangement by selecting a configuration from the pop-up menu at the bottom left.

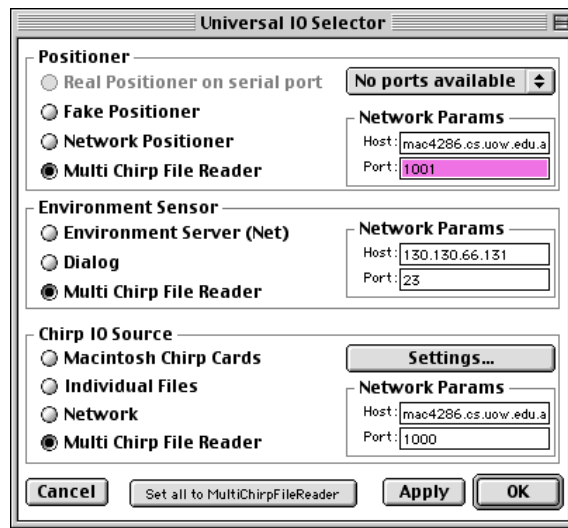


Figure 5.14: The universal IO selector dialogue, which allows the user to select the driver sub-system used for the positioner, the environment sensor and the chirp IO sub-system. All sub-systems have drivers for physical devices as well as drivers which can read input from a file or connect remotely over a network. The MultiChirpFile option refers to the file which stores results from all three IO sub-systems for later playback.

5.4.3 Hardware Abstraction

The chirp generation and capture software, previously described in Section 5.1.3, is implemented using a hardware abstraction layer which allows the driver to be replaced by a file interface or a network connection. There are three reasons for implementing such an abstraction layer.

1. The file interface allows sonar systems to be tested and debugged off-line.
2. The network interface allows the use of a faster machine to perform on-line signal processing. As the chirp cards were built using NuBus interfaces which were discontinued, the computer containing these cards could not be upgraded. The network interface allowed a faster, current model computer to read sonar data from the chirp cards in the slow computer. This system performed better than running the entire system on the slow computer.
3. The file interface also allowed one set of sonar echoes to be recorded and processed by different sonar systems or configurations for evaluation under identical conditions.

The abstraction layer developed for the chirp IO sub-system was found to be useful and was extended to the positioner control system and the environment sensor. In particular, readings from all of the IO sub-systems comprising sonar echoes, locations of the positioner, and readings of the environmental parameters, could all be recorded into a single file for later playback and analysis.

The various options for the three IO sub-systems used by the chirp application program are shown in Figure 5.14.

5.5 Auxiliary Equipment

Some additional equipment which was used is mentioned for completeness.

5.5.1 High Voltage Power Supply

A high voltage DC power supply, Delta Elektronika E0300-0.1, was used to provide a stable and reliable bias voltage to the transducers. There is a picture in the top right of Figure 5.1. This power supply is continuously adjustable from 0 to 300V, and has a built in current limit adjustable up to 0.1A.

5.5.2 Ultrasonic Microphone

A Brüel & Kjær 1/4 inch condenser microphone, type 4135, and associated microphone power supply type 2804 (Figure 5.15a), were used for measuring frequency responses of ultrasonic transducers.

5.5.3 Calibrated Source

A Brüel & Kjær sound level calibrator, type 4230 (Figure 5.15b), was used to normalise the measurements taken using the ultrasonic microphone.

5.5.4 Vaisala HM-34

A Vaisala hand-held temperature and humidity probe, model HM-34, was used for making various measurements and as a reference instrument for calibrating humidity (Section 12.4.2) is shown in Figure 5.15c. This instrument is calibrated by Vaisala.

5.6 Conclusion

This chapter has described the experimental system supporting the work of this thesis. The host computer executes the chirp application program, which coordinates three other sub-systems to perform the experiment. The three sub-systems are

1. the signal path through the chirp IO system and associated amplifiers, interconnections and the sensor head.
2. the positioner.
3. the environment sensor.

The transducers used within the sensor head will be described in Chapter 6, and the ultrasonic sensor head itself will be described in Chapter 7. The other auxiliary equipment is used either within the signal path or for calibrating other sub-systems. Together these sub-systems underpin the experimental work of this thesis.



Figure 5.15: Auxiliary equipment used for experiments. (a) Brüel & Kjær Ultrasonic Microphone and pre-amplifier, (b) Brüel & Kjær calibrated source, and (c) Vaisala HM-34 hand-held temperature and humidity probe.

5.6.1 Acknowledgements and Contributions

The large amount of equipment required to perform the experimental work could not possibly be constructed or maintained by one person. The equipment was built up by a small group over several years. Those responsible for each part are named here.

1. The chirp IO cards described in Section 5.1 were designed by Phillip M^cKerrow and Michael Milway, as credited in the text. The original cards were constructed by Michael Milway. A further card was constructed by the author. The cards were later upgraded by the author to improve the signal to noise ratio and various extra timing and synchronisation features. This author performed all necessary design, construction, testing and fault-finding.
2. The precision positioner described in Section 5.2 was designed by Phillip M^cKerrow and constructed by an external contractor. The electrical control system was constructed by Michael Milway. The initial version of the software was written in Modula2 by Neil Harper. The author of the thesis translated the Modula2 version into C++, with numerous modifications and extensions to improve flexibility and reliability.
3. The environment sensor described in Section 5.3 was initially designed by Michael Milway, but the design provided inadequate sampling resolution and had to be revised by the author. The sensor interface box was constructed by support staff. Modifications were performed by this author. The software was written entirely by this author.
4. The software executing within the host computer described in Section 5.4 was written by this author. This software is the embodiment of the major contributions of this thesis. The software also makes use of standard libraries. These include:
 - (a) BLAS [34] and LAPACK [4] for performing linear algebra.
 - (b) FFTW [41] for performing the Fast Fourier Transform (FFT).
 - (c) Metrowerks PowerPlant [92] for creating the graphical user interface.
5. The auxiliary equipment described in Section 5.5 was all purchased from external sources, and is mentioned only for completeness.

Chapter 6

Transducers

The operation of a sonar system critically depends on the quality and the properties of the ultrasonic transducers which are used. For example, the properties of the transducer determine the exact shape of the impulse response which is used in Kleeman's template matching system, discussed in Section 2.5.6. The properties of the transducer provide the frequency response as a function of bearing which allowed Yata [142] to measure the target angle from the transducer normal with a single transducer, as discussed in Section 2.5.5. Thus the transducers are one of the critical components around which the rest of the system is built. In this chapter, the principles of transducer operation are described and published models of beam-forming are reviewed, using the Kay transducer which was adopted for this study.

The Kay transducer was developed for use in the KASPA system described in Section 1.3.2. The Kay transducers were specifically developed to have a wide bandwidth supporting the one octave sweeps used by CTFM sonar systems, as described in Chapter 3.

The final parts of this chapter review the performance of the mass-produced Polaroid transducer which is descended from Kay's transducers (Kay [74]). The performance of this transducer is contrasted to the Kay transducers.

6.1 Construction of the Kay Transducer

The design parameters and method of construction of electrostatic ultrasonic transducers is discussed at length by Martin [90]. The transducers in our units are slightly different from Martin's, but the theory of operation is the same.

The transducers are of the solid dielectric type. They are constructed with a backing plate, a mylar film with a vacuum-deposited gold conductor on one side, and a frame to hold the mylar film tight on the backing plate (Figure 6.1). The mylar is glued to the diaphragm mounting ring, and then clamped in place by the case. The backing plate is pressed against the back of the diaphragm so that it pushes it forward by 2 thousandths of an inch (thou). This places the diaphragm under tension. The backing plate (Figure 6.2) has 10thou (0.254mm) wide circular grooves in it, with 10 thou spacing in between. The grooves are 20thou (0.508mm) deep. The diaphragm is supported by microscopic high points on the surface of the backplate. There is a small layer of air in between the diaphragm and the backplate. The air provides the restoring force on the diaphragm, rather than the tension in the diaphragm as might be supposed. The grooves are used to increase the compliance of the air layer. They provide a reservoir of air which allows the diaphragm to




Figure 6.1: The construction of the Kay transducer described in [71, 72, 90]. The backing plate forms one electrode, while the evaporated gold surface on the front of the mylar diaphragm forms the other electrode. The dielectric is the diaphragm itself, which is tensioned by the backplate. The grooves supply a reservoir of air, which allows the mylar film to move more than it would if the grooves were not present. The electrostatic effect operates in the region at the top of the grooves.

move more freely than without the grooves. The smoothness of the top surface of the backing plate is one of the most important variables in transducer construction, determining the sensitivity of the transducer. Martin [90] states that more reliable results are obtained with heavily scratched backplate surfaces. However, disassembly and re-assembly of the same transducer can affect the sensitivity by 10%. The roughness of the backing plate controls the size of the air gap present between the backing plate and the diaphragm. The broad bandwidth of these transducers is achieved by having irregular scratches with random spacing on the backing plate, such as that produced by hand-rubbing with emery cloth. Backing plates with regularly spaced scratches produce sharp resonances in the frequency response, with correspondingly low bandwidth.

The components must be machined to a fraction of a thou for the transducer to work optimally.

6.1.1 Dimensions

A plan of the Kay transducers used in this thesis is shown in Figure 6.3. A photograph may be seen in Figure 7.11. The active surface of the transducer has a diameter of $5/16$ in or ~ 8 mm. The transducer inserts (brass part) have an outer diameter of $1/2$ in. The transducer diaphragm is recessed from the front of the insert by 16 thou.

6.2 Operation of the Kay Transducers

Ultrasonic transducers may use any of several physical phenomena to transform electrical energy into sound energy and vice versa. These include piezo-electricity, magnetostriction, electrostatic,




Figure 6.2: The backing plate of the Kay transducer. The grooves are machined in rings for ease of manufacture. The top of the rails is scratched in a random manner to produce very small grooves. This backing plate is used in one of the transducers of the KASPA system [74, Figure 3], shown in Figure 1.1.




Figure 6.3: The Kay transducer insert used by the sonar system described in this thesis. Dimensions in Inches.




Figure 6.4: The movement of the transducer diaphragm as the voltage across the electrodes is increased. The diagrams show the diaphragm being pulled into the scratches on the backing plate at the top of the rails. The scratches are also visible in Figure 6.2. The grooves do not play a part in the electrostatic operation of the transducer, and serve mainly as an air reservoir (Martin [90]).

electro-magnetic and optical interferometry. As the Kay transducers are of the electrostatic type, we shall concentrate on those here.

When there is no potential difference between the two transducer electrodes, the diaphragm assumes a position shown in Figure 6.4a because it is under tension. Before using the transducer, it must be biased with an offset voltage. This causes the diaphragm to be attracted to the backing plate (Figure 6.4b). As the voltage is increased, the diaphragm will stretch more and enter further into the groove (Figure 6.4c). As the voltage is lowered, the diaphragm will relax (Figure 6.4b). Martin [90] found that the restoring force is supplied mainly by the air pressure behind the diaphragm. The transducer cannot push the mylar outside the straight position - it can only pull the diaphragm toward the backing plate. When the voltage between the plates is removed, the diaphragm will return again to the flat position (Figure 6.4a). In order to use the transducer to generate a sound wave, an oscillating voltage must be applied between the two electrodes. It is desirable to have a linear relationship between the applied voltage and the displacement of the diaphragm to maintain the LTI properties of the system mentioned in Section 4.2. The linear relationship is not maintained when the applied voltage becomes small, so the minimum drive voltage must remain above zero.

The upper limit on the drive voltage is provided by the breakdown voltage of the mylar film. Martin [90] calculated this to be 600 V for 0.15 thou film. However, the stresses around the supporting points will lead to a lower breakdown voltage, so Martin suggests 200 V to 250 V as a working maximum.

The signal processing requires that the signal be passed through the transducer without harmonic distortion. Therefore experiments were conducted to determine the linear region of operation of the transducer. It was found that the linear region is bounded at both low voltages and high voltages. However, the actual voltage applied to the transducer element was not measured, as the oscilloscope connection would have changed the load and distorted the measurement. Both the voltage at the input of the transmit amplifier and the bias voltage applied to the transducer were measured. The single sided amplitude used for MF signals was 0.15 V, while the single sided amplitude used for CTFM signals was 0.1 V. A bias voltage of 150 V was applied in both cases. The gain of the transmit amplifier was not measured.

The signal voltages applied to the transducer in this work will be described in Section 6.6.

6.3 Models of Transducer Operation

There are two main parts to modelling a transducer. The first is to model the conversion from electrical energy into acoustical energy directly in front of the transducer. The second is to model how the acoustic field radiates from the transducer, producing a dynamic pressure field $p(r, \theta)$. This describes the beam-forming of the transducer.

The conversion of electrical to acoustical energy is addressed by a transduction model of the transducer, which operates in the frequency domain. The beam-forming aspect is addressed by two alternative models. The first is the radiating plane piston model (Kinsler [78, Section 8.8]), which operates in the frequency domain. The second is the Kuc-Siegel impulse model [85], which operates in the time domain. The two different beam-forming models are shown to be equivalent through the Fourier transform.

6.3.1 Transduction Model

The transduction model exploits the similarity in the differential equation representation of electrical, mechanical and acoustical systems. This similarity allows us to draw the equations from these three separate domains into a single circuit which describes the operation of the transducer as a whole. For a general description of the transduction model, see Hunt [63]. The process is described in detail in Appendix A.1.

The transduction model demonstrates that the transducer has a characteristic frequency response due to its design, construction, and fundamental operating principles. Of major importance is the bandwidth of the transducer, and its ability to reproduce wide-bandwidth CTFM signals. The frequency response is also influenced by the transmit amplifier circuit, so it is also required to have a wide bandwidth.

6.3.2 Radiating Plane Piston Model

The radiating plane piston model considers a single frequency signal being radiated from a circular aperture embedded in an infinite baffle, and describes the amplitude of the pressure wave at a point $\{r, \theta\}$ in the far field of the transducer. The angle θ is from the transducer normal to the ray from the transducer centre-point to the point at which the pressure is evaluated (Figure 6.5). The full derivation of the model is given in Appendix A.2, with the solution as equation (A.8). The pressure field predicted by (A.8) is plotted in Figure 6.6 for several different frequencies. It is to be noted that the radiating plane piston model incorporates the far field approximation, so the near field shown in Figure 6.6 is unrealistic. However, at ranges $r > a^2/\lambda = a^2 f/c$, the far field approximation is valid, and the radiating plane piston model makes useful predictions about the pressure field. The most important feature is the formation of lobes with nulls in between them. The nulls appear in Figure 6.6 as a line of nodes, where no pressure variations occur. The angles at which the nulls occur become smaller as the driving frequency becomes higher. The main lobe, being directly in front of the transducer, becomes narrower as the frequency increases.




Figure 6.5: The radiating plane piston model. The transducer radius is labelled a . The angle between the transducer normal and the ray to the point $\{r, \theta\}$ is denoted θ . The point $\{r, \theta\}$ is at radius r from the centre of the transducer. The point $\{r, \theta\}$ is constrained, without loss of generality, to lie in the x - z plane.




Figure 6.6: The sound pressure field calculated according to the radiating plane piston model (A.8) at several frequencies. The piston has the same radius $a = 4$ mm as the Kay transducer used in this thesis. The plot axes are x and y Cartesian coordinates, centred at the transducer. The transducer is shown as a red line, while the wavelength of the sound is shown in green. Note the phase reversal between adjacent lobes, and the formation of radial lines without wave information (grey). The angles of these lines depend on the wavelength of the sound and the radius of the transducer. The far field approximation has been used, so only the far field is correct. For correct plots of the near field, see Zhu [147].




Figure 6.7: Beam pattern calculated according to the radiating plane piston model (6.1), for a transducer with radius $a = 4$ mm, being the same size as the Kay transducer used in this thesis. Other parameters $c = 346 \text{ m} \cdot \text{s}^{-1}$, $f = 100 \text{ kHz}$. This produces a transducer radius to wavelength ratio of $a/\lambda = 1.156$. The beam pattern is calculated in decibels, and normalised to 0 dB at the transducer normal. This plot uses the same parameters as are plotted for the sound pressure field in Figure 6.6c.

The lobe structure is caused by the part of (A.8) enclosed in square brackets. The pressure wave and the spreading loss components may be neglected by evaluating only the part in the square brackets, thus obtaining the bearing dependence as given by Kinsler [78]:

$$p_{\theta}(\theta, f) = \begin{cases} \left[\frac{2J_1(k_w a \sin \theta)}{k_w a \sin \theta} \right] & \text{if } 0 < |\theta| \leq \frac{\pi}{2} \\ 1 & \text{if } \theta = 0. \end{cases} \quad (6.1)$$

Here, $J_1(x)$ is a Bessel function of the first kind, and k_w is the inverse wavelength in $\text{rad} \cdot \text{m}^{-1}$ ($k_w = 2\pi/\lambda = 2\pi f/c$). The frequency of the signal is denoted f in Hz, the transducer radius a in m, the the speed of sound c in $\text{m} \cdot \text{s}^{-1}$ and the wavelength λ in m. As the factor

$$k_w a = 2\pi \frac{a}{\lambda}$$

includes wavelength, λ , and the transducer radius, a , the important factor in determining the beam pattern is the ratio a/λ . Thus a 4 mm radius transducer operating at 202.5 kHz will have the same beam pattern as a 16.2 mm radius transducer operating at 50 kHz, as the ratio a/λ is 2.34 in both cases. The beam narrows as the ratio a/λ increases.

The beam pattern (6.1) of the transducer is plotted, for a single frequency signal of 100 kHz, in Figure 6.7. The beam pattern obtained clearly shows the main lobe extending to $\pm 30^\circ$, with secondary lobes out to $\pm 70^\circ$, and tertiary lobes after that. The secondary lobes are 18 dB lower




Figure 6.8: The beam spectrogram calculated according to the radiating plane piston model, for a transducer with a radius of $a = 4$ mm, the same size as the Kay transducers used in this thesis. This function only includes the frequency response as a function of bearing, and assumes that the transducer itself has an infinite bandwidth. The speed of sound is assumed to be $c = 346 \text{ m} \cdot \text{s}^{-1}$. The colour indicates the response in decibels. The linear increase in response as a function of frequency, mentioned in Appendix A.2, is neglected in this plot to emphasise the beam formation. The beam pattern of Figure 6.7 may be read from this graph by tracing an arc at a radius of 100 kHz. Another view of the beam spectrogram may be seen in Figure 6.14.

than the main lobe, and do not change their level with either wavelength or transducer radius according to the radiating plane piston model.

Although this model only explicitly considers the case of a signal containing a single frequency, the wave equation permits any linear combination of solutions to be added together to obtain a new solution. Thus broad-band signals may be considered by evaluating the model at various single frequencies in parallel, and summing the results. Thus the radiating plane piston model provides the transducer beam pattern in the frequency domain. A plot showing the response as a function of bearing θ and frequency f is shown in Figure 6.8. This plot shows the response as a function of bearing and frequency. The bearing is represented by the angle on the polar plot, while the frequency is indicated by the radius, as marked. Thus the plot of Figure 6.7 is equivalent to an arc at a radius corresponding to 100 kHz on this plot. The increase in beam power shown in Figure 6.7 is seen instead as a change in colour in Figure 6.8.

Figure 6.8 shows that the beam width is quite large at low frequencies, having a wavelength much longer than the dimensions of the transducer. At this wavelength, the transducer acts like a baffled simple source. As the frequency increases, so does the radius to wavelength ratio a/λ , and hence the beam pattern changes as well. The first null appears when the frequency reaches 52 kHz, having a radius to wavelength ratio of $a/\lambda = 0.6$. The second null appears at a frequency of 97 kHz, having a radius to wavelength ratio of $a/\lambda = 1.12$. At this frequency, the first null occurs at $\pm 32^\circ$, and the width of the main lobe has shrunk considerably. The width of the main

lobe continues to shrink as the frequency, and hence the radius to wavelength ratio, continues to increase. Thus the beam becomes more and more focused as the radius to wavelength ratio increases.

Equation (A.8) includes a term $ck_w = 2\pi f$, indicating that the power of the pressure wave will depend upon f^2 . This indicates that the radiated power will be constant if the velocity amplitude, U_0 , of the transducer surface, is constant. As the width of the beam becomes smaller, the power radiated into that area must increase in order to keep the total power radiated constant. This has not been taken into account in producing Figure 6.8.

Modified Annular Radiating Plane Piston Model

Anke [5] derived a modified radiating plane piston model. He assumed that only the regions of the transducer diaphragm lying atop the rails, shown in Figure 6.2, actually move. Thus the moving regions are a set of concentric annuli. However, the angular response derived using this model is nearly identical to the radiating plane piston model. Measurements of the movement of the transducer surface by Martin [90] do not indicate that the annular regions over the grooves stay still, as assumed by Anke's model. As the radiating plane piston model is simpler, that model was adopted instead.

6.3.3 Kuc-Siegel Impulse Model

Kuc and Siegel [85] provide a time domain derivation of the transducer beam pattern. A modified version of their derivation is reproduced in Appendix A.3. The final result is the time domain response of the transducer to an impulse function, working either as a transmitter or as a receiver. The response can be transformed from the time domain to the frequency domain, where it may be compared with the result obtained by the radiating plane piston model. The angular dependency produced by both methods is the same.

The modification to the Kuc-Siegel impulse model was to remove a $\cos \theta$ term, as shown in Appendix A.3. The $\cos \theta$ term is appropriate where the transducer actuates air velocity, as it selects the rectilinear motion component. It is inappropriate to use it where the transducer is actuating air pressure, which is non-directional in a fluid (Halliday [51]). Removal of the $\cos \theta$ term from the Kuc-Siegel [85] impulse model reconciles it with the standard radiating plane piston model [78] when the two are related through the Fourier transform. Furthermore, the removal of the $\cos \theta$ term provides a much improved match with the experimental results for angular dependency shown in Sections 6.7 and 6.8.

Both the radiating plane piston model and the impulse model provide the beam-forming response of the transducer, and are shown to be equivalent through the Fourier transform, as demonstrated in Appendix A.3.

6.4 Total Effect of a Transducer

We have now examined the transduction model and two beam-forming models, being the radiating plane piston model and the impulse model. It can be seen that the total effect of sending a signal

through a transducer can be described using an impulse response model with three parts.

1. The operation of the transducer itself, which is modelled by the transduction model. While the transduction model operates primarily in the frequency domain, its time domain impulse response, $h_f(t)$, may be found through the inverse Fourier transform.
2. The sound radiating and interfering at a particular angle, which is equivalently modelled by either the radiating plane piston model or the Kuc-Siegel impulse response model. We denote its impulse response by $h_\theta(\theta, t)$, with the spreading loss and propagation delay component of the models removed.
3. The spreading loss and propagation delay, with an impulse response denoted by $h_{rad}(r, t)$. This component forms part of both the radiating plane piston model and the Kuc-Siegel impulse model, but is drawn out separately here.

Thus the effects of the signal as it is transformed from electrical energy to a pressure wave at a point $\{r, \theta\}$ may be modelled by

$$h(r, \theta, t) = h_f(t) * h_\theta(\theta, t) * h_{rad}(r, t). \quad (6.2)$$

The first two parts of this impulse response may be attributed to the transducer. The third part is common to any propagating spherical wave, as found in the far field of the transducer. Thus $h_{rad}(r, t)$ is the same as that defined in Section 4.2.3, while the terms $h_{trans}(\theta_T, t)$ and $h_{rec}(\theta_R, t)$, representing the response of the transmitter and receiver, are individually equivalent to the convoluted product $h_f(t) * h_\theta(\theta, t)$ in (6.2).

6.5 Beam Directivity

The beam spectrogram shown in Figure 6.8 shows that the shape of the beam changes with frequency. At low frequencies, the beam is widely spread, while at high frequencies the beam is narrower. Therefore the transmitted sound energy must be spread over a larger area at low frequencies than at higher frequencies. This is demanded by conservation of energy.

A measure of the amount of beam focusing is the directivity (Kinsler [78, Section 8.9]), which is found by integrating the radiating plane piston model beam pattern (6.1) over the spherical solid angle [64, Section 17.7-8]:

$$D(f) = \frac{4\pi}{\int_{\theta=0}^{\theta=\pi} \int_{\phi=0}^{\phi=2\pi} |p_\theta(\theta, f)|^2 \sin \theta d\phi d\theta} \quad (6.3)$$

The beam function $p_\theta(\theta, f)$ is rotationally invariant around the transducer's principal axis, so the integration over ϕ may be separated and performed immediately:

$$D(f) = \frac{2}{\int_{\theta=0}^{\theta=\pi} |p_\theta(\theta, f)|^2 \cos \theta d\theta} \quad (6.4)$$

The remaining integration in θ is the angle from the transducer's principal axis. Performing the integration numerically and plotting as a function of frequency yields Figure 6.9. The high fre-




Figure 6.9: Beam directivity, $D(f)$, calculated for transducers of three different radii (marked a) using (6.3), as a function of frequency. (The value plotted is called the directivity index as it has been converted to decibels.) Larger values of $D(f)$ indicate that the beam energy is concentrated into a smaller area. The low frequencies, having a wide beam width in Figure 6.8, are the least concentrated. The beam is more focused at higher frequencies, so the directivity is higher there. Furthermore, larger transducers have higher directivity than small transducers.

quencies are more tightly focused by the beam-forming function than the low frequencies.

Figure 6.9 can also be regarded as a beam-focusing gain function. Microphone measurements made on-axis must be corrected for this directivity factor.

The transducer directivity function, $D(f)$, will be used when the experimental transducer response is determined in Sections 6.7 and 6.8.

6.6 Circuits

The circuits which drive the transmitter and amplify the signal from the receiver are built into the sonar head. The design of the sonar head will be described in Chapter 7.

As discussed in Section 6.2, the transducers require a bias voltage to be applied before they can be used as either a transmitter or a receiver. The bias voltage was originally self-generated by the transmit circuit, which assumed that only CTFM signals would be used. However, we required an arrangement which would also support the use of Barker coded signals, which are much shorter. The short signal produces an unsatisfactory bias voltage, so the transmit circuit was modified to use an external bias. The bias voltage is also used by the receive amplifier circuit, which is described below. Furthermore, the circuits were modified to reduce the electrical crosstalk from the transmitter to the receiver, but this will be discussed further in Section 7.3.1.

6.6.1 Transmit Circuit

The bias voltage in the original circuit was produced as a by-product of the amplified drive signal by a clever design (Figure 6.10). The bias voltage, being the wire connected to the bottom of R16,

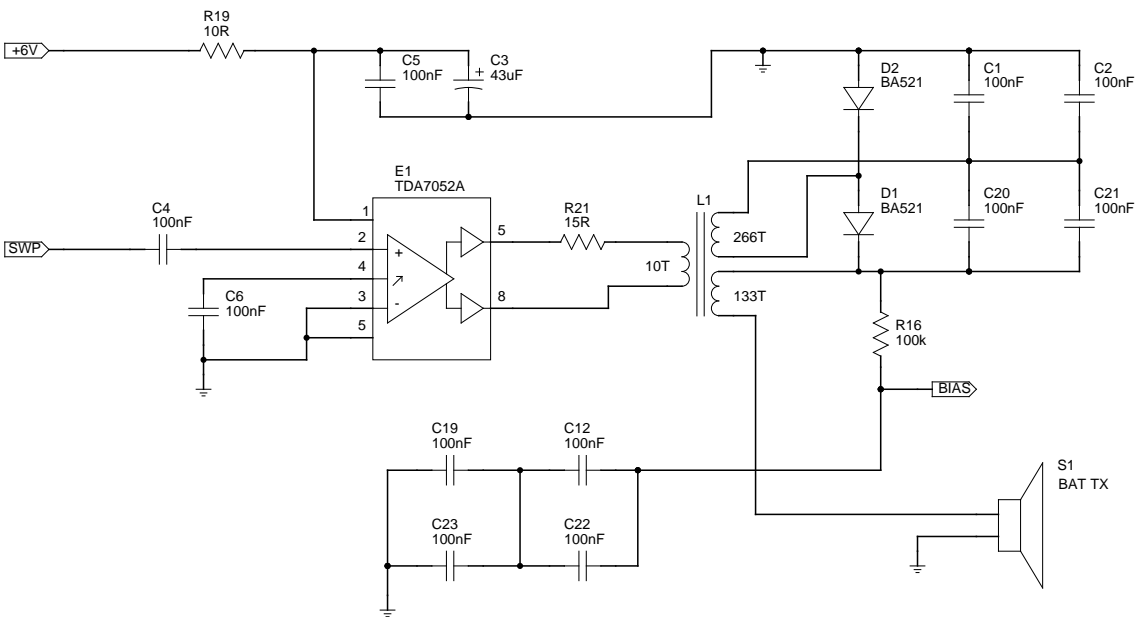


Figure 6.10: The original transmit circuit built into the Kay sonar head. The bias voltage is also connected to the receive circuit, which is shown in Figure 6.11.

Table 6.1: The measured self-generated bias using the original transmit circuit from Figure 6.10. The amplitude of input signal column refers to the single sided amplitude of the voltage presented at SWP in Figure 6.10.

Amplitude of Input Signal Volts	DC Bias Volts
0.1	169.5
0.05	90.0
0.025	51.0

is generated by rectifying the amplified transmitted signal. The amplified transmit signal is then added to the bias signal.

Values obtained for the bias voltage, as a function of the amplitude of the transmit signal at the amplifier input, are shown in Table 6.1. Experiments were performed to find the optimal input voltage for the transmit circuit. It was found that single-sided input voltages over 0.125 V produced clipping of the waveform at the 10T input to the transformer L1. The maximum swing voltage was 8 V, due to the double-sided drive employed by the TDA7052A integrated circuit.

6.6.2 Receive Circuit

The receive amplifier circuit, shown in Figure 6.11, is used for each of the four receivers in the sonar head. It uses the bias voltage generated by the transmitter section, which is applied to the transducer for proper operation. The AC component of the transducer voltage, corresponding to the received signal, is amplified and made available at the rec out connection.

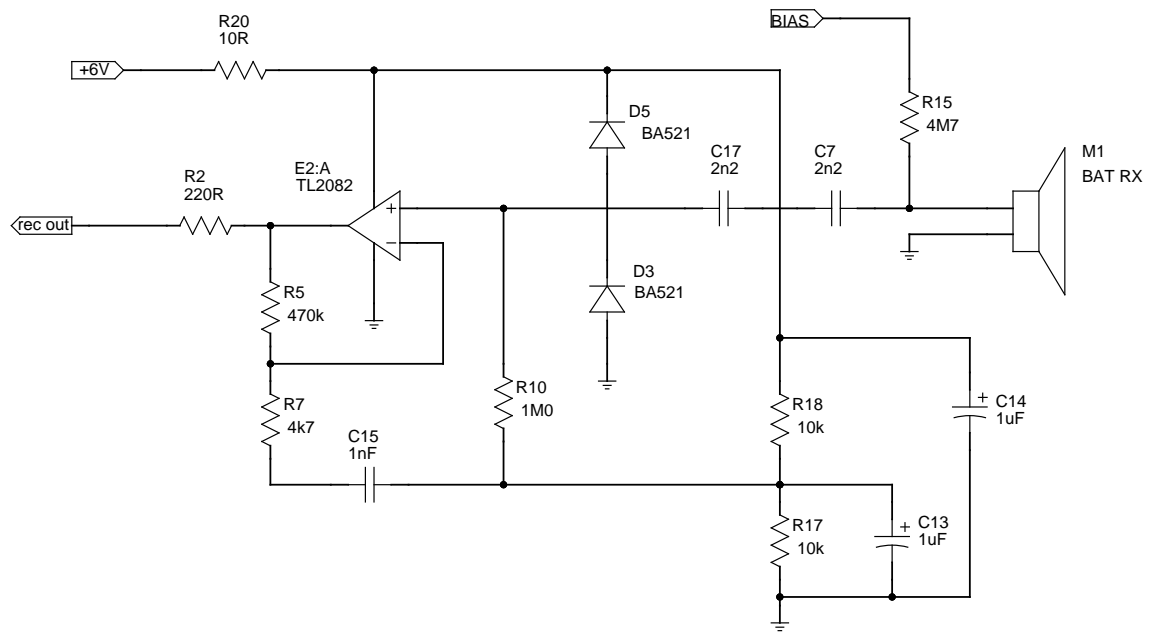


Figure 6.11: The receive circuit built into the Kay sonar head. This circuit is used for each of the four receivers. The BIAS pin is connected to the pin of the same name in the transmit circuit of Figure 6.10.

A modification was made to the receive circuit to reduce crosstalk from the transmit circuit. The modification consisted of installing shielded cable between the receiver and the circuit board. This modification can be seen in Figure 7.15.

6.6.3 Transmit Circuit with External Bias

The transmit circuit was changed to use an external bias voltage, derived from a high voltage power supply (Section 5.5.1). The circuit was modified as shown in Figure 6.12. The primary modification was to remove the connections to the 266T coil of the L1 transformer, and to connect an external bias voltage at the point labelled EXTB. The bias voltage was supplied using the Delta Elektronika power supply described in Section 5.5.1. This modification also reduced harmonics which were present when 50kHz signals were produced using the self-generated bias. This is thought to be caused by inadequate bias voltage being self-generated by low frequency signals, which subsequently caused the diaphragm to bottom out when a high amplitude transmit signal was applied. These problems were avoided by using the externally generated bias.

The transmit circuit was also modified to reduce crosstalk. Two modifications were performed. The first was to place a brass shield around the high voltage circuits to reduce electro-static coupling. This shield will not reduce electro-magnetic coupling. The second modification was to replace both of the wires to the transmitter with shielded cables to prevent electrostatic radiation. These modifications are visible in Figure 7.15.

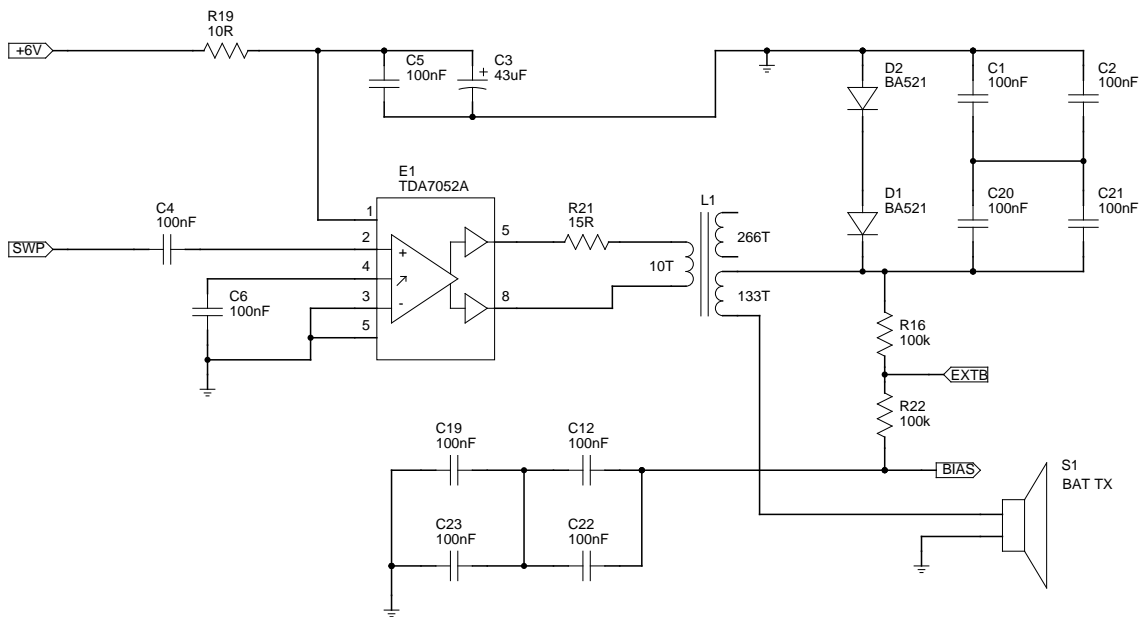


Figure 6.12: The transmit circuit, modified to use an externally supplied bias voltage, connected at EXTB.

6.7 Experimental Kay Transducer Response

Section 6.4 demonstrated that the transducer can be expected to have an inherent frequency response, as well as a separate beam-forming response. The inherent frequency response of the transducer has too many variables to be estimated or calculated (see Appendix A.1).

The transducer response was measured using the Brüel & Kjær 1/4 inch microphone described in Section 5.5.2. The front cover was removed from the microphone to improve the high frequency response. The microphone was placed 0.5m in front of the sonar head, and a frequency sweep from 10kHz to 150kHz was transmitted. The sound recorded by the receiver was digitised, and the Fourier transform was taken. The reconstruction filter mentioned in Section 5.1.1 was removed for this experiment.

The experiment described above will involve some of the elements of the signal path described in Chapter 4. These are the transducer response described by transduction, the air absorption of Section 4.3.3, and the directivity gain described by (6.3). A proper assessment of the transducer characteristics requires that these be compensated. Such a compensation is shown in Figure 6.13. The experimental response shows a peak at 85 kHz.

The directivity curve slopes gently up, at 12 dB per 100kHz. The gain is only moderate due to the wide beam width of the small transducer. The transduction response is obtained by subtracting out the beam-forming gain in decibels, or, equivalently, dividing by the beam-forming gain in a non-decibel domain. The transduction response displays the same peak at 85 kHz as the on-axis response. This appears to be the response of a resonant circuit in the transmit amplifier circuit.

A complete determination of the transduction response itself would require that the spectrum output by the transmit amplifier, shown in Figure 6.12, also be subtracted from the measured




Figure 6.13: The frequency response of a Kay transducer (Head 1) being used in transmit mode. The curve labelled experimental on-axis is measured using the microphone. The curve labelled directivity is the on-axis gain calculated in Section 6.5. The air absorption term is calculated according to Appendix E.4. The curve labelled transduction is calculated by adding the air absorption and subtracting the directivity from the experimentally measured response, and represents the spectrum of the signal after the transduction process, but before beam-forming and propagation. Thus this curve represents the acoustic field at the surface of the transducer. The zero level of the decibel scale is arbitrary for the experimental and transduction curves, but is normalised for zero frequency on the beam-forming curve.




Figure 6.14: Experimentally determined Kay transducer transmit response including both on-axis frequency response factor and the bearing frequency response factor. (This differs from Figure 6.8 which does not include the on-axis frequency response factor.) This graph shows the normalised signal amplitude as a function of frequency f and the bearing θ of the microphone from the principal direction of the transmitter. The colours indicate the signal power in decibels, and match those used in Figure 6.15.

response. However, measurement of the amplifier response, with an attached transducer, is a difficult task, due to the perturbation of the circuit parameters by the instrument probe. No such measurement has been performed.

Beam Spectrograms

The beam spectrogram was measured by rotating the sonar head about the centre-point of the transmitter in increments of 1° . At each position the frequency response was measured using the procedure previously described for the on-axis response. The software described in Section 5.4.1 performs this procedure automatically. Taking the Fourier transform of the received echo, and plotting as a function of frequency and bearing yields the visualisation shown in Figure 6.14.

The same experimental data is shown in Figure 6.15a, using a polar representation and using colour to represent the signal strength. When the experimental beam spectrogram is compared to the radiating plane piston model shown in Figure 6.8, it is seen that there is a good match in the bearing dependency, but a poor match in frequency. A more realistic hybrid model may be constructed by using the experimentally determined on-axis frequency response $H_f(\omega)$ in conjunction with the modelled angular response $H_\theta(\theta, \omega)$. The result of this calculation is shown

in Figure 6.15b. The hybrid model of Figure 6.15b is seen to be a better match for the experimental beam spectrogram shown in Figure 6.15a than the radiating plane piston model shown in Figure 6.8. The null patterns predicted by the radiating plane piston model are present in the experimental response, although the nulls are not as sharply pronounced. The main beam and side lobes are quite similar to the composite model. The experimental response has a large amount of high frequency response at large angles, which is not predicted by the composite model. The differences between the experimental response and the composite model may be due to non-uniform movement of the diaphragm surface. Martin [90] showed that the movement varies considerably across the surface of an individual transducer.

6.8 Response of Polaroid Transducers


This thesis used the small hand-made Kay transducers for many reasons, some of which will be addressed in Chapter 7. The radiating plane piston model shows that a small transducer has a much wider beam width than a large transducer at the same operating frequency.

A common alternative approach to custom designed and built units is the use of cheap commercial alternatives. One of these is the Polaroid transducer. Since the heritage of these units stems from those designed by Kay there are no fundamental design or operational differences. Rather the major differences pertain to the application and the cost of this component in the overall item cost. As the Polaroid transducers are commonly used in robotics, their properties were measured for comparison.


The on-axis frequency response of the Polaroid transducer was measured using the same method as for the Kay transducer, as described in Section 6.7. A custom-built transmit amplifier, with a relatively flat response over the ultrasonic frequency range, was used to drive the transducer. The transducer response was measured in two configurations. The first was as supplied by the manufacturer, as shown in Figure 2.5, with its aluminium grille in place. The second configuration was with the grille removed. More precisely, a second transducer was dismantled, and the central part of the outer casing was cut out, and the transducer was then re-assembled. The reconstruction filter mentioned in Section 5.1.1 was removed for this experiment. The results for the on-axis frequency response experiment are shown in Figure 6.16.

The results of the experiment show that the cover significantly affects the frequency response, causing an increased frequency response around 60kHz and 105kHz, and a reduced response in between. The increased response may be due to single and double resonance within the cavity between the diaphragm and the grille, as suggested by Martin's transduction model in Appendix A.1. The reduced response may be due to destructive interference between the direct and reflected waves within the cavity. Thus, depending upon the working frequency and the application, it may be beneficial to remove the cover, obtaining a flatter on-axis response. However, for the original application, using a single frequency 50kHz sine-wave, the grille provides a +5dB gain in the output.

The transduction response was calculated in the same manner as for the Kay transducer. Comparison of the transduction curves shows that the Polaroid transducer has less efficient transduction than the Kay transducer. Furthermore, we will see in Section 7.2.1 that the Polaroid transducer has




(a) Experimental angular and frequency response




(b) Modelled angular response and experimental frequency response

Figure 6.15: Beam spectrograms for a 4 mm radius Kay transducer used in transmit mode (Head 1). (a) Experimental data for both angular and frequency components. This is the same data set as is plotted in Figure 6.14. (b) Hybrid model, using the experimental on-axis frequency response combined with the theoretical radiating plane piston model to obtain the angular response. The product $H_f(\omega) \cdot H_\theta(\theta, \omega)$ is shown. The power level in these spectrograms has been normalised so that the maximum is at 0 dB.



(a) with grille



(b) grille removed

Figure 6.16: The measured on-axis frequency response of a Polaroid transducer. The power scales are the same as shown in Figure 6.13, so the responses may be compared. The curve labelled experimental on-axis is measured using the microphone. The curve labelled directivity is the beam-forming gain as calculated in Section 6.5, using a radius of 16.2 mm. Note the directivity is different from the Kay transducer shown in Figure 6.13 due to the larger transducer radius. The curve labelled air absorption is calculated according to Appendix E.4. The curve labelled transduction is calculated by adding the absorption and subtracting the directivity from the experimentally measured response, and represents the spectrum of the signal after the transduction process, but before beam-forming and propagation. Thus this curve represents the acoustic field at the surface of the transducer. These measurements were taken using a 150 V bias.

a 22dB gain over the Kay transducer, due to increased surface area. This increased surface area generates additional gains due to improved beam focusing. However these gains over the smaller Kay units are lost due to poor transduction. This poor transduction is attributed to the Polaroid backing plate with its machined v-grooves, instead of the rails used in Kay's design, resulting in a smaller surface area for electrostatic interaction.

Other experimental determinations of the Polaroid frequency response, such as Peremans [108], show a peaked response, in contrast to the comparatively flat responses seen in Figure 6.16. The peaked response is attributed to the standard amplifier supplied by Polaroid, which has a tuned circuit. As these measurements were made using an amplifier with a flat response, being a separate circuit from Figure 6.12, the power spectra observed in Figure 6.16 are attributed to the operation of the transducer.

Beam Spectrograms

The beam spectrogram experiment, which was described for the Kay transducers, was repeated for the Polaroid transducer. In this case the hybrid beam spectrogram, shown in Figure 6.15b, was not computed. The Polaroid transducer has a radius of $a = 18.25$ mm. When the beam spectrogram (6.1) is calculated using this radius, a poor match is obtained. However, when the beam pattern is evaluated using a transducer radius of $a = 16.2$ mm, a good match is found. M^cKerrow [95, Chapter 10] indicates that this is due to a difference between the radius of the transducer and the radius of the active area.


The beam spectrogram was measured using the two configurations previously described for the frequency response measurement. The results are plotted in Figure 6.17.

Comparing the two figures, the grille does not seem to significantly affect the shape of the beam. This could be due to the dismantling of the transducer. The main effect of the grille is a narrowing of the main beam around 90kHz.


When these beam spectrograms of the Polaroid transducer in Figure 6.17 are compared with the ones from the Kay transducer in Figure 6.15, we see that the beam is much narrower overall than for the Kay transducer. This is due to the increased size of the Polaroid transducer, which causes the beam to be more focused overall. The width of the main lobe at 60kHz at the -10 dB point is about $\pm 10^\circ$ for the Polaroid transducer, as compared to the Kay transducer which is $\pm 30^\circ$ when it is -10 dB down from its value at 0° . Thus the Kay transducer has a much wider beam width than the Polaroid transducer, and can insonify and receive echoes from a much larger area.

6.9 Conclusion

The ability of the transducer to support wide bandwidth signals, such as those used by CTFM, is of critical importance to the operation of a sonar system. Both the Polaroid transducer and the Kay transducer have the bandwidth to support CTFM and other wide bandwidth applications. The Polaroid transducer should have a strong power response, due to its large surface area and tight beam-forming capability. However, its additional potential power over the Kay transducer is lost due to poor transduction.



(a) with grille



(b) grille removed

Figure 6.17: The measured beam spectrogram of two different Polaroid transducers of the same model. In (a), the original grille of the transducer was retained. In (b), the grille was removed from the transducer. The power levels in these spectrograms has been normalised so that the maximum is at 0 dB.

The Kay 4mm transducer may be best utilised with a wide bandwidth signal in the range 50kHz to 110kHz. This harnesses the peak power of the transducer and provides a wide beam width. Inspection of the beam spectrogram in Figure 6.15a shows that any other choice would limit the transmitted power and beam width. The Polaroid transducer may be best used over the range 40kHz to 130kHz, and obtains a flatter response with the grille removed. However, inspection of the beam spectrogram in Figure 6.17b reveals that the beam is much narrower than that of the Kay transducer, and especially narrow at high frequencies, as determined by (6.1) and reflected by the directivity shown in Figure 6.9. This must be taken into account when designing an ultrasonic sensing system.

This chapter will avoid drawing conclusions about which transducer is best for this study, instead leaving that task to the next chapter where all of the design issues are considered together.

This chapter has demonstrated that the transducer response may be represented as a sequence of three filters. The first filter corresponds to the transduction process of converting the electrical energy into sound pressure wave energy at the transducer surface. The second filter corresponds to the formation of the beam. The third filter corresponds to the delay while the sound travels through the air, and also incorporates the spreading loss. When viewed in the frequency domain, these effects are multiplicative, which serves as a convenient representation. We attribute the first two effects to the transducer, and the third to the propagation of sound through the air.

6.9.1 Acknowledgements and Contributions

This chapter has drawn together pre-existing information from the literature and combines it with some original contributions.

1. The description of the construction of the Kay transducers was taken from Martin [90].
2. The description of the operation of the Kay transducers was constructed from information in physics textbooks [51, 47], as well as information in Martin [90] and Hunt [63].
3. The radiating plane piston model material was taken from Kinsler and Frey [78].
4. The Kuc-Siegel impulse model was taken from Kuc and Siegel [85].
5. The modification of the Kuc-Siegel impulse model and the comparison with the radiating plane piston model is the work of this author.
6. The plotting of the beam pattern as a function of frequency and bearing as a beam spectrogram (e.g. Figure 6.8) is claimed as the invention of this author. The code to perform this plotting was custom-developed in Mathematica [139].
7. The application of directivity to in-air transducer measurements is a contribution of this author.
8. The circuits used to drive the transmit transducers and the amplifiers after the receivers were a standard design produced by Kay, and were provided with the sonar head.

-
9. Modifications to the circuits which drive the transmit transducers and the amplifiers after the receivers to reduce cross-talk and to make the transmit circuit suitable for short-time signals are claimed as the work of this author.
 10. The measurement of the transduction response of the Kay and Polaroid transducers, and the associated comparison with theoretical results from the radiating plane piston model and the Kuc-Siegel impulse model, is the work of this author.

Chapter 7

Sensor Head Design

This chapter addresses the integration of a set of transducers into an operational sonar head which is capable of addressing the goals of Section 1.2. The chapter first reviews some tasks of the *robotic agent*, and determines the sensing requirements to perform the tasks. Subsequently, the chapter reviews the interaction between the size of the transducers and the placement of the transducers. The impact of these decisions upon the operation of the sonar system and its ability to meet the requirements is then discussed and design decisions are made. Finally, the construction of the sonar head is described.

7.1 Requirements of the *Robotic Agent*

This section considers some of the tasks that the *robotic agent* will be required to perform, and from this derives a specification for the sensor system to perform these tasks. Four aspects of the sensory problem are considered. The first is the precision and resolution required to perform adequate sensing. The second and third aspects consider the horizontal and vertical field of view, or area of insonification, over which the sensor must operate to obtain adequate performance. The fourth and final aspect considers the required update rate to provide adequate performance of the *robotic agent* when performing sensing based tasks.

7.1.1 Precision and Resolution

The requirements for precision and resolution vary depending upon the application. The following discussion considers the problem of a *robotic agent* navigating through a doorway. Precise determination of both range and bearing are required to solve the problem. We assume that the sensor is accurate and without bias.

Budenske [15] found that it was difficult for a *robotic agent* equipped with a sonar ring to navigate through a doorway. This may be primarily attributed to the poor bearing precision provided by typical sonar rings, which only determine the bearing inasmuch as it lies within the beam of the ultrasonic transducer. Kleeman [80] used a scanned monaural sonar system which measured the amplitude of the echo and fitted a Gaussian function to the observed amplitude data to solve the problem. However, we consider here the requirements of a sonar system which can provide sufficient information to solve the problem without scanning.

When the *robotic agent* is approaching the doorway along a centre line, both of the door posts will appear to be at about the same range. Sonar sensors are typically centrally mounted on a

robotic agent, therefore placing the sonar sensor equidistant from both the left and right door posts. While the door posts are well separated in bearing, they will have very similar ranges. This is a problem for a sonar system with a wide beam width, as both door posts will return echoes to the same sensors, and they will probably be unresolvable in this configuration. The echoes must be resolvable before the bearing can be computed. Thus a system with good range precision and resolution is an absolute requirement.

There are three ways to solve this problem. The first is to ensure that the sensor has good range resolution. The second is to program the *robotic agent* to approach the door on a skewed angle, so that the echoes from the left and right door posts may be easily resolved. If we assume that echoes from targets separated by ± 10 mm in range may be resolved, then a skew angle of $\pm 4^\circ$ will ensure that the door posts may be resolved until the sensor is 0.15 m from the centre of a 0.8 m wide door. The third alternative is for the *robotic agent* to pass through the door along an off-centre path, or to mount the sensor off-centre on the robot. All of these methods should allow a *robotic agent* with a binaural bearing sensor of sufficient beam width to pass through a doorway without scanning.

Navigating through a doorway can be a tricky operation if the *robotic agent* is nearly the same size as the doorway. The sensor must observe both the left and right door posts simultaneously to ensure that no collision occurs. In the case of the Titan wheelchair robot used at the University of Wollongong [117], there is a clearance of 50 mm on each side. This means that the sensor must be able to measure the lateral position of the door posts with a precision of ± 10 mm. In the case of the Labmate robot, which is almost the same width as the doorway, a precision of ± 1 mm is demanded.

When the *robotic agent* is lining up its approach to the doorway, for example at a range of 0.5 m, the required lateral precision of ± 10 mm translates to a bearing precision of $\pm 1.1^\circ$. As the *robotic agent* approaches and enters the doorway, the range precision becomes more important. In this configuration, the required lateral precision translates to a range precision directly.

The doorway problem also requires a sensor with a wide beam width if we are to avoid scanning. A sensor with a beam width of $\pm 30^\circ$ or 60° allows both of the door posts to be seen until the sensor is 0.7 m from the doorway. Therefore scanning or two sensors are required.

The problem of navigation within a large room requires that the sensor have a large maximum range. We desire a maximum range of 5 m.

From this discussion we can see that it is desirable to achieve a range precision of ± 1 mm, a resolution of ± 10 mm, a maximum range of 5 m, and a bearing precision of $\pm 1^\circ$.

7.1.2 Horizontal Beam Width

The *robotic agent* needs to be able to observe the area in front of it for obstacle avoidance, and the area to its sides for landmark navigation. A wide sensing area is desirable. With a narrow beam sensor, the sensor must be scanned, which slows down the information gathering rate of a sensor.

The sensor can only determine the bearing of targets for which it detects echoes on both receivers. We shall define this region to be the field of binaural audition.

A *robotic agent* which is 0.5 m wide will be able to see its full width at a range of 0.73 m without scanning if it has a sensor with a $\pm 20^\circ$ field of binaural audition. A sensor with a $\pm 30^\circ$




Figure 7.1: Plan view geometry of the sensor system mounted on a mobile robot, showing the extent of horizontal sensing extended by panning.

field of binaural audition will sense the full width of the *robotic agent* at 0.5 m. Therefore, unlike sensors which cannot measure the target bearing, we desire a beam width as wide as possible, reducing the need to scan and increasing the information rate.

Wide Robots

The Labmate robot is 0.8 m wide. A sensor with a $\pm 30^\circ$ field of binaural audition does not satisfactorily cover the width of such a *robotic agent*, so a sensor configuration must be designed. This is determined by the desired horizontal field of audition around the robot. As *robotic agents* have physical area, more than one sensor is usually required. One possible sensor arrangement to achieve coverage of the area in front of and beside the robot is shown in Figure 7.1. This figure shows the sensors mounted upon pan and tilt units at the front corners of a Labmate robot. While this decision is a trade off between cost and sensing time, it has three benefits. Firstly, it minimises the chance of crosstalk between the transmitters. Secondly, it enables the sensors to work together to detect wedge shaped obstacles directly in front of the robot. Thirdly, as the sensors can be panned to detect each other's transmitted pulse, automatic calibration is possible on a regular basis. Because of the large physical separation of the sensors, they can obtain different views of the same nearby object. For all the applications mentioned earlier, the ultrasonic sensing system must be able to detect objects in front of the robot and on both sides. Thus a horizontal field of audition greater than 180° is required. In fact, the axes of the sensors must be able to point at walls to either side. As shown in Figure 7.1, this field of audition can be achieved with a pan of less than 180° .

7.1.3 Vertical Beam Width

A typical environment for a mobile *robotic agent* is shown in Figure 7.2. McKerrow [96] categorises the objects in the environment according to height, as:

- being low enough for the robot to drive over (electrical cables, carpet & tile edges, door sills),
- occupying the same height region as the robot (plant, chair, wall), and




Figure 7.2: Obstacles and landmarks in the environment of a *robotic agent*. Note that some objects occupy the same vertical space as the robot, and therefore form a navigational hazard. Other objects are higher or lower than the extent of the robot, allowing the robot to pass over or underneath them.

- being high enough to drive underneath (table, doorway, ceiling).

Such a classification scheme depends upon the height of the *robotic agent* itself. Two of the classifications indicate that the object is not hazardous to robotic navigation. Thus it is considered an advantage if the sensor used by the *robotic agent* can determine the height or elevation angle of an obstacle in addition to its bearing.

Having a wide beam width allows the sensor to observe close objects which lie away from the principal axis of the sensor, but still within the vertical height of the robot, for example those on the floor in front of the robot. However, having a wider beam will make it more difficult to distinguish separate targets, as the echo signal can become quite cluttered. As shown in Figure 7.3, a $\pm 30^\circ$ field of binaural audition enables the robot to view the floor one metre in front of it and the ceiling 3100 mm in front of it when the sensors are located 500 mm above the floor. This provides the right balance between target visibility and excess clutter for this work.

7.1.4 Update Time

The operation of a *robotic agent* performing a mapping or navigation task within a static environment is limited by the speed with which it can accomplish sensing tasks (Heale [60]). Heale also states that the time taken by a *robotic agent* to perform many tasks is determined by the time taken to accomplish the necessary sensing. Furthermore, sensors which have a narrow field of view must be scanned, adding to the time taken. Therefore the sensor must complete many measurements per second to obtain adequate mapping and navigation performance.




Figure 7.3: The vertical sensing area of a sensor with a $\pm 30^\circ$ beam. Such a sensor can observe the wall-ceiling intersection at about 3 m, and observe the floor in front of the robot at 0.87 m.

A more difficult problem is faced by a *robotic agent* operating within a dynamic environment. Such a *robotic agent* must be capable of detecting moving obstacles such as people or other *robotic agents*. The *robotic agent* must be able to detect these moving obstacles in time to avoid a collision either by stopping or by changing its path. In such an environment, rapid update is required of the sensor. To obtain smooth sensing, update is required at rates similar to television frame rates (25 updates per second).

7.2 Transducer Selection and Arrangement

The width of the binaural field of audition depends upon the size of the transducers and the splay angles of the receivers. The size of the transducers determines the beam width and the beam power, while the splay angle of the receivers determines the overlap of the two beams and the calculation of bearing by [Inter-aural Power Difference \(IPD\)](#). These parameters are now determined to achieve the desired binaural field of audition.

7.2.1 Transducer Size

The power radiated by the transducer depends upon the surface area, as determined by the radiating plane piston model ([A.8](#)). A plot of the relationship is shown in [Figure 7.4](#). This plot demonstrates that, as long as the transduction remains the same (and hence the speed amplitude term U_0 in [\(A.8\)](#) remains constant), a much more powerful signal may be obtained by using a larger transducer.

The effect of transducer size upon beam width is demonstrated in [Figure 7.5](#). As the radius of the transducer increases, the beam becomes more focused and narrow. The beam patterns in




Figure 7.4: Plot showing the effect of the transducer size (radius, in mm) upon the beam power, determined from (A.8). The beam power has been normalised against the power emitted by a 4 mm radius transducer. This plot demonstrates the substantial benefits of using a larger transducer.




Figure 7.5: Diagram showing the effect of the transducer radius upon the beam width. The red line is for a transducer with a radius of 4 mm, the green line for an 8 mm radius, and the blue line is for a 16 mm radius, all operating at 75 kHz in air, with an assumed speed of sound of $346 \text{ m} \cdot \text{s}^{-1}$. The radius to wavelength ratios are 0.87, 1.7, and 3.5 respectively. The beam patterns were evaluated using the radiating plane piston model (6.1), and were normalised to 0 dB at 0° . This plot demonstrates the benefits of having a small transducer, as it produces a wide beam.




Figure 7.6: Responses of several different sized transducers (black curves) to a CTFM signal, swept from 100 kHz down to 50 kHz. These curves were generated by integrating the beam forming response over these frequencies. After integration, the curves were normalised by dividing by the signal bandwidth. The effect of the transducer frequency response was neglected. The effect of the side lobes are greatly reduced by using a wide bandwidth signal and a matched filter or CTFM demodulation process. The grey curves are Gaussian functions fitted to the central part of the integrated data. The b parameters denote the -8.69 dB width of the Gaussian functions.

Figure 7.5 are normalised to 0 dB on axis. If we were to normalise the beam levels to maintain a constant total beam power as the beam width is varied, then the central level for narrow beams would be much higher than that of wider beams. This is the effect of the directivity discussed in Section 6.5. Thus the directivity increases as the transducer is made larger, as can be seen in Figure 6.9.

The use of a CTFM or other broadband signal will also introduce a variation in the beam width as the signal sweeps. This may be seen in Figure 6.15a. A high frequency produces a narrower and more focused beam than a low frequency. Thus an object at one edge of the beam pattern will return an echo containing more low frequencies than high frequencies. The effective beam width of a CTFM signal may be found by integrating the beam function over the frequencies used in the sonar signal (Figure 7.6). The CTFM system effectively performs the integration in the spectral analysis stage. The integration tends to remove the nulls in the beam pattern.

It is convenient for the work in the next section on the orientation of the receivers if we fit a Gaussian function to the integrated beam width shown in Figure 7.6. The Gaussian function is related to the probability density function of the normal distribution used in statistics, described by Abramowitz [1, Section 26.2.9]. However, the normalisation requirements of statistics do not apply here. The fitted function is of the form $\exp\left[-(\theta/b)^2\right]$, where b defines the width of the beam at the -8.69 dB point. The fitted Gaussian functions are shown in grey in Figure 7.6. Thus it is possible to associate a beam width, corresponding to the broad bandwidth CTFM signal, with each different transducer radius. The width of the fitted Gaussian functions can be seen to correspond closely with the width of the central lobe calculated at the median sweep frequency,

shown in Figure 7.5.

7.2.2 Receiver Orientation

Previous work on binaural ultrasonic sensors for *robotic agents* included work by Kleeman and Kuc [82, 108] in which the receivers were not angled. Other work by Kuc [84] which made the receivers independently steerable, enabled them to fixate upon the target. This work is based on the design principles used by Leslie Kay's ultrasonic glasses [71], which suggest that angling the receivers apart can provide bearing measurement information. The idea dates back to the early sonar and radar work done before and during World War II (Hackmann [48] and von Kroge [134]). This is the **Inter-aural Power Difference (IPD)** method of measuring bearing.

The operation of **IPD** may be seen by examining the effect of simultaneously measuring an incoming echo with two receivers, each with a splay angle α_r from the transmitter, as shown in Figure 7.7a. The left and right receivers are each splayed by an angle $\alpha_r = 16^\circ$ from the primary axis of the sensor, leading to a 32° difference between them.

The power which will be detected by the left and right receivers can be modelled by multiplying together the transmitter beam function with each of the receiver beam functions. As the beam functions have been converted to decibels, this is equivalent to addition. The resulting composite beam functions, for the left and right paths, are shown in Figure 7.7b.

Appendix H.4 describes the combination of Gaussian beam patterns. The appendix demonstrates that the width, b , of the beam is reduced by a factor of $1/\sqrt{2}$ when the composite beam function is formed. Furthermore, the peak beam power, in decibels, is reduced by


$$-20\alpha_r^2 / \{2\ln(10)b^2\},$$

a function of both the splay angle and the beam width. The composite beam will have its centre at a splay angle of $\alpha_r/2$.


Using the composite beam functions for the left and right paths, we can see that the signal from a target located at 20° will be measured by the left receiver with a power of -4dB , while the right receiver will measure the same signal with a power of -18dB . This gives a difference of 14dB between the power measured by the left and right receivers, which is drawn as a black line in Figure 7.7b. As the difference is caused by the orientation of the receivers, it is independent of the spreading loss, air absorption and the reflecting strength of the target.

Inter-aural Power Difference


The **IPD** which will be measured for a target at a given angle is found by dividing the left received power by the right received power. As the curves have been converted to decibels, this is equivalent to subtraction. The result is shown as the black line in Figure 7.7c. This black curve may be calibrated for a given set of transducers and used to calculate the bearing of targets as previously described.



(a) Gaussian beams of the transmitter and two receivers



(b) Product of transmitter and receiver beam patterns



(c) Inter-aural Power Difference

Figure 7.7: (a) Relationship between beam patterns of the transmitter (red), left -16° receiver (green), and right 16° receiver (blue). (b) Shows the beam patterns obtained by taking the product of the transmitter beam function with each of the receiver beam functions in turn. This process ignores the spreading loss and air absorption, which are assumed to be common to both channels. The difference in detected power on the two channels is attributed to the beam patterns of the receivers. Thus a target located at 20° will have an detected power difference of about 14dB, as indicated by the black line. (c) Shows the IPD curve (black) which we may expect to measure as a target changes bearing. Note that the scale has been shifted to accommodate the IPD curve.




Figure 7.8: Derivation of the extent of the field of binaural audition. The red curve corresponds to the integrated transmitter response for a 4mm radius transducer, as shown in Figure 7.6. The blue curve is the response of the right receiver with a splay angle of -16° . (The sense of left and right are reverse by the clockwise positive angle convention.) The purple curve represents the composite response of the transmitter and the receiver, which are added. The grey curves are the fitted Gaussian functions, as in Figure 7.6. The maximum bearing at which IPD may be used to calculate target bearing is defined by the point of separation, which also defines the field of binaural audition. The composite response point of separation is defined by the left receiver point of separation, which is shifted from the transmitter point of separation by the splay angle. Therefore the composite point of separation may be found by taking the point of separation for one transducer and subtracting the splay angle.

Design Considerations

Now that the operation of IPD bearing calculation has been explained, we turn to some of the finer points of the design of such a system.

The extent over which the target bearing may be computed by IPD is limited by:

1. the breakdown of the Gaussian approximation to the beam power at large angles, as seen in Figure 7.6,
2. the noise floor of the receiver, which limits the smallest detectable signal, which in turn limits the maximum angle at which a target may be detected, and
3. the need to detect echoes from both the left and right receivers.

The breakdown of the Gaussian approximation is the primary limitation upon the operation of IPD at large bearings. The extent may be quantified by measuring the angle of separation in Figure 7.6, where the Gaussian approximation departs from the frequency integrated response. This is the extent due to one receiver. The extent defined by two receivers separated by a splay angle α_r may be found by subtracting the splay angle from the angle of separation. This is demonstrated in Figure 7.8. Table 7.1 shows the field of binaural audition which can be expected from a given

Table 7.1: The predicted widths of the field of binaural audition, as a function of the transducer radius and the splay angle. The separation angle is determined from Figure 7.6. The half-width of the field of binaural audition is determined by subtracting the splay angle from the separation angle.

Transducer Radius mm	Separation Angle Degrees	Width of cone of binaural audition Degrees			
		Splay angle α_r			
		4°	8°	12°	16°
4	36	± 32	± 28	± 24	± 20
8	17	± 13	± 9	± 5	± 1
16	10	± 6	± 2	-	-

transducer radius and splay angle configuration.

Kay [71] indicates that it is also important to obtain a sufficient slope on the IPD curve to ensure reliable operation. If the IPD curve is too shallow, the measurement noise will lead to a correspondingly large bearing uncertainty. The slope of the IPD curve can be made steeper by increasing the splay angle α_r , as shown in Figure 7.9, or by using transducers with a narrower beam width. Thus if we are able to measure the echo power to a precision of σ_p , a steeper IPD curve will allow bearing to be estimated more precisely than a shallow one, all other things being equal.

7.2.3 Selection of Transducers and Splay Angles

The properties of ultrasonic transducers discussed in Chapter 6, together with the requirements of the *robotic agent* discussed in Section 7.1 and the operation of the IPD method of measuring bearing, now provide us with sufficient information to select the ultrasonic transducers and the splay angles.

It is clear that a wide field of binaural audition is required, so that objects may be detected and localised within a wide area. Inspection of Table 7.1 demonstrates that choosing one of the larger transducers has a severe impact upon the width of the field of binaural audition. The Polaroid transducer with radius 16.2mm, which was reviewed in Section 6.8, clearly has inadequate beam width for this application. Therefore the trade off between transducer beam width and radiated power, which are related through directionality, is resolved in favour of beam width. It is intended that the improved signal to noise ratio made available by the CTFM and the Barker coded matched filter processes will be used to compensate for the low power available with small transducers. Of the transducers available off the shelf, the 4mm radius transducer provided the desired characteristics.

The splay angle α_r was subsequently selected to be 16° as this provides the best IPD resolution, while retaining a $\pm 20^\circ$ field of binaural audition.

Experimental results of bearing determination by IPD will be discussed in Chapter 14, and may be seen in Figure 14.20.




Figure 7.9: The IPD curves obtained for various splay angles with 4 mm radius transducers. The black lines are the difference between the right and left composite integrated frequency response, and the grey lines are the difference between the Gaussian models of the frequency response, in decibels. The grey line in this plot corresponds to the black curve in Figure 7.7c. The angle of separation is seen to decrease as the splay angle is increased, as seen in Table 7.1. The IPD slope also increases, allowing improved IPD precision. The experimental results corresponding to this calculation are shown in Figure 14.20.

7.2.4 Distance Between Receivers

The distance between receivers in a binaural system is an important design consideration. It has several effects.

1. Echoes from the left and right (and possibly the top and bottom) receivers need to be correctly associated with each other for each target before direction angle can be calculated. This is the correspondence problem, which will be discussed in Chapter 15. A sonar system with a short baseline will be shown to have a shorter search space, and hence a faster search, than one with a long baseline.
2. Bearing can be calculated by Inter-aural Distance Difference (IDD), which will be described in Section 14.1. The IDD bearing precision is improved by lengthening the baseline.
3. Different reflector types have different path geometries. A short baseline reduces the error caused by ignoring these differences (Section 14.1.2).
4. Local air turbulence affects ultrasonic range and amplitude measurements. If the two receivers are located close together, then the sonar paths of the signals will also be close together, and likely suffer the same turbulence. This leads to increased cross-correlation between the two channels, which is shown in Chapter 14 to improve bearing precision by both IDD and IPD methods.

When all of these effects were considered, it was decided to place the receivers as close together as possible. We intend to address the issue of IDD bearing precision by making precise measurements




Figure 7.10: The design of the sensor head, including the angles of the transducers. Measurements taken from Head 3. There is some variation between the sonar heads - in particular, the transmitter of Head 1 is aligned with the centre of the mounting hole. The central transducer is used as a transmitter, while the surrounding four transducers are receivers.

of the distance of flight.

7.3 Manufactured Sonar Head

The design decisions described in Sections 7.2.3 and 7.2.4 culminated in the design shown in Figure 7.10. Three of these sonar heads were manufactured by Leslie Kay's company Bay Advanced Technology. Each sonar head was hand made, causing some variation of the dimensions between each head. The measurements shown in Figure 7.10 were taken from Head 3. Photographs of the complete sonar head are shown in Figure 7.11.

Note that the mounting hole of Head 3 is centred 1.74 mm behind the centre of the transmitter. This causes rotation of the transmitter centre when the sonar head is mounted upon the turret of the precision positioner. It would have been better to have the mounting hole directly underneath the transmitter, as on Head 1. However, the tool offset of the positioner control software can be used to compensate for the offset of the mounting hole.

The dimensions of the sensor head which are crucial to the operation of the sonar are identified in Figure 7.12. The point P_{HC} identifies the head centre point, and occurs half way between the centre points of the two receivers. The distance from the transmitter to the point P_{HC} is denoted



Figure 7.11: The custom-built sensor head, showing the Kay transducers.




Figure 7.12: Geometry of the sensor head. P_{HC} is the head centre point, half-way between the left and right receivers. d_T is the distance of the transmitter in front of P_{HC} , while d_R is the distance of the receivers from P_{HC} . The splay angle of the receivers, α_r , is also equal to the angle between the principal directions of the transmitter and receiver.

d_T , and is equal to 2.65 ± 0.05 mm. The half-distance between the two receivers is denoted d_R , and is equal to 15.75 ± 0.05 mm. The angle of the receivers is denoted by α_r and is 16° . These measurements were taken from Head 3.

7.3.1 Crosstalk

There is acoustic crosstalk between the transmitter and the receiver (Figure 7.13). The figure is labelled with approximate distances corresponding to the apparent echo position, using $c = 343 \text{ m} \cdot \text{s}^{-1}$. These distances may be compared with some obvious paths on the sonar head itself (Figure 7.14). There are two kinds of crosstalk echo on each channel,

1. echoes caused by direct acoustic transmission from the transmitter to the receivers, at a slightly greater distance of flight than the distance from the transmitter to the receivers, and
2. echoes reflected from the foot of the unit.

The signal reflected from the foot was reduced by building up a fillet with blu-tac along the foot. No remedy was found to reduce the direct signal travelling from the transmitter to the receiver.

Reduction of Electrical Crosstalk

Measures were taken to eliminate direct electrical crosstalk from the transmit circuits to the receive circuits by installing extra shielding around the transmit transformer (L1 in Figure 6.12), and by using separately shielded cable for both the high voltage and ground return connections to the transmit transducer. Photographs of the modifications are shown in Figure 7.15.

As it was impossible to eliminate the direct path crosstalk, it was not possible to evaluate the effectiveness of the internal electrical shielding.

7.4 Conclusion

This chapter has described the requirements of the *robotic agent* in terms of sensing capacity, and derived the requirements of the sensing volume, the range precision and resolution, and the bearing precision. A sensor which can measure range to a precision of ± 1 mm, with range resolution of 10 mm and a maximum range of 5 m is desired. In addition, it is desired to have a bearing precision of $\pm 1^\circ$.

The selection of the transducers and the receiver splay angles were considered together, and a configuration consisting of a 4 mm radius transducer together with a splay angle of 16° for each receiver was selected to provide a field of binaural audition of $\pm 20^\circ$. The low sensitivity implications of selecting small transducers are to be overcome by utilising high gain signal processing techniques.

The transducers were placed as close together as possible to favour the cross-correlation of the distance of flight and the echo power measurements, which helps to produce more precise bearing measurements by both IDD and IPD. This decision also makes the solution of the correspondence problem simpler and reduces the effects of different reflector geometries.




Figure 7.13: Crosstalk signals recorded using the Barker coded matched filters, which will be described in Chapter 8. The echo time is ascribed to the peak in the output of the matched filter, as marked. There are two principal echoes visible on each channel, corresponding to the path lengths shown in Figure 7.14.




Figure 7.14: Diagram showing the crosstalk paths and their lengths. Dimensions in millimetres.

7.4.1 Acknowledgements and Contributions

The principal design elements of the sonar head are due to my supervisor, Phillip M^cKerrow. Figure 7.7 is also a device used by M^cKerrow.

The comparison between the ideal Gaussian beam function and the actual beam pattern averaged over the bandwidth of the signal (Figure 7.6), and its subsequent application to the determination of the binaural operational region of the sensor is this authors original work.

The sensor head was manufactured by Professor Leslie Kay's company, Bay Advanced Technologies.

The identification of the acoustical crosstalk paths and the efforts to improve the shielding and reduce the electrical cross-talk within the sonar head are due to this author.




Figure 7.15: Photos showing the internal construction of the sonar head, and the additional shielding which was fitted to the electronics. Head 3 is shown. The brass transformer shield is visible at the right of the middle photograph. The lower photograph shows the shielded cable connected to the rear of each of the five transducers. There are two shielded wires to the transmitter carrying the forward and reverse current. The back of the transmit transducer has also been shielded using brass plate.

Part III

Sonar DOF Sub-System

Chapter 8

Matched Filter System

This chapter builds an understanding of [Matched Filter \(MF\)](#) sonar systems. It examines how the selection of the transmit signal's shape determines the power in the sonar signal, the precision with which the target can be located, and the minimum separation at which two targets can be resolved. Comparison of several available signals leads to the selection of a Barker coded signal as the best signal to achieve the goals set out in [Section 1.2](#). The chapter discusses the methods of implementing the matched filter, and the extraction of an envelope signal ready for the detector. The operation of the detector is described, followed by the interpolation of the precise time and amplitude of the detected peaks. The chapter ends with the application of the [time-bandwidth product](#) theory from [Section 4.6](#) to estimate the resolution of this type of sonar system.

The work described in this chapter builds upon and extends the work of Peremans [[108](#)].

8.1 Matched Filters

The sonar system transmits a signal, $s_{trans}(t)$, which follows a path as described in [Chapter 4](#). The actions of each element of the path have a cumulative filtering effect, $h_{path}(t)$, upon the signal, as described in [Section 4.2.3](#). (While [Section 4.2.3](#) includes the receive processor in the path, it is omitted here.) Furthermore, noise, $n(t)$, is added to the signal by various elements. Thus the received signal is

$$s_{rec}(t) = h_{path}(t) * s_{trans}(t) + n(t), \quad (8.1)$$

being a filtered version of the transmit signal with additive noise. The path filter, $h_{path}(t)$, may be decomposed into two parts. The first is a filter component, $h_{path-filt}(t)$, which has no delay, making it non-causal. The second is a delay component, $\delta(t - \tau)$. Thus the path response is

$$h_{path}(t) = \delta(t - \tau) * h_{path-filt}(t). \quad (8.2)$$

The received signal ([8.1](#)) may now be written as

$$s_{rec}(t) = \delta(t - \tau) * h_{path-filt}(t) * s_{trans}(t) + n(t). \quad (8.3)$$

This is simplified by combining the transmitted signal, $s_{trans}(t)$, with the path filter component, $h_{path-filt}(t)$, to obtain $s_{tpf}(t)$. Thus we may write

$$s_{rec}(t) = \delta(t - \tau) * s_{tpf}(t) + n(t) = s_{tpf}(t - \tau) + n(t). \quad (8.4)$$

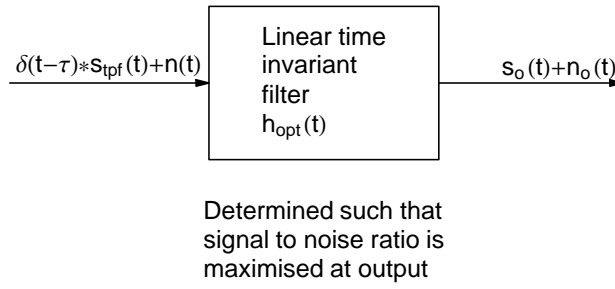


Figure 8.1: The problem of the matched filter is to maximise the power of the output signal, $s_o^2(\tau)$, at the arrival time, τ , relative to the expected noise power, $E\{n_o^2(t)\}$, for a known input signal $\delta(t - \tau) * s_{tpf}(t)$. The method of finding the optimal filter, $h_{opt}(t)$, is derived in Appendix B.1.

The received signal (8.4) is then filtered to obtain the best **Signal to Noise Ratio (SNR)**, as shown in Figure 8.1. The operation of the optimal filter, $h_{opt}(t)$, may be written as

$$s_o(t) + n_o(t) = h_{opt}(t) * [\delta(t - \tau) * s_{tpf}(t) + n(t)]. \quad (8.5)$$

Then $s_o(t)$ is the filtered signal component, while $n_o(t)$ is the filtered noise component.

The matched filter problem is to find the optimal filter, $h_{opt}(t)$, which will maximise the output **SNR** at the echo arrival time, τ , given the received signal, $s_{rec}(t)$, as input. The expected noise power, $E\{n_o^2(t)\}$, may be estimated by averaging the noise power over a time interval,

$$E\{n_o^2(t)\} = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} n_o^2(t) dt,$$

where t_1 and t_2 delimit the time period of interest. The **SNR** at the echo arrival time, τ , may then be evaluated by $s_o^2(\tau) / E\{n_o^2(t)\}$. The filter which produces the best **SNR** is derived in Appendix B.1. It depends upon the expected input signal, $s_{tpf}(t)$, and the noise spectrum. When the noise signal, $n(t)$, is white, the optimal filter is given by

$$h_{opt}(t) = k'_{mf} \cdot s_{tpf}(-t), \quad (8.6)$$

where k'_{mf} is the matched filter scaling constant. This may be used to normalise the output amplitude if desired. A method of calculating k'_{mf} is given in Appendix B.1. The peak in the matched filter output will occur at the delay time, τ , of the received signal, due to the delay function in (8.4). It is to be noted that the impulse response of the matched filter is the temporal reversal of the expected input signal. This expected input signal is the transmitted signal, $s_{trans}(t)$, modified by the filter component of the path, $h_{path-filt}(t)$.

Equation (8.5) may be decomposed into two parallel equations,

$$s_o(t) = h_{opt}(t) * \delta(t - \tau) * s_{tpf}(t) \quad (8.7)$$

$$n_o(t) = h_{opt}(t) * n(t), \quad (8.8)$$

acting upon the signal and the noise separately. Turning to the signal component, the filter output signal (8.7) may be expanded to

$$s_o(t) = k'_{mf} s_{tpf}(-t) * s_{tpf}(t - \tau). \quad (8.9)$$

When the travel delay time, τ , is equal to zero, the output signal, $s_o(t)$, is the autocorrelation, of the expected received signal $s_{tpf}(t)$. That is,

$$s_{tpf}(-t) * s_{tpf}(t) = s_{tpf}(t) \star s_{tpf}(t),$$

where \star denotes correlation.

The interested reader is referred to Appendix B.2 for a discussion of various types of signals and their autocorrelation properties, which shows that the Barker coded signal has some desirable and useful properties which recommend it above many others. The remainder of this chapter examines Barker codes and their implementation in a [matched filter](#) sonar system

8.2 Barker Code

In order to use the Barker code with a sonar system it is necessary to modulate the code onto a sinusoidal carrier. The modulation may be carried out by phase reversal, amplitude modulation or phase modulation. Only phase modulation is considered here, as the effects of other types of modulation are less convenient and produce similar results. Both the envelope and modulation functions of the Barker code are impressed upon the carrier by multiplying the carrier with the Barker code signal. For more information on Barker codes, see Berkowitz [11, Part IV, Chapter 4].

The code itself is a sequence, $a[n]$, composed of the elements $\{+1, -1\}$. Any sequence may be chosen, but for this work a sequence is chosen whose autocorrelation, $C[k]$, consists of only a central peak of magnitude N , and small, unit magnitude sidelobes:

$$C[k] = \sum_{n=0}^{N-n-k} a[n]a[n+k] = \begin{cases} -1, 0, +1 & \text{for } k = 1, 2, \dots, N-1 \\ N & \text{for } k = 0 \end{cases} \quad (8.10)$$

Sequences with this property are called perfect words. There are known perfect words with length $N = 1, 2, 3, 4, 5, 7, 11$, and 13. However, Berkowitz [11, Part 4 Section 4.3] references an alluded proof (Storer and Turyn [125]) that no such perfect words exist for odd N between 13 and 101. For this work a perfect word of length 13 is chosen,

$$\{1, 1, 1, 1, 1, -1, -1, 1, 1, -1, 1, -1, 1\}.$$

There are three additional perfect words of length 13, being the sign reversal, the temporal reversal and the combined sign and temporal reversal of the one shown above. The perfect word given above is plotted, along with its autocorrelation, in Figure 8.2.

The implementation of a Barker coded matched filter system is shown in Figure 8.3. This implementation will now be described by addressing the major components of Figure 8.3.

8.3 Barker Coded Transmit Signal Generator

The Barker coded transmit signal generator provides a signal which is suitable for transmission and reception by the sonar system. To match the signal carrier frequency and bandwidth to that




Figure 8.2: The Barker code and its autocorrelation, which satisfies (8.10). The processing gain is 13, and the maximum sidelobe height is 1. The ratio of the maximum height to the sidelobe height is the [Signal to Sidelobe Ratio \(SSR\)](#).

of the Kay transducers (Figure 6.13), the carrier frequency was selected to be $f_c = 83.3\text{ kHz}$. This frequency is near the peak response of the Kay transducer. Each bit of the code word was used to modulate three cycles of the carrier. Using a small number of cycles increases the bandwidth of the signal and, importantly, decreases the width of the autocorrelation function.

An additional reason for choosing the carrier frequency to be 83.3 kHz is the quadrature demodulation process described in Section 8.6.2. Quadrature demodulation requires a sampling frequency, f_s , four times the carrier frequency, or $f_s = 333.3\text{ kHz}$. This sampling frequency coincides with one of the available sampling rates on the chirp cards, described in Section 5.1, obtained by dividing the NuBus 10 MHz clock frequency by 30. Thus the sampling rate and carrier frequency are determined exactly by the relationships

$$f_s = \frac{10\text{ MHz}}{30} = 333.3\text{ kHz}, \quad f_c = \frac{f_s}{4} = 83.3\text{ kHz}. \quad (8.11)$$

The coded word is used to phase modulate the carrier, by direct multiplication. This signal is shown in Figure 8.4a, together with its autocorrelation function in Figure 8.4b. Changing values of the code bit causes a phase change of 180° .

The spectrum of the transmitted signal is shown in Figure 8.4c. The basic pattern is a sinc function, which Section 8.5 will demonstrate to be linked to the spacing between the code bits. While the main lobe of the signal is within the bandwidth of the transducer, the high side lobes are outside the transducer bandwidth, and will be lost. The effects of this will be seen in the next section.

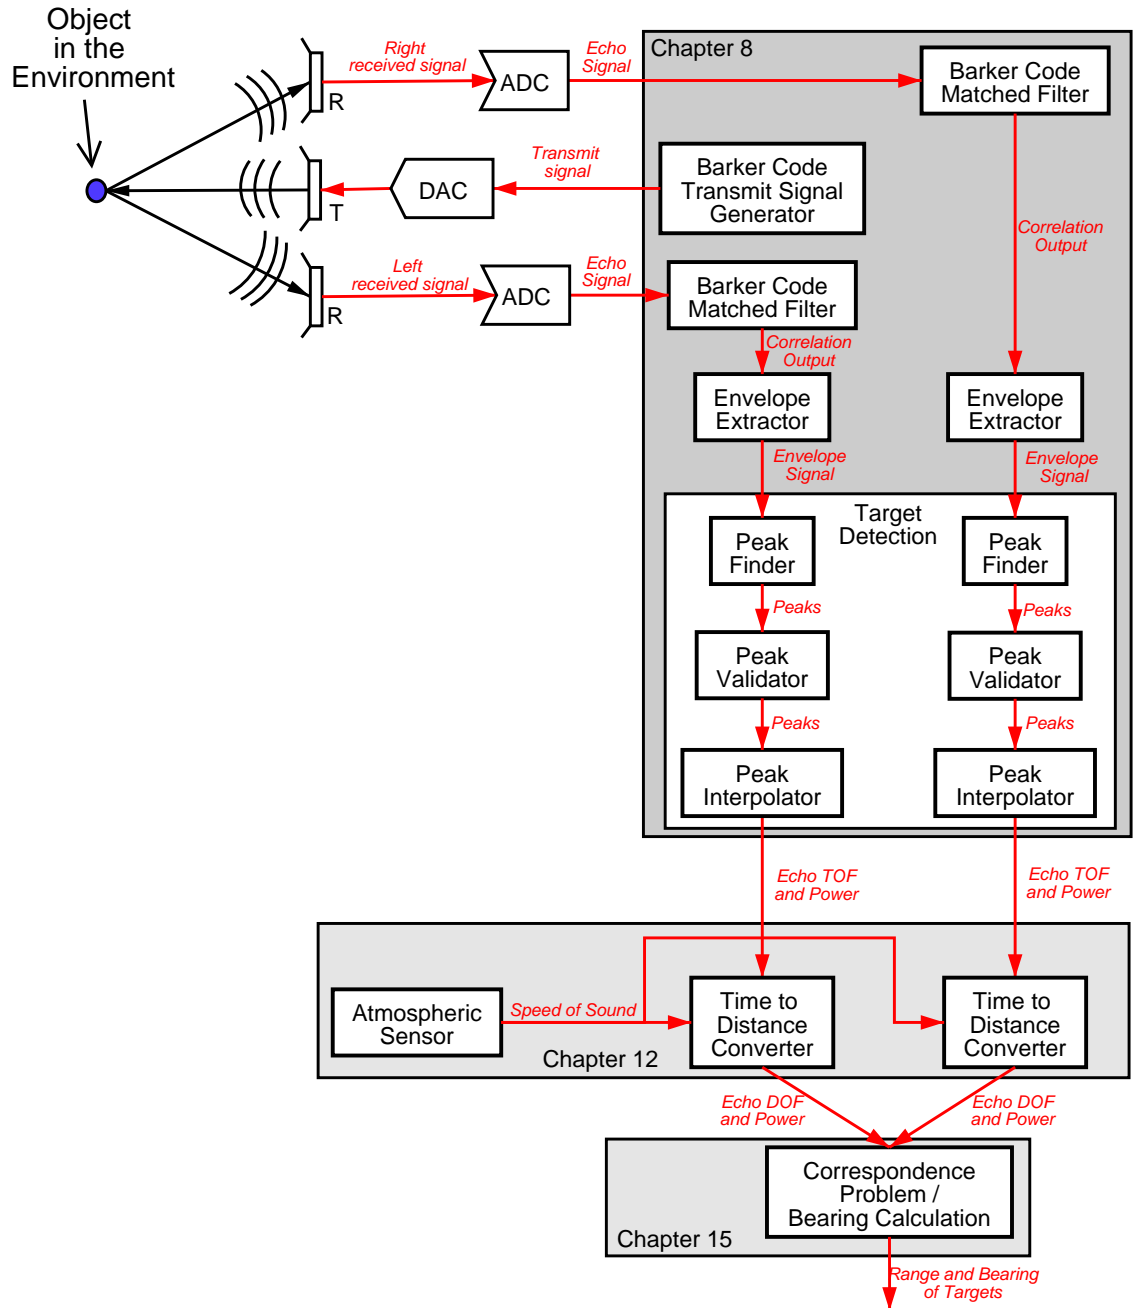





Figure 8.3: Overview of the Barker code matched filter sonar system. This diagram shows two receiver channels making a binaural system.



(a) transmitted signal



(b) transmit signal autocorrelation



(c) transmit signal power spectrum

Figure 8.4: The signal transmitted by the Barker code matched filter sonar. (a) The time domain representation of the transmit signal. (b) Transmit signal autocorrelation. The timescale in this plot indicates the length of signal used in the sonar system. The packet length is $468\mu\text{s}$ or 162mm . However, the total width of the autocorrelation peak is only $72\mu\text{s}$ or 25mm . (c) Frequency power spectrum of the transmit signal. The main lobe of the spectrum is spread over 50kHz . The two sidelobes on the left are stronger than those on the right due to additive combination with the sidelobes from the negative half of the spectrum. The sidelobes fit a sinc function, and are caused by the square edges of the modulation function. They are at least 10dB down in power relative to the main lobe.

8.4 ADC Output

The received echo, as detected by the receiver and processed by the ADC, is filtered by the transmission path, $h_{path}(t)$. Figure 8.5a shows a real echo that was transmitted as shown in Figure 8.4a. The transmit signal has been changed by significantly reducing the out of band power, as may be seen by comparing Figures 8.4c and 8.5c. A significant proportion of these changes can be ascribed to the transmitter and receiver frequency responses, shown in Figure 6.13. This has the effect of rounding the edges of the time domain signal, which may be seen by comparing Figures 8.4a and 8.5a.

The missing parts of the signal spectrum influence the response at the output of the matched filter. The filter is not matched to the transducer and air effects. As the path angle at the receiver must be known a-priori to form the proper matched filter, and a different angle used for each target, this approach was not used. Instead, the degradation in the filter output is measured, and used to estimate the bearing. This topic will be further addressed in Chapter 14.

Figure 8.5b shows the matched filter output. The signal to sidelobe ratio in this plot is 9. The reduction in performance compared to the reference filter output shown in Figure 8.2b is due to the loss of the high frequency components of the signal.

8.5 Barker Code Matched Filter


The principal function of the Barker code matched filter component is to perform the optimal filtering operation, $h_{opt}(t)$, upon the received signal. Equation (8.6) defined the optimal filter to be the time reversed expected echo signal, $s_{tpf}(-t)$, including the filtering effects of the path. The filter is applied through the convolution operator, whose definition includes the time reversal of one of the arguments. This argument is assumed to be the optimal filter, without loss of generality. Thus the filtering operation is equivalent to a correlation, involving no time reversal, between the received signal and the expected received signal. As the actual echo signal and the expected echo signal are the same, the autocorrelation of the echo signal is obtained.

However, calculation of the optimal filter requires a-priori knowledge of the signal path, including the transmit and receive angles. As this is unknown, the path filtering component, $h_{path-filt}(t)$, is omitted when computing $h_{opt}(t)$. Thus instead of $h_{opt}(t)$ a sub-optimal filter $h_{mf}(t)$ is used, defined by


$$h_{mf}(t) = s_{trans}(-t). \quad (8.12)$$

While this produces a sub-optimal matched filter, it allows the computation of the filtering operation to be optimised. Furthermore, the effects of the receiver upon the signal, which are used in the computation of bearing by IPD, are passed through.


The transmit signal may be viewed, in addition to the modulated carrier model, as a convolution between a coded impulse train, a one bit impulse sequence, and one cycle of the carrier signal, as shown in Figure 8.6. The time domain convolution decomposition was chosen as it corresponds to a frequency domain multiplicative decomposition, which in turn leads to a decibel power domain additive decomposition. This is instructive for describing the source of the various components of the spectrum.



(a) received echo signal



(b) received echo signal matched filter output



(c) received echo signal power spectrum

Figure 8.5: A sample experimental received Barker coded echo signal. The signal transmitted by the Barker code matched filter sonar. (a) The time domain representation of the echo signal. (b) The matched filter output. The width of the autocorrelation peak is about $80\mu\text{s}$ or 28mm. The signal to sidelobe ratio is 9. (c) The frequency power spectrum of the echo signal. The main lobe of the spectrum remains similar to that of the transmit signal, shown in Figure 8.4b. However, the upper sidelobes have been lost, and the lower sidelobes attenuated. This causes the time domain signal to lose amplitude near the phase changes.

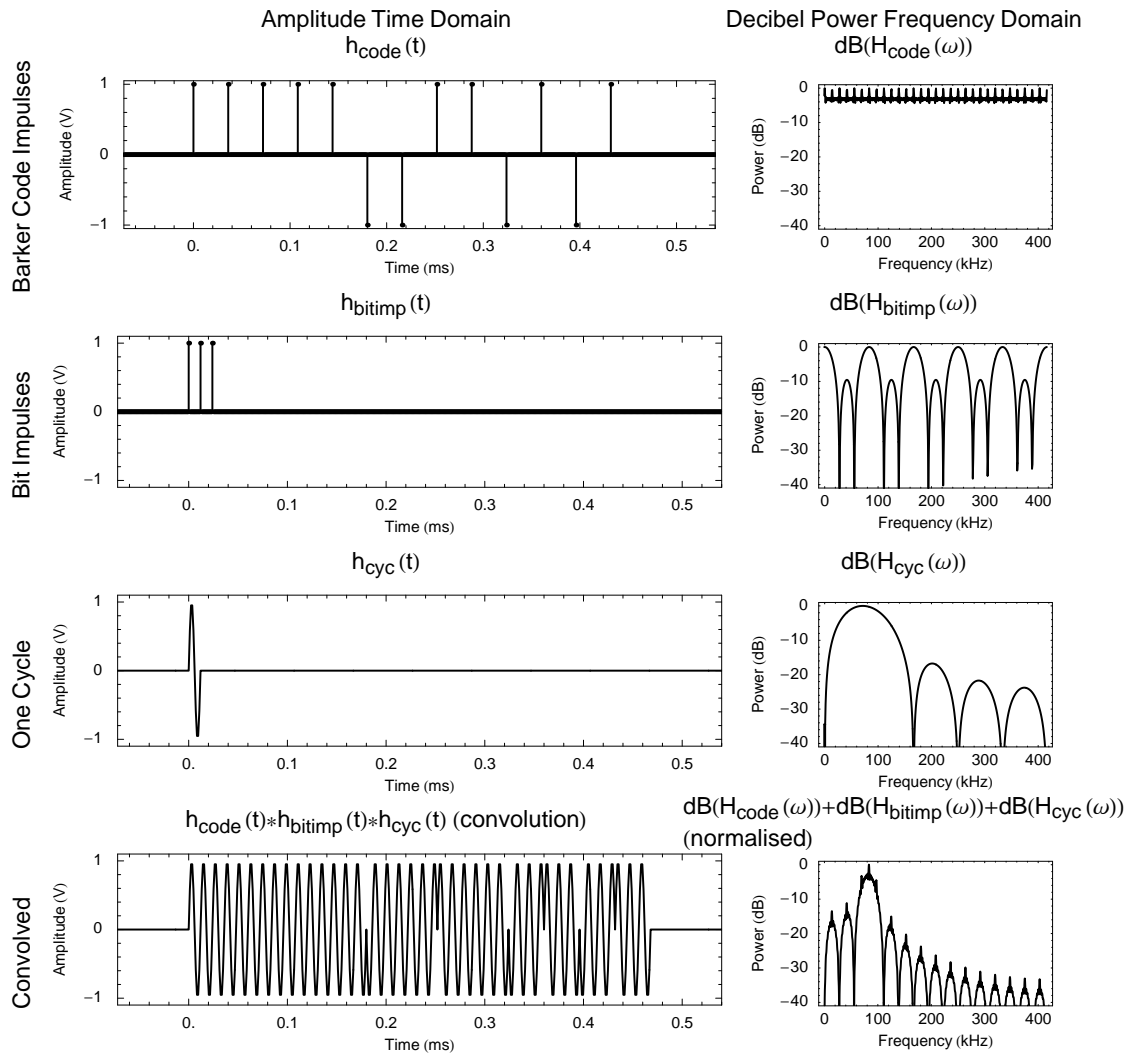


Figure 8.6: The ideal computer synthesised Barker coded transmit signal may be viewed as a convolution between a coded impulse train, a one bit impulse sequence, and one cycle of the carrier signal. The first three panels in the left column are combined by convolution to form the waveform at the bottom left. Convolution of a signal with the impulse function reproduces the signal (see Section 4.2.3 or Bracewell [14, Chapter 5]), with a delay if the impulse is offset. The result after convolution is the same as the amplitude modulation technique mentioned in Section 8.2. The right column shows the (normalised) spectra of each of the signals, and demonstrates that the small spikes come from the coded impulse train. The spectra from the first three panels in the right column are combined by addition in the Decibel domain to form the spectra in the fourth panel.

Utilising the time domain convolution decomposition, the transmit signal, $s_{trans}(t)$, may be written as

$$s_{trans}(t) = h_{code}(t) * h_{bitimp}(t) * h_{cyc}(t).$$

The matched filter, $h_{mf}(t)$ from (8.12), is therefore

$$h_{mf}(t) = h_{code}(-t) * h_{bitimp}(-t) * h_{cyc}(t),$$

where the scaling constant has been ignored. The filtered received signal is

$$s_o(t) = h_{mf}(t) * s_{rec}(t) = h_{code}(-t) * h_{bitimp}(-t) * h_{cyc}(t) * s_{rec}(t).$$

The convolution of $h_{code}(-t)$, $h_{bitimp}(-t)$ and $h_{cyc}(t)$ may be performed in any order as convenient. A diagram of the three-step matched filter procedure is shown in Figure 8.7. A matched filter implemented in this manner may be implemented with 4 multiply-add instructions per output point for $h_{cyc}(-t)$, another 4 multiply-add instructions for $h_{bitimp}(-t)$, and a further 13 multiply-add instructions per output point for $h_{code}(-t)$, giving a total of 21 multiply-add instructions per output point. This is to be contrasted with the direct implementation, which requires a total of 156 multiply-add instructions per output point. The optimised computation reduces the workload by 86%.

A further computational optimisation may be obtained by removing the multiplication by ± 1 and instead coding the additions and subtractions explicitly. This computational simplification may also be carried over to the sine-wave packet by choosing the samples $\{1, 1, -1, -1\}$, as shown in Figure 8.8. The sine-wave convolution may then be coded directly in terms of additions and subtractions. The samples chosen in Figure 8.8 introduce a shift of $1/(8f_c)$ relative to the sine-wave. The software used in this thesis attaches a zero offset field to every signal record, indicating the distance between the first array element and zero time. The shift is compensated simply by modifying the offset field.

Only some of the optimisations described above are implemented in the software. The unit sine-wave samples are used, but the Barker code convolution is implemented using general purpose code, so that the Barker code and the width of each bit may be changed.

At this point the cross-correlation between the received signal and the transmit signal has been calculated, and the signal marked correlation output in Figure 8.3 has been obtained. This signal is now passed to the envelope extractor.

8.6 The Envelope Extractor

The output of the matched filter operation is shown in Figure 8.5b. Since the signal retains the carrier frequency at this point, it is important to extract the envelope before performing detection so that the position and magnitude of the peak, and therefore the time of flight and echo power of the returned echo, may be accurately obtained. The reasons for this are explained in detail in Section 8.6.1.

There are three principal processes which may be used to extract the envelope.

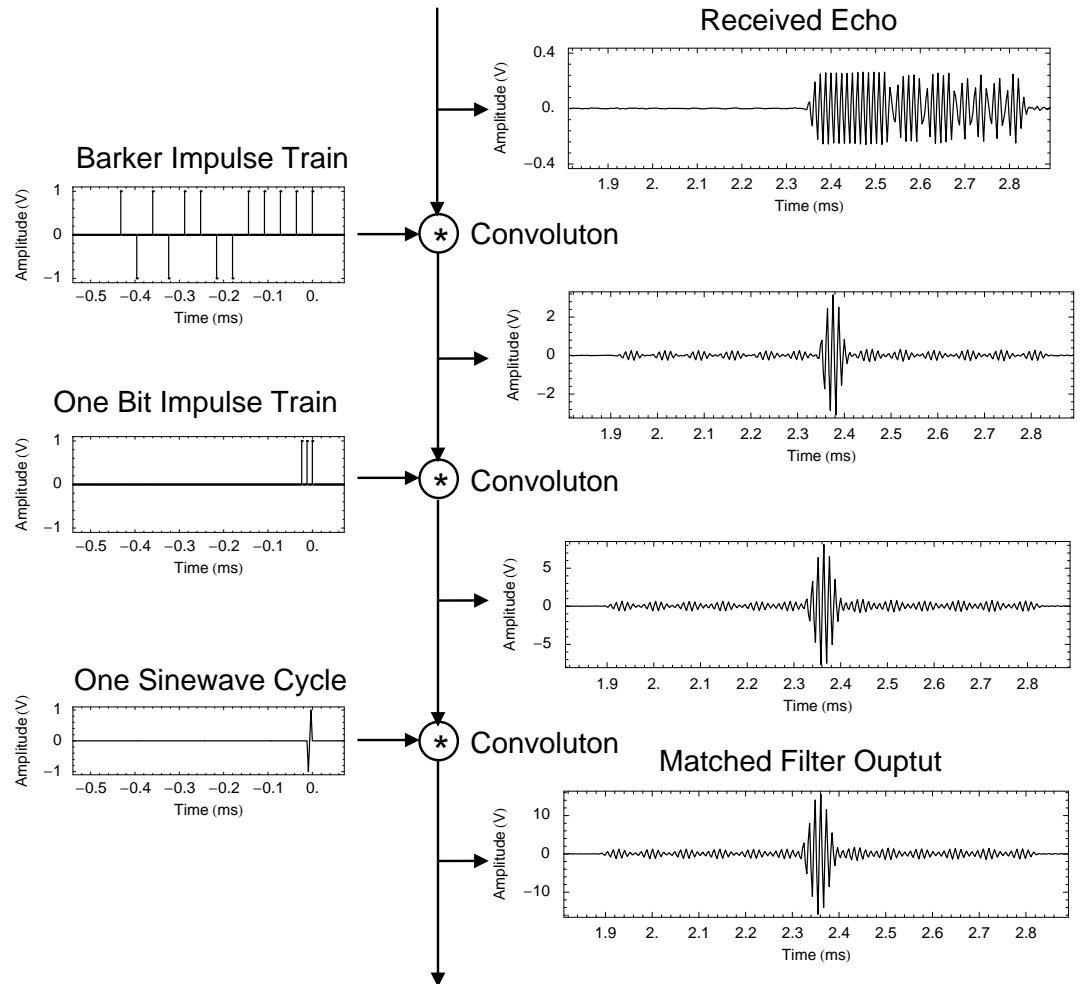


Figure 8.7: The operation of the three-step matched filter procedure. The first convolution includes only 13 non-zero points, while the second and third include only four non-zero points each. While the intermediate stages might appear to provide suitable output to perform peak finding, these stages will provide poor noise suppression. The second and third convolutions are required to realise the noise rejection benefits of matched filtering, and also to move the peak to the correct position. The convolution operation includes a temporal reversal of one of the signals. Thus the signals of Figure 8.6 are reversed to obtain a cross-correlation between them and the received signal. The amplitude scales of the signals on the right increase as the various convolution stages are carried out.




Figure 8.8: Unit sampled one-cycle sine-wave. The sine-wave amplitude is $\sqrt{2}$, and the samples have unit amplitude. There is a time offset of $1/(8f_c)$ which must be taken into account. The unit samples allow for efficient evaluation. The convolution output will require normalisation by $1/\sqrt{2}$ if it is to be compared with other methods.

Quadrature demodulation: The signal is demodulated by taking the root-mean-square sum of adjacent samples. This only works if the signal is sampled using exactly $f_s = 4f_c$, as defined in (8.11).

Unsynchronised demodulation: The signal is rectified and passed through a suitable low-pass filter.

Synchronised demodulation: The signal is multiplied in the time domain by the carrier signal, to obtain a sum frequency and a difference frequency. Then a low pass filter is used to remove the sum frequency. This method requires the demodulation carrier to be phase and frequency synchronised with the signal carrier (Oppenheim et al [104, Chapter 7]). Difficulties will be encountered if the channel has modified the signal spectrum (by filtering) so as to alter the mean or carrier frequency (equation (4.19), see also Rihaczec [118, Section 2.1]).

Synchronised demodulation is difficult, due to the need for a synchronised carrier. Additionally, the carrier frequency phase will vary for each target, requiring demodulation to be performed separately for each target. Determining the phase accurately in the case of overlapping echoes would be difficult. The methods of quadrature and unsynchronised demodulation are practical to implement, and will be described in Sections 8.6.2 and 8.13 respectively.

8.6.1 The Need for Envelope Extraction

From the output of the matched filter shown in Figure 8.5b, it can be seen that it may be possible to perform detection directly upon the matched filter output. The signal path may be represented as a sequence of filters (Section 4.2.4), and each of these filters may affect the mean frequency of the signal (4.19) by modifying its spectrum. The mean frequency of the signal will be observed as the carrier frequency (Rihaczec [118, Section 2.6]). If the carrier frequency is changed, then the position of the peak within the autocorrelation function will necessarily move, thereby destroying the ideal relationship shown in Figure 8.4b. Thus the effect of filters in the signal path is to

change the relationship between the carrier peak and the envelope peak. This has been verified experimentally, but the results are not shown. Since the echo time is associated with the peak of the matched filter output rather than the peak of the carrier, it is necessary to extract the envelope of the signal for detection rather than performing detection directly upon the matched filter output.

8.6.2 Quadrature Demodulation

Quadrature Demodulation (QD) appears in Peremans' [108] implementation of the Barker coded matched filter sonar system. It assumes that the input signal is sampled at $f_s = 4f_c$. This condition is met by the selection of the carrier and sampling frequencies in Section 8.3. As the quadrature demodulation is carried out on discrete samples, a discrete notation, as used by Oppenheim and Schaffer [105], is adopted for this sub-section. The integer n denotes the array index.

Quadrature demodulation assumes that the sampled signal, $x[n]$, is composed of a carrier frequency, f_c , and an amplitude modulation, $A[n]$, with frequencies much lower than f_c , as seen in the matched filter output of Figure 8.5. Thus the signal $x[n]$ may be represented by

$$x[n] = A[n] \cdot \sin\left(2\pi f_c \frac{n}{f_s} + \phi\right),$$

where ϕ is an arbitrary phase offset. Substituting $f_s = 4f_c$ and simplifying yields

$$x[n] = A[n] \cdot \sin\left(\frac{\pi n}{2} + \phi\right). \quad (8.13)$$

To recover $A[n]$ from $x[n]$ and $x[n+1]$, adjacent samples are combined in quadrature. Substituting (8.13) yields

$$\sqrt{x^2[n] + x^2[n+1]} = \sqrt{A^2[n] \sin^2\left(\frac{\pi n}{2} + \phi\right) + A^2[n+1] \sin^2\left(\frac{\pi(n+1)}{2} + \phi\right)}$$

Assuming that the amplitude of the modulation does not change significantly between samples (i.e., $A[n+1] \approx A[n]$), the right hand side may be simplified to

$$\sqrt{A^2[n] \left\{ \sin^2\left(\frac{\pi n}{2} + \phi\right) + \sin^2\left(\frac{\pi n}{2} + \phi + \frac{\pi}{2}\right) \right\}}.$$

Using the trigonometric identity $\sin\left(x + \frac{\pi}{2}\right) = \cos x$ and moving the amplitude term outside the square root yields

$$\sqrt{x^2[n] + x^2[n+1]} \approx A[n] \sqrt{\sin^2\left(\frac{\pi n}{2} + \phi\right) + \cos^2\left(\frac{\pi n}{2} + \phi\right)},$$

which reduces to

$$\sqrt{x^2[n] + x^2[n+1]} \approx A[n].$$

A typical demodulated echo is shown in Figure 8.9. Note that quadrature demodulation produces a ragged envelope due to the approximation $A[n+1] \approx A[n]$. This has been found to adversely affect the process of accurately extracting the peak positions (Section 8.7).

The output sampling period, f_s , being linked to the carrier frequency, f_c , is indirectly determined by the frequency response of the transducer. The values provided in (8.11) provide a sampling period of $1/f_s = 3 \mu s$, and a DOF quantisation of 1 mm. The precision of measurements will be improved by the interpolation scheme of Section 8.7.3.




Figure 8.9: The various stages of **quadrature demodulation**. The black line is the matched filter output, the red line is after the upsampler, and the green and blue lines are the positive and negative envelopes. The upsampled signal is included for comparison with the determined envelope, but is not used in the quadrature demodulation process. The data is the same as that shown in the matched filter output of Figure 8.5b, zoomed in to the peak region.

8.6.3 Unsynchronised Demodulation

The four steps comprising **Unsynchronised Demodulation (UD)** in this implementation are summarised in Figure 8.10. While the first block shown in this diagram is the four times upsampler, it is instructive to start with the rectification.

The effect of rectification is related to synchronised demodulation. Consider performing two parallel synchronised demodulations upon an input signal, $s_i(t)$, using two synchronised oscillators, as shown in Figure 8.11. The two oscillators have a $\pi/2$ phase difference between them, but are not synchronised with the received signal. Now combine the two demodulated signals in quadrature. The algebraic representation of these operations is

$$s_o(t) = \sqrt{\{s_i(t) \cdot \sin(2\pi f_{base}t)\}^2 + \{s_i(t) \cdot \cos(2\pi f_{base}t)\}^2},$$

where $s_o(t)$ is the output signal, which is equivalent to

$$s_o(t) = |s_i(t)| \sqrt{\sin^2(2\pi f_{base}t) + \cos^2(2\pi f_{base}t)} = |s_i(t)|.$$

Therefore the rectification is equivalent to multiplication by two local oscillators being 90° out of phase. The implication in the frequency domain is that a sum and difference frequency are produced by the multiplication with the local oscillator, and a further set of sum and difference frequencies are produced when the squaring procedure takes place, for a maximum frequency of four times the original carrier frequency. The signal must therefore be upsampled before carrying out the absolute value procedure so that the sum and difference frequencies will not be aliased. Failing to up-sample before rectification produces unreliable results, due to the aliasing.

The signal is upsampled by four using standard techniques, as described by Proakis and Manolakis [116], to obtain a new sampling rate f_{su} . There are three steps. Firstly, multiply the

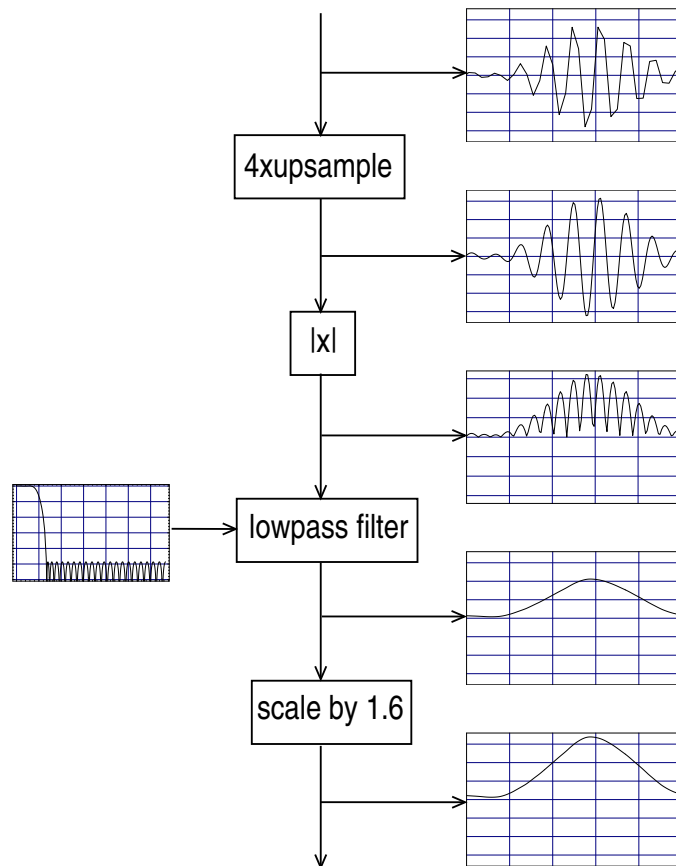


Figure 8.10: Block diagram of **unsynchronised demodulation**. All of the waveforms are plotted on the same scale. The input signal is sampled at four times the carrier frequency, so the samples are visible. The first stage is to up-sample the signal by multiplying the samples by 4, inserting three additional samples between every existing sample, and applying a smoothing filter. The upsampling procedure is described by Proakis and Manolakis [116]. The rectification step takes the absolute value of the signal, and a low pass filter is applied. This provides an envelope function, which is a little too small. This is corrected by applying a scaling factor of 1.6, which was empirically determined to provide a good match to the input signal.




Figure 8.11: The process of dual synchronised demodulation. The two oscillators run at the same frequency as the carrier, but are not phase matched to the carrier signal. Instead, they are $\pi/2$ out of phase with each other.




Figure 8.12: The spectrum of the signal before and after the time domain rectification operation. The movement of the carrier frequency, visible in (a), to baseband and harmonics in (b), may be seen graphically in the time domain in Figure 8.10, where the doubling of the carrier frequency is clearly visible.

amplitude of each sample by four. Secondly, insert three zeros between each sample of the signal. Thirdly, low pass filter the signal to maintain the spectrum of the original signal. The sampling rate of the output is $f_{su} = 4f_s = 1.333\text{ MHz}$, where f_s was defined in (8.11).

The rectification of the signal, carried out by the absolute value function $|x|$, has two effects upon the signal, which are shown in Figure 8.12. The first effect is to move the amplitude modulation down to base band or zero frequency. The second effect is to double the carrier frequency and introduce harmonics.

The carrier and its harmonics may be removed by using a low pass filter, which is shown in Figure 8.13. After processing the rectified signal with this filter, the envelope is clearly visible (Figure 8.10), although its amplitude is too small. This is due to the filter averaging the signal, which occurs in a predictable way, and may be corrected by multiplying the filter output by a constant. The appropriate value was empirically determined to be 1.6. The correction factor is




Figure 8.13: The filter used to remove the (doubled) carrier frequency and its harmonics. The filter type is equiripple FIR, passband $0-0.05 f_{su}$, ripple 0.1, stop-band $0.1-0.5 f_{su}$, and 50dB attenuation. The filter order is 48. The filter is used on data with a sampling rate of $f_{su} = 1.333 \text{ MHz}$.

incorporated into the low pass filter coefficients, so no separate stage is necessary. The resulting envelope is plotted, together with the matched filter output and the upsampled signal, in Figure 8.14. The sampling period of the envelope, $1/f_{su} = 0.75 \mu\text{s}$, provides a DOF quantisation of 0.26mm.

Two methods of envelope extraction have been described. Either one may be used to provide input for the next stage, which performs target detection.

8.7 Target Detection

The envelope signals, shown in Figures 8.9 and 8.14, are suitable for target detection. The peak position indicates the time of flight to the target and back, and the peak amplitude indicates the power of the echo.

There are three steps in the target detection process, shown within the white box of Figure 8.3. The first step is the peak finder, which locates indexes in the waveform array containing elements which are larger than several adjacent points on either side. The second step is the peak validator, which verifies that the peak which was found is not a sidelobe of a larger peak, as shown in Figure 8.2. The third step, the peak interpolator, performs curve fitting to estimate the peak position and peak amplitude with sub-sample precision, thereby providing the echo time of flight and echo power.

8.7.1 Peak Finder

Peaks are found by examining the envelope samples. Monotonically increasing values occur up to the peak and monotonically decreasing values thereafter. An elegant method of performing peak




Figure 8.14: Comparison of some of the signals from various stages of **unsynchronised demodulation**. The resulting envelope smoothly interpolates the peak points of the upsampled matched filter output.

finding in Matlab is the program `pkpicker.m`, by Burrus et al. [19, Appendix A]. The `pkpicker.m` program works well on the unsynchronised demodulation envelope, shown in Figure 8.14. However, it does not work well on the **QD** envelope shown in Figure 8.9 as the monotonic requirement is violated. Using the `pkpicker.m` algorithm on this signal results in an excessive number of false peaks. The peak finder for **QD** uses a window of three points on each side to validate each peak. A sample point is deemed to be a peak if it is larger than the three points on either side. This reduces the number of false peaks reported for noisy signals, and for quadrature demodulation signals in particular. The window width is adjustable.

Peaks are only reported if they exceed the noise threshold. This is adjusted to eliminate the majority of noise peaks constituting the noise floor of the system.

The result of this stage is a list of indexes into the envelope array denoting the position of the peaks. Because sidelobes must be expected, the peak index list is passed to a peak validator.

8.7.2 Peak Validator

The matched filter output shown in Figure 8.5 contains six sidelobes on each side of the main peak. The list of peak indexes produced by the peak finder will include these sidelobes. As these peaks do not indicate targets, they must be removed from the peak output. To be valid, a peak must not fall within the sidelobe region, shown as grey in Figure 8.15, of a more powerful peak. The peak validation is carried out as shown in Algorithm 8.1.

Ideally, the **Signal to Sidelobe Ratio (SSR)** (a/b in Figure 8.15) will be 13, as defined in Section 8.2. However, the filtering of the signal path causes some degradation to occur, and the **SSR** will be reduced. A practical value which produces reliable results is 6.5, which is half the original value. The actual value used is adjustable through the user interface (see Figure 8.17).




Figure 8.15: Peak validation. This is the Barker code autocorrelation, re-drawn from Figure 8.2b. Any short peaks falling within the grey region, defined in relation to a large peak already found, are removed. The width of the grey region is $2N + 1$, where N is the length of the code.

Algorithm 8.1 Method of validating peaks and removing sidelobes. The algorithm uses list semantics.

Require: list of peaks including peak index and amplitude

Ensure: that all sidelobe peaks are removed

SortPeaksByDescendingAmplitude(peakList)

currentPeak \leftarrow first(peakList)

sidelobeAmplitude \leftarrow amplitude of currentPeak / 6.5

while currentPeak \neq end(peakList) **do**

 searchPeak \leftarrow next(currentPeak)

while searchPeak \neq end(peakList) **do**

if amplitude of searchPeak < sidelobeAmplitude **then**

if searchPeak is within sidelobe width **then**

 invalidPeak \leftarrow searchPeak

 searchPeak \leftarrow next(searchPeak)

 remove(peakList, invalidPeak)

else

 searchPeak \leftarrow next(searchPeak)

end if

else

 searchPeak \leftarrow next(searchPeak)

end if

end while

 { All peaks between first(peakList) and currentPeak inclusive are valid }

 currentPeak \leftarrow next(searchPeak)

end while




Figure 8.16: The curve fitted to the envelope from unsynchronised demodulation. Fitted curve shown in red. The blue line joins the original sample points of the matched filter output. The \times marks the detected peak.

8.7.3 Peak Interpolator

The peak interpolator, the third step in the target detection process, receives a validated list of peak positions (indexes), for the determination of time of flight and echo amplitude. These are estimated to sub-sample precision using a polynomial curve fit. The fit is performed using a standard least squares technique, see Appendix I, on a quadratic polynomial,

$$f(t) = \beta_0 + \beta_1 t + \beta_2 t^2.$$

The fit is performed using 4 or 6 points on either side of the central peak, depending upon the demodulation method used. The peak is extracted using

$$t_{peak} = \frac{-\beta_1}{2\beta_2}, \beta_2 < 0.$$

If $\beta_2 \geq 0$, then the peak shape is wrong and the peak is discarded. An example fitted peak is shown in Figure 8.16.

8.8 Range Resolution

The resolution of the sonar system is evaluated in three different ways in this section. The first method derives the theoretical resolution limit in terms of the system bandwidth, as given by equation (4.12). Section 4.6.3 demonstrated that the CTFM system has the same theoretical resolution limit. The second method is to measure the time width of a practical matched filter signal, as presented to the detector. This section also evaluates the time-bandwidth product of four signals, being a reference transmit and matched filtered signal, and a measured received and matched filtered signal. This provides insight into the operation of the MF processor. The third method is

to measure the bandwidth of the signal, and to use that to predict the limiting width. Finally, the three resolution results are compared, and the [time-bandwidth product](#) results are discussed.

8.8.1 Theoretical System Resolution

The theoretical system resolution was determined in Section 4.6.3 to depend upon the system bandwidth. This section will utilise the transducer frequency response measured in Section 6.7 to estimate the system bandwidth and hence calculate the resolution. The response of a single transducer is shown in Figure 6.13. The response of two transducers can be obtained by squaring the response of one transducer. Measuring the frequency width using (4.22) produces an estimate of 27 kHz. Application of (4.12) produces a resolution limit of 3.2 mm. This limit is the same whether a [MF](#) or [CTFM](#) signal is used, provided the signal covers the system bandwidth. The resolution limit will be degraded by any of the reasons given in Section 4.6.3.

8.8.2 Time-Bandwidth Measurements

[Time-bandwidth product](#) calculations are applied at two points in the sonar system (Figure 8.3):

1. the transmit/receive signal, which gives an indication of the signal power available, and
2. the output from the matched filter, which is used to predict the sonar resolution.

The calculations are performed upon both an ideal and an experimental set of signals, for comparison. The ideal set of signals assumes that the transducers and the air path have perfect impulse responses, or equivalently, that no transducers or air path are used at all. The results of the computations are shown in Table 8.1.

The results shown in Table 8.1 indicate that the ideal transmit signal has a moderate [time-bandwidth product](#), but that it is significantly degraded by the process of transmission and reception through real transducers. The time width is nearly unchanged by the real transducers, but the bandwidth is reduced by more than half. This indicates that there is significant power in the upper frequency sidelobes (see bottom right panel of Figure 8.6) which is filtered out by the [ADC](#) reconstruction filter and the transducers. The signal characteristics were selected to match the main lobe with the channel response, but it is impossible to faithfully transmit such a wide-band signal under limited bandwidth conditions.

The [MF](#) output signals show that the signal bandwidth has been reduced. In the case of the ideal signal, the reduction in bandwidth may be seen to be caused by the squaring of the power spectrum when performing the autocorrelation. This corresponds to the cross-correlation between the real signal and the reference signal, which is represented as the product of the two spectra in the frequency domain. The squaring of the spectrum causes the reduction in bandwidth.

The measure of the equivalent duration of the [MF](#) output signals in Table 8.1 is considerably wider than the central peaks of the [MF](#) output. This is due to the width computation including the effect of the correlation side-lobes in the time domain. If the computed equivalent widths are replaced with the observed half-height width of the central peak, then much better [time-bandwidth products](#) are obtained.

Table 8.1: Summary of time-bandwidth product computations for ideal and realistic transmit/receive and matched filtered signals. See Section 8.8 for full details.

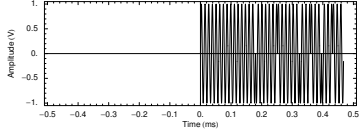
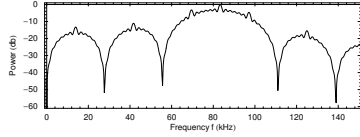
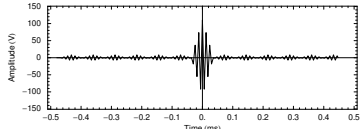
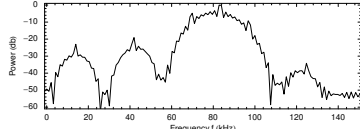
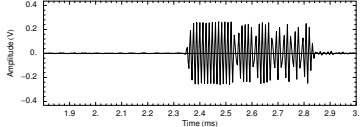
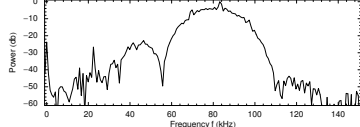
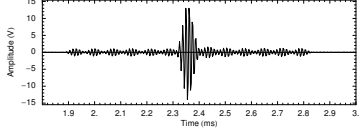
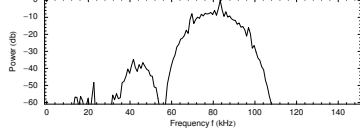
Time Domain Signal	Frequency Domain Signal	Duration	Bandwidth	TB Product
Ideal Transmit/Receive Signal 		$t_0 = 234 \mu s$ $\sqrt{t^2} = 270 \mu s$ $\Delta t = 338 \mu s$	$f_0 = 81.3 \text{ kHz}$ $\sqrt{f^2} = 84.1 \text{ kHz}$ $\Delta f = 54.5 \text{ kHz}$	$\Delta t \cdot \Delta f = 18.4$
Ideal MF Output Signal 		$t_0 = 0 s$ $\sqrt{t^2} = 73 \mu s$ $\Delta t = 183 \mu s$	$f_0 = 82.1 \text{ kHz}$ $\sqrt{f^2} = 82.5 \text{ kHz}$ $\Delta f = 19.9 \text{ kHz}$	using equivalent width $\Delta t \cdot \Delta f = 3.6$ using central width = $36 \mu s$ $\Delta t \cdot \Delta f = 0.7$
Real Received Signal 		$t_0 = 2562 \mu s$ $\sqrt{t^2} = 2566 \mu s$ $\Delta t = 335 \mu s$	$f_0 = 80.6 \text{ kHz}$ $\sqrt{f^2} = 81.0 \text{ kHz}$ $\Delta f = 19.9 \text{ kHz}$	$\Delta t \cdot \Delta f = 6.7$
Real MF Output Signal 		$t_0 = 2361 \mu s$ $\sqrt{t^2} = 2362 \mu s$ $\Delta t = 204 \mu s$	$f_0 = 81.8 \text{ kHz}$ $\sqrt{f^2} = 82.0 \text{ kHz}$ $\Delta f = 15.0 \text{ kHz}$	using equivalent width $\Delta t \cdot \Delta f = 3.0$ using central width = $40 \mu s$ $\Delta t \cdot \Delta f = 0.6$

Table 8.2: Results of theoretical MF resolution calculations.

Method	Resolution (mm)
Theoretical System Resolution	3.2
Time Width Resolution	7
Frequency Width Resolution	5.8

8.8.3 Theoretical Resolution Using Time Width

The resolution of the sonar is predicted by the width of the autocorrelation output. The widths computed by the procedure of Section 4.6.4 are clearly quite large when compared with the width of the central autocorrelation peak. This is seen to be a problem with the width computation. With this in mind, the half maximum height width of the central peak is computed to be $\Delta t = 40 \mu\text{s}$ in the case of the experimental data shown. The range resolution is then computed to be

$$c \cdot \frac{\Delta t}{2} = 7 \text{ mm.} \quad (8.14)$$

8.8.4 Theoretical Resolution Using Bandwidth

The range resolution limit may be computed from the signal bandwidth using Gabor's uncertainty relation (4.11). The *limit* is due to assuming that the signals used are Gabor's elementary signals. The difference between the practical signals and Gabor's elementary signals means that we will not be able to match the limit in practice. The relationship between (Gabor's measure of) bandwidth, Δf , and range resolution, Δr , is found by combining Gabor's relation (4.11) with the relationship between range and TOF. Applying this to the bandwidth of the experimental matched filter output signal, $\Delta f = 15 \text{ kHz}$ (from Table 8.1), yields

$$\Delta r = \frac{c \cdot \Delta t}{2} = \frac{c}{4 \cdot \Delta f} = 5.8 \text{ mm.} \quad (8.15)$$

8.8.5 Summary of Resolution Results

Three different measures of resolution were computed. These are summarised in Table 8.2. The best is the theoretical system resolution, which makes the most assumptions. The frequency width resolution makes fewer assumptions, while the time width resolution is the closest to that expected in experimental work, as suggested by Rihaczec [118]. There is a variation of a factor of two between these results, suggesting that there is room for improvement in the design to extract the potential resolution provided by the transducers. It should be noted that the MF design squares the spectrum, necessarily reducing the bandwidth and hence the resolution. The reduction in resolution of the practical signals compared to the theoretical system resolution is attributed to the reduced bandwidth of the practical signals compared to the estimated system bandwidth of 27 kHz. However, it should also be noted that the practical signals are affected by air absorption and other effects not considered in the calculation of the theoretical system resolution. A list of the possible effects is given in Section 4.6.3.

The range resolution is measured experimentally in Section 13.3. The result given here for the frequency width resolution of 5.8 mm is remarkably close to the experimental result.

8.9 Software Implementation

The overview diagram of Figure 1.2 indicates that two different matched filter sonar systems are investigated in this thesis. The difference between the two systems lies chiefly in the envelope extractor stage shown in Figure 8.3. The sonar system using the [quadrature demodulation](#) envelope extractor is referred to as the MF-QD sonar, while the sonar system using the unsynchronised demodulation scheme is referred to as the MF-UD sonar.

The software interface for the matched filter sonar system is shown in Figure 8.17 .

Most of the settings are common to both sonar systems. The initial stage in both cases is the matched filter system, corresponding to the boxes marked Barker coded transmit signal and Barker code matched filter in Figure 8.3. The initial stage, in Figure 8.17, shows the setting for the carrier frequency, f_c , the number of sine-wave cycles used to encode one bit of the Barker code, the amplitude of the transmit signal as supplied to the transmit amplifier, the duration of the capture, and the Barker code.

Processor stage 1, shown in Figure 8.17, performs the envelope extraction, peak finding, peak validation and peak interpolation steps of Figure 8.3. This stage has settings for the noise threshold, the loss used to compute the [SSR](#) relative to the ideal value for the code specified in the initial stage, and the width of the window used to check that a peak dominates over the local area, before performing sidelobe suppression. Additionally, the type of envelope extractor is set by a pop-up menu, and the filters used for the [UD](#) method are set here by loading from a user-selected file.

Processor stage 2 shows the conversion of time of flight information to distance of flight information using the speed of sound calculated by the atmospheric sensor, shown in Figure 8.3. This stage will be discussed further in Chapter 12.

The “Finish Here” radio button on each stage determines where processing stops, and hence the type of data which is returned as the result. For example, if only the initial stage is run, the raw matched filter output is plotted.

8.10 Conclusion

This chapter has described the principles governing the design and operation of a matched filter sonar system, and the implementation of a Barker code matched filter sonar system. The matched filter allows the signal energy to be increased, thereby improving the range performance of the sonar system, without sacrificing range resolution. The increased signal energy is realised by the processing gain of the matched filter, which compresses the signal along the time axis into an impulse. The matched filter output must be demodulated to provide a smooth envelope signal for precise target detection.

The matched filter transforms the target detection task from a thresholding problem to a peak finding problem, as the matched filter output produces peaks at the onset of each echo. The peaks are subsequently located, validated, and interpolated to provide the time of flight and echo amplitude as the output.

The [UD](#) technique was seen, in Figure 8.14, to provide a smoother envelope signal than the [QD](#) method, shown in Figure 8.9. With hindsight, the [QD](#) output could also have been smoothed with a

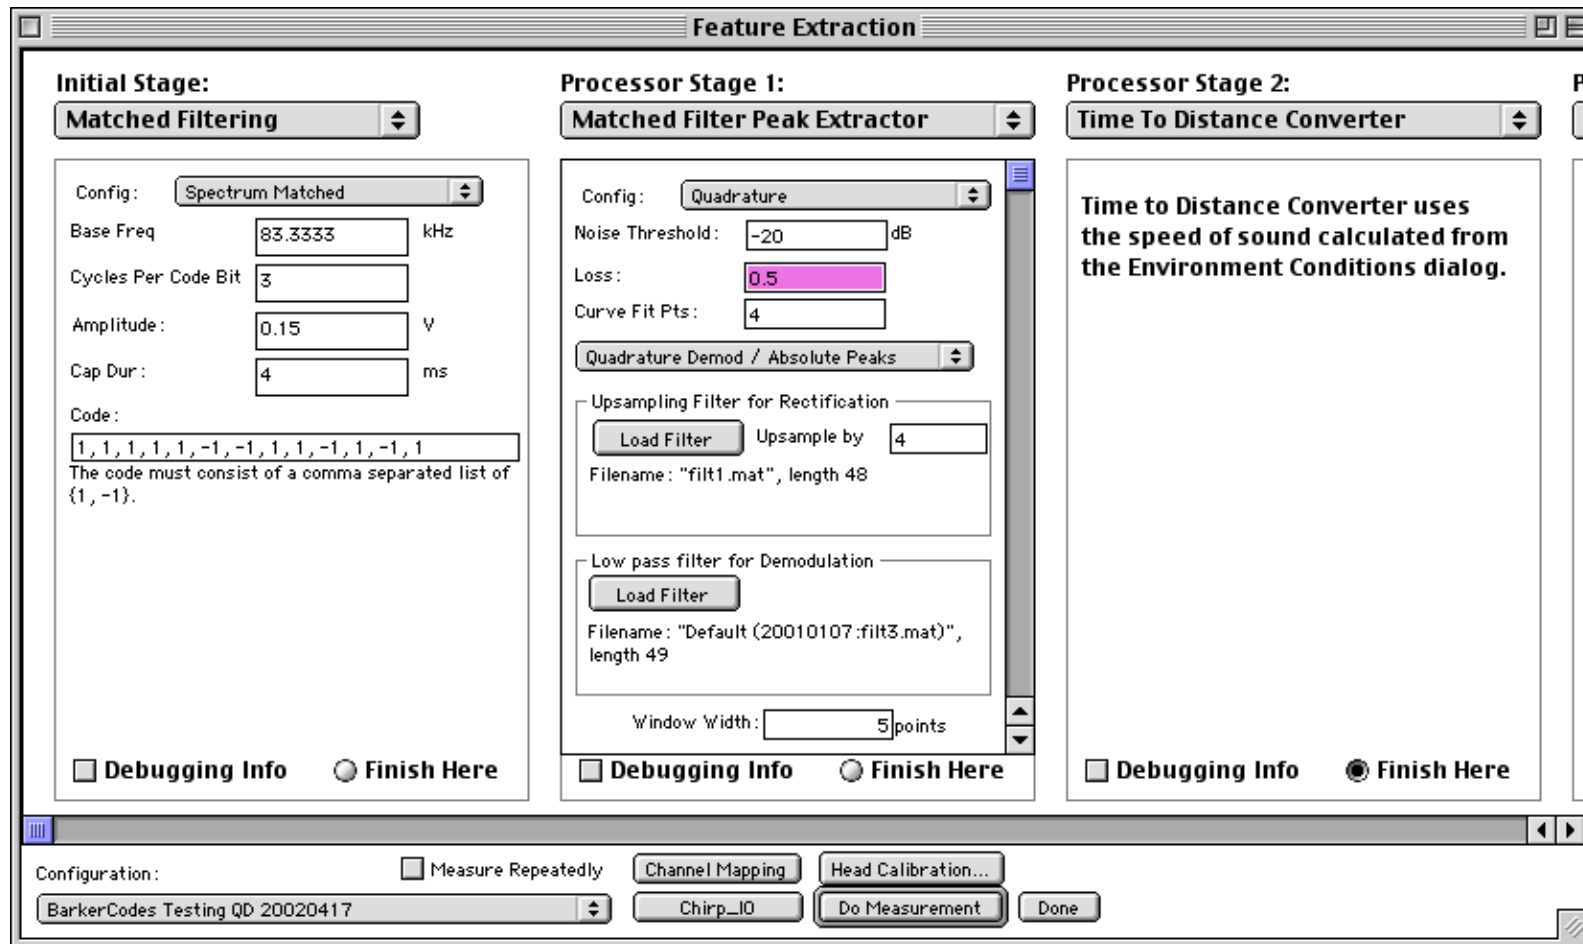


Figure 8.17: Settings for the matched filter quadrature demodulation (MF-QD) sonar system. The matched filter unsynchronised demodulation (MF-UD) sonar system uses the same dialogue. A pre-configured MF-UD sonar system may be found by selecting BarkerCodes Testing AM 20010108 from the bottom left configuration menu. There are three differences in the settings. The first is that the Quadrature Demod / Absolute Peaks control is changed to Rectify and Filter. The second is that the window width is changed to 10 points, and the third is that the Curve Fit Points is changed from 4 to 6.

filter, resulting in a more efficient demodulation process. However, in the current implementation, the UD output is expected to produce better results. The differences between these two methods are measured in Chapters 13 and 14, where it is demonstrated that the UD method is more robust and precise than the QD method.

This chapter also demonstrated how the time-bandwidth product may be applied to calculate the theoretical limit of resolution of the system. Three different methods were applied:

1. the system resolution, independent of whether an MF or CTFM sonar is used, was calculated using the bandwidth of the transducers as an estimate of the system bandwidth.
2. the resolution was estimated from the time-width of the matched filter output signal.
3. the time-bandwidth product relationship was applied to the matched filter output signal bandwidth to compute the theoretical resolution.

These results are compared to experimental results in Section 13.3. However, Rihaczec's [118] comments on resolution are that it is best assessed by examining the width of the autocorrelation signal in the axis where the detector operates. The results seen in Section 8.8 uphold Rihaczec's comments in the case of application to the signal bandwidth at the detector. However, the use of Gabor's uncertainty principle provides acceptable results for computing the range resolution.

The time-bandwidth product of each of the transmit, receive and matched filter output signals were computed, which showed that the matched filter acts as a pulse compressor while largely preserving the bandwidth. The results of this computation clarify the effect of the matched filter processor, and also the effect of the signal path upon the bandwidth.

8.10.1 Acknowledgements and Contributions

The work described by this chapter builds upon that done by Peremans [107, 108, 109], which was itself based upon standard radar techniques as described by Berkowitz [11]. The signal was modified by this author to suit the bandwidth of the Kay transducers used in these experiments. The application of the unsynchronised demodulation to this problem is a contribution of this author, as was the software development of the signal processing and user interface to drive a MF system with either QD or UD envelope extraction. The computation of the time-bandwidth product of the various signals as they pass through the sonar system is also the contribution of this author.

Chapter 9

Implementation of CTFM

The literature review of Chapter 2 indicated that the [Continuous Tone Frequency Modulated \(CTFM\)](#) sonar system had potential for application to a *robotic agent* in addition to the [Matched Filter \(MF\)](#) system. Further investigation in Chapter 3 revealed that two types of CTFM sonar were available: the single demodulation system and the dual demodulation system. As the dual demodulation system provides a demodulated output with no blind time, it was selected for application in this work.

An overview diagram of the complete CTFM sonar system is shown in Figure 9.1. This chapter describes the implementation of the signal processing of the dual demodulation CTFM sonar system described in Chapter 3. The description of the three different spectral analysers and their associated peak detectors is left to Chapters 10 and 11.

The CTFM signal processing was implemented digitally using software. As the signals have different frequencies and bandwidths, a multi-rate signal processing system was implemented, as shown in Figure 9.2. This figure describes only one channel of the sonar system shown in Figure 9.1, and stops at the point of the tonal echo signal, leaving the other elements to the following chapters. The sampling rates of the different parts of the system are marked by areas of different background shading, with approximate sampling rates as marked.

The chapter first describes, in summary form, the parameters used for the various parts of the system, as all of the parts are inter-related. Secondly, the selection of the sampling rates of the various sections shown in Figure 9.2 is explained. Thirdly, each block shown in Figure 9.2 is described with the processed signal being handed from one component to the next. Fourthly, the [time-bandwidth product](#) theory is applied to the signals and the theoretical resolution of the system is calculated.

9.1 System Parameters

The system operation is specified by several parameters and their dependent variables. The function of most of these variables and the relationships between them are described in Chapter 3.

The parameter values which were used in this system are shown in Table 9.1. The original parameters are chosen by the user to inter-operate with the properties of the transducer and of the Fourier transform. Various minor adjustments are made to the original parameters in Sections 9.2 and 9.3 to ensure that the transmit and receive systems are synchronised, and that the transmit signal is free of glitches.

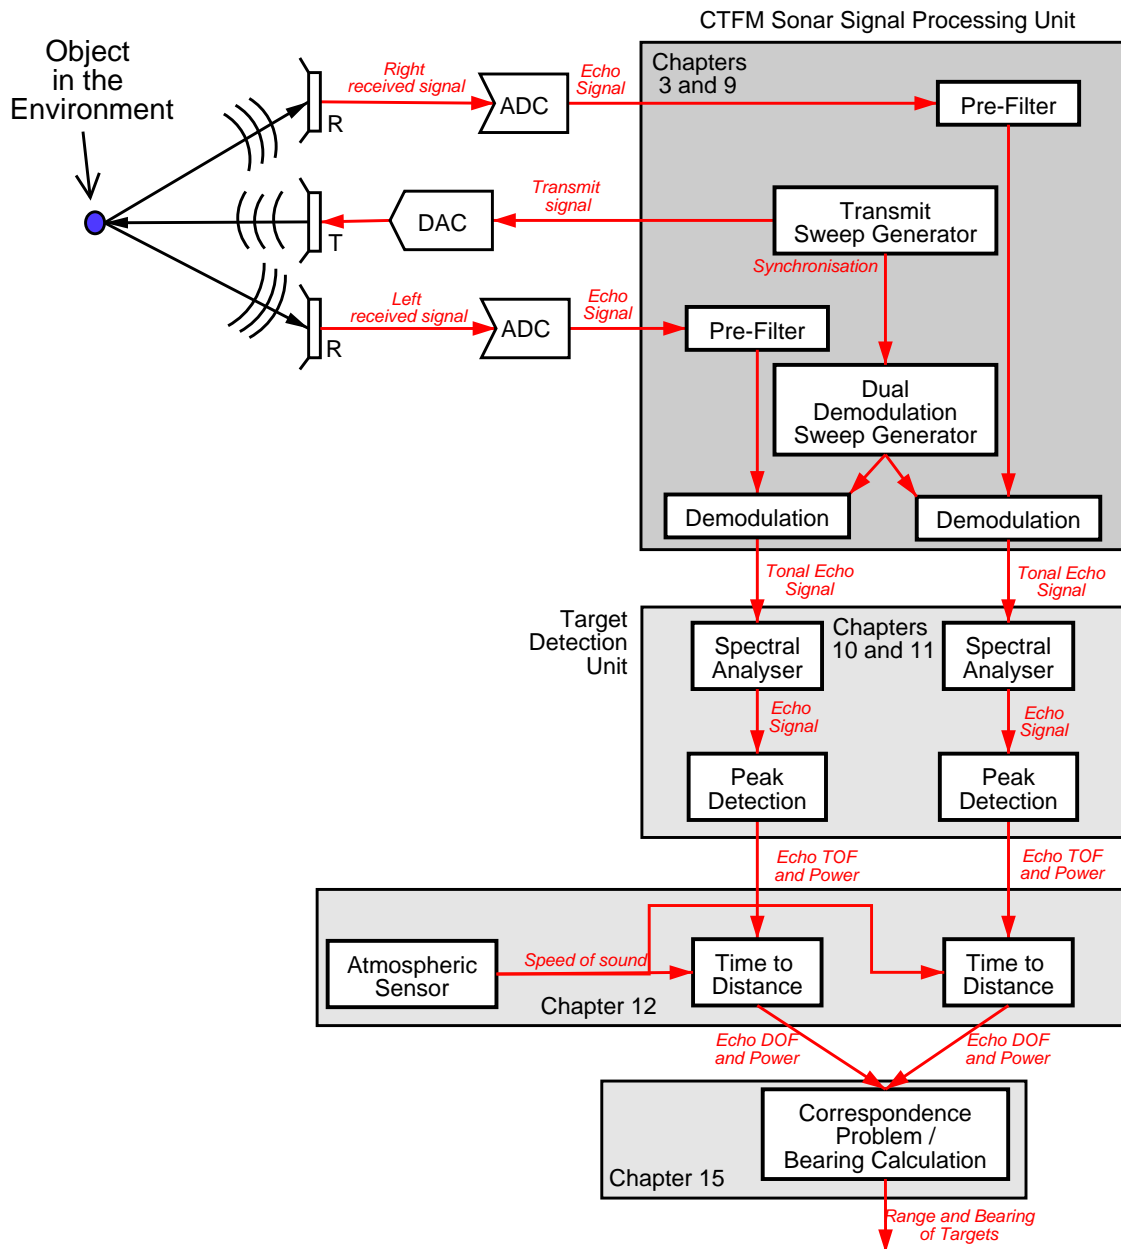


Figure 9.1: Overview of the CTFM sonar system.

Table 9.1: Parameters used in this implementation of dual demodulation. The adjusted values are explained in Sections 9.2.3 and 9.3. The items marked ‘-’ are not adjusted.

Description	Variable	Original Value	Adjusted Value	Adjustment Equation
Transmit sweep start freq	f_0	100 kHz	-	-
Transmit sweep end freq	f_1	50 kHz	49.9984 kHz	(9.12) and (9.13)
1st demod sweep start freq	f_{00}	139.6552 kHz	-	-
Sweep time	T_{sw}	160 ms	159.935 ms	(9.9)
Maximum range ratio	m_{rr}	0.068	-	-
Speed of Sound	c	346 ms ⁻¹	-	-

The sweep frequency range was selected to be compatible with the most powerful operating region of the transducer, as measured in Section 6.7. The value used for the first demodulation sweep start frequency, f_{00} , is selected for compatibility with the downsampling filter, which will be described in Section 9.7.3. The selection of the sweep time, T_{sw} , and the maximum range ratio, m_{rr} , will be discussed in Section 10.2.4.

In addition to the parameters of Table 9.1, we require the frequencies f_{mind} and f_{maxd} corresponding to the minimum, 0, and maximum, d_{max} , distances of flight of the sonar system. These may be computed by inserting 0 and d_{max} into (3.44):

$$f_{mind} = f_{00} - f_0 \quad (9.1)$$

$$f_{maxd} = f_{00} - f_0 + \frac{\mu d_{max}}{c}, \quad (9.2)$$

where μ is the sweep rate. The maximum DOF, d_{max} , is computed using (3.24). The frequencies found by (9.1) and (9.2) define the minimum and maximum DOF frequencies respectively which will be seen at (5) in Figure 9.2.

The resulting values of the dependent variables are shown in Table 9.2. While the speed of sound and the actual maximum range may change slightly, the minimum and maximum DOF frequencies remain fixed. The dependent variable d_{max} shown in Tables 9.1 and 9.2 is the maximum DOF, so the maximum range is about 1.88 m, being half d_{max} .

Table 9.2: Dependent variables used in our implementation of dual demodulation. Items marked ‘-’ do not change after adjustments.

Description	Variable	Original Value	Adjusted Value	Defining Equation
Sweep rate	μ	$-312\,500\text{Hz}^2$	$-312\,636.8\text{Hz}^2$	(3.2)
Sweep bandwidth	B	50 kHz	50.001 6 kHz	(3.26)
1st demod sweep end freq	f_{01}	89.655 2 kHz	89.655 363 3 kHz	(3.29)
2nd demod sweep start freq	f_{10}	89.655 2 kHz	89.655 363 3 kHz	(3.28)
2nd demod sweep end freq	f_{11}	39.655 2 kHz	39.655 206 6 kHz	(3.30)
Maximum DOF	d_{max}	3.764 48 m	3.762 95 m	(3.24)
Minimum DOF freq	f_{mind}	39.655 2 kHz	-	(9.1)
Maximum DOF freq	f_{maxd}	36.255 2 kHz	36.255 07 kHz	(9.2)

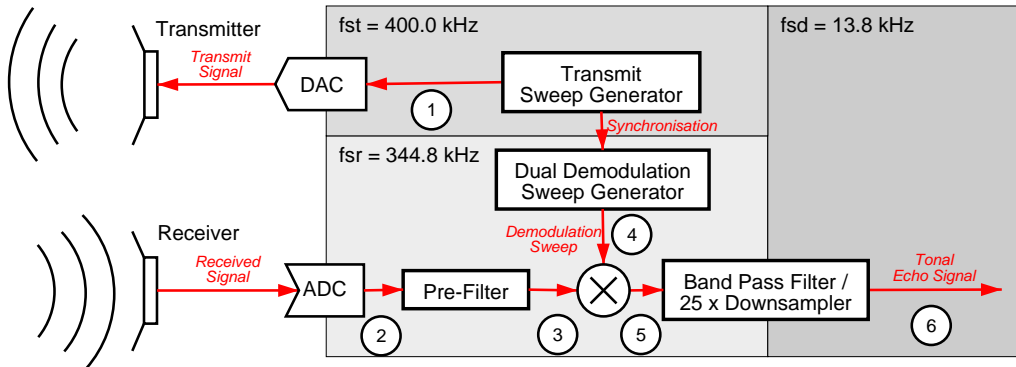


Figure 9.2: Digital implementation of the signal processing for the dual demodulation CTFM sonar system. The numbers in circles identify the important signals for reference from the text. The circle containing \times denotes time domain multiplication of two sampled series. The background blocks delimit parts of the system which use different sampling rates, being the transmit sampling frequency, f_{st} , the receive sampling frequency, f_{sr} , and the downsampled sampling frequency, f_{sd} .

9.2 Sampling Rates

High transmit sampling rates do not place any computational burdens on the CPU as the output signal is assumed to be constant and this is synthesised only once. However, the received signal needs to have the lowest possible sampling rate so that computational burdens are minimised. This minimisation also maximises the length of the echo signal that can be stored in the capture card FIFO. The lower limit of the receive sampling frequency is defined by the highest frequency present in the received signal and intermediate stages of processing.

Thus the transmit and receive signals use different sampling frequencies. As the sweep time, T_{sw} , must be precisely the same in both the transmit and receive systems, the transmit and receive sampling rates must be related by a rational number.

The synchronisation is achieved by utilising the NuBus clock, $f_{bus} = 10\text{MHz}$, which is shared by both cards. Both the transmit, f_{st} , and receive, f_{sr} , sampling rates are generated by division of the bus clock frequency. Let n_{dt} be the number of f_{bus} cycles per cycle of the transmit clock, f_{st} , and let n_{dr} be similarly defined for the receiver clock, f_{sr} . Let n_{st} be the number of samples transmitted in one sweep period, T_{sw} , and n_{sr} be the number of samples received in one sweep period. Let T_{swo} be the original sweep period before any adjustments.

There are three steps to calculating a consistent set of sampling rates and sweep times.

1. Determine the receive sampling rate, f_{sr} .
2. Determine the transmit sampling rate, f_{st} .
3. Determine the adjusted sweep period, T_{sw} .

9.2.1 Receive Sampling Rate

The receive sampling frequency f_{sr} is a design parameter, as it is used to design all of the filters and cannot be adjusted. The receive sampling rate is related to the bus frequency by the receive clock division, n_{dr} , such that

$$f_{sr} = \frac{f_{bus}}{n_{dr}}. \quad (9.3)$$

Thus n_{dr} is a design parameter determining the receive sampling rate f_{sr} .

The initial sampling rate, f_{sr} , was chosen to be $10\text{MHz}/29 = 344.8\text{kHz}$ as it satisfies the aliasing requirements outlined in Section 9.6. However, it may be possible to lower this frequency even further, as there is significant clearance between the aliased frequencies and the frequency range of interest (Figure 9.10a).

9.2.2 Transmit Sampling Rate

The transmit sampling rate, f_{st} , is calculated to be the highest available sampling rate which is not greater than the original transmit sampling rate, f_{sto} . The original rate is calculated from the original length of the sweep, T_{swo} , and the size of the transmit FIFO, by

$$f_{sto} = \frac{65536}{T_{swo}}.$$




Figure 9.3: Diagram showing how the sweep time T_{sw} is adjusted to ensure that the transmit and receive sweeps are exactly co-periodic. The number of samples indicated is smaller than used in a real system for clarity of the diagram. The last sample of a repeating sequence, such as those at T_{sw} , is deemed to belong to the next repetition of the sequence, so is drawn in grey instead of black. The counts on the braces indicate the number of intervals enclosed. The number of samples enclosed by the braces are equal to the number of enclosed intervals when the end sample is omitted. The largest number of receive samples falling within the original sweep time, T_{swo} , is denoted n_{sr1} . A similar quantity is defined for the transmit samples (not shown), denoted n_{st1} .

Then the transmit clock divisor, n_{dt} , is calculated by finding the next highest available frequency to the desired transmit sampling rate f_{sto} . Using $\lceil \cdot \rceil$ to denote upward-rounding or the ceiling operation (that is, $\lceil 5.1 \rceil = 6$), this may be written

$$n_{dt} = \left\lceil \frac{f_{bus}}{f_{sto}} \right\rceil, \quad f_{st} = \frac{f_{bus}}{n_{dt}} = 400.0\text{kHz}, \quad (9.4)$$

providing the adjusted value for the transmit sampling rate, as shown in Figure 9.2.

9.2.3 Adjusted Sweep Time

Now that the receive and adjusted transmit sampling rates have been chosen, the sweep time T_{sw} is chosen to be the largest interval smaller than the original sweep time, T_{swo} , which obtains synchronised sample end-points, as shown in Figure 9.3. The value for the user-chosen original sweep time, T_{swo} , is chosen in Section 10.2.4. There are three steps to the calculation of the adjusted sweep time T_{sw} .

The first step is to calculate the number of samples falling within the original sweep period, T_{swo} . This is calculated by dividing the sweep period by the sample period and rounding down to the nearest integer, denoted by the $\lfloor \cdot \rfloor$ operator (that is, $\lfloor 5.9 \rfloor = 5$). This yields the number of transmit samples, n_{st1} , falling within the original sweep time, T_{swo} ,

$$n_{st1} = \lfloor f_{st} \cdot T_{swo} \rfloor = \left\lfloor \frac{f_{bus}}{n_{dt}} \cdot T_{swo} \right\rfloor, \quad (9.5)$$

and the similarly defined number of receive samples, n_{sr1} , falling within the original sweep time, T_{swo} ,

$$n_{sr1} = \lfloor f_{sr} \cdot T_{swo} \rfloor = \left\lfloor \frac{f_{bus}}{n_{dr}} \cdot T_{swo} \right\rfloor, \quad (9.6)$$

falling within the original sweep period. The lengths of these sequences are not co-periodic.

The second step is to find the greatest number of transmit samples which produces a sweep period which is co-periodic with the number of receive samples,

$$n_{st} = n_{dr} \cdot \left\lfloor \frac{n_{st1}}{n_{dr}} \right\rfloor = n_{dr} \cdot \left\lfloor \frac{f_{bus}}{n_{dr} \cdot n_{dt}} \cdot T_{swo} \right\rfloor. \quad (9.7)$$

The equivalent calculation in terms of the number of receive samples is

$$n_{sr} = n_{dt} \cdot \left\lfloor \frac{n_{sr1}}{n_{dt}} \right\rfloor = n_{dt} \cdot \left\lfloor \frac{f_{bus}}{n_{dr} \cdot n_{dt}} \cdot T_{swo} \right\rfloor. \quad (9.8)$$

The final step is to calculate the adjusted sweep period, T_{sw} . This may be calculated from the transmit cycle period,

$$T_{sw} = \frac{n_{st}}{f_{st}} = \frac{n_{dr} \cdot n_{dt}}{f_{bus}} \cdot \left\lfloor \frac{f_{bus}}{n_{dr} \cdot n_{dt}} \cdot T_{swo} \right\rfloor, \quad (9.9)$$

or from the receive cycle period,

$$T_{sw} = \frac{n_{sr}}{f_{sr}} = \frac{n_{dr} \cdot n_{dt}}{f_{bus}} \cdot \left\lfloor \frac{f_{bus}}{n_{dr} \cdot n_{dt}} \cdot T_{swo} \right\rfloor. \quad (9.10)$$

Comparison of (9.9) and (9.10) reveals that they are identical, as required. Thus the transmit and received signals are exactly co-periodic, with period T_{sw} , as shown in Figure 9.3, even though they have different sampling rates, as shown in Figure 9.2.

9.3 Transmit Sweep Generator

The transmit sweep generator, shown in Figure 9.2, synthesises the transmit signal, which was described in Section 3.2, culminating in the definition of the frequency sweep formula (3.8).

The sweep end frequency, f_1 , must be adjusted to ensure that the transmit signal is continuous at T_{sw} , where one sweep ends and the next begins. If adjustments are not made, then the signal may be discontinuous or have a large discontinuity in the derivative, as shown in Figure 9.4. The discontinuity may be avoided by requiring that the phase of the transmit signal, $\phi_t(t)$, defined in (3.7), complete an integer number of cycles, l . Thus it must obey the condition

$$\phi_t(T_{sw}) = 2\pi l. \quad (9.11)$$

This may be accomplished by slightly adjusting the (already adjusted) sweep period, T_{sw} , or the sweep end frequency, f_1 . Doing so results in only a small discontinuity in the derivative of the signal (and in the derivative of the signal phase, see (3.16)) at the join between periods, instead of a discontinuity in the signal itself.




Figure 9.4: Example of a frequency sweep with a mis-matched join between two consecutive sweeps. The time scale is in units of T_{sw} . We require that the signal be continuous, and have a minimum discontinuity in the derivative. The signal shown here has a discontinuity in both its value and its derivative. This plot uses very low frequencies for clarity.

The sweep period, T_{sw} , was adjusted to obtain co-periodic transmit and receive signals in Section 9.2.3, and cannot be changed again. Therefore the sweep end frequency, f_1 , is adjusted instead. Substituting (3.7) into (9.11) and solving for f_1 yields

$$f_1 = \frac{2l}{T_{sw}} - f_0. \quad (9.12)$$

That is, the adjusted sweep end frequency, f_1 , is a simple function of l . The integer l is found by solving (9.12) and selecting the nearest integer, yielding

$$l = \text{round} \left(\frac{T_{sw}}{2} (f_0 + f_1) \right). \quad (9.13)$$

The calculation procedure is to first use the nominal sweep end frequency, f_1 , to calculate l using (9.13), and then to find the adjusted sweep end frequency, f_1 , using (9.12). The sweep start frequency, f_0 , and the sweep time, T_{sw} , are not adjusted by this procedure. The final values used, after all adjustments have been made, are shown in Tables 9.1 and 9.2.

After the transmit sweep is digitally synthesised it is loaded into the generator card FIFO and transmitted as previously described in Section 9.2. This is the signal marked ① in Figure 9.2, and is shown in Figure 9.5.

Note that the signal shown in Figure 9.5 is the ideal digitally synthesised signal before conversion to an analogue signal. When it is converted and subsequently transmitted, it has a frequency response imposed upon it by the transducer. The on-axis response of the transducer was shown in Figure 6.13. If the target is off-axis, then we may determine the bearing-dependent filter response by referring to Figure 6.15a. Thus the signal which emerges from the transmitter is a modified form of that shown in Figure 9.5. The response at the extremities of the bandwidth is reduced compared to the response at the transducer's most efficient frequency, which is 86kHz (Figure 6.13).




Figure 9.5: Spectrogram of the digitally synthesised transmit signal. Signal ① in Figure 9.2. Time measured in seconds, frequency in Hz. The colour scale is in decibels relative to 1 Volt amplitude. The colour map comes from a suggestion by Rogowitz and Treinish [119]. The colour change is used to indicate the noise floor. All of the spectrograms in this chapter use the same scale for the colour map for easy comparison.




Figure 9.6: Spectrogram of the digitally synthesised Dual Demodulation signal. Signal ④ in Figure 9.2. Time measured in seconds, frequency in Hz. The colour scale is in decibels relative to 1 Volt amplitude.

9.4 Dual Demodulation Sweep Generator

The dual demodulation sweep generator, shown in Figure 9.2, operates at the receive sampling rate, f_{sr} . It pre-computes the sweeps defined by (3.36) and (3.37). The values used for the various parameters are shown in Tables 9.1 and 9.2. The two signals are added together and are stored as a single signal for use by the demodulation system, which operates by multiplying the received signal with the demodulation signal in the time domain, as described in Section 3.3.2.

We have now obtained the signal marked ④ in Figure 9.2, which is shown in Figure 9.6.

9.5 Pre-Filtering

The received signal pre-filter, shown in Figure 9.2, operates upon the received echo signal, marked ② in Figure 9.2 and shown in Figure 9.7, and removes any noise outside of the range 50-100kHz. Any signals outside this range are either Doppler shifted echoes or noise. Significant environmental noise in the range 0-20kHz may be expected due to talking, music, machinery etc. Evaluation of the system showed that there was significant noise near 0Hz. If such frequencies are present in the signal as it is demodulated, the output of the demodulator will contain some potentially interfering signals. In particular, a signal at 0Hz will cause the demodulation signal to be reproduced in the demodulator output. This interferes with the demodulated echoes. The easiest way to remove

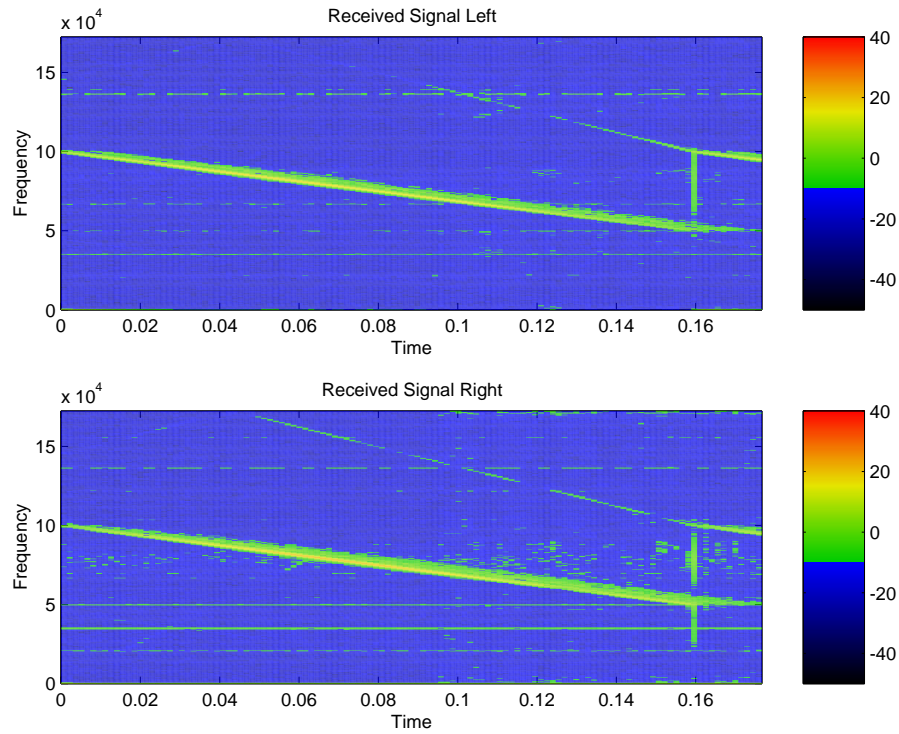


Figure 9.7: Spectrogram of a typical received signal. Signal ② in Figure 9.2. Time measured in seconds, frequency in Hz. The colour scale is in decibels relative to 1 Volt amplitude.

this possibility is to pre-filter the received signal. The filter used is shown in Figure 9.8. This is a linear phase **Finite Impulse Response (FIR)** filter (see Oppenheim and Schaffer [105]) designed in Matlab [130] using the `remez` command. The specifications were $f_{sr} = 10\text{MHz}/29$, stop-band from 0 to 30 kHz, passband from 50 kHz to 100kHz, and stop-band from 120 kHz to $f_{sr}/2$. The filter was designed with 51 zeros, and has 52 time domain coefficients.

Care must be taken in implementing the filter to avoid shifting the time of the signal as it passes through. This is because the time of flight information contained in the signal is recorded in the time domain at this stage. It is transformed into frequency when the multiplication shown in Figure 9.2 is carried out.

The filter could have alternatively been implemented using an **Infinite Impulse Response (IIR)** filter as this type of filter can achieve the required filtering function using fewer terms than a **FIR** filter. However, such filters do not have linear phase characteristics, which are required to preserve the temporal integrity of the signal.

9.5.1 Cancelling Filter Delay

The filter implementation cancels the **FIR** filter delay by compensating the start offset of the signal by the group delay of the filter. Proakis and Manolakis [116, Section 4.5.1] give a formula for the group delay $\tau_g(\omega)$ in terms of radian frequency, $\omega = 2\pi f$, as

$$\tau_g(\omega) = -\frac{1}{f_{sr}} \frac{d\Theta(\omega)}{d\omega}, \quad (9.14)$$




Figure 9.8: The filter used to pre-filter received echoes before being demodulated. This filter was designed for use at $f_{sr} = 344.82$ kHz.

where $\Theta(\omega)$ is the filter phase characteristic, which is different for each filter. For a linear phase filter with a symmetric impulse response such as that used above, the filter phase is (see Proakis and Manolakis [116, Section 8.2.1])

$$\Theta(\omega) = \begin{cases} -\omega \left(\frac{M-1}{2} \right) & \text{if } H(\omega) \geq 0 \\ -\omega \left(\frac{M-1}{2} \right) + \pi & \text{if } H(\omega) < 0 \end{cases} \quad (9.15)$$

for both odd and even filter lengths M . Noting that

$$\frac{d\Theta(\omega)}{d\omega} = \frac{M-1}{2} \text{ for all } H(\omega), \quad (9.16)$$

we can substitute (9.16) into (9.14) to calculate the delay time of the symmetrical filter to be

$$\tau_g(\omega) = \frac{1}{f_{sr}} \frac{M-1}{2} \text{ for all } H(\omega). \quad (9.17)$$

This frequency independent nature of the group delay is a property of linear phase filters. This important result allows us to exactly cancel out the delay of the pre-filter.

While large FIR filters may be efficiently computed using the Fast Fourier Transform (FFT) to perform the convolution (see Proakis and Manolakis [116, Section 5.3]), this implementation used a time domain convolution. The relative benefits of using the FFT filter algorithm were not evaluated, due to the additional complexity of the block overlap implementation required.

We have now obtained the signal marked ③ in Figure 9.2, which is shown in Figure 9.9.

9.6 Multiplication

The time domain multiplication step, shown in Figure 9.2, is the first part of a two step demodulation procedure. The second step is to filter the signal to remove the unwanted parts, which will be described in Section 9.7. The demodulation procedure operates as described in Section 3.3.2.

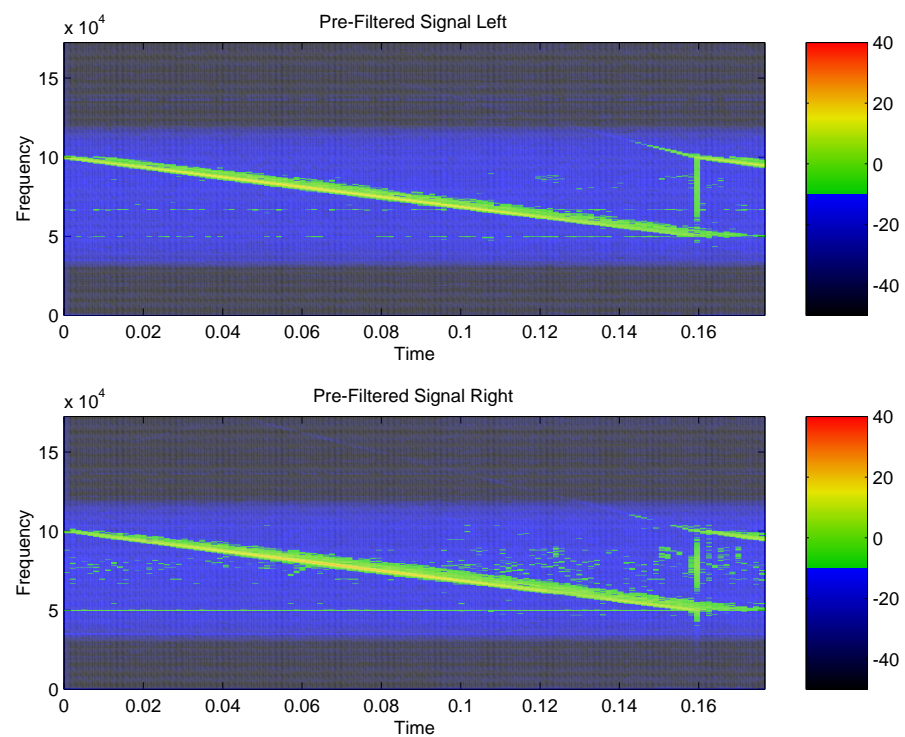


Figure 9.9: Spectrogram of a typical received signal after pre-filtering. Signal ③ in Figure 9.2. Time measured in seconds, frequency in Hz. The colour scale is in decibels relative to 1 Volt amplitude.

The multiplication input signals, labelled ③ and ④ in Figure 9.2, are sampled at $f_{sr} = 10\text{MHz}/29 = 344.8\text{kHz}$. The signals are then digitally multiplied, sample by sample. The multiplication output contains sum frequencies which exceed the Nyquist rate, $f_{sr}/2$, for some of the sweep period. Aliasing causes these frequencies to be reflected about the Nyquist rate, as shown in Figure 9.10a. Note that the lower difference frequency has also been aliased, having been reflected around 0Hz.

The system has been designed to include aliasing because it allows a lower sampling frequency to be used, thus lowering the computational burden of the signal processing. The aliasing is acceptable so long as it does not overlap the frequencies of interest. In this design, the frequencies of interest after demodulation are between f_{mind} and f_{maxd} , corresponding to zero to maximum DOF. Table 9.2 shows that f_{mind} and f_{maxd} are 39.7kHz and 36.2kHz respectively, rounding toward the outside of the range. Figure 9.10a shows that no aliased signals enter this region at maximum range. Graphical analysis of the frequencies of the aliased signals shows that they remain outside the frequencies of interest, as required.

We have now obtained the signal marked ⑤ in Figure 9.2, which is shown in Figure 9.10b. The signal is composed of four tones, two of which change over time. It is now sent to the downsampler.

9.7 Band Pass Filter and Downsampler

We now apply a carefully designed downsampling system which filters out the signals that are not required and simultaneously reduces the sampling rate while retaining the signals of interest. Since the downsampling or decimation affects the signal of interest, and this effect can be mitigated by correctly choosing the band-pass filter, it is instructive to describe the decimation before the band-pass filter. In most implementations, including the system described by this thesis, these processes are implemented in a single step.

The decimation system is required because of the different frequency ranges occupied by the multiplication (Section 9.6) product signals and the demodulated output signals (Section 9.7.6).


As the output of the downsampler system is the final demodulated output signal, this section also treats the problem of computing the target range from the signal frequency (and vice versa) in Section 9.7.4, and examines some of the trade-offs chosen when designing such a demodulation and decimation system, including the need for a decimation system, in Section 9.7.5. Finally, the output signals are described in Section 9.7.6.

9.7.1 Decimation

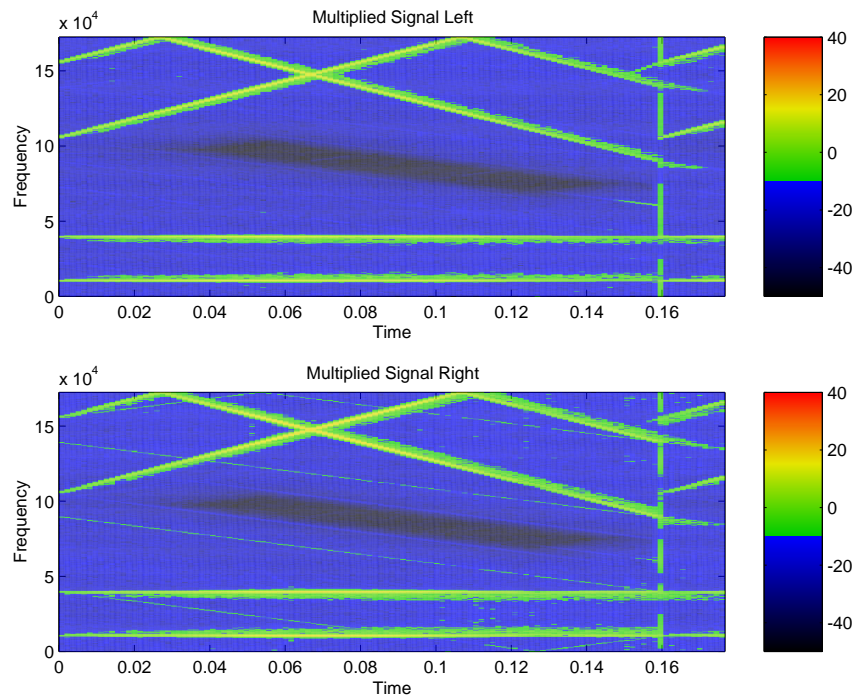
The process of decimation by N_d consists of retaining one data point and discarding the other $N_d - 1$ data points from each block of N_d points. This procedure causes the sampling frequency to be reduced to

$$f_{sd} = \frac{f_{sr}}{N_d}. \quad (9.18)$$

It also reduces the Nyquist frequency, thereby causing many frequencies to be aliased. The aliased output frequency, $f_{aliased}$, may be computed for each input frequency, f , and a given sampling



(a) model of signal frequencies after multiplication



(b) Experimental spectrogram after multiplication

Figure 9.10: Signal frequencies for digital dual demodulation, after multiplication. See the middle part of Figure 3.6 for analogue comparison. (a) Shows modelled values, while (b) shows experimental data, corresponding to signal (5) in Figure 9.2. There is aliasing of the sum and difference signals around $f_{sr}/2 = 172.4\text{kHz}$ and 0Hz respectively, where the frequency has been folded around the Nyquist frequency $f_{sr}/2$. The frequency range of interest, corresponding to zero and maximum DOF, is between 39.7kHz and 36.2kHz on the vertical axis, and includes the constant, continuous tone at 37kHz . The colour scale is in decibels relative to 1 Volt amplitude.

frequency, f_s , using a function given by Jacobs [38]:

$$f_{\text{aliased}} = \left| f \cdot \left\{ 1 - f_s \cdot \text{round} \left(\frac{f}{f_s} \right) \right\} \right|. \quad (9.19)$$

In the case of decimation, the sampling frequency f_s represents the decimated sampling rate, f_{sd} . Equation (9.19) is illustrated for the case of decimation by four in Figure 9.11.

The decimation causes the four frequency bands to be transformed into a single frequency band. Thus signals having frequencies of $3f_s/32$, $5f_s/32$, $11f_s/32$ and $13f_s/32$ within the input signal are all aliased to the frequency $3f_s/32 = f_{sd}/8$ in the output signal. Thus the four input signals become indistinguishable. This drawback may be overcome by filtering to remove frequencies from all but one band before decimation is applied. Then frequencies in the output signal space will have a unique reverse mapping to frequencies in the input signal space.

The aliasing occurs in a predictable way. Odd bands retain their forward orientation, but the even bands are frequency reversed by the decimation process. This must all be taken into account when calculating the frequencies corresponding to minimum and maximum range.

9.7.2 Bandpass Filtering

The problem of selecting a unique passband when downsampling a signal, as examined in the previous section, is typically solved using a low-pass filter to select frequencies in the first band (see Figure 9.11). However, it is equally valid to use a band-pass filter to select some other band, for example band 2. The frequencies contained within that band will be frequency reversed by the decimation procedure, but this may be taken into account.

The filter calculation procedure may be optimised in the case where a FIR filter is used, by only calculating those output points that will not be discarded. This reduces the computational burden of the filtering by the same order as the decimation, N_d . This is discussed in full by Proakis and Manolakis [116, Section 10.2].

9.7.3 Decimation of CTFM Signals

Moving from the decimation by four example to the multiplication output signal (5) in Figure 9.2, the frequency range of interest, being 3.4kHz wide, covers only 0.02 of the frequency range contained in signal (5), which has a Nyquist frequency of 172.4kHz. Therefore the signal may be decimated by $N_d = 50$, obtaining a new Nyquist frequency of 3.448kHz. However, such a scheme requires a filter with infinitely sharp transition regions. It is impossible to build such a filter. Therefore we decimate by $N_d = 25$ instead, which provides an output Nyquist frequency of 6.896kHz, double the required bandwidth. This provides one quarter band both above and below the required output frequency range where the filter transition region may be placed (Figure 9.12). The half band transition region allows for the limitations of a practical filter.

The resulting output signal spectra will be shaped by the downsampling filter as shown in Figure 9.13. This shows the central region containing the expected tones from echoes, and also the garbage bands on either side which contain signals that may have been aliased during the decimation process. These bands may be made narrower, but at the cost of longer filters.

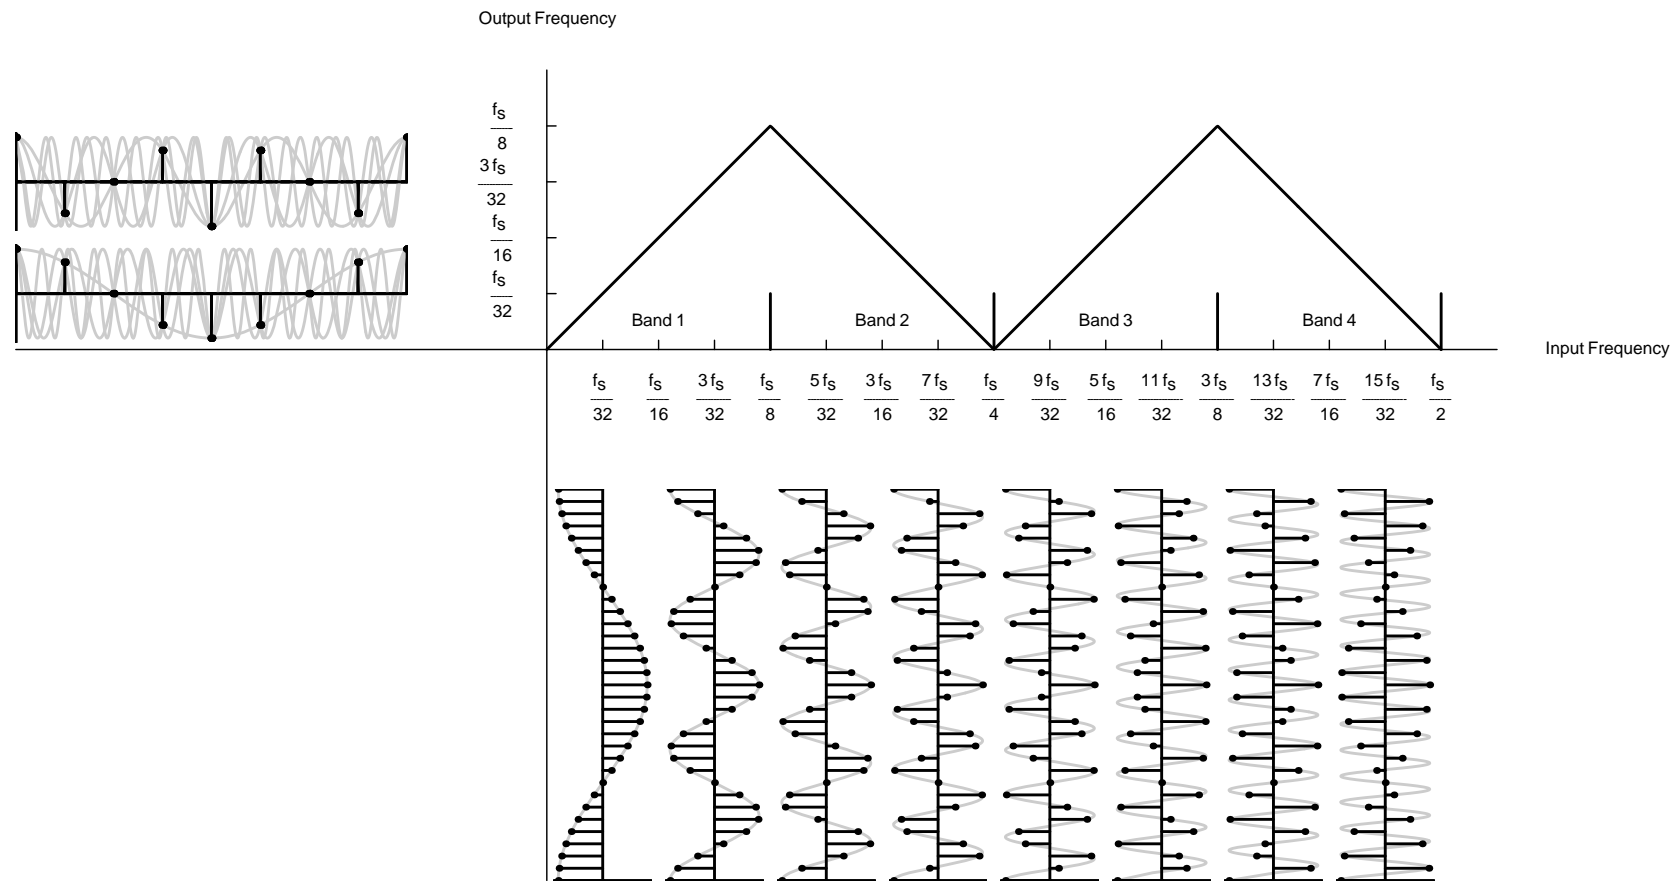


Figure 9.11: Diagram illustrating the aliasing that occurs when a signal is decimated by a factor of four, i.e., by discarding three out of every four samples. The frequency bands shown along the horizontal axis are superimposed into the output frequency space, shown on the vertical axis, according to the triangle function shown. The samples on the signals shown at the left are all that remain after the downsampling has occurred. A close inspection of the grey signal lines reveals that all of the signals pass through all of the remaining sample points. The signals in even bands have their spectra reversed. Any frequencies lying outside of band 1 are aliased into the frequency range $0 \dots f_s/8$.




Figure 9.12: Decimation filter band design. The large tick marks delimit bands, while the minor tick marks delimit the 1/4 and 3/4 band positions. The pass-band is thus half a band wide. The stop-bands end at half a band width from the pass band, allowing one half band width transition region, which will be aliased into 1/4 band of the output frequency space. The output frequency space, illustrated by the vertical axis of Figure 9.11, is only one band wide.




Figure 9.13: Map of the downsampled output frequency space and the range to frequency mapping. The maximum range frequency is just below the corner frequency of the filter. The difference is minimised during the design of the downsampling system.




Figure 9.14: The $25\times$ downsampling filter used after demodulation. This filter has 247 (symmetrical) terms.

The Anti-Aliasing Downsampling Filter

The frequency range of interest falls into filter band 6 when the signal spectrum is divided into 25 bands. A suitable downsampling filter, designed according to the pass-band and stop-bands shown in Figure 9.12, is shown in Figure 9.14. When the filter of Figure 9.14 is used for decimation, it may be regarded as 25 equally spaced frequency bands up to the Nyquist frequency $f_{sr}/2$. The signal is digitally downsampled by $N_d = 25$ using this filter.

The frequency offset of the demodulation signals must be selected so that minimum range corresponds to one of the edge frequencies of the passband of the filter. The width of the filter passband must be wider than the demodulation bandwidth (3.27), in order to pass all of the signal. Thus, the design of the decimation filter is an integral part of the design of the digital demodulation system.

The actual corner frequencies of the decimation filter are shown in Table 9.3. The input sampling rate used was $f_{sr} = 10\text{MHz}/29 = 344827.6\text{Hz}$. The minimum DOF frequency, f_{mind} , has been designed to correspond exactly with the top edge of the passband of the decimation filter, while the maximum DOF frequency, f_{maxd} , remains above the start of the passband.

The frequency range 39.6...36.2 kHz will become 1.7...5.2 kHz after downsampling. The reversal is due to the frequencies of interest lying in band 6, which is even. The downsampled sampling rate becomes

$$f_{sd} = \frac{f_{sr}}{N_d} = 13793\text{Hz}. \quad (9.20)$$

The output of the decimation filter consists of a 1/4 band of unspecified signal, then 1/2 band of demodulated CTFM output, and another 1/4 band of unspecified signal before the Nyquist rate $f_{sd}/2$ is reached.

Table 9.3: Corner frequencies for the decimation filter in Figure 9.14.

Description	Frequency	Formula
End stopband	32 758.6 Hz	$\frac{5-1/4}{N_d} \frac{f_{sr}}{2}$
Start passband	36 206.9 Hz	$\frac{5+1/4}{N_d} \frac{f_{sr}}{2}$
End passband	39 655.2 Hz	$\frac{6-1/4}{N_d} \frac{f_{sr}}{2}$
Start stopband	43 103.4 Hz	$\frac{6+1/4}{N_d} \frac{f_{sr}}{2}$

9.7.4 Frequency and Distance of Flight

One remaining problem is how to calculate the relationship between frequency and distance of flight after decimation has occurred. The effect of decimation with the specified filter is to re-map the frequency space:

$$f_{out} = 3 \cdot f_{sd} - f_{in}, \quad 36\,206.9 \text{ Hz} \leq f_{in} \leq 39\,655.2 \text{ Hz}. \quad (9.21)$$

Combining (3.44) and (9.21), the output frequency for a given distance of flight is found to be

$$f_{out} = 3 \cdot \frac{f_{sr}}{N_d} - \left(f_{00} - f_0 + \frac{\mu d}{c} \right).$$

Recalling the value used for f_{00} (Table 9.1), and using the formula for the upper edge of the filter passband from Table 9.3, this can be simplified to

$$\begin{aligned} f_{out} &= 3 \cdot \frac{f_{sr}}{N_d} - \frac{6 - \frac{1}{4}}{2} \frac{f_{sr}}{N_d} - \frac{\mu d}{c} \\ &= \frac{1}{8} \cdot \frac{f_{sr}}{N_d} - \frac{\mu d}{c}. \end{aligned} \quad (9.22)$$

In this case, noting that μ is negative, we find that an increase in range brings about an increase in frequency.

By choosing an odd band of the decimation filter, instead of the even band described above, an increase in distance will bring about a negative change in frequency. In such a case, the lower edge of the filter passband would be selected for adjusting the frequency f_{00} before decimation.

9.7.5 Design Considerations

Section 9.2.1 mentioned that the selection of the receive sampling rate, f_{sr} , must provide sufficient bandwidth that the aliased components of signal (5) in Figure 9.2 do not enter the pass-band of the filter. Only the frequencies corresponding to minimum and maximum range should fall within the passband of the filter. This condition is satisfied through the selection of the demodulation

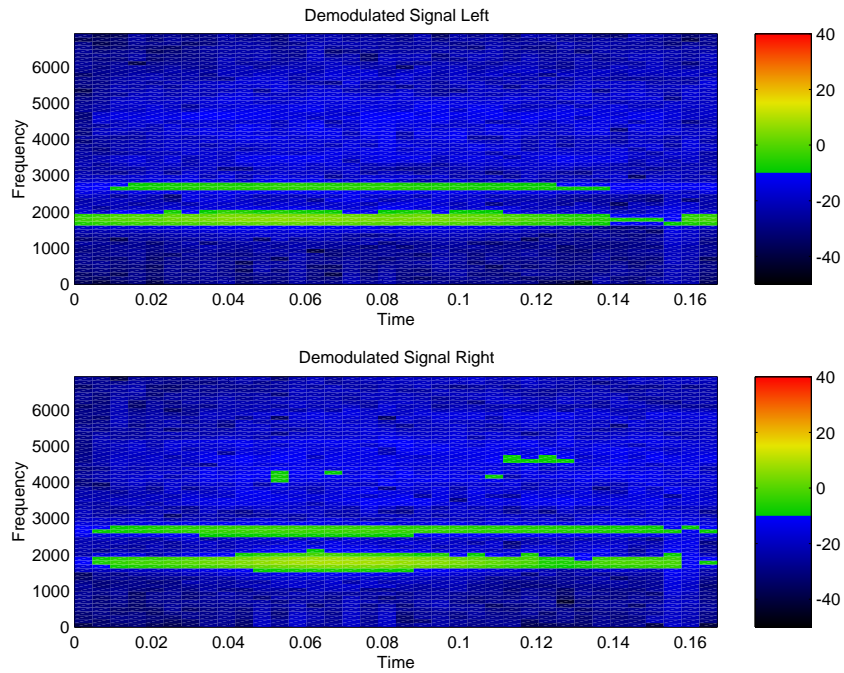


Figure 9.15: Spectrogram of a typical demodulated, tonal echo signal, marked ⑥ in Figure 9.2. Time measured in seconds, frequency in Hz. The colour scale is in decibels relative to 1 Volt amplitude. The lower frequency signal, about 1900Hz, is cross-talk, while the upper frequency, about 2800Hz, is the echo from a 30mm diameter pole target placed 500mm in front of the sonar head. The other green squares are noise or echoes from a weak target.

sweep start frequency, f_{00} . The criterion used for selecting the downsampling number, N_d , is the bandwidth of the output signal compared with the bandwidth required to prevent aliasing during the multiplication stage. The main consideration in the selection of the filter band is to prevent the upper or lower aliased signals (Figure 9.10) of the demodulation multiplication from entering the passband.


Echoes arriving after T_{max} will be demodulated to a frequency outside the primary echo detection frequency band. The the demodulation scheme described in this chapter allows echoes from up to twice the maximum range, d_{max} , to be isolated from the echo detection band.

9.7.6 Output Signals


We have now obtained the tonal echo signal, marked ⑥ in Figure 9.2, which is also shown in Figure 9.15. It is also instructive to examine the tonal echo signal in the time domain, as shown in Figure 9.16.

9.8 Software Implementation

The user interface controlling the CTFM dual demodulation sonar system is shown in Figure 9.17. This module implements all of the elements shown in Figure 9.2, and provides the tonal output



(a) Demodulated output signals



(b) Demodulated output signals, filtered to remove crosstalk

Figure 9.16: CTFM tonal echo signals after demodulation (same signal as Figure 9.15). (a) Shows signal ⑥ in Figure 9.2. Time is measured in seconds, while the vertical scale indicates the signal amplitude (V). The two channels have very different amplitude profiles, due to variations between the left and right transducers. The amplitude profiles are dominated by the crosstalk signal. Their shapes are determined by the frequency response curve of the transducers (Figure 6.13) as the signal frequency sweeps from 100 kHz down to 50 kHz over the interval 0 s to 0.16 s. (b) shows the same signal as (a), but with the crosstalk frequency (about 1900 Hz) filtered out in Matlab using a notch filter. The shape is quite different from (a) due to the air absorption.

Initial Stage:

CTFM Double Demod

Configuration: 344kHz 25 downsample short ...

Start Freq: 100 kHz

End Freq: 50 kHz

Amplitude: 0.1 V

Demod Max Freq: 39.6552 kHz

Sweep Time: 160 ms

Max Range Ratio: 0.068

Capt. Samp. Freq: 344.827 kHz

Cap. Duration: 180 ms

☐ Output Frequency Reversed

Load Filter 1 Downsample: 25

Filename 1: "pb6ds25n451.filt", filter length 451

☐ Enable Filter 2

Load Filter 2 Downsample: 5

Filename 2: "pb5ds5n85.filt", filter length 85

Load Pre-filter ☒ Enable Pre-filter

Pre-Filter Filename: "ultrFilt52.344k.filt", filter length 52

☐ Debugging Info ☒ Finish Here

Figure 9.17: The module for the **CTFM** dual demodulation stage, from the feature extraction dialogue shown in Figure 5.13. The fields are the transmit sweep start frequency, f_0 , the transmit sweep end frequency, f_1 , the one-sided amplitude of the transmit signal, the maximum range frequency after demodulation, corresponding to the upper edge of the downsampling filter passband, the sweep time, T_{sw} , the maximum range ratio, m_{rr} , the capture sampling frequency, f_{sr} , and the capture duration. The other controls set the downsampling count, N_d , and allow the filters to be used for pre-filtering and downsampling to be selected. The Filter 2 section allows for the downsampling to be carried out in two stages where this is more efficient.

signal as input to the next module.

9.9 Range Resolution

This section firstly evaluates the **time-bandwidth product** of several signals, both reference and experimental. This further explains the operation of the **CTFM** sonar system and provides insight into the resolution properties. Secondly, the theoretical resolution of the **CTFM** system is evaluated using the frequency width technique. Thirdly, the theoretical resolution is evaluated using the time width technique. Finally, the results are compared with the theoretical system resolution limit calculated in Section 8.8.1.

9.9.1 Time-Bandwidth Measurements

It is instructive to measure the time and frequency widths of the signals appearing at various points through the **CTFM** system shown in Figure 9.2. Table 9.4 shows the results of doing this for both ideal and experimental received and demodulated signals. Applying the **time-bandwidth product** computations to the received signal shows the amount of compression and also signal to noise ratio improvement that is available. The frequency width of the demodulated output signal will be used to predict resolution.

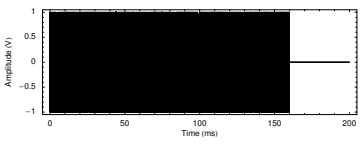
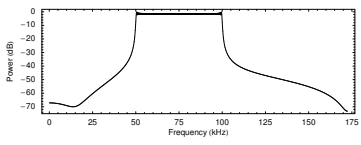
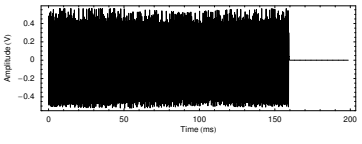
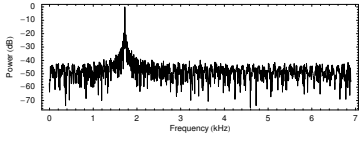
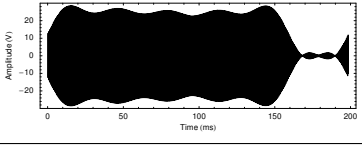
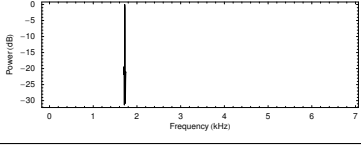
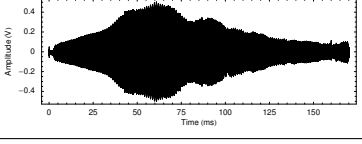
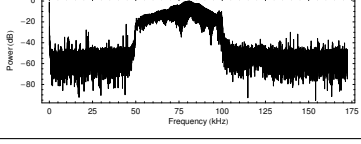
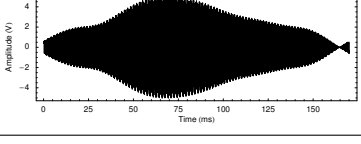
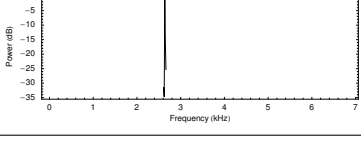
The results shown in Table 9.4 are now discussed row by row. The first row shows the reference transmit/receive signal, which ignores the effect of the elements of the signal path. The measured time width is significantly narrower than the specified $T_{swd} = 160\text{ms}$, and the frequency width is also considerably narrower than the specified bandwidth $B = 50\text{kHz}$. This is due to the method of measuring the width. Nevertheless, the time-bandwidth product is seen to be quite large.

The second row shows the ideal received signal after demodulation. The noise floor is clearly visible. However, the measured frequency width of 731 Hz is significantly wider than the main peak, which is 10 Hz wide at the -6dB point. The width measurement is significantly influenced by the noise floor.

The third row shows the result of masking out the frequencies which are not part of the demodulated signal. The masking was performed in the Fourier domain by zeroing out the frequency samples which were not required. The justification for doing this is by analogy with the time domain measurement of Section 8.8, where a section of the time domain signal containing the echo of interest is selected and used for computing the time-bandwidth product. In this case, the selection is performed in the frequency domain. The passband of the blanking operation was 60 Hz wide. The time width of the masked signal is computed by performing the inverse Fourier transform and measuring the width in the time domain. The results of performing this procedure are shown in the third row of Table 9.4. The frequency width reflects the actual width of the peak, and the centre frequency, f_0 , is now free of bias toward the centre of the frequency axis at $f_{sd}/4$. The time-bandwidth product, $\Delta t \cdot \Delta f$, is also much reduced and approaching the limit given by Gabor (see Section 4.6.1).

Turning now from reference signals to real signals, the fourth row of Table 9.4 shows the real received echo signal from a target placed at 500 mm from the sonar head. The signal also includes cross-talk and double and triple bounce returns. Thus it is not directly comparable with the ideal

Table 9.4: Summary of time-bandwidth product computations for reference and actual transmit/receive and demodulated CTFM signals. The equivalent duration and bandwidth were computed using the technique described in Section 4.6.4. See Section 9.9 for full details.

Time Domain Signal	Frequency Domain Signal	Duration	Bandwidth	TB Product
Reference Transmit/Receive 		$t_0 = 80.0 \text{ ms}$ $\sqrt{t^2} = 92.3 \text{ ms}$ $\Delta t = 115.7 \text{ ms}$	$f_0 = 75001 \text{ Hz}$ $\sqrt{f^2} = 76377 \text{ Hz}$ $\Delta f = 36185 \text{ Hz}$	$\Delta t \cdot \Delta f = 4188$
Reference Demodulated 		$t_0 = 79.5 \text{ ms}$ $\sqrt{t^2} = 91.8 \text{ ms}$ $\Delta t = 115.2 \text{ ms}$	$f_0 = 1745 \text{ Hz}$ $\sqrt{f^2} = 1769 \text{ Hz}$ $\Delta f = 731 \text{ Hz}$	$\Delta t \cdot \Delta f = 84.3$
Masked Reference Demodulated 		$t_0 = 80.3 \text{ ms}$ $\sqrt{t^2} = 92.2 \text{ ms}$ $\Delta t = 113.5 \text{ ms}$	$f_0 = 1723 \text{ Hz}$ $\sqrt{f^2} = 1723 \text{ Hz}$ $\Delta f = 12 \text{ Hz}$	$\Delta t \cdot \Delta f = 1.3$
Experimental Received 		$t_0 = 69.3 \text{ ms}$ $\sqrt{t^2} = 74.7 \text{ ms}$ $\Delta t = 70.0 \text{ ms}$	$f_0 = 79026 \text{ Hz}$ $\sqrt{f^2} = 79496 \text{ Hz}$ $\Delta f = 21651 \text{ Hz}$	$\Delta t \cdot \Delta f = 1515$
Masked Experimental Demodulated 		$t_0 = 76.7 \text{ ms}$ $\sqrt{t^2} = 82.3 \text{ ms}$ $\Delta t = 74.8 \text{ ms}$	$f_0 = 2637 \text{ Hz}$ $\sqrt{f^2} = 2637 \text{ Hz}$ $\Delta f = 10 \text{ Hz}$	$\Delta t \cdot \Delta f = 0.75$

signals described previously which only contained one echo return. The time and frequency widths are smaller than those of the ideal signal shown in the first row, due to the filtering action of the transducers and other elements in the signal path. The time-bandwidth product, at 1515, is significantly less than the ideal 4188. However, the time-bandwidth product of this CTFM signal is much larger than that of the MF signal seen in Table 8.1. The effect of this will be seen in the DOF precision experimental results in Section 13.2.1, where it will be seen that the maximum range achieved by the MF system is lower than that achieved by the CTFM system (due to the higher energy level of the signal).

The real demodulated signal was masked to select the echo of interest, and a similar procedure to that already described for the ideal demodulated signal was performed. The results are shown in the fifth row of Table 9.4. The time width is not as great as that of the ideal received signal. However, the frequency width, at 10Hz, is slightly narrower than that of the ideal signal at 12Hz. The time-bandwidth product is also smaller at $\Delta t \cdot \Delta f = 0.75$. Comparing this with the time-bandwidth product of the received signal, it can be seen that the demodulation process has compressed the TB product almost down to limit imposed by (4.11). Unlike the MF processor described in Chapter 8, the CTFM sonar system has achieved this chiefly by compressing the bandwidth of the signal rather than by compressing its duration.

When comparing the ideal and experimental demodulated signals, it is seen that the experimental signal achieves a lower time-bandwidth product than does the ideal signal. This is primarily due to significant narrowing of the time width (due to the transducer frequency response) rather than to the marginal narrowing of the frequency width.

9.9.2 Theoretical Resolution Using Frequency Width

The frequency width, Δf , is used to directly predict the resolution of the CTFM sonar. The relationship (9.22) between frequency and DOF may be utilised to calculate the range increment, Δr , corresponding to the change in frequency, Δf , yielding

$$\Delta r = \frac{\Delta d}{2} = \frac{-c}{\mu} \frac{\Delta f}{2} = 5.5 \text{ mm.} \quad (9.23)$$

This is the predicted minimum range separation between two targets at which the CTFM sonar system will be able to resolve them.

9.9.3 Theoretical Resolution Using Time Width

This section calculates a theoretical limit for the range resolution using the time width of the CTFM signal. The frequency width is calculated from the time width using Gabor's uncertainty relation (4.11), in a technique complementary to that employed in Section 8.8.4. The relationship between (Gabor's measure of) signal width, Δt , and range resolution, Δr , is found by combining Gabor's time-bandwidth product relation (4.11) with the relationship between range and frequency, (9.22). Applying this to the time width of the experimental demodulated signal, $\Delta t = 115.7 \text{ ms}$ (from Table 9.4), yields

$$\Delta r = \frac{c \cdot \Delta t}{2} = \frac{c}{-4 \cdot \mu \cdot \Delta t} = 3.7 \text{ mm.} \quad (9.24)$$

Table 9.5: Summary of range resolution results for the CTFM system

Method	Resolution (mm)
Theoretical System Resolution	3.2
Frequency Width Resolution	5.5
Time Width Resolution	3.7

9.9.4 Summary of Range Resolution Results

The theoretical range resolution results for the CTFM sonar system are summarised in Table 9.5.

The time width result is somewhat better than the experimental result shown in Section 13.3, and also much better than the result predicted by the frequency width calculation. However the difference is not so large as to invalidate the calculation. The time width resolution compares favourably with the theoretical system resolution limit from Section 8.8.1. The more favourable results make more assumptions about the system. The frequency width resolution technique makes the fewest assumptions, and is therefore expected to be the closest to experimental results, which are given in Section 13.3.

9.10 Conclusion

A CTFM dual demodulation sonar system has been constructed as a synthesis of hardware and software, with the majority of the signal processing being carried out digitally in software.

The output signal consists of one tone per target. The relationship between frequency and distance of flight is determined by the CTFM demodulation parameters. The extraction of the target echoes must be performed in the frequency domain. Several precautions have been taken to remove glitches from the CTFM echo signal in order to provide suitable input to the spectral estimator. Three different spectral estimators are described in Chapters 10 and 11. These chapters will also discuss the detection of targets, and the estimation of their time of flight and echo power. Thus the complete CTFM sonar system will provide similar information to the matched filter sonar system.

This chapter also utilised time-bandwidth product theory to calculate the theoretical limit of resolution of the system. Two different methods were applied:

1. The resolution was estimated directly from the bandwidth of the demodulated signal.
2. The time-bandwidth product relationship was applied to the CTFM output signal time width to compute the bandwidth and hence the theoretical resolution.

The time width technique produces a result which compares favourably with the theoretical system resolution limit calculated in Section 8.8.1. These results will be compared with experimental results and subsequently discussed in Section 13.3.

The time-bandwidth product of each of the transmit, receive and demodulated output signals was computed, and showed that the CTFM processor acts as a frequency compressor while preserving the bandwidth of the signal.

9.10.1 Acknowledgements and Contributions

The dual demodulation scheme was taken from descriptions in the literature, as described in Chapter 3. The digital implementation described here, including the system design shown in Figure 9.2 and its implementation in C++, is the sole work of this author. The computation of the [time-bandwidth product](#) of the various signals as they pass through the sonar system is also the contribution of this author.

Chapter 10

DFT Spectrum Estimation

The signals output by the [CTFM](#) sonar system described in Chapter 9 provide a frequency which is proportional to the time of flight, and hence to the range to the target. Thus the range measurement problem becomes a frequency estimation problem when a [CTFM](#) sonar system is used. The traditional solution to this problem is to apply the [Discrete Fourier Transform \(DFT\)](#) (Kay [73]) to the demodulated echo signal. This chapter summarises the properties of the [DFT](#) which are relevant to this problem, and relates them to the [CTFM](#) sonar characteristics. In particular it is shown that there is a practical lower limit upon the [Distance Of Flight \(DOF\)](#) quantisation by the [DFT](#).

The spectral estimate produced by the [DFT](#) process requires that the time samples be properly spaced according to the Nyquist criteria. Furthermore, any practical time sample data must necessarily be finite in extent. This is equivalent to using a rectangular window on the signal (Kay and Marple [75]). The effects of a rectangular window are discussed by Harris [57], chief of which is the high side-lobe level in the spectral estimate. Furthermore, the [DFT](#) method of producing the spectral estimate, comparable to the periodogram method (Kay and Marple [75]), presumes that the signal is extended by periodic repetition in the time domain. The periodic extension is a symptom of the sampling in the frequency domain. However, most of the damage is done separately by the rectangular windowing (Harris [57]). The effect of these assumptions upon the quality of the spectral estimate will be seen in the experimental results sections of Chapters 13 and 14.

10.1 The Fourier Transform

The Fourier transform maps a real (or complex) function of time into a complex function of frequency. Bracewell [14] gives the definition for a continuous signal, $g(t)$, as

$$\mathcal{F}\{g(t)\} = G(f) = \int_{-\infty}^{\infty} g(t) e^{-j2\pi ft} dt \quad (10.1)$$

where t is time measured in seconds and f is frequency measured in Hertz. \mathcal{F} is used to denote the Fourier transform. The constant j denotes the imaginary number $\sqrt{-1}$.

There is also a discrete version, called the [Discrete Fourier Transform \(DFT\)](#). This is usually referred to as the [Fast Fourier Transform \(FFT\)](#), after a particular implementation by Cooley and Tukey [24, 25] which reduced the computational complexity of the [DFT](#). The implementation of the [FFT](#) is described in many texts, for example Oppenheim and Schaffer [105]. Aside from reduced computational complexity, its numerical properties are identical to those of the [DFT](#).




Figure 10.1: The relationship between time and frequency domain data points using the [Discrete Fourier Transform \(DFT\)](#), for the case where the time domain samples have only real values. The front left box face shows a time domain signal, which is composed of the cosine signals shown within the body of the box. The samples are separated by units of Δt . The front right box face shows frequency domain representation of the same signal, representing the amplitude of each of the cosine signals. The samples are separated by units of Δf . The time domain signal shown here has a nominal frequency of $2.2\Delta f$. This is not periodic in T_{DFT} , so its [DFT](#) representation distorts at the end points. The last (ninth) point, shown in grey, has the same value as the first point at 0, and is regarded as the first point in the next repetition of the sequence. Thus the last point is not included in the [DFT](#) input. The frequency domain points are usually complex, but the phase of the signal shown here has been specially selected to obtain real values for clarity.

Using array notation, the equation defining the [DFT](#) is

$$G[k] = \frac{1}{\sqrt{N_{DFT}}} \sum_{l=0}^{N_{DFT}-1} g[l] e^{-j2\pi lk/N_{DFT}}. \quad (10.2)$$

10.1.1 Time and Frequency Sampling

The relationship between the samples of the time and frequency domains is shown in Figure 10.1. The number of samples taken in the time domain is N_{DFT} . The output of the [DFT](#) is, in general, complex, and of the same length as the input. However, for real input, the output will be Hermitian,

$$G(-f) = G^*(f), \quad (10.3)$$

where $*$ denotes the complex conjugate. Therefore only the positive frequency components need to be computed, as the negative frequency components may be obtained from (10.3) if needed.

An important consideration when using the **DFT** is to try to select the number of sample points to be $N_{DFT} = 2^k$, where k is a positive integer, as this provides superior performance over other values nearby (Bracewell [14, Chapter 18]), even when radix division algorithms are used.

The number of spectral lines N_{spec} computed for the positive frequency domain is

$$N_{spec} = \frac{N_{DFT}}{2}. \quad (10.4)$$

This is only valid when N_{DFT} is even, as used in this study due to the 2^k constraint. The sampling rate of the time domain signal is denoted f_{sd} and is measured in Hertz. The duration of the sample sequence in the time domain, T_{DFT} , is measured in seconds and is defined by

$$T_{DFT} = \frac{N_{DFT}}{f_{sd}}, \quad (10.5)$$

where f_{sd} is the sampling frequency of the downsampled output defined by (9.20). This is the time interval from the first sample in one periodic sequence to the first sample in the next periodic sequence, which is not included in the **DFT** input. The spacing, Δt , between samples in the time domain is

$$\Delta t = \frac{1}{f_{sd}}, \quad (10.6)$$

making the time $t(l)$ of sample $g[l]$ equal to

$$t(l) = l \cdot \Delta t = \frac{l}{f_{sd}}, \quad (10.7)$$

measured in seconds. The spacing, Δf , between spectral lines in the frequency domain is

$$\Delta f = \frac{f_{sd}}{N_{DFT}} = \frac{1}{T_{DFT}}, \quad (10.8)$$

making the frequency $f(k)$ of spectral line $G[k]$ equal to

$$f(k) = k \cdot \Delta f = \frac{k f_{sd}}{N_{DFT}} = \frac{k}{T_{DFT}}, \quad (10.9)$$

measured in Hertz. The spectral line index is bound by the limits $0 \leq k \leq N_{spec}$, making the maximum available frequency equal to N_{spec}/T_{DFT} . The maximum available frequency may be written $f_{sd}/2$, which is equal to the Nyquist frequency.

The **DFT** assumes that the input signal repeats periodically, with period T_{DFT} . The sample which occurs at T_{DFT} is not part of the input data array, as it is assumed to be the replica of the sample at $t = 0$. Thus Figure 10.1 shows the end point greyed out.

The first and last points in the **DFT** output (frequency domain) are real. This may be derived from the **DFT** of real and imaginary sequences given by Oppenheim and Schaffer [105]. Thus some **FFT** packages pack these points into a single complex storage unit.

10.2 Relationship Between the DFT and CTFM

The **CTFM** sonar system requires that the tonal echo signal, obtained as the result of Chapter 9, be analysed in the frequency domain. The **DFT** system equations are now combined with the **CTFM** system equations from the previous chapter.

10.2.1 Signal Time

The length of echo signal available for spectral analysis is now determined. If the Fourier transform were taken over the entire sweep period, then some initial transients would be included in the signal. The Fourier transform of the transients has the effect of broadening the peak in the frequency domain (Bracewell [14]). Thus the signal from before T_{max} , defined by (3.22), is omitted from the input to the Fourier transform.

The single demodulation scheme, described in Section 3.3.1, will only provide an echo signal up to the sweep time, T_{sw} . While the dual demodulation scheme described in Section 3.3.2 extends the amount of data available, the echo signal will, in general, include a phase change at $T_{sw} + T_{max}$ which will introduce fine structure into the frequency domain.

To avoid peak broadening and fine structure, only the signal from the maximum range echo arrival time, T_{max} , to the end of the transmit sweep time, T_{sw} , is used as input to the Fourier transform. This is the quantity T_{sig} , defined by (3.23). This limits the time over which the DFT may be taken,

$$T_{DFT} = \frac{N_{DFT}}{f_{sd}} \leq T_{sw} \cdot (1 - m_{rr}), \quad (10.10)$$

where m_{rr} is the maximum range ratio, defined by (3.22).

The sonar system discussed in this thesis uses $T_{sig} = 0.158$ s, and a downsampled sampling rate of $f_{sd} = 13\,793$ Hz. Using (10.6), there are 2184 samples available. Therefore N_{DFT} is chosen to be 2048, being the largest power of 2 that is consistent with the data.

10.2.2 Spectral Line to Distance of Flight

The second connection between the CTFM equations and the DFT distance of flight corresponding to a particular spectral line may be calculated by combining the frequency of a spectral line (10.9) with the relationship between frequency and DOF (9.22) to give

$$d(k) = \frac{c}{\mu} \left(\frac{1}{8} - \frac{k f_{sd}}{N_{DFT}} \right), \quad (10.11)$$

where k is the spectral line number, μ is the CTFM sweep rate defined by (3.2), and c is the speed of sound. The time of flight is obtained by omitting the speed of sound from (10.11).

10.2.3 DOF Quantisation

The distance of flight quantisation may be calculated by finding the change in distance, Δd , from one spectral line to the next, as defined by (10.11):

$$\Delta d = d(k+1) - d(k) = \frac{-c f_{sd}}{\mu N_{DFT}}. \quad (10.12)$$

Using the sweep rate, μ , defined by (3.2), the sweep bandwidth, B , defined by (3.26), and the DFT sample time, T_{DFT} , defined by (10.5), equation (10.12) can be re-arranged to yield

$$\Delta d = \frac{c T_{sw}}{B T_{DFT}}. \quad (10.13)$$

Assuming that (10.10) is an equality, using the maximum available signal for the DFT, then

$$\Delta d = \frac{c}{B \cdot (1 - m_{rr})}, \quad (10.14)$$

where the sweep time is seen to cancel out. The maximum range ratio is minimally constrained by $0 \leq m_{rr} \leq 1$, but also has other constraints, which will be discussed further in Section 10.2.4.

10.2.4 Minimisation of DOF Quantisation

The computation of bearing by Inter-aural Distance Difference (IDD), which is discussed in Chapter 14, requires precise measurements of the distance of flight. The precision obtainable is limited by the Distance Of Flight (DOF) quantisation calculated by (10.14). Therefore the DOF quantum defined by Equation (10.14) must be minimised. As the speed of sound, c , cannot be controlled, one or more of the following courses of action must be followed to obtain smaller quantisation:

1. Increase the bandwidth B .
2. Reduce the maximum range ratio, m_{rr} . This implicitly affects the sweep time, T_{sw} , as discussed below.

Increase Bandwidth

The usable bandwidth of the CTFM signal is limited by the bandwidth of the signal path, described in Chapter 4. The low frequency limit is defined by the transducer response, shown in Figure 6.13, while the high frequency limit is defined by both the transducer response and the absorption of the air (Figure E.1). These limits give us no scope to increase the system bandwidth beyond the parameters already achieved by the CTFM system in Tables 9.1 and 9.2. Therefore the value used for the bandwidth was 50 kHz.

Reduce Maximum Range Ratio

The maximum range ratio, m_{rr} , is not only important in determining the DOF quantisation, but also determines the maximum distance of flight by (3.24) and determines the maximum range by (3.25). The maximum range is increased by increasing the maximum range ratio, m_{rr} . Thus maximum range and DOF quantisation are competing requirements. This conflict is balanced in this implementation, by selecting a large sweep time, T_{sw} , and selecting the maximum range ratio, m_{rr} , to provide a moderate maximum range and a slightly higher than minimum DOF quantisation. The values used are $T_{sw} = 160$ ms and $m_{rr} = 0.068$, providing a maximum range of 1.88 m. Reducing the maximum range ratio, m_{rr} , will decrease this maximum range.

Summary

While (10.14) appears to allow arbitrary selection of Δd , the best obtainable value of the DOF quantisation, Δd , is 7.42 mm. This is equivalent to a range quantisation, Δr , of 3.7 mm. The large quantisation degrades the precision which may be obtained. The quantisation may only be reduced by using very wide bandwidth transducers or by sacrificing the maximum range.

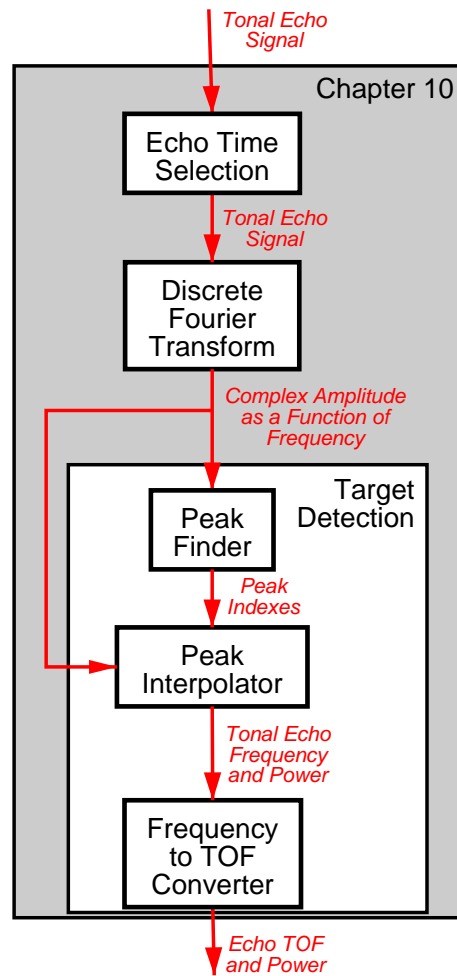


Figure 10.2: Overview of the procedure for extraction of echo information from the tonal echo signal using a DFT. The entire DFT output is sent to both the peak finder and the peak interpolator.

10.3 Echo Extraction

The output of the CTFM sonar system described in Chapter 9 is a tonal echo signal. The method of extracting the echo parameters from the tonal echo signal is outlined in Figure 10.2. The remainder of this section covers the major boxes of this diagram.

10.3.1 Echo Time Selection

The echo time selection process selects the portion of the signal which will contain echo tones from all of the targets. The lead time before the arrival of the echo from maximum range at T_{max} (see Section 10.2.1 or Figure 3.5), is discarded. Similarly, any portion after the sweep time, T_{sw} , is also discarded, as described in Section 10.2.1. The output of this stage is still a tonal echo signal, and is passed on to the discrete Fourier transform stage.

10.3.2 Discrete Fourier Transform

The **DFT** was implemented using the **Fastest Fourier Transform in the West (FFTW)** [41, 42], a freely available optimised implementation of the **FFT** algorithm.

The **FFTW** package divides the computation of the **FFT** into two phases: a planning phase, which must be performed once for any given data size, and a subsequent computation phase. The planning phase finds the optimum way of performing the **FFT** given the computing architecture and the data size. The computational cost of the planning phase is recovered over many **FFT** computations, which execute more efficiently when compared with a non-optimal method.

The chirp program manages these two phases using a C++ front end to **FFTW**, which stores plans for each data size used within a program run. A request for a **FFT** first checks whether a plan has been pre-computed; if none exists then a plan is generated and stored. The **FFT** is subsequently computed using the plan. The C++ class which does this is called **FFTDriver**.

A typical **DFT** of **CTFM** echoes is shown in Figure 10.3. Of interest in the context of peak finding is the large number of peaks. The **DFT** output contains a large amount of noise, which is not reduced by increasing the length of the data record (Kay and Marple [75]).

Welch Method

The standard technique for reducing the noise, the Welch method (Welch [138], Therrien [131, Section 10.1.3]), is to divide the data set into several overlapping segments, to apply a taper window to each one, and to take the Fourier transform of each segment. The Fourier transforms are then averaged, reducing the magnitude of the noise. However, the reduction in noise comes at a price. In the process of dividing the data set into shorter segments, the time interval, T_{DFT} , over which the Fourier transform is taken, is reduced. In turn, this also increases the frequency quantisation (see (10.8)), and hence the distance of flight quantisation as well. Thus the Welch method is incompatible with the quantisation requirements of a precision sonar system, and is not used. No noise reduction is performed on the **DFT** output.

The output of the **DFT** stage, which is an array of complex numbers representing the spectral lines, is passed to both the peak finder and to the peak interpolator. The peak finder executes first.

10.3.3 Peak Finder

The operation of the peak finder comprises four steps, which are executed in order.

1. The magnitudes of the complex data points are found, and are used throughout the rest of the procedure.
2. Each point is compared to the detection threshold, and is discarded if this is not met.
3. Points not discarded by the detection threshold are found to be peaks if they are larger than the points on either side.
4. As there are a significant number of noise peaks, a peak must be larger than the neighbouring peaks of three samples on either side. This serves to reduce the effect of noise on the sides of other larger peaks.

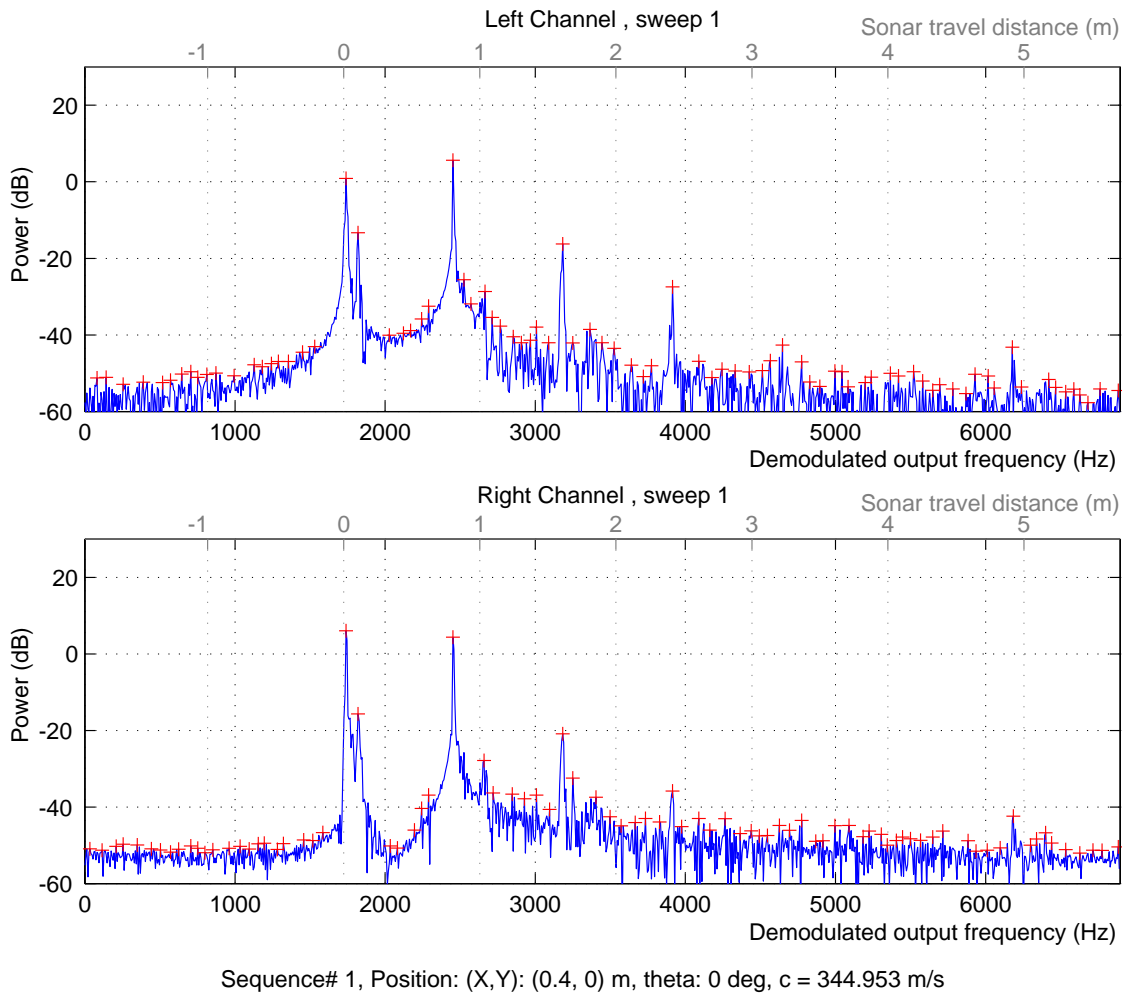


Figure 10.3: Plot of the entire DFT output of an experiment with a glass panel at 400mm range from the sensor. The plane of the glass was normal to the principal direction of the transmitter. Frequency is shown along the bottom axis, while the distance of flight is shown along the top (grey). The vertical scale is power in decibels, obtained by taking the square of the magnitude of the complex DFT output. Peaks are found using the method described in Section 10.3.3, and are shown with a cross +. Only the interval from 0...3.76m corresponds to distance of flight. The peaks near 0m are acoustical crosstalk. The peak at 0.8m is the target, while the peaks at 1.6m and 2.4m are multiple reflections of the target echo.

The indexes of the peaks which are found are then passed on to the peak interpolator.

10.3.4 Peak Interpolator

The peak interpolator requires both the peak indexes from the peak finder and the complex amplitudes from the [DFT](#).

The peak interpolator's function is to provide a more precise estimate of the frequency of the tonal echo than is possible using the discrete indexes of the peak finder which are tied to system characteristics. The interpolator is required because the spectral line step quantisation of $\Delta f = 6.73 \text{ Hz}$ converts to a [DOF](#) quantisation which is 7.4mm here, compared to a worst case of 1 mm in the matched filter systems of Chapter 8. Thus some form of interpolation is clearly necessary.

The literature indicates that there are two ways in which interpolation can be implemented.

1. Zero padding in the time domain has the effect of increasing the length of the sample, which narrows the quantisation in the frequency domain. This can be seen from equation (10.8), where N_{DFT} is increased while retaining the same sampling rate, f_{sd} , producing a smaller value of Δf . Note that this does not narrow the width of the sinc function (see equation (10.17) below) that is produced by the Fourier transform of the data, and therefore cannot improve the resolution of any technique based upon the [DFT](#).
2. Fit a function, in the frequency domain, to the complex amplitude peak value and its left and right off-peak values (Davies [28]).

An initial investigation of Davies' work suggested that improved interpolation algorithms requiring less [CPU](#) resources than the zero padding technique were now available. In particular, Davies' method E [28], complex exact interpolation, showed clear superiority over five other techniques also evaluated by Davies [28], and thus we implemented Davies' method E.

Implementation

Davies' method E [28] assumes that the input signal is of the form

$$s(t) = a \sin(2\pi ft + \phi), \quad (10.15)$$

where f is the frequency of the signal, in Hertz, and ϕ is a phase offset, being constrained by $0 \leq \phi < 2\pi$. The signal is sampled, $s[n] = s(n \cdot \Delta t)$, and the [DFT](#) is taken. Davies then interpolates to find the peak frequency, f , using the largest sample point, with index k_m , and the points on its left and right, with indexes $k_m - 1$ and $k_m + 1$, with the formula

$$f^2 = \frac{1}{T_{DFT}^2} \cdot \frac{Ck_m - 1^2 - 2C[k_m]k_m^2 + Ck_m + 1^2}{C[k_m - 1] - 2C[k_m] + C[k_m + 1]}, \quad (10.16)$$

where $C[k]$ refers to the k th complex spectral line provided by the [DFT](#) output, and T_{DFT} is the period over which the [DFT](#) input data extends, and is given by $N_{DFT} \cdot \Delta t = N_{DFT} / f_{sd}$.

The interpolation formula (10.16) only includes the central three points because this is where the signal energy is concentrated. The points outside the central peak have lower signal energy

and therefore a lower signal to noise ratio. Thus the effect of including more points would be to include more noise in the computation (Davies [28]).

Equation (10.16) produces a real frequency when applied to the [discrete Fourier transform](#) of a signal such as (10.15). However, it tends to give a complex frequency when experimental data is used to generate the spectral lines, $C[k]$. Since frequency, f , is defined to be a real quantity, only the real part is used. In this implementation, if the complex part is too large, the peak is assumed to be badly formed and is discarded. A threshold of 15 for the imaginary part was found to be sufficient to avoid rejecting valid targets.

The Discrete Fourier transform of (10.15), found by the methods of Bracewell [14], is

$$C[k] = \frac{jaT_{DFT}}{2} e^{-j\frac{k\phi}{fT_{DFT}}} \{\text{sinc}(k + fT_{DFT}) - \text{sinc}(k - fT_{DFT})\}. \quad (10.17)$$

This includes terms for both positive and negative frequencies, which can only be distinguished when the input is complex. As the positive and negative frequency components are quite far apart, the negative frequency sinc function in (10.17) may be ignored. Furthermore, the normalisation for T_{DFT} is computed by the [DFT](#) code, so it will not be considered further here. Thus (10.17) simplifies to

$$C[k] = \frac{-ja}{2} e^{-j\frac{k\phi}{fT_{DFT}}} \text{sinc}(k - fT_{DFT}). \quad (10.18)$$

Furthermore, when the magnitude, $|C[k]|$, is taken, the complex exponential drops out, leaving

$$|C[k]| = \frac{a}{2} \text{sinc}(k - fT_{DFT}).$$

Davies' interpolation method does not provide an estimate for the amplitude, a , of the sine-wave. As this sonar system requires the amplitude, a , it is found using the value of the maximum spectral line,

$$a = \frac{2|C[k_m]|}{\text{sinc}(k_m - fT_{DFT})}.$$

Thus the sinc function will pass exactly through the maximum spectral line. The power is subsequently found from a^2 .

A typical interpolated peak is shown in Figure 10.4.

The interpolation scheme outlined in this section makes three assumptions upon the input signal:

1. that the signal is of the form given by equation (10.15),
2. that the negative frequency sinc function (see (10.17)) is far enough away that it has no effect upon the positive frequency sinc function,
3. that the target is an isolated one with no interference from adjacent echoes, and
4. that the three points used for interpolation contain negligible amounts of noise.

In practice, these assumptions have varying degrees of applicability. Assumption 1 is slightly invalidated by the time domain envelope of practical echo signals, as the envelope is modulated rather than flat (see Figure 9.16b). Assumption 2 is made valid by the duration of the signals used.




Figure 10.4: Closeup of **DFT** spectrum around the main target. The geometrical path length is 0.8028m. The comb plot is the values $|z|^2$ computed by the **DFT** upon the standard input data set, containing 2048 samples in the time domain. The peak finder has selected the highest point in this series and marked it with a magenta +. The green curve consists of the values $|z|^2$ computed by a **DFT** with zero padding for a total of 131072 samples in the time domain. The peak finder has selected the highest peak, marked with a green \times . The zero padded curve is included for comparison only, and is not used by the sonar system. The red curve is a sinc function fitted to three points, being the spectral line marked with a magenta + and the two peaks on either side. The maxima of the sinc function, which is used by the sonar system as an estimate of the tonal echo frequency and power, is indicated by a red \times .

Assumption 3 is invalid when targets are adjacent, as when performing the resolution experiment. The breakdown of this assumption will be examined practically in Section 13.3.3. Assumption 4 is supported by only selecting the central three points, which have the best available SNR.

The position of the peak found by the complex interpolation scheme is quite similar to that found by the more computationally expensive zero padding technique. Thus the interpolation scheme is used to provide the estimate of the tonal echo frequency and power.

At this point, the frequency and power of the tonal echo signals have been determined. Target detection is completed with the frequency to TOF converter.

10.3.5 Frequency to TOF Converter

The tonal echo frequencies are converted to Time Of Flight (TOF) using (9.22) and (3.9), where the speed of sound, c , has been cancelled out:

$$\tau = \frac{1}{\mu} \left\{ \frac{1}{4N_d} \frac{f_{sr}}{2} - f_{out} \right\}. \quad (10.19)$$

Using the parameters and dependent variables of Tables 9.1 and 9.2, (10.19) becomes

$$\tau = -5.515 \times 10^{-3} + 3.198600 \times 10^{-6} \cdot f_{out}.$$

The conversion of TOF, τ , to DOF, d , will be performed in Chapter 12, and is performed by a separate module, as shown in Figure 9.1.

The final output of the sonar system has now been obtained, providing similar information to both of the matched filter sonar systems seen in Chapter 8.

10.3.6 User Interface

The sonar system described in this chapter is designated the CTFM-DFT sonar system throughout the remainder of this thesis. The user interface for the CTFM component was shown in Figure 9.17. The user interface corresponding to the block diagram of Figure 10.2 is shown in Figure 10.5.

10.4 Conclusion

The DOF quantisation of the DFT spectral estimator is tightly constrained by its interaction with the CTFM sonar system. The only practical way to improve the DOF quantisation, Δd , from the current value of 7.4mm is to reduce the maximum range of the sonar system. While at first it would seem that increasing the sampling rate, f_s , in the time domain would improve the DOF quantisation, Section 10.1.1 demonstrates that this does not improve Δf , and by extension, Δd remains constant throughout this operation. The additional samples instead increase the maximum sample available in the frequency domain.

As the DOF quantum is quite large, especially when compared to the quantum in the matched filter output of Chapter 8, interpolation of the peak position is of crucial importance. A complex interpolation scheme due to Davies and Creecraft [28] was implemented to perform this task efficiently. The effectiveness of the system as a whole will be evaluated in Chapters 13 and 14.

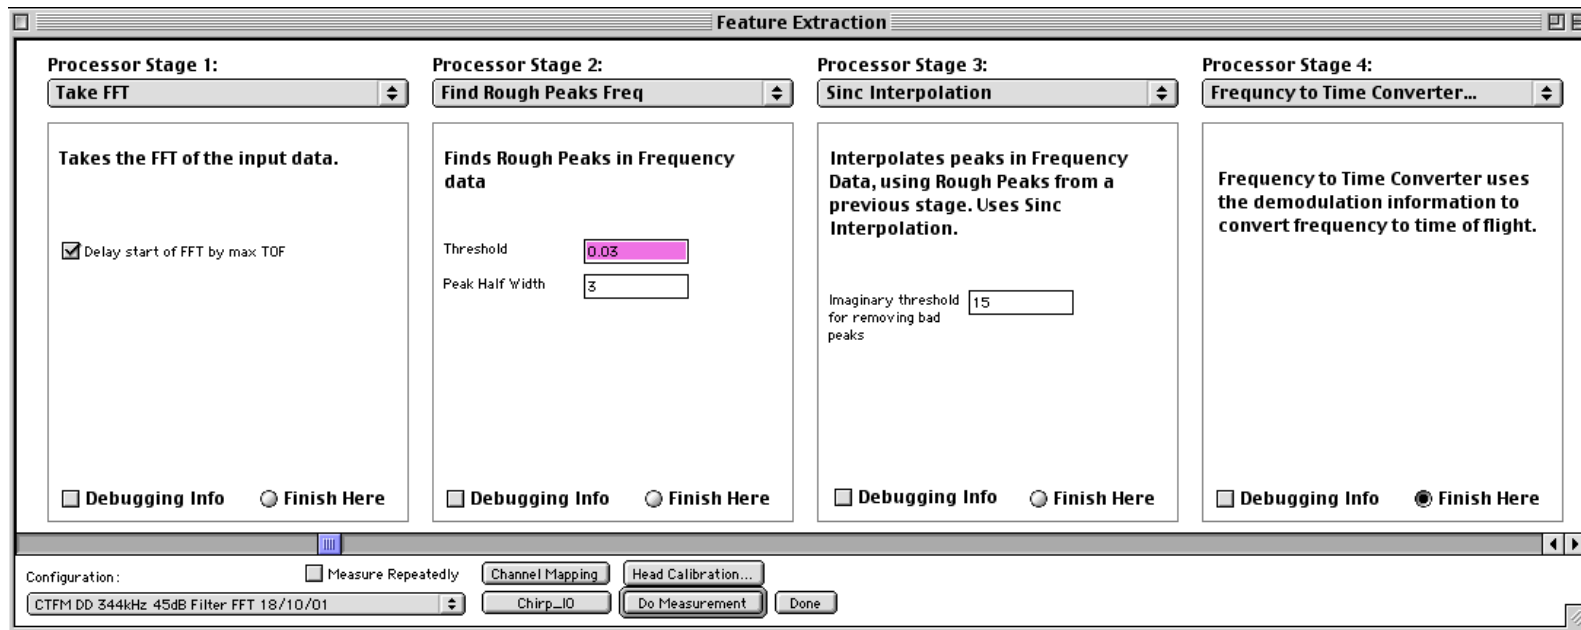


Figure 10.5: Settings for the [DFT](#) and spectral estimation modules used as part of the CTFM-DFT sonar system. Processor Stage 1 implements the first two stages of Figure 10.2. The check box controls whether the Echo Time Selection stage is performed. Processor Stage 2 implements the peak finder stage of Figure 10.2, and returns peak indexes. The settings control the detection threshold and the width of the window described in Section 10.3.3. Processor Stage 3 implements the peak interpolator, and includes a setting for the threshold for rejecting complex peaks. Processor Stage 4 implements the Frequency to [TOF](#) converter described in Section 10.3.5, using the sweep information from the [CTFM](#) stage.

10.4.1 Acknowledgements and Contributions

The analysis of a [CTFM](#) signal using the [DFT](#) has been performed elsewhere (Davies [\[28\]](#)). However, the analysis of the design parameters to optimise the range step, Δd , and the implementation of this system, are the work of this author.

Chapter 11

Auto-Regressive Spectral Estimation

The traditional or classical power spectrum estimators belong to the non-parametric class and are based on the periodogram. Chapters 4, 6, 8 and 10 are based on these methods which have their origins in the now classical work of Bartlett (1948) [8], Blackman and Tukey (1958) [13], and Welch (1967) [138]. Much of this work is incorporated into the work of Bracewell (1965) [14] which has underpinned much of this thesis.

These classical approaches are characterised by a frequency resolution quantisation that is tied to the length of the data span. In the case of Chapter 10, where a CTFM signal is analysed with such a classical DFT, the level of quantisation is a poor 7.4mm. The broad nature of the peaks in the DFT output is caused by the assumed periodic extension of the available input data points. The literature suggests a number of approaches, chief among which is the parametric class of spectral estimators. These estimators utilise additional knowledge about the process producing the signal to extend the data points past the available data span without making the assumption of periodicity. This provides a higher resolution estimate of the Power Spectral Density (PSD) than the DFT method.

The parametric class of spectral estimators is well described in the seminal tutorial paper of Kay and Marple (1981) [75]. This paper suggests that the following three Auto-Regressive (AR) class spectral estimators offer advantages over classical DFT methods:

1. Autoregressive (AR) Yule-Walker (YW) version.
2. Autoregressive (Burg) algorithm.
This estimator was specifically rejected as it is known to have reduced performance in the presence of moderate to high noise levels. It is also known to have frequency biases.
3. Autoregressive (AR) least squares or forward-backward linear prediction version. This is called the Least Squares Modified Yule-Walker (LSMYW) method in this thesis.

The discussion in Kay and Marple [75] is heavily biased toward resolving two close frequencies, the resolution problem of Section 2.1, compared to accurately and precisely determining the peak frequency which is related to range precision. It is shown in Chapters 13 and 14 that range precision, synonymous in this case with the precision of the frequency determination, is a significant constraint on the determination of bearing.

The fundamental difference between the DFT and the AR spectral estimates is in the assumptions made of the data beyond the samples supplied from the experimental data. Whereas the DFT

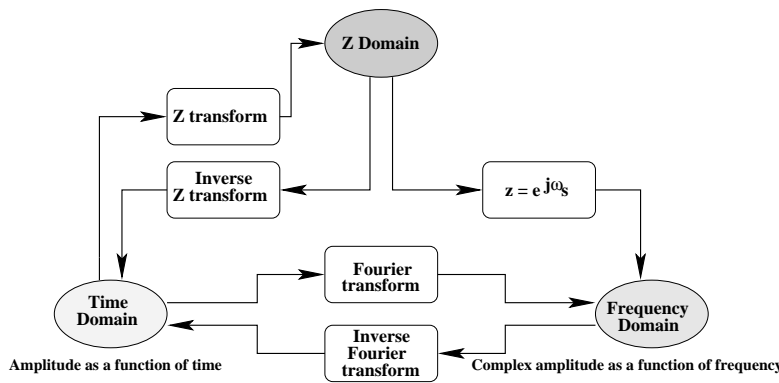


Figure 11.1: The relationship between the sampled time series, the Z domain, and the Fourier frequency domain. There is no method of transforming from the frequency domain to the Z domain, as the Z domain contains a superset of the information present in the frequency domain.

assumes a periodic extension, the [AR](#) spectral estimators extend the data by assuming that the statistical properties of the data remain constant, that is the data is stochastic, and that a model order is known a-priori. While there are techniques to estimate the model order (Kay and Marple [75]), these are of questionable reliability, and have not been applied here.

It should be noted that the performance of modern spectral estimation techniques is highly dependent upon the data presented, and techniques which work well on one data set may perform poorly on another data set. Thus, it is possible that one of the other methods listed by Kay and Marple [75] may provide better results than either of the two [AR](#) models tested in this thesis.

The following section develops the basic signal processing knowledge which is necessary to apply these two methods. The chapter then describes the operation of both the [Yule-Walker](#) and the [least squares modified Yule-Walker](#) models and finally describes their implementation and integration into the chirp application program used in this thesis.

11.1 Theory

11.1.1 The Z Transform

Modern digital signal processing techniques utilise the Z transform for design and analysis [116, Chapter 3], as it simplifies computations involving convolutions between sampled series. It also allows the transfer function of a digital system to be analysed in terms of its frequency response, due to the close relationship between the Z transform and the Fourier transform. The relationships between the time domain samples, the Z transform, and the frequency domain are shown in Figure 11.1.

The Z transform, denoted \mathcal{Z} , of a sampled sequence $x[n]$ is

$$\mathcal{Z}\{x[n]\} = X(z) = \sum_{n=-\infty}^{\infty} x[n]z^{-n}, \quad (11.1)$$

where z is a complex number, and $X(z)$ represents a continuous complex function in the complex plane.

If the sequence $x[n]$ is zero outside the interval $0 \leq n < N_x$, the limits of summation may be adjusted to

$$\mathcal{Z}\{x[n]\} = X(z) = \sum_{n=0}^{N_x-1} x[n] z^{-n}. \quad (11.2)$$

This is quite similar to the definition of the Fourier transform (10.2). If the z argument is equated to $\exp(j2\pi k/N_z)$, the term which appears in (10.2), then (11.2) and (10.2) become equivalent. The equivalence only holds around the unit circle in the complex plane. Furthermore, using equation (10.9) which relates the frequency, f , and the spectral line number, k , the z parameter may be related to frequency:

$$z = e^{j2\pi k/N_z} = e^{j2\pi f/f_s} = e^{j\omega_s}. \quad (11.3)$$

The variable f_s is the sampling frequency, and ω_s is scaled angular frequency, defined by $\omega_s = 2\pi f/f_s$, and is introduced for convenience of notation.

Given that the power spectral density, $\mathcal{P}(f)$, for a Fourier transform is calculated by $|F(\omega_s)|^2$, the relationship between the Z transform and the Fourier transform allows the PSD of a Z transform to be calculated in a similar way:

$$\mathcal{P}(f) = |X(e^{j\omega_s})|^2, \quad (11.4)$$

where it is understood that the frequency conversion, from f to ω_s , is carried out by $\omega_s = 2\pi f/f_s$ from (11.3).

11.1.2 Rational Model for Spectral Estimation

Systems in the form of a rational model are computationally convenient, and may be used to approximate many other types of system, given sufficient order. Systems conforming to the rational model in the time and Z domains will now be described, and the use of the factored representation of the Z transform of the system function will be explained.

A system described by the rational model processes input samples, $x[n]$, into output samples, $y[n]$, according to the relationship

$$y[n] = \sum_{k=1}^p a_k y[n-k] + \sum_{k=0}^q b_k x[n-k]. \quad (11.5)$$

This computes the new output sample point, $y[n]$, in terms of the previous output sample points, $y[n-k]$, and the current and previous input sample points, $x[n-k]$. The a_k and b_k parameters are real. Oppenheim and Schaffer [105] show that the Z transform of (11.5) is

$$S(z) = \frac{b_0 + b_1 z^{-1} + \dots + b_q z^{-q}}{1 + a_1 z^{-1} + \dots + a_p z^{-p}} = \frac{\sum_{k=0}^q b_k z^{-k}}{1 + \sum_{k=1}^p a_k z^{-k}}, \quad (11.6)$$

where z is the transform variable. The polynomials in the numerator and denominator of (11.6) will have the same number of complex roots as their polynomial order. Equation (11.6) may be factored to obtain

$$S(z) = b_0 \cdot z^{p-q} \frac{(z-q_1)(z-q_2)\dots(z-q_q)}{(z-p_1)(z-p_2)\dots(z-p_p)} = b_0 \cdot z^{p-q} \cdot \frac{\prod_{k=1}^q (z-q_k)}{\prod_{k=1}^p (z-p_k)}, \quad (11.7)$$

where q_k represents the k th complex root of the numerator of $S(z)$, and p_k represents the k th complex root of the denominator. Zeros in the denominator cause poles in the system function $S(z)$, so the p_k are referred to as poles. The poles and zeros must occur in complex conjugate pairs to obtain real time domain coefficients, a_k and b_k . Equation (11.7) may be combined with the Fourier transform relationship, (11.4), to analyse the frequency response of the rational system.

The role of the numerator of the rational model, (11.6), is now discussed. Setting all the a_k equal to zero yields

$$S(z) = b_0 + b_1 z^{-1} + \dots + b_q z^{-q}. \quad (11.8)$$

The time domain equivalent of (11.8) is

$$y[n] = \sum_{k=0}^q b_k x[n-k]. \quad (11.9)$$

This is a **Finite Impulse Response (FIR)** filter system, also known as a **Moving Average (MA)** system. The factored Z-transform function, derived from (11.7), is now composed of zeros. The frequency response of such a system will consist of low points near where the zeros occur. If the zeros are placed upon the unit circle, then the frequency response will contain a zero. Cadzow [20] demonstrates that the **DFT** fits a **MA** system to the data in order to estimate the spectrum. The order q of the **MA** system is equal to the number of input samples when estimated by the periodogram method. Kay and Marple [75] show that **MA** and **DFT** systems have broad peaks and function poorly when modelling spectra with peaks, requiring many zeros, and hence many time samples, to compensate.

Turning now to the denominator of (11.6), setting all but the first of the b_k parameters from (11.6) equal to 0 yields

$$S(z) = \frac{b_0}{1 + a_k z^{-1} + \dots + a_p z^{-p}}. \quad (11.10)$$

The time domain version of (11.10) is

$$y[n] = \sum_{k=1}^p a_k y[n-k] + b_0 x[n]. \quad (11.11)$$


Equation (11.11) computes the output point, $y[n]$, as a function of the previous output points, $y[n-k]$, and the current input point. Such a system has an **Infinite Impulse Response (IIR)**, as the input is recycled by the filter. This filter structure is also called an **Auto-Regressive (AR)** system.

The Z transform of the system function, (11.10), may also be factored. Recognising that (11.10) has no zeros, and only poles, setting $q = 0$ in (11.7) yields


$$S(z) = \frac{b_0 \cdot z^p}{(z - p_1)(z - p_2) \dots (z - p_p)} = \frac{b_0 \cdot z^p}{\prod_{k=1}^p (z - p_k)}, \quad (11.12)$$

showing that the **AR** system contains p poles, and a multiple zero of order p at $z = 0 + 0j$.


An example **AR** system, having poles at $z = 0.98 \exp(\pm j3\pi/8) = 0.375 \pm 0.905j$, is shown in Figure 11.2. The form of (11.12) requires that a double zero is placed at $z = 0 + 0j$. The positions of the poles and zeros is shown in Figure 11.2b. This can be seen to have a direct effect upon the magnitude of the Z transform, as shown in Figure 11.2a. The numerator roots cause zeros in the Z



(a) Z-plane decibel magnitude



(b) Z-plane pole-zero plot



(c) spectrum determined by Fourier transform

Figure 11.2: An example of an **auto-regressive** (all-pole) system. (a) Shows the decibel squared magnitude of the Z transform of the system function. The unit circle is drawn on the surface in red. (b) Shows the location of the two conjugate poles, p_k , marked \times , in the complex plane of the Z transform, causing the peaks to appear at those locations in (a). The poles must occur within the unit circle for the system to be stable (Proakis and Manolakis [116]). This plot also shows the double zero located at the origin, marked with \odot . (c) Shows the frequency response spectrum of the system, obtained from the Z transform by (11.4). Only two poles are required to produce sharp peaks. The locations of the peaks on the spectrum plot, at $\omega_s = \pm 3\pi/8$, are controlled by the positions of the poles. The poles are located at $z = 0.98 \exp(\pm j3\pi/8) = 0.375 \pm 0.905j$.




Figure 11.3: Time domain input and output of a [linear predictive filter](#). The integer n is related to time by $t(n) = n/f_s$. The vertical scale indicates the sample amplitudes. The problem is to predict, or estimate, $\hat{x}[n]$ given p past samples of the input sequence, $x[n-p], \dots, x[n-1]$.

transform magnitude, while the denominator roots cause poles in the Z transform magnitude. The frequency response of the system, which may be evaluated using (11.4), is shown in Figure 11.2c. The frequency response is the same as the magnitude of the Z transform evaluated around the unit circle, which is shown in red on Figure 11.2a.

The multiple zeros located at $z = 0 + 0j$, introduced by the factorisation of (11.12), have no effect upon the frequency response of the system. The zeros are introduced by the term $S_z(z) = z^{-2}$ from (11.12). Evaluating the PSD via (11.4), yields $\mathcal{P}(f) = |e^{-2j\omega_s}|^2 = 1$, showing that the additional zeros have no effect upon the spectrum.

Figure 11.2 shows that each pole produces a sharp peak in the system frequency response. Thus it is hypothesised that [AR](#) systems will be good at modelling power spectra containing peaks.

Having examined the rational model, and identified the [AR](#) system type as having desirable properties for estimating spectra with peaks, [AR](#) models must now be constructed to represent a particular time domain data sequence. This problem is addressed by the [Yule-Walker](#) and [least squares modified Yule-Walker](#) spectral estimators in the next two sections.

11.2 Yule-Walker Spectral Estimator

The [Yule-Walker \(YW\)](#) spectral estimator operates by fitting an [AR](#) system to the input data, in the sense that the [AR](#) system will be able to predict and extend the sequence, and then use the spectral properties of the system to estimate the spectrum of the input signal. The system is modelled using a [linear predictive filter](#).

11.2.1 Linear Predictive Filter

Therrien [131, Section 7.2] describes how a signal may be statistically modelled with a [Linear Predictive Filter \(LPF\)](#). The [LPF](#) predicts sample $x[n]$ given the p samples $x[n-p], \dots, x[n-1]$ which have preceded it, as shown in Figure 11.3. The [linear predictive filter](#) problem is to find negative [MA](#) filter coefficients, $-a_1, -a_2, \dots, -a_p$, which may be used in equation (11.9) to estimate

the next point, $\hat{x}[n]$, in the sequence

$$\hat{x}[n] = -a_1x[n-1] - a_2x[n-2] - \dots - a_px[n-p], \quad (11.13)$$

where $\hat{}$ denotes an estimated value. The negative notation has been adopted for later convenience. The solution to (11.13) is constrained by the requirement to minimise the error, denoted $\varepsilon[n]$. The error is calculated by the difference between the actual value of the next sample, $x[n]$, and the estimated value $\hat{x}[n]$. Thus,

$$\varepsilon[n] = x[n] - \hat{x}[n]. \quad (11.14)$$

The error sequence, $\varepsilon[n]$, may be statistically characterised by its variance, σ_ε^2 , which is calculated using the expectation value operator, $E\{\}$,

$$\sigma_\varepsilon^2 = E\{|\varepsilon[n]|^2\} = E\{|x[n] - \hat{x}[n]|^2\}, \quad (11.15)$$

Thus the **linear predictive filter** problem is solved by finding a set of $a_1 \dots a_p$ which minimises the variance, σ_ε^2 , of the prediction error sequence.

Equation (11.13) may also be written in the equivalent form

$$\hat{x}[n] = \sum_{k=1}^p -a_k x[n-k]. \quad (11.16)$$

This has the same form as (11.9), showing that $\hat{x}[n]$ is produced as the output of a **linear time-invariant FIR** filter. This explains the use of the term filter in the name **linear predictive filter**.

It is also possible to express the error process, $\varepsilon[n]$, as a **FIR** filter, by defining $a_0 \equiv 1$. Thus,

$$\varepsilon[n] = x[n] - \hat{x}[n] = \sum_{k=0}^p a_k x[n-k]. \quad (11.17)$$

This is the **Prediction Error Filter (PEF)**. The a_k parameters were selected to be negative to support the construction of (11.17). They cancel the subtraction of the estimate, $\hat{x}[n]$, in (11.17), making the **PEF** a positive sum. The relationship between the **LPF** and the **PEF**, which are modelled by (11.16) and (11.17) respectively, is shown in Figure 11.4.

The method of finding the parameters, a_k , which satisfy the linear prediction problem is found using the orthogonality principle (Therrien [131]). Two sequences are said to be orthogonal when their cross correlation, $E\{x[n-l]y[n]\}$, is equal to zero. This is equivalent to saying that the signals are uncorrelated. The orthogonality principle for optimal filtering states that the filter is optimal when the error output is orthogonal to the input signal:

$$E\{x[n-l]\varepsilon[n]\} = 0, \quad l = 1, 2, \dots, p. \quad (11.18)$$

The orthogonality principle only operates at offsets, l , between the two sequences of one or more samples, as the predictor does not yet know the next point, $x[n]$, in the sequence, and can only be orthogonal to the p points it has already seen. The error in the estimate provides the expected value of the equivalent of (11.18) at zero offset l ,

$$E\{x[n]\varepsilon[n]\} = \sigma_\varepsilon^2. \quad (11.19)$$

This is the minimum prediction error, which is obtained when condition (11.18) is met. If condition (11.18) is not met, the error will be larger than given by (11.19).

The a_k parameters may be derived, using (11.18), from the previously observed data. Equation (11.19) will be used to find the parameter b_0 which sets the overall level of the spectral estimate.

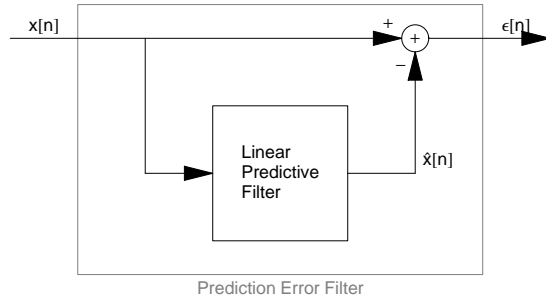


Figure 11.4: The relationship between the **Linear Predictive Filter (LPF)**, which predicts $\hat{x}[n]$ given the previous p samples of the sampled sequence $x[n]$, and the **Prediction Error Filter (PEF)**, which calculates the error $\epsilon[n]$ in the prediction $\hat{x}[n]$. The diagram depicts the flow of one sample through the filter system. The diagram is executed once for each input sample. The **LPF** maintains a memory of the previously seen samples for computing its output.

Finding the a_k Parameters

Starting with (11.18) and substituting for $\epsilon[n]$ using (11.17) yields

$$\sum_{k=0}^p a_k E \{x[n-l]x[n-k]\} = 0, \quad l = 1, 2, \dots, p,$$

which may be re-written as

$$\sum_{k=0}^p a_k R_x[k-l] = 0, \quad l = 1, 2, \dots, p \quad (11.20)$$

where $R_x[l]$ is the autocorrelation of the sequence $x[n]$. Discussion of how to calculate the autocorrelation is deferred to Section 11.2.2.

In the case of $p = 3$, Equation (11.20) may be expanded into three equations. Remembering that $a_0 \equiv 1$, the equations are

$$\begin{aligned} R_x[-1] + a_1 R_x[0] + a_2 R_x[1] + a_3 R_x[2] &= 0 \\ R_x[-2] + a_1 R_x[-1] + a_2 R_x[0] + a_3 R_x[1] &= 0 \\ R_x[-3] + a_1 R_x[-2] + a_2 R_x[-1] + a_3 R_x[0] &= 0 \end{aligned}$$

These are known as the Yule-Walker equations. The three equations may be used to solve for the three unknowns $a_1 \dots a_3$. The equations may be re-written in matrix form for $p = 3$ by subtracting the $R_x[l]$ terms with unit coefficients to the right hand side:

$$\begin{bmatrix} R_x[0] & R_x[1] & R_x[2] \\ R_x[-1] & R_x[0] & R_x[1] \\ R_x[-2] & R_x[-1] & R_x[0] \end{bmatrix} \cdot \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} -R_x[-1] \\ -R_x[-2] \\ -R_x[-3] \end{bmatrix}. \quad (11.21)$$

This may be solved by the standard methods of linear algebra (Anton [6]). More generally, for any

number of filter coefficients p the matrix form may be written as

$$\begin{bmatrix} R_x[0] & R_x[1] & \dots & R_x[p-1] \\ R_x[-1] & R_x[0] & \dots & R_x[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ R_x[1-p] & R_x[2-p] & \dots & R_x[0] \end{bmatrix} \cdot \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = \begin{bmatrix} -R_x[-1] \\ -R_x[-2] \\ \vdots \\ -R_x[-p] \end{bmatrix}. \quad (11.22)$$

The computational complexity of solving (11.22) directly is $O(p^3)$.

Once the a_k are known, equation (11.19) may be used to easily calculate the variance of the prediction error, σ_ε^2 :

$$\sigma_\varepsilon^2 = \sum_{k=0}^p a_k R_x[k]. \quad (11.23)$$

Thus the a_k parameters from the autocorrelation of the sampled sequence $x[n]$ have been determined, together with the prediction error variance σ_ε^2 . Methods of calculating the autocorrelation $R_x[l]$ will be considered next.

11.2.2 Autocorrelation of a Sequence

The autocorrelation of a sequence is

$$R_x[l] = E\{x[n]x[n-l]\}, \quad (11.24)$$

where $E\{\}$ is the expectation value operator. There are several methods available for computing the estimate $\hat{R}_x[l]$. Cadzow [20] discusses the selection of the autocorrelation estimator and notes that it is very important to select an estimator appropriate to the problem at hand. Two formulae using the available data $x[k]$, $k = 1, 2, \dots, N_x$ are discussed here. The first is the unbiased estimator,

$$\hat{R}_x[l] = \frac{1}{N_x - |l|} \sum_{k=1}^{N_x} x[k]x[k-l], \quad (11.25)$$

which tends toward the true autocorrelation as the number of sample points, N_x , becomes large. The second formula is the biased estimator,

$$\hat{R}_x[l] = \frac{1}{N_x} \sum_{k=1}^{N_x} x[k]x[k-l], \quad (11.26)$$

which is equivalent to the unbiased estimate, given by (11.25), multiplied by a triangular window function, $N_x / (N_x - |l|)$. Kay and Marple [75] indicate that the biased estimator (11.26) will have less mean square error than (11.25) for many practical data sets. The biased estimator (11.26) was used in this work. In both formulae, (11.25) and (11.26), any use of $x[k]$, $k \notin \{1, 2, \dots, N_x\}$ is deemed to be 0 for convenience.

The sums in (11.25) and (11.26) have a computational complexity of order $O(N_x^2)$. A speedup is discussed in Section 11.4.2. The negative lags are computed using the symmetry relationship,

$$R_x[-n] = R_x[n]. \quad (11.27)$$

This reduces the required computation by half.

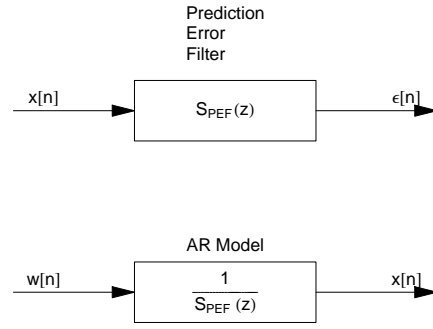


Figure 11.5: Comparison between the [prediction error filter](#) (whitening filter) and the [AR](#) model used for spectral estimation.

11.2.3 Spectral Estimation

Having discussed the estimation of the autocorrelation function, the [linear predictive filter](#) may now be applied to the problem of spectral estimation.

Consider the [prediction error filter](#) depicted at the top of Figure 11.5. If the [prediction error filter](#) of Section 11.2.1 is constructed with sufficient order p , then according to (11.18) and (11.19) the [prediction error filter](#) output $\varepsilon[1], \varepsilon[2], \dots, \varepsilon[n]$ will be pure white noise, representing a noise process. Since a white noise signal has a flat spectrum, this implies that the linear predictive filter must have the inverse spectrum of the signal, and filters it out. Therefore the inverse system of the [prediction error filter](#) must have a spectral response which matches that of the signal. The system depicted in the lower part of Figure 11.5 is the inverse of the [prediction error filter](#). It shows a white noise source driving a filter with the same spectral response as the sampled sequence, $x[n]$. The output of the filter then produces a signal which models the sequence $x[n]$ in a spectral sense. It is this inverse system which is used for performing spectral estimation.

Going over this again in more detail, the Z transform of the [prediction error filter](#) is

$$S_{PEF}(z) = \sum_{k=0}^p a_k z^{-k},$$

with $a_0 \equiv 1$. The inverse of the [prediction error filter](#) is

$$S_{IPEF}(z) = \frac{1}{\sum_{k=0}^p a_k z^{-k}},$$

If a white noise source or noise process with variance $\sigma_w^2 = \sigma_\varepsilon^2$ is used to drive this filter, a signal $\hat{x}[k]$ with the same statistical and spectral properties as the original input signal $x[k]$ will be obtained. Thus the system

$$S_{est}(z) = \frac{\sigma_\varepsilon^2}{\sum_{k=0}^p a_k z^{-k}} = \frac{b_0}{\sum_{k=0}^p a_k z^{-k}} \quad (11.28)$$

is used to estimate the spectrum of the input signal $x[n]$. The form including b_0 demonstrates that (11.28) is equivalent to (11.10).




Figure 11.6: Example of using the minimum number of data points (2) for a linear least squares fit. This provides an extremely poor fit to the data, as the two points in question have a considerable error. The errors in the estimated line fit become less significant as more points are used. From Cadzow [20].

To obtain a spectrum in terms of scaled angular frequency ω_s instead of z , equation (11.3) may be used to write

$$\mathcal{P}_{AR}(f) = |S_{est}(e^{j\omega_s})|^2. \quad (11.29)$$

This allows the AR spectral estimate to be found from the signal data by the Yule-Walker method.

Assumptions

The YW spectral estimation technique assumes that the prediction error sequence, $\varepsilon[n]$ as given by (11.17), is uncorrelated with the signal (see (11.18)). This assumption is only valid if the order of the LPF is sufficient to model the signal. The assumption may be tested for a particular filter order by computing (11.18) and verifying that the cross-correlation is low. A high cross-correlation indicates a breakdown of the orthogonality principle due to insufficient filter order.

The practical implication of utilising an insufficient filter order is that parts of the signal escape the prediction filter and remain in the prediction error sequence. The impact of this upon spectral estimation is that some spectral components are absent from the spectral estimate. Close spectral peaks may be modelled by a single pole instead of using one pole each, as is necessary for faithful reproduction. Thus it is important to choose the correct filter order to obtain good results. There are some automated methods of doing this (Kay and Marple [75]).

11.3 Least Squares Modified Yule-Walker Spectral Estimation

The second method of calculating the a_k parameters which is used in this thesis is the Least Squares Modified Yule-Walker (LSMYW) method. This method recognises that the weakness in the AR method is the accuracy of the autocorrelation estimate $\hat{R}_x[p]$, and uses an additional number of autocorrelation lags to solve for the a_k parameters in the least squares sense. The motivation for doing this is similar to the straight line fit process outlined in Figure 11.6. The method of finding the a_k parameters is due to Cadzow [20].

The **LSMYW** method uses more than the minimal number of autocorrelation points, $\hat{R}_x[p]$, needed to solve an extended set of **Yule-Walker** equations (11.20). This makes the system over-determined. The addition of this extra data for $k > p$ introduces an extra parameter in addition to the number of poles, p . This extra parameter is the number $t - p$ of additional autocorrelation points, $\hat{R}_x[p]$, used to solve the **Yule-Walker** equations. The extra $t - p$ points improve the statistical robustness of the estimated **PSD** when there are errors present in the autocorrelation estimate.

A set of equations is formed by t (instead of p) applications of (11.20),

$$\begin{bmatrix} R_x[-1] & R_x[0] & R_x[1] & \cdots & R_x[p-1] \\ R_x[-2] & R_x[-1] & R_x[0] & \cdots & R_x[p-2] \\ R_x[-3] & R_x[-2] & R_x[-1] & \cdots & R_x[p-3] \\ R_x[-4] & R_x[-3] & R_x[-2] & \cdots & R_x[p-4] \\ R_x[-5] & R_x[-4] & R_x[-3] & \cdots & R_x[p-5] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ R_x[-t] & R_x[1-t] & R_x[2-t] & \cdots & R_x[p-t] \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (11.30)$$

These equations are over-determined. The coefficients of $a_0 \equiv 1$, forming the first column, may be moved to the right hand side analogously to (11.21). This eliminates both the 1 at the top of the \mathbf{a} vector and the left column of the autocorrelation matrix:

$$\begin{bmatrix} R_x[0] & R_x[1] & \cdots & R_x[p-1] \\ R_x[-1] & R_x[0] & \cdots & R_x[p-2] \\ R_x[-2] & R_x[-1] & \cdots & R_x[p-3] \\ R_x[-3] & R_x[-2] & \cdots & R_x[p-4] \\ R_x[-4] & R_x[-3] & \cdots & R_x[p-5] \\ \vdots & \vdots & \ddots & \vdots \\ R_x[1-t] & R_x[2-t] & \cdots & R_x[p-t] \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = \begin{bmatrix} -R_x[-1] \\ -R_x[-2] \\ -R_x[-3] \\ -R_x[-4] \\ -R_x[-5] \\ \vdots \\ -R_x[-t] \end{bmatrix}. \quad (11.31)$$

This may also be written in matrix form

$$\mathbf{R}_1 \mathbf{a} = -\mathbf{r}.$$

This problem cannot be solved directly, as the matrix on the left hand side is not square. However, a least squares solution may be found by constructing a Moore-Penrose pseudo-inverse (Weisstein [137]):

$$\mathbf{R}_1^T \mathbf{R}_1 \mathbf{a} = -\mathbf{R}_1^T \mathbf{r}, \quad (11.32)$$

where T denotes the matrix transposition. At this point we note that $\mathbf{R}_1^T \mathbf{R}_1$ is square and the inverse exists, since it is positive definite. Thus

$$\mathbf{a} = -(\mathbf{R}_1^T \mathbf{R}_1)^{-1} \mathbf{R}_1^T \mathbf{r}. \quad (11.33)$$

While this formulation demonstrates that the \mathbf{a} coefficients may be found, superior numerical results may be found by solving (11.32) with a library such as LAPACK [4].

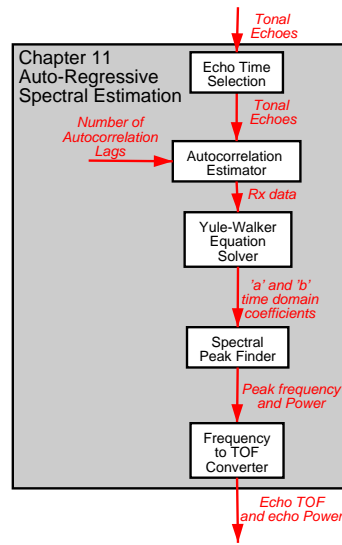


Figure 11.7: Outline of the steps performed in **YW** spectral estimation and peak finding. The **YW** and **LSMYW** methods follow the same sequence of steps, but use a different mechanism to solve the Yule-Walker equations. The number of autocorrelation lags is selected according to the Yule-Walker solver in use, and the (user) setting for p and t .

The variance of the noise process, σ_ε^2 , which is needed to normalise the estimator as described by (11.28), is obtained in the same way as for the **YW** method, by (11.23). The spectrum is calculated using (11.28) and (11.29), with the \mathbf{a} parameters found by (11.32).

The assumptions made by the **YW** spectral estimator also apply similarly to the **LSMYW** spectral estimator. In particular, poor results will be obtained if an insufficient filter order is selected.

11.4 Implementation

The **YW** and **LSMYW** spectral estimation techniques described earlier in this chapter are implemented within the C++ chirp application program described in Section 5.4. The operation of both the **YW** and **LSMYW** spectral estimators follows the same sequence of steps, which are shown in Figure 11.7.

The tonal echo signals produced as output by the **CTFM** sonar system in Chapter 9 are used as input to the spectral estimation system. The first stage is echo time selection.

11.4.1 Echo Time Selection

The echo time selection stage selects the interval of the echo signal once the echoes have all arrived, and performs identically to the block described for the **DFT** system in Section 10.3.1.

The output of this stage is a tonal echo signal, which is passed to the autocorrelation estimator, as shown in Figure 11.7.

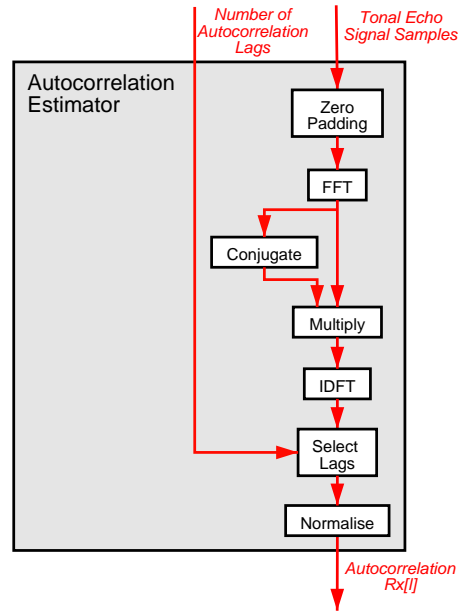


Figure 11.8: Block diagram of the operation of the autocorrelation estimator of Figure 11.7.

11.4.2 Autocorrelation Estimator

The autocorrelation estimator computes the biased estimate of equation (11.26), as described in Section 11.2.2. However, the computation is made more efficient by performing the computation using the FFT. The procedure used is outlined in Figure 11.8.

The first stage appends zeros to the end of the time amplitude array. This is necessary as the DFT computes circular convolution, rather than the required linear convolution (Oppenheim and Schaffer [105, Section 8.9]). Linear convolution may be simulated, for signals of finite extent, by appending zeros to the end of the input array, making the total input size

$$N_{\text{FFT}} = 2N_x - 1,$$

where N_x is the size of the input array. Thus $N_x - 1$ zeros are appended to the input data.

Bracewell [14] shows that the autocorrelation may be computed from (4.6) by

$$g(t) \star g(t) = g(-t) * g(t) = \mathcal{F}^{-1} \{ (\mathcal{F} \{g(t)\})^* \cdot \mathcal{F} \{g(t)\} \}. \quad (11.34)$$

The \star denotes correlation, the $*$ denotes convolution, \mathcal{F} denotes the Fourier transform, and $*$ denotes complex conjugation. The blocks FFT, Conjugate, Multiply, and Inverse Discrete Fourier Transform (IDFT), of Figure 11.8 compute equation (11.34). The computational complexity of (11.26) when computing p lags is $O(N_x p)$, whereas the work in computing (11.34) is $O(2N_{\text{FFT}} \log_2 N_{\text{FFT}})$, independent of how many lags are required (Figure 11.9).

The second last stage of Figure 11.8 selects the number of autocorrelation lags required by the spectral estimation method. The YW method requires p autocorrelation lags, while the LSMYW method requires t autocorrelation lags. The number of lags is passed as an input argument.

The last stage of Figure 11.8 performs the normalisation indicated in (11.26), multiplying each autocorrelation lag by $1/N_x$.




Figure 11.9: The computational complexity of the direct (green) and DFT (blue) methods of computing the autocorrelation. It was assumed that $p = 128$ lags were required when computing the graph.

The autocorrelation estimate, $R_x[l]$, has now been obtained to the required number of lags. This estimate is now passed on to the Yule-Walker Equation Solver, as shown in Figure 11.7.

11.4.3 Yule-Walker Equation Solver

The YW and LSMYW spectral estimators solve different sets of Yule-Walker equations to obtain the a_k and b_0 coefficients.

Yule-Walker

The YW estimator solves (11.22) using the Durbin algorithm, which exploits the symmetry present in the matrix on the left hand side. The Durbin algorithm is described in Appendix C.2.

Least Squares Modified Yule-Walker

The LSMYW estimator solves (11.32) using LAPACK [4]. The method used is described in Appendix C.3.

At this point the a_k and b_0 coefficients have been obtained. They are now passed over to the spectral peak finder.

11.4.4 Spectral Peak Finder

The operation of the spectral peak finder is somewhat different from the peak finders used in the matched filter and CTFM-DFT sonar systems described in Chapters 8 and 10, due to the availability of a continuous spectral function which may be evaluated at any frequency. A block diagram of the process used is shown in Figure 11.10. Each of the blocks are now described in turn.

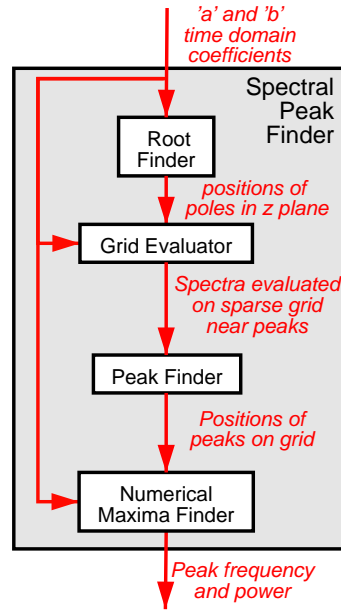


Figure 11.10: Block diagram of the steps performed by the spectral peak finder. This is used by both the YW and LSMYW systems, and corresponds to the Spectral Peak Finder box of Figure 11.7.

Root Finder

The YW and LSMYW methods provide estimates of the AR coefficients $a_1 \dots a_p$. These coefficients can be used to calculate a spectrum as described in Section 11.1.2. Figure 11.2 shows that the spectral peaks will be near the poles of (11.10), which are in turn determined by the roots of the polynomial

$$1 + a_1 z + a_2 z^2 + \dots + a_p z^p = 0. \quad (11.35)$$

The roots may be found efficiently by constructing a companion matrix and finding its eigenvalues, as described by Golub and Van Loan [45, Section 7.4]. The code which performs the root finding may be found in the file `roots.h` of the chirp application program. The roots are then used as starting points for a more precise numerical search, which is performed by the grid evaluator.

Grid Evaluator

The grid evaluator computes the spectrum using the function

$$\mathcal{P}(f) = \left| \frac{1}{1 + a_k e^{-j\omega_s} + \dots + a_p e^{-jp\omega_s}} \right|^2, \quad (11.36)$$

which is derived from (11.10) using the relationship between the Z and Fourier transforms (11.4). The scaled radial frequency, ω_s , is related to the frequency f in Hertz as described in Section 11.1.1. The function (11.36) is evaluated on a fixed grid, equivalent to 16384 points between 0 and $f_{sd}/2$ inclusive. The peaks are only evaluated for ± 30 points on either side of a pole (see Figure 11.11). In general, the peak is not located where a radius line drawn through the pole

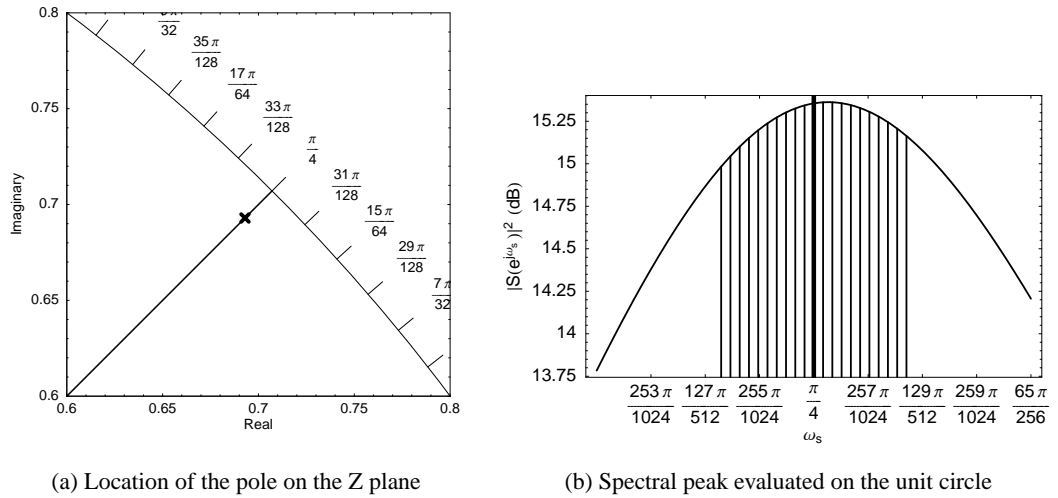


Figure 11.11: Diagrams showing how the grid evaluator works. A system having four evenly spaced poles, and their four complex conjugates, is used. The poles are placed near the unit circle. (a) Shows a closeup of the position of one of the poles. A radius line is drawn from $0 + 0j$ through a pole location to the unit circle. (b) Shows a sample grid of points where the spectrum is evaluated. The grid shown here is not as fine as the one used by the program. The spectral peak does not coincide with the position of the pole, which is marked by a longer vertical line than the others.

intersects with the unit circle, due to interaction with the other poles.

The grid evaluation procedure is performed once for each root. The grid points are stored into a sparse array. The C++ code uses the STL `map` class, described by Josuttis [66], to store the points, associated with their index positions on the grid. Thus the `map` ends up containing a set of points clustered around the positions of the poles. If two poles occur very close together, then some points may be evaluated more than once. This is not considered to be a problem for performance.

The sparse grid of the power spectral density is passed on to the peak finder.

Peak Finder

A peak finder is then used upon the data in the sparse array. This finds the actual positions of the peaks, to a precision within the granularity of the grid. The peak finder is the same as that described in Section 8.7.1. The peak positions are then passed to the numerical maxima finder.

Numerical Maxima Finder

The numerical maxima finder exploits the continuous power spectral density function (11.36) to find the position of the peak without requiring interpolation. The numerical maxima finder utilises the Brent method described by Press [114], and requires three points as input. These three points are initialised to be the maxima found by the peak finder, and the two grid points on either side. The numerical maxima finder evaluates (11.36) and recursively determines the position of the

peak much more precisely than the precision of the power spectral density function itself. Thus the quantum used in the grid to find the peaks has no bearing on the final peak output by the maxima finder.

The spectral peak finder produces as output a list containing the frequency of each peak, and its associated peak power. This concludes the operation of the spectral peak finder component of Figure 11.7. The list is passed to the block which converts frequency to TOF.

11.4.5 Frequency to TOF Converter

Once the peak frequencies and powers have been obtained, each frequency is converted to TOF by the same method as is used for the DFT estimator in Section 10.3.5. The TOF and echo power are passed on to the TOF to DOF converter, as shown in Figure 9.1. The TOF to DOF converter will be described in Chapter 12.

11.4.6 User Interface

The combination of CTFM and a YW spectral estimator will be referred to as the CTFM-YW sonar system for the remainder of this thesis. Sometimes the notation CTFM-YW(p) is used, where p indicates the number of \mathbf{a} coefficients used when solving (11.22). The user interface controlling the CTFM-YW sonar system is shown in Figure 11.12.

The combination of CTFM and a LSMYW spectral estimator will be referred to as the CTFM-LSMYW sonar system for the remainder of this thesis. Sometimes the notation CTFM-LSMYW($p, t - p$) is used, where p indicates the number of \mathbf{a} coefficients, and t indicates the number of autocorrelation lags, used when solving (11.32). The user interface for the CTFM-LSMYW sonar system has a user interface which is very similar to that shown in Figure 11.12, and a pre-configured system may be obtained by selecting an LSMYW configuration using the pop-up menu at the bottom left of the dialogue. There is one additional field, the Number of Extended YW Equations, which controls $t - p$.

11.5 Conclusion

Two methods of parametric spectral estimation have been described and implemented. These methods are expected to have better precision than the DFT spectral estimator, due partly to the lack of quantisation in the output, but also due to the different method used to extend the series outside the relatively short time domain data record available. The type of information provided by the CTFM-YW and the CTFM-LSMYW sonar systems is no different in its nature to that provided by the matched filter sonar system of Chapter 8, and the CTFM-DFT sonar system of Chapter 10, providing estimates of the TOF and echo power. The TOF will be converted to DOF in Chapter 12.

The YW and LSMYW spectral estimators are both parametric. They require as input the number of zeros to use in the PEF model of the signal, which is equivalent to the number of poles used in the AR model which is used to generate the spectral estimate. If the number of poles is not matched to the signal, and in particular is insufficient, then a poor estimate results.

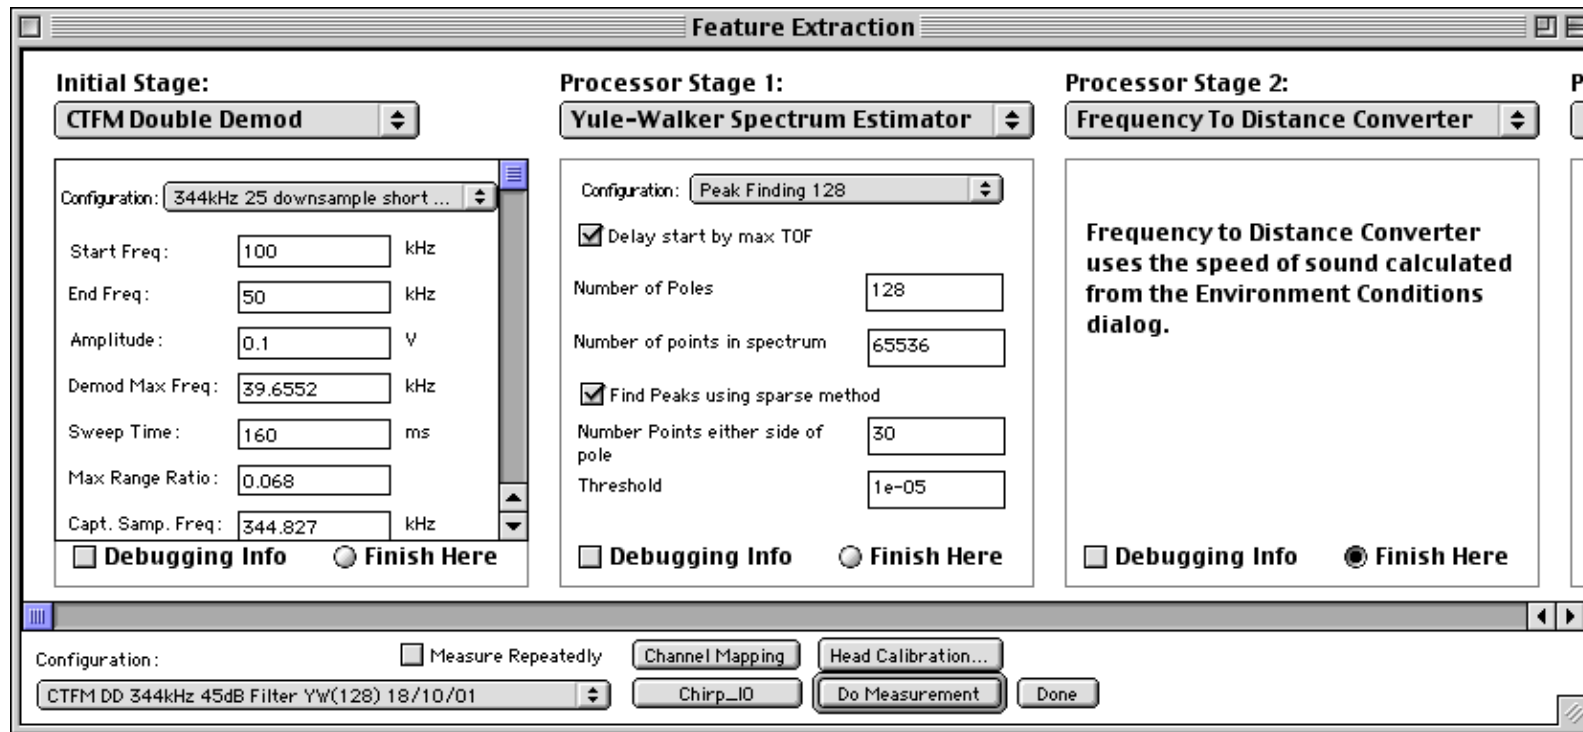


Figure 11.12: Settings for the CTFM sonar and YW spectral estimation and peak finding modules used as part of the CTFM-YW sonar system. The Initial Stage is the same as seen in Figure 9.17. Processor Stage 1 implements all of the elements of the YW spectral estimator and peak finder shown in Figure 11.7. The delay start by max TOF check box allows the echo time selection stage of Figure 11.7 to be bypassed. The number of poles edit box controls p . The number of points in spectrum box is only used if the sparse method of peak finding is not selected. The Find peaks using sparse method check box controls whether the sparse method described in Section 11.4.4 is used, or the grid is simply evaluated over the entire spectrum. The number of points on either side of pole box controls the width of the sparse grid of Section 11.4.4. The threshold field determines the detection threshold, as described in Section 4.2.2. Processor Stage 2 combines the functionality of the frequency to TOF converter of Figure 11.7 and the TOF to DOF converter of Figure 9.1.

11.5.1 Acknowledgements and Contributions

The underlying theory of the Z transform, the rational model for spectral estimation, the YW and the LSMYW methods of spectral estimation are standard techniques described in the literature. The application of these methods to the problem of DOF estimation using CTFM sonars, and the implementation of the software permitting this estimation in a single coherent system, is the unique contribution of this author.

Chapter 12

The Speed of Sound

Many sciences and technologies have undergone revolution in the last few decades. Metrology, the science of weights and measures, has been part of this revolution. No longer is the standard of length defined in terms of a bar or rod held in a particular national institution. Rather it is

...the length of the path travelled by light in a vacuum during a time interval of $1/299\,792\,458$ of a second. (BIPM [17])

The definition of the second is

The duration of 9 192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom. (BIPM [17])

That is, length is defined indirectly through the definitions of the speed of light and time. Clearly an error in the definition of the speed of light results in an error in the definition or determination of length.

While not of the same fundamental importance, a similar situation applies to sound which can be viewed as just another propagating wave even though it requires a compressible medium for propagation. That is, the determination of distance or length is dependent upon a precise knowledge of the speed of propagation under a wide range of environmental conditions.

The aim of this chapter is to validate our knowledge of the speed of sound and how the environment, as described by temperature, pressure and humidity, modifies this speed. This allows us to minimise the errors occurring in the deduced distances which are used by a *robotic agent* to perform navigation or generate a map.

Given the importance of the speed of sound for computing the correct [DOF](#), it is now pertinent to review the formulae available for calculating it, appraise their claimed precision, and describe and calibrate equipment for measuring the environmental parameters upon which it depends.

The material in this chapter is supported by [Appendix E](#), which presents coherent review of the properties of air and the equations which may be used to model the speed of sound and the absorption.

12.1 Literature

This section briefly reviews different formulae that appear in the literature for calculating the speed of sound and compares their forms. The formulae will be revisited with a view to practical

application in Section 12.2.

Wong [140] lists some historical investigations of the speed of sound from 1919 to 1985 and summarises their results. Wong also points out that with modern measurements of gas constants, the speed of sound can be determined indirectly from the thermodynamic properties of the air (Wong [140]),

$$c = \sqrt{\frac{\gamma R T_a}{M}}. \quad (12.1)$$

Equation (12.1) is most often seen as

$$c_1 = 20.05 \sqrt{T_a} \quad (12.2)$$

(e.g. Poole [113]) where the numerical values of γ , R and M are assumed. Wong [140] states that the specific heat ratio, γ , is a dominating factor in the computation of the speed of sound, and that values ranging from $331.297 \text{ m} \cdot \text{s}^{-1}$ to $331.641 \text{ m} \cdot \text{s}^{-1}$ are possible with published values of γ for air at 273.15 K and 1013.25 kPa (Standard Temperature and Pressure (STP)). This variation is overlooked for the sake of simplicity by forms such as (12.2) and tables such as those commonly attributed to Kaye and Laby [76].

An expansion of (12.1) using the equation of state for a real gas up to the second virial coefficient B is (Cramer [26], Wong [140])

$$c = \sqrt{\frac{\gamma R T_a}{M} \left(1 + \frac{2 p_s B}{R T_a} \right)} \quad (12.3)$$

where p_s is the static air pressure. A common polynomial approximation to (12.3) to account for the partial pressure of atmospheric water vapour is (due to Kleeman [81])

$$c_2 = 20.05 \sqrt{T_a} + h_r \left(\frac{1.0059 \cdot 10^{-3}}{+1.7776 \cdot 10^{-7} \cdot (T_a - 275.37)^3} \right). \quad (12.4)$$

That is, there is no modification of the specific heat ratio γ .

Both Poole and Kleeman provide forms which include modifications for pressure and relative humidity. This form is

$$c_3 = 20.05 \sqrt{\frac{T_a}{1 - 0.00379 \cdot h_r \cdot \frac{p_{sat}}{p_s}}} \quad (12.5)$$

which is equivalent to

$$\sqrt{\frac{\gamma R T_a}{M} (1 + \beta)}.$$

Compared to (12.3), this indicates that specific heat ratio γ is modified. The major difference between the forms of Wong and Cramer and those of Kleeman and Poole is the extent and manner in which the accommodation is performed.

There appears to be a clear preference in the robotics literature for the forms of Kleeman and Poole rather than the forms of Cramer and Wong. It is suggested that this is primarily due to the difficulty of choosing γ (see Wong [141]), and accounting for all atmospheric constituents including carbon dioxide.

Another method of estimating the speed of sound is available to a mobile *robotic agent* which is performing sonar sensing continuously in its environment. Kleeman [79] includes the speed of sound as a parameter in a Kalman filter model incorporating the positions of targets measured by the sonar system and odometry information from the mobile robot platform, and estimates it from the sonar measurements. This method is not available in this work as the fundamental properties of the sonar system must be determined.

12.2 Formulae for Calculation

Formulae are now presented in a form suitable for calculation. Their precision is analysed and the implications upon the determined DOF are described. The errors which are considered in this section are only those due to the approximations inherent in the formulae, and the experimental errors in the data from which the formulae were derived. There will be additional error terms to be incorporated from measurements of the environmental parameters.

12.2.1 Poole's Formulae

Temperature Only

The first approximation c_1 to the speed of sound in $\text{m} \cdot \text{s}^{-1}$ using only temperature is (from Poole [113])

$$c_1(T_a) = 20.05\sqrt{T_a} \quad (12.6)$$

where T_a is the absolute air temperature in Kelvin (K). There is some confusion in the literature about how to convert temperatures in degrees Celsius ($^{\circ}\text{C}$) to Kelvin (see BIPM [18]). The International Temperature Scale of 1990 [115] defines temperatures in $^{\circ}\text{C}$ to be offset from temperatures in K by the ice point $T_{ice}=273.15\text{ K}$. It is common to see the triple point of water $T_3 = 273.16\text{ K}$ used instead, as in Poole [113].

Poole claims that (12.6) is accurate to 1% for most conditions. This translates to $\sigma_c = \pm 3\text{ m} \cdot \text{s}^{-1}$ at STP. The error per metre, σ_d , in the measured DOF due to the uncertainty, σ_c , in the speed of sound may be calculated from $\sigma_d = \sigma_c/c$. Choosing to use a value of $c = 343.4\text{ m} \cdot \text{s}^{-1}$, corresponding to dry air at STP, produces an error of $\sigma_d = \pm 10\text{ mm} \cdot \text{m}^{-1}$. This is clearly inadequate for use in a precision sonar system.

Temperature and Humidity

A better estimate for the speed of sound, c_2 , which uses the relative humidity, h_r , (the ratio between the partial pressure of water vapour and the saturation pressure (as a percentage) is provided by Kleeman [81], converted from a formula in Fahrenheit by Poole [113]):

$$c_2(T_a, h_r) = c_1(T_a) + h_r \cdot \left(\frac{1.0059 \cdot 10^{-3}}{+1.7776 \cdot 10^{-7} \cdot (T_a - 275.37)^3} \right). \quad (12.7)$$

This formula is an approximation which may be used when the static pressure p_s is not available. Poole claims that this formula is accurate to 0.1%, which corresponds to $\sigma_c = \pm 0.3\text{ m} \cdot \text{s}^{-1}$ and $\sigma_d = \pm 1\text{ mm} \cdot \text{m}^{-1}$. This is marginal for use in a precision sonar system.

Temperature, Pressure and Humidity

When measurements of the pressure p_s , in kPa, are available in addition to the temperature and humidity, the speed of sound may be calculated by (due to Poole [113])

$$c_3(T_a, h_r, p_s) = 20.05 \sqrt{\frac{T_a}{1 - 0.00379 \cdot h_r \cdot \frac{p_{sat}}{p_s}}}. \quad (12.8)$$

A formula for the saturation vapour pressure p_{sat} as a function of temperature may be found in Appendix E.2. Poole does not provide a statement of precision for (12.8), claiming only that it is the best available.

12.2.2 Cramer's Formula

The most recent published data available on the speed of sound is Cramer [26]. Cramer's paper includes a polynomial approximation to (12.3) which incorporates the effects of the 13 most important constituents of the standard atmosphere (ISO 2533-1975[65]). Cramer's formulation is summarised in Appendix E.5 with a formula for the speed of sound, c_0 , for pressure waves with zero frequency being provided as (E.11).

While the proportion of CO_2 may be fixed in the standard atmosphere, the amount varies seasonally by 9 ppm due to plant uptake (Kimball [77]). Cramer includes the mole fraction of CO_2 as a parameter in his formulation to account for this. The standard atmosphere specifies that the mole fraction of CO_2 is $x_c = 314$ ppm or 0.0314 %. As the change in the speed of sound caused by the change in CO_2 levels is quite small (see Figure E.2), the standard value has been used in the computations of this work.

Cramer claims that the uncertainty in his calculations is ± 300 ppm or $\sigma_c = \pm 0.10 \text{ m} \cdot \text{s}^{-1}$, giving a DOF uncertainty of $\sigma_d = \pm 0.3 \text{ mm} \cdot \text{m}^{-1}$ for (standard) dry air at STP. Wong [141], however, claims that Cramer neglected some sources of error, and that the formulae are only accurate to ± 545 ppm or $\pm 0.18 \text{ m} \cdot \text{s}^{-1}$, giving a DOF uncertainty of $\sigma_d = \pm 0.5 \text{ mm} \cdot \text{m}^{-1}$ for dry air at STP.

Cramer's formulation for the speed of sound has been chosen over the form of Poole for use in this thesis, as Cramer's work is more recent and was published in peer reviewed journals.

12.2.3 Dispersion

The speed of sound c_0 given by Cramer applies to a pressure wave with zero frequency. For use with sound waves with a frequency f the speed must be corrected by a dispersion factor which takes the frequency into account. The formulae are reproduced for convenience in Appendix E.6 in (E.14). The symbol for the speed of sound corrected for dispersion is c_ϕ .

The variation from the zero frequency speed c_0 due to dispersion, determined by examining Figure E.2, is between $+0.1 \text{ m} \cdot \text{s}^{-1}$ and $+0.15 \text{ m} \cdot \text{s}^{-1}$, which is accounted for by using (E.14) instead of (E.11). This correction reproduces the effect seen in tabulated values of the speed of sound (Evans [39]).

The variation in the speed of sound over the frequency range of 50-100 kHz used by the sonar signals in this thesis is $0.05 \text{ m} \cdot \text{s}^{-1}$ (see Figure E.2). Noting that it increases monotonically over

the range 50-100 kHz, the error may be minimised by evaluating c_ϕ at the central frequency of 75 kHz. This is the non-dispersive assumption, which causes an error of only $\pm 0.025 \text{ m} \cdot \text{s}^{-1}$ over the frequency band of the signal.

Cramer does not give uncertainties for the derivation of the dispersion correction formula (E.2). For this reason their contribution to the total error budget will be neglected.

12.3 Error Analysis

The errors, σ_c , in the speed of sound c calculated by Cramer's formula (E.11) due to errors σ_{T_c} , σ_{h_r} and σ_{p_s} in the measurement of the environmental parameters will now be determined. A differential analysis of (E.14) is prohibitively complex. Therefore a numerical simulation is performed for conditions commonly encountered in our laboratory, located in a coastal area. The central values for the analysis are

$$p_s = p_{s0}, T_c = 23^\circ\text{C}, h_r = 65\%. \quad (12.9)$$

As no instruments were available for measuring the amount of CO_2 present, the amount specified by the standard atmosphere was assumed, i.e. $x_c = 314 \text{ ppm}$. By varying each of the parameters in turn, it was found that

- an error in temperature, T_c , of $+0.62^\circ\text{C}$ produces a $+0.4 \text{ ms}^{-1}$ error in the speed of sound,
- an error in relative humidity, h_r , of $+28\%$ produces a $+0.4 \text{ ms}^{-1}$ error in the speed of sound,
- an error in static pressure, p_s , of $+84 \text{ kPa}$ produces a -0.4 ms^{-1} error in the speed of sound, and
- an error in the mole fraction of CO_2 , x_c , of $+4472 \text{ ppm}$ produces a -0.4 ms^{-1} error in the speed of sound.

The typical range of variations through the course of one experiment may be seen in Figure 13.6. While these quantities may not change much from moment to moment, there are marked differences from week to week or from location to location. (The air pressure, for example, is typically near 100 kPa at sea level near Sydney, but decreases to around 80 kPa at Cramer's laboratory on a mountain top in South Africa.) The list above is ranked by importance to the error calculations.

12.3.1 Relative Importance of Error Terms

The above analysis indicates that the speed of sound is most sensitive to errors in temperature. The humidity and the pressure are of secondary importance. The effect due to the molar fraction of CO_2 is small. The effect shown here is 14 times the value assumed in the standard atmosphere. The effect of each term may also be seen in Figure E.2.

12.3.2 Combining Errors

The errors from the various sources, including the approximations of the formulae discussed in Section 12.2.2, combine quadratically. Assuming that the derivatives of c_ϕ (E.14) with respect




Figure 12.1: Graph of the error in the speed of sound, as calculated by (12.10), as a function of the temperature error σ_{T_c} , the dominant component. The non-zero y-intercept is caused by the error attributed to the formulation. The second most important error term, σ_{hr} , causes no perceptible difference from the line shown over the range $\sigma_{hr} = 0$ to 1%. The other sources of errors make negligible contributions, which is reinforced by inspection of the magnitudes of the derivatives in (12.10).

to each parameter remain sufficiently linear over the region of interest (see Figure E.2 for validation), the derivatives may be evaluated numerically from the variation data previously discussed. Assuming that the measurement errors from each source are statistically independent, the formula for combining the errors from the various components used by (E.14) is

$$\sigma_{c\phi}^2 = (0.18)^2 + (0.645\sigma_{T_c})^2 + (0.014\sigma_{hr})^2 + (-4.8 \times 10^{-3}\sigma_{ps})^2 + (-89.4 \times 10^{-6}\sigma_{xc})^2, \quad (12.10)$$

where the formulation assumes that σ_{ps} is measured in kPa and that σ_{xc} is measured in units of ppm. The first term is due to the approximations of the formulation itself. The terms are ordered by their contribution to the total error (the contribution of typical temperature errors is smaller than the formula error). A graph of (12.10) is shown in Figure 12.1. The errors described by 12.10 are only valid for environmental parameters within the vicinity of those defined in (12.9).

12.4 Calibrating the Environment Sensor

The temperature and humidity, being the most important parameters in Cramer's model for the speed of sound, must be measured by an accurate instrument. The precision must be quantified so that the error in the speed of sound may be determined. The contributions made by errors in pressure and CO_2 concentration are negligible, in this context, according to (12.10), and are therefore neglected.

12.4.1 The Temperature Scale

It is clear that temperature is the major component. Temperature, a non-dimensional quantity, is one of the seven basic physical quantities maintained by the BIPM. The thermodynamic scale, Kelvin, is more fundamental than the Celsius scale. However, temperatures measured in Kelvin may be converted to Celsius by subtracting $T_{ice} = 273.15 \text{ K}$.

The International Temperature Scale-1990 [115], which covers the region of interest of this study, is defined in terms of platinum resistance thermometers. Unfortunately there are a number of practical difficulties in using platinum resistance thermometry in sonar experiments.

- A calibrated linear output of sufficient sensitivity must be determined or obtained. This problem is most commonly overcome by contracting a standards laboratory for a schedule of calibrations of the total system. This was not possible in this project.
- The platinum resistance element must be protected from light.
- The response time must be reduced so that observed temperatures accurately track the variations in atmospheric temperature.

In order to overcome these difficulties a calibration procedure was developed which used a meteorological quality mercury in glass thermometer, Dobbie Instruments part number 526.10946, with 0.1°C graduations. (The thermometer was provided with a calibration certificate.) The index for the scale was validated using a de-mineralised ice bath which has a known value of 0°C if the small effects of pressure (Sostmann [123]) are neglected. The scale of the mercury in glass thermometer was assumed to be correct. It was also assumed that the lag or delay of this thermometer, when placed in a draught, was essentially zero. Thus a temperature scale could be transferred to temperature transducers with high precision by direct comparison of observations from the de facto standard mercury in glass thermometer and the output of the environmental sensor.

We tested two different probes for measuring temperature:

1. a Vaisala HMP-35A probe (see Figure 5.10b) which uses a 100Ω platinum resistor (PT100), and
2. a National Semiconductor LM335 zener diode temperature transducer (see Figure 12.2).

The platinum resistance thermometer in the HMP-35A was found to have inadequate response characteristics, due to the location of the thermometer in a housing which trapped a pocket of air. The LM335 zener diode was found to have superior response characteristics, especially when it was placed in a suitable draught. The LM335 zener diode output was converted to a digital signal using an ADC card described in Section 5.3. This digital output had a quantisation of 0.01°C . The calibration of this output was by direct comparison. In this technique the output from both the mercury in glass thermometer and the digital ADC were recorded while the temperature fluctuated in the environment. A typical set of these observations is shown in Figure 12.3. A linear regression was performed to model the ADC voltage output as a function of temperature. The standard deviation of the temperature error was determined to be $\sigma_{T_c} = 0.3^\circ\text{C}$, which encompasses 68% of the results. As the thermometer was accurate to 0.1°C (according to manufacturer's specification), the step size has not impacted upon the precision obtained.




Figure 12.2: The LM335 zener diode temperature transducer mounted in the draught of a standard computer fan. The transducer is held in place by wires. The black box with a white sticker is a Vaisala PTA427 pressure transducer.




Figure 12.3: The LM335 calibration graph of the voltage measured by the [ADC](#) and the temperature recorded by the reference mercury in glass thermometer. The line is the least squares fit to the data points. The values near 7.1 and 8.1 are important in constraining the linear relationship.




Figure 12.4: The relative humidity sensor calibration graph of the voltage measured by the [ADC](#) vs the relative humidity measured by the Vaisala HM-34. The line is the least squares line of best fit to the data points.

12.4.2 Humidity

The relative humidity, h_r , is the ratio of the partial pressure of water vapour to the partial pressure of saturated air. It is most commonly expressed as a percentage. Equation (12.10) indicates that while the effect of humidity is not as large as the effect of temperature on the speed of sound it is nevertheless an important factor. The original goals for determining humidity were similar to the original goals for temperature. That is, a calibrated standard, a Vaisala HM-34 probe, can be used to calibrate the working instrument, a Vaisala HMP-35A. Both of these instruments use humicap transducers. As was the case in the temperature calibration, direct observations of the [ADC](#) output of the HMP-35A and the output of the HM-34 were made, in a fluctuating environment, and a regression analysis performed. Figure 12.4 shows a typical set of observations. It is estimated (from the linear regression) that the precision of the determined humidity is $\sigma_{hr} = 0.4\%$ at the 68% confidence level.

12.4.3 Pressure

The requirement to observe pressure to any moderate precision is not evident from (12.10). However we also observed pressure with a precision of at least 0.1 kPa using a Vaisala PTA427 precision pressure transducer (Figure 12.2) which was supplied with a calibration certificate. The rationale for observing pressure in this way was partly to ensure a quality description of the environment in which the sonar experiments were to take place and partly to provide diagnostic parameters for unexpected outcomes.

12.4.4 Experiments

It is known from (12.10) that temperature must be precisely determined in order to precisely model the speed of sound. In Section 12.4.1 we developed a calibration technique that used a mercury in glass thermometer as the de facto standard rather than the Vaisala Platinum thermometer. The principal reason for this change was the level of induced lag and filtering that was built into the design of the Vaisala unit. The influence of thermal lags became apparent during the calibration experiments when we attempted to resolve differences between model determined values for the speed of sound and experimentally determined values using a purpose designed experimental setup. The need for a system which is sensitive to the fluctuations in temperature and is able to respond quickly to these fluctuations is shown in Figure 12.5. This figure shows that agreement between model and experimental values for the speed of sound requires simultaneous observation of the environment with each sonar measurement. Furthermore the probes must be precise and have minimum lag characteristics.

The design, construction and operation of the calibration device used to determine the speed of sound is given in Appendix F.

12.4.5 Uncertainty

Combining the experimental errors of the environmental parameter measurements using (12.10) yields a value of $\sigma_{c\phi} = 0.27 \text{ m} \cdot \text{s}^{-1}$ for the one standard deviation or 68% confidence level. The two largest contributions are from the formulation error and the temperature error.


It must be recognised that this error figure defines a normal distribution of errors, and that the probability of obtaining errors larger than this is about 32%, according to the usual laws governing normal distributions (Mendenhall [91]).

12.5 Conclusion


The precision of the speed of sound used to convert TOF to DOF is crucial to the precision of the computed DOF. The precision of the speed of sound depends upon using a precise model which accurately reflects the properties of the air in which the sonar is to be used. The model derived by Cramer [26] is such a model. The precision of the speed of sound computed by such a model depends in turn on the precision with which the environmental parameters are measured, with the temperature having the most impact upon the result. The adequacy of the measurement of the environmental parameters is verified by comparing the modelled speed of sound computed from such measurements with the speed of sound measured directly by the sonar. By doing so, deficiencies in the measurement of the environmental parameters were detected and rectified. The calibration device may be used periodically to ensure that the precision of the environmental parameter measurements is maintained.

12.5.1 Acknowledgements and Contributions

The formulae for calculating the speed of sound were brought together from various sources as referenced in Section 12.1. The unification of the formulae in Appendix E is the contribution of



(a) case mounted LM335



(b) fan mounted LM335

Figure 12.5: The variation in the speed of sound, measured by the CTFM-YW sonar (c_s , marked with +) and calculated from the measurements of the environmental parameters (c_ϕ , marked with \times). The transducer configuration was (a) mounted on the aluminium case and (b) mounted in the draught of a fan. The method of determining the speed of sound from the sonar measurements is described in Appendix F. In (a), the speed of sound determined from the model can be seen to lag behind that determined from the sonar, indicating that the response time of the temperature transducer is inadequate. In (b), the variations in the speed of sound measured by both sources agree, indicating that the response time is sufficient.

this author. The design of the calibration equipment, as described in Appendix [F](#), is the work of this author. Thanks are due to Mr. Kettlewell for his patient assistance with the construction of this device.

Chapter 13

DOF and Power Measurement Performance

This chapter measures the performance of the five distance of flight sub-systems, shown in Figure 1.2. The measurement properties of the distance of flight and echo power are quantified for both the left and right sonar channels. The primary measures of performance, which were discussed in Chapter 2, are accuracy, precision and resolution. One additional measure of performance, the correlation between the left and right channels, was also assessed. In all, six different quantities were assessed for each sonar DOF sub-system. These are:

1. DOF Accuracy: The difference between the sonar measured distance of flight and the ruler measured or geometrical distance of flight.
2. DOF Precision: The variation within a set of measurements of the distance of flight.
3. Echo Power Precision: The variation within a set of measurements of the echo power.
4. DOF Cross-correlation: The amount of random variation which is common to both the left and right channel distance of flight measurements.
5. Echo Power Cross-correlation: The amount of random variation which is common to both the left and right channel echo power measurements.
6. Noise Level and Signal to Noise Ratio: The noise level at the output of each sonar system is important for determining the [signal to noise ratio](#), and hence the detection threshold and the maximum operable range.
7. Range Resolution: The minimum distance by which two objects must be separated before their echoes can be distinguished.

Items 1, 2 and 7 directly characterise range measurements made by the sonar system. Item 6 is used to normalise the power measurement results and assess the [SNR](#). Items 3, 4 and 5 are used as intermediate characterisations which determine the bearing measurement properties. These three measures will be used in Chapter 14.

These seven different performance measures are assessed using only two different experimental configurations, which are shown in Figure 13.1. Configuration (a) provides data from which the first five performance measures are assessed. Configuration (b) provides data from which the sixth quantity, range resolution, is assessed. In these experiments, the principal characteristics of the experimental setup can be defined by

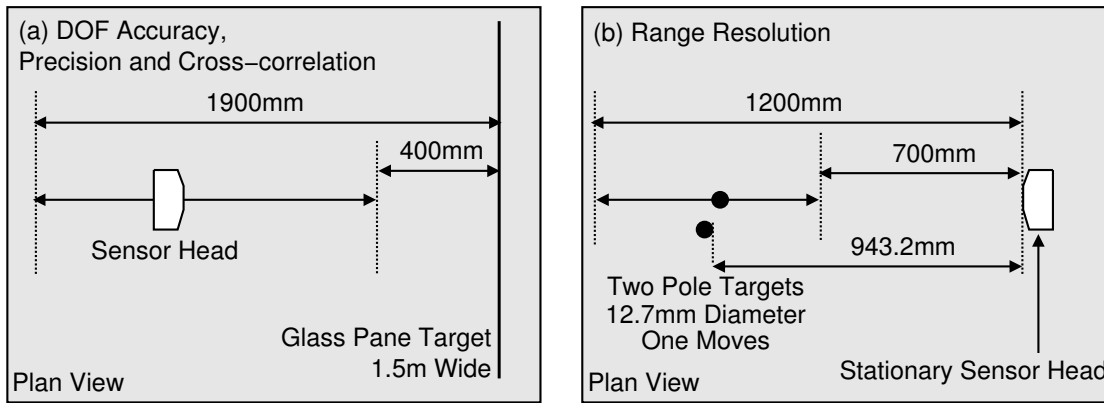


Figure 13.1: Two experimental configurations designed to characterise the measurement performance of the distance of flight sub-system. (a) Depicts a sonar head with the principal direction of its transmitter being normal to a plane. The sonar head is moved perpendicularly to the plane. (b) Depicts a fixed sonar head with two targets, one of which is moved in range while the other is fixed.

1. the distance, along the perpendicular to the plane, between the plane and the transmitter of the sonar head (Figure 13.1a), and
2. the range difference between two targets (Figure 13.1b).

A significant terminology problem is usually associated with these experiments due to the fact that all sonar systems, like all non-transponder radars, measure a distance of flight which is generally about twice the range, the Euclidean distance from the transmitter to the target. These DOFs are subsequently transformed into ranges in Chapter 14. However, the concept of range resolution, directly analogous to distance of flight resolution, is often used in preference to DOF resolution, as it can be readily interpreted as the distance separating two objects (with similar bearings).

Section 13.1 briefly describes the two experiments and the data acquired. Section 13.2 deals with the DOF accuracy, precision and cross-correlation experiment (Figure 13.1a), while Section 13.3 deals with the range resolution experiment (Figure 13.1b). Finally Section 13.4 summarises the performance of the five sub-systems against the six assessment items.

13.1 The Experiments

Sensor head 1 was used to perform all the experiments described in this chapter. The sonar head has one transmitter and four receivers, but only two receivers, the left and right, were used for these experiments (Section 5.1.2).

13.1.1 Five Types of DOF Sub-system

Table 13.1 shows the two experimental configurations and both the sonar classes and their associated sub-classes that make up the performance table. The table shows that 30 items can be

Table 13.1: Experiments required to quantify **DOF** and echo power measurement performance. Each of the six performance measures are assessed against each of the five different **DOF** sub-systems through a total of four experimental data sets, MFa, MFb, CTFMa and CTFMb.

Experimental Configuration	Measurement Type	Class Sub-Type	MF		CTFM		
			QD	UD	DFT	YW	LSMYW
(a)	1. DOF Accuracy		MFa	MFa	CTFMa	CTFMa	CTFMa
	2. DOF Precision		MFa	MFa	CTFMa	CTFMa	CTFMa
	3. Echo Power Precision		MFa	MFa	CTFMa	CTFMa	CTFMa
	4. DOF Cross-correlation		MFa	MFa	CTFMa	CTFMa	CTFMa
	5. Echo Power Cross-correlation		MFa	MFa	CTFMa	CTFMa	CTFMa
	6. Noise Level and SNR		MFa	MFa	CTFMa	CTFMa	CTFMa
(b)	7. Range Resolution		MFb	MFb	CTFMb	CTFMb	CTFMb

evaluated with just 4 experiments, since sub-classes within the same class use identical transmit signals, and hence the echo signals are interchangeable. The raw echo signals from the four experiments shown in Table 13.1 are stored in computer files. The **MF** data sets are subsequently processed by the MF-QD and MF-UD sonar types, which were defined in Chapter 8. The CTFM data sets are subsequently processed by the **CTFM** dual demodulation sonar system described in Chapter 9, and the demodulated output signal is analysed by three different spectral estimation systems. These are the **DFT**, discussed in Chapter 10, and the **Yule-Walker** and **least squares modified Yule-Walker** estimators discussed in Chapter 11. The processing of the same sonar record by the different sonar systems is made possible by the file storage mode of operation, which also allows direct performance comparisons to be made between different sonars within the same class.

13.1.2 Configurations

Figure 13.2 shows the experimental setup using a glass pane, corresponding to configuration (a) in Figure 13.1. This configuration provides data for the six performance measures listed under (a) in Table 13.1. The glass pane was selected as the plane target due to its flatness and availability. This experiment and its results are described in Section 13.2.1.

The configuration shown in Figure 13.3, corresponding to configuration (b) in Figure 13.1 and Table 13.1, was used only to investigate range resolution. This experiment and its results are described in Section 13.3.

13.1.3 The Observations

During the **DOF** precision and range resolution experiments, the following information was recorded for each observation.

1. The position and orientation of the positioner stage, from which the range to the target and the geometrical **DOF** may be deduced.
2. The environmental parameters that affect the velocity of sound and the absorption of the air as discussed in Chapter 12.




Figure 13.2: The configuration used for range and bearing precision experiments. The target is a glass pane which has been attached to a white board for support. The sonar head has been mounted upon the positioner so that the distance and angle between the sonar head and the plane can be controlled. The range movement is used in this chapter. This experimental configuration can also be used to measure bearing, by using only the rotational movement. This will be described in Chapter 14.




Figure 13.3: The experimental setup for the range resolution experiment, showing the sonar head mounted upon a fixed end of the positioner, one fixed pole, and one movable pole mounted upon the positioner stage. The fixed pole is suspended from the ceiling by a string (invisible in the photograph). If left this way, the pole vibrates, necessitating the use of a clamp for damping.

3. The received sonar echoes from which, using a given sonar sub-system, the echo TOF and power may be extracted. The speed of sound c_ϕ from Section 12.2 was used to convert TOF to DOF.

For experimental configuration (a), only the echo closest to the ruler or geometrical DOF was stored. Where there was more than one echo to choose from, the most powerful echo was selected. In configuration (b), all echoes within an interval were stored for later analysis.

13.2 The DOF Precision Experiment

The aim of the DOF experiment is to assess the six measurement types marked (a) in Table 13.1. This data is then used to determine several key datums as listed below:

1. range accuracy and precision (derived in Section 14.1.3 of the next chapter),
2. bearing precision, by both IDD and IPD methods (derived in Sections 14.1.3 and 14.2.4 of the next chapter),
3. the background noise level, and
4. the maximum DOF of each system.

The DOF experiment was conducted using the configuration shown in Figure 13.2. The sonar head was initially placed so that its transmitter was 400mm from the plane. During the experiment it was moved to 1900mm in steps of 10mm. At each of the 151 positions, 10 observations were taken.

The two sets of echo measurements were subsequently processed by the five different sonar systems to extract the DOF and echo power information. The precision of the DOF is discussed in Section 13.2.1, while the precision of the power measurement is discussed in Section 13.2.2. The cross-correlation between the two channels of both the echo DOF and power fit residuals provides us with a measure of the variations in both left and right channels due to turbulence and noise in the the air path. The cross-correlation of the residuals of the DOF and echo power fit data will be discussed in Section 13.2.3.

The underlying noise level determines the maximum range of the sonar system, as seen in Figure 4.3. It also affects the precision with which we can expect to make measurements, as mentioned in Section 4.7. Therefore the noise level is measured in Section 13.2.4.

13.2.1 DOF Precision Analysis

Geometric DOF

The geometric DOF, $d_g(x)$, is computed according to

$$d_g(x) = \sqrt{(2x + d_T)^2 + d_R^2}, \quad (13.1)$$

where the variables are

- x , the distance of the sonar head transmitter from the glass pane, as reported by the positioner, and
- d_T , the distance of the receiver centres behind the transmitter, as shown in Figure 7.12, and
- d_R , half the distance between the two receivers, as shown in Figure 7.12.

The term $2x$ is due to the virtual source construction. The path lengths to the two receivers are assumed to be equal as the sonar head was aligned to point along the normal to the window pane.

The values used for the sonar head dimensions were $d_T = 2.7$ mm and $d_R = 13.5$ mm, for reasons which are made clear in Section 14.1.4.

Acoustic DOF

The acoustic DOF is provided by the sonar DOF sub-system under test. The DOF sub-system will produce many echoes, but in this experiment only the echo corresponding to the glass pane is required. This is assumed to be the strongest echo within ± 20 mm of the geometric DOF. The DOF, d_i , corresponding to observation number i was calculated from $c_i \tau_i$. The speed of sound, c_i , was calculated using the frequency corrected speed of sound formula c_ϕ , defined in Section 12.2, using the environmental parameters recorded during observation number i .

DOF Precision Experimental Results

We expect that a plot of the acoustic DOF, $c_i \tau_i$, against the geometric DOF, $d_g(x_i)$, would yield a straight line with a slope of 1. However, the errors would be unresolvable on such a plot. Therefore the tuple

$$\{d_g(x_i), c_i \tau_i - d_g(x_i)\}$$

is graphed instead, allowing residual analysis. The residual plots are shown as dots in the left column of panels in Figures 13.4 and 13.5 for the left and right channels respectively.

The line in the left column of panels, drawn in red, is the least squares line of best fit having slope m and y-intercept b ,

$$y = md + b. \quad (13.2)$$

The residuals about this line are plotted in the middle column of panels. The third column of panels shows a histogram of the middle panel residuals. The mean and the standard deviation of the residuals are shown on the histogram. This plot may be used to determine whether the residuals are normally distributed.

The environmental parameters which were recorded during the matched filter and CTFM experiments are shown in Figures 13.6 and 13.7.

Converting Line Parameters to Offset Time and Speed of Sound Difference

The average speed of sound \bar{c}_s observed by the sonar experiment is related to the average speed of sound \bar{c}_ϕ computed from Cramer's model, described in Section 12.2, by

$$\bar{c}_s = \bar{c}_\phi + \Delta c. \quad (13.3)$$

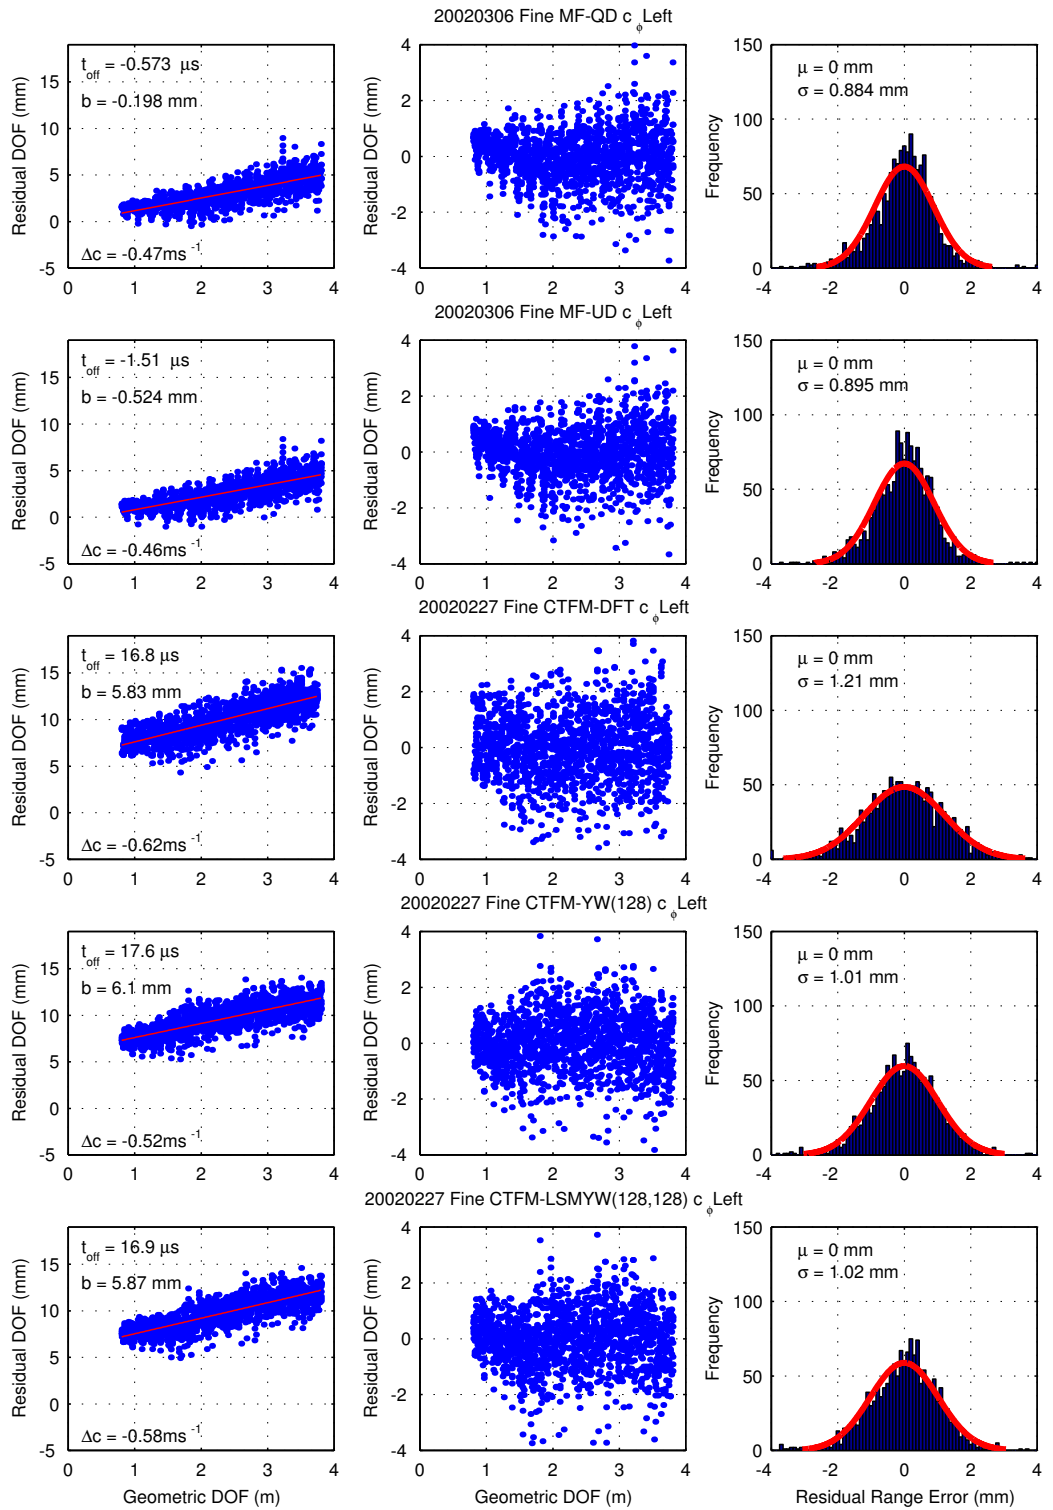


Figure 13.4: Results of the DOF precision experiment for the left channel. The title of the middle column of graphs indicates the experimental data set and the sonar system which was used. The first column is the acoustic DOF minus the geometrical DOF. The second column is the residuals around the red line. The third column is a histogram of the residuals in the second column. See text for more information.

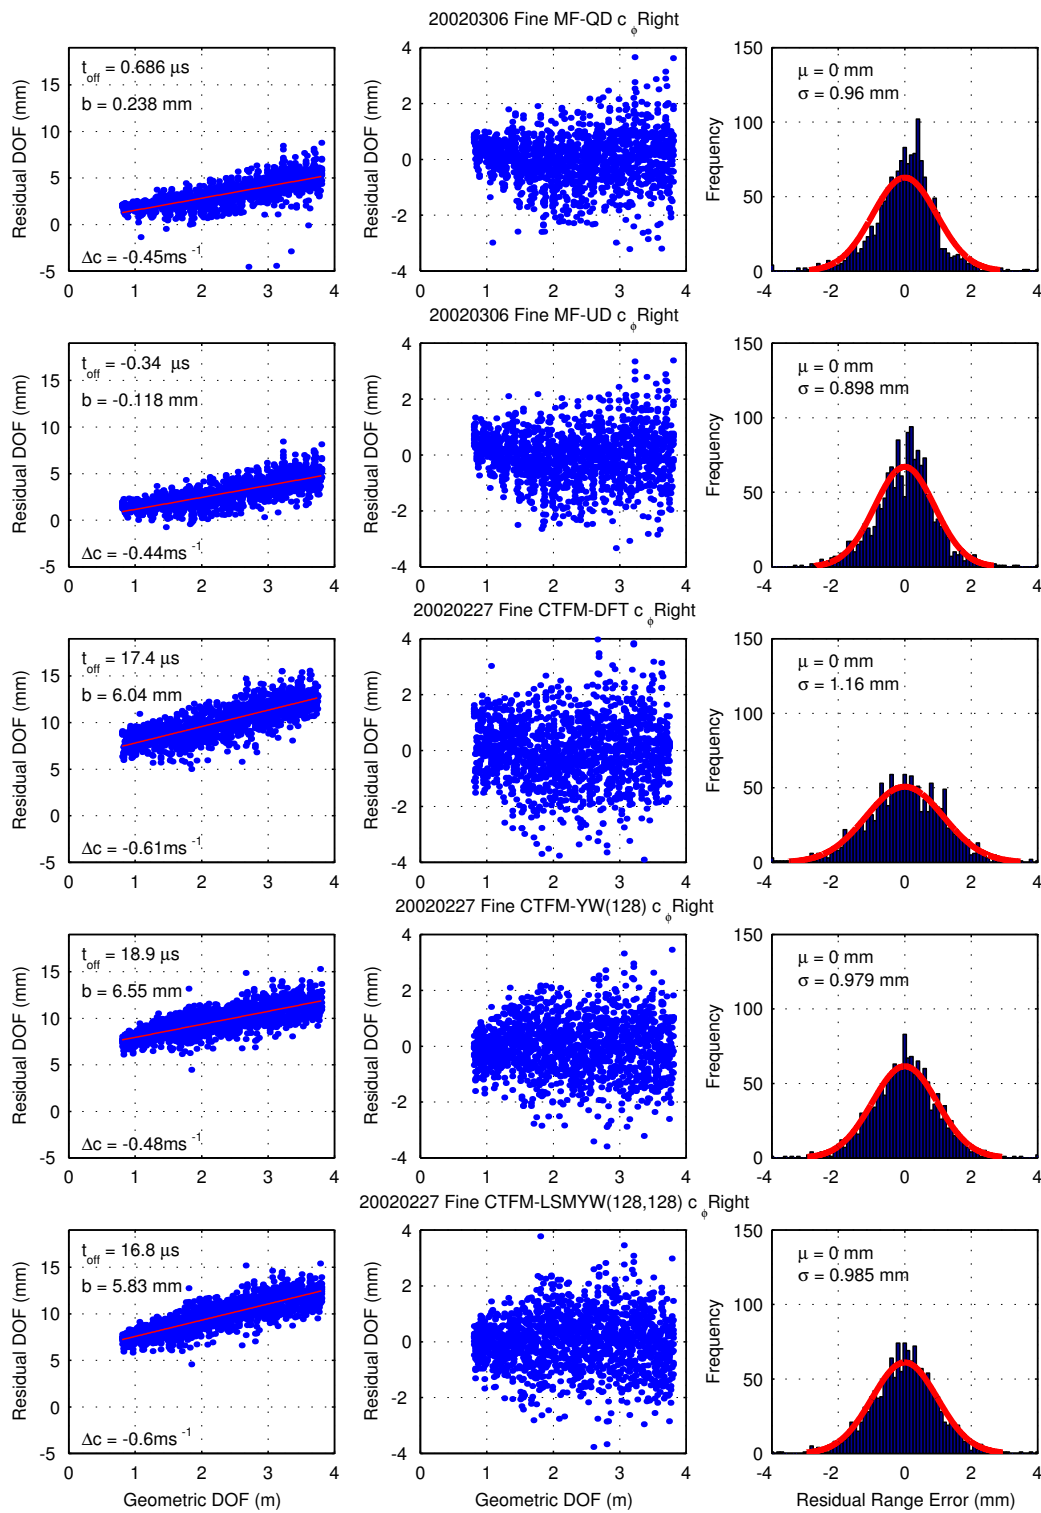


Figure 13.5: Results of the DOF precision experiment for the right channel. The title of the middle column of graphs indicates the experimental data set and the sonar system which was used. The first column is the acoustic DOF minus the geometrical DOF. The second column is the residuals around the red line. The third column is a histogram of the residuals in the second column. See text for more information.




Figure 13.6: Environmental parameters recorded during the matched filter range precision experiment. As the experiment was carried out from $d_g = 800\text{ mm}$ to $d_g = 3800\text{ mm}$, the time scale of this plot may be compared to the geometric DOF scales in Figure 13.4 and Figure 13.5. The speed of sound is calculated using c_ϕ described in Section 12.2.




Figure 13.7: Environmental parameters recorded during the CTFM range precision experiment. As the experiment was carried out from $d_g = 800$ mm to $d_g = 3800$ mm, the time scale of this plot may be compared to the geometric DOF scales in Figure 13.4 and Figure 13.5. The speed of sound is calculated using c_ϕ described in Section 12.2. The time scale on this figure is different from that in Figure 13.6 because the CTFM sonar system takes longer to run.

Table 13.2: The offset times t_{off} obtained with the different sonar types.

Sonar System	t_{off} (μ s)	
	Left	Right
MF-QD	-0.57	0.68
MF-UD	-1.51	-0.34
CTFM-DFT	16.8	17.4
CTFM-YW(128)	17.6	18.9
CTFM-LSMYW(128,128)	16.9	16.8

The value of Δc is determined by its relationship to the slope m of the linear fit,

$$\Delta c = \frac{-m\bar{c}_\phi}{m+1}. \quad (13.4)$$

The y-intercept b , measured in mm, may be converted to an offset time, t_{off} , according to

$$t_{off} = \frac{b}{\left(\bar{c}_\phi\right)}. \quad (13.5)$$

The values obtained for Δc and t_{off} are drawn on each plot together with b .

The effect of the choice of the speed of sound model upon the slope of the residuals m may be seen in Figure 13.8. This demonstrates graphically the relationship of (13.4) between m and Δc .

Offset Time

The measured values of the offset time, t_{off} , are recorded for each sonar system in Table 13.2. The table indicates that the offset is class dependent. Specifically, the offset is small for the **matched filter** class and large for the **CTFM** class. The differences between the left and right channels represent a combination of experimental uncertainty and the differences in the physical mounting of the two channels. The experimental uncertainties appear to be the larger of these two effects. At first this may appear to be due to a problem in the implementation of one of the sonar systems, but system tests using simulated echo signals demonstrate that both classes of sonar system provide accurate time of flight measurements. Furthermore, the changes in t_{off} between each of the spectral estimators are quite similar. It is hypothesised that the cause is not attributable to the software but to the signal and the mode of handling the received signal in the receiver, as the effect is only a function of the signal class and independent software systems were implemented for each of the sub-classes.

Since the offset time, t_{off} , provides only a small bias in the case of **matched filter** sonars, but makes a much larger contribution to the total bias in **CTFM** systems, it is important to calibrate t_{off} and remove it if accurate distance measurements are required when using the **CTFM** class. The range precision experiment may be used as a calibration method, although it takes some time to execute. The offset time, t_{off} , may be also be obtained directly using the calibration device shown in Appendix F.

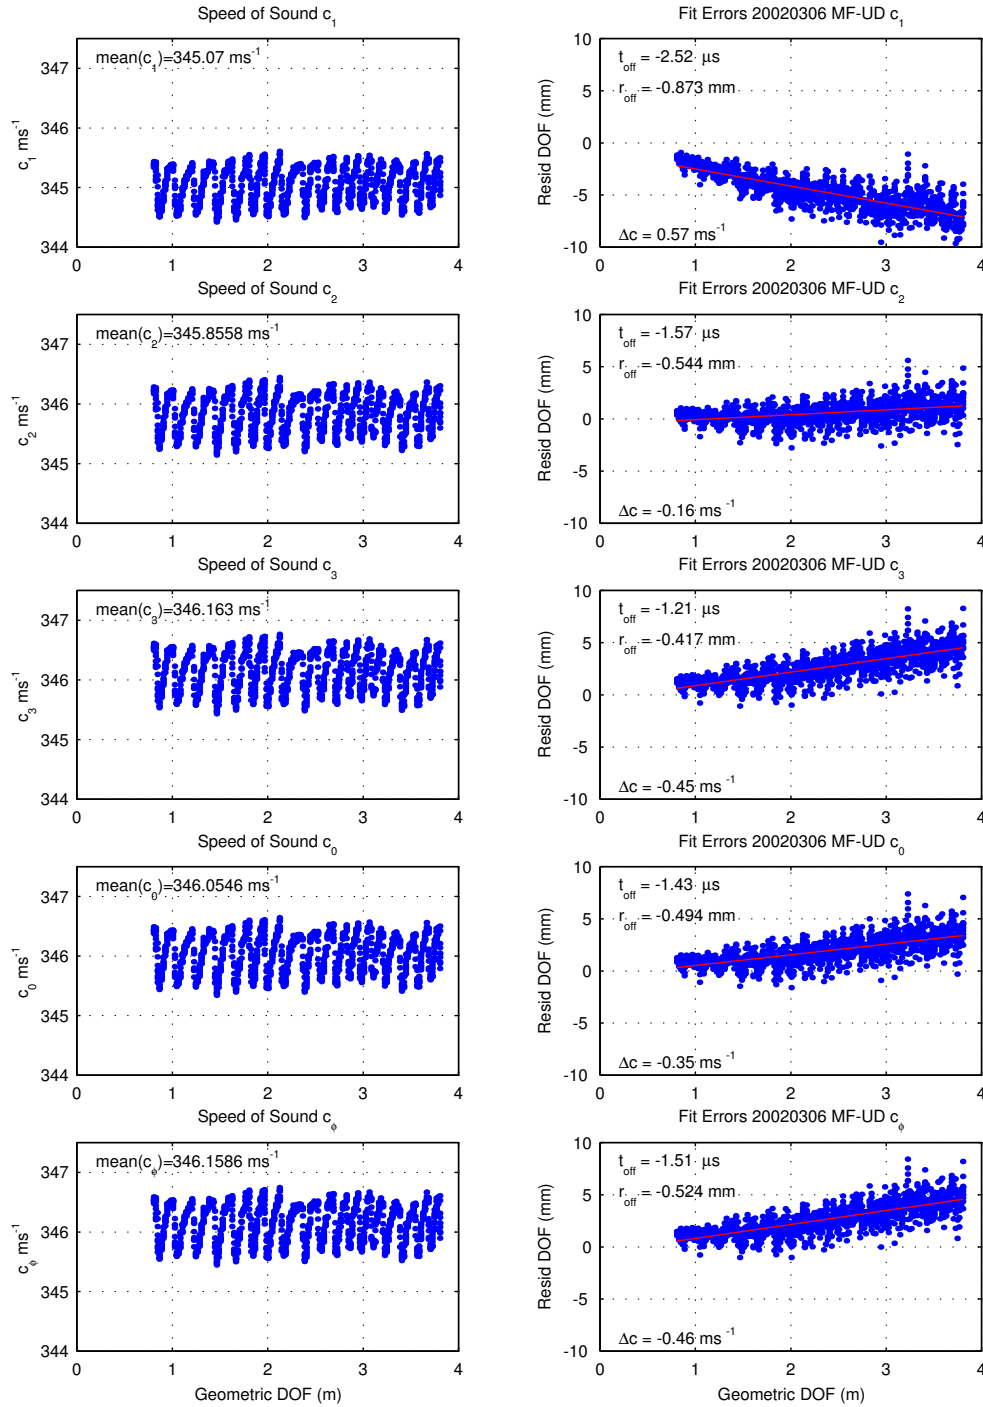


Figure 13.8: The effect of the choice of speed of sound calculation upon the measured DOF. The graph on the left is the speed of sound calculated by the formulae (12.6), (12.7), (12.8), (E.11) and (E.14) from the experimentally recorded temperature, humidity and pressure. The graph on the right is the difference between the measured DOF, calculated by multiplying the measured TOF by the speed of sound shown in the left column, and the geometric DOF.

Table 13.3: Measured differences from the speed of sound, c_ϕ . These values indicate that the speed of sound measured acoustically is approximately 0.15% slower than that computed from Cramer’s model.

Sonar System	Δc ($\text{m} \cdot \text{s}^{-1}$)	
	Left	Right
MF-QD	-0.47	-0.45
MF-UD	-0.46	-0.44
CTFM-DFT	-0.62	-0.61
CTFM-YW(128)	-0.52	-0.48
CTFM-LSMYW(128,128)	-0.58	-0.60

Speed of Sound Differences

The values obtained for Δc , a difference from the model value of the speed of sound, c , are tabulated in Table 13.3. These corrections imply that either the model value is over-estimated, or the acoustic DOF is larger than the geometric DOF. The values are statistically significant due to the large number of observations used in the regression analysis. However, replication is needed to be able to state with confidence that there is a class or sub-class effect.

In the MF class the errors in Δc dominate, because the offset errors are small, reaching values near 8mm with a DOF of 4m. In the CTFM class, the Δc errors grow in contribution as the DOF increases, reaching a 50% contribution level for a DOF of 4m. Figures 13.4 and 13.5 suggest that the residuals are also a function of DOF. Statistical tests, at the 95% confidence level, do not support this suggestion. However, the behaviour is predicted by the theory of Section 4.7.

While the offset time, t_{off} , can be ascribed, with confidence, to system hardware, it is much more difficult to assign possible causative effects to the need to adjust the model value for the speed of sound by the order of $-0.5 \text{ m} \cdot \text{s}^{-1}$. Some possible causative effects and their approximate contributions are listed below.

1. Sampling clock: The sampling frequency used by the DAC and ADC cards, described in Sections 5.1.1 and 5.1.2, is derived from the 10MHz NuBus clock signal. This clock signal is not used to record the time of day, and does not need to be precise for normal operation of the computer. An error of +50kHz or 0.5% on the clock signal would produce a $-0.5 \text{ m} \cdot \text{s}^{-1}$ error in the speed of sound.

The clock frequency has since been checked with a Hewlett Packard 5315A Universal Counter. The NuBus clock frequency was found to be 10000047Hz. This would, in turn, produce an error of $-0.0016 \text{ m} \cdot \text{s}^{-1}$ in the measured speed of sound. Thus this possible source of error is excluded from consideration.

2. Impulse response of air: Kleeman and Kuc [82] describe a method of calculating the impulse response of air as a function of the (frequency dependent) absorption. The absorption is a function of distance (see Appendix E.4). Since the absorption curve can be used to determine the impulse response, it provides a delay additional to the effect of the speed of sound. As the air impulse delay is a function of absorption which increases with distance,

the delay also increases with distance. Preliminary calculations indicate that the delay of an air impulse response, at $T_c = 23^\circ\text{C}$, $h_r = 60\%$ and $p_s = 101.55\text{kPa}$, is $6\mu\text{s}$ at 2m and $12\mu\text{s}$ at 4m. (These measurements only include the rise time of the impulse. The group delay may be longer.) This corresponds to an additional measured distance of 2mm at 2m DOF and 4mm at 4m DOF. This translates to a $-0.4\text{m}\cdot\text{s}^{-1}$ variation in the speed of sound. These rough calculations indicate that this effect may account for 80% of the discrepancy.

3. Model uncertainty: The model value must be corrected by about $-0.5\text{m}\cdot\text{s}^{-1}$. It was shown previously that Cramer's estimated precision [26] of $\pm 0.1\text{m}\cdot\text{s}^{-1}$ has been estimated by Wong [141] to be nearer $\pm 0.18\text{m}\cdot\text{s}^{-1}$. This is about a third of the required value. Thus while it may be an important component it is not capable of resolving the discrepancy.
4. Wind effect: The experiment was carried out underneath an air-conditioning duct which could not be disabled. This produced a downward air motion in the experimental area. The turbulent air current can change the acoustic path length that the sound must follow to reach the target, as the sound vibration is travelling relative to a moving medium (Osborne [106]). This would lengthen the acoustic DOF, leading in turn to a longer TOF and a lower apparent speed of sound. Modelling of such a turbulent air movement is difficult. However, the effect could be removed by conducting an experiment within a small anechoic chamber which excludes air movement. This has not been done.
5. Plane geometrical misalignment: It can be shown that a 1° error in the alignment of the normal of the glass pane causes a -0.5mm error at 1m range and -1mm error at 2m range on the left channel. On the right channel the errors are -0.06mm and -0.5mm . Rotating the plane in either direction causes the DOF to be shortened. Thus a 1° misalignment of the glass pane would cause an error in the measured speed of sound of $+0.16\text{m}\cdot\text{s}^{-1}$, and is positive for both directions of rotation. This error has the opposite sign of that required to solve the discrepancy.
6. Sonar head geometrical misalignment: A 1° error in the principal direction of the sonar head causes a $+0.23\text{mm}$ error in the left DOF and a -0.23mm error in the right DOF, independent of target range. The independence of range indicates that the speed of sound difference, Δc , due to misalignment of the sonar head is zero. Thus rotation of the sonar head will not contribute to the error budget.
7. Positioner calibration: The positioner is uncalibrated. The manufacturer's specifications for the stepper motors and the pitch of the lead screws were used to determine the number of steps per mm of travel. No feedback systems were used. Stepper motors sometimes skip steps when driven at their resonant frequency, but this was not observed. Comparisons of the positioner with a steel ruler indicated the positioning was accurate when the controlling software was written. The positioner would have to accrue a $+2\text{mm}$ error over its 1.5m length of travel during the experiment to provide a $-0.5\text{m}\cdot\text{s}^{-1}$ error in the speed of sound. Comparison of the positioner movement with a ruler indicates that the positioner is accurate to within $\pm 0.5\text{mm}$ along the 1m ruler length.

8. Unrepresentative environment parameters: The environmental parameters were measured by a transducer at a single point, while the sonar signal propagates through a volume of air. If the air is not mixed properly, then the environment sensor measurements may be unrepresentative of the parameters affecting the sonar signal. This effect is considered unlikely as the experiment was situated underneath an air conditioning duct which stirred the air continuously.
9. Change in environmental calibration: It is possible that the electrical components in the environment sensor are themselves sensitive to changes in the environmental parameters, leading to a change in the calibration of the environment sensor when the experiments were performed. The experiment described in Appendix F was not performed on either of the days that the range precision experiment was carried out.
10. Error in CO₂ mole fraction: The speed of sound observed by the sonar system would require a CO₂ mole fraction of approximately 5000ppm. While the CO₂ mole fraction has increased by 20% over the last 30 years to 370ppm in 1999 (Kimball [77]), a value of 5000ppm is inconceivable.

Of the ten possible causes enumerated here, only items 2 and 3 are capable of explaining a significant proportion of the observed discrepancy. The contribution of the fourth item remains unquantified at this point.

Variation of Δc between Sonar Classes

The variation in the observed Δc is smaller within each sonar class. The variation in Δc between sonar classes may be attributed to either differences in the environmental conditions on the different days of the MF and CTFM experiments, or differences between the MF and CTFM signals themselves. A change in the environmental conditions would also lead to a change in the air impulse response due to absorption, which has not been compensated. The variation in Δc may be attributed to either the signal type or to a different air impulse response delay.

The variation between the left and right channels, typically $0.02 \text{ m} \cdot \text{s}^{-1}$, is an indication of the experimental error, common to both channels, in the determination of the speed of sound.

DOF Precision Results

The precision of the DOF determination is summarised in Table 13.4. The histograms in the right column of Figures 13.4 and 13.5 demonstrate that the residuals are normally distributed with a mean of zero. There is no statistical difference between the left and right channels. However, there is a statistical difference between the matched filter, MF, values and the CTFM values, indicating greater variation in the CTFM system.

The worst system with respect to precision is the CTFM-DFT sonar system, which has lower statistical precision than either of the CTFM-AR systems. This can be attributed to the large quantum between spectral lines used in the DFT spectral estimate (Section 10.2.3). The sinc interpolation, discussed in Section 10.3.4, is only partially effective at interpolating the large quanta. The

Table 13.4: Results for the **DOF** precision of the different sonar systems, quoted to one more place than is significant for computational convenience.

Sonar System	σ_d (mm)		
	Left	Right	Average
MF-QD	0.88	0.96	0.92
MF-UD	0.89	0.90	0.90
CTFM-DFT	1.21	1.16	1.19
CTFM-YW(128)	1.01	0.98	1.00
CTFM-LSMYW(128)	1.02	0.98	1.00

MF sonar systems have a **DOF** quantum of 1 mm and 0.26 mm before interpolation (Sections 8.6.2 and 8.6.3), while the two **YW** based systems have no quantisation (Section 11.4.4).

The regular oscillations which are visible in the residuals are related to the variations in the environmental parameters. This may be verified by comparing the residuals to the speed of sound graphed against time in Figure 13.9. The data shown in the figure demonstrate that correcting for the variations in the speed of sound significantly reduces the standard deviation of the residuals. The small amount of cross-correlation that remains is considered statistically insignificant. While a large amount of the remaining variation is random noise, some of it is correlated to the variation in the speed of sound, and hence to the environmental parameters. We cannot reject the hypothesis that there is a second order effect connecting the environmental parameters with the **TOF**. A hypothesised connection is through the impulse response of the air, which is sensitive to changes in temperature and humidity, described in Item 2 on page 251.

13.2.2 Power Precision Analysis

Power Model

The sonar signal suffers spherical spreading loss, as described in Chapter 4, and absorption due to interaction with the air through which it propagates, as summarised in Appendix E.4. Thus the signal amplitude may be modelled approximately by

$$a = \frac{a_1}{d} e^{-\alpha d}, \quad (13.6)$$

where

a_1 is the measured source level of the signal when it has travelled a distance, d , of 1 m, ignoring the effect of absorption.

α is a mean air absorption, in Nepers per metre, over the frequency band of the signal.

Note that the data has been normalised to the noise level, so the measured source level is not absolute. The procedure is described in Section 13.2.4.

The virtual source concept (see Figure 4.5) has been used to apply (13.6) to this experiment. Thus we reflect the transmitter in the plane to form a virtual source and consider a spherical

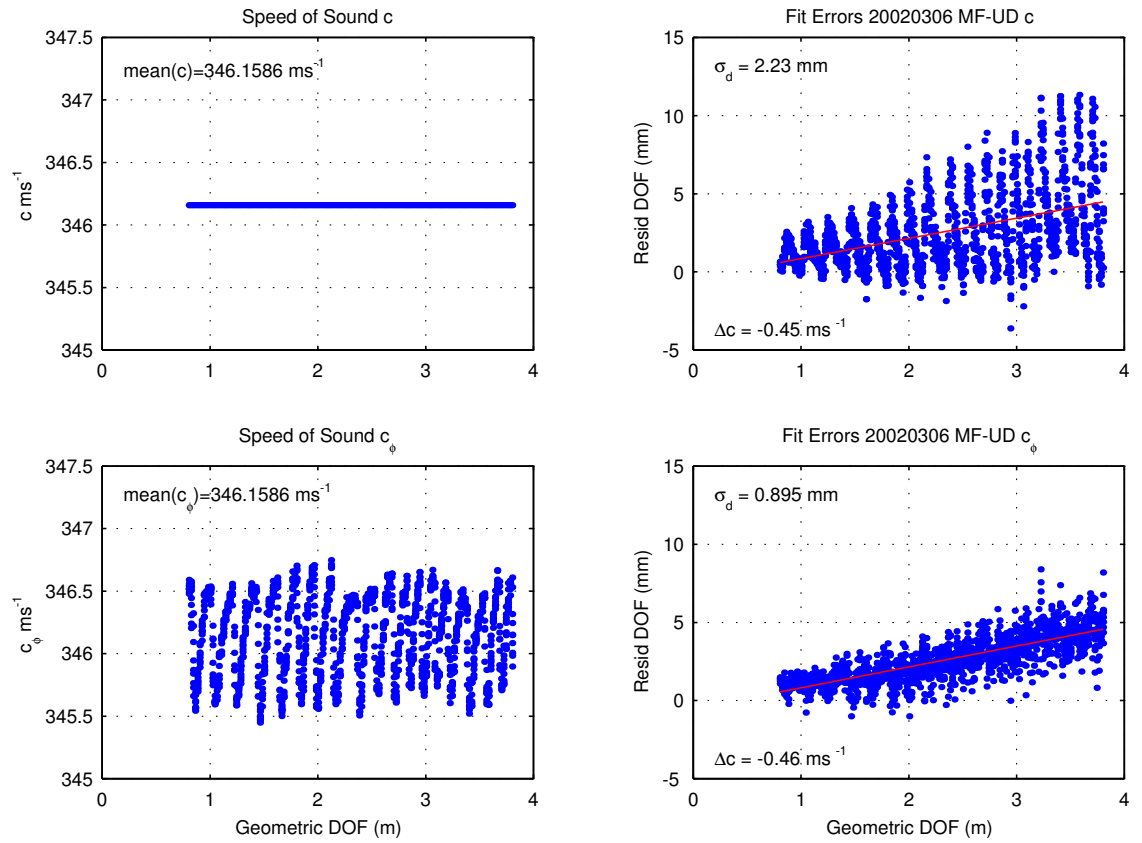


Figure 13.9: Analysis of speed of sound compensation. The left column shows the values used for the speed of sound. The right column shows the residual DOF , $c_i \tau_i - d_g(x_i)$, as calculated for the main body of results. The straight line fit is also shown as a red line. The first row shows the effect of using a constant speed of sound, or equivalently, not compensating for variation in the speed of sound. The second row uses Cramer's model for the speed of sound, c_ϕ . The standard deviation of the residuals, σ_d , measured after subtracting out the fitted line, is reduced by 40% by compensating for the speed of sound. The cross-correlation between the modelled speed of sound, c_ϕ , shown bottom left, and the residuals for a constant speed of sound, shown upper right, is -0.83 . The correlation coefficient between c_ϕ and the residuals for the compensated speed of sound, shown lower right, is 0.23 . The residuals in the top right panel are inverted relative to the modelled value of the speed of sound, c_ϕ , as indicated by the negative cross-correlation. The residuals in the lower right panel have the same phase as the modelled value c_ϕ , as indicated by the positive correlation coefficient.

wavefront travelling directly from the virtual source to the receiver. The received power, P_r , in decibels is then

$$P_r = 20\log_{10}(a_1) - 20\log_{10}(d) - \frac{20}{\ln 10}\alpha d = A_1 - 20\log_{10}(d) - \alpha' d, \quad (13.7)$$

where

α' is the absorption in decibels per metre and is related to α by (E.10); and

A_1 is the (normalised) measured source level in decibels at $d = 1$ m if the absorption, α' , were 0.

For a target at 1 m DOF, equation (13.7) for received power may then be simplified to $A_1 - \alpha'$.

The model (13.7) was fitted to the data using a least squares technique, as described in Appendix I.

Power Precision Experimental Results

The experimental data set from Section 13.2.1 was re-used to provide amplitude measurements. The detected echo amplitude was converted to decibels of power and plotted in the left column of panels in Figures 13.10 and 13.11 for the left and right channels.

The power was normalised relative to the 3σ level of the normally distributed background noise. The normalisation was carried out using per-channel and per-system noise measurements. The noise level was estimated using a single representative signal, selected to be observation record 101 where the target is at a range of 500 mm. The method of measuring the noise level will be described in Section 13.2.4. With this in mind, the power plotted in the left column of Figures 13.10 and 13.11 is seen to be the signal to noise ratio of the signal from a target at the given DOF.

The modelled signal power, P_r from (13.7), was fitted to the data using the techniques of Appendix I and is drawn in the left column of panels as a red line. The values used for the fitted parameters, A_1 and α' , are also drawn on the plots. The parameter A_1 is the normalised measured source level at $d = 1$ m, while α' is the absorption in decibels per metre. The power of the received echo is expected to be a curve, due to the logarithmic function of DOF.

The second column of panels shows the residuals, which have been generated by subtracting the modelled data, P_r , from the experimental echo power.

The third column of panels shows a histogram of the residuals shown in the second column. The mean μ and standard deviation σ are indicated on the panels. This plot may be used to determine whether the residuals are normally distributed. Non-normally distributed residuals indicate a model error or other perturbation.

As this analysis uses amplitude data from the range precision experiment, the measured environmental parameters from that experiment also apply, and may be found in Figures 13.6 and 13.7.

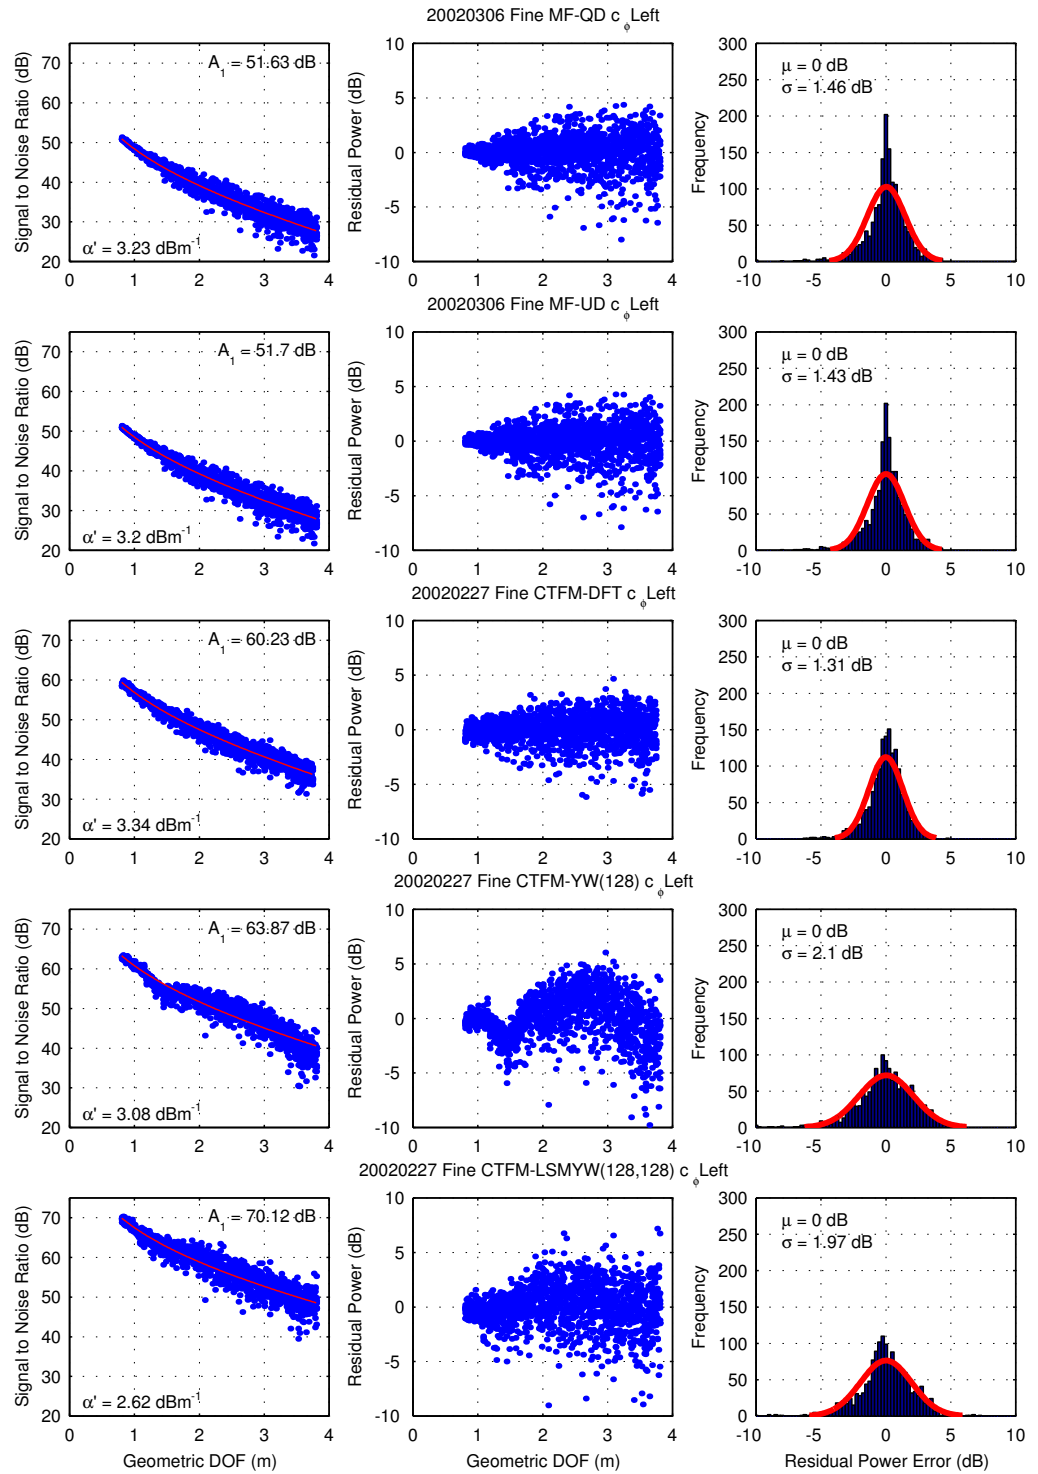


Figure 13.10: Results of the echo power precision measurements for the left channel. The left column of plots is the received **signal to noise ratio**, plotted as a function of **DOF**. The red line is the fitted model P_r . The power levels are normalised to the 3σ or 99.7% confidence level of the normally distributed system noise. The middle column of plots shows the residuals after subtracting the model from the measured data. The third column of plots shows a histogram of the residuals.

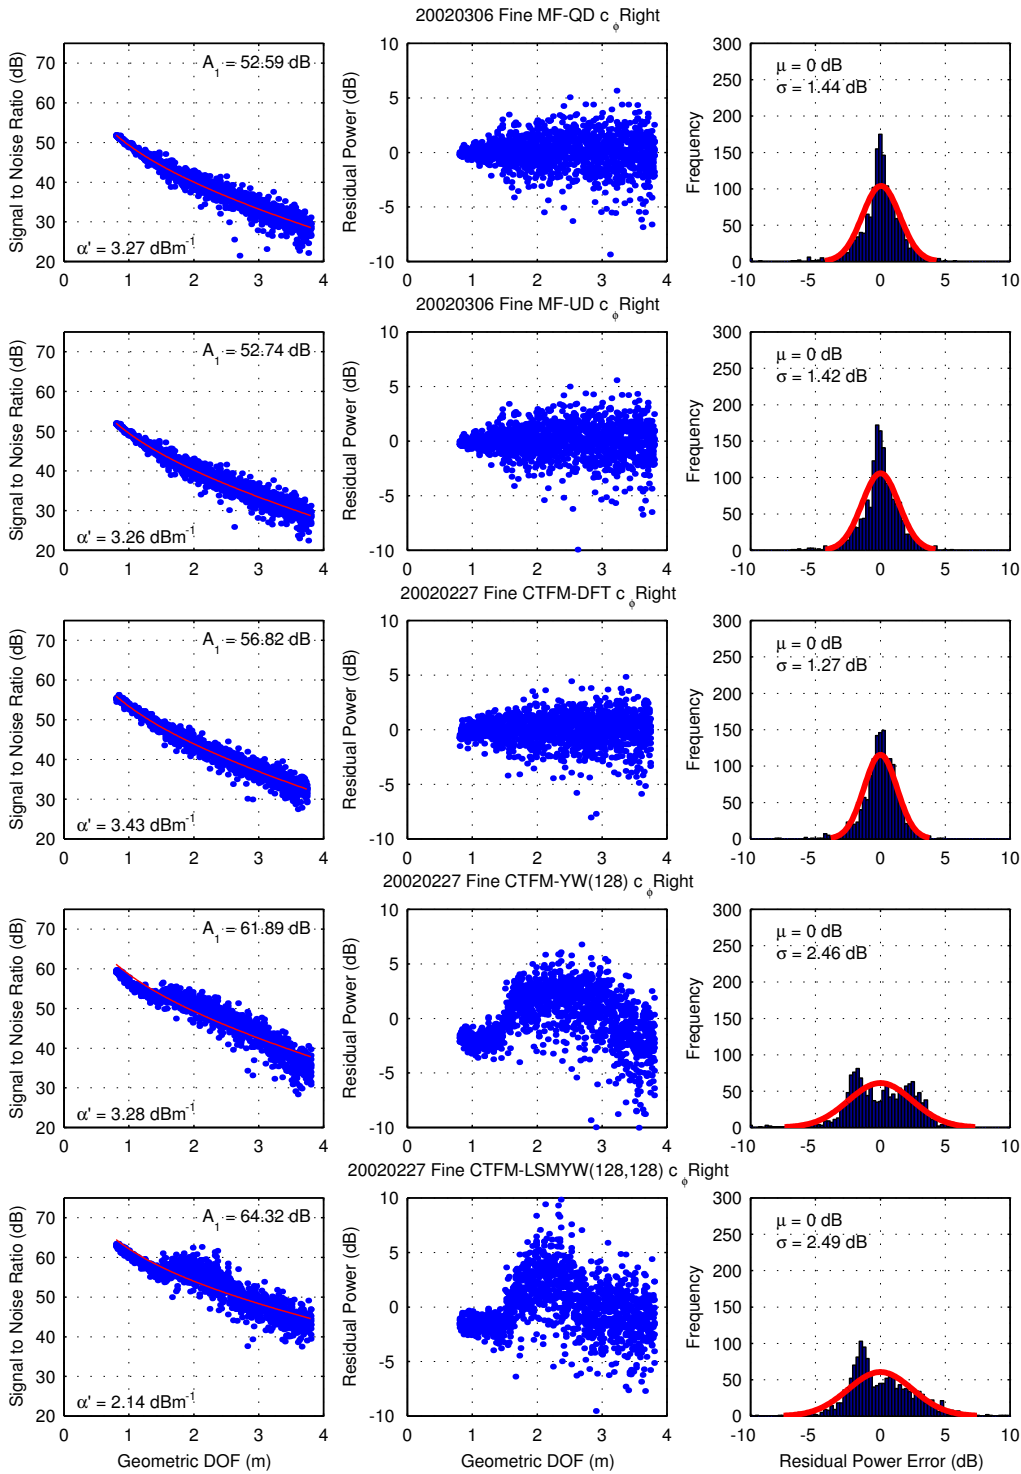


Figure 13.11: Results of the echo power precision measurements for the right channel. The left column of plots is the received **signal to noise ratio**, plotted as a function of **DOF**. The red line is the fitted model P_r . The power levels are normalised to the 3σ or 99.7% confidence level of the normally distributed system noise. The middle column of plots shows the residuals after subtracting the model from the measured data. The third column of plots shows a histogram of the residuals.

Table 13.5: Received echo power precision, taken from Figures 13.10 and 13.11.

Sonar System	σ_p (dB)		
	Left	Right	Average
MF-QD	1.46	1.44	1.45
MF-UD	1.43	1.42	1.43
CTFM-DFT	1.31	1.27	1.29
CTFM-YW(128)	2.10	2.46	2.28
CTFM-LSMYW(128,128)	1.97	2.49	2.23

Discussion of Precision of Received Power

The precision of the received echo power estimation of each of the five sonar systems is presented in Table 13.5. In each of these systems the noise and other unmodelled effects such as the impulse response of the air increase as a function of DOF.

The MF-QD and the MF-UD systems are statistically identical for measuring echo power. Both systems utilise the same transmit signal and matched filter system, and operate at the same sampling rate up to the envelope extractor. The only difference between the two systems is in the envelope extraction and peak finding stages, which utilise [quadrature demodulation](#), Section 8.6.2, and [unsynchronised demodulation](#), Section 8.6.3. While these two systems have different output sampling rates and hence different DOF quantisation, they both utilise polynomial interpolation to find the peaks. As the demodulation technique is the only difference between the two systems, and the two systems provide equivalent results, it is concluded that the MF-QD and MF-UD envelope extractors perform equivalently with respect to echo power estimation.

The CTFM-DFT system has the best precision for echo power. This suggests that the position of the interpolated peak is repeatable in amplitude but not, according to the DOF results, in frequency. This performance dichotomy is attributed to the shape of the peak; see Figure 10.4 for an example in which the width of the peak makes the position of the peak maximum ill-defined, but allows the amplitude to be reliably determined.

The [auto-regressive CTFM](#) sub-systems, CTFM-YW and CTFM-LSMYW, have the largest uncertainties. A significant reason for this increase is the poor modelling of the observed echo power, which is shown by the residual structure in these panels in the middle column of Figures 13.10 and 13.11, and by the multi-modally distributed residuals which may be seen in the third column of Figure 13.11. An analysis of the spectra (see Figure 13.12) indicates that as the power of adjacent peaks increases, the level of the minimum between these peaks also increases, thereby affecting the modelling and accuracy of the received power level while leaving the DOF determination unaffected. It is seen from Figures 13.10 and 13.11 that a major change in characteristics takes place near a DOF of 1.4m. A detailed investigation of spectra shows that this is where a secondary echo from an off-axis corner begins to have significant power (see also Figure 13.12).

An additional but secondary issue is the effect of noise upon [auto-regressive](#) systems. Kay and Marple [75] indicate that AR systems are unable to function satisfactorily in a noisy environment, and that [Auto-Regressive Moving Average \(ARMA\)](#) systems are to be preferred. They have not been applied in this work due to their added complexity.

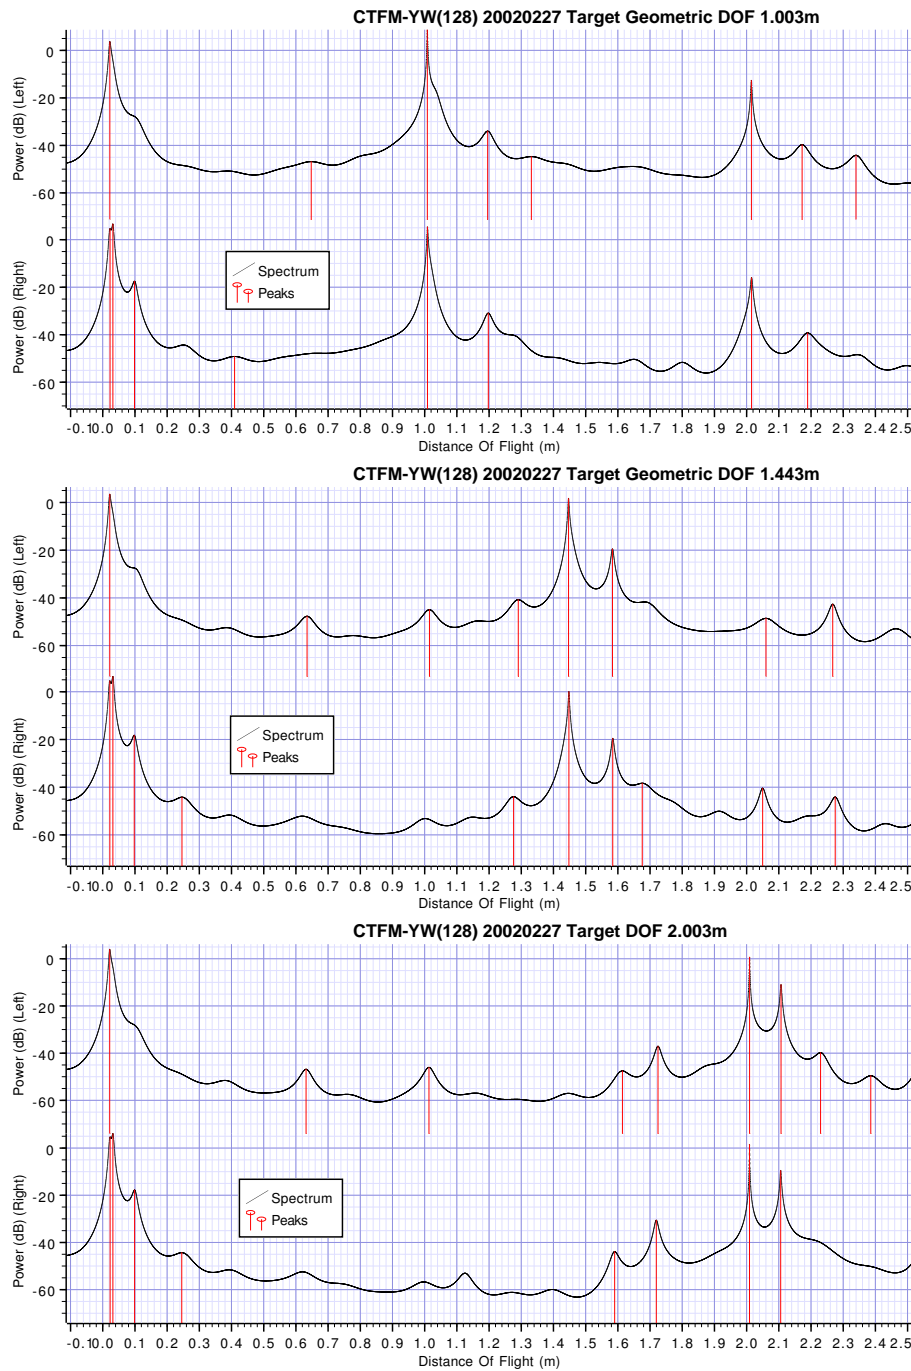


Figure 13.12: Returned echo power estimated by the CTFM-YW(128) sonar system for various target distances. The target is at 1 m **DOF** in the first plot, 1.45 m in the second, and 2 m in the third. The echo visible at 2 m in the first plot is an echo which has travelled between the sonar head and the wall 4 times. The secondary echo is visible at 1.2 m, 1.58 m and 2.1 m in the sequence of three plots. The secondary peak gets closer to the main peak as the target **DOF** increases. The secondary peak power also increases, as the off-axis corner reflector enters the beam. The additional power in the secondary peak, in close proximity to the primary peak, causes the estimated power in the primary peak to be perturbed.

Table 13.6: Normalised measured source level at 1 m, A_1 . The values shown are measured in decibels relative to the 3σ level of the normally distributed system noise.

Sonar System	A_1 (dB)	
	Left	Right
MF-QD	51.6	52.6
MF-UD	51.7	52.7
CTFM-DFT	60.2	56.8
CTFM-YW(128)	63.9	61.9
CTFM-LSMYW(128,128)	70.1	64.3

Discussion of Normalised Measured Source Level

The normalised measured source level, A_1 , incorporates several factors:

1. The elements of the signal path outlined in Figure 4.2,
2. The sonar processing,
3. The energy in the transmit signal, and
4. The normalisation due to the system noise level.

The contribution of each of these items to the measured normalised source level is only partially separable. Item 1 is controlled by utilising the same data set for the analysis of each sonar class. Some variation in the signal path may occur between the experiments for each sonar classes as the two experiments were conducted on different days. Items 2 and 3 currently remain uncalibrated. Item 4 is tabulated in Table 13.8. The main reason for the normalisation is to provide the basis for a comparison between the five sonar systems, based upon [signal to noise ratio](#).

The values obtained for the measured normalised source level, A_1 , are listed in Table 13.6. These values indicate a clear superiority of about 8dB of the [CTFM](#) class systems over the [MF](#) class systems. The left channel of the CTFM-LSMYW system is specifically noted as an outlier, due to the 6dB difference between the two channels. This is primarily due to the low value measured for the noise floor, which is discussed in Section 13.2.4. The perturbation of the [LSMYW](#) measurements by the secondary echo, discussed previously, is a secondary effect. With the noted exception of the CTFM-LSMYW results, the data in Table 13.6 show that the left and right channels are quite similar.

Discussion of Absorption

The fitted values of the absorption, α' , in decibels per metre are summarised in Table 13.7. The absorption is determined by the air properties and by the bandwidth of the signal. The air properties (see Figures 13.6 and 13.7) changed between the experiments for [MF](#) and [CTFM](#) sonar classes with a 1°C change in temperature, and a 13% change in relative humidity. The absorption (see Figure E.1) can be expected to remain constant over this particular region of humidity parameter variation. Furthermore, the bandwidth of the two classes of sonar system, shown in Tables 8.1

Table 13.7: Air absorption α' in $\text{dB} \cdot \text{m}^{-1}$. The values shown here represent the mean absorption over the bandwidth of the signal.

Sonar System	α' ($\text{dB} \cdot \text{m}^{-1}$)	
	Left	Right
MF-QD	3.2	3.3
MF-UD	3.2	3.3
CTFM-DFT	3.3	3.4
CTFM-YW(128)	3.1	3.3
CTFM-LSMYW(128,128)	2.6	2.1

and 9.4, is seen to be 19.9kHz for the MF class, and 21.6kHz for the CTFM class. The mean frequencies are 80kHz for the MF class and 79kHz for the CTFM class. Thus the frequency band occupied by the two signals is nearly identical. The similar air properties and almost identical bandwidth explain the experimental similarity between the absorption in the two sonar classes, with the exception of the CTFM-LSMYW system.

The absorption values obtained by the CTFM-YW(128) system are close to those obtained by the MF and CTFM-DFT systems, despite the power disturbance by the close echo. The CTFM-LSMYW(128,128) result has been badly affected by the close echo problem described previously, causing a clear difference in the estimated air absorption. This leads to the conclusion that the absorption estimates are invalid in the case of the CTFM-LSMYW system.

Discussion of Power Residuals

The standard deviation of the residuals statistically increases over the 3m DOF variation presented in Figures 13.10 and 13.11. This is also reflected in the histograms, which are sharper in the middle than their fitted normal distributions. This is attributed to the sonar signal accumulating airflow fluctuations in proportion to the path length through the propagating medium.

The histograms shown in the right columns of Figures 13.10 and 13.11 also show that the residuals are not normally distributed, which is attributed to the change in the distribution of the residuals with range, and to the unaccounted range dependent structure.

The variation of the received echo power is quite large when compared with the absorption and spreading loss, which makes echo power unsuitable for determining the DOF.

13.2.3 Cross-correlation Analysis

The cross-correlation of DOF and echo power data output by the detector from the left and right channels or transducers is important for the determination of bearing and bearing precision. This is shown in Sections 14.1.3 and 14.2.4. However, since the principal data products of a receiver are DOF and echo power, which are the topic of this chapter, it is appropriate to present these correlations here.

The correlations between the left and right channel detector residual data are shown in Figure 13.13. These plots are generated by taking the residuals corresponding to a single measurement and drawing a point corresponding to that measurement. In the first column, the horizontal axis




Figure 13.13: Correlation between the left and right channels for DOF and power residual measurements output by the detector. The correlation coefficient is marked on each plot. See Section 13.2.3 for full details.

represents the residual of the left channel DOF measurement, while the vertical axis represents the residual of the right channel DOF measurement, in mm. In the second column, the axes represent the echo power residuals, in decibels. The correlation coefficients are calculated by

$$c_{12} = \frac{\sigma_{12}}{\sqrt{\sigma_{11}\sigma_{22}}},$$

where

σ_{12} denotes the covariance between two different quantities, and

σ_{11} and σ_{22} represent the variance of their respective quantities.

The variances and covariances appear in the Variance CoVariance (VCV) matrix Σ , which is described in Appendix I.4.

Discussion of DOF Results

The DOF residuals show moderate levels of correlation, with the CTFM-LSMYW sonar having the highest correlation coefficient. While the two matched filter systems do not have statistically different DOF standard deviations (see Table 13.4), the MF-UD system does have a noticeably higher correlation between the left and right DOF measurements.

The CTFM-DFT system displays tram-tracks in its DOF residual results. An obvious cause would be the underlying discrete nature of the DFT spectral estimation process. However, the spacing between the tram-tracks is $\Delta d = 2.5$ mm, which is quite different from the DOF quantisation of $\Delta d = 7.42$ mm found in Section 10.2.4. There is no tram-track at this distance. Indeed, the tram track spacing of 2.5 mm is about one third of the required DOF. Therefore the tram-tracks cannot be explained by the discrete nature of the DFT.

Preliminary experiments utilising a zero padded DFT in place of Davies' [28] complex interpolation technique (see Section 10.3.4) indicate that the tram track phenomenon is completely removed. This demonstrates that the tram track fault lies completely within Davies' complex interpolation technique. The interpolation technique was selected due to its low computational complexity. The zero padded technique appears to remove the problem, albeit with significantly higher computational complexity.

A possible explanation of the tram-track phenomenon is as follows. The interpolation procedure outlined in Section 10.3.4 assumes that the power spectrum is composed of a sinc function with a phase component (see equation (10.18)). This implies that the two samples on either side of the central peak are very small when the peak lines up exactly with a spectral line. This can be seen in the red curves in the top panel of Figure 10.4. However the real sample points will always lie upon lines exemplified by the green curve in Figure 10.4 and therefore the two side points will diverge greatly from the required zero model at these points. The divergence is partly due to additive noise and partly due to the extended low-level sinc sidelobes of other targets. Thus when the central peak is aligned with a spectral sample, the effect of these components will be to shift the estimated peak frequency slightly away from the central peak frequency. In relation to the tram-track phenomenon, it is postulated that the narrowness of the zeroes of the sinc function

(see Figure 10.4) is what causes the localisation of the tram-track, rather than a continuously distributed error. The 2.5 mm width between the clusters is postulated to be linked to the background level (of the green curve), and the width at which the fitted curve (red) achieves this level. This is supported by Figure 10.4.

The tram-track phenomenon is expected to cause problems with the estimation of bearing using CTFM-DFT data and the IDD technique, as the bearing calculation, performed in Section 14.1, relies upon the difference between the two DOF measurements. The correlation between the left and right measurements will cause bearing offsets.

It may be possible to remedy the tram-track problem by adding additional code to determine when the two side peaks have a low energy, and to use a different technique to perform the interpolation. The primary strength of the interpolation technique is its ability to accurately interpolate the position of the peak when it falls near the mid-point between two spectral lines.

The DOF measurements are used to compute both range and bearing. Section 14.1.3 shows that positive correlations near 1 in the DOF results lead to highly precise bearing estimates, but less precise range estimates. On the other hand, correlations near -1 lead to precise range estimates but imprecise bearing estimates. In this instance the type of correlation is determined by the properties of the air column, which are nearly identical between the left and right channels. This leads to high positive correlations, which in turn lead to precise bearing estimates.

Discussion of Power Results

Cross-correlation for power in the left and right channels is variable, but the level is from values that are usually not statistically significant, $|c_{12}| < 0.5$, to levels that are moderately significant, $|c_{12}| < 0.8$. As the bearing calculated from IPD utilises power difference (see Section 14.2.3), correlations near 1 will lead to more precise bearing estimates by IPD. Since range is not computed from echo power, there is no trade-off as there was with the DOF correlations.

Cross Correlation of DOF and Power

The cross correlation between the DOF and echo power measurements (not plotted) has a maximum value of 0.47 over all sonar systems. This indicates that the DOF and power estimates are essentially uncorrelated. This result will be used in Chapter 14 to show that the two methods of calculating bearing, the IDD and IPD methods, will provide independent results.

13.2.4 Noise Level Analysis

Each of the five different sonar types produces an output signal with a different level. For the purposes of comparison, it is therefore necessary to normalise the levels by some common reference. Two reference points are possible:

1. The level of an echo from a target at a standard distance under standard conditions.
2. The level of the background noise, or some offset from that noise level.

The second option was chosen, as it allows ready evaluation of the signal to noise ratio of the echo signals. This in turn affects both the precision of the detection results and the maximum range of the sonar system. The second option also removes the restrictions of standard conditions, which are satisfied by using one experimental data set for each sonar class, but are not satisfied between different sonar classes and data sets.

The background noise level was calculated using a sample echo recording from each sonar type, extracted from the experimental results from the range experiment configuration described in Section 13.1. The echo was from a plane target placed at 500mm from the sonar head. For each sonar type, data was obtained in decibels of power as a function of range, and a quiet area containing no echoes was selected. The RMS noise level, V_{RMS} , was then computed using

$$V_{RMS} = \sqrt{\frac{1}{n} \sum_{i=1}^n 10^{(\frac{x_i}{10})}}, \quad (13.8)$$

where the data points, x_i , are decibels of power. Equation (13.8) must be applied in the domain in which the echoes can be distinguished from the noise. The formulation is valid whether the data represents time power samples or frequency power samples, due to Parseval's theorem (Bracewell [14]). If the noise is assumed to be Gaussian, then this level may be described as a single standard deviation, and will encompass 68.3% of the corresponding amplitude sample points.

Detection and the Probability of False Alarm

The signal is examined by the detector which decides whether or not a contact is present, and then determines its location (Van Trees [132, Chapter 2]). There are four possible outcomes of the detection process, ignoring the determination of location. These are:

1. There is actually a target, and it is correctly detected.
2. There is actually a target, but it is not detected.
3. There is actually no target, but one is detected. This is a false alarm.
4. There is actually no target, and no target is detected.

These four configurations represent conditional probability statements. Thus the likelihood of statements 1 and 2 add to a probability of 1, and the likelihood of statements 3 and 4 also add to a probability of 1.

The probability of each of the statements listed above may be related, under ideal conditions, to the level of the detection threshold with regard to the system noise. It is standard engineering practice to select the **Probability of False Alarm (PFA)** and then compute the required detection threshold, using the system noise level as a parameter. A small PFA implies a detection threshold which is set far above the noise floor, which will, in turn, prevent the detection of weak targets. The selection of the detection threshold is a design issue, requiring the designer to select between detection of weak reflectors and a low probability of false alarm.

In this work, the detection threshold was selected empirically to obtain a small but acceptable number of false alarms. The **probability of false alarm** corresponding to each detection threshold level was then calculated after the noise floor had been measured.

The noise levels determined by the method described above are summarised in Table 13.8. The results of each system will now be considered.

Matched Filter with Quadrature Demodulation Noise Level Results

The noise level of the MF-QD sonar system is shown in Figure 13.14. Equation (13.8) was computed over the region indicated in grey. This data set includes a quiet area near the end of the recording where there were no echoes present. The RMS noise level is indicated, as well as the percentage of the normally distributed noise that it contains. The two and three σ noise levels are also included, along with the percentages of the noise distribution that they contain.

The peaks detected by the peak finding algorithm (described in Section 8.7.1) are shown in red, while the sidelobe exclusion regions (described in Section 8.7.2) are indicated by the blue lines on either side of the detected peaks.

The detection threshold labelled in Figure 13.14 was selected empirically as described above. The noise level differs by less than 2 dB between the two channels. This difference in noise levels leads to a difference in the PFA, as listed in Table 13.8. The large PFA difference is attributed to the sensitivity of the cumulative probability distribution. However, the levels achieved should provide for reliable operation with relatively few false alarms. The SNR measured for the two channels is in good agreement, with only 0.3 dB between the two channels.

Matched Filter with Unsynchronised Demodulation Noise Level Results

The noise level of the MF-UD sonar system is shown in Figure 13.15. The features on this plot are the same as for the MF-QD, and the noise level, the detection threshold and the PFA are not significantly different from those of the MF-QD system. This is to be expected due to the similarity of these two designs (see Chapter 8).

CTFM-DFT Noise Level Results

The noise level of the CTFM-DFT sonar system is shown in Figure 13.16. This sonar system does not display a constant noise floor as the MF systems do. This is due to the sidelobes of the sinc function which is introduced for every target (see equation (10.17)).

The detection threshold labelled in Figure 13.16 was selected empirically as described above. The very high setting of the detection threshold for this system, at 45.6σ and 33.4σ for the left and right channels respectively, implies a very low PFA of 10^{-453} and 3×10^{-245} for the left and right channels under ideal circumstances. Consequently, the maximum operating range of the system is reduced while the detection threshold remains at this setting. An example of non-ideal circumstances where the PFA is lowered is in the side-lobes of the main peaks, where it can be seen that many smaller peaks reside. The application of a shaped sidelobe removal system, similar to that implemented for the MF system in Section 8.7.2, would remove the need for such a high detection threshold. This proposal remains as future work.

The noise level results for the left and right channels differ by 2.7 dB, while the signal to noise ratio differs by 4 dB. The sign of the difference is consistent with the other systems in the CTFM class shown in Table 13.8, but is opposite to the sign of the difference in the two systems in the

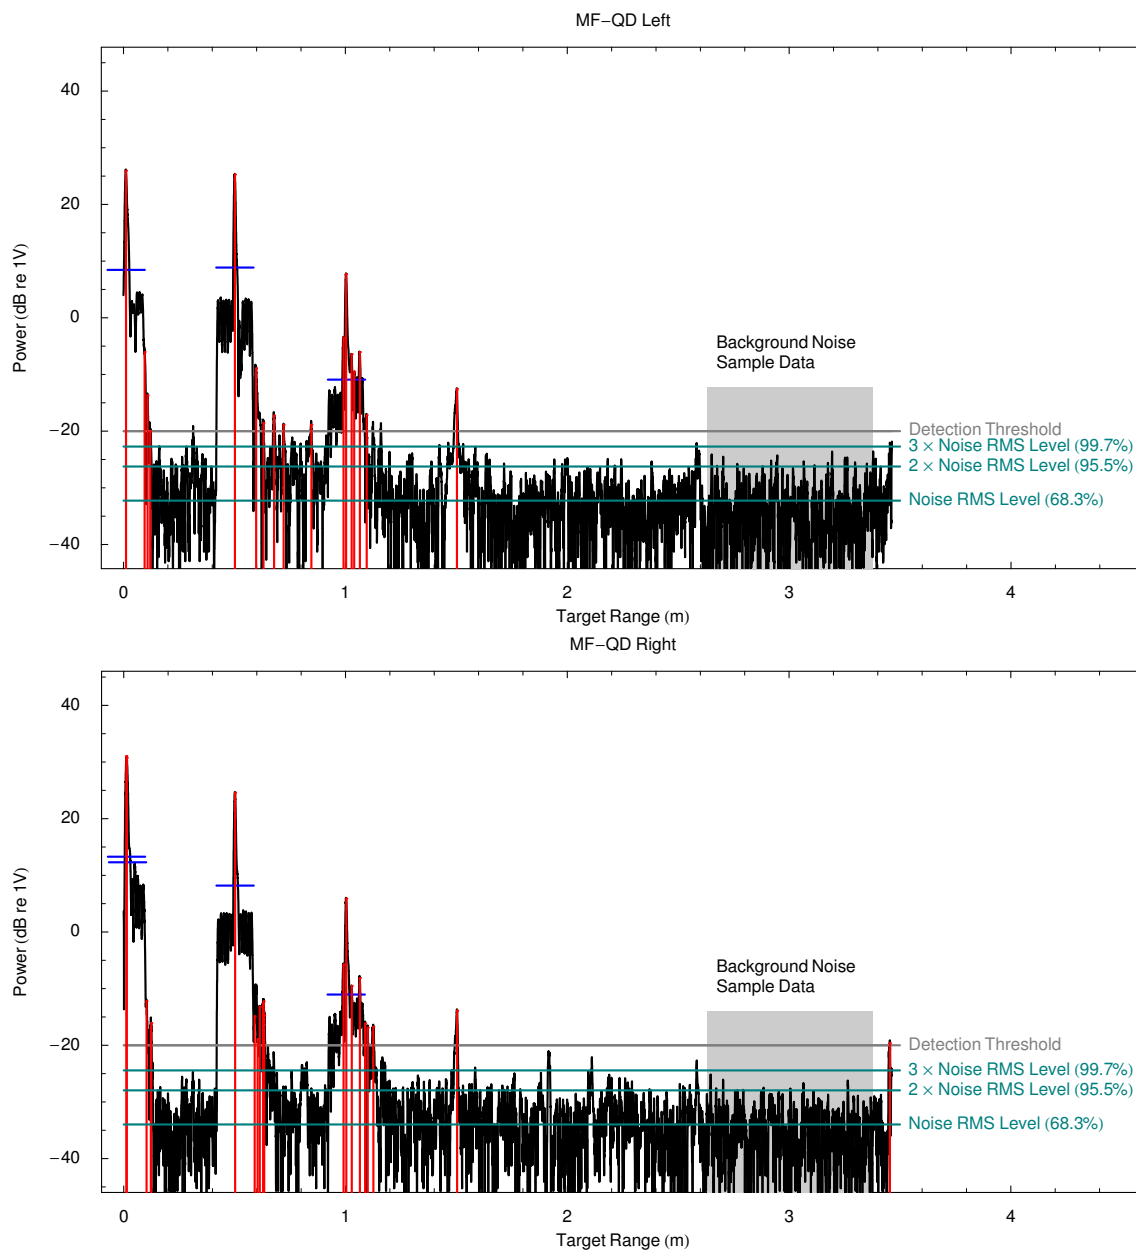


Figure 13.14: Determination of the noise level for the MF-QD sonar. See p.267 for discussion.

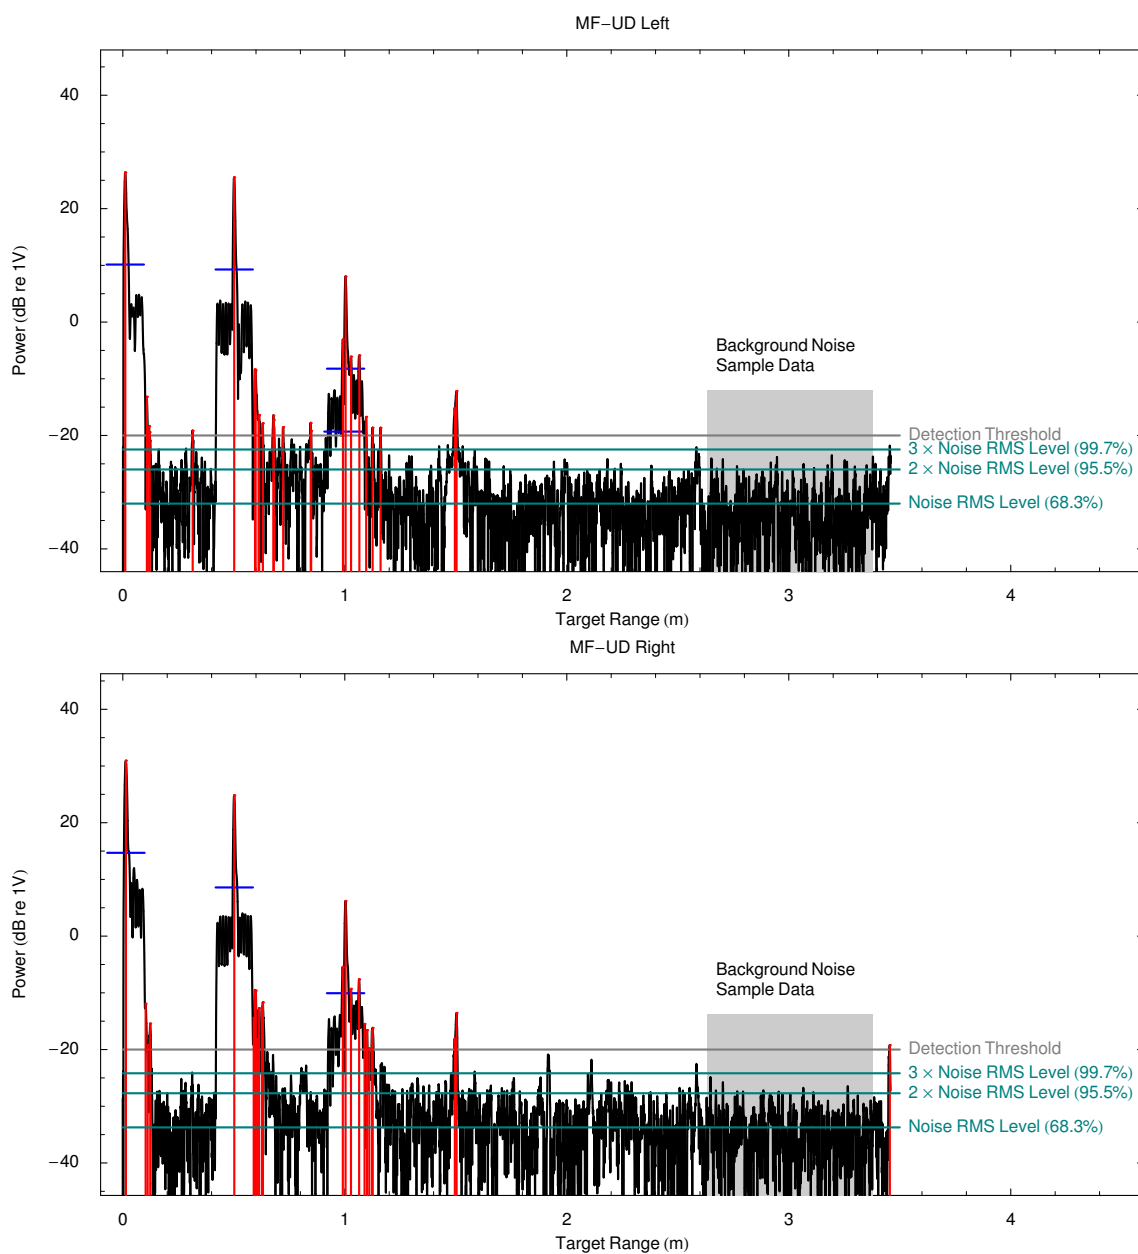


Figure 13.15: Determination of the noise level for the MF-UD sonar. See [p.267](#) for discussion.

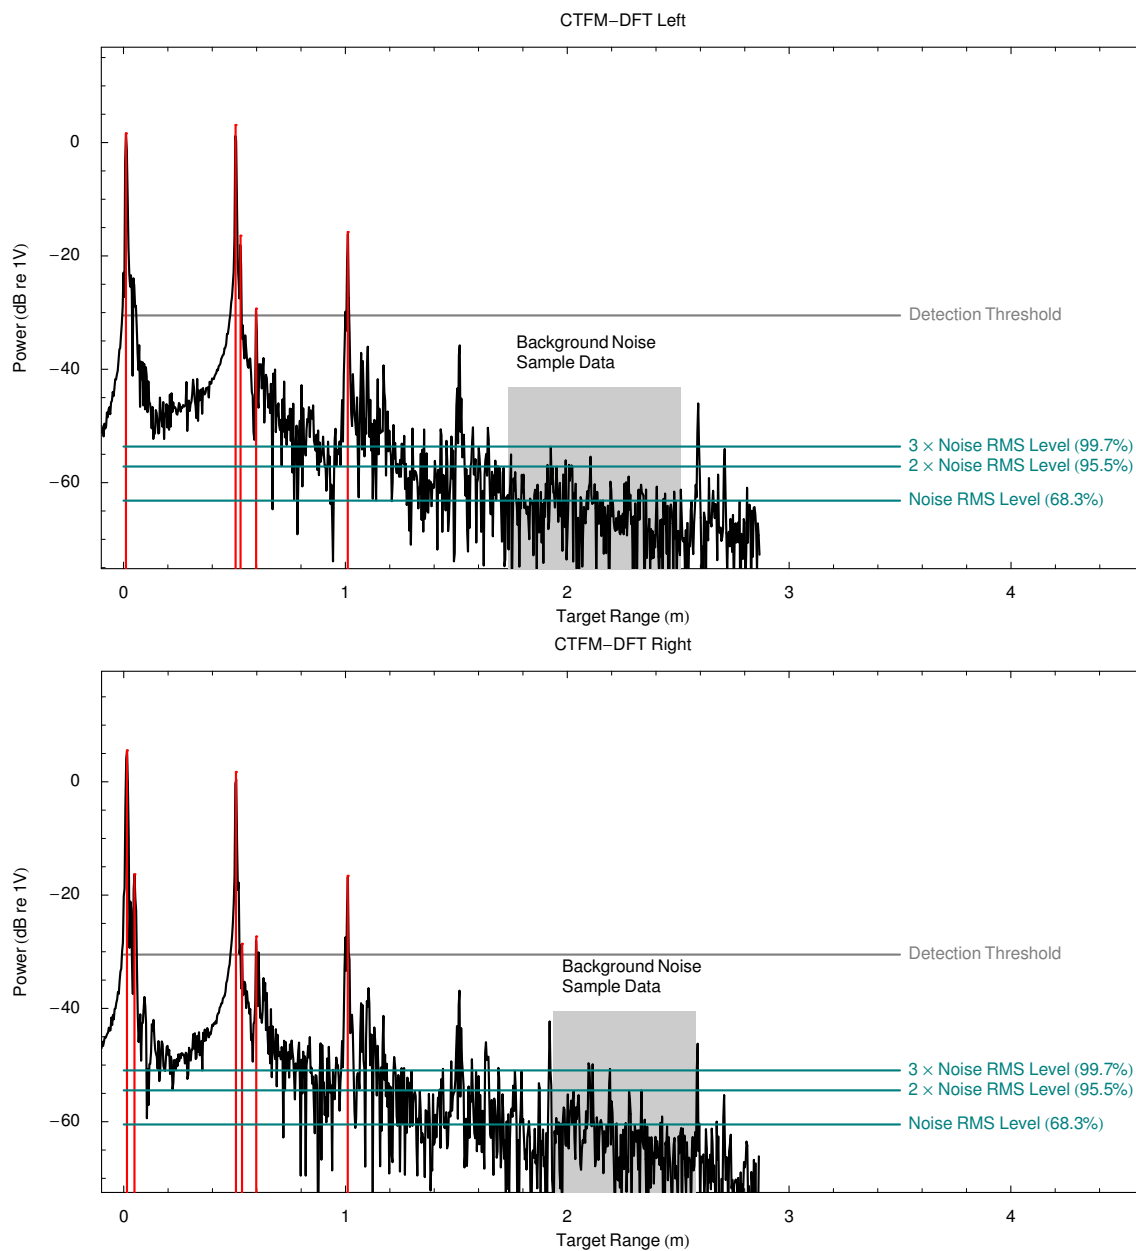


Figure 13.16: Determination of the noise level for the CTFM-DFT sonar. See p.267 for details.

MF class. This may point to a difference in the underlying datasets, as the experiment for each class was performed on a different day.

CTFM-YW Noise Level Results

The noise level of the CTFM-YW sonar system is shown in Figure 13.17. The YW sonar systems directly model the peaks of the spectrum using poles, as discussed in Section 11.2.3, but the poles do not directly model the noise floor. It is necessary to use a model with zeroes, such as an ARMA model (mentioned in Section 11.1.2), to accurately model the noise floor. For this reason, the spectral estimate appears quite structured in the region of the noise. Furthermore, there is a slope, which is also present in the CTFM-DFT data. This appears to be the sidelobes of the peaks in the spectral estimate.

The detection threshold labelled in Figure 13.17 was selected empirically as described above. The detection threshold is 3.82σ for the left channel, corresponding to a probability of false alarm of 1.3×10^{-4} under ideal circumstances. An example of non-ideal circumstances is, again, the side-lobes of the larger targets, where several additional detections may be seen. However, the possibility that these targets may be genuine has not been excluded by inspection of the experimental setup.

CTFM-LSMYW Noise Level Results

The noise level of the CTFM-YW sonar system is shown in Figure 13.18. The noise level for the CTFM-LSMYW sonar is significantly lower than that of the CTFM-YW sonar. The LSMYW spectral estimator has modelled the noise region quite differently from the YW spectral estimator. The theory of the YW based spectral estimators, in Section 11.2.3, shows that it is the peaks which are modelled rather than the noise. Therefore it is not surprising that the noise floor may be modelled inconsistently and perhaps incorrectly. Comparison of the CTFM-LSMYW results with the CTFM-DFT and CTFM-YW systems shows that the noise level is too low. This has, in turn, affected the normalised level of the power level results of Section 13.2.2. However, the power standard deviation has not been affected, as the same noise calculation was used to normalise all of the observations.

The detection threshold labelled in Figure 13.18 was selected empirically as described above. Upon comparison with the theory, the detection threshold is 7.52σ and 5.2σ for the left and right channels, corresponding to a PFA of 5×10^{-14} and 2×10^{-7} for the left and right channels under ideal circumstances. However, this spectral estimator produces many small peaks on the sidelobes of large peaks, which will not be removed by a static threshold system. These sidelobe peaks could be masked out to obtain fewer false detections.

Comparison of Noise Levels

The noise levels, signal to noise ratios, detection thresholds and PFAs measured in this section are summarised in Table 13.8.

The signal to noise ratio is measured using the target at a range of 500mm. The 3σ level is used as the noise reference for the signal to noise ratio calculation. The detection threshold is not

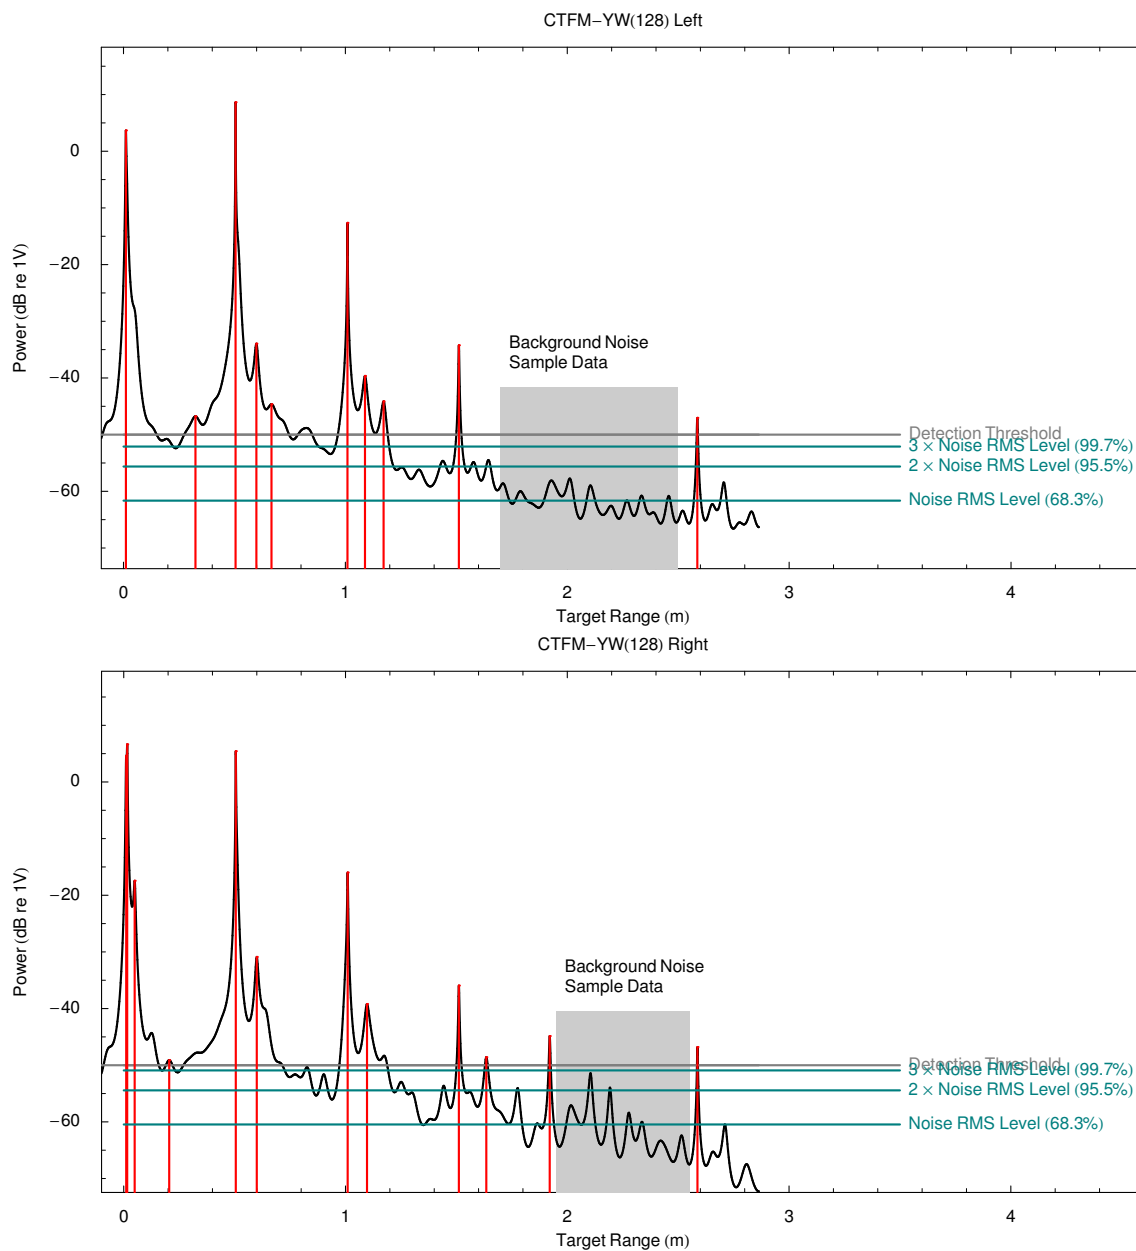


Figure 13.17: Determination of the noise level for the CTFM-YW sonar. See p.271 for discussion.

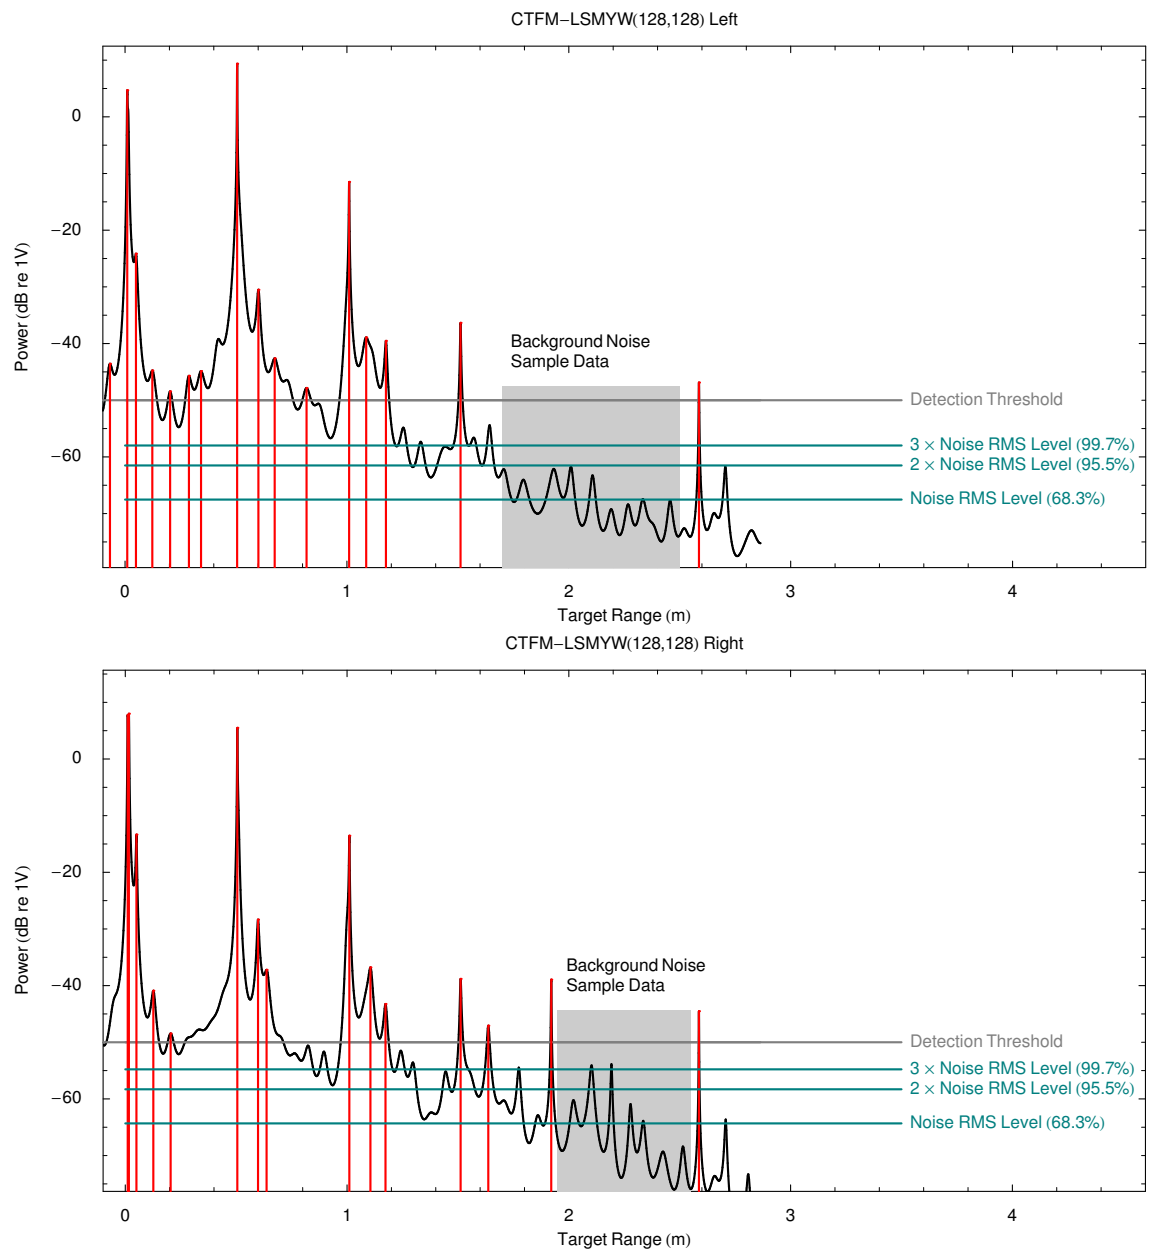


Figure 13.18: Determination of the noise level for the CTFM-LSMYW sonar. See p.271 for discussion.

Table 13.8: Summary of noise level results for the five different sonar systems. The [SNR](#) column is for a plane target at 500mm, and is measured using 3σ noise as the reference level. The noise level is tabulated in decibels referenced to a signal with a single-sided amplitude of 1 Volt.

Sonar System	Detection Threshold dB re 1 Volt	Left				Right				Results Discussed on Page
		Noise level (dB re 1V)		SNR dB	PFA (probability)	Noise Level (dB re 1V)		SNR dB	PFA (probability)	
		1σ (68.3%)	3σ (99.7%)			1σ (68.3%)	3σ (99.7%)			
MF-QD	-20	-32.3	-22.7	48.7	4×10^{-5}	-34.0	-24.4	49.0	6×10^{-7}	267
MF-UD	-20	-32.0	-22.5	48.0	7×10^{-5}	-33.7	-24.2	49.0	1.2×10^{-6}	267
CTFM-DFT	-30	-63.2	-53.6	57.0	10^{-453}	-60.5	-50.9	53.0	3×10^{-245}	267
CTFM-YW	-50	-61.6	-52.1	60.7	1.3×10^{-4}	-60.5	-50.9	56.3	8×10^{-4}	271
CTFM-LSMYW	-50	-67.5	-58.0	67.4	5×10^{-14}	-64.3	-54.8	60.3	2×10^{-7}	271

Table 13.9: The predicted maximum range of each sonar type. See text on p.275 for full details, particularly regarding the unusual CTFM-LSMYW result.

Sonar System	Maximum DOF		Maximum Range	
	Left m	Right m	Left m	Right m
MF-QD	9.4	9.5	4.7	4.7
MF-UD	9.4	9.6	4.7	4.8
CTFM-DFT	11.2	10.1	5.6	5.0
CTFM-YW	12.9	11.7	6.4	5.8
CTFM-LSMYW	16.7	17.5	8.3	8.7

normalised in any way. The PFA, P_{fa} , is computed from

$$x = 10^{(DT-\sigma)/20} \quad (13.9)$$

$$P_{fa} = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt, \quad (13.10)$$

where σ is the RMS noise level expressed in decibels, and (13.10) is the integral of the normal distribution of mean $\mu = 0$ and standard deviation $\sigma = 1$. This function is called the complementary error function or erfc (Abramowitz [1, Chapter 7]).

The CTFM systems all have a much larger signal to noise ratio than the MF systems. This is predicted by the time-bandwidth product theory described in Section 4.6, and by the TB measurements made in Tables 8.1 and 9.4. This result also indicates that we may expect the CTFM systems to have a longer maximum range than the MF sonar systems. The maximum DOF of each system will be evaluated in the next sub-section.

Maximum DOF

The maximum DOF at which each sonar system is capable of detecting a plane target may be calculated using (4.2) and (4.3), by finding the DOF which causes the echo excess to reach zero. The calculations were performed using the echo power information tabulated in Tables 13.6 and 13.7. Thus the calculations assume a plane reflector with the normal reflection point which lies on the acoustic axis of the sonar head. Targets having a lower sonar cross section or which are placed off-axis will suffer a lower maximum range. A PFA of 10^{-7} was assumed for calculating the detection threshold. The calculation is performed using an inverse form of (13.10), which is provided by Mathematica [139]. Thus the detection threshold used for determining the maximum range was 1.976 dB above the 3σ noise level, which is 11.51 dB above the RMS noise level. The results are shown in Table 13.9. Before discussing the results in general, it is worth looking into the unusually high value obtained for the maximum DOF of the CTFM-LSMYW system. This has come about for several reasons:

1. The estimate for the air absorption in Table 13.7 is unusually low, perhaps due to the spectral estimate being perturbed by a second target nearby.

2. The noise level estimate from Figure 13.18 is lower than that of all the other CTFM systems which were applied to the same data set. The noise level in the signal cannot change. The low value measured by the LSMYW estimator may be because the YW system does not specifically model the noise floor.

This leads to the conclusion that the CTFM-LSMYW maximum DOF computation is invalid. More conclusive results may be obtained by an experiment that measures the maximum DOF directly.

The calculated maximum DOF data shown in Table 13.9 demonstrate that the CTFM systems are expected to achieve an extra metre of range over the MF systems in noise limited operation. This is due to the higher SNR achieved by the CTFM systems.

13.2.5 Summary of DOF Precision Results

The DOF and echo power precision experiment has provided data for the six attributes labelled (a) in Table 13.1. The six attributes are:

1. DOF Accuracy. The DOF precision experiment of Section 13.2.1 was designed to allow the accuracy to be evaluated. The parameters governing the accuracy are the offset time, t_{off} , and the difference between the calculated and measured speed of sound, Δc . These are summarised in Tables 13.2 and 13.3. However, the difference between the calculated speed of sound and the measured speed of sound caused the accuracy to be poor. Therefore the absolute DOF accuracy has not been rigorously assessed.
2. DOF Precision, measured in Section 13.2, and summarised in Table 13.4. The DOF precision is a factor in determining the precision of the range and the IDD bearing estimates, as explained in Section 14.1.3.
3. Echo Power Precision, measured in Section 13.2.2, and summarised in Table 13.5. The echo power precision of the left and right channels determines the IPD bearing precision, as explained in Section 14.2.4.
4. DOF Cross-correlation, measured in Section 13.2.3, and summarised in Figure 13.13. The DOF cross correlation is a factor in determining the range and IDD bearing precisions, as explained in Section 14.1.3.
5. Echo Power Cross-correlation, measured in Section 13.2.3, and summarised in Figure 13.13. The echo power cross correlation is a factor in determining the IPD bearing precision, as explained in Section 14.2.4.
6. Noise Level and Signal to Noise Ratio, measured in Section 13.2.4, and summarised in Table 13.8. The noise level at the output of each sonar system, together with the power levels determined in the echo power precision experiment in Section 13.2.2, was then used to estimate the maximum range of each sonar type. These results are summarised in Table 13.9.

Several important observations were made in the course of the discussion of the DOF precision experiment results. These observations may be divided into problematic, neutral and beneficial categorisations. The following results are regarded as problematic:

1. A significant difference was observed between the calculated and the observed speed of sound, as tabulated in Table 13.3. This prevented the assessment of the absolute DOF accuracy. Solving for a correction factor and understanding that this factor is sub-class dependent is considered an important result.
2. The CTFM-YW and CTFM-LSMYW echo power results, seen in Figures 13.10 and 13.11, display a perturbation from the ideal echo power model. It was determined that this was due to interference from other close echoes, which the CTFM-YW and CTFM-LSMYW systems could not resolve. The fact that the CTFM-DFT system was not affected indicates that the fault lies in the YW and LSMYW spectral estimators. This will be seen again in the resolution experiment in Sections 13.3.4 and 13.3.5. This result impacts upon the usefulness of the YW based CTFM sonar systems.
3. The CTFM-DFT DOF cross-correlation results display tram-tracks, as seen in Figure 13.13. This is due to the behaviour of the interpolation scheme when the actual peak falls near a spectral line.
4. The two YW systems, and CTFM-LSMYW in particular, have difficulty in estimating the background noise levels, as seen in Section 13.2.4.
5. All of the experiments display small differences between the left and right channels. This is due to the manufacturing tolerances of the hand-made DAC and ADC cards (described in Sections 5.1.1 and 5.1.2), transducers (Chapter 6), and amplifiers (Section 6.6).

The following observation has a neutral effect upon the operation of the sonar system:

1. An offset time was measured for all sonar systems, as summarised in Table 13.2.

The following results are regarded as beneficial:

1. The compensation of the DOF measurements has showed a significant reduction in the residuals, as shown in Figure 13.9. However, there is some additional effect which is not compensated.
2. All systems, with the exception of the CTFM-DFT system, have achieved good DOF precision results, around ± 1 mm. The CTFM-DFT system was noted to have problems due to its interpolation scheme.
3. The MF-QD, MF-UD and CTFM-DFT systems have achieved good power precision results, as summarised in Table 13.5.

4. The left and right DOF residuals show high positive correlation (see Figure 13.13), indicating that the errors originate in the common air path. This will improve the precision of bearing calculations, at the expense of range precision. However, the bearing calculations depend more critically upon DOF precision than range calculations do, so the trade-off is seen as a net benefit in this case.

This concludes the range precision experiment. The range resolution experiment, will be discussed next.

13.3 The Range Resolution Experiment

The range resolution experiment corresponds to experimental configuration (b) in Table 13.1. The configuration is shown in Figures 13.1b and 13.3, and is aimed at determining the ability of each sonar system to resolve two pole targets as being separate. The poles, visible in Figure 13.3, both have a diameter of 12.7 mm. The coordinates of the front face of the fixed pole were $\{x, y\} = \{943.2, 138.7\}$ mm, giving it a range of 953.3 mm. The movable pole was moved along a straight line between the coordinates $\{700, 90\}$ mm and $\{1200, 90\}$ mm in steps of 1 mm. At each of the 501 positions, 3 observations were recorded, as described in Section 13.1.3. However, instead of recording just the echo which was closest to the expected DOF, all of the echoes within the range 0.7 m to 1.2 m were recorded.

The results for each of the various sonar types are shown in Figures 13.21, 13.24, 13.26, 13.27 and 13.29. Each dot in the figures represents a detected echo. These dots are intensity coded by colour, according to the legend shown in Figure 13.20. The horizontal axis of each plot indicates the range from the sonar to the reflecting face of the moving pole. This moving pole is initially in front of the fixed pole and at the end is behind the fixed pole. Thus range resolution is half of the range difference from the first range of a single echo to the last range of a single echo. In this experiment, simple range is considered to be half of the DOF. A proper formulation of range is developed in Chapter 14. The vertical axis of the resolution results figures indicates the measured DOF of the echoes. The decibel power of each echo is plotted as the colour of a dot according to the colour scale shown in Figure 13.20.

The echo from the fixed pole is observed as a horizontal line of dots having a large power at constant DOF, $d = 1.92$ m. The echo from the movable pole appears as the lowermost upward sloping line with strong power.

While there were only two principal targets, echoes have been recorded from other targets as well. In particular, there are strong secondary echoes coming from a target which has a motion similar to the moving target. The most likely reflectors are the mounting stage of the positioner and the traverse motor. The change in the secondary DOF relative to the principal echo is explained by considering that the difference in radii between the pole reflecting point and a reflection from the positioner base, as shown in Figure 13.19, becomes smaller as the positioner moves further away from the sensor; it will eventually have a smaller range than the echo from the moving pole target. There are two secondary echoes visible in the plots, coming from the moving positioner base. There is also a fixed secondary echo, visible in the MF experiments, which is presumed to come from the retort stand shown in Figure 13.3.




Figure 13.19: Secondary reflections coming from the positioner base have a larger range than the reflections from the pole at short ranges. At longer ranges than shown here, the secondary reflection will have a shorter path length than the primary reflection.

The results for each type of sonar are now presented and discussed in turn.

13.3.1 Matched Filter QD Resolution Results

The results from the matched filter sonar with quadrature demodulation, shown in Figure 13.21, display a good ability to resolve close targets. However, there are a huge number of weak echoes bearing no resemblance to physical targets. This is noise in the sonar system, not eliminated by the noise threshold, which was set to -20 dB re 1 V amplitude, as shown in Table 13.8. The noise level in this data set varies with each record, generally increasing as the experiment progresses, as shown in Figure 13.22. The theory discussed in Section 13.2.4 shows that the PFA is greatly reduced by this fault. The false detections may be eliminated by increasing the detection threshold,




Figure 13.20: Legend indicating the signal strength of echoes, in decibels. This legend applies to all of the range resolution plots in Section 13.3.

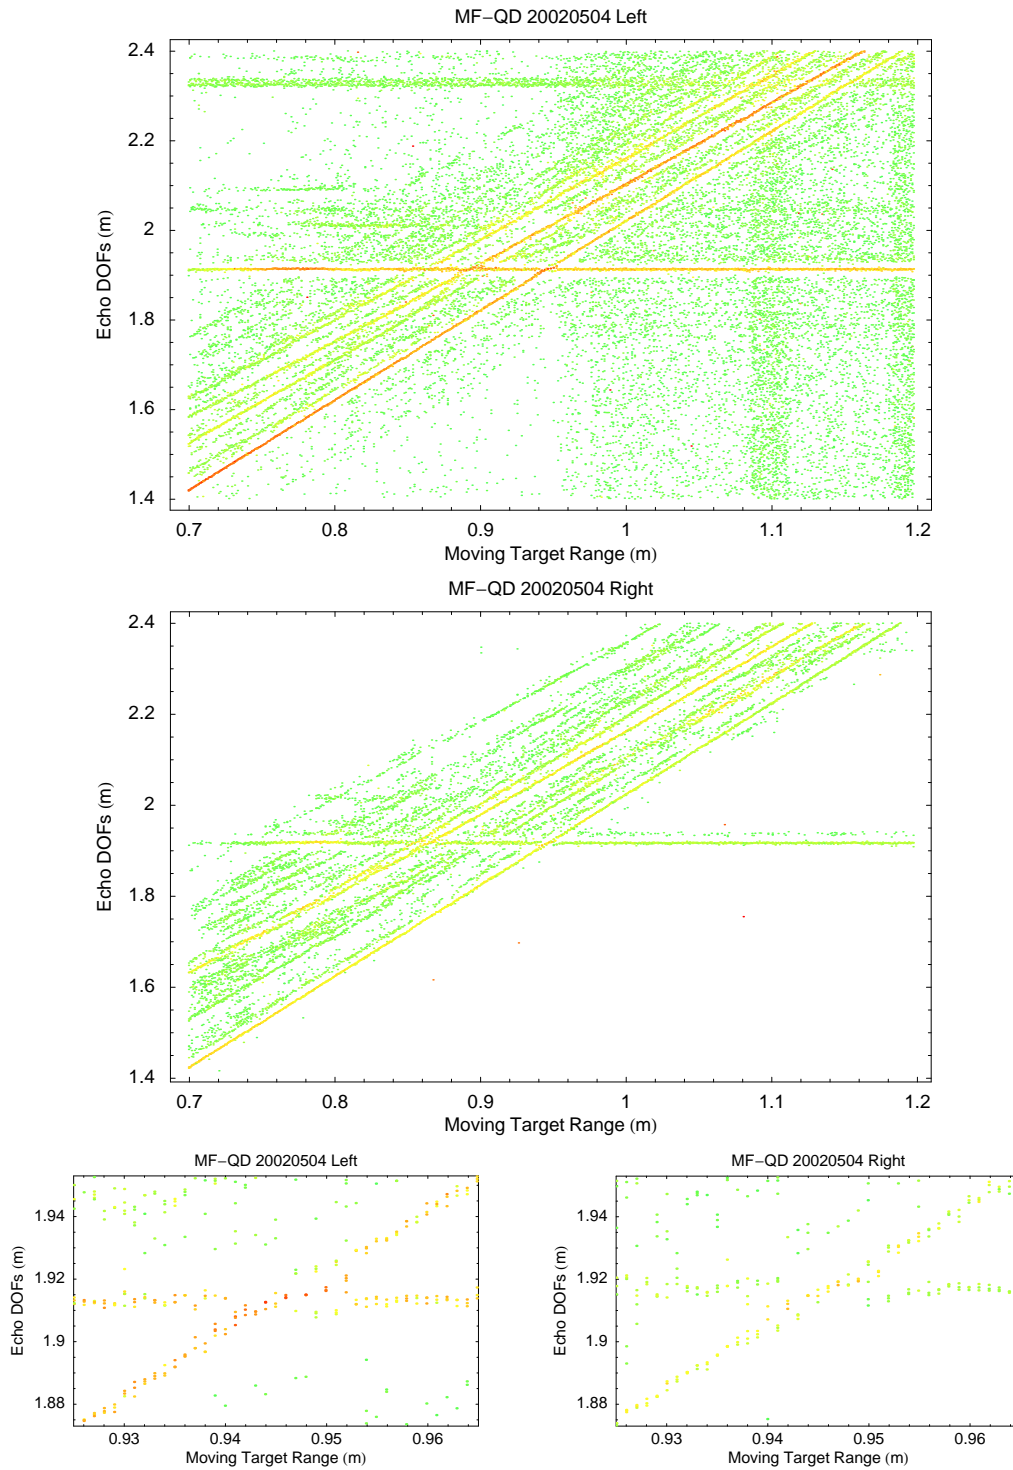


Figure 13.21: Resolution results for the matched filter quadrature demodulation sonar system. The entire set of results is shown in the upper two panels, while closeups of the crossover regions are shown in the lower two panels. Each dot represents an echo, with the DOF indicated by the vertical position, and the power indicated by the colour, according to the legend shown in Figure 13.20. The fixed target has a DOF of 1.913 m, while the moving target has a DOF of approximately twice its range (the horizontal coordinate).

but this also eliminates the echoes from the fixed target on the right channel. The weakness of the echo from the right channel can be seen in Figure 13.21 as it is plotted in green, rather than the red/yellow colour appearing on the left channel. The difference in received echo strengths is attributed to the different pointing angles of the receivers, and will be discussed more during the computation of bearing by inter-aural power difference in Section 14.2.

Putting aside the false detections, the resolution of the system can still be determined from this data set. Closeups of echo detections in the crossover region, shown in Figure 13.21, show that the resolution of this sonar system is 6mm in both channels. This figure also shows that the echoes are not significantly perturbed from their true positions until the target resolution criterion is reached. The recorded target distance appears to be the average of the two target distances within the unresolvable zone.

Raw Received Echo

Figure 13.23 shows the raw echo before any processing as it is received from the two pole targets. The right echo is barely above the noise floor, and the sampling quantisation level is visible in the plot. It appears that some of the noise is in the same proportion to the signal on both channels, indicating that the noise source could be before an amplifier stage which has an unequal gain between the two channels.

The echo power, in decibels, received from a pole target is described by

$$P_r = A_1 - 40\log_{10} r - 2\alpha' r,$$

where r is the range to the target, A_1 is the normalised measured source level, and α' is the mean absorption over the signal bandwidth, as defined in Section 13.2.2. The main difference from the echo power received from a plane target, described by (13.7), is the $40\log_{10} r$ term. This difference is caused by the spreading loss on the outward path followed by the spreading loss from a point reflector on the return path, as discussed in Chapter 4. This term is responsible for the echo being so much weaker than the plane reflection studied in Section 13.2.2.

Demodulated Waveform

The demodulated output waveform may be seen in Figure 13.25a for the case where the two targets are resolved.

13.3.2 Matched Filter UD Resolution Results

The **matched filter** sonar with unsynchronised demodulation, shown in Figure 13.24, gives very similar results to the MF-QD sonar. The similarity is to be expected as the only difference between the sonar systems is in the envelope extractor. The MF-UD sonar produces peaks with a slightly higher amplitude. There appear to be more false positives in this data set than in the MF-QD data set, which is attributed to the difference in the peak extractor. The underlying data set used is the same as was used for the MF-QD experiment, and contains the same noise fault as was discussed in that experiment. The detection threshold was set to -20dB re 1V amplitude, as shown in

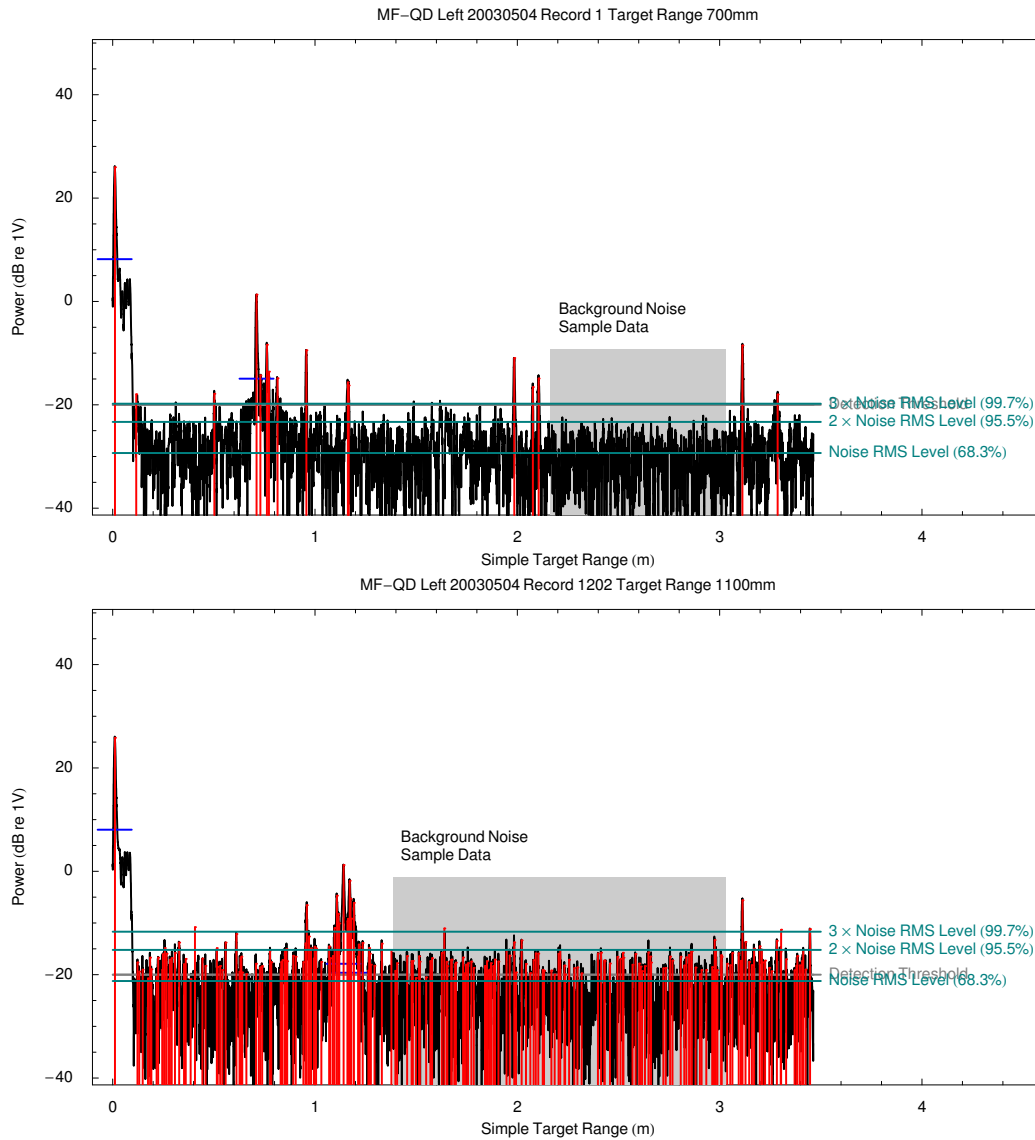


Figure 13.22: Investigation of the noise level in the MF-QD sonar system for the resolution experiment. The top plot corresponds to the measurements drawn at the left of the top panel (left channel) of Figure 13.21, while the lower plot corresponds to the measurements drawn in the middle panel (right channel) of Figure 13.21. The background noise at the 3σ level for record 1 is -19.8 dB re 1V , while the equivalent measurement for record 1202 is -11.7 dB re 1V . Both of these noise levels are higher than those recorded in Table 13.8. The detection threshold is set to -20 dB re 1V , and remains constant throughout the experiment. The PFA reduces from 3.4×10^{-3} for record 1 to 0.25 for record 1202. The increase in noise level indicates a fault in the equipment.




Figure 13.23: Raw echo waveform for matched filter sonar, measured in volts. Movable target range is 0.934 m, fixed target range is 0.953 m. The approximate TOFs are 5.40 ms and 5.51 ms.

Table 13.8. However, putting aside the false detections, the resolution of the system can still be determined from this data set.

The closeups at the bottom of Figure 13.24 show that the resolution is 6 mm for both the left and right channels, which is the same as the quadrature demodulation case.

Demodulated Waveforms

Figure 13.25b shows the same measured data set as Figure 13.25a, but processed by the MF-UD sonar system. The behaviour of these systems is quite similar as the extracted envelopes are quite similar. The main difference is that the QD envelope is slightly rougher than the UD envelope.

Figure 13.25c shows the demodulated echoes where separation between the targets is only 4.3 mm. The echoes have fused together, rendering separate detection impossible. Note that the echo powers have added to form a peak larger than either of the individual peaks in Figure 13.25b.

13.3.3 CTFM-DFT Resolution Results

The results for this sonar are shown in Figure 13.26. This sonar system has slightly improved resolution over the MF sonar system, despite having a lower DOF precision (see Table 13.4). The reduced DOF precision is also visible in all plots of Figure 13.26 as a wider track of dots than is produced by the other sonar systems. The detection threshold was set to -30 dB re 1 V amplitude, as shown in Table 13.8.

The resolution of this sonar system can be measured from the closeup plots in Figure 13.26 to be 5 mm for both the left and right channels.

An important feature of this sonar system is that the pole targets were detected without lowering the detection threshold. It remained at its standard setting of -30 dB re 1 V. A quick check of the data presented to the detector over multiple data sets shows that the noise level is consistently similar to that measured in Table 13.8. Thus the PFA is maintained, and is reflected in the clarity of Figure 13.26, which lacks the noise seen in the matched filter results shown in Figures 13.21

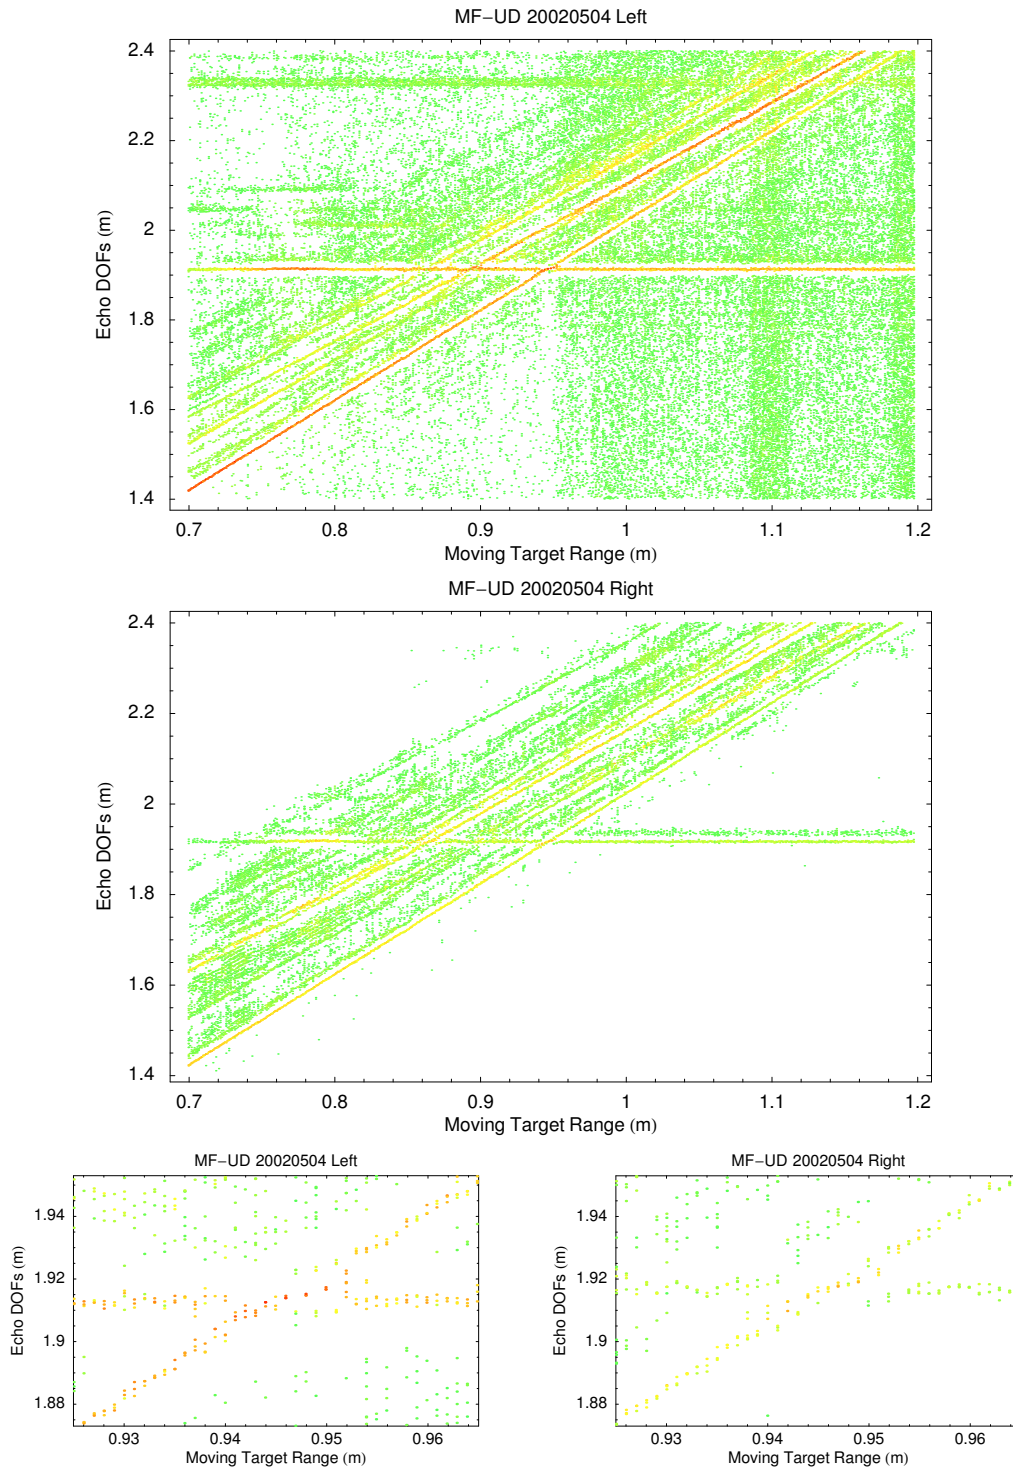





Figure 13.24: Resolution results for the matched filter unsynchronised demodulation sonar system. The entire set of results is shown in the upper two panels, while a closeup of the crossover region is shown in the lower two panels. Each dot represents an echo, with the DOF indicated by the vertical position, and the power indicated by the colour, according to the legend shown in Figure 13.20. The fixed target has a DOF of 1.913 m, while the moving target has a DOF of approximately twice its range (the horizontal coordinate).



(a) MF-QD output, target range 0.944m



(b) MF-UD output, target range 0.944m



(c) MF-UD output, target range 0.949m

Figure 13.25: Output of the **matched filter** sonar, in volts. In (a), the MF-QD sonar was used, and the target range was 0.944m, while the fixed target was at 0.953m. The corresponding **TOFs** are about 5.53ms and 5.48ms. With the two targets separated by 9mm the sonar can resolve them. In (b) the MF-UD sonar was used, with the same target ranges and **TOFs** as before. The two targets are similarly resolved. In (c) the MF-UD sonar was used, with a movable target range of 0.949m, giving a separation of 4mm. The echoes are indistinguishable at 5.52ms. The echo amplitude is also increased by comparison with the amplitudes of the separate echoes in (b). The echoes detected by the right receiver are weaker due to the difference in the angles subtended by the echo at the receivers.

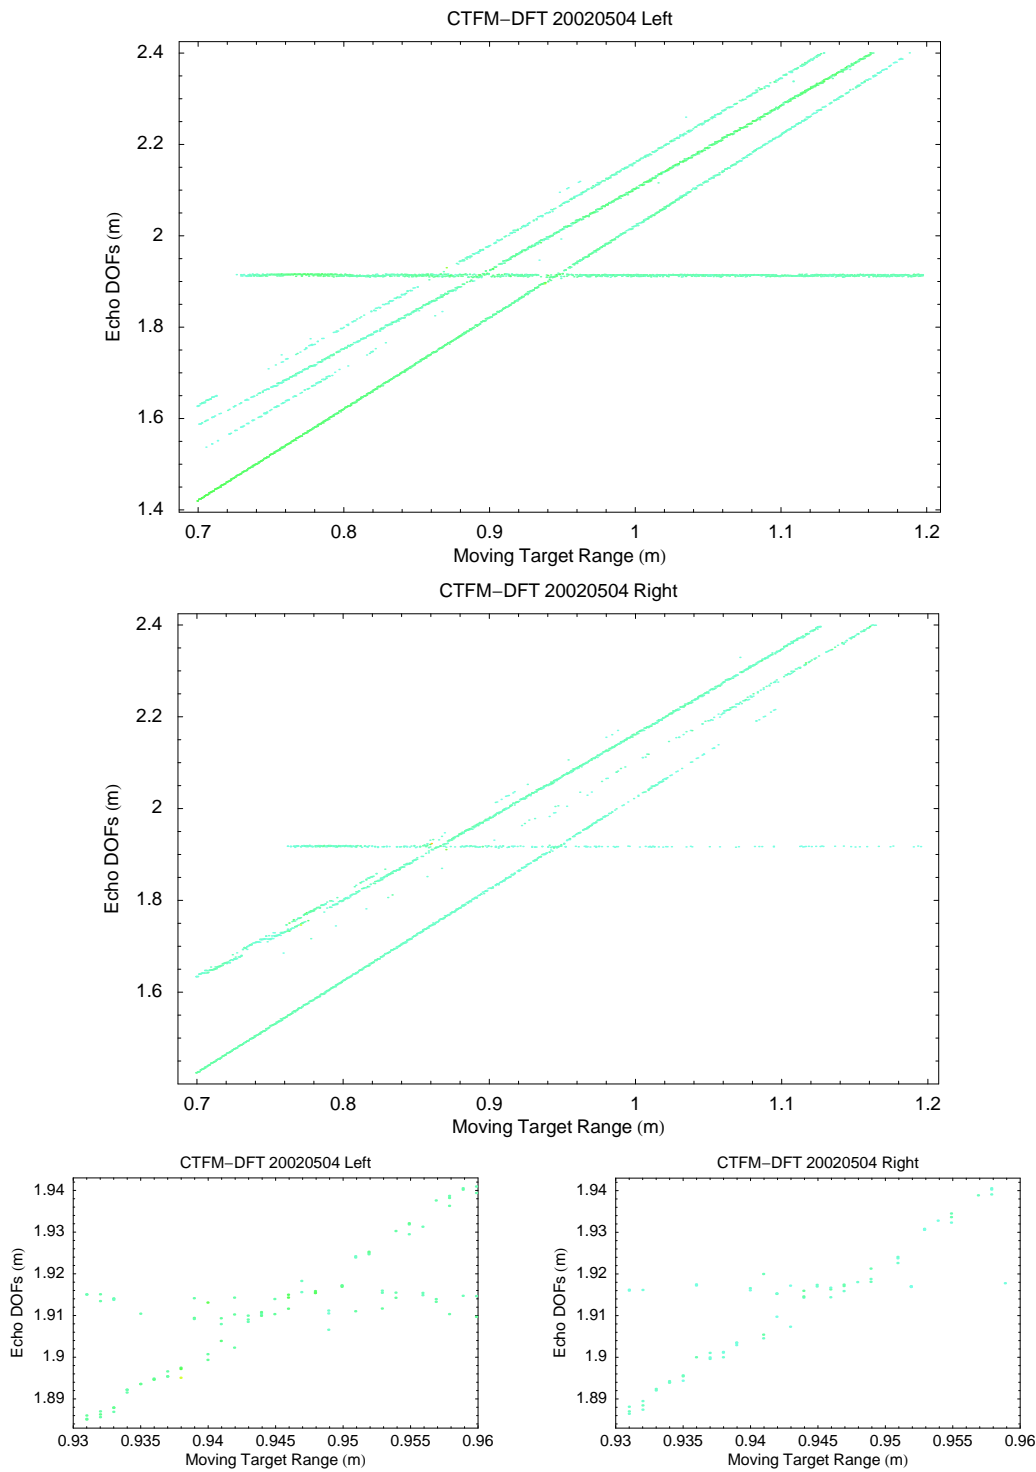


Figure 13.26: Resolution results for the CTFM sonar system with a DFT spectral estimator. The entire set of results is shown in the upper two panels, while a closeup of the crossover region is shown in the lower two panels. Each dot represents an echo, with the DOF indicated by the vertical position, and the power indicated by the colour, according to the legend shown in Figure 13.20. The fixed target has a DOF of 1.913 m, while the moving target has a DOF of approximately twice its range (the horizontal coordinate).

and 13.24.

Spectra

The output of the CTFM-DFT sonar system is shown in Figure 13.28a. The movable target causes a peak at 3475 Hz, while the fixed target peak is located at 3455 Hz. The peaks are clearly resolved on both channels. The coarseness of the DFT is clearly evident in this plot.

Examination of a sequence of several DFTs corresponding to observations taken with the same experimental setup reveals that the spectra change considerably from moment to moment. The variation of low power levels is attributed to noise which is commonly eliminated by using the Welch [138] method, but using the Welch method on this data would further reduce the available precision. The variation at high power levels may be due to variations in the environmental parameters, causing slight changes in the observed TOF, which are translated into large changes in the DFT output because of the quantisation of the spectral lines. The effect of varying environmental parameters may be gauged by comparing the results from one of the other sonar systems. As all of the other sonar systems have a better precision than this sonar, this suggests that nearly all of the variation in the peaks detected by DFT is due to the quantisation of any noise present.

13.3.4 CTFM-YW Resolution Results

The CTFM Yule-Walker sonar system results, shown in Figure 13.27, display a significant perturbation in the measured DOF as the two targets approach each other. This effect extends beyond the point where the two targets are resolvable, extending out to targets separated by as much as 25 mm. Secondary echoes have a significant impact. This perturbing effect has the potential to destroy bearing precision for targets which are separated in range, although the effect upon the bearing estimate may be reduced by the fact that both channels are affected approximately equally. The detection threshold was set to -50 dB re 1 V amplitude, as shown in Table 13.8.

The resolution of this sonar system, being 20 mm, is the worst of all the systems tested, but like the CTFM-DFT system, it produces noise free output.

Spectra

The output of the CTFM-YW(128) sonar system is shown in Figure 13.28b, for the same target separation (11 mm) as shown in the CTFM-DFT sonar in Figure 13.28a. The smoothness of the estimate compared to the DFT is immediately apparent. However, with this smoothness comes a lack of resolving power. When the peak corresponding to the moving target finally emerges, at a target separation of 26 mm, it is a very weak peak (Figure 13.28c), and is only visible on the right channel.

The Yule-Walker spectral estimator operates by placing poles within the z -plane, as described in Chapter 11. The lack of precision observed in these results indicates that the capacity of the spectral estimator to precisely separate the effects of two closely spaced frequencies is limited. This translates to a reduced ability to accurately place two poles close together.

The Yule-Walker spectra shown in Figures 13.28b and 13.28c display considerable variation between the left and right channels. It is hypothesised that this is caused by the same mechanism







Figure 13.27: Resolution results for the CTFM sonar system with a Yule-Walker spectral estimator. The entire set of results is shown in the upper two panels, while a closeup of the crossover region is shown in the lower two panels. Each dot represents an echo, with the DOF indicated by the vertical position, and the power indicated by the colour, according to the legend shown in Figure 13.20. The fixed target has a DOF of 1.913 m, while the moving target has a DOF of approximately twice its range (the horizontal coordinate).



(a) CTFM-DFT sonar, target range 0.964m



(b) CTFM-YW sonar, target range 0.964m



(c) CTFM-YW sonar, target range 0.979m

Figure 13.28: Output of the CTFM sonar. In (a), the DFT spectral analyser was used, and the movable target range is 0.964m, while the fixed target range is 0.953m. The corresponding frequencies are 3475Hz and 3455Hz respectively. The two targets are clearly resolved by this sonar system at this separation of 11mm. In (b) the YW(128) spectral analyser was used, for the same position of the movable target as in (a). The two targets are not resolvable by the YW(128) spectral analyser at this separation of 11mm. In (c) the movable target range is 0.979m, with the target resolved on the right channel at 3525Hz, but not on the left. The targets are separated by 26mm at this point. The other peaks present in the plots are the secondary echoes described previously.

which causes the position of the estimated peak to be perturbed when there are two close targets.

It is possible that the performance of the YW estimator could be improved by using more than 128 poles, but this has not been investigated. Kay and Marple [75, Section IIE] state that using too many poles causes false peak splitting, which would be detrimental.

13.3.5 CTFM-LSMYW Resolution Results

The CTFM least squares modified Yule-Walker sonar system results, shown in Figure 13.29, contain a much smaller DOF perturbation than is present in the CTFM-YW sonar, affecting measurements of targets separated by a maximum of 17.5 mm. The estimated echo power is increased out to a maximum separation of 42 mm, at which point it returns to the normal level. Unlike the other sonar systems, the echo power is not significantly increased in the irresolvable region. The detection threshold was set to -50 dB re 1 V amplitude, as shown in Table 13.8.

The resolution of this sonar system can be measured from the closeup plots in Figure 13.29 to be 11.0 mm for the left channel and 16.5 mm for the right channel.

This sonar system produces enough signal that echoes from weak targets exceed the noise floor and are reliably detected, without detection of noise.

Spectra

The output of a CTFM-LSMYW sonar system is shown in Figure 13.30a for a target separation of 11 mm. The echoes have fused together at 3465 Hz. The peak on the left channel has significantly broadened, but does not display a dip and is not separable. When the targets are separated by 21 mm (Figure 13.30b), two peaks are clearly visible at 3450 Hz and 3496 Hz. The LSMYW spectra shown in Figure 13.30 display less variation between the left and right channels than those estimated by the YW estimator in Figure 13.28.

The improved resolution behaviour of the least squares modified Yule-Walker spectral estimator over the YW estimator is attributed to the use of 128 extra values of the autocorrelation of the signal to perform the least squares operation (Cadzow [20]). However, it seems that the inclusion of these extra autocorrelation points leads to a degraded power precision. Cadzow [20] indicates that the autocorrelation estimates with larger lags are averaged over fewer samples, producing less precise estimates. It is hypothesised that this is causing the variation of the estimated echo power which is seen in the results.

13.3.6 Summary of Resolution Results

The resolution results for each of the five sonar systems are summarised in Table 13.10, along with theoretical resolutions predicted using the time-bandwidth product theory introduced in Section 4.6.

Resolution

The experimental resolution can be determined by identifying the area where the two echoes are fused together, that is, only one echo is detected by the sonar system, usually with increased power.






Figure 13.29: Resolution results for the CTFM sonar system with a least squares modified Yule-Walker spectral estimator. The entire set of results is shown in the upper two panels, while a closeup of the crossover region is shown in the lower two panels. Each dot represents an echo, with the DOF indicated by the vertical position, and the power indicated by the colour, according to the legend shown in Figure 13.20. The fixed target has a DOF of 1.913 m, while the moving target has a DOF of approximately twice its range (the horizontal coordinate).



(a) target range 0.964m



(b) target range 0.974m

Figure 13.30: Output from the CTFM sonar with LSMYW(128,128) spectral analyser, power in decibels. In (a) the movable target range is 0.964m, while the fixed target range is 0.953m. The targets are separated by 11 mm, and the echoes have fused together at 3465 Hz. In (b) the movable target range is 0.974m, being separated from the fixed target by 21 mm. The targets produce well resolved peaks at 3450 Hz and 3495 Hz.

Table 13.10: Results for the minimum separation at which the echoes from two targets may be resolved. The left, right and average columns indicate the minimum range difference at which two targets may be resolved. Results from the various theoretical methods are also included for comparison.

Sonar System	Range Resolution (mm)						Results
	Experimental			Theoretical Method			Discussed in Section
	Left	Right	Average	1	2	3	
Refer to Section				8.8.1	8.8.3 9.9.2	8.8.4 9.9.3	
MF-QD	6	5	6	3.2	7	5.8	13.3.1
MF-UD	6	6	6	3.2	7	5.8	13.3.2
CTFM-DFT	5	5	5	3.2	5.5	3.7	13.3.3
CTFM-YW(128)	20	20	20	3.2	5.5	3.7	13.3.4
CTFM-LSMYW(128,128)	11	17	14	3.2	5.5	3.7	13.3.5

The width of this region is divided by two to obtain the minimum resolvable distance between two targets, which is shown in Table 13.10 in the left, right and average columns. As noted before, differences between the left and right channels should be small. The differences are thought to be primarily due to the individual manufacture of the transducers. These differences do not manifest themselves in the experimental resolution results shown in Table 13.10.

The echo power in the overlap region is stronger than outside the overlap region in all systems except the CTFM-DFT system. This is attributed to the echoes adding together constructively as they overlap.

The MF and CTFM-DFT systems are approximately comparable, but the YW based systems are greatly inferior at resolving close targets. Moreover these systems have a significantly inferior performance compared to theoretical expectation. This is primarily due to the inability of the YW based spectral estimators to reduce the spectral level between the two peaks (poles) by using a zero. Other models e.g. ARMA and Prony (see Kay and Marple [75]) may have performed better. This has not been tried here due to the additional complexity.

Perturbed Echoes

The CTFM auto-regressive based sonar systems produce perturbed DOF and echo power output as the two targets approach each other. The minimum distance of separation at which the DOF and echo power outputs were accurate is 25mm for the CTFM-YW system, and 18mm for the CTFM-LSMYW system. These are the minimum separation distances which will produce accurate position measurements.

Signal to Noise Ratio

Section 13.3.1 showed that there was a fault with the sonar system during the data gathering of the MF experiment which increased the noise level above that measured in Section 13.2.4. Therefore it is inappropriate to compare the number of false detections between the MF and CTFM sonar

systems for this experiment.

Time-Bandwidth Product

The theoretical resolution of all five sonar systems was predicted in Sections 8.8 and 9.9, using the [time-bandwidth product](#) theory introduced in Section 4.6. The experimental results may now be compared with the theoretical results. Both sets of results are summarised in Table 13.10. Three theoretical methods were applied, labelled methods 1 to 3. Method 1 estimated the resolution limit from the system bandwidth. It does not allow for practical effects which occur in the signal processing, and so produces a very optimistic result. Method 2 directly measured the width of the function presented to the detector. The results obtained using this method were adjusted in the case of the [MF](#) system, as the sidelobes of the autocorrelation output artificially expanded the width. Method 3 measured the width of the transform of the signal presented to the detector, and subsequently applied Gabor's relation to estimate the minimum possible width. The roles of the two domains are reversed between the [MF](#) and [CTFM](#) sonar systems, as discussed in Section 4.6. The application of Gabor's time-bandwidth product relation assumes that the signal is the limiting case of a Gabor elementary signal (described in Section 4.6). The methods may be ordered in the assumptions they make about the system. Method 1 makes the most simplifying assumptions, method 3 makes an intermediate number of assumptions, while method 2 estimates the detector signal width function most directly and takes into account the most practical information available. The divergence between the various techniques is an indication of how far the implementation is from the ideal.

Method 2 has produced results which are larger than the measured resolution in the case of the [MF](#) sonar systems. The result for the [CTFM](#) systems is a very close match to the experimental result in the case of the [CTFM-DFT](#) sonar, but the two [YW](#) based systems provide much worse resolution results than the [time-bandwidth product](#) theory indicates is possible. As the underlying data and the demodulation processing are the same, the resolution failure is blamed upon the [YW](#) systems. This is due to the rational model being unable to model the signal with sufficient accuracy to resolve the signals as they get quite close.

Method 3 has produced results which are lower than the measured resolution figures in all cases except the right channel of the [MF-QD](#) sonar system. This appears to be an anomaly, perhaps due to the difficulty of reading the plot in Figure 13.21. The theoretical resolution limits obtained by method 2 are only a little below the measured resolution in the case of the [MF](#) class of sonar systems, indicating that this class has achieved state-of-the-art resolution. The [CTFM-DFT](#) system achieves a resolution which is 1.3mm above the theoretical limit. It appears that the [DFT](#) based system is in fact limited by the coarseness of the underlying spectral lines, which corresponds to 3.7mm in range (see Section 10.2.4). Resolution of two targets closer than two spectral line widths, or 7.4mm, is not possible in a single measurement. However, the experimental results show that resolution is possible, due to the air noise, by averaging over several measurements.

The close agreement between the theoretical resolution results for the [MF](#) sonar systems and the [CTFM-DFT](#) based system indicates that these systems have are close to their resolution limit. The reduced resolution of the [AR](#) based systems compared to the resolution capability of the [CTFM](#) signal is attributed to the [YW](#) and [LSMYW](#) spectral estimators.

13.4 Conclusion

This chapter has described two experiments, the precision experiment and the resolution experiment, as detailed in Table 13.1. These experiments have evaluated the performance of the five sonar systems using just two experimental configurations, making a total of four data sets. This has allowed the evaluation of a total of 35 direct parameters, and many more indirect parameters, thereby enabling the inter-comparison aims of the experiments to be fulfilled.

The DOF and power precision results found in Sections 13.2.1 and 13.2.2 of this chapter determine the limit with which the bearing may be determined in Chapter 14. The resolution results of Section 13.3, while only performed as range measurements, will also determine the minimum resolvable bearing, which will be discussed in Chapter 15. The correlation results play an important secondary role in determining the bearing precision. It will be discussed in detail in Chapter 14.

The best DOF precision results are obtained by the matched filter class of sonar systems, the unsynchronised demodulation sub-class being superior to the quadrature demodulation sub-class. However, the maximum range of these systems is less than that of the CTFM class.

The MF class has, in general, about 11 dB less signal to noise ratio (see Table 13.8) than the CTFM class. This, in turn, is reflected in the higher predicted maximum range (Table 13.9) of the CTFM class of sonar systems. However, this increase in range over which the sensor can be used has been won at the expense of DOF precision and resolution.

It could be argued that the reduction in precision of DOF from values less than 1 mm in MF systems to values greater than 1 mm in CTFM systems is offset by the increased versatility of the CTFM systems.

The resolution results of Table 13.10 indicate that the YW based sonar systems have difficulty in resolving closely spaced targets. Additionally, the LSMYW system also has difficulty estimating the correct noise floor, as seen in Section 13.2.4. However, this is not seen as an issue during normal system operation, as the noise floor only needs to be established during system characterisation and configuration.

The theoretical resolution of each class of sonar system was compared with the measured resolution, and it was found that the MF class of systems was performing very close to the theoretical limit. The CTFM-DFT system was found to be performing near the theoretical limit, but the two YW based systems performed poorly in comparison to the theoretical limit.

Two other important results should be noted. The first is the greater significance of the role played by the time of flight offset, t_{off} , in the CTFM system compared with the MF systems. The second is the fact that the modelled velocity of sound does not agree with that measured experimentally in this chapter. Furthermore, precise modelling requires additional parameters to be determined and second order effects to be included. An alternative to this modelling is a self-calibration procedure, involving precisely known targets at fixed ranges.

A full comparison between the results obtained here and those published in the literature and reviewed in Chapter 2 will be performed in Section 16.2.1.

This concludes the characterisation of the performance of each of the five DOF sub-systems in terms of DOF and echo power measurement performance, and in terms of echo resolution. Now that the DOF sub-systems have been described and characterised, they will be applied to the

measurement of target bearing. This is the subject of Part [IV](#) of this thesis.

13.4.1 Acknowledgements and Contributions

The work described in this section is entirely the work of the author. Significant new analysis techniques, detailed in Appendix [I](#), were applied to fit the theoretical formulae to the experimental results and to validate both theoretical constructs and the operation of the sonar system. Furthermore, rigorous statistical analysis has been applied to the experimental results to obtain precision results which will in turn determine the precision of the range and bearing calculations in Chapter [14](#). This innovation and attention to detail has led to a more precise quantification of the performance of the sonar system than has been previously reported for any of the sonar systems reviewed in Chapter [2](#).

Part IV

Range and Bearing Sub-System

Chapter 14

Range and Bearing Calculation

The goals of this thesis, defined in Section 1.2, require that the sonar system should output the range and bearing of the target. The reliable calculation of the target range utilising the measured DOFs is relatively straightforward. It is detailed in Section 14.1 of this chapter. However, the reliable determination of bearing utilising a small number of receivers is not as straight forward. The literature review of Chapter 2 indicated that the following three methods are available for calculating the bearing of a target:

1. Utilising the peak frequency of an echo transmitted and received by a single transducer (Yata, Section 2.5.4).
2. Utilising the DOF measured by two or more separate receivers to calculate bearing by **Inter-aural Distance Difference (IDD)**, as done by Yata (Section 2.5.5), Kleeman and Kuc (Section 2.5.6), Peremans (Section 2.5.7), Davies (Section 2.5.8) and Kay (Section 2.5.9).
3. Utilising the difference in decibel echo power to calculate the bearing by **Inter-aural Power Difference (IPD)**, as has been done by Davies (Section 2.5.8) and Kay (Section 2.5.9).

The calculation of bearing utilising a **CTFM** signal combined with the first technique is problematic, requiring filtering to be able to separate the signals and then peak detection to be able to estimate the peak frequency. Furthermore, this method has a sign ambiguity in the bearing if the results are considered in the plane, and a cone ambiguity if considered in three dimensions. For these reasons, it will be discounted from further consideration. The remaining two techniques, being the calculation of bearing by **IDD** and by **IPD**, will be studied further in Sections 14.1 and 14.2.

The calculation of bearing assumes that each target generates a unique echo and that this unique echo is detected by both receivers. The problem of isolating the single or unique echo associated with a target in a multi-target environment is called the correspondence problem. This chapter assumes a single-target environment, or equivalently, in the case of a multi-target environment, that the correspondence problem has already been solved. In practice, the solution of the correspondence problem also depends upon the calculation of the target bearing, as will be described in Chapter 15. The circularity of the solution requires that both problems are solved simultaneously by the same code module.

The two techniques of calculating bearing which are explored in this chapter are built upon the **DOF** sub-systems which were developed in Part III of this thesis. The precision of the **DOF**




Figure 14.1: The experimental configuration used to characterise the bearing measurement performance of each of the five sonar systems. The sonar head is rotated with a fixed plane object within the field of view.

and echo power measurements, and their correlations, which were rigorously assessed in Chapter 13, are further utilised in this chapter to calculate the range and bearing precision, for both the **IDD** and **IPD** bearing calculation techniques. The low correlation of 0.47 between **DOF** and echo power residuals, demonstrated in Section 13.2.3, shows that the two methods of calculating bearing can be expected to provide independent results. This information is used in Chapter 15 to gain confidence in the pairing of the echoes between the left and right channels, or to reject a pairing where the bearings calculated by the two techniques would disagree.

In addition to calculating the bearing precision using the **DOF** and echo precision results of Chapter 13, this chapter also experimentally assesses bearing measurement performance in terms of the major criteria of accuracy and precision. Their performance in terms of a minor criterion, the correlation between the two bearing measurements, is also assessed. A total of four performance assessments are carried out:

1. **IDD** bearing accuracy,
2. **IDD** bearing precision,
3. **IPD** bearing precision, and
4. correlation between **IDD** and **IPD** bearing.

These four performance measures are first calculated using the **DOF** and echo power precision and correlation results from Chapter 13. They are also assessed experimentally using the experimental configuration shown in Figure 14.1. Unlike the **DOF** experiments in Chapter 13, there is only one experimental configuration. However, a scheme similar to that documented in Table 13.1 is used to assess the four performance measures on each of the five **DOF** sub-systems using only two data sets. The experimental bearing accuracy, precision and correlation are then compared with the values calculated from the **DOF** and echo power precision and correlation results from Chapter 13.




Figure 14.2: The three types of sonar reflector considered in this section. Note that in the case of the edge reflector shown in (b), unlike the other target types, there is only one path from the transmitter to the target.

Since the **IDD** bearing is computed from fundamental geometrical principles, accuracy and precision can be determined. However the **IPD** bearing requires calibration data or replicated experiments in order to determine its accuracy. Unfortunately neither calibration nor replication data is available and hence only precision can be assessed.

14.1 Inter-Aural Distance Difference

The **inter-aural distance difference** method of determining bearing relies upon the measured **DOFs** and the different paths taken by the sonar signal from the transmitter to the left and right receivers via an object with a unique geometry. To understand this, three common geometries are reviewed and developed. As it is not possible to reliably differentiate between these target types, the effect of ignoring the target type is evaluated, and found to be small. The range and bearing precision are shown to be linked to the **DOF** precision and correlation. The final section (Section 14.1.4) shows the results of a bearing calibration experiment utilising the **IDD** method.

The term **IDD** used in this thesis is similar to the terms **Inter-aural Time Difference (ITD)** and, in the case of a **CTFM** sonar system, **Inter-aural Frequency Difference (IFD)**, which are used by Rowell [120] and others.

14.1.1 Target Types

Peremans [108] and Kleeman and Kuc [82] have identified several types of sonar target which have different geometrical path lengths or different observational properties. Three common targets are:

1. plane reflectors (Figure 14.2a),
2. edge or point reflectors (Figure 14.2b), and
3. corner reflectors (90°) (Figure 14.2c).

In addition to these three basic types, there are also non- 90° corners, concave and convex curved surface reflectors, and surfaces of continuously varying curvature. Some properties of




Figure 14.3: Geometry of a plane reflector, in which the normal of the transmitter and the reflecting plane are not co-aligned. The arcs at T_V have centres at R_1 and R_2 , and radii d_1 and d_2 , being the measured distances of flight. Thus T_V is readily located from the sonar measurements. The outline of the sonar head is shown in grey.

non-90° corners are described in Appendix D.5. Reflection from curved surfaces is treated geometrically by Peremans [108, Chapter 4] and numerically by Zhu [147]. These additional reflector types are not considered further here.

There are two reasons for identifying the target type. The first is that the path geometry is different for each reflector (see Figure 14.2), and therefore the model for computing the range and bearing from the sonar path lengths (DOFs) is different. The second reason is that it provides more information to the user of the sonar system. When performing localisation it is useful to know if a particular feature is a corner or a wall.

Plane Reflectors

Figure 14.3 is an expansion upon Figure 14.2a. The virtual transmitter T_V is constructed by reflecting the transmitter T at the planar reflecting surface (shown in blue). Simplified sonar paths may then be drawn from T_V to the respective receivers R_1 and R_2 . The point half way between R_1 and R_2 is designated the head centre point P_{HC} . The points R_1 , R_2 , P_{HC} and T are fixed by the construction of the sonar head. The normal of the reflecting plane which passes through the transmitter T is designated P_T . The reflecting object, and hence T_V , are free to move in the plane.

The variables shown on the diagram are listed below.

d_1 is the DOF between T_V and R_1 .

d_2 is the DOF between T_V and R_2 .

$2r$	is the distance from T to T_V . Note r is the distance from T to P_T and also from P_T to T_V .
θ	is the bearing of plane target relative to the normal of transmitter T .
$2r'$	is the distance from P_{HC} to T_V .
θ'	is the bearing from P_{HC} to T_V .
d_R	is the half distance between the two receivers R_1 and R_2 . It is assumed that d_R is the same for both receivers, i.e., that the transmitter is central.
d_T	is the distance of T in front of P_{HC} .
d_{TR}	is the geometric distance from the transmitter T to R_1 or R_2 .
ϕ	is the angle $P_{HC}R_1T$, which is equal to the angle $P_{HC}R_2T$. This is half the receiver splay angle.

The polar coordinate system $\{r', \theta'\}$ has its centre at P_{HC} , while the polar coordinate system $\{r, \theta\}$ has its centre at T . The line from P_{HC} to T_V , with bearing θ' , is not normal to the plane target. The only normal in the geometry is the normal of the plane target through the transmitter T . Therefore T is adopted as the origin and $\{r, \theta\}$ as the principal coordinate system.

The forward problem, to compute the DOFs for a target with a given position $\{r, \theta\}$, is solved using the law of cosines and a trigonometric identity. The solutions are

$$d_1 = \sqrt{d_{TR}^2 + 4r^2 - 4d_{TR}r \sin(\theta - \phi)} \quad (14.1)$$

and

$$d_2 = \sqrt{d_{TR}^2 + 4r^2 + 4d_{TR}r \sin(\theta + \phi)}, \quad (14.2)$$

where d_{TR} is given by

$$d_{TR} = \sqrt{d_T^2 + d_R^2} \quad (14.3)$$

and ϕ is the angle $P_{HC}R_1T$, which is given by

$$\phi = \tan^{-1} \left(\frac{d_T}{d_R} \right). \quad (14.4)$$

The reverse problem of calculating the coordinate $\{r, \theta\}$ of the target with respect to T in terms of the measured DOFs $\{d_1, d_2\}$ is solved by first calculating $\{r', \theta'\}$ with respect to P_{HC} . Secondly the coordinates are transformed to $\{r, \theta\}$ with respect to T .

The coordinate $\{r', \theta'\}$ may be found by applying the law of cosines to the triangles $R_1P_{HC}T_V$ and $R_2P_{HC}T_V$, using the angles $90^\circ - \theta'$ and $90^\circ + \theta'$ respectively, and applying a trigonometric identity. The range is

$$r' = \frac{1}{2} \sqrt{\frac{d_1^2 + d_2^2 - 2d_R^2}{2}}, \quad (14.5)$$

while the angle is

$$\theta' = \sin^{-1} \left(\frac{d_2^2 - d_1^2}{8d_R r'} \right). \quad (14.6)$$




Figure 14.4: Geometry for an edge or point reflector, shown in blue. The outline of the sonar head is drawn in grey.

The range and bearing $\{r, \theta\}$ to T_V from T may be found by applying the law of cosines to the triangle $P_{HC}TT_V$. The first application uses the angle θ' , while the second uses the angle $180^\circ - \theta$. This procedure yields

$$r = \frac{1}{2} \sqrt{d_T^2 + (2r')^2 - 4d_T r' \cos \theta'} \quad (14.7)$$

and

$$\theta = \cos^{-1} \left(\frac{(2r')^2 - d_T^2 - (2r)^2}{4d_T r} \right). \quad (14.8)$$

Note that (14.8) will not return sign information, because cosine is an even function. The sign can be obtained from (14.6) where the odd function sin is used instead. The sign of the bearing may also be obtained from the **DOFs** by $\text{sign}(d_2 - d_1)$.

The orientation of the plane reflector can be determined from the reported coordinate $\{r, \theta\}$. It is known that the reflection is specular, so the range line TT_V will be perpendicular to the plane.

The bearing computed by the method outlined in this section will be referred to in the remainder of this thesis as the **IDD** bearing θ_{IDD} .

Edge Reflectors

The edge or point reflector is shown in the expanded representation of Figure 14.4. The virtual source construction of the plane is not appropriate as the reflections from the planes shown in Figure 14.4 will not strike the receivers. Rather, the construction is based upon diffraction principles and the re-radiation of the diffracted signal.

Referring to the diagram, the line P_1P_2 is an extension of the transmitter surface. The variables are listed below.

r_1	is the distance from R_1 to P_R .
r_2	is the distance from R_2 to P_R .
r	is the range from T to P_R .
θ	is the bearing of the edge reflector P_R , as measured from the normal of the transmitter.
d_R	is the half distance between the two receivers R_1 and R_2 . It is assumed that d_R is the same for both receivers, i.e. that the transmitter is central.
d_T	is the distance of T in front of P_{HC} .
d_{TR}	is the geometric distance from the transmitter T to R_1 or R_2 .
ϕ	is the angle $P_{HC}R_1T$, which is equal to the angles $P_{HC}R_2T$, R_1TP_1 and R_2TP_2 .

The DOFs d_1 and d_2 recorded by the receivers for an edge at a given $\{r, \theta\}$ relative to the transmitter T are

$$d_1 = r + r_1 \quad (14.9)$$

$$d_2 = r + r_2. \quad (14.10)$$

The unknown distances r_1 and r_2 may be found by applying the law of cosines to the triangles R_1TT_V and R_2TT_V , using the angles $90^\circ - \theta + \phi$ and $90^\circ + \theta + \phi$, and simplifying with a trigonometric identity. Thus the solutions to the forward problem are

$$d_1 = r + \sqrt{d_{TR}^2 + r^2 - 2d_{TR}r \sin(\theta - \phi)} \quad (14.11)$$

and

$$d_2 = r + \sqrt{d_{TR}^2 + r^2 + 2d_{TR}r \sin(\theta + \phi)}. \quad (14.12)$$

The position of the reflecting edge, $\{r, \theta\}$, may be found relative to T by applying the law of cosines to the triangles TR_1P_R and TR_2P_R . Solving for $\{r, \theta\}$ in terms of the known DOFs d_1 and d_2 obtains the solution to the reverse problem,

$$r = \frac{-(d_1 + d_2)d_{TR}^2 + d_1^3 + d_2^3 + \cos(2\phi)(d_1 + d_2)(d_1d_2 - d_{TR}^2) - a \sin(2\phi)}{2d_1^2 + 2d_2^2 - d_{TR}^2 + 4d_1d_2 \cos(2\phi) + d_{TR}^2 \cos(4\phi)} \quad (14.13)$$

and

$$\theta = -\sin^{-1} \left[\frac{(d_1 - d_2) \left\{ \left[d_{TR}^2 (d_1 - d_2)^2 - 2d_{TR}^4 + d_1d_2(d_1^2 + d_2^2) \right] \cos(\phi) + a(d_1 + d_2) \sin(\phi) \right\}}{d_{TR} [2d_{TR}^4 + d_1^4 + d_2^4 - 2d_{TR}^2(d_1^2 + d_2^2) + 2(d_{TR}^2 - d_1^2)(d_{TR}^2 - d_2^2) \cos(2\phi)]} \right], \quad (14.14)$$

with the common factor a defined by

$$a = \sqrt{(d_{TR}^2 - d_1^2)(d_{TR}^2 - d_2^2) [2d_{TR}^2 - (d_1 - d_2)^2 + 2d_{TR}^2 \cos(2\phi)]}. \quad (14.15)$$




Figure 14.5: Corner and plane reflectors, with their virtual source constructions. The 90° corner produces the same observable DOFs d_1 and d_2 as a plane which forms a perpendicular bisector of a line joining the transmitter and the virtual source.

These formulae allow the coordinates $\{r, \theta\}$ of an edge reflector to be computed from the observed DOF measurements d_1 and d_2 . The orientation of the edge reflector cannot be determined from this measurement. The orientation of the edge reflector may only be found by observing the wall segments themselves, which would normally occur as the robot moves around during a map building phase.

Ninety Degree Corner

The virtual source construction for a 90° corner and for a plane reflector are shown in Figure 14.5. Both of these reflectors utilise the same virtual source T_V . This may be seen by the following procedure. First, reflect the transmitter T in the extension of the plane W_1 , which allows the virtual source V_{T1} to be obtained. Secondly, reflect V_{T1} in the extension of the plane W_2 , yielding T_V . A similar procedure, reflecting T , first in W_2 and then W_1 , derives V_{T2} and T_V again. The corner reflector has a virtual source at T_V , which is also the position of the virtual source of a plane reflector W_3 . Thus the path lengths which are detected for a corner reflector are the same as those detected for a plane reflector, and the two are indistinguishable without further information.

It should be noted that as the sensor is moved, the plane observed in the place of a corner will appear to pivot about the corner point. Thus it is possible to discriminate a plane from a corner by making observations from multiple positions (Peremans [108, Chapter 6]). Kleeman and Kuc [82] utilise two separate transmitters to carry out the necessary observations without moving the sonar




Figure 14.6: The sonar paths for an edge reflector, drawn together with the pseudo source construction for a plane reflector, having the same DOFs d_1 and d_2 as the edge reflector.

head.

Given that the geometry for a 90° corner is identical to that for a plane, the DOFs d_1 and d_2 which will be observed are given by (14.1) and (14.2). The range and bearing to the corner reflector target may also be calculated using equations (14.5) to (14.8).

The orientation of the corner cannot be determined from the echoes from the corner. If observations of the target from different locations establish that it is a corner, the sensor can be rotated to one side to find the adjoining plane, thus establishing the orientation of the corner.


14.1.2 The Effect of Ignoring the Target Type

This section explores the effect of ignoring the true target type, and instead applying the plane reflector model to planes, corners and edges. This is done by examining the error in the range and bearing estimates which are incurred relative to the true values. As the corner reflector provides DOFs which are identical to the plane reflector, there can be no error caused by treating a corner reflector as a plane. However, the geometry and formulae for an edge reflector are different from those of a plane reflector.


The error analysis proceeds by using (14.11) and (14.12) to compute d_1 and d_2 which would be observed for a physical edge reflector located at the polar coordinate $\{r, \theta\}$. Using these DOFs, a pseudo source corresponding to the plane reflector geometry is constructed, as shown in Figure 14.6. Using the pseudo source, the range and bearing are computed using equations (14.5) to (14.8). The results are then compared to the original $\{r, \theta\}$ specified for the edge reflector. The errors are plotted in Figure 14.7. These plots use $d_R = 15$ mm and $d_T = 2$ mm.

The plots show that the error in bearing is only significant when the target range is less than 0.4 m, and even then the errors are still within the bounds desired, as described in Section 7.1.1. The range error is more serious, being between +0.8 mm and 1.2 mm depending upon the bearing and range to the target. It is to be noted that the range and bearing errors become large for ranges less than 0.1 m.

While these errors are undesirable, the design of this sonar system does not currently permit the classification of targets prior to computing range and bearing. These errors only apply to edge



(a) error in range



(b) error in bearing

Figure 14.7: The error in (a) range and (b) bearing due to considering a point reflector as a plane reflector. The average systematic range error is about 1 mm. The bearing error vanishes as the range is increased, and is small even at close range.

and point reflectors. Corner reflectors will have their position computed correctly by the plane reflector formulae.

It is possible to reduce the magnitude of these range and bearing errors through a careful re-design of the sonar head. The details of the computation are left to Appendix G. However, the range error is essentially reduced to zero.

14.1.3 Range and Bearing Precision

The formulae for computing the polar coordinate tuple, $\{r, \theta\}$ from the observed distances of flight were presented in Section 14.1.1. In particular this section discussed the implication upon the precision of the results when the plane target geometry is used it in the place of point target geometry. Other factors that control the precision of the estimated tuple, especially the bearing parameter, are the precision of the range observations and the correlation that exists between the left and right channels. The following discussion firstly derives an expression for the precision of the bearing computation model, and then looks at how the correlation of range observations affects the precision of range and bearing.

Formulae

The variances and covariances between the range measurements, which were studied in Sections 13.2.1 and 13.2.3, may be assembled into a **Variance CoVariance (VCV)** matrix, Σ_{DOF} . A similar **VCV** matrix, Σ_{pol} , may be constructed to describe the variations in the polar coordinate tuple $\{r, \theta\}$. The contents of the **variance covariance** matrices are defined by

$$\Sigma_{DOF} = \begin{bmatrix} \sigma_{d1d1} & \sigma_{d2d1} \\ \sigma_{d1d2} & \sigma_{d2d2} \end{bmatrix} \text{ and } \Sigma_{pol} = \begin{bmatrix} \sigma_{rr} & \sigma_{r\theta} \\ \sigma_{r\theta} & \sigma_{\theta\theta} \end{bmatrix}. \quad (14.16)$$

The polar **VCV** matrix Σ_{pol} may be calculated in terms of the **DOF VCV** matrix, Σ_{DOF} , using the techniques of Appendix I. The calculation requires a design matrix which is

$$\mathbf{G} = \begin{bmatrix} \frac{\partial r}{\partial d_1} & \frac{\partial r}{\partial d_2} \\ \frac{\partial \theta}{\partial d_1} & \frac{\partial \theta}{\partial d_2} \end{bmatrix}. \quad (14.17)$$

The equations defining r and θ are (14.5), (14.6), (14.7) and (14.8). The computation of the derivatives was performed symbolically using Maple [136], and the forms shown here employ extensive back substitution to obtain simpler forms. The derivatives are

$$\frac{\partial r}{\partial d_1} = \frac{d_1 (2r'd_R C' - d_T d_R - 2r'd_T S')}{16rr'd_R C'} \quad (14.18)$$

$$\frac{\partial r}{\partial d_2} = \frac{d_2 (2r'd_R C' - d_T d_R + 2r'd_T S')}{16rr'd_R C'} \quad (14.19)$$

$$\frac{\partial \theta}{\partial d_1} = \frac{d_1 (2r'd_R C C' - d_T d_R C - 2r'd_T C S' - 2rd_R - 4rr'S')}{16r^2 r'd_R S C'} \quad (14.20)$$

$$\frac{\partial \theta}{\partial d_2} = \frac{d_2 (2r'd_R C C' - d_T d_R C + 2r'd_T C S' - 2rd_R + 4rr'S')}{16r^2 r'd_R S C'}, \quad (14.21)$$

where d_R and d_T are defined in Figure 14.3, and r' can be computed in terms of d_1 and d_2 using (14.5). The remaining variables, C , S , C' and S' , can be defined through trigonometry or in terms of the DOF observations:

$$S' = \sin \theta' = \frac{d_2^2 - d_1^2}{8d_R r'} \quad (14.22)$$

$$C' = \cos \theta' = \sqrt{1 - S'^2} \quad (14.23)$$

$$C = \cos \theta = \frac{4(r'^2 - r^2) - d_T^2}{4d_T r} \quad (14.24)$$

$$S = \sin \theta = \sqrt{1 - C^2}. \quad (14.25)$$

The forms involving square roots have numerical stability problems around $\theta = 0$, making the trigonometric versions preferable. If necessary, the angles θ' and θ , used by the trigonometric versions, can be computed from observed distances of flight using (14.6) and (14.8).

The complexity of the parameters of the design matrix \mathbf{G} prevents Σ_{pol} from being computed algebraically. As the design matrix \mathbf{G} depends upon the coordinates of the target $\{r, \theta\}$, it must be evaluated for each target coordinate used.

There is a singularity in the derivatives of θ when $d_1 = d_2$ ($\theta = 0$), which may be avoided by adding a micrometre to d_2 before computing \mathbf{G} .

Uncorrelated DOF Measurements

Since the polar VCV matrix can only be evaluated numerically, it is convenient to define some simulation values and propagate these through the appropriate models.

Consider a plane target located at $\{d_1, d_2\} = \{1.0028 \text{ m}, 1.0028 \text{ m}\}$ or $\{r, \theta\} = \{0.500 \text{ m}, 0^\circ\}$. The left panel of Figure 14.8 shows a 2D density frequency plot of $\{d_1, d_2\}$ with added uncorrelated Gaussian noise of the form $n(0, 10^{-6})$ where the variance is equivalent to a standard deviation of $\pm 1 \text{ mm}$, and the off-diagonal elements of Σ_{DOF} are zero. This level of precision is equal to that recorded for the CTFM-YW and CTFM-LSMYW systems. The frequency plot is plotted using a grey scale index.

The right panel shows a 2D density frequency plot of $\{r, \theta\}$ which is the transform of $\{d_1, d_2\}$. It is interesting to note that the circular pattern of uncorrelated points in the $\{d_1, d_2\}$ domain is transformed to an arc-like domain in $\{r, \theta\}$. The range precision in the $\{r, \theta\}$ coordinate system is $\sigma_r = \pm 0.35 \text{ mm}$, while the bearing precision is $\sigma_\theta = \pm 3.0^\circ$. Thus the area of the confidence figure is maintained under the transformation from DOF coordinates to range and bearing coordinates.

Correlated DOF Measurements

The sonar paths which produce the measurements d_1 and d_2 are very close together, and any changes in the air column will most probably affect both measurements. Additionally both paths share the same transmitter and the unit is manufactured as a whole. Thus rather than uncorrelated observations predominating, Section 13.2.3 demonstrated that the observations are correlated. The effect of admitting correlated range measurements is demonstrated in Figure 14.9. As the correlation between the DOF measurements increases, the variance of the bearing estimate decreases.

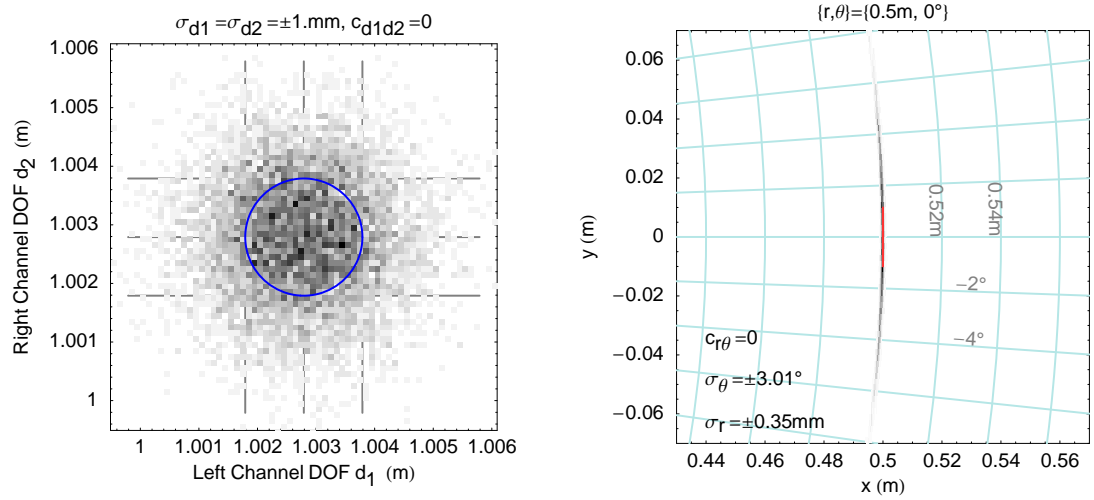


Figure 14.8: The **DOF** measurement shown on the left panel was perturbed by additive, synthetic uncorrelated noise. The right panel shows the effect of the noise upon the precision of the resulting $\{r, \theta\}$ estimate. The position of the target is indicated by a red line on the right panel, having a tangent point at the $\{r, \theta\}$ indicated. The grid in the right panel is marked in polar coordinates, with the sensor located at the origin. The uncorrelated noise in the left panel has a circle of standard deviation, drawn in blue. The lines mark the bounds of this circle. The two plots have different scales to properly show the shape of the noise distribution. When the **DOF** measurements with noise are transformed point by point into polar coordinates (right panel), the distribution of the points is also transformed. Poor bearing precision is obtained for the case of uncorrelated noise as shown here.

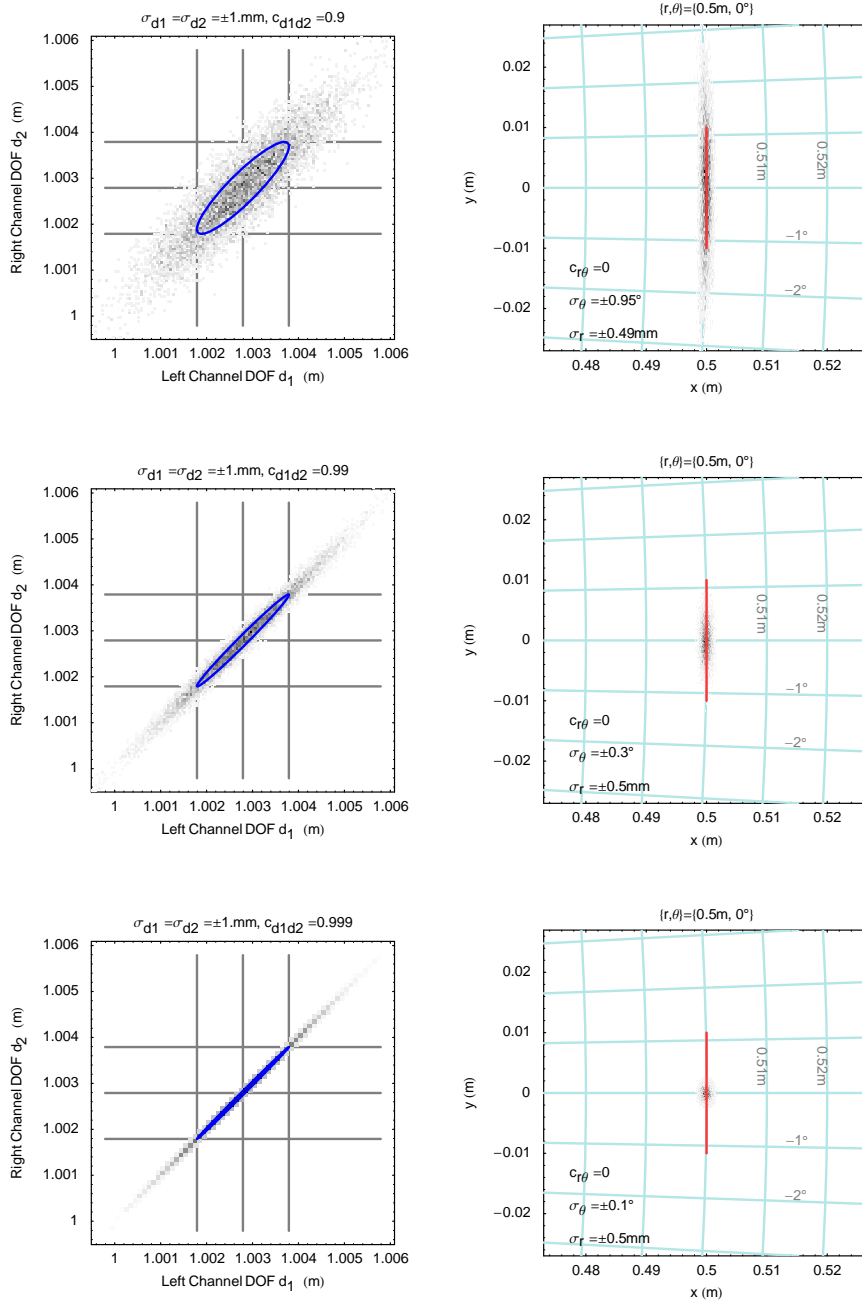


Figure 14.9: A sequence of 2D density frequency plots showing the effect of the correlation coefficient c_{d1d2} upon the standard deviation of the bearing σ_θ . The target, at the coordinates $\{r, \theta\} = \{0.5, 0^\circ\}$, is marked as a red line in the right panel. This sequence demonstrates that the standard deviation σ_θ of the computed bearing is dependent on the correlation coefficient c_{d1d2} between the measured DOFs, and that highly correlated DOF measurements reduce the variance in the bearing measurements. The top row of panels, with $c_{d1d2} = 0.9$, most closely resembles the observed data from Sections 13.2.1 and 13.2.3.

Table 14.1: Range precision, σ_r , as estimated from the DOF precision and correlation coefficient measurements of Sections 13.2.1 and 13.2.3 using the method of Section 14.1.3.

Sonar System	σ_r (mm) Range Precision
MF-QD	0.44
MF-UD	0.43
CTFM-DFT	0.57
CTFM-YW(128)	0.48
CTFM-LSMYW(128,128)	0.49

Highly correlated DOF measurements with a correlation coefficient of $c_{d1d2} = 0.999$ can be seen to give $\sigma_\theta = \pm 0.1^\circ$. This corresponds to a tangential distance of 1 mm at $r = 0.5$ m, which is an order of magnitude decrease or 30 times less than the value for uncorrelated DOF measurements with the same variance (see Figure 14.8). Correlation values encountered in this work vary from 0.79 to 0.89, as may be seen in Section 13.2.3.

The variance of the range measurement marginally increases when the DOF noise is correlated. The standard deviation increases from $\sigma_r = \pm 0.35$ mm to ± 0.5 mm since the error propagation law changes. The benefits of averaging the two measurements are lost.

Predicted Range and Bearing Precision

The DOF standard deviation, σ_d , and the DOF correlation coefficient, c_{d1d2} , which were measured in Sections 13.2.1 and 13.2.3, may be used to calculate the IDD bearing precision, $\sigma_{\theta\text{IDD}}$, as shown in previous sub-sections. The results of doing so are graphed in Figure 14.10. The best bearing precision is indicated by having a point $\{\sigma_d, c_{d1d2}\}$ toward the left or the top of the graph. By examining the graph, the most efficient method of improving bearing resolution may be determined by choosing either to improve the DOF standard deviation σ_d or to increase the correlation coefficient c_{d1d2} so as to cross the most isopleths toward the top or left of the graph.

The range standard deviation, σ_r , is plotted in Figure 14.11 as an isopleth plot. This isopleth plot indicates that increasing the DOF correlation coefficient incurs a small penalty in range precision.

The points marked on the graphs in Figures 14.10 and 14.11 show the expected range and bearing precision, given the DOF precision. The values obtained for the range precision σ_r are shown in Table 14.1. The MF-UD sonar provides the best precision, followed closely by the MF-QD sonar. The two CTFM-AR based systems perform better than the CTFM-DFT system, which may be attributed to the quantisation present in the DFT.

The expected bearing precision will be compared with actual measurements of the bearing precision in Section 14.1.4.

Examination of $c_{r\theta}$ (see Figures 14.8 and 14.9) reveals that there is a small and insignificant covariance between range and bearing. The range-bearing correlation coefficient $c_{r\theta}$ decreases toward zero as the DOF correlation coefficient c_{d1d2} increases. This correlation coefficient occurs in proportion to target bearing θ .




Figure 14.10: Isopleth plot of the bearing standard deviation, σ_θ , as a function of the DOF standard deviation, σ_d , measured in Section 13.2.1, and the DOF correlation, c_{d1d2} , measured in Section 13.2.3, for a target at $\{r, \theta\} = \{0.5\text{m}, 0^\circ\}$. The bearing variance increases toward the bottom right of the plot. The isopleths change for targets at different positions, particularly for different bearings.




Figure 14.11: Isopleth plot of the standard deviation, σ_r , of the target range as a function of the DOF standard deviation, σ_d , measured in Section 13.2.1, and the DOF correlation coefficient, σ_{d1d2} , measured in Section 13.2.3, for a target at $\{r, \theta\} = \{0.5\text{m}, 0^\circ\}$. The range variance increases toward the top right of the plot. The isopleths change for targets at different positions, particularly for different bearings.

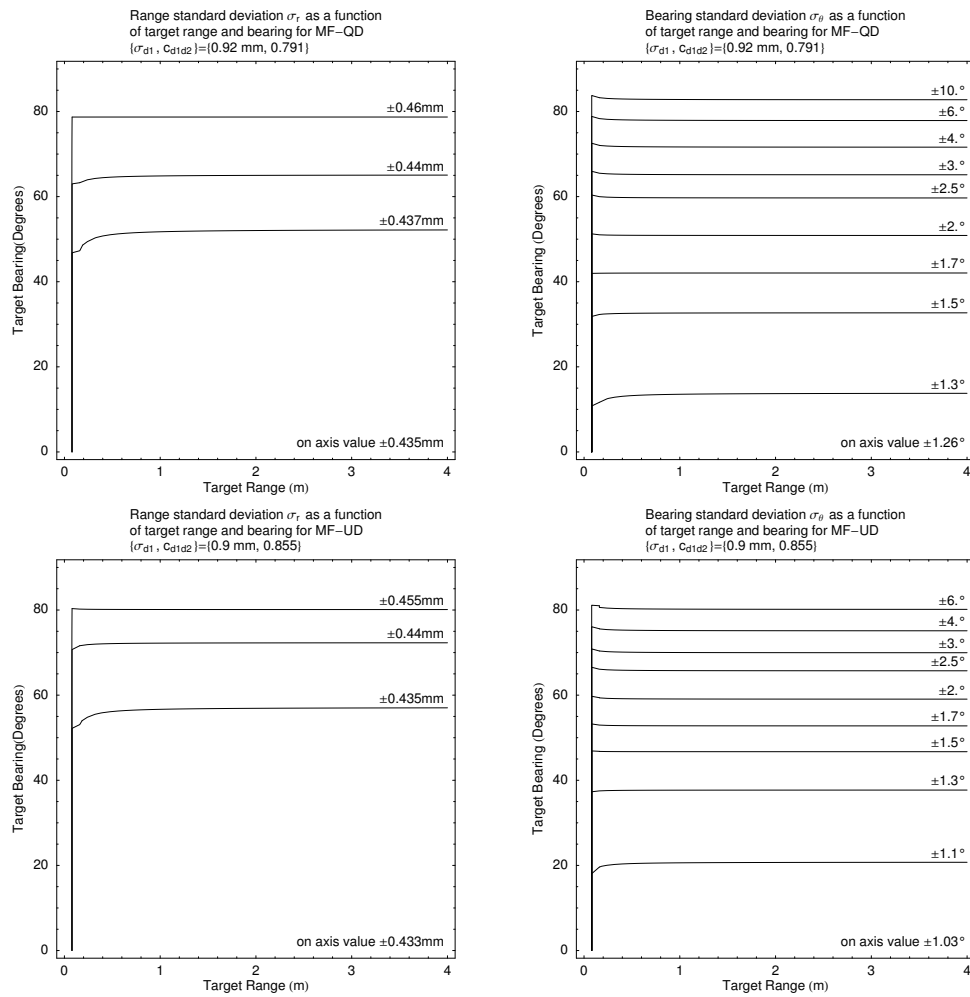


Figure 14.12: Range and bearing precision of the MF class of sonar systems as a function of range and bearing.

The isopleth plots of bearing and range precision in Figures 14.10 and 14.11 only apply at the fixed range and bearing of $\{r, \theta\} = \{0.5\text{ m}, 0^\circ\}$. The plots in Figures 14.12 and 14.13 show the dependence of the range and bearing precision upon the range and bearing values. The plots show that once outside the near field, the range and bearing precision are essentially independent of range. However, there is a strong bearing dependence, particularly as the bearing approaches 90° . The plots are symmetrical for negative bearings, so only positive bearings are shown. The low precision (or high standard deviation) obtained at large bearing angles is of little concern, as the operational bearing range of the sonar is limited by the cone of binaural audition (see Section 7.2.2).

The isopleth plots in Figures 14.10 and 14.11 may be used to provide an engineering specification for the DOF ranging sub-system in terms of the required range and bearing precision, while allowing for the additional uncertainty brought about by off-axis targets as shown in Figures 14.12 and 14.13.

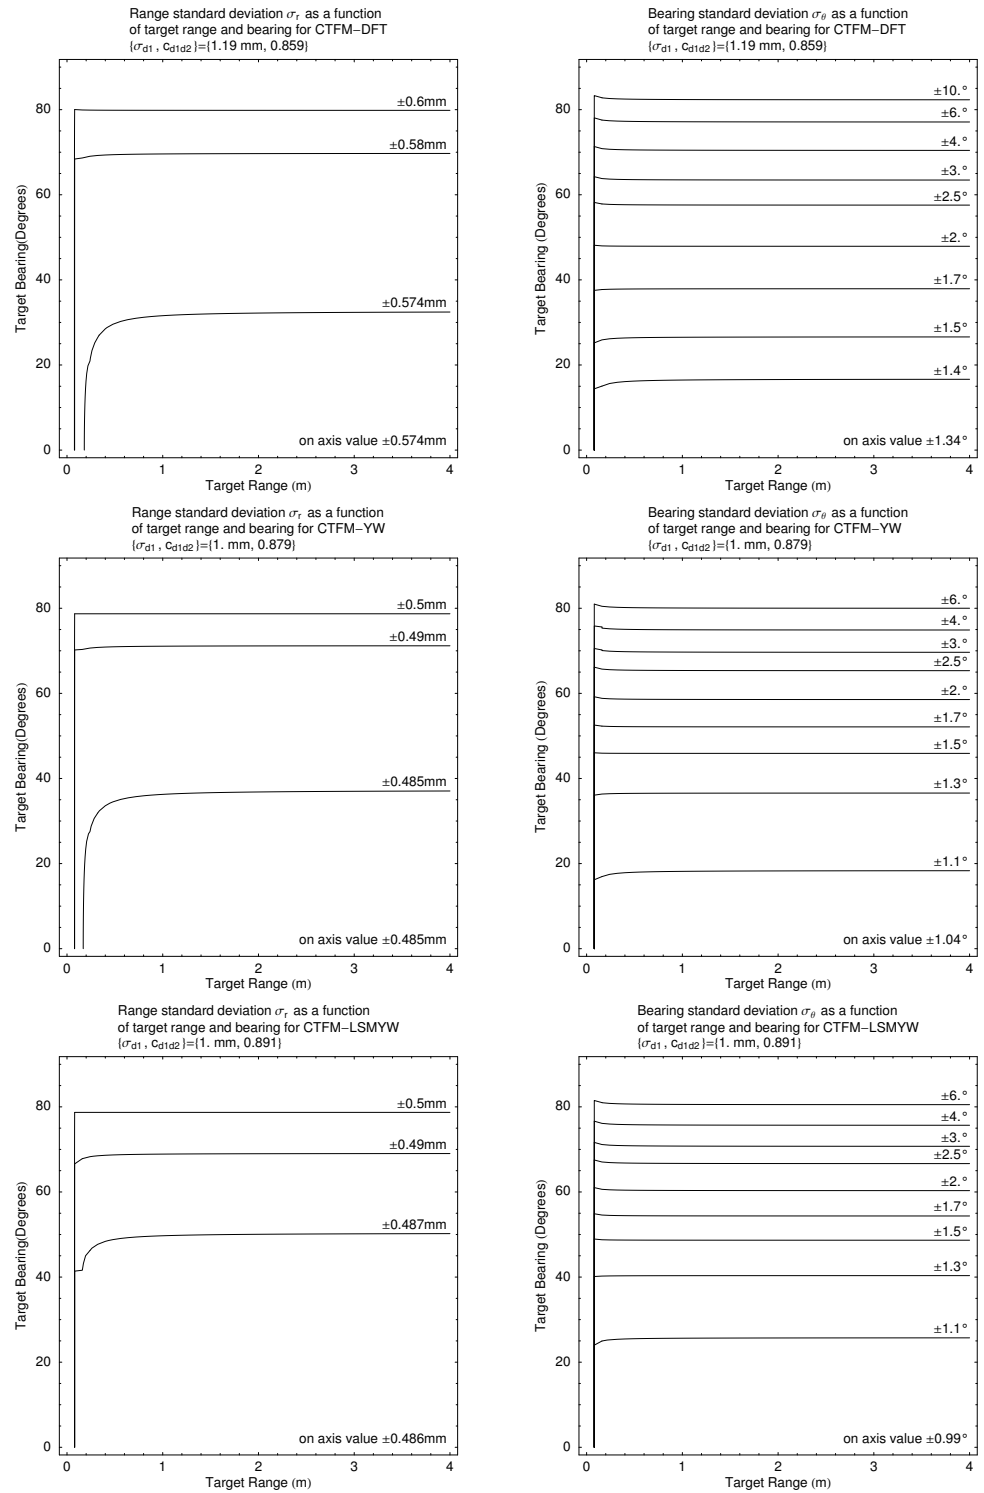


Figure 14.13: Range and bearing precision of the CTFM class of sonar systems as a function of range and bearing.




Figure 14.14: The configuration used in the bearing precision experiment. The range r was kept constant at 0.5 m, while the positioner rotation $-\theta$ was varied between -45° and 45° in steps of 1° .

14.1.4 IDD Bearing Precision Experiment

The experiment described in Figure 14.1 was performed to test the bearing precision of the sonar system. The sensor head (Head 1) was placed upon the movable stage of the precision positioner described in Section 5.2. A glass pane was set up to act as a target, in the same way as shown in Figure 13.2. The transmitter of the sensor was mounted above the point of rotation on the turret of the positioner. The experiment consisted of rotating the sensor so that, in the frame of the sensor, the target appears to rotate in the opposite direction, as shown in Figure 14.14. The sonar head was rotated from -45° to 45° in steps of 1° . At each position 10 observations were taken, as described in Section 13.1.3. The speed of sound was calculated using c_ϕ , as described in Section 12.2. The experiment was performed twice, once for the matched filter sonar class and once for the CTFM sonar class, as described in Chapter 13. The two data sets were processed by the five different types of sonar described in Figure 1.2.

Results

The results are presented in Figure 14.15. Each row of panels represents results from one type of sonar system. The panels in the first column are the result of applying equations (14.5) to (14.8) to the measured DOFs d_1 and d_2 , using the head parameters $d_R = 15$ mm and $d_T = 2$ mm. These experimental data are graphed as dots, while the expected bearing, calculated using the positioner angle, is shown as a red line.

The second column shows the residual in the measured bearing, calculated from

$$\theta_{\text{residual}} = \theta_{\text{measured}} - \theta_{\text{positioner}}.$$

These residuals should be as close to zero as possible. All of the sonar systems display large residuals outside the range $-20^\circ \leq \theta \leq 20^\circ$. Only some of the causes of the large residuals outside the $\pm 20^\circ$ range are understood, and these will be discussed below.

The third column of graphs displays the histograms of θ_{residual} where $-20^\circ \leq \theta_{\text{positioner}} \leq 20^\circ$.

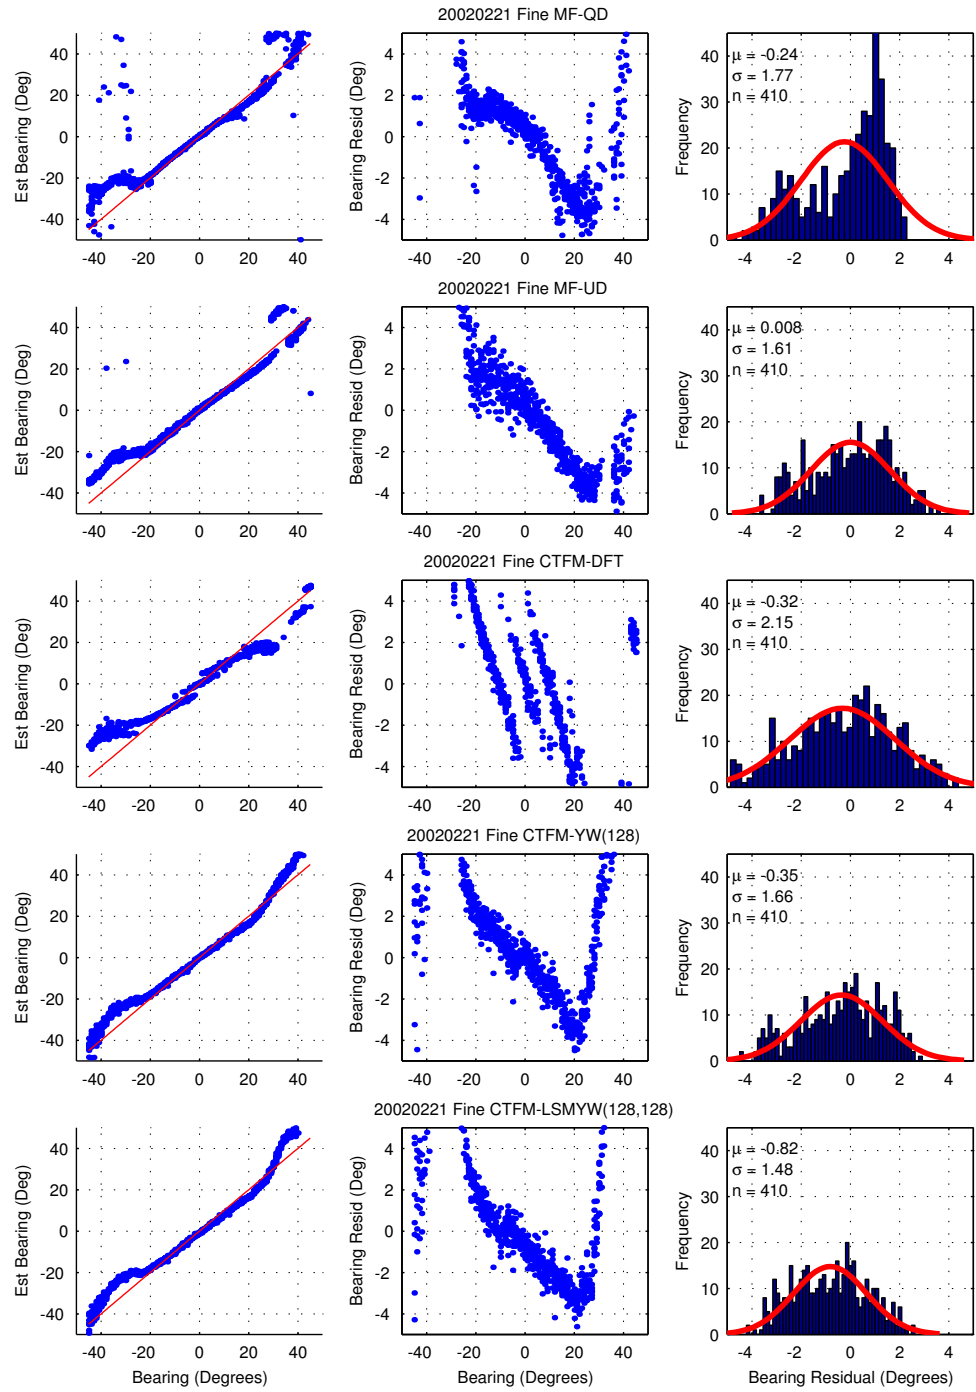


Figure 14.15: Summary of results for bearing calculation by **IDD** with $d_R = 15$ mm and $d_T = 2$ mm. Results from each sonar system are in rows. The first column of panels shows the bearing computed by (14.8) from the sonar measured **DOFs**. The horizontal bearing is determined from the positioner. The red line shows the expected bearing. The second column shows the residuals between the points and the line. The third column shows histograms of the residuals in the second column. The mean and standard deviation are marked. Many of these results show a lack of fit to the normal distribution. This is discussed on p.320. The target was a glass pane at a range of 0.5 m.

A normal distribution is also shown, having mean μ and standard deviation σ calculated from the relevant values of $\theta_{residual}$.

The $\theta_{residual}$ plot shows a large skew. Furthermore, the histograms match poorly with the normal distribution, indicating poor modelling. Close examination of the raw and expected data in the left panels of Figure 14.15 reveals a lack of fit.

The slope of the points in the left column is controlled by the distance, d_R , between the receivers. The transmitter distance, d_T , has minimal effect upon the bearing computation. The receiver distance $d_R = 13.5$ mm was found empirically to provide a good match to the data and minimises the skew of the residuals. At the same time the transmitter distance was changed to $d_T = 2.7$ mm. The data graphs re-computed with these values are shown in Figure 14.16. These data sets have histograms which have excellent agreement with their fitted normal distributions. Furthermore, the residuals are small within the range $-20^\circ \leq \theta \leq 20^\circ$. The range of fitting was restricted to fit inside the field of binaural audition, as discussed in Section 7.2.2. Outside this range, the echoes from the glass pane target may fall beneath the detection threshold (see Section 13.2.4 and Figure 7.8). However, the echo selection algorithm simply selects the echo nearest the expected DOF, so an incorrect echo may be selected. It is simplest to exclude data outside the range $-20^\circ \leq \theta \leq 20^\circ$ from further consideration.

The reason for the mismatch between the measured value of $d_R = 15$ mm and the value which all of the bearing measurement experiments point to, $d_R = 13.5$ mm, is unknown. It seems that the geometrical centre point of the transducer is not the most important point when the transducer is turned. This is consistent with known antenna theory where the electrical centre is not coincident with the mechanical centre.

There is also some curvature in the MF residuals. This cannot be removed by adjusting d_R or d_T . The cause will be discussed in the MF-QD section.

The next five subsections discuss the results of each of the sonar systems in turn.

Matched Filter Quadrature Demodulation

This data set exhibits some residual patterning after d_R has been adjusted. The bowing cannot be removed by adjusting d_R or d_T . The fact that it does not occur in the CTFM results indicates that this effect is unique to the matched filter sonar. The patterning causes a distortion in the histogram plot, which diverges from the normal distribution plotted over it. There is also a significant number of outliers.

This method measures bearing to a precision of $\pm 1.06^\circ$. The small bias which is measured must be subtracted out to remove the bias from the sonar calculated bearing measurements.

Matched Filter Unsynchronised Demodulation

This method has produced a residual histogram which fits much better with the normal distribution than the MF-QD sonar, although the bearing residuals still reveal a similar pattern. This method utilises an interpolated signal to provide a smoother envelope signal for peak finding than that produced by the quadrature demodulation method. The extra precision gained is shown in these results, which achieve a one-sigma bearing precision of $\pm 0.86^\circ$. The measurement is also

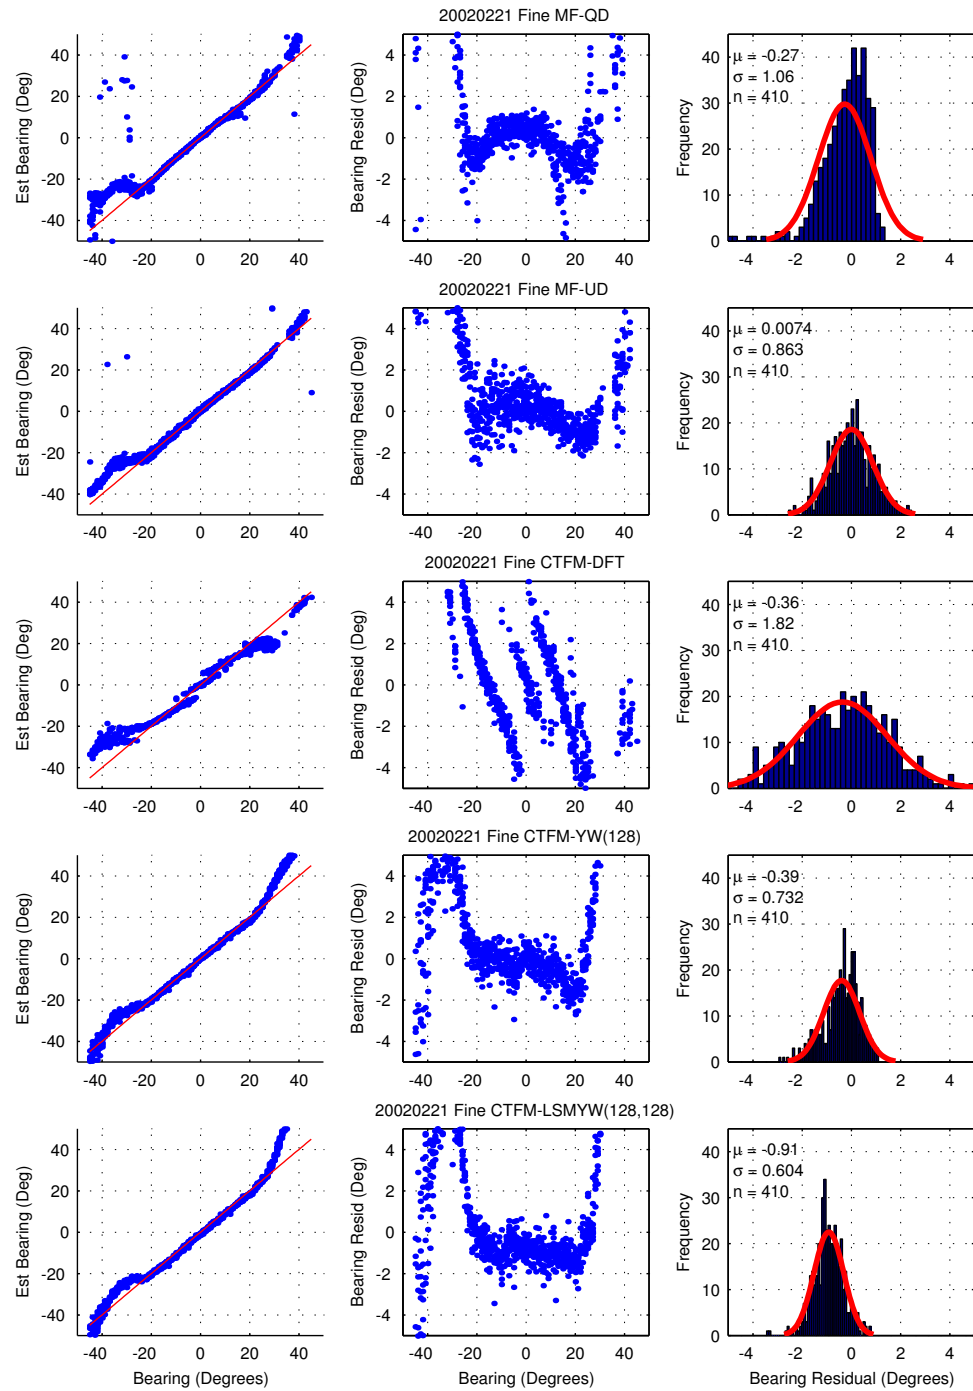


Figure 14.16: Summary of results for bearing calculation by IDD with $d_R = 13.5$ mm and $d_T = 2.7$ mm. Results from each sonar system are in rows. The first column of panels shows the bearing computed by (14.8) from the sonar measured DOFs. The horizontal bearing is determined from the positioner. The red line shows the expected bearing. The second column shows the residuals between the points and the line. The third column shows histograms of the residuals in the second column. The mean and standard deviation are marked. The results shown here conform more closely to the normal distribution than those shown in Figure 14.15. The target was a glass pane at a range of 0.5 m.

unbiased.

The MF-UD method provides a 23% improvement in bearing precision over the MF-QD method. As both methods utilise the same parabolic interpolation scheme to find the peak, the difference indicates that the UD demodulation method is superior to the quadrature demodulation method.

The MF-UD sonar displays the same residual patterning which was seen in the MF-QD results. The cause is unknown.

CTFM-DFT

The DFT results display the tram-track phenomenon observed and explained in the cross correlation results of Section 13.2.3. This phenomenon is due to Davies' [28] sinc interpolation technique malfunctioning when the central peak lines up with a spectral line, as discussed in Section 13.2.3.

The histogram fits the estimated normal distribution well, but the patterning of the errors indicates that this method has serious problems when it comes to calculating bearing by IDD. The precision of the bearing measurement is $\pm 1.82^\circ$, which is the worst of the methods reviewed here.

CTFM-YW

The Yule-Walker spectral estimator has good residuals with only slight patterning within the $-20^\circ \leq \theta \leq 20^\circ$ range, and only a small amount of skew. The bearing precision is $\pm 0.73^\circ$, which is a 61% improvement on the DFT, and an 18% improvement over the MF-UD method.

The YW method has measured a -0.4° bearing offset, which is statistically significant. When the TOF offsets which were measured in Section 13.2.1 are subtracted from the TOF measurements, the bearing offset is not reduced. The precision of the offset times is insufficient to correct the bearing offset. The measured bearing offset could be due to a misalignment of the sonar head before the experiment was performed. The bearing offset is seen consistently in all of the CTFM sonars, but is quite different in the CTFM-LSMYW sonar. The offsets measured by the CTFM-DFT and CTFM-AR sonars statistically agree, which may support this case.

CTFM-LSMYW

The least squares modified Yule-Walker method provides even less patterning in the residuals than the Yule-Walker method. The one-sigma precision is $\pm 0.60^\circ$, which is a 16% improvement upon the YW method. The bearing offset of -0.9° is statistically significant, and is quite different from that provided by the YW and DFT methods. The difference between these results must be attributed to the 128 additional Yule-Walker equations used in the LSMYW estimate. The estimator may be confused by another nearby target perturbing the echo, as demonstrated by the range resolution results in Section 13.3.5.

Bearing Precision Summary

The methods of Section 14.1.3 allow the prediction of the IDD bearing precision of a sonar system based upon its DOF precision and DOF correlation measurements from Sections 13.2.1 and

Table 14.2: Bearing accuracy and precision, $\sigma_{\theta\text{IDD}}$. The first column is taken from the μ values for the histograms shown in Figure 14.16. The second column is calculated using the models of Section 14.1.3 from the DOF precision and correlation data of Sections 13.2.1 and 13.2.3. The third column is taken from the bearing precision experiment from Figure 14.16.

Sonar System	Accuracy Deg	Bearing Precision $\sigma_{\theta\text{IDD}}$ (°)	
		Computed	Observed
MF-QD	-0.27	1.27	1.06
MF-UD	0.01	1.03	0.86
CTFM-DFT	-0.36	1.34	1.82
CTFM-YW(128)	-0.39	1.05	0.73
CTFM-LSMYW(128,128)	-0.91	0.99	0.60

13.2.3. These predictions were plotted as points in Figure 14.10. The experiment of Section 14.1.4 measured the IDD bearing precision directly. A comparison of the two sets of results appears in Table 14.2. The bearing calibration indicates that, for most of the sonar types, the bearing can be measured more precisely than is indicated by the DOF precision experiment. However, it is necessary to take note of the ranges over which the two different measurements were made. The DOF precision experiment was made over the range 0.4m to 1.9m, while the bearing precision experiment was carried out at the fixed range of 0.5m. The possibility that the DOF precision, σ_d , is range dependent cannot be excluded, and therefore the standard deviation of the DOF precision experiment results is larger than appropriate for targets at 0.5m, where the bearing precision experiment was carried out.

The CTFM-DFT result displays the largest variation between the computed and the observed values. It is also the only system in which the observed bearing precision was worse than the computed bearing precision. This is thought to be due to the significant residual structure in the DOF correlation results of Section 13.2.3 and in the residuals of the bearing precision experiment shown in Figure 14.16. The residual structure is not modelled by the statistical theory utilised to calculate the computed bearing precision, and thus leads to the significant difference in the two results.

14.2 Inter-Aural Power Difference

This section demonstrates the method of calculating bearing by using the inter-aural power difference between the left and right receivers, as initially described in Section 7.2.2. Firstly, a simplified model of the beam is created, which describes the power of the received signal as a function of bearing after being detected by the sonar system. Secondly, the difference between these two powers is computed, which produces a function with a linear region which is useful for performing bearing estimation. Thirdly, bearing estimates and calibrations are presented using echo power taken from the same experimental data sets which were used in Section 14.1.4.

The inter-aural power difference method is also known as the Inter-aural Amplitude Difference (IAD) method in Kay [71] and Inter-aural Intensity Difference (IID) method in Kay [74].




Figure 14.17: The variation in echo amplitude with target angle. The continuous and dashed lines represent two Gaussian functions which have been fitted to the left and right channel data (dots). The parameters which are drawn on the plot are for (14.26). The nature of left and right is reversed due to the anti-clockwise positive angle convention. The data was gathered using the configuration of Figure 14.14 in the same experiment as for the IDD calibration.

14.2.1 Target Types

It is assumed that the different types of target mentioned in Section 14.1.1 all reflect an equal amount of sonar energy to the left and right receivers, as they are placed so close together. Therefore the target type is ignored when working with the IPD model.

14.2.2 Beam Power Modelling

The echo power which is detected by the sonar system varies with target angle, but is not directly described by the radiating plane piston models of Sections A.2 and 6.3.2. These models operate in terms of a single frequency, and display nulls in their angular response function. However, when broadband signals such as the Barker coded signal or CTFM signal are used, and the echo is subsequently processed so as to concentrate the energy of an echo into a single peak, the nulls are largely eliminated. This may be seen from experimental data shown in Figure 14.17.

Each channel of the data has been fitted with a Gaussian function (Abramowitz [1, Section 26.2.9]). This procedure was discussed in Section 7.2.1 and is compared to the beam pattern (integrated over the bandwidth of the signal) in Figure 7.6. The Gaussian function provides a good fit in the central area of the beam pattern, but diverges at a point called the point of separation, at about 37° in the case of our transducers.

The Gaussian function is usually used in statistics. The normalisation requirements of statistics

do not apply here, so a more convenient form is adopted. The function used is

$$f_{\text{gaus}}(\theta, a, \theta_0, b) = a \exp \left\{ - \left(\frac{\theta - \theta_0}{b} \right)^2 \right\}, \quad (14.26)$$

where the parameters are:

- θ , the bearing of the target,
- a , the peak value of the fitted echo amplitude,
- θ_0 , the central angle, where the peak echo amplitude occurs, and
- b , the width of the peak. The points of inflexion (change in direction of curvature) occur at $\theta = \theta_0 \pm \sqrt{2}b$, and have an amplitude of a/e .

The values of a , θ_0 and b used to fit (14.26) to the left and right experimental data are displayed on the plot in Figure 14.17. The subscripts L and R are used to denote the channel. Only data points with amplitude greater than 3 V were used to perform the fit. The method used to obtain the fit parameters a , θ_0 and b is described in Appendix H.1, but is also quite similar to that outlined by Kleeman [80].

It can be seen from Figure 14.17 that the Gaussian function is a good fit for the experimental data except at the edges, which validates the use of the Gaussian function.

The Gaussian functions fitted to the left and right channels in Figure 14.17 combine the effects of the two transducers in each signal path. A model which separates the effect of each transducer may be found in Appendix H, but this model was not found to be necessary in this work.

Conversion to Decibels

The variation in the amplitude measurements increases in proportion to the amplitude, and it is impossible to see what happens at small amplitudes. Therefore a log plot is required, as shown in Figure 14.18. The amplitude information is converted into decibels of power using the relationship

$$P = 20 \log_{10}(A), \quad (14.27)$$

where A is amplitude and P is power. Usually a decibel is relative to some reference level, but the adjustment has not been carried out here. From this point on, converted power P_L and P_R are referred to instead of the amplitude A_L and A_R .

When the Gaussian representation is converted to decibels, the width parameter, b , refers to the -8.69 dB point relative to the peak.

The plot in decibels (Figure 14.18) shows the scatter of the data as being approximately constant over the entire data set, and allows closer inspection of the low power regions which were difficult to analyse in the previous plot (Figure 14.17). The fitted curves match the data well in this plot, and it is necessary to perform the curve fitting in the dB space to obtain an accurate fit.


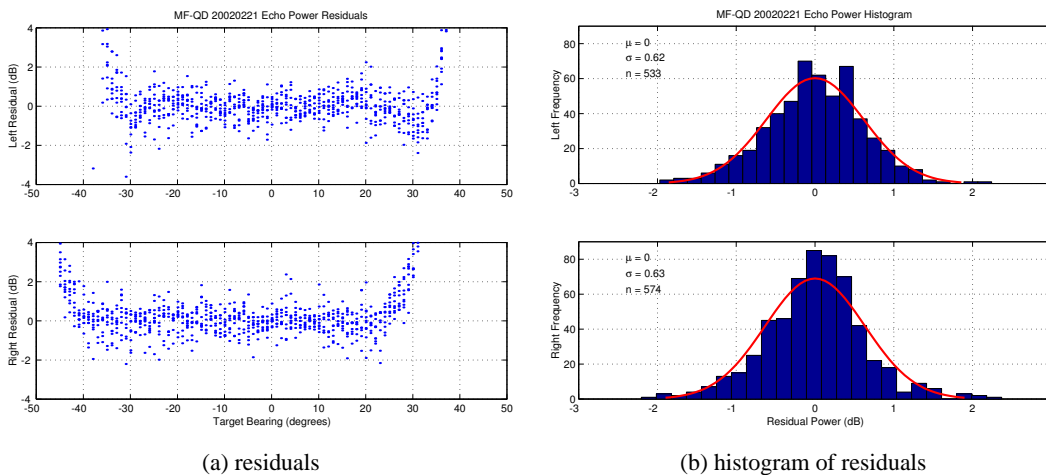


Figure 14.18: The variation in echo power with target angle, in decibels. The fitted Gaussian functions are drawn as dashed and solid lines for the left and right channels, while the experimental data is drawn as dots. This plot contains the same data as is presented in Figure 14.17, except that the vertical scale has been converted to decibels. The data points break away from the fitted curve at the point of separation (see Section 7.2.2). The actual points of separation are different from those predicted in Figure 7.8, due to the individual beam patterns of each of the transducers.



(a) residuals

(b) histogram of residuals

Figure 14.19: Analysis of the power residuals, in decibels, from fitting exponential functions to the received echo power on the left and right channels. (a) The residuals and (b) histogram of residuals with echo power greater than 3 V. The histograms show that the residuals are normally distributed.

Residuals

The residuals between the data points and the fitted curves are shown in Figure 14.19, which also shows a histogram of the residuals which were used for performing the fit. The lack of patterning in the residuals and the agreement between the histogram and the fitted normal distribution demonstrates that the Gaussian function (14.26) with the parameters shown in Figure 14.18 is a good fit for the observed power data.

14.2.3 Log Power Difference

This section shows how to obtain a fitted line, which may subsequently be used to compute the bearing from the observed echo power difference.

Section 7.2.2 suggested that a useful function for estimating bearing using IPD is obtained by subtracting the right power from the left power, in decibels. The subtraction is written

$$P_{diff} = P_R - P_L, \quad (14.28)$$

where P_L and P_R are the power of the echoes from the left and right channels respectively, and P_{diff} is the power difference in dB. This technique is called the **Inter-aural Power Difference (IPD)** method due to this formula. This relationship may also be expressed in terms of the echo amplitudes as

$$P_{diff} = 20 \log_{10} \left(\frac{A_R}{A_L} \right). \quad (14.29)$$

We require a function providing θ in terms of P_{diff} , calculated from (14.28). The function f_{gaus} which was fitted for the left and right channels may be used to compute the power difference, yielding

$$P_{diff} = 20 \log_{10} \left\{ \frac{f_{gaus}(\theta, a_R, \theta_{0R}, b_R)}{f_{gaus}(\theta, a_L, \theta_{0L}, b_L)} \right\}. \quad (14.30)$$

Substituting (14.26) into (14.30) yields

$$P_{diff} = 20 \left(\log_{10} \left[a_R \exp \left\{ - \left(\frac{\theta - \theta_{0R}}{b_R} \right)^2 \right\} \right] - \log_{10} \left[a_L \exp \left\{ - \left(\frac{\theta - \theta_{0L}}{b_L} \right)^2 \right\} \right] \right). \quad (14.31)$$

Re-writing this equation as a polynomial in θ makes it possible to solve for θ in terms of P_{diff} :

$$\begin{aligned} P_{diff} = & -\frac{20}{\ln 10} \cdot \frac{b_L^2 \theta_{0R}^2 - b_R^2 \theta_{0L}^2 + b_L^2 b_R^2 \ln \left(\frac{a_L}{a_R} \right)}{b_L^2 b_R^2} \\ & - \frac{20}{\ln 10} \cdot \frac{2(b_R^2 \theta_{0L} - b_L^2 \theta_{0R})}{b_L^2 b_R^2} \theta - \frac{20}{\ln 10} \cdot \frac{b_L^2 - b_R^2}{b_L^2 b_R^2} \theta^2. \end{aligned} \quad (14.32)$$

Figure 14.20 shows what happens when these techniques are applied to the data described in the previous section. The fitted line is curved. Thus the coefficient of θ^2 in (14.32) is not zero. This can only happen when the widths b_L and b_R are unequal, as they are in this example. When the transducers are perfectly matched it may be expected that these terms should be equal and a straight line fit would emerge.

The residuals and histogram of the fit are shown in Figure 14.21. The residuals show some minor patterning, but the histogram shows that the residuals are normally distributed.




Figure 14.20: The log power difference between the right and left channels, showing the power difference from experimental data calculated using (14.28) and the line calculated according to (14.32). The data breaks away from the fitted line at the left and right because the echo power of one of the channels hits the noise floor at that point. The parameters β_i will be defined in (14.33). The experimental data in this plot may be compared with the reference plot for ideal transducer beam patterns in Figure 7.9. The differences are due to the variations in the beam patterns of the real transducer from the ideal patterns calculated in Section 6.3.2.

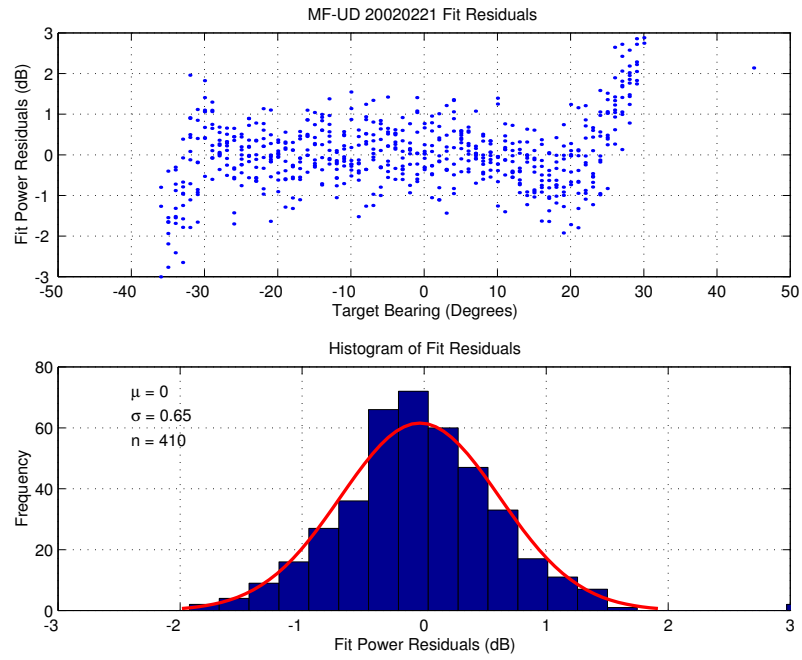


Figure 14.21: The residuals from the fit shown in Figure 14.20. The histogram is only computed for data points in the interval $-20^\circ \leq \theta \leq 20^\circ$. The histogram shows that the residuals are normally distributed.

Direct Fitting

In practice, it is not necessary to fit for the six parameters a_L , θ_{0L} , b_L , a_R , θ_{0R} and b_R . Instead, the 3 polynomial coefficients of powers of θ may be solved for directly, using a least squares fit procedure as outlined in Appendix I. The equation which is fitted is

$$P_R - P_L = \beta_0 + \beta_1 \theta + \beta_2 \theta^2. \quad (14.33)$$

As there are fewer parameters being fitted, this form is more computationally convenient. However, if Gaussian fits to each channel have been made, the Gaussian parameters may be used to calculate the β_i by equating coefficients of θ^n in (14.32) and (14.33).

Bearing Computation

A given pair of echoes, having power P_L and P_R , together with calibrated values for β_0 , β_1 and β_2 , are necessary in order to solve for the target bearing θ . The solution for θ may be found using the quadratic function

$$\theta = \frac{-\beta_1 \pm \sqrt{\beta_1^2 - 4\beta_2(\beta_0 - (P_R - P_L))}}{2\beta_2}. \quad (14.34)$$

Equation (14.34) is numerically unstable if $\beta_2 \approx 0$. This is indeed the case, as β_2 , being equal to the coefficient of θ^2 in (14.32), is computed as the difference between two similar values. Press [114, Section 5.6] advises that in such situations (14.34) should not be used, as subtracting the




Figure 14.22: The three possible root configurations encountered when solving (14.33) for bearing θ by IPD. Roots are shown as red dots. The measured target angles, having bearings between -45° and 45° , occur close to the origin $\theta = 0$ on these graphs. When β_2 is (a) negative, there are two distinct roots. The desired root is closest to $\theta = 0$. When β_2 is (b) zero, the fit is a straight line, and there is a root near $\theta = 0$. When β_2 is (c) positive, there are again two distinct roots. The desired root is again closest to $\theta = 0$.

nearly equal quantities causes a loss of precision. Press advises that it is best to first calculate

$$q = \frac{-1}{2} \left[\beta_1 + \text{sign}(\beta_1) \sqrt{\beta_1^2 - 4\beta_2(\beta_0 - (P_R - P_L))} \right], \quad (14.35)$$

where the function $\text{sign}(x)$ is defined by

$$\text{sign}(x) = \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases}. \quad (14.36)$$

The roots θ_1 and θ_2 are then given by

$$\theta_1 = \frac{q}{\beta_2} \quad (14.37)$$

and

$$\theta_2 = \frac{\beta_0 - (P_R - P_L)}{q}. \quad (14.38)$$

The fitted parabola has two roots, but only one is required. It is necessary to choose the appropriate root. The three situations which may occur are illustrated in Figure 14.22. In all three situations the desired root has the smallest magnitude. When β_2 is small, the root given by (14.37) will have a large magnitude, so the smaller root will always be given by (14.38). Therefore the bearing computed by IPD in terms of the echo amplitudes is

$$\theta_{\text{IPD}} = \frac{2(P_R - P_L - \beta_0)}{\beta_1 + \text{sign}(\beta_1) \sqrt{\beta_1^2 - 4\beta_2(\beta_0 - (P_R - P_L))}}. \quad (14.39)$$

14.2.4 IPD Bearing Precision

The precision of the measured bearing θ_{IPD} may be computed in terms of the precision of the left and right echo power measurements σ_{PL} and σ_{PR} using the methods of Appendix I. The echo

power **VCV** matrix is

$$\mathbf{\Sigma}_{power} = \begin{bmatrix} \sigma_{PLPL} & \sigma_{PLPR} \\ \sigma_{PLPR} & \sigma_{PRPR} \end{bmatrix}, \quad (14.40)$$

and the design matrix is

$$\mathbf{G} = \begin{bmatrix} \frac{\partial \theta_{IPD}}{\partial P_L} & \frac{\partial \theta_{IPD}}{\partial P_R} \end{bmatrix}. \quad (14.41)$$

The two derivatives in the design matrix are

$$\frac{\partial \theta_{IPD}}{\partial P_L} = -\alpha \text{ and } \frac{\partial \theta_{IPD}}{\partial P_R} = \alpha, \quad (14.42)$$

where

$$\alpha = \frac{2(\beta_1 + \text{sign}(\beta_1) \frac{b}{a})}{(\beta_1 + \text{sign}(\beta_1) a)^2}$$

and

$$a = \sqrt{\beta_1^2 - 4\beta_2(\beta_0 - (P_R - P_L))},$$

$$b = \beta_1^2 - 2\beta_2(\beta_0 - (P_R - P_L)).$$

The **IPD** precision is then found to be

$$\sigma_{\theta_{IPD}}^2 = \alpha^2 (\sigma_{PLPL} - 2\sigma_{PLPR} + \sigma_{PRPR}). \quad (14.43)$$

A plot of $\sigma_{\theta_{IPD}}$ appears in Figure 14.23. This figure shows that **IPD** bearing precision is improved by having precise power measurements and a high correlation between the power measurements.

14.2.5 IPD Calibration Experiment

The amplitude values A_L and A_R were taken from the same experimental data as the **IDD** precision experiment described in Section 14.1.4. Each sonar **DOF** sub-system provides its own method of estimating the power of the echo, as described in the relevant chapters. The polynomial fit parameters of (14.33) were found using the least squares techniques of Appendix I. The results appear in Figure 14.24. The first column of plots is equation (14.28), along with the polynomial fit (14.33). The fit only used the data points where the target bearing is in the range $-20^\circ \leq \theta \leq 20^\circ$. The fitted parameters are drawn on each plot.

The second column of plots is the difference between the bearing estimated using (14.38) and the target position derived from the positioner coordinates. This quantity is the fit residual, but has been transformed from decibels into bearing so that the magnitude of the variations may be easily understood. These plots have different units from Figure 14.21.

The third column of plots is the histogram of the residuals in the second column. Only residuals for target bearings in the range $-20^\circ \leq \theta \leq 20^\circ$ are included in the histogram.

The remainder of this section discusses the results of each sonar system appearing as a row in Figure 14.24.




Figure 14.23: An isopleth plot of **IPD** bearing precision as a function of power precision and correlation coefficient. The plot is calculated from (14.43). The points marked use the data from Table 13.5 and Figure 13.13. The CTFM-DFT sonar system is predicted to have the most precise **IPD** bearing measurements. This graph was evaluated for a target at $\{r, \theta\} = \{0.5\text{ m}, 0^\circ\}$.

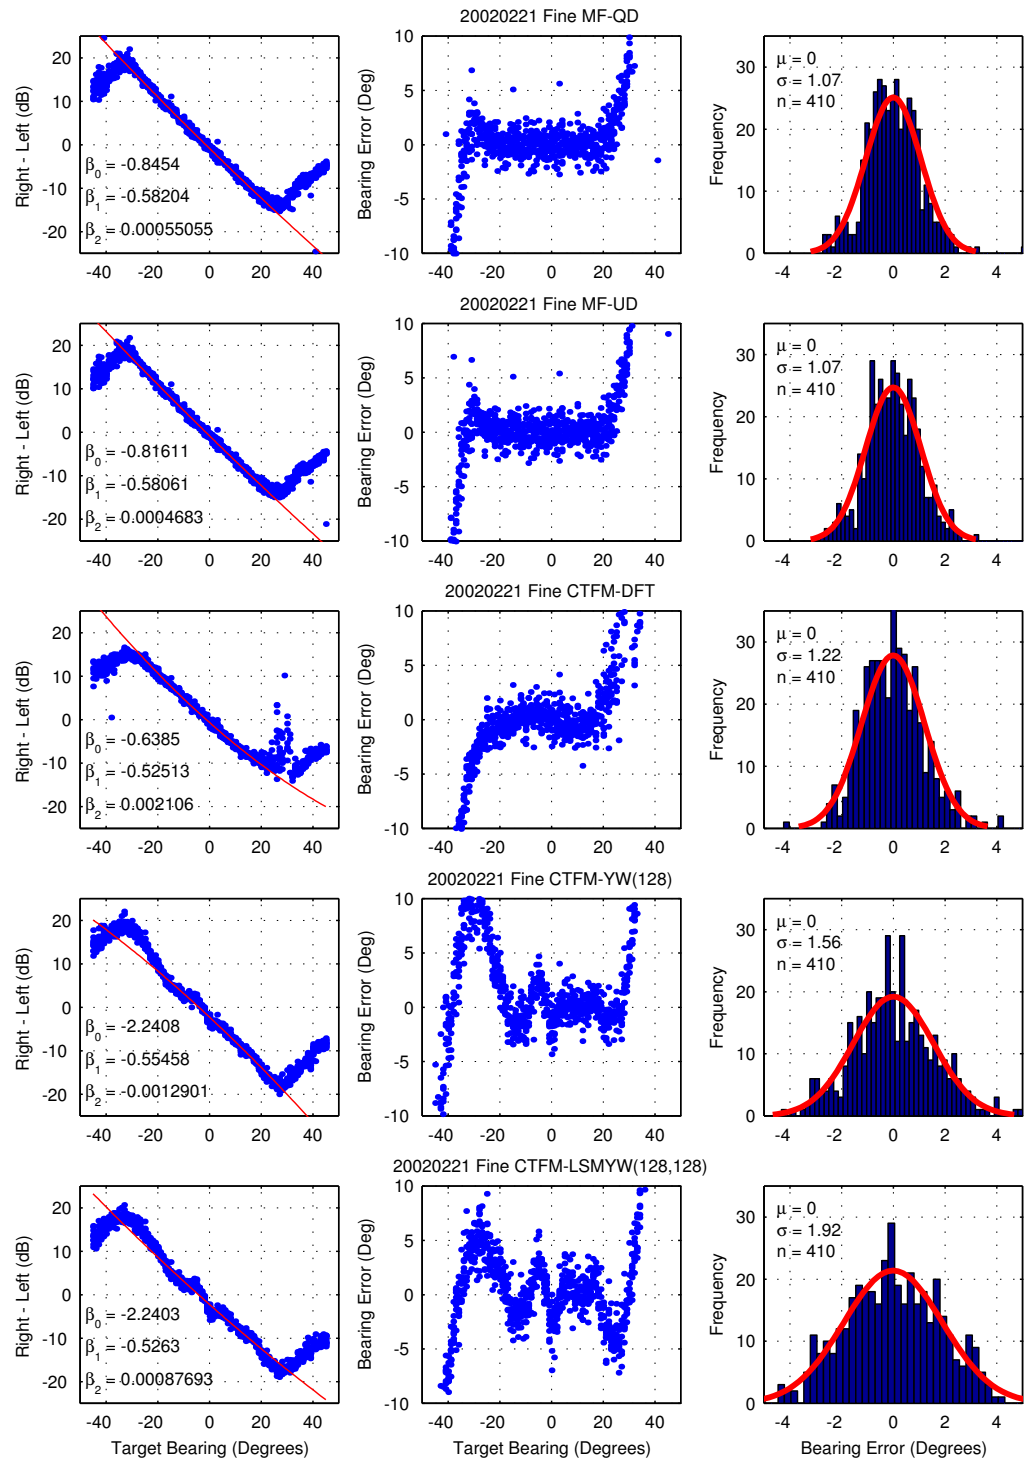


Figure 14.24: Summary of results for IPD bearing calculation. The results for each kind of sonar system are presented as a row. The first column is the measured right minus left channel echo power, plotted as points. The fitted polynomial of (14.33) is shown as a red line, and the values of the fit parameters β_i are indicated. The second column is the residuals formed by subtracting the line from the data points in the first column. The third column contains histograms of the residuals in the second column.

Matched Filter Quadrature Demodulation

The matched filter system with quadrature demodulation has a narrow set of unpatterned residuals within the range $-20^\circ \leq \theta \leq 20^\circ$. The histogram shows that the results are normally distributed and free from bias, and that a one-sigma bearing precision of $\pm 1.1^\circ$ is obtained. This indicates that the quadrature demodulation system combined with a parabolic interpolation scheme as used by this sonar system provides precise amplitude estimates.

The results obtained here contradict those obtained by the power precision experiment (see Table 13.5), but validate the interpretation of the power residuals from that experiment.

Matched Filter Unsynchronised Demodulation

The matched filter system with unsynchronised demodulation provides results which are statistically equivalent to the matched filter system with quadrature demodulation. The polynomial fitted to the IPD curve is equivalent, and the histogram indicates the same normal distribution applies.

The MF-QD and MF-UD sonar systems both utilise the same parabolic interpolation scheme to estimate the position and height of the peak. The two methods have produced almost identical results for IPD bearing estimates.

CTFM-DFT

The CTFM sonar combined with a DFT spectral estimator and a sinc interpolator produces a system with some residual patterning of the residuals. The patterning could be eliminated by introducing a $\beta_3 \theta^3$ term to compensate.

The raw echo powers for the left and right channels of this sonar system is shown in Figure 14.25. The breakdown of the power difference data on the left channel near $\theta = 25^\circ$ is caused by a lack of fit in the data. It is impossible to correct this lack of fit using the Gaussian model of (14.26). The lack of fit at this point does not appear in other data sets, which indicates that the problem is unique to the DFT estimator and the sinc interpolation system. Examination of the echoes detected by this system indicate that a secondary echo is being confused with the echo from the primary target, and the peak picking subroutine is selecting the secondary one.

The bearing estimates by DFT and sinc interpolator appear good within the range $-20^\circ \leq \theta \leq 20^\circ$, being nearly as good as the best estimator, the matched filter system. This is in contrast to the IDD results, where it is apparent that the sinc interpolation has provided poor results which exhibit tram-tracks. These results indicate that the sinc interpolator is effective at estimating signal power but not at estimating frequency (TOF).

The CTFM-DFT sonar system can measure IPD bearing to $\pm 1.2^\circ$ with negligible bias.

CTFM-YW

The CTFM sonar with a Yule-Walker spectral estimator exhibits patterning in its residuals, and these are so large as to be clearly visible in the power difference plots. The underlying cause of these problems is a lack of fit in the echo power detected by each receiver (see Figure 14.26). These disturbances do not appear in the DFT estimate, which indicates that they are a property of




Figure 14.25: Fits for the received echo amplitude for the CTFM-DFT sonar. Note in particular the lack of fit in the right channel data around $\theta = 25^\circ$.




Figure 14.26: Fits for the received echo power for the CTFM-YW sonar. Note in particular the lack of fit toward the left of the left channel data, which causes the very large residuals in that region. Also note the kinks in the amplitude around the peaks, which cause the kinks in the residuals in the central area.

Table 14.3: Results for measurement of **IPD** bearing precision. The first column representing accuracy has not been calculated for reasons which are explained in the text. The second column is calculated using (14.43), utilising the power precision results from Section 13.2.2. The third column summarises the observed bearing precision results, σ , shown in Figure 14.24. The calculated and observed results come from experiments using different range and bearing intervals.

Sonar System	Accuracy Deg	Bearing Precision $\sigma_{\theta\text{IPD}}$ ($^{\circ}$)	
		Calculated	Observed
MF-QD	-	2.59	1.07
MF-UD	-	1.92	1.07
CTFM-DFT	-	1.65	1.22
CTFM-YW(128)	-	2.98	1.56
CTFM-LSMYW(128,128)	-	3.44	1.92

the **Yule-Walker** spectral estimator. The same disturbances appear in data analysed by the Matlab implementation of the Yule Walker spectral estimator. The precise cause of the disturbances is unknown, but it is expected that the experimental data triggers a weakness in the **YW** spectral estimator design.

The large patterning in the residuals leads to the bearing estimate being compromised in quality. The one-sigma bearing precision is $\pm 1.5^{\circ}$.

CTFM-LSMYW

The results for this system appear quite similar to those for the CTFM-YW system, but there are two important differences. The bearing estimation error near -28° has been reduced, but the amplitude of the patterning in the central region has increased. Any differences between this estimator and the **YW** estimator must be due to the extra 128 points which were used in the computation. The **LSMYW** estimator is also known to be poor at frequency resolution (see Section 13.3.5), and is possibly confused by another nearby target. The precision of this estimator is the worst of the five, providing a one-sigma bearing precision of $\pm 1.9^{\circ}$.

Summary of Results

The **IPD** bearing precision results are summarised in Table 14.3. The first column, **IPD** bearing accuracy, was not calculated as an accuracy measurement must be carried out separately from a calibration procedure, which was performed here. Measurement of accuracy requires a second data set which is not available. The second column of Table 14.3 is calculated using (14.43) applied to the precision data in Table 13.5 and the correlation data in Figure 13.13. The third column summarises the precision results observed in the **IPD** calibration experiment shown in Figure 14.24. The two columns of precision results shown in the table are surprisingly different, but there is only limited scope for comparison due to the different methods used to estimate each uncertainty.

The large apparent difference in the MF-QD bearing precision is attributed to the incorrect

power precision being calculated in the power precision experiment in Section 13.7.

IPD Conclusion

The most precise bearing measurements by IPD are given by the matched filter systems. Both matched filter systems have consistently smaller residuals resulting in narrower histograms. The YW and LSMYW estimators both suffer from glitches in the detected echo power, the causes of which are unknown. As all three CTFM experiments are carried out using the same data set, the variation between these three systems must be attributed to the ability of the spectral estimators to extract the echo power. The DFT results have suffered by selecting the wrong echo in some parts.

The calibration procedure outlined above cannot be used to assess the accuracy of bearings measured by IPD. The calibration results shown in Figure 14.24 show that the mean is zero for all of the sensor systems. This is caused by the curve fitting procedure carried out in the process of calibration. Whereas the IDD bearing measurement process relies on the geometric properties of the sonar head, the IPD method relies upon calibration from experimental data. Therefore, to assess the IPD bearing accuracy, a separate experimental data set must be obtained, and the mean error computed.

14.2.6 Targets Lying Outside the Horizontal Plane

When a target does not lie in the horizontal plane of the sonar sensor and is displaced vertically, the angles subtended by the sonar path to the normals of the transmitter and the receivers will no longer obey the previously assumed relationship. The IPD calibration curve will change depending upon the vertical angle submitted to the sensor. These effects have not been considered in this work, but will need to be considered when the system is extended to full 3D.

This problem does not apply to bearings estimated by IDD.

14.3 Correlation Between IDD and IPD

The two methods of measuring bearing, IDD and IPD, are considered to be independent and therefore uncorrelated. The uncorrelated or correlated nature of these bearing measurements is important to the correspondence problem of Chapter 15.

The level of correlation present between the IDD and IPD estimates can be readily evaluated using the same methodology used in the earlier simulation studies, that is, to look at a scatter plot of the estimates. Such a scatter plot is shown in Figure 14.27.

The correlation coefficient, $c_{\theta_{IDD}\theta_{IPD}}$, is marked as corr on each plot. All of the sonar types shown in this figure display a lack of correlation. This result is used when combining the bearing estimates in Section 15.4.

14.4 Conclusion

Other sonar systems have required the target type to be determined before bearing can be computed, due to the different reflection geometry of each target type. Section 14.1.2 has demonstrated




Figure 14.27: Plots of the residuals of θ_{IDD} and θ_{IPD} , showing the correlation between the two bearing estimates. The correlation coefficient, c_{IDDIPD} , is marked on each plot. In all cases it is considered to be small and insignificant.

that the target type need not be identified before calculating range and bearing, if some small errors can be tolerated. Higher precision may be obtained if the design of the sonar head is changed to place the transmitter in line with the two receivers, as described in Appendix G.

Two different methods were developed to measure the bearing to targets, the [inter-aural distance difference](#) method and the [inter-aural power difference](#) method, found in Sections 14.1 and 14.2 respectively. These two methods rely on statistically independent properties of the detected echo (Section 13.2.3), and the resulting computed bearing estimates are also statistically independent (Section 14.3). The precision of the range and both methods of bearing determination have also been assessed. In the case of the target range, this has only been assessed through calculation from the [DOF](#) results, as shown in Table 14.1. The target bearing has been assessed both computationally in terms of the [DOF](#) and echo power precisions measured in Chapter 13, and also by direct experiment. The results of these two independent determinations are shown in Tables 14.2 and 14.3. The computed precision results were obtained by rigorous application of the error analysis techniques outlined in Appendix I.

The calculated bearing precision results for the CTFM-DFT sonar system were the only results which were more precise than the associated experimental results. The large difference between the calculated and the measured precision of this system is attributed to the breakdown of the statistical assumption that the bearing residuals are normally distributed. The other four systems, having normally distributed residuals for both the [DOF](#) residuals and the experimental bearing residuals, agree more closely between the calculated and experimental results, as shown in Tables 14.2 and 14.3. There remains some disagreement, which is attributed to a slightly larger value for the measured [DOF](#) standard deviation, which appears to increase slightly with increasing target range.

The geometry of the sonar head was carefully measured and used to locate the physical centres of the transducers. However, when these measurements were used in the calculation of bearing by [IDD](#), poor results were obtained, as shown in Figure 14.15. The model distance between the receivers was subsequently adjusted to obtain an improved fit to the experimental data, as shown in Figure 14.16. Even so, the geometry utilised by the [IDD](#) bearing calculation still assumes that the sonar head is perfectly symmetrical.

The calculation of bearing by [IPD](#) makes fewer assumptions than the [IDD](#) method, as it is calibrated to the experimental data with less modelling. The only assumption made in this case is that the Gaussian fit is valid in the central region of the beam pattern. This assumption was validated in Section 7.2.2 by comparison to the theoretical beam patterns, and also within Section 14.2.2 by the good agreement with experimental data over the central beam region.

While the method of calculating bearing by [IPD](#) has been applied previously, this work has decomposed the beam pattern into separate Gaussian curves and subsequently fitted them to the experimental data. This method takes into account differences in the beam patterns of the left and right signal paths, incorporating the beam patterns of two of the three transducers. The result of this fitting process is the curved [IPD](#) curve as shown in Figure 14.20. Previous experimenters have simply fitted a straight line to this data. Thus the separate and rigorous fitting of the Gaussian curves is seen as a new contribution.

All five of the sonar systems have achieved a $\pm 20^\circ$ field of binaural audition, providing the

wide sensing area as designed in Chapter 7.

When the five different sonar systems are compared by their ability to measure bearing by **IDD** and **IPD**, the most precise measurement is obtained by the **CTFM least squares modified Yule-Walker** system in the **IDD** mode. All of the sonars produce more precise **IDD** measurements than **IPD** measurements, except for the CTFM-DFT sonar system. The **IDD** precision of the CTFM-DFT system is the second worst over all, which may be linked to the problem (see Section 13.2.3) with the complex interpolation scheme (see Section 10.3.4). The **CTFM Yule-Walker** based methods produce excellent **IDD** precision, but poor **IPD** precision. Thus no single sonar system comes out best on both **IDD** and **IPD**.

A full comparison between the results shown here and comparable results reported by other researchers in the literature will be performed in Section 16.2.2.

14.4.1 Acknowledgements and Contributions

This chapter describes several contributions to knowledge made by this author, including:

1. The rigorous propagation of error analysis through the range and bearing calculations, by both the **IDD** and **IPD** techniques.
2. The rigorous experimental assessment of the bearing measurement errors of five different types of sonar system.
3. The recommendation to change the design of the sonar head so that all of the transducers are mounted in-line, as described in Appendix G, to reduce and eliminate the error due to not classifying the target time.
4. The range and bearing isopleth plots shown in Figures 14.10 and 14.11, which may be used to obtain an engineering requirement for the **DOF** precision in terms of a specified range and bearing precision.
5. The fitting of the **IPD** curve by separate Gaussian fits to the amplitude profile of each echo channel, which is seen as an extension of previous work and therefore a contribution.

The assistance of Dr. Ian Lisle in simplifying and factoring the formulae of the range and **IDD** bearing error analysis and the Gaussian curve combination theory of Appendix H.3 is greatly appreciated.

Chapter 15

Correspondence Problem

The correspondence problem is concerned with choosing an echo, in say the left receiver channel, and the corresponding echo from the same target in another channel, typically the right receiver channel. This problem, which is illustrated in Figure 15.1 for three plane targets, must be precisely and reliably solved in order to estimate bearing.

The first criterion in limiting the choice of echoes is the DOF window. The DOF window restricts the choice of echoes to those which occur closer together than the distance between the receivers, $2d_R$. This restriction alone vastly simplifies the problem.

When more than one echo falls within the DOF window, we must find reason to prefer a pairing with one echo over another. The two separate methods of calculating bearing, which were established in Chapter 14, are used to solve this problem.

When an echo pairing has been identified, the two bearing results may be combined to obtain a more precise bearing estimate. Methods to detect incorrect pairings and the consequences of selecting an incorrect pairing are also discussed.

There are two steps to solving the correspondence problem. The first is to find all echo pairings between the left and right channels which satisfy the DOF window test. The second step is to reject those pairings which do not correspond to real targets, using both methods of computing bearing described in Chapter 14. There are two methods of rejecting bad pairings. They are called the simple correspondence solver and the statistical correspondence solver. A road-map of the process is shown in Figure 15.2.

Solutions to the correspondence problem were tested using synthetic echoes derived from both a geometrical model and experimental data. The simulated beam pattern and the additive random noise, for both DOF and echo power, as well as resolution effects, were modelled upon the experimental results of the MF-UD sonar described in Chapters 13 and 14. The effects of noise and resolution were introduced and evaluated separately, allowing the properties of the correspondence solver to be investigated independently of echo resolution. Four simulated experiments were performed using the geometrical model. The simulated experiments were evaluated by their ability to correctly identify targets, which relies upon accurate echo detection and a robust solution to the correspondence problem. The simulated experiments show that the region where correct target identification is impossible is determined by the ability of the DOF sub-system to resolve two targets. By comparison, the region where the correspondence problem is insoluble is much smaller. Thus improvement of the ability of the sonar system to correctly identify targets requires

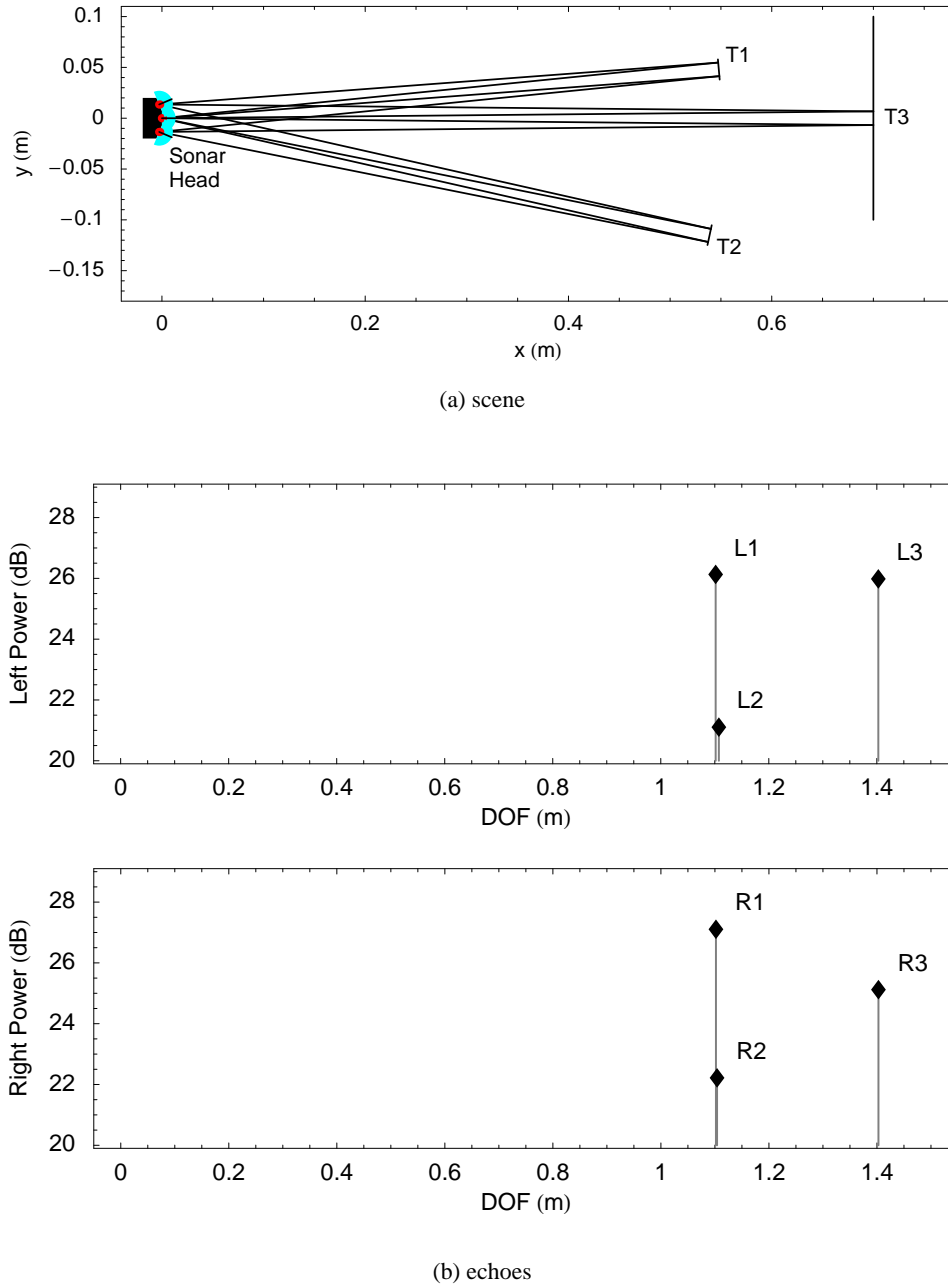


Figure 15.1: (a) A scene with three plane targets, showing the sonar paths. Target T1 is located at $\{r, \theta\} = \{0.550\text{m}, 5^\circ\}$, target T2 at $\{0.551\text{m}, -12^\circ\}$ and target T3 at $\{0.700\text{m}, 0^\circ\}$. (b) The echoes from the scene. The horizontal scale represents distance of flight, and is approximately double the range. The tags identify the echoes, in order of their **DOF**, not their corresponding target. The vertical scale shows the detected power of each echo. The letters L and R identify the receiver on which the echo was detected. The correspondence problem is to identify the target which produces each echo. The echoes L3 and R3 obviously come from target T3, but the echoes from targets T1 and T2 are harder to identify.

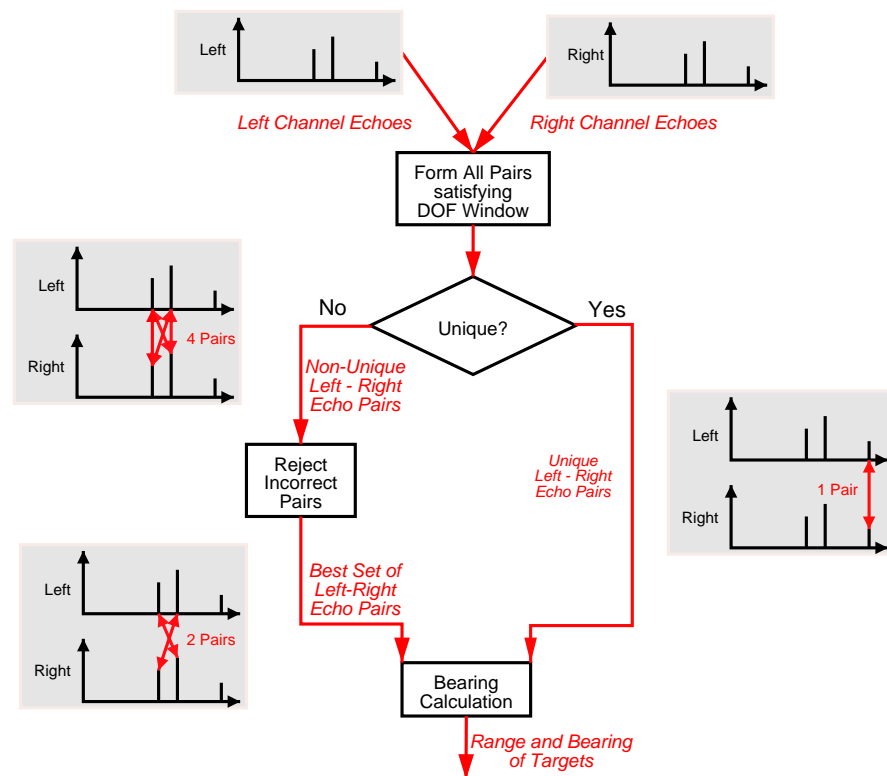


Figure 15.2: Road-map of the solution of the correspondence problem. The first stage forms all echo pairs which satisfy the **DOF** window. The pairs formed by this stage are then divided into unique pairs, and groups of duplicate pairs which share a single echo between more than one pairing. These duplicate pairs require further processing to reject the incorrect pairings. Some incorrect pairs may not be rejected, leaving a duplicate pair relationship in the output. The pairs are subsequently used to compute the range and bearing of targets, using the techniques of Chapter 14.




Figure 15.3: The difference in **DOF** between right and left channels, from (15.1). The plot used a transmitter-receiver spacing of $d_{TR} = 13.7\text{mm}$ and a transmitter-receiver angle of $\phi = 11.3^\circ$ (both defined in Section 14.1.1), and a target bearing of $\theta = 20^\circ$. The difference $d_2 - d_1$ is essentially constant at working ranges.

improvement of the range resolution, thereby lowering the region over which the target identification fails. The range resolution, in turn, is determined by the **time-bandwidth product** properties of the signal, as discussed in Section 4.6.

15.1 DOF Window

The construction of the sonar head makes it impossible for the echoes from a particular target, as detected by the left and right receivers, to be separated by more than the distance between the receivers. This section utilises this information to formally define a search window, from which the echo pair can be found. The method used is due to Kleeman [82].

From (14.1) and (14.2), the difference between the right and left **DOFs** is

$$d_2 - d_1 = \sqrt{d_{TR}^2 + 4r^2 + 4d_{TR}r \sin(\theta + \phi)} - \sqrt{d_{TR}^2 + 4r^2 - 4d_{TR}r \sin(\theta - \phi)}. \quad (15.1)$$

A graph of this function is shown in Figure 15.3. The graph indicates that the dependence upon range r is insignificant, allowing simplification. Taking the limit of (15.1) as $r \rightarrow \infty$ and simplifying using (14.3) yields

$$d_{sep} = \lim_{r \rightarrow \infty} d_2 - d_1 = 2d_R \sin \theta. \quad (15.2)$$

This equation is graphed in Figure 15.4. This provides a useful search window, as no real target can produce echoes separated by more than $d_{sep} = 2d_R = 27\text{mm}$. When the maximum usable bearing of the sonar is limited, this separation distance may be reduced by evaluating (15.2) directly with the maximum usable bearing θ of the sonar.

Clearly (15.2) is odd for negative angles, so for an echo detected by the left channel at d_1 , the corresponding echo on the right channel at d_2 must be contained within the interval

$$d_1 - d_{sep} \leq d_2 \leq d_1 + d_{sep}. \quad (15.3)$$




Figure 15.4: The difference in [DOF](#) between right and left channels, from (15.2), plotted for $d_R = 13.5\text{mm}$.

Thus we write d_{sep} with a \pm symbol. A diagram showing the search window, as applied to the previous example, is shown in Figure 15.5. Chapter 14 established that all five of the sonar systems are able to determine bearing over the range $\pm 20^\circ$, so we shall use a [DOF](#) window of $\pm 9.2\text{mm}$.

15.2 All Pairings

The [DOF](#) window is 18.4mm wide, which is wider than the [DOF](#) resolution of 12mm of several of the sonar systems (obtained by doubling the range resolution listed in Table 13.10). Thus it is possible that more than one echo may be detected within the [DOF](#) window, and hence a correspondence solver will be required. The correspondence solver must decide which way to pair the echoes from the two channels.

Returning to the previous example shown in Figure 15.5, it is seen that there are four possible pairings of the echoes, $\{\{L1, R1\}, \{L1, R2\}, \{L2, R1\}, \{L2, R2\}\}$, for only two targets. That is only two pairs are associated with targets. Two methods of solving this problem are discussed in Section 15.3. The targets described by all of the echo pairs satisfying the [DOF](#) window constraint are shown in Figure 15.6. The rejection of the two fictitious targets is the work of the correspondence solver.

15.3 Correspondence Solvers

This section describes two methods of identifying and rejecting incorrect pairs. The first one is the simple method, while the second is based upon rigorous statistical methods.

15.3.1 Simple Method

Chapter 14 described two different methods of calculating bearing from a single pair of echoes. These methods were the [IDD](#) and [IPD](#) methods. It is possible to apply both of these methods to all pairings of the echoes. The results for the duplicate pairs discussed in Section 15.2 are shown in




Figure 15.5: Two **DOF** windows for different beam widths, centred at echo L1, for the echoes from the targets T1 and T2 in Figure 15.1. Note that the window is applied to the channel not including L1. As the window is centred upon the echo L1, it may be paired with either R1 or R2, which both fall inside the window. The sonar systems described in this thesis have a maximum usable bearing of $\pm 20^\circ$, so the ± 9.2 mm **DOF** window applies. The echo L2 requires a separate application of the **DOF** window (not shown). When this is performed, it is shown that L2 may also be paired with either R1 or R2.




Figure 15.6: Targets found using good and bad echo pairs. Target T3, whose echo pairing $\{L3, R3\}$ is uniquely isolated by the **DOF** window, is correctly identified (and coloured black). The echoes from targets T1 and T2 are not uniquely isolated by the **DOF** window. The blue pairings produce targets corresponding to T1 and T2, while the red pairings are fictitious targets. The red and blue targets together constitute a set of duplicate pairs.

Table 15.1: Calculation of bearing by **IDD** and **IPD** for all possible pairings of echoes falling within the **DOF** window.

Method	Pairing			
	$\{L1, R1\}$	$\{L1, R2\}$	$\{L2, R1\}$	$\{L2, R2\}$
θ_{IDD}	0.70°	5.00°	-12.05°	-7.65°
θ_{IPD}	-3.15°	5.02°	-11.94°	-3.38°
Difference $ \theta_{\text{IDD}} - \theta_{\text{IPD}} $	3.81°	0.02°	0.11°	4.27°

Algorithm 15.1 Simple method of solving the correspondence problem.

Require: pairList

Require: bearingDiscrepancyList calculated according to pairList

Ensure: no echo will be used in more than one pair

SortPairsByAscendingBearingDiscrepancy(pairList, bearingDiscrepancyList)

while ContainsDuplicates(pairList) **do**

RemoveLastPair(pairList)

end while

Table 15.1. This indicates that the bearings calculated by **IDD** and **IPD** have the best agreement in the case of the second and third pairings, i.e. $\{L1, R2\}$ and $\{L2, R1\}$, while $\{L1, R1\}$ and $\{L2, R2\}$ are a poorer match. An algorithm which exploits this to solve the correspondence problem is given in Algorithm 15.1. Using this algorithm upon the pairs described in Table 15.1 results in the removal of $\{L2, R2\}$, followed by $\{L1, R1\}$. The remaining set of pairs, $\{\{L1, R2\}, \{L2, R1\}\}$, use no echo more than once, so the algorithm terminates at this point. These two pairs correctly describe the situation of Figure 15.1, as the computed bearings closely match the true bearings of the targets.

15.3.2 Statistical Method

Statistical theory (Crow [27, Section 2.4.2]) provides a method for testing whether two observations are statistically similar or different. Firstly, compute the bearings θ_{IDD} and θ_{IPD} from a single sonar observation of a target. Secondly, test whether the two bearings are statistically equivalent. Assuming that they are drawn from populations with known standard deviations $\sigma_{\theta_{\text{IDD}}}$ and $\sigma_{\theta_{\text{IPD}}}$, which are taken from the bearing calibration experiments of Chapter 14, then the statistical test may be performed as follows. Compute

$$z = \frac{\theta_{\text{IDD}} - \theta_{\text{IPD}}}{\sqrt{\frac{\sigma_{\theta_{\text{IDD}}}^2}{n} + \frac{\sigma_{\theta_{\text{IPD}}}^2}{n}}}, \quad (15.4)$$

where n is the number of sonar samples used to compute the bearing estimate. Setting n to 1 reflects the fact that the bearing is computed from a single sonar measurement. Thus,

$$z = \frac{\theta_{\text{IDD}} - \theta_{\text{IPD}}}{\sqrt{\sigma_{\theta_{\text{IDD}}}^2 + \sigma_{\theta_{\text{IPD}}}^2}}.$$

Table 15.2: Values of the computed minimum bearing error rejection criterion for the statistical correspondence solver. The standard deviation of **IDD** bearing estimates are taken from Figure 14.16, while the standard deviation of **IPD** estimates are taken from Figure 14.24, and the minimum bearing difference such that there is only a 5% chance that the pairing is incorrect. The rejection difference column is computed using (15.5). Any pairings for which $|\theta_{\text{IDD}} - \theta_{\text{IPD}}|$ is larger than the value shown in the right column are rejected by the statistical correspondence solver.

Sonar System	Minimum Rejection Difference	
	$\sigma_{\theta_{\text{IDD}}}$	$\sigma_{\theta_{\text{IPD}}}$
MF-QD	1.06°	1.07°
MF-UD	0.86°	1.07°
CTFM-DFT	1.82°	1.22°
CTFM-YW(128)	0.73°	1.56°
CTFM-LSMYW(128,128)	0.60°	1.92°

Now, θ_{IDD} and θ_{IPD} are statistically equivalent with a confidence level of 95% if $|z| < 1.960$. The relationship may be re-written to provide a maximum acceptable bearing difference:

$$|\theta_{\text{IDD}} - \theta_{\text{IPD}}| < 1.960 \sqrt{\sigma_{\theta_{\text{IDD}}}^2 + \sigma_{\theta_{\text{IPD}}}^2}. \quad (15.5)$$

Using the experimental values for $\sigma_{\theta_{\text{IDD}}}$ and $\sigma_{\theta_{\text{IPD}}}$ from Chapter 14, the maximum acceptable bearing differences can be computed for each sonar type. The results of this computation are shown in Table 15.2. This table shows that the MF-UD method provides the tightest rejection criterion, and is therefore able to reject more false pairings. The CTFM-DFT method provides the worst bearing rejection criterion. A larger bearing rejection criterion results in more bad pairings passing through than a small bearing rejection criterion would allow.

Example

Continuing with the example from Section 15.2, we apply the statistical correspondence solver to the pairings shown in Table 15.1 with an MF-UD sonar system. The statistical test, comparing the bearing difference to 2.69°, indicates that each of the pairs $\{L1, R1\}$ and $\{L2, R2\}$ has only a 5% chance of being correct. The two remaining pairs, $\{L1, R2\}$ and $\{L2, R1\}$, with a 95% probability of being correct, accurately describe the targets which were used to generate the echoes.

Plot

A plot of (15.5) is shown in Figure 15.7. Using this plot, we can determine the performance improvement required to obtain a specified rejection criterion. For example, if we desire a bearing difference rejection criterion of 2°, then this may be achieved by a sonar having $\{\sigma_{\theta_{\text{IDD}}}, \sigma_{\theta_{\text{IPD}}}\}$ which lies along the 2° isopleth, such as $\{0.74, 0.74\}$ or $\{0.50, 0.93\}$.




Figure 15.7: Isopleth plot of the 95% confidence bearing difference rejection criterion in terms of $\sigma_{\theta|IDD}$ and $\sigma_{\theta|IPD}$, computed using (15.5). The isopleth levels indicate the minimum bearing error which must occur before a pairing may be rejected with a 95% confidence. The points corresponding to the **IDD** and **IPD** bearing precisions measured in Chapter 14 are also shown for each sonar system. This plot allows different sonar systems to be compared based upon their **IDD** and **IPD** bearing precision. A smaller rejection criterion, obtained by a sonar with the best bearing precision, will appear toward the bottom left of the plot.

Analysis of the plot indicates that improvement of the bearing rejection criterion of the CTFM-YW and CTFM-LSMYW sonar systems requires improvement of the **IPD** bearing precision. Similarly, improvement of the CTFM-DFT rejection criterion requires improvement of the **IDD** bearing precision. The MF-UD sonar system provides the tightest bearing rejection criterion of all of the systems tested, requiring that the bearings computed by the two methods should not differ by more than 2.69° if the pairing is to be accepted. This is to be compared to the worst performer, the CTFM-DFT, where the rejection level is 4.29° , being almost twice the minimum value.

Discussion

The statistical correspondence solver, which rejects all pairs having a difference in their computed bearings of larger than the rejection criterion, may fail to remove all pairs which use a given echo more than once. When there are two pairs using the same echo which have bearing differences of less than the rejection criterion, we have no statistical basis to prefer one over the other. In this case the sonar system must report some fictitious targets.

15.4 Combining IDD and IPD Bearings

In experimental work, it is common to combine two or more differently derived results to reduce random errors, inherent in the experimental process, that affect the accuracy and precision of the derived results.

A single sonar measurement and echo pairing will result in bearing estimates θ_{IDD} and θ_{IPD} . The experiments of Chapter 14 have established the precision which can be expected of these observations, in the form of the standard deviations $\sigma_{\theta_{\text{IDD}}}$ and $\sigma_{\theta_{\text{IPD}}}$ which were observed during the calibration experiments. The bearing estimated from a single sonar measurement is therefore assumed to be drawn from a population having the same standard deviation as the calibration experiment.

As the relative precisions of the bearing estimates are known, they can be combined using the weighted averaging technique described in Appendix I.8. The combined bearing estimate, computed by weighting the two bearing estimates according to their precision, is (using (I.24))

$$\theta_c = \frac{\sigma_{\theta_{\text{IDD}}}^2 \theta_{\text{IDD}} + \sigma_{\theta_{\text{IPD}}}^2 \theta_{\text{IPD}}}{\sigma_{\theta_{\text{IDD}}}^2 + \sigma_{\theta_{\text{IPD}}}^2}. \quad (15.6)$$

The correlation between the bearing residuals measured by the **IDD** and **IPD** methods was demonstrated to be small in Section 14.3, but not small enough to be neglected. The standard deviations of the **IDD** and **IPD** methods are thus combined using (I.27) to yield

$$\sigma_{\theta_c} = \sqrt{\frac{\sigma_{\theta_{\text{IDD}}}^2 \sigma_{\theta_{\text{IPD}}}^2 (\sigma_{\theta_{\text{IDD}}}^2 + 2\sigma_{\theta_{\text{IDD}}\theta_{\text{IPD}}} + \sigma_{\theta_{\text{IPD}}}^2)}{(\sigma_{\theta_{\text{IDD}}}^2 + \sigma_{\theta_{\text{IPD}}}^2)^2}}. \quad (15.7)$$

The combination of bearing estimates using statistical weighting is only useful if it significantly increases the precision of the final result, that is, the level of random error is reduced. To test this, equation (15.7) was used to compute the standard deviation that would be obtained by

Table 15.3: Standard deviations of combined **IDD** and **IPD** bearing estimates for each sonar type tested in the bearing calibration experiment. The $\sigma_{\theta_{\text{IDD}}}$ are extracted from Figure 14.16, while the $\sigma_{\theta_{\text{IPD}}}$ are extracted from Figure 14.24. The covariance $\sigma_{\theta_{\text{IDD}}\theta_{\text{IPD}}}$ is calculated from the correlations in Figure 14.27. The σ_c column is computed using (15.7). Note that in the case of the CTFM-LSMYW system the combined standard deviation is worse than that of the **IDD** bearing. This is mostly due to the imprecise **IPD** measurement, with a small contribution from the covariance.

Sonar System	Bearing Precision			
	$\sigma_{\theta_{\text{IDD}}}$	$\sigma_{\theta_{\text{IPD}}}$	$\sigma_{\theta_{\text{IDD}}\theta_{\text{IPD}}}$	σ_{θ_c}
MF-QD	1.06	1.07	0.245	0.83
MF-UD	0.86	1.07	0.234	0.75
CTFM-DFT	1.82	1.22	0.788	1.17
CTFM-YW(128)	0.73	1.56	0.080	0.68
CTFM-LSMYW(128,128)	0.60	1.92	0.272	0.61

combining the **IDD** and **IPD** bearing calculations for each of the five methods presented in Chapter 14. The results are shown in Table 15.3. This table indicates that the most precise results are obtained from the combined CTFM-LSMYW sonar system, followed closely by the CTFM-YW and MF-UD systems.

An important insight into the operation of the weighted average formula can be obtained from the isopleth plot shown in Figure 15.8. The covariance was set to zero to obtain a function of two variables. The graph can be used to analyse how to improve the combined bearing estimate performance of the sonar system. Values of the combined bearing precision σ_{θ_c} may be found as isopleths. The values of $\sigma_{\theta_{\text{IDD}}}$ and $\sigma_{\theta_{\text{IPD}}}$ required to achieve the desired output precision σ_{θ_c} may then be found by reading the coordinates $\{\sigma_{\theta_{\text{IDD}}}, \sigma_{\theta_{\text{IPD}}}\}$ from the graph.

Using the graph, we can see that the CTFM-DFT precision may be significantly improved by modestly improving $\sigma_{\theta_{\text{IPD}}}$. However, a much larger improvement in $\sigma_{\theta_{\text{IDD}}}$ is necessary to achieve the same result.

The graph also shows that the major contributor to the precision of the CTFM-LSMYW result is the low value of $\sigma_{\theta_{\text{IDD}}}$ obtained by this system. It would take a very large improvement in $\sigma_{\theta_{\text{IPD}}}$ before a noticeable improvement in σ_{θ_c} would be realised. However, any improvements in $\sigma_{\theta_{\text{IPD}}}$ would be almost directly reflected in σ_{θ_c} .

The greatest benefit of weighted averaging is obtained when the standard deviations of the two measurements being combined are approximately equal. Combining two uncorrelated bearing measurements with standard deviations of $\sigma_{\theta_{\text{IDD}}} = \sigma_{\theta_{\text{IPD}}} = 1^\circ$ results in $\sigma_{\theta_c} = 0.71^\circ$, a significant improvement. However, when one of the standard deviations is large, as with the CTFM-LSMYW system, the smaller of the two dominates the contribution to the combined precision. This is because there is no significant reduction in the random errors present, due to the large uncertainty in the complementary component of the pair.

Table 15.3 shows that the CTFM-LSMYW sonar is the only system which loses precision by combining the two bearing measurements. This only occurs when correlation between the bearing estimates is taken into account.




Figure 15.8: Isopleth plot of the standard deviation of the combined **IDD** and **IPD** bearing measurements, computed using (15.7). The results determined for each sonar type are also marked, neglecting correlation. As the covariance between the measurements was set to zero to obtain this plot, the values obtained for σ_{θ_c} are slightly smaller than those in Table 15.3. The CTFM-LSMYW can be seen to be the most precise, while the CTFM-DFT system has the worst precision.

15.5 Simulation Experiments

Four simulation experiments were carried out to test the effectiveness of the simple and statistical correspondence solvers. The simulated echoes were modelled to provide both the **DOF** and echo power which were subsequently interpreted to yield the **IDD** and **IPD** estimates of bearing.

The four experiments used:

1. The simple correspondence solver with noise-free echo data from the model,
2. The simple correspondence solver with realistic noise added to the echo data from the model,
3. The statistical correspondence solver with realistic noise added to the echo data from the model, and
4. The statistical correspondence solver with realistic noise added to the echo data from the model, with echoes closer together than the **DOF** resolution being fused together.

This section now discusses the modelling used for the four experiments before describing the experiments and the principal results. The section concludes with a discussion of the results.

15.5.1 Modelling

The experiment was conducted using the 2D geometrical model described in Appendix D. The model does not simulate the sonar signal, but directly computes the path distances. The model produces various pieces of information about the path, such as which transmitter and receiver form its endpoints. The model can return the angle at which the path exits the transmitter and the angle at which it enters the receiver. Using the experimental data shown in Figure 14.18, an **IPD** model was constructed which fitted Gaussian curves representing the responses of the transmitter and left and right receivers, similar to Figure 7.7. The method used to construct the **IPD** model is described in Appendix H. Thus the geometrical model can return similar information to that which is extracted from the sonar **DOF** sub-system, as described in Part III of this thesis.

The geometrical and **IPD** models have been used to test the solution to the correspondence problem without interference from practical sonar effects such as interference between overlapping echoes, which are discussed in Section 13.3. As the model was only required to provide sufficient information to test the correspondence solver, it was not required to accurately simulate the proper echo power which would be detected by the sonar system, but only to simulate the correct power ratio. Therefore the physical properties of spherical spreading loss and the relative strengths of different kinds of reflectors have not been included in this model.

The model was built to simulate the operation of the **matched filter unsynchronised demodulation** sonar system. The **IDD** bearing results were simulated by setting the sonar head dimensions to $d_T = 2.7$ mm and $d_R = 13.5$ mm in the geometrical model. The **IPD** results were imitated by using the left and right channel modelled beam powers from Figure 14.17, obtained from the MF-UD sonar system, to calculate the beam functions of the three transducers. The **IPD** calibration curve was taken from Figure 14.20. Thus the model could produce simulated echo **DOF** and echo power

given a target range and bearing. The bearing system outlined in Chapter 14 could then calculate the range and bearing of a target given this echo information.

The model additionally provides the ID number of the object which produced each echo. This information can be used to evaluate whether a given pairing is accurate, by testing whether the object ID numbers corresponding to each of the echoes forming a pair are one and the same.

IPD Model Verification

The IPD modelling, which uses the echo power component, was validated by placing a target at $r = 0.5$ m and then varying the bearing over the range -40° to 40° . A maximum discrepancy of 0.7° was found, which was attributed to the zero order approximation used to linearise the problem of Appendix H.4. This maximum discrepancy indicated that the modelling undertaken to generate the synthetic data was both valid and realistic.

15.5.2 Noise-Free Echoes with Simple Solver

The first experiment used noise free echo data and two plane targets. Target A was assumed fixed at $\{x, y\} = \{0.55 \text{ m}, 0.05 \text{ m}\}$. Target B was swept past target A, such that the normal of both targets always pointed toward the centre of the transmitter at $\{x, y\} = \{0, 0\}$. The model used in this experiment was able to resolve two targets located at the same DOF, normally not possible due to finite resolution.

The number of correct pairs was independently assessed by comparing the echo pairings to an object ID which was attached to each echo by the model. A correct pairing will include echoes from the left and right channels with the same object ID. The correspondence solver does not require or have access to this information.

The aim of this experiment was two fold:

1. to validate the experimental approach, and
2. to determine regions where the correspondence solver fails to identify both of the targets.

Since the two correct echo pairings are unique, it is expected that the correspondence solver will find them. However, there are two regions where two correct pairings are not presented.

The first, in close proximity to the fixed target, is due to both occlusion of target A by target B, and vice versa. This area, grey in Figure 15.9, is where only one of the two required pairs is determined.

The second area is where no correct pairs, black in Figure 15.9, are determined. This area is limited to points having the same range but different bearings. The small 0.7° discrepancy in the IPD calibration is sufficient to confuse the simple correspondence solver into discarding the correct pairs. This experiment validates the approach taken. In particular, the principal region where the correspondence solver fails was determined. The same range, different bearing problem is expected to be present in all experiments.




Figure 15.9: Plot of the cumulative number of pairs correctly identified by the simple correspondence solver, when used with noise-free measurements. Black = 0, grey = 1, white = 2 correct pairs. The sonar head was located at $\{x, y\} = \{0, 0\}$. The plot was generated using two targets, one fixed (drawn in red) and one movable. The movable target is placed at a location, the model is used to compute echoes, and the simple correspondence solver is used to produce echo pairs. These echo pairings are then checked against the information provided by the model and the number of correct pairs is computed. This number is then plotted at the centre point of the movable target. The procedure is repeated at each position on the diagram.

15.5.3 Noisy Echoes with Simple Solver

The second experiment, the same as the first experiment but with added random noise, was aimed at making a more realistic data set with which to test the correspondence solver.

The **DOF** additive noise was modelled by a multi-normal distribution having $\sigma_{d1} = \sigma_{d2} = 0.90$ mm and correlation coefficient $c_{d1d2} = 0.858$. The power additive noise was also modelled by a second multi-normal distribution with $\sigma_{p1} = \sigma_{p2} = 1.45$ dB and correlation coefficient $c_{p1p2} = 0.705$. The additive noise models are those determined from the **matched filter unsynchronised demodulation** sonar system, as reported in Tables 13.4 and 13.5. They are therefore representative values. Because white additive noise was used in this experiment, it was replicated ten times. The replication is consistent with the way the physical experiments in the previous chapters were run and ensures that these random errors do indeed behave in a random manner. The number of correct pairs was determined using the same object ID technique which was used in the first experiment. In all there are up to 2 possible correct echo pairings from each of 10 trials.

The cumulative results of the experiment are shown in Figure 15.10a. These results have many of the same features as the previous experiment. Overlapping targets produce a grey region near the fixed target, where one or more echoes are occluded. Most of the rest of the space corresponds to correctly identified pairs. However, the difference in this experiment is that the spotty arc of Figure 15.9a has become a solid locus with a thickness dimension, consisting of incorrect pairings. The simple correspondence solver always fails when there is more than one target with similar **DOFs**.

The width of the failure region is of interest. This was investigated by graphing the number of incorrect pairings as a function of the target range, at several values of the target bearing, in Figure 15.10b. The same experimental configuration was used as for Figure 15.10a, but the movable target B was only moved along a radial line, and on a denser grid. The number of incorrect pairs was counted. Thus these plots are histograms of the failure range of the correspondence problem. The incorrect pair counts are normally distributed in range, and have an average standard deviation ± 0.76 mm. As 95% of the incorrect pairings are included in a window of 1.96σ , we assume that the failures occur within a window of ± 1.5 mm.

While the sonar resolution has not been modelled in this experiment, it is instructive to compare the measured resolution with the width of the correspondence failure region. Table 13.10, which lists the range resolution results, shows the minimum range resolution of any sonar system is 6 mm. Thus the correspondence failure region is smaller than the region where echoes cannot be resolved. Therefore it is expected that the simple correspondence solver will perform adequately when presented with real echo data.

15.5.4 Noisy Echoes with Statistical Solver

The third experiment used the same setup and additive noise model as was used in the second experiment, producing perfectly resolvable echoes. The difference is that rather than using the simple correspondence solver of Section 15.3.1, the statistical correspondence solver of Section 15.3.2 was implemented. The statistical correspondence solver, unlike the simple correspondence solver, outputs duplicate pairs when it cannot completely solve a set of inter-related echoes. Thus for

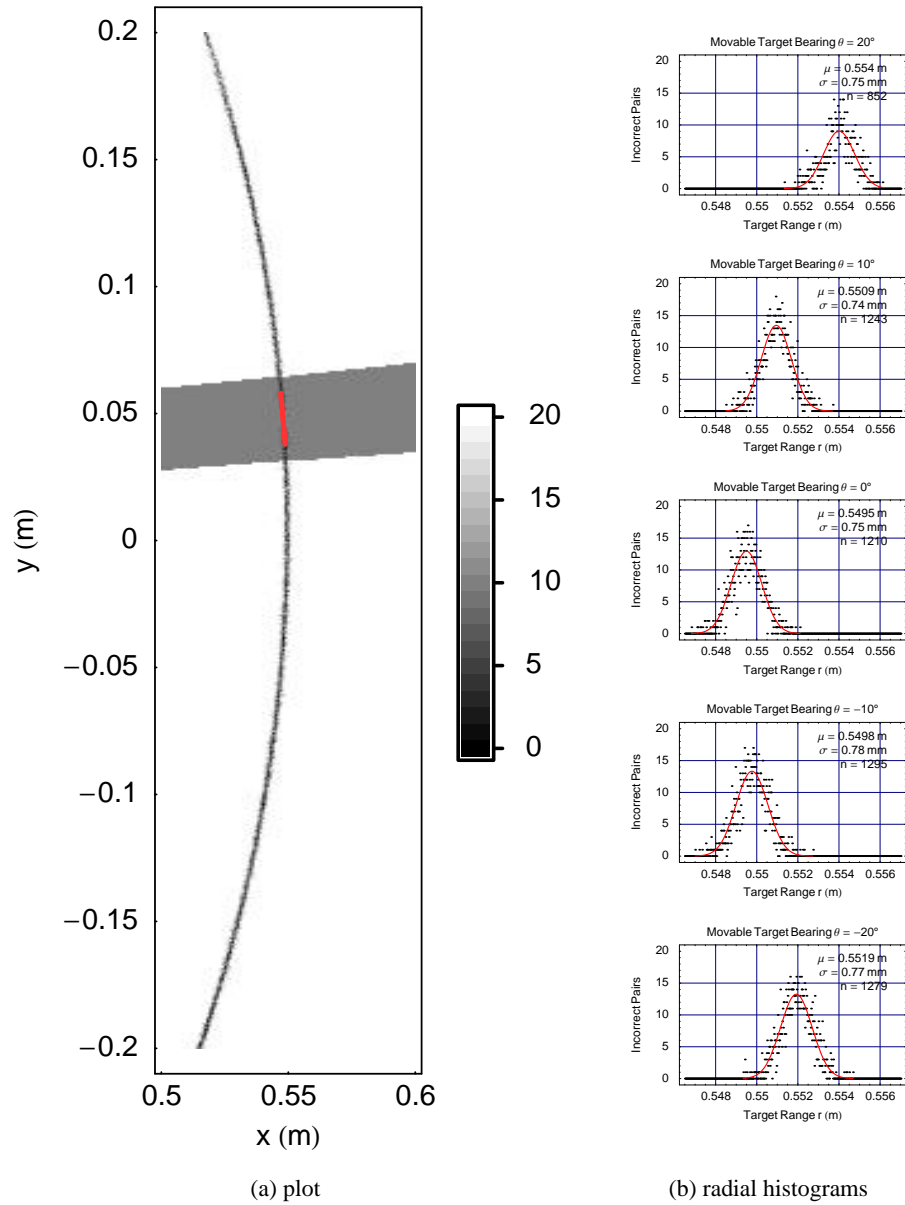


Figure 15.10: (a) Plot of the cumulative number of pairs correctly identified by the simple correspondence solver. This plot was carried out by adding noise to the echo measurements before computing bearings by **IDD** and **IPD** for the simple correspondence solver. The number of pairs identified correctly over ten trials was counted and plotted as a grey level (see legend) drawn at each location of the movable target. (b) Histograms of the range of incorrect pairings produced by the simple correspondence solver when presented with noisy data. Two targets were used, in the same configuration as the experiment shown in (a). The plots are cross sections of the failure region shown in (a). The red curve indicates a normal distribution which has been fitted to the n data points with the indicated mean and standard deviation. The data consists of the range of each incorrect pair, so a dot at height 12 indicates that the range is replicated 12 times in the data set, leading to the high value of n . The standard deviations measured at each bearing are statistically equivalent.

regions where a pair is missing due to occlusion, no pair is reported. Thus a different evaluation mechanism is adopted. In this case it is appropriate to focus on the number of reported unresolved pairs.

Figure 15.11a displays the results of this experiment. The first difference between the results in Figure 15.10a and Figure 15.11a is that the occlusion region is gone. This is due to counting the number of duplicate pairs output by the statistical correspondence solver, not because the correspondence solver is able to resolve occluded objects.

The shape of the area in which the statistical correspondence solver returns duplicate errors is quite similar to the area in which the simple correspondence solver returns incorrect echo pairings. For completeness, histograms of the radii where the duplicates occur are presented in Figure 15.10b. The range of the duplicates area changes with bearing. The average standard deviation of the duplicate region is ± 0.64 mm, giving a 95% confidence region of ± 1.25 mm. Due to the number of trials involved, this can be seen to be significantly smaller than the region where the simple correspondence solver outputs false echo pairings.

This experiment has shown that the statistical solver produces duplicate echo pairings over a smaller area than that for which the simple solver produces incorrect pairings. The duplicate echo pairing region has a smaller width than the resolution criteria, which is 6 mm. Thus we expect that the failure of the statistical correspondence solver will be masked by the inability of the sonar system to resolve echoes from objects closer together than 6 mm.

15.5.5 Unresolved Noisy Echoes with Statistical Solver

The fourth experiment used the same geometrical model and noise model as was used previously, but also simulated sonar resolution. The resolution limit was simulated by fusing echoes which would be unresolvable in a real sonar system, determined by the results shown in Table 13.10. As the computer model is based upon the [matched filter unsynchronised demodulation](#) system, a [DOF](#) resolution of 12 mm was used. The fusion was implemented by averaging the two [DOFs](#) and adding the two echo amplitudes to obtain a new echo record. Only echoes occurring on the same channel are fused.

As the fused echoes no longer have a one to one correspondence with a target, a new method of evaluating the sonar performance was devised. The sonar computed range and bearing is compared with the model target range and bearing. The two are compared using the standard deviation of the range and bearing measurements as obtained in Tables 14.1 and 15.3. If a target is found which has a difference of less than 1.96σ from the sonar measured value, that is a 95% confidence interval, then the sonar target is deemed to be correct. The criteria were measured based upon the range and bearing differences in a similar manner to the drawing of the [VCV](#) ellipsoid of standard deviation in Appendix I.4. If no real target is found matching the criterion, then the sonar target is defined to be incorrect.

The aim of this experiment was to test the effect of echo resolution upon the solution of the correspondence problem. Specifically, we wish to determine if the [DOF](#) resolution is indeed the limiting factor in determining the size of the region where targets cannot be correctly identified, as predicted from previous experiments.

The results for the experiment which counted the number of correctly identified targets show




Figure 15.11: Results of the experiment using noisy echoes to test the statistical correspondence solver. The same experimental configuration was used as in Figure 15.9. Panel (a) shows the number of duplicate pairs. Ten trials were performed at each position on the plot, and the number of pairs within the duplicate group was counted. The sum over the ten trials is the number plotted. The fixed target is shown as a red line. Panel (b) shows radial histograms of the number of duplicate pairings. The histograms show the number of duplicate pairs evaluated along a radial line placed at five different bearings. Thus the plots are cross sections of the duplicate region in (a). The red curve indicates a normal distribution which has been fitted to the range of the duplicate count data. The mean and standard deviation are indicated. These plots were produced using DOF and echo power noise modelled from the MF-UD results. The range standard deviations measured at each bearing are statistically equivalent.

three main regions. The first, being the 20 count medium grey region near the stationary target A in Figure 15.12a, coincides with the occlusion region seen in experiments one and two. The second, being the 35 count grey region, contains mostly correctly identified target positions. The identification noise present in this region is attributed to the method of assessing the results and not to the incorporation of echo fusion, which does not operate in this region. The third, being the circular 0-20 count black-grey region, is attributed to the effect of target resolution. The 20 count grey subregion is attributed to the generous bearing criterion used to test for correct targets. The width of the resolution failure region is about 0.01 m. This is about twice the resolution criterion.

The false sonar target results shown in Figure 15.12b have two important features. The first is an edge effect, corresponding to the edges of the failure region in Figure 15.12a. The cause of this is unknown. As it is of a similar strength to the background level of undetected targets in Figure 15.12a, it is suspected that the same mechanism is causing both phenomena. The second region where false sonar targets are reported is near the coordinate $\{0.53\text{ m}, -0.12\text{ m}\}$. This is attributed to the requirement to accept a wide range of bearings, and the effect of the false pairings (as shown in Figure 15.6) will place the false sonar targets outside the range of the matching code described earlier. Note that this region complements the grey region in Figure 15.12a.

15.5.6 Discussion

The four computer based experiments have demonstrated that:

1. the sonar system cannot resolve two targets when one is occluded by the other, thus blocking one or both of the sonar paths,
2. the simple correspondence solver produces incorrect echo pairings when the movable target is placed within an arc like region of width $\pm 1.5\text{ mm}$,
3. the statistical correspondence solver produces duplicate echo pairings when the movable target is placed within an arc like region of width $\pm 1.25\text{ mm}$, and
4. the range resolution criterion, being 6mm, causes the width of the region where targets cannot be identified to increase to about 10mm. The width has approximately doubled, as targets are unresolvable whether they are 6mm in front or 6mm behind.

Target identification failure may be caused by resolution failure or correspondence solver failure. Of these, it is the sonar resolution which is shown to be the defining quantity in the sonar system tested. Shrinking the identification failure region demands that the sonar resolution be improved, and is not practically limited by the correspondence solver.

It is instructive also to compare the size of the DOF window, which is $\pm 9.2\text{ mm}$ for the sonar systems tested here, with the width of the identification failure region. The DOF window may be thought of as a range window of $\pm 5.1\text{ mm}$, which is of about the same width as the resolution failure region, which is about 10mm wide. This indicates that, in general, the resolution criteria will prevent more than one echo from being detected within the DOF window, and thus will remove the occurrence of duplicate pairs from the system. It is expected, therefore, that there will be little call for pair rejectors unless the DOF resolution is improved.




Figure 15.12: Plot showing the results of testing the statistical correspondence solver with noisy echo data and fused echoes where the echoes would not be resolvable by the MF-UD sonar system. Ten trials were performed at each position. In (a), the number of targets correctly identified by the sonar system was counted and added over all of the trials. In (b), the number of sonar targets having no corresponding modelled target was counted.

15.6 Bearing Resolution

It was shown in Section 14.1.3 that the precision of bearing measured by **IDD** is determined by the precision and correlation of the **DOF** measurements. It was similarly shown in Section 14.2.4 that the precision of bearing measured by **IPD** is determined by the precision and correlation of the power measurements. The ability to measure the **DOF** and echo power of two close targets is further limited by the interference between the two echoes, producing a fundamental **DOF** resolution limit, which was investigated in Section 13.3.

Consider first the case where two targets are not contained inside the 2D horizontal plane, such as the targets of the calibration device shown in Figure F.1. This device has four targets, all with the same bearing, but each with a different elevation angle. The sonar results shown in Figure F.4 demonstrate that the sonar system is capable of detecting four targets having the same bearing when they do not occlude each other and have different ranges. Note that targets lying far outside the horizontal plane will obtain poor **IPD** bearing results, as discussed in Section 14.2.6.

Secondly, consider the case where the two targets are restricted to lie within the same horizontal plane. The bearing resolution of such a system is constrained by two cases. The first case is where one target occludes one or both of the echoes from the other target, as may be clearly seen in the results of the correspondence solver in Figures 15.9, 15.10a, and 15.12a. The second case is the limitation imposed not by the correspondence problem but by the range resolution problem. When the echoes from the two targets cannot be distinguished, their bearings cannot be computed. More likely, a fused echo will be detected, and a false sonar target reported.

Thus the mechanisms limiting bearing resolution are due to occlusion and the range resolution. The region over which targets cannot be resolved is quite small when compared with the whole sensing volume. No separate experiment was performed to measure the bearing resolution.

15.7 Conclusion

Two methods of solving the correspondence problem have been discussed. These are the simple correspondence solver and the statistical correspondence solver. Both techniques compare the results of the two different methods of computing bearing. The statistical correspondence solver was shown to have a slightly smaller region where incorrect or inconclusive results are obtained than the simple correspondence solver. The statistical correspondence solver provides a pairing rejection threshold based upon the bearing difference between the **IDD** and **IPD** bearing results. The rejection threshold, shown in Table 15.2, shows that the MF-UD sonar system has the tightest criteria leading to the most reliable pairing.

To robustly deal with the situation where duplicate pairs cannot be eliminated, the correspondence solver must report all possible targets, as indicated by the pairings which were not rejected by the statistical test. The false targets must subsequently be distinguished from the real targets by moving the sonar head and observing which targets remain fixed, and which targets move around or disappear. As the correspondence problem is most acute when there is more than one target at approximately the same range, it is advantageous to move the sonar head so that the range to the targets can be differentiated. This can be achieved with a sideways movement of the sonar head.

If two sonar heads are available, as shown in Figure 7.1, then the output from both sonar heads may be fused to solve the problem and no robot motion is necessary.

However, of more importance in a practical sonar system is the impact of the range resolution which was measured in Section 13.3. The experiments of Section 15.5.5 showed that the inability to resolve targets extends over a similar region to that where the correspondence problem is relevant. Therefore it is postulated, based on the experiments with computer models performed in this chapter, that there is little need for a pair rejector, as the ability to form pairs is tightly coupled to the range resolution. The resolution results obtained by the MF-UD system, which was used as the basis for the computer model in this chapter, are near the limits of resolution predicted by the *time-bandwidth product* theory of Sections 4.6 and 8.8; further improvement in range resolution seems unlikely. Consequently there is a lower bound upon the size of the region where targets cannot be identified from the echoes.

When the echo pairing corresponding to a target has been identified, the bearings computed by the methods of *IDD* and *IPD* may be combined (see Table 15.3). This provides better bearing precision than using only one of the bearing measurements, except when one of the bearing measurements has a large error. In that case, it is better not to combine the bearing measurements, and to instead use the more precise measurement directly.

The bearing resolution is found to have several important cases, depending upon whether or not the targets are constrained to lie within the horizontal plane. When several targets lie outside the plane, at separate ranges, do not obscure each other, and have identical bearings, they may be resolved and have their bearings measured. However, objects having similar ranges or obscuring each other may not have their bearings measured. The region where the range resolution problem prevents bearing from being measured is small when compared with the sensing region.

15.7.1 Acknowledgements and Contributions

The *DOF* window of Section 15.1 is the work of Kleeman and Kuc [82]. The two correspondence solvers, described in Section 15.3, the combination of the *IDD* and *IPD* bearings, described in Section 15.4, and the computer experiments to test and validate the solutions to the correspondence problem, described in Section 15.5, are the contribution of this author.

Part V

Conclusion

Chapter 16

Conclusion

16.1 Summary of Aims and Results

Chapter 1 stated that the aims of this thesis were to probe and investigate the use of two binaural sonar classes, the [Matched Filter \(MF\)](#) sonar system and the [Continuous Tone Frequency Modulated \(CTFM\)](#) sonar system, as a technique to sense the environment, enabling a *robotic agent* to function in the place of a *living agent*. These aims were to be achieved through four goals:

1. *To build a [MF](#) binaural ultrasonic sensing system.* The construction and operation of a [MF](#) binaural ultrasonic sensing system is discussed in Chapter 8. That chapter concludes that the MF-UD system is expected to have superior performance characteristics to the MF-QD system as the demodulation is superior.
2. *To build a [CTFM](#) binaural ultrasonic sensing system.* The construction of a [CTFM](#) binaural ultrasonic sensing system is discussed in Chapter 9. That chapter concludes that a [CTFM](#) system capable of producing tonal echo signals as its output signal was successfully built. Three spectral estimators were then built in software, to analyse the output of the [CTFM](#) system. The first, utilising the [discrete Fourier transform](#), is described in Chapter 10 which reaches the conclusion that the poor level of quantisation inherent in the method will limit the applicability of this approach. The second and third estimators, which belong to the [auto-regressive](#) class, are discussed in Chapter 11 which reaches the conclusion that this [AR](#) class has better precision than the [DFT](#) class of Chapter 10.
3. *To investigate the ability of the [MF](#) and [CTFM](#) sonar systems to measure range and bearing and to quantify their performance.* Chapter 2 defined several performance criteria which may be used to quantify sonar performance. Chapters 13 and 14 discuss the range and bearing performance of the five sonar systems built. The significant results extracted from these chapters are tabulated in Tables 16.1 and 16.2.

Table 16.1 and the discussion in Chapter 13 indicate that the best precisions and resolutions for range are achieved by the [MF](#) systems. These precisions and resolutions are statistically different from those achieved for the [CTFM](#) systems due to the large number of degrees of freedom ($n = 1500$) in the data sets. On the other hand Table 16.2 and the discussion in Chapters 14 and 15 indicate that the best precision for bearing is the [IDD](#) approach using a [CTFM-LSMYW](#) system.

For bearing resolution the MF-UD system is clearly the best performing system studied in this thesis. The differences between the principal sonar systems, MF, CTFM-DFT and CTFM-AR are significant due to the large ($n = 400$) degrees of freedom that were used to generate the comparative information.

The **Signal to Noise Ratio (SNR)**, measured for a plane target at $r = 0.5$ m, determines the maximum range (see Section 13.2.4). A noisy environment will increase the noise level and reduce the SNR. Therefore it is important to have some excess SNR if the system is to operate in a noisy environment. Table 16.1 shows that the CTFM class has an advantage over the MF class of sonar systems when the SNR is compared.

4. *To determine the best system or systems to emulate sight as a sensor for a robotic agent.* Numerical indicators were introduced in Chapter 2 to facilitate this determination. These numerical indicators have been used to generate the information in Tables 16.1 and 16.2. In addition to these numerical indicators it is also necessary to consider the following three factors:

- (a) The complexity of each sonar system. Clearly the *robotic agent* needs to be able to process information, but it must simultaneously perform other tasks such as waveform analysis and mapping of its environment.
- (b) The dependency of the system on external calibrations and whether these calibrations will remain stable. The necessity to calibrate the offset time is discussed in Chapter 13, while transducer beam power, width and biases for bearing calculation by IPD are discussed in Chapter 14.
- (c) The ability to accurately solve the correspondence problem, as measured by the bearing rejection criterion of Chapter 15.

The performance of the sonar systems as measured by these three factors is summarised in Table 16.3.

Thus selecting the *best* system is a complex task of optimising the results from Tables 16.1, 16.2 and 16.3 for the application at hand. In general, selecting a MF based system would yield a high quality system having stable characteristics over a wide range of environmental conditions, including the harsh and dangerous conditions found in environments such as coal mines and nuclear reactors. The major drawback of the class is its lower SNR, and its lesser ability to insonify a large area and to map this space from a single vantage point.

On the other hand, selecting a CTFM based system will significantly improve the bearing determination without significantly degrading the range precision. CTFM systems have been shown to be capable of insonifying a larger area from a fixed position which is an advantage, but these systems require either external or continuous calibration to perform at their peak capacity.

This thesis suggests that a CTFM system is the best system with which to equip our *robotic agent* if it is to emulate the *living agent*. In making this choice the SNR and the ability to insonify a large area from a single position were considered to be important components of the emulation of

Table 16.1: Results for the performance criteria applied to range measurements, and the results from the comparable work of other researchers as reviewed in *Chapter 2*. The results are compared in Section 16.2.

Sonar System	Range					
	Accuracy mm	Precision mm	Resolution mm	SNR at $r = 0.5$ m dB	Maximum Range m	
	Section 13.2.5	Table 14.1	Table 13.10	Table 13.8	Table 13.9	Source
MF-QD	-	0.44	6	48.9	4.7	
MF-UD	-	0.43	6	48.5	4.7	
CTFM-DFT	-	0.57	5	55.0	5.0	
CTFM-YW	-	0.48	20	58.5	5.8	
CTFM-LSMYW	-	0.49	14	-	-	
						See Section
<i>Polaroid</i>	-	1.0	-	-	5	2.5.1
<i>Kleeman and Kuc</i>	0.8	0.1 – 0.2	10	-	8	2.5.6
<i>Peremans</i>	-	1.0	-	-	4	2.5.7
<i>Kay Scanning Array</i>	-	-	6	-	-	2.5.9
<i>Akbarally Combined Visual</i>	0.8	0.15	-	-	6	2.6

Table 16.2: Results for the various performance criteria applied to the two methods of measuring bearing, the [Inter-aural Distance Difference \(IDD\)](#) method and the [Inter-aural Power Difference \(IPD\)](#) method. The [IPD](#) accuracy results are not quoted as this would have required a second experiment, as explained in [Section 14.24](#), which was not performed. Comparable work of other researchers, as reviewed in [Chapter 2](#), is also included. The results are compared in [Section 16.2](#).

Sonar System	Bearing							
	Accuracy Degrees	IDD Precision Degrees	Resolution Degrees	Accuracy Degrees	IPD Precision Degrees	Resolution Degrees	Combined Precision Degrees	Beam Width Degrees
	Figure 14.16	Table 14.2	Section 15.6	Section 14.2.5	Table 14.3	Section 15.6	Table 15.3	Source
MF-QD	-0.27	1.06	N/A	-	1.07	N/A	0.83	±20
MF-UD	0.01	0.86	N/A	-	1.07	N/A	0.74	±20
CTFM-DFT	-0.36	1.82	N/A	-	1.22	N/A	1.17	±20
CTFM-YW	-0.39	0.73	N/A	-	1.56	N/A	0.68	±20
CTFM-LSMYW	-0.91	0.60	N/A	-	1.92	N/A	0.60	±20
					Other			See Section
<i>Yata Single Transducer</i>				-	0.66	-	-	±4
<i>Yata Ring Sensor</i>	-	0.8	-				-	360
<i>Kleeman and Kuc</i>	0.2	0.1	-				-	±11
<i>Peremans</i>	-	0.6	-				-	±12
<i>Kay Scanning Array</i>				-	0.4	0.8	-	±15
<i>Akbarally Combined Visual</i>	-	0.05	-				-	±11

[2.5.4](#)
[2.5.5](#)
[2.5.6](#)
[2.5.7](#)
[2.5.9](#)
[2.6](#)

Table 16.3: Results for the qualitative performance criteria applied to range measurements. Complexity indicates the rank of each sonar system, with 1 representing the simplest and 5 the most complex. Calibration Required indicates whether the DOF sub-system must be calibrated to establish the offset time, t_{off} , which is used for compensating range measurements. The bearing rejection criterion listed in the last column, taken from Table 15.2, is indicative of the ability of the correspondence solver to reject false pairs.

Sonar System	Qualitative Indicators		
	Complexity	Calibration Required	Bearing Rejection
MF-QD	1	No	Good
MF-UD	2	No	Superior
CTFM-DFT	3	Yes	Poor
CTFM-YW	4	Yes	Acceptable
CTFM-LSMYW	5	Yes	Acceptable

sight. The refinement of precise coordinates can be achieved by approaching the object of interest just as a *living agent* normally does.

Unknowns that need to be addressed to back this up are described in Section 16.3.

16.2 Comparison of Results with Other Systems

The review of robotic location sensing technology in Chapter 2 indicated that a number of systems have potential for our *robotic agent*, and defined quantitative methods of comparison. Two classes with a small number of sub-classes were chosen. The MF class was well described by Peremans [108, 110]. Indeed much of the work in this class was a validation of his work. However, the extensions made in Chapter 8 have been significant as shown in Table 16.1 where this study has achieved precisions half that determined by Peremans. The CTFM class was much less defined. Kay's [71, 72] practical work in designing and constructing the CTFM sensors has only seen significant evaluation in the work of Davies [28], but that work presented results in a form which prevented significant comparisons from being made.

The results presented in Tables 16.1 and 16.2 clearly show that both the range and bearing measurements achieved in this thesis, for all five systems, are comparable with other state-of-the-art systems in accuracy, precision, resolution and limiting values.

16.2.1 Range

Range Accuracy

Table 16.1 does not list accuracy estimates for the five systems investigated in this thesis preferring to defer to the discussion in Section 13.2.5. This is primarily due to the fact that accuracy, as defined in Section 2.1, has not been determined. If a relaxed definition, one that includes a range component, is accepted, then the precision of the slope determination is an estimate of accuracy. This is, according to Table 13.3, about -0.55 ms^{-1} , leading to range estimates which are 6 mm

too long for targets at a range of 2 m.

Only two other authors have addressed the topic of sonar range accuracy. The accuracy quoted by Akbarally for his combined visual and sonar sensor utilises Kleeman and Kuc's sonar sensor for the chief contribution to its range measurement. Therefore that work is focused on here. Kleeman and Kuc report ± 0.8 mm accuracy for target ranges up to 4 m. Specifically, Kleeman and Kuc determined the impulse response function of air at 20°C, which is one of the components discussed in Section 13.2.1 as having a range dependent nature. Unfortunately they also used a linear fit value for the speed of sound, as discussed in Section 2.5.6, that is incapable of revealing systematic errors. Thus some cloud also rests over the validity and applicability of the Kleeman and Kuc accuracy determination.

Range Precision

Several of the systems reviewed in Chapter 2 have DOF precisions at the 1 mm level or better, as listed at the bottom of Table 16.1. There are two systems which claim to achieve a precision of ± 0.1 mm. Akbarally's combined ultrasonic and visual sensor uses Kleeman and Kuc's ultrasonic sensor, and is considered to be identical with that system for the purposes of range measurement. Thus Kleeman and Kuc's system is considered.

Kleeman and Kuc only claim the quoted precision of ± 0.1 mm in still air with a calibrated speed of sound. They reduce the quoted precision to ± 0.4 mm:

"...the standard deviation of the distance of flight is 0.4 mm and has been found to depend on air flow conditions in the laboratory – the results were collected with computer fans and air conditioning vents operating as might be the case in a typical office environment."

Kleeman and Kuc [82], p. 309.

They further go on to demonstrate that this variation or repeatability is not dominated by thermal noise and trigger jittering but by variations in the TOF of the pulses through the air. The variations in the TOF are, in turn, ascribed to variations and turbulence in the air path. The experimental work described in this thesis was carried out under similar circumstances. However, the variations and turbulence in the air path are not quantified in either work, making further comparison difficult. The best precisions obtained in the work described by this thesis are comparable with that of Kleeman and Kuc under similar circumstances. It is interesting to note that this result has been achieved with a sampling rate which is only $\frac{1}{3}$ of that used by Kleeman and Kuc (see Table 16.4). In addition, the receive capture trigger in this work has been designed to trigger synchronously with the transmit pulse (see Section 5.1.2), removing trigger jitter as a source of error.

Kleeman and Kuc note a correlation of 0.997 between the echo arrival times recorded by two receivers. Typical values obtained in this work (see Figure 13.13) are in the range 0.8 to 0.9. Comparison between these figures is limited due to their different nature of measurement. Kleeman and Kuc [82] measured their correlation using 350 samples from a single plane target at a range of 3 m. The correlation results from this work were calculated using 10 samples from each position at 151 ranges varying from 0.4 m to 1.9 m in steps of 0.01 m. However, the high degree of correlation observed supports the hypothesis that the majority of the noise is in the air TOF rather than being thermal amplifier noise.

Table 16.4: Signal durations and length in air of the various sonar systems.

System	Rectangular Duration	Rectangular Length in air	Receive Sampling Rate	Gabor Width at Detector	Gabor Range Width at Detector
Barker Coded Matched Filter	468 μ s	162 mm	333.3 kHz	40 μ s	6.9 mm
CTFM	160 ms	55 m	344.8 kHz	10 Hz	5.5 mm
Kleeman and Kuc Impulse	50 μ s	17.3 mm	1 MHz	13 μ s	2.3 mm
Kay CTFM	250 ms	86 m	-	-	-

The figures summarised in Table 16.1 suggest that range precisions of better than ± 1 mm are routinely achievable by ultrasonic sensors.

Range Resolution

The problem of range resolution is addressed by relatively few researchers. Kay appears to have a good understanding of the topic, and provides a calculation in terms of [time-bandwidth product](#) theory. Kleeman and Kuc provide experimentally measured results. The results obtained in this work are comparable with the state-of-the-art figures achieved by Kay, in line with the predictions of the [time-bandwidth product](#) theory.

The range resolution of 10 mm quoted by Kleeman and Kuc is directly related to the length of their signal, which is 17.3 mm. The theory of Section 4.6.3 may be applied to Kleeman and Kuc's system. Assuming a bandwidth of 46.9 kHz (measured from one of the signals published in [82]), the theoretical resolution limit is 1.8 mm. There is a significant gap between this result and the 10 mm resolution quoted. The relatively poor range resolution results of the Kleeman and Kuc template matching system are attributed to the lack of a processing stage in the receiver. The template matching system is constructed to reject overlapping echoes. Thus the resolution of the Kleeman and Kuc system is limited by the duration of the signal, rather than by the available bandwidth. The narrow Gabor width of the 50 μ s signal, compared with that of the [matched filter](#) or the [continuous tone frequency modulated](#) signals employed in this thesis (see Table 16.4), has not translated into improved resolution due to the overlap rejection. On the other hand, Table 13.10 shows that the experimental resolution of the [matched filter](#) class, and the CTFM-DFT system, achieved a result comparable to the theoretical system resolution limit.

Maximum Range

It was shown in Section 13.2.4 that the maximum range is dependent on the [SNR](#), a specified [probability of false alarm](#), and the sonar equation as developed in Section 4.2. Several techniques are available for increasing the signal strength, as listed below:

1. Increasing the transmit voltage (limited by transducer breakdown, see Section 6.2).
2. Increasing the signal duration and utilising a pulse or frequency compression system to retain resolution capability (see Sections 4.6, 8.1 and Chapter 3).

3. Using larger transducers to obtain a narrower beam and hence a larger directivity or beam-forming gain (see Figure 6.9).

The Kleeman and Kuc system simultaneously utilises techniques 1 and 3, while the systems in this thesis concentrate on utilising technique 2. Technique 3 was specifically rejected so as to increase the usable beam width, as discussed below. It is to be noted from the results summarised in Table 16.1 that the maximum range values obtained in this study are similar to those reported in the literature.

16.2.2 Bearing

The determination of bearing is dependent upon the geometry of the target and the geometry of the sensor head as discussed in Section 14.1.1. Due to this dependence some authors, e.g. Kleeman and Kuc and Peremans, have designed their analysis systems to pre-classify the targets under surveillance before determining their bearings. The alternative and less demanding strategy of not pre-classifying the targets and accepting an additional small error was adopted in this thesis. However it has also been shown, in Appendix G, that there exists a head design which minimises the small error. Additionally, restricting the experimental configuration to plane targets, as adopted in the range and bearing precision experiments of Section 13.2 and Chapter 14, also eliminates this error.

Bearing Accuracy

Bearing accuracy estimates can be determined from the experiments made in this thesis due to the absolute nature of the positioner. Of the other systems reported in the literature, only Kleeman and Kuc report bearing accuracy estimates.

Of the five tested systems there is a good agreement between the work of this thesis in the **IDD** method and the work of Kleeman and Kuc. The poor performance of the CTFM-LSMYW system was previously attributed to interference from another target, as described in Section 14.1.4. The **LSMYW** spectral estimator has poor frequency resolution, as documented in Section 13.3.5.

The accuracy of the bearings computed using the **IPD** method is not reported as this method of computing bearing requires calibration, which was performed using the only available data set. Assessment of the accuracy would have required a second data set.

Bearing Precision

Bearing precision determinations are much more variable than bearing accuracies, ranging from a precision estimate of $\pm 0.05^\circ$ (Akbarally) to values near $\pm 1^\circ$ (Yata). The results determined by Akbarally are incidental to his experimental work rather than being determined by an experiment expressly designed to assess precision. Therefore Akbarally's results are discounted in favour of those published by Kleeman and Kuc, who achieved bearing precisions of $\pm 0.1^\circ$ for the **IDD** technique. This result is a factor of 6 to 8 better than that achieved in this thesis. The **IDD** bearing precision discussion in Section 14.1.3 clearly states that the bearing precision is controlled by the precision of the **DOF** measurement and the correlation that exists between the left and right

channels. The superior precision of Kleeman and Kuc's DOFs and their higher correlations of 0.997 compared to 0.8 – 0.9 have in this case added to yield the superior results of Kleeman and Kuc. A more detailed comparison between the Kleeman and Kuc system and the system of this thesis is not possible without a more detailed understanding of the implications of the choice of system parameters. This is more appropriate for further work.

Bearing Resolution

None of the non-beamforming systems are able to resolve targets in bearing when they lie at approximately the same range (see Section 15.5). This is due to the overlap of the echoes from the different targets, which is a range resolution problem. Only beam-forming systems such as Kay's (see Section 2.5.9) have the capability of resolving multiple targets at the same range with different bearings. These beam-forming systems operate by scanning a narrow beam over the sensing region. This improvement in performance is accompanied by a penalty in the time taken to perform the scan.

In the case where the two targets are separated in range by more than the range resolution, the bearing resolution is determined simply by occlusion. This limitation is common to all systems which do not perform beam-forming.

Beam Width

The sonars described in this thesis all have an operational bearing range of $\pm 20^\circ$, as described in Table 16.2. This beam width is larger than that of any other system described in Table 16.2, with the exception of Yata's omni-directional ring sensor. The wide beam widths obtained by the systems described in this thesis were obtained by selecting small transducers with a wide beam pattern, and were designed to provide the mobile robot with a wide sensing area. This reduces the need to scan the operating volume, as described in Section 7.1.2. Thus the design presented in this thesis provides a clear advantage over other systems such as Yata's single transducer (with a $\pm 4^\circ$ beam width) and that of Kleeman and Kuc (with a $\pm 11^\circ$ beam width) or of Peremans (with a $\pm 12^\circ$ beam width).

16.2.3 Spectral Estimates

The range results of Table 16.1 show that the AR class of spectral estimator does indeed have a significantly better performance than the classical DFT approach for range precision, whereas resolution performance is degraded by comparison with the classical DFT approach. It is clearly seen that precision is strongly linked to the regularity of the peak and weakly to the sharpness of the peak whereas resolution is strongly linked to the sharpness of the peak. This is in accordance with Section 2.1. This distinction is not emphasised by Kay and Marple [75].

The bearing results of Table 16.2 again show the significantly improved performance of the AR class compared with the classical DFT approach. What is not immediately apparent is the reason for the improved performance of the LSMYW technique compared with the YW technique, given the similarity of range precision in Table 16.1 (see also Figures 13.4 and 13.5). The difference is attributed to the improved correlation that exists between the left and right channels when

the **LSMYW** technique is used (see Figure 13.13). The improved performance of the **LSMYW**, which uses an over-determined set of **Yule-Walker** equations, is in line with the reported results of Cadzow [20].

While the use of an **AR** based spectral estimator has led to improved precision, Table 16.1 shows that the **AR** based spectral estimators have poor resolution. This was not appreciated in this study until quite recently, when the differing abilities of the various spectral estimators described by Kay and Marple [75] were examined in the light of experimental results.

The differences between the **YW** and the **DFT** spectral estimators are clear. As discussed in Chapter 11, the **YW** estimators assume that the process producing the samples is stochastic, and that its statistical properties do not change as a function of time. This assumption is used to extend the sequence past the finite number of experimental data points available. In contrast, the **DFT** estimator, discussed in Chapter 10, assumes a periodic extension of the data points. Thus the different techniques each make different assumptions about the data, and fail in different ways when the data fails to satisfy the different assumptions. It is noted that the **CTFM** data used in this thesis have an amplitude profile (see Figure 9.16), and that this will violate the stochastic assumption of the **YW** estimator.

16.2.4 Summary

The comparison of the results, mainly accuracies and precisions of range and bearing, of the five systems studied in this thesis and those to be found in the literature, indicates that this study is near the cutting edge of the science. The comparative study has invariably been with the impulse system of Kleeman and Kuc, who forsake the common radar and sonar mantra of constructing a signal with a large **time-bandwidth product** so as to obtain the required **signal to noise ratio**. Instead, they use a signal which is as short as possible. The required **signal to noise ratio** is obtained by driving the transmit transducer with a very large step voltage, and utilising larger transducers with a higher directivity (see Section 6.5). These design choices have come at the price of a reduced sensing angle when compared with the systems investigated in this thesis.

De Roos [29] clearly states that the accuracy of frequency of the echo in a **CTFM** signal is determined by the properties of the spectrum analysis and not the transmitted signal or the received signal. This statement neglects the interfering effects of air or water turbulence and noise in the electronics to focus on the signal type, being time or frequency domain. It is this need for accuracy that suggests that more work needs to be undertaken in understanding the performance characteristics of spectrum analysis.

The similarity of the results in this thesis for the **MF** and **CTFM** class of sonars, and their close agreement with their theoretical values as deduced using the **time-bandwidth product** theory of Section 4.6.3, is also something that de Roos noted. De Roos explicitly states that the differences are caused by practical limits on the realisation of the system components rather than theoretical considerations.

16.3 Limitations

The experiments revealed several limitations of the sonar systems under investigation.

The resolution of the CTFM sonar systems with Yule-Walker and least squares modified Yule-Walker spectral estimators was found to be poor compared with the matched filter and the CTFM-DFT sonar systems. This is clearly due to a deficiency in the spectral estimator. Steven Kay and Lawrence Marple [75, Section IIE] suggest the cause is that Yule-Walker estimators fail for sinusoids in noise, as the YW estimator cannot model the noise. Furthermore, targets which are separated by much more than the resolution limit also interfere with the estimation of echo power.

The work in Chapters 12 and 13 demonstrates that the variability of the speed of sound is a significant factor for precise sonar systems. Chapter 12 demonstrates that the speed of sound can vary considerably over a relatively short space of time in a typical well ventilated air conditioned environment. Chapter 13 discusses several effects, chief of which is the impulse response of air due to absorption, which also affect precise sonar systems. The disagreement between modelled values of the speed of sound and those measured by precise sonar systems currently limits the accuracy obtainable by the precise sonar systems. It is hypothesised that a self-calibrating system would negate these effects but this is left for future work.

16.4 Discussion of Contributions

The aim of this thesis was to build an air sonar sensor system to enable a *robotic agent* to function in place of a *living agent*. In meeting the aim, this thesis makes several contributions to knowledge, which can be summarised as follows:

- A robust statistical quantification and comparison of the accuracy, precision and resolution of five sonar systems in range and bearing has been performed. Part IV of this thesis, the Range and Bearing Sub-System, determines both the mean and standard deviation of all experimentally determined quantities. These experimentally determined values are then rigorously propagated to determine the precision of the estimated range and bearing. This rigorous propagation has shed new light on the relative importance of both IDD and IPD estimates of bearing and the weighted average of these estimates.
- A robust and reliable method of solving the correspondence problem has been demonstrated. The development of a rigorous statistical basis for solving the correspondence problem in addition to the DOF window approach, as undertaken in Chapter 15, is considered essential for a sonar with multiple receivers operating in a real environment. It is expected that the need to solve the correspondence problem will increase with the size of the area insonified.
- An in-depth study was made of both classical and auto-regressive spectral analysis techniques as applied to CTFM sonar systems for the extraction of distance of flight and echo power data. Previous studies have concentrated on extracting echoes from CTFM signals using DFT spectral analysis. Chapter 10 of this thesis demonstrates that the DFT is unsuited to analysing CTFM signals and is particularly unsuited to the estimation of target bearing. The YW and LSMYW methods of spectral estimation, discussed in Chapter 11, have been shown to provide more precise results and better bearing estimates, improving upon those available using matched filter sonar systems.

- A method of self-calibrating the speed of sound continuously from temperature observations was developed. This important contribution made in Chapters 12 and 13 improves both the accuracy and precision of all the studied sonar systems. Significantly, it allows all systems to operate with greater reliability and precision under a wide variety of conditions.

16.5 Further Work

Some items of further work remain. These include:

1. Establishing the cause of the difference between the modelled speed of sound and the sonar measured speed of sound.
2. Reducing the errors in the range and bearing, described in Section 14.1.2, due to not classifying the target type. This may be achieved by re-designing the sonar head so that the transmitter is in-line with the receivers. The design is detailed in Appendix G.
3. Improving the [matched filter IDD](#) bearing precision by reducing the patterning in the bearing residuals. The patterning is caused by [MF DOF](#) perturbation with target bearing, and it should be possible to fit and compensate for this.
4. Improving the [CTFM IDD](#) bearing precision by improving the range precision and the correlation between the left and right channels.
5. Changing the [Yule-Walker](#) spectral estimator for another type which can handle noise. Steven Kay and Lawrence Marple [75, Section IIE] indicate that [Yule-Walker](#) spectral estimators, which apply an [auto-regressive](#) model to the data, do not give accurate results when applied to sinusoids in noise. A [moving average](#) component is required to model the noise. Thus an [auto-regressive moving average](#) model would be required. Such a model is more complex than the [Yule-Walker](#) technique undertaken in this thesis. However, it also has poor resolution properties.

Part VI

Appendices

Appendix A

Transducer Models

This appendix discusses three different methods of modelling the transducer. The first is the transduction model, which models the physical transformations of energy from one form to another with equivalent circuits, which may then be analysed using the standard techniques of electrical engineering. This model is difficult to apply and thus only a brief overview is presented. The second model is the radiating plane piston model, which relates the pressure wave in the far field to the movement at the surface of the transducer. The third model is the Kuc-Siegel impulse model, which is derived similarly to the radiating plane piston model, but provides results in the time domain rather than the frequency domain. Upon taking the Fourier transform, the impulse model is shown to be equivalent to the radiating plane piston model.

A.1 Transduction Model

The transduction model relates the output sound pressure wave to the input electrical signal by modelling the intermediate stages of the energy conversion process as an electrical circuit.

The input of the electrical circuit will be in electrical units. The transducer then converts this electromotive power into mechanical energy, so there is a circuit in mechanical units. Finally, the transducer couples the vibrations into the air, so there is a circuit in acoustic units, with pressure and air velocity. The transduction model unifies these three separate domains into one single model, an equivalent circuit, which may be analysed using the well-known electrical circuit techniques.

The transduction model shown here treats the transducer diaphragm and the air individually as point masses, whereas in reality they are distributed over an area. Thus, the model is called a lumped parameter model. This is a simplifying assumption which allows the circuit analogy to apply.

The analogies between the different domains are shown in Table A.1. Martin [90, Chapter 11] derives a generalised transduction model for grooved backplate transducers. Martin goes into considerable detail concerning the derivation. Martin does not present any experimental results to indicate that he has tested the model as a complete unit, although he mentions some experimental results through the derivation. The main difficulty in applying this model is in attributing numerical values to all the various components of the resulting circuit (Figure A.1). The variables and their units are described in Table A.2.

Martin describes the gauze screen that he was using on his transducers to mechanically protect

Table A.1: Analogies between electrical, mechanical and acoustic domains in transduction analysis.

Electrical	Mechanical	Acoustic
voltage V	force N	sound pressure Pa
current A	velocity $\text{m} \cdot \text{s}^{-1}$	volume velocity $\text{m}^3 \cdot \text{s}^{-1}$
capacitance F	compliance $\text{s}^2 \cdot \text{kg}^{-1}$	acoustic compliance $\text{m}^3 \cdot \text{Pa}^{-1}$
inductance H	mass kg	acoustic inertance $\text{kg} \cdot \text{m}^{-4}$
resistance Ω	viscous resistance $\text{kg} \cdot \text{s}^{-1}$	viscous acoustic resistance $\text{Pa} \cdot \text{s} \cdot \text{m}^{-3}$

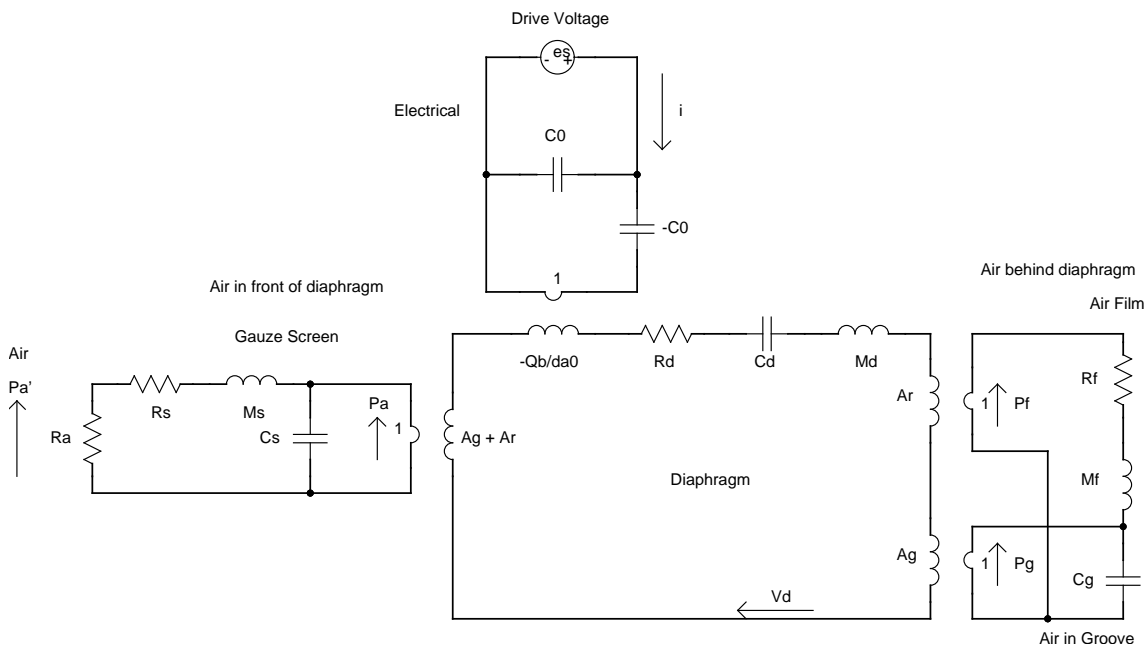


Figure A.1: The complete equivalent circuit of the grooved backplate transducer. The top circuit is the electrical circuit. The middle circuit describes the mechanical movement of the diaphragm. The right circuit describes the air behind the diaphragm, while the left circuit describes the air in front of the diaphragm. This circuit also includes the effect of a gauze screen protecting the transducer, as seen in Figure 2.7. The coupled coils are ideal ratio changers, with ratios as marked. Arrows indicate the direction of conventional current through a component. The variables and units are described in Table A.2.

Table A.2: Description of the variables used in the transduction equivalent circuit of a grooved backplate transducer (Figure A.1).

Variable	Meaning	Units
P'_a	Air pressure outside screen	Pa
R_a	Air load viscous acoustic resistance	$\text{Pa} \cdot \text{s} \cdot \text{m}^{-3}$
R_s	Screen viscous acoustic resistance	$\text{Pa} \cdot \text{s} \cdot \text{m}^{-3}$
M_s	Screen acoustic inertance	$\text{kg} \cdot \text{m}^{-4}$
C_s	Compliance of air in cavity	$\text{m}^3 \cdot \text{Pa}^{-1}$
P_a	Pressure of air in cavity	Pa
A_g	Area of diaphragm over grooves	m^2
A_r	Area of diaphragm over rails	m^2
Q_b	Charge per unit area due to bias voltage	$\text{C} \cdot \text{m}^{-2}$
d_{a0}	Equivalent air gap distance	m
R_d	Diaphragm viscous resistance	$\text{kg} \cdot \text{s}^{-1}$
C_d	Inverse diaphragm restoring force (tension)	$\text{m} \cdot \text{N}^{-1}$
M_d	Mass of diaphragm	kg
V_d	Velocity of diaphragm	$\text{m} \cdot \text{s}^{-1}$
R_f	Air film (on rail) viscous acoustic resistance	$\text{Pa} \cdot \text{s} \cdot \text{m}^{-3}$
M_f	Diaphragm acoustic inertance	$\text{kg} \cdot \text{m}^{-4}$
P_f	Air pressure in air film above rail	Pa
C_g	Air in groove acoustic compliance	$\text{m}^3 \cdot \text{Pa}^{-1}$
P_g	Air pressure in groove	Pa
C_0	Electrical transducer static capacitance with air gap of d_{a0}	F
e_s	Drive voltage	V
i	Drive current	A

them, as seen in Figure 2.7. His modelling indicates that if the screen density and distance from the diaphragm is chosen correctly, then a 2 – 4 dB increase in output can be obtained, due to resonance in the cavity between the diaphragm and the screen. Note that there is no gauze on the transducers used in this thesis, so R_s , M_s and C_s will be zero.

Reversibility of the transduction model

If we replace the voltage source e_s with a high-resistance load, we may model the voltage generated when a sound pressure wave is applied to the transducer. The equivalent circuit shown in Figure A.1 is therefore completely reversible. Since the transduction model is based upon the physical operation of the transducer, we can say that the transducer may also be used as a microphone.

Consequences of the transduction model

Examination of the equivalent circuit in Figure A.1 will show that the transducer will not pass all frequencies equally. The frequency response of this circuit will model the frequency response due to the transducer's operation.

No attempt has been made to apply this model in this thesis, primarily due to the difficulty of finding values for all of the elements in the equivalent circuit of Figure A.1.

A.2 Radiating Plane Piston Model

The pressure field at a distance from the transducer may be modelled using the radiating plane piston model, described by Kinsler [78, Section 8.8]. The model assumes that the entire surface of the transducer moves rigidly and in parallel by the same amount. This is not necessarily true for an electrostatic transducer, as the diaphragm will not move at all toward the edges where it is pinned, nor over the large ventilation grooves. However, the model predicts features observed by experiment, and is useful as a tool for understanding the formation of the beam. The model also assumes that the transducer surface is embedded within an infinite baffle. The baffle need not be infinite; it is sufficient that it be much larger than the wavelength of the sound being modelled. The baffle forces the sound field to be radiated entirely in a forward direction, and doubles the amplitude of the pressure field and the power output (Kinsler [78, Section 8.4]).

This derivation is for continuous single frequency signals. We wish to find the pressure at position $\{r, \theta\}$, at time t (Figure A.3). The dynamic pressure, p , at the position, $\{r, \theta\}$, is the sum of the contributions from all the points on the transducer surface. Each point on the surface of the transducer may be regarded as a simple baffled source. The simple source is assumed to have a displacement amplitude which is much smaller than the wavelength of the sound wave in the medium.

The mathematics of the derivation is simplified by working in the complex domain. Thus we define a complex pressure, p_z . The dynamic physical pressure, p_d , shall be obtained from the (dynamic) complex pressure, p_z , by taking the real part, i.e. $p_d = \text{Re}\{p_z\}$. The complex pressure radiation pattern at the position $\{r, \theta\}$, due to a simple baffled source located at the origin, is




Figure A.2: Sound waves radiating from a baffled point source. The axes are x and y , in metres. The plot is evaluated from $p_d(r, \theta, t) = \text{Re}\{p_z(r, \theta, t)\}$, using (A.1). The parameters were $f = 50\text{kHz}$ and $c = 346\text{m} \cdot \text{s}^{-1}$, producing a wavelength of 7 mm. As the source radiates eternally, the field already exists at $t = 0\mu\text{s}$. The spreading loss component of (A.1) is seen as a decrease in wave amplitude toward the right of the plot. The plot is saturated near the source. The travelling wave component produces the movement in the wave-fronts from one frame to the next.

defined by

$$p_z(r, t) = j\rho_0 c \frac{Qk_w}{2\pi r} e^{j(\omega t - k_w r)}. \quad (\text{A.1})$$

The components of this equation are listed below.

- j denotes the imaginary number $\sqrt{-1}$.
- ρ_0 is the volume density of the material ($1.21\text{ kg} \cdot \text{m}^{-3}$ for air).
- c is the speed of sound in the medium in $\text{m} \cdot \text{s}^{-1}$.
- Q is the source strength, which is the volume of fluid displaced by the sound source per second, measured in $\text{m}^3 \cdot \text{s}^{-1}$.
- k_w is the inverse wavelength in $\text{rad} \cdot \text{m}^{-1}$ ($k_w = 2\pi f/c$).
- r is the radius in metres, from the simple source, to the point where the complex pressure, p_z , is evaluated.
- ω is the frequency of the sound in $\text{rad} \cdot \text{s}^{-1}$ ($\omega = 2\pi f$).
- t is the time, in seconds.
- f is the frequency of the sound wave in Hz.

There are two important components in (A.1). The first is the spherical spreading loss, of the form $1/r$. The second is the travelling wave component, $\exp\{j(\omega t - k_w r)\}$. The remaining terms control the amplitude of the wave. The real valued dynamic sound pressure sound field, $p_d(r, \theta, t)$, radiated from a simple source is shown in Figure A.2. The field only exists in the right half plane, as the baffle prevents it from moving to the left.




Figure A.3: Geometry used in deriving the far-field radiation characteristics of a circular plane piston. The angle between the transducer normal and the ray to the point $\{r, \theta\}$ is denoted θ . The point $\{r, \theta\}$ is at radius r from the centre of the transducer. The point $\{r, \theta\}$ is constrained, without loss of generality, to lie in the x - z plane. The dark region on the transducer surface denotes an infinitesimal strip, whose centre point is a distance r' from the point $\{r, \theta\}$. The infinitesimal strip lies parallel to the y axis, and the end of it subtends an angle ϕ to the z axis.

We wish to calculate the far-field sound pressure at a point $\{r, \theta\}$ due to a vibrating circular piston, as shown in Figure A.3. Without loss of generality, we restrict the point at which we are calculating the pressure to have no y component, so that it lies in the x - z plane.

The pressure at the point $\{r, \theta\}$ is calculated by considering the influence of each part of the transducer surface in turn. For computational convenience this is done by dividing the face of the piston into infinitesimal strips, themselves composed of infinitesimal simple baffled sources as defined by (A.1). Each such strip source is centred on the z axis, and has an infinitesimal source strength of

$$dQ = 2U_0 a \sin(\phi) dz, \quad (\text{A.2})$$

where

- U_0 is the speed amplitude of the vibrating surface in $\text{m} \cdot \text{s}^{-1}$,
- a is the radius of the transducer,
- ϕ is the angle between the infinitesimal element and the z axis, and
- dz is the height of the infinitesimal strip, as shown in Figure A.3.

The incremental pressure from each such infinitesimal strip source can be found by substituting the infinitesimal source strength (A.2) into the equation for the simple baffled source (A.1):

$$dp_z = j\rho_0 c \frac{U_0}{\pi r'} k_w a \sin(\phi) e^{j(\omega t - k_w r')} dz \quad (\text{A.3})$$

where r' is the distance from the point $\{r, \theta\}$ to each infinitesimal element making up the strip.

The exact form for r' is

$$r' = \sqrt{(r \cos \theta)^2 + (a \sin \phi)^2 + (r \sin \theta - a \cos \phi)^2} = \sqrt{a^2 + r^2 - 2ar \sin \theta \cos \phi}.$$

For points $\{r, \theta\}$ in the far field, defined by $r > a^2/\lambda$ (where $\lambda = c/f$, measured in metres), Kinsler [78] gives the approximate form

$$r' \approx r + \Delta r = r \cdot \left(1 - \frac{a}{r} \sin \theta \cos \phi\right).$$

Integrating the contributions from each of the infinitesimal strips (A.3) to form the entire contribution to the pressure field at a distant point $\{r, \theta\}$ yields

$$p_z(r, \theta, t) = j\rho_0 c \frac{U_0}{\pi r} k_w a e^{j(\omega t - k_w r)} \int_{-a}^a e^{jk_w a \sin \theta \cos \phi} \sin \phi dz.$$

This has been simplified by applying the far field approximation $r' \approx r$ in the spreading loss term, but retaining r' containing the phase information in the exponential.

Changing the variable of integration to eliminate the vertical component, z , using the relationships $z = a \cos \phi$ and $dz = -a \sin \phi d\phi$, yields

$$p_z(r, \theta, t) = j\rho_0 c \frac{U_0}{\pi r} k_w a^2 e^{j(\omega t - k_w r)} \int_0^\pi e^{jk_w a \sin \theta \cos \phi} \sin^2 \phi d\phi. \quad (\text{A.4})$$

The minus sign has been cancelled by a reversal of the range of integration. At this point it is convenient to write $k_w a \sin \theta = b$. The integral part of (A.4), requires special treatment, so it is defined separately

$$p_z(r, \theta, t) = j\rho_0 c \frac{U_0}{\pi r} k_w a^2 e^{j(\omega t - k_w r)} p_z(\theta) \quad (\text{A.5})$$

$$p_z(\theta) = \int_0^\pi e^{jb \cos \phi} \sin^2 \phi d\phi \quad (\text{A.6})$$

Turning first to the imaginary part of (A.6),

$$\text{Im}\{p_z(\theta)\} = j \int_0^\pi \sin(b \cos \phi) \sin^2 \phi d\phi,$$

we split the range of integration in half to give

$$\text{Im}\{p_z(\theta)\} = j \int_0^{\frac{\pi}{2}} \sin^2 \phi \cdot \sin(b \cos \phi) d\phi + j \int_{\frac{\pi}{2}}^\pi \sin^2 \phi \cdot \sin(b \cos \phi) d\phi.$$

Now, changing variables $\phi = \pi - \phi'$ in the second integral yields

$$\text{Im}\{p_z(\theta)\} = j \int_0^{\frac{\pi}{2}} \sin^2 \phi \cdot \sin(b \cos \phi) d\phi - j \int_{\frac{\pi}{2}}^0 \sin^2 \phi' \cdot \sin(-b \cos \phi') d\phi'.$$

Taking the negation out of the sin function (which is odd) and reversing the range of integration yields

$$\text{Im}\{p_z(\theta)\} = j \int_0^{\frac{\pi}{2}} \sin^2 \phi \cdot \sin(b \cos \phi) d\phi - j \int_0^{\frac{\pi}{2}} \sin^2 \phi' \cdot \sin(b \cos \phi') d\phi'.$$

This forms the difference between two equal terms, and is equal to zero. Thus the imaginary part of the exponential integral (A.6) is also zero.

Turning now to the real part of (A.6), integral tables (Abramowitz and Stegun [1, 9.1.20]) show that

$$\text{Re}\{p_z(\theta)\} = \int_0^\pi \cos(b \cos \phi) \sin^2 \phi d\phi = \frac{\pi J_1(b)}{b}, \quad (\text{A.7})$$

where $J_1(x)$ is a Bessel function of the first kind. Equation (A.7) does not evaluate formally at $b = 0$, but evaluating the limit of (A.7) as $b \rightarrow 0$ yields the value 1 as the result. Substituting $b = k_w a \sin \theta$ again, the final equation for the complex pressure, $p_z(r, \theta, t)$, in the far field is

$$p_z(r, \theta, t) = \begin{cases} j\rho_0 c \frac{U_0}{r} k_w a^2 e^{j(\omega t - k_w r)} \left[\frac{2J_1(k_w a \sin \theta)}{k_w a \sin \theta} \right] & \text{if } 0 < |\theta| \leq \frac{\pi}{2} \\ j\rho_0 c \frac{U_0}{r} k_w a^2 e^{j(\omega t - k_w r)} & \text{if } \theta = 0 \end{cases} \quad (\text{A.8})$$

The real dynamic pressure, $p_d(r, \theta, t)$, is obtained from the real part, $\text{Re}\{p_z(r, \theta, t)\}$.

A close examination of (A.8) shows a factor of $ck_w = 2\pi f$, implying that the amplitude of the pressure wave increases with frequency.

While the transducer beam pattern has been derived by treating the transducer as a transmitter, the principle of acoustic reciprocity (Kinsler [78, Section 8.3]) allows us to say that the same transducer, when operated as a receiver, will have the same beam pattern.

Further implications of (A.8) are discussed in Section 6.3.2.

A.3 Kuc-Siegel Impulse Model

The impulse model is derived by Kuc and Siegel [85]. The version reproduced here largely follows Kuc and Siegel's work, but a factor of $\cos \theta$ has been removed by this author, for reasons discussed later in this section. The derivation is from the point of view of a receiver. However, the principle of acoustic reciprocity guarantees that a transmitter will operate in the same way.

We seek to model the transducer and the air that it couples to as a linear system. That is, for an input signal $s(t)$, we wish to derive the dynamic pressure, $p_d(r, \theta, t)$, at the point $\{r, \theta\}$. According to linear system theory, we should expect a result of the form

$$\begin{aligned} p_d(r, \theta, t) &= \int_{-\infty}^{+\infty} h(r, \theta, \tau) \cdot s(t - \tau) d\tau \\ &= h(r, \theta, t) * s(t). \end{aligned} \quad (\text{A.9})$$

where $h(r, \theta, t)$ is the impulse response of the system.

When deriving this relationship, it is convenient to work with impulses, $\delta(t)$. Impulses have the following properties:

$$\delta(t) = 0 \text{ for all } t \neq 0$$

$$\int_{-\infty}^{+\infty} \delta(t) dt = 1$$

The value of $\delta(t)$ at $t = 0$ is undefined, and the integral only exists as the limit of a series of generalised functions which approximate the impulse function, $\delta(t)$ (Bracewell [14]). The Fourier transform of an impulse contains all frequencies - it is an infinite bandwidth signal. It is useful to consider the case where the input signal $s(t)$ is an impulse $\delta(t)$, as the transducer system response is

$$h(r, \theta, t) * \delta(t) = \int_{-\infty}^{+\infty} h(r, \theta, \tau) \cdot \delta(t - \tau) d\tau = \int_{-\infty}^{+\infty} h(r, \theta, t - \tau') \cdot \delta(\tau') d\tau' = h(r, \theta, t),$$

where a change of variables, $\tau' = t - \tau$, has been used. Using the definition of the impulse, the integral can be removed leaving us with just the beam-forming response. This is the reason that $h(r, \theta, t)$ is called the impulse response - it is the response of the system to an impulse input.




Figure A.4: A point transmitter, located at the end of the x axis, and a single point receiver, located at the origin. The graph indicates the received signal when an impulse is generated by the transmitter and travels a distance r to the receiver.

Huygen's principle is valuable in deriving the impulse response of a transducer's radiation pattern. Huygen's principle (from Hecht [61]) states that

... every point on a primary wavefront serves as the source of spherical secondary wavelets, such that the primary wavefront at some later time is the envelope of the wavelets. Moreover, the wavelets advance with a speed and frequency equal to those of the primary wave at each point in space.

Huygen's principle is just that - a principle. It does not exactly satisfy the travelling wave equation. However, it suits the current derivation, by allowing us to divide the surface of the transducer into individual elements each much smaller than the wavelength of sound which is being considered. Since we are considering an impulse as our signal, which contains infinitely large and infinitely small wavelengths, we shall use differential elements in the derivation.

A.3.1 Single Point Receiver

The simplest case to begin with is that of two small point transducers, with one transmitting to the other. These transducers lie along the x axis. (Figure A.4.) When an impulse is generated from the transducer at the left, it will be received some time later by the transducer at the origin. The travel time is the travel distance divided by the speed of sound, r/c . The received signal may be modelled by

$$s(r, t) = \frac{A}{r} \delta\left(t - \frac{r}{c}\right), \quad (\text{A.10})$$

where A is the signal amplitude and has units of $\text{V} \cdot \text{m}$.




Figure A.5: A point transmitter and two point receivers. The graph indicates the received signal when an impulse is generated by the transmitter, located at a distance r along the x axis, and travels to the receivers centred about the origin.

A.3.2 Rotated Two Point Receiver

Now let us develop this model by using two point transducers instead of one. Let these point transducers be separated by a distance $2a$ and rotated from the y axis by θ degrees, as shown in Figure A.5. An electrical amplifier connected to both of these elements in parallel will detect two impulse signals, as shown on the graph in the figure. Note that this computation has assumed that $r > a^2/\lambda$, so that the distance from the transmitter to the receiver can be approximated by the distance from x to the origin minus the x component of the receiver's tilt angle. This is the far field approximation again. Thus, the receiver will display two impulses at the times labelled on the graph.

The detected signal will be

$$s(r, \theta, t) = \frac{A}{r} \left\{ \delta \left(t - \frac{r - a \sin(\theta)}{c} \right) + \delta \left(t - \frac{r + a \sin(\theta)}{c} \right) \right\}, \quad (\text{A.11})$$

where A is the signal amplitude, and has units of $\text{V} \cdot \text{m}^{-1}$ (ignoring the acoustical units of the radiating plane piston model, and instead concentrating upon the signal). The far field approximation ($x \gg a$) has been used in the spreading loss term, and A is a constant representing the sensitivity of the transducer.

A.3.3 On-Axis Transducer

The entire surface of a transducer can be considered to be composed of these point elements. We can find their combined effect by integrating over all the elements. To begin with, we consider a transducer which is aligned with the x axis (Figure A.6). We will assume that the transmitter and




Figure A.6: A disk, with its normal aligned to the x axis, and its surface divided into infinitesimal elements. The graph indicates the impulse which is received from a single transmitter element, located a distance r along the x axis, in the far field.

receiver are sufficiently far apart that the wave is flat across the receiving elements. Integrating over the surface of the transducer, where the variables y and z correspond to the axes of the same names, yields

$$s(r, t) = \frac{A}{r} \int_{-a}^a \int_{-\sqrt{a^2-y^2}}^{\sqrt{a^2-y^2}} \delta\left(t - \frac{r}{c}\right) dz dy. \quad (\text{A.12})$$

Performing the integral yields

$$s(r, t) = \frac{A}{r} \pi a^2 \delta\left(t - \frac{r}{c}\right). \quad (\text{A.13})$$

The units of A must be converted to $\text{V} \cdot \text{m}$ to provide the correct output. Here we see that a circle of infinitesimal elements produces a signal which is stronger in proportion to its area than a single point receiver would be.

A.3.4 Rotated Transducer

When the normal of the receiver is rotated by an angle θ from the x axis (Figure A.7), the derivation can be adapted as follows. The impulse will sweep across the surface of the receiver, contacting the transducer along a vertical line at each point in time. Incorporating this feature and integrating over the elements of the surface yields the signal detected by the receiver,

$$s(r, \theta, t) = \frac{A}{r} \int_{-a}^a \int_{-\sqrt{a^2-y'^2}}^{\sqrt{a^2-y'^2}} \delta\left(t - \frac{r + y' \sin(\theta)}{c}\right) dz dy', \quad (\text{A.14})$$

where the z variable runs vertically across the surface, and the y' variable runs horizontally across the transducer surface.

Kuc and Siegel [85] include the factor $\cos \theta$ in their equivalent of equation (A.14). This term is intended to represent the directional property of each differential element of the flat transducer




Figure A.7: A disk with its normal at an angle θ to the x axis, and its surface divided into infinitesimal elements. There is an angle θ between the transducer and the y axis. The graph shows the signal that will be detected when an impulsive signal is transmitted by the point transmitter.

surface. This is inappropriate as we are calculating the pressure at a point due to the change in pressure at a point source. Fluids are unable to sustain directional pressures or shear forces (Halliday [51]), so the inclusion of this term in a pressure calculation is inappropriate. Furthermore, including the $\cos \theta$ factor produces a result which disagrees with the radiating plane piston model, and also with experiment. This may be seen by examining the experimentally determined beam pattern in Figure 6.15a. The $\cos \theta$ term would require the beam pattern to go to zero as the angle goes to $\pm 90^\circ$, thereby contradicting the experimental data. For these reasons, the $\cos \theta$ factor was removed.

Performing the integration over dz in equation (A.14) yields

$$S(r, \theta, t) = \frac{2A}{r} \int_{-a}^a \sqrt{a^2 - y'^2} \cdot \delta \left(t - \frac{r + y' \sin(\theta)}{c} \right) dy'. \quad (\text{A.15})$$

We now perform a substitution of variables, $y' = (c \cdot (t - t') - r) / \sin \theta$, which will serve to simplify the expression inside the impulse function. The derivative of the substitution is $dy/dt' = -c / \sin \theta$, and the new limits of integration must be substituted using $t' = t - (r + y' \sin \theta) / c$. Performing these substitutions yields

$$\begin{aligned} s(r, \theta, t) &= \frac{2A}{r} \int_{t - \frac{r+a \sin \theta}{c}}^{t - \frac{r-a \sin \theta}{c}} \sqrt{a^2 - \left(\frac{c(t-t')-r}{\sin \theta} \right)^2} \delta \left(t - \frac{r + \left(\frac{c(t-t')-r}{\sin \theta} \right) \sin \theta}{c} \right) \frac{-c}{\sin \theta} dt' \\ &= \frac{2Aac}{r} \frac{1}{\sin \theta} \int_{t - \frac{r+a \sin \theta}{c}}^{t - \frac{r-a \sin \theta}{c}} \sqrt{1 - \left(\frac{ct - ct' - r}{a \sin \theta} \right)^2} \delta(t') dt'. \end{aligned}$$

The integral has 3 important cases:

- 1: $t < \frac{r-a \sin \theta}{c}$
- 2: $\frac{r-a \sin \theta}{c} \leq t \leq \frac{r+a \sin \theta}{c}$
- 3: $\frac{r+a \sin \theta}{c} < t$

In cases 1 and 3, the range of the integral does not intersect with the position of the impulse, so the integral becomes 0. This corresponds to a time when the impulse, travelling through the air, does not meet with the surface of the receiver at all. In case 2, however, the impulse intersects with the receiver. The integral can be evaluated using the sifting property of the δ function at $t' = 0$:

$$s(r, \theta, t) = \frac{2Aac}{r} \frac{1}{\sin \theta} \sqrt{1 - \left(\frac{ct - r}{a \sin \theta} \right)^2},$$

where the signal is zero outside the time limits imposed by case 2 above. The amplitude term has units of $V \cdot m^{-1}$. When compared to the results of Kuc and Siegel [85], the results shown here have an extra factor of πa^2 . This is because this formula takes into account the area of the transducer, whereas Kuc and Siegel have normalised the area of the receiver to 1.

Appealing to the principle of reciprocity, we may now state the impulse response of the pressure at a point due to an impulse on the surface of the transducer:

$$h(r, \theta, t) = \begin{cases} \frac{2Aac}{r} \frac{1}{\sin \theta} \sqrt{1 - \left(\frac{ct - r}{a \sin \theta} \right)^2} & \text{if } \frac{r - a \sin \theta}{c} \leq t \leq \frac{r + a \sin \theta}{c}, 0 < |\theta| \leq \frac{\pi}{2} \\ \frac{A}{r} \pi a^2 \delta\left(t - \frac{r}{c}\right) & \text{if } \theta = 0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.16})$$

At this point it is useful to separate the impulse response into two parts,

$$h(r, \theta, t) = h_{rad}(r, t) * h_{\theta}(\theta, t), \quad (\text{A.17})$$

where $h_{rad}(r, t)$ represents the wave travelling through the air, and $h_{\theta}(\theta, t)$ represents the interference pattern of the transducer. These components are defined by

$$h_{rad}(r, t) = \frac{1}{r} \delta\left(t - \frac{r}{c}\right) \quad (\text{A.18})$$

$$h_{\theta}(\theta, t) = \begin{cases} \frac{2Aac}{r} \frac{1}{\sin \theta} \sqrt{1 - \left(\frac{ct}{a \sin \theta} \right)^2} & \text{if } \frac{a}{c} \sin \theta \leq t \leq \frac{a}{c} \sin \theta, 0 < |\theta| \leq \frac{\pi}{2} \\ \frac{A}{r} \pi a^2 \delta(t) & \text{if } \theta = 0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.19})$$

This separation may be verified by substituting (A.18) and (A.19) into (A.17) and performing the convolution. The bearing component, $h_{\theta}(\theta, t)$, is non-causal but that shall always be compensated by using $h_{\theta}(\theta, t)$ with $h_{rad}(r, t)$, and by realising that (A.16) is only valid in the far field.

A.3.5 Comparison with the Radiating Plane Piston Model

It is useful to compare (A.16) to the radiating plane piston model, which is similar to this derivation. While the radiating plane piston model operates in the frequency domain, the Kuc-Siegel impulse model operates in the time domain. To compare the two, we utilise the Fourier transform, as shown in Figure 4.1. We treat only the first case of the bearing component in equation (A.19),




Figure A.8: Diagram showing graphically the geometry used in the change of variables. The plane is the impulse impinging upon the receiver at time ct . The angle ϕ is the angle to the intersection of the plane and the transducer circumference.

as it converges to the correct answer in the limit as $\theta \rightarrow 0$. Taking the Fourier transform of the bearing component of the impulse response,

$$H_{\theta}(\theta, \omega) = \int_{-\infty}^{+\infty} h_{\theta}(\theta, t) \cdot e^{-j\omega t} dt \quad (\text{A.20})$$

$$= 2Aac \frac{1}{\sin \theta} \int_{-\frac{a}{c} \sin \theta}^{\frac{a}{c} \sin \theta} \sqrt{1 - \left(\frac{ct}{a \sin \theta} \right)^2} \cdot e^{-j\omega t} dt. \quad (\text{A.21})$$

Now we change variables according to $t = \frac{a}{c} \sin \theta \cos \phi$ (Figure A.8). Substituting $dt = -\frac{a}{c} \sin \theta \sin \phi d\phi$ yields

$$H_{\theta}(\theta, \omega) = 2Aac \frac{1}{\sin \theta} \int_{\pi}^0 -\frac{a}{c} \sin \theta \sin \phi a \sqrt{1 - \left(\frac{c}{a \sin \theta} \frac{a}{c} \sin \theta \cos \phi \right)^2} e^{-jk_w a \sin \theta \cos \phi} d\phi.$$

Performing the substitution $-k_w a \sin \theta = b'$ and reversing the limits of integration cancels the negation to give

$$H_{\theta}(\theta, \omega) = 2Aa^2 \int_0^{\pi} \sin \phi \sqrt{1 - \cos^2 \phi} \cdot e^{jb' \cos \phi} d\phi, \quad (\text{A.22})$$

and simplifying further yields

$$H_{\theta}(\theta, \omega) = 2Aa^2 \int_0^{\pi} \sin^2 \phi \cdot e^{jb' \cos \phi} d\phi. \quad (\text{A.23})$$

This integral is the same as (A.6). Substituting the answer and substituting $b' = -k_w a \sin \theta$ yields

$$H_{\theta}(\theta, \omega) = A\pi a^2 \left[\frac{2J_1(-k_w a \sin \theta)}{-k_w a \sin \theta} \right]. \quad (\text{A.24})$$

The Bessel function, $J_1(x)$, is odd, so we may move the negation outside the function and cancel it with the one in the denominator. Although $H_{\theta}(\theta, \omega)$ does not formally exist at $\theta = 0$, the value 1 may be found by evaluating the limit as $\theta \rightarrow 0$. Ignoring the terms which are independent of θ , the frequency response of the bearing component of the Kuc-Siegel impulse model is

$$H_{\theta}(\theta, \omega) = \begin{cases} \left[\frac{2J_1(k_w a \sin \theta)}{k_w a \sin \theta} \right] & \text{if } 0 < |\theta| \leq \frac{\pi}{2} \\ 1 & \text{if } \theta = 0 \end{cases}. \quad (\text{A.25})$$

This result clearly has the same dependency upon θ as the radiating plane piston model (6.1). This shows that the time domain Kuc-Siegel impulse model and the frequency domain radiating plane piston model have equivalent bearing dependencies when the $\cos \theta$ factor is removed from the Kuc-Siegel impulse model.

The travelling wave portion of the radiating plane piston model, which is $\cos(\omega t - k_w r)$ when the real part of the complex pressure, p_z , is taken, is modelled by the travelling impulse signal, $\delta(t - r/c)$, in the Kuc-Siegel impulse model. The spreading loss component, $1/r$, is the same in both models. The main difference is the different normalisations and units which are used. The radiating plane piston model, taken from Kinsler [78], has units of pressure, while the Kuc-Siegel impulse model only examines the signal detected by the transducer. This particular component remains unreconciled between the two models.

Appendix B

Matched Filter Theory

This appendix, composed of two sections. The first provides a proof of the optimal matched filter criteria, while the second discusses the desirable properties of a signal when designing a [matched filter](#) system.

B.1 Derivation of Matched Filter Criteria

Radar, sonar and ultrasonic sensing have many features in common. One critical common feature is the detection of the signal in an environment that has a fixed or variable background noise. Another is the interpretation of this signal to obtain the echo [TOF](#), and the echo power.

This section derives a filter which is tuned to amplify the signal but reject as much noise as possible. Thus it maximises the [Signal to Noise Ratio \(SNR\)](#). The filter is specifically designed to recognise the expected received signal. The [SNR](#) will have its maximum value at the signal arrival time.

While the matched filter criteria was originally demonstrated by North [101], the derivation presented here is due to DiFranco and Rubin [32, Chapter 5].

We wish to find a filter, $h_{opt}(t)$, which will maximise the output signal to noise ratio at time $t = \tau$, the signal arrival time. The input waveform, $y_i(t)$, is composed of signal $s_i(t)$, delayed by the travel time, τ , and modified by the elements of the signal path, as described in Chapter 4. Furthermore, the signal contains additive noise, $n_i(t)$, incurred at various stages of the path. Thus the received signal before processing is

$$y_i(t) = s_i(t - \tau) + n_i(t). \quad (\text{B.1})$$

The travel time, τ , is related to the distance travelled, d , by $\tau = d/c$. The signal output by the matched filter is $y_o(t) = s_o(t) + n_o(t)$, and is found by filtering the input waveform, $y_i(t)$, with the optimal filter, $h_{opt}(t)$, using the convolution formula,

$$\begin{aligned} y_o(t) &= h_{opt}(t) * y_i(t) = \int_{-\infty}^{\infty} h_{opt}(t') y_i(t - t') dt' \\ &= \int_{-\infty}^{\infty} h_{opt}(t') \cdot \{s_i(t - t' - \tau) + n_i(t - t')\} dt', \end{aligned} \quad (\text{B.2})$$

where t' is a variable of integration. The signal and noise outputs of (B.2) may be found separately. The signal output is

$$s_o(t) = \int_{-\infty}^{\infty} h_{opt}(t') s_i(t - t' - \tau) dt' \quad (\text{B.3})$$

and the noise output is

$$n_o(t) = \int_{-\infty}^{\infty} h_{opt}(t') n_i(t-t') dt'. \quad (\text{B.4})$$

The maximum signal to noise power ratio, which we wish to occur at the signal arrival time, τ , is given by

$$\chi_m(\tau) = \frac{s_o^2(\tau)}{E\{n_o^2(t)\}}, \quad (\text{B.5})$$

where $E\{\}$ represents the expectation value of a random variable. The matched filter problem is to find $h_{opt}(t)$ so as to maximise $\chi_m(\tau)$.

To facilitate the derivation, let us define a non-optimal filter, $h_{nopt}(t)$, composed of the optimal filter $h_{opt}(t)$ and a component, $g(t)$,

$$h_{nopt}(t) = h_{opt}(t) + g(t), \quad (\text{B.6})$$

The function $g(t)$ is defined to be orthogonal to the received message signal, $s_i(t)$,

$$\int_{-\infty}^{\infty} g(t') s_i(-t') dt' = 0. \quad (\text{B.7})$$

Thus the amplitude function $g(t)$ is not unique, and any scaled version of $g(t)$ is sufficient to satisfy (B.7). We shall assume that there exists a $g(t) \neq 0$ which satisfies (B.7). Consider the non-optimum filter, $h_{nopt}(t)$, and its non-optimum signal to noise power ratio, $\chi'_m(\tau)$. The maximum signal to noise power ratio, $\chi'_m(\tau)$, will be less than the optimum signal to noise ratio of the optimal filter,

$$\chi'_m(\tau) < \chi_m(\tau). \quad (\text{B.8})$$

The response of a non-optimum filter, $h_{nopt}(t)$, to an input signal $y_i(t)$, will be

$$y'_o(t) = h_{nopt}(t) * y_i(t) = s'_o(t) + n'_o(t),$$

where the output signal has again been divided into its (non-optimal) signal and noise components. The signal to noise ratio, $\chi'_m(\tau)$, at the signal arrival time, τ , of the non-optimal filter, $h_{nopt}(t)$, is found similarly to (B.5),

$$\chi'_m(\tau) = \frac{s_o'^2(\tau)}{E\{n_o'^2(t)\}}. \quad (\text{B.9})$$

The signal component, $s'_o(t)$, of the output at the signal arrival time, τ , will be

$$s'_o(\tau) = \int_{-\infty}^{\infty} [h_{opt}(t') + g(t')] s_i(\tau - t' - \tau) dt' \quad (\text{B.10})$$

$$\begin{aligned} &= \int_{-\infty}^{\infty} h_{opt}(t') s_i(-t') dt' + \int_{-\infty}^{\infty} g(t') s_i(-t') dt' \\ &= s_o(\tau), \end{aligned} \quad (\text{B.11})$$

where the integral involving $g(t)$ vanishes due to the orthogonality relationship between $g(t)$ and $s_i(t)$, defined by (B.7). Thus the non-optimal filter provides the same signal output as the optimal filter. The noise component of the output is

$$\begin{aligned} n'_o(t) &= \int_{-\infty}^{\infty} [h_{opt}(t') + g(t')] n_i(t-t') dt' \\ &= \int_{-\infty}^{\infty} h_{opt}(t') n_i(t-t') dt' + \int_{-\infty}^{\infty} g(t') n_i(t-t') dt' \\ &= n_o(t) + \int_{-\infty}^{\infty} g(t') n_i(t-t') dt'. \end{aligned} \quad (\text{B.12})$$

Comparing the noise output of the non-optimal filter, $n'_o(t)$, to $n_o(t)$ of the optimal filter (B.4), we see that an additional noise component is transmitted through the non-optimal filter. The mean square value of the output noise signal, $n'_o(t)$, as required by (B.5), is found by

$$\begin{aligned} E\{n_o'^2(t)\} &= E\{n_o^2(t)\} + E\left\{2n_o(t) \int_{-\infty}^{\infty} g(t') n_i(t-t') dt'\right\} \\ &\quad + E\left\{\left[\int_{-\infty}^{\infty} g(t') n_i(t-t') dt'\right]^2\right\}. \end{aligned} \quad (\text{B.13})$$

Using the relationship (B.8) between the optimum and non-optimum signal to noise ratios, we may write

$$\chi'_m(\tau) = \frac{s_o^2(\tau)}{E\{n_o'^2(t)\}} < \chi_m(\tau) = \frac{s_o^2(\tau)}{E\{n_o^2(t)\}}. \quad (\text{B.14})$$

Equation (B.11) demonstrated that $s'_o(t) = s_o(t)$, so (B.14) implies that

$$E\{n_o'^2(t)\} > E\{n_o^2(t)\}. \quad (\text{B.15})$$

Combining (B.15) with (B.13) yields

$$E\left\{2n_o(t) \int_{-\infty}^{\infty} g(t') n_i(t-t') dt'\right\} + E\left\{\left[\int_{-\infty}^{\infty} g(t') n_i(t-t') dt'\right]^2\right\} > 0. \quad (\text{B.16})$$

The function $g(t)$ has an arbitrary positive or negative amplitude, and so may cause the first integral in (B.16) to be either positive or negative. For (B.16) to be true for any function $g(t) \neq 0$, the first term of (B.16) must be equal to 0. Expanding this out in full, along with the complete definition of $n_o(t)$ from (B.4) gives

$$E\left\{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(t') h_{opt}(\sigma) n_i(t-t') n_i(t-\sigma) dt' d\sigma\right\} = 0, \quad (\text{B.17})$$

where σ is a variable of integration representing time. The expectation value operator is averaged over the independent variable t , and may thus be confined to the noise factors, $E\{n_i(t-t') n_i(t-\sigma)\}$, which represent the autocorrelation of the input noise,

$$E\{n_i(t-t') n_i(t-\sigma)\} = \phi_{nn}(t'-\sigma). \quad (\text{B.18})$$

This allows (B.17) to be simplified to

$$\int_{-\infty}^{\infty} g(t') \left[\int_{-\infty}^{\infty} h_{opt}(t'-\sigma) \phi_{nn}(t'-\sigma) d\sigma \right] dt' = 0 \quad (\text{B.19})$$

By comparing (B.19) with (B.7), we can see that the bracketed part, being the whole inner integral, corresponds to $s_i(-t)$ in (B.7). Furthermore, (B.19) continues to hold if a scaling constant, k_{mf} , is used, as this does not effect the zero result. Thus the inner integral of (B.19) may be re-written

$$\int_{-\infty}^{\infty} h_{opt}(\sigma) \phi_{nn}(t-\sigma) d\sigma = k_{mf} s_i(-t). \quad (\text{B.20})$$

Equation (B.20) is a convolution between the optimal filter and the noise autocorrelation. Taking the Fourier transform of this equation yields

$$H_{opt}(\omega) \cdot \Phi_{nn}(\omega) = k_{mf} S_i^*(\omega),$$

where $H_{opt}(\omega)$ is the Fourier transform of $h_{opt}(t)$, $\Phi_{nn}(\omega)$ is the transform of the autocorrelation of the stationary noise process, and $S_i^*(\omega)$ represents the complex conjugate of the transform of the input signal $s_i(t)$. The Wiener-Khintchine relation (Therrien [131]) states that $\Phi_{nn}(\omega) = G_{nn}(\omega)$, the spectral density of the stationary noise process. We may now write the spectrum of the optimal filter as

$$H_{opt}(\omega) = \frac{k_{mf} S_i^*(\omega)}{G_{nn}(\omega)}. \quad (\text{B.21})$$

This shows that the optimal filter, viewed in the frequency domain, allocates more weight to the frequencies where the signal to noise ratio is highest. In a system where the noise is white and the noise power is equal to $N_0/2 \text{ W} \cdot \text{Hz}^{-1}$, the noise spectral density will be $\phi_{nn}(t') = N_0 \delta(t')/2$. Thus (B.21) becomes

$$H_{opt}(\omega) = \frac{2k_{mf}}{N_0} S_i^*(\omega) \quad (\text{B.22})$$

or, in the time domain

$$h_{opt}(t) = \frac{2k_{mf}}{N_0} s_i(-t) = k'_{mf} s_i(-t), \quad (\text{B.23})$$

where the constants $2k_{mf}/N_0$ have been combined into a single constant k'_{mf} .

The result given in (B.23) can be related to (B.11), where it was shown that any component of the filter function which was orthogonal to the expected message signal would be removed. Therefore $h_{opt}(t)$ must be linearly related to the temporal reversal of the message signal $s_i(-t)$, as (B.23) shows.

The filter impulse response given in (B.23) can be seen to be non-causal, since it is the temporal reversal of a message signal which is zero for all times less than zero. This may be overcome in a digital system by pre-pending zeros to the output record and using an impulse response record which begins at some negative time for performing the convolution. We shall not concern ourselves further with such details here.

The constant k_{mf} can be evaluated by substituting (B.20) into (B.3), and evaluating at time τ :

$$s_o(\tau) = \frac{1}{k_{mf}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h_{opt}(t') h_{opt}(\sigma) \phi_{nn}(t' - \sigma) dt' d\sigma.$$

Substituting the definition of $\phi_{nn}(t' - \sigma)$ from (B.18), we may expand the expected value operator to cover both of the integrals, as the other terms are constant with respect to the expected value operation. Doing this yields

$$s_o(\tau) = \frac{1}{k_{mf}} E \left\{ \int_{-\infty}^{\infty} h_{opt}(t') \int_{-\infty}^{\infty} h_{opt}(\sigma) n_i(t - t') n_i(t - \sigma) dt' d\sigma \right\}.$$

The integrals may now be separated, by moving components which are constant in one integral into the other, yielding

$$s_o(\tau) = \frac{1}{k_{mf}} E \left\{ \int_{-\infty}^{\infty} h_{opt}(t') n_i(t - t') dt' \cdot \int_{-\infty}^{\infty} h_{opt}(\sigma) n_i(t - \sigma) d\sigma \right\}.$$

The integrals may be recognised as the output noise, $n_o(t)$, given in (B.4). Simplifying yields

$$s_o(\tau) = \frac{1}{k_{mf}} E \{ n_o(t) \cdot n_o(t) \}.$$

Thus k_{mf} is

$$k_{mf} = \frac{E \{n_o^2(t)\}}{s_o(\tau)}. \quad (\text{B.24})$$

Thus we have determined k_{mf} , and using (B.23) we may compute the matched filter scaling constant k'_{mf} , which completes the definition of $h_{opt}(t)$.

The implication of equations (B.23) and (B.24) is that the optimal filter is formed by using the expected signal and time-reversing it, then shifting it so that its left hand side matches the origin. Equation (B.22) shows that the gains of the matched filter are entirely due to suppressing the frequencies where the noise is highest, so as to turn coloured noise into white noise.

The matched filter is supposed to provide a maximum signal to noise power ratio at the echo arrival time, τ . This may be verified by substituting the definition of $h_{opt}(t)$, being (B.23), into the expression for the output signal, $s_o(t)$, from (B.3), obtaining

$$s_o(t) = k'_{mf} \int_{-\infty}^{\infty} s_i(-t') s_i(t - t' - \tau) dt'. \quad (\text{B.25})$$

When $\tau = 0$, the output signal $s_o(t)$ is the autocorrelation of the input signal $s_i(t)$, and will always have a maximum at $t = 0$. Introducing a non-zero travel time, τ , has the effect of moving the autocorrelation peak to the position $t = \tau$. Thus the maximum of the output of the matched filter denotes the precise arrival time of the echo, as desired. Furthermore, the output signal amplitude, $s_o(t)$, is linearly related to the strength of the echo echo power through the scaling constant k'_{mf} .

B.2 Signal Selection

We saw in equation (8.9) that the matched filter output signal may be approximated by the autocorrelation or self correlation of the transmit signal. We now assess the autocorrelations of several different types of signals, and their properties as they relate to the detection of sonar targets. This is done with a view to choosing the best signal.

Some example autocorrelation functions for ideal signals are shown in Figure B.1. The most important feature of the autocorrelation is the central peak. This is used to detect the echoes. The effect of increasing travel time, τ , upon the convolution output, is to shift the autocorrelation patterns seen in Figure B.1 to the right, in proportion to the travel time, τ . In this manner, the travel time of the message signal is determined by finding the position of the peak in the convolution output.

The received noise, $n_i(t)$, is also processed by the optimal filter, $h_{opt}(t)$. The filtered output noise, $n_o(t)$, which is derived in Appendix B.1, will also produce peaks. The magnitude of these peaks determines the noise floor of the system, and is used to set the detection threshold, as described in Section 4.2.

When there are multiple targets present in the environment, the echo signals obtained from each target combine additively in the time domain, as discussed in Section 4.4. As the matched filter is a **linear time-invariant** system, the effect of multiple targets is equivalent to adding the time-shifted autocorrelation patterns. Thus the minimum distance of separation at which two targets of equal power may be resolved is determined by the width of the central peak of the correlation output (Section 4.6, Rihaczec [118]). Thus narrower peaks are better.

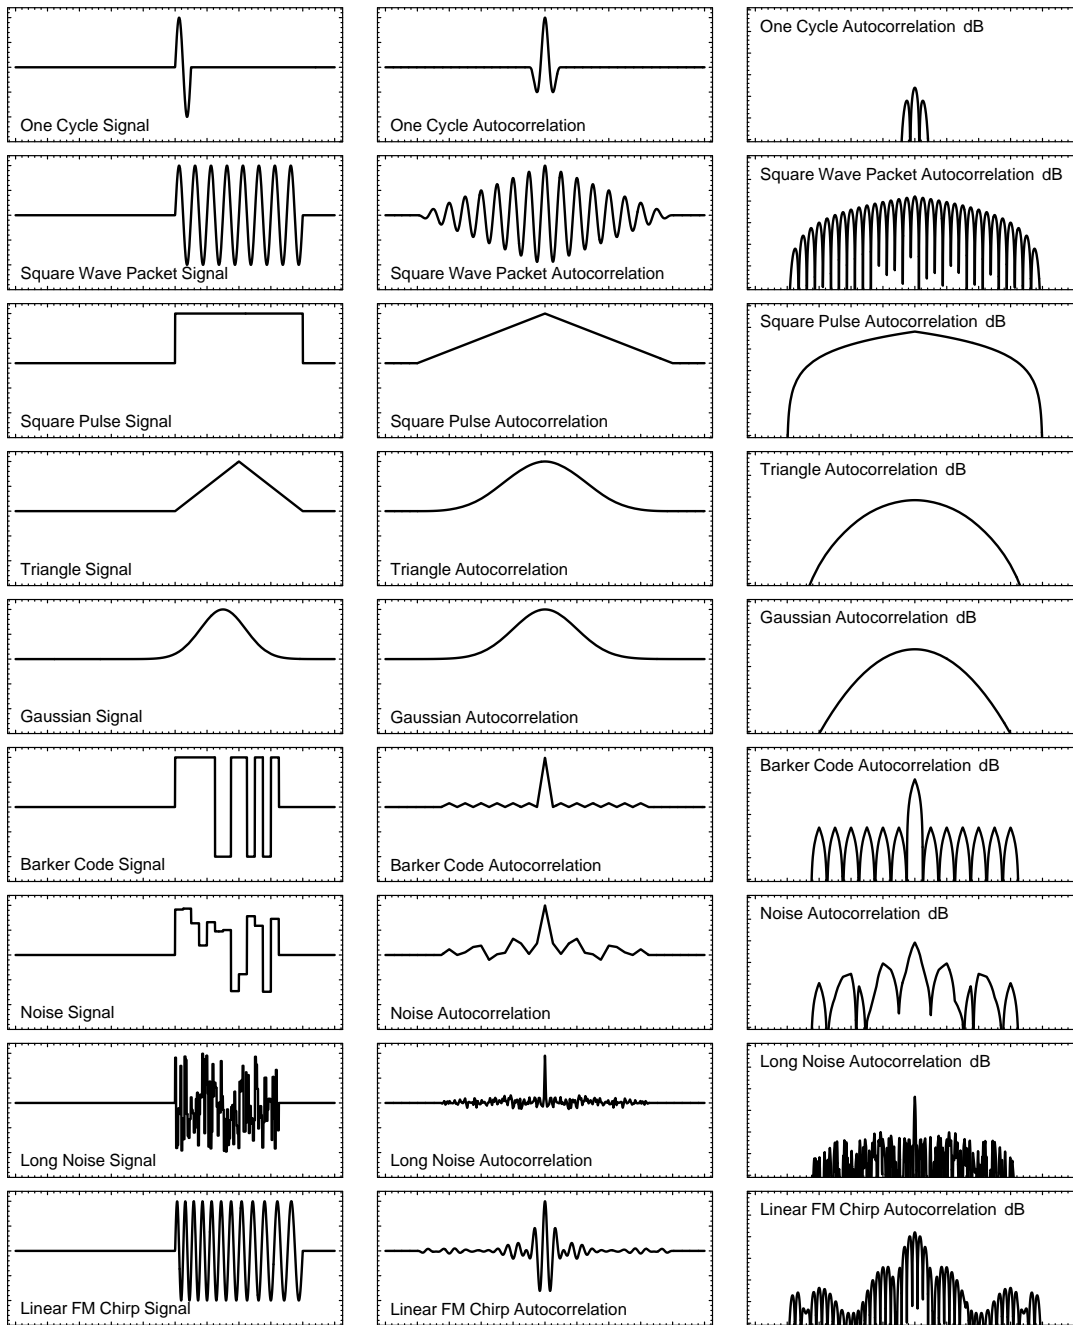


Figure B.1: A graphical dictionary of functions with their autocorrelations. The horizontal axes of all columns represents time, with the same scale. The first column shows the signals, in units of volts, which all begin at $t = 0$. The second column shows the autocorrelation, normalised to one volt at the peak, which always occurs at $t = 0$. The second column, like the first, uses an amplitude vertical scale. The third column shows the autocorrelation in decibels, and shows the correct relative strength of each signal. The major tick marks on the vertical axis of the third column represent 10dB.

Section 4.2 discussed the limitations placed upon the maximum range of the sonar system, including the maximum transmit signal power and the processing gain. If the transmit signal is selected to be very short, as shown for the one cycle signal in Figure B.1, then good precision and resolution may be achieved. This can be seen by the narrow autocorrelation function. However, there is no processing gain with this system (see the decibel plot in Figure B.1). Furthermore, such a system is limited by the maximum instantaneous transmit power, which is a function of the mylar breakdown voltage in the case of our transducers (Section 6.2). The only way to improve the signal power in this case is to increase the length of the transmit signal. If this is done by simply transmitting more cycles, the autocorrelation function becomes wider. The multiple sidelobes must be screened out if false detection is to be avoided. The additional width also implies a loss of resolution.

The two signals considered thus far, being the one cycle signal and the square wave packet signal, may both be regarded as a sine-wave multiplied by a square envelope signal. The sine-wave may also be regarded as a carrier signal. The autocorrelation of a sine-wave is a cosine-wave of the same frequency (Cadzow [20]), so the removal of the sine-wave from the signal also removes the cosine-wave from the autocorrelation. While real signals require a carrier to propagate, the carrier may be removed by demodulation. Therefore we now consider only the envelope of the signal. Comparing the square pulse signal to the square wave packet signal, we see that the autocorrelation of the signal envelope matches the envelope of the square wave packet autocorrelation, as expected. There is some gain introduced by removing the carrier, which can be seen by the different levels in the decibel plot. However, it is easier to compare the envelope functions than the modulated carrier signals.

The width of the autocorrelation function may be reduced slightly by using triangular or Gaussian signals, but the autocorrelation functions demonstrate that these signals carry less energy than the square pulse of a similar width. Increasing the width of the triangular and Gaussian functions brings up the signal energy, but also widens the autocorrelation and degrades resolution.

The Barker coded signal provides an autocorrelation with a very narrow central peak. This provides good target resolution abilities, while still allowing a long signal to carry more energy. This signal provides good peak autocorrelation power as well. While there are many sidelobes, their power is significantly lower than the main peak, allowing closely spaced targets to be resolved. Since the energy of this signal has been squashed into the central peak, matched filter systems are also known as pulse compression systems. When this signal is applied to a carrier, it becomes both an envelope, defining the signal's extent, and also a modulation, defining the phase of the signal.

The noise signal, composed of 13 random numbers, displays a central peak, but the sidelobe structure is not as good as the Barker coded signal. A long noise signal composed of 100 random numbers displays a very narrow central peak, and has good sidelobe properties as well. This signal cannot be compared with the other sequences, as the modulation function contains much higher frequencies. However, when long noise sequences are used, the autocorrelation becomes more like a delta function. Such signals form the basis of spread spectrum systems.

The final signal, the linear FM chirp, is included for completeness, as it is commonly used in radar systems. The autocorrelation strength is good, but the sidelobes in this configuration are

not ideal, as they can readily be mis-interpreted, leading to erroneous detection. Furthermore, the central peak is wider than that of the Barker coded system.

All of the signals shown in Figure B.1 are ideal, and many assume that the signal path has a very wide bandwidth. In reality this is not the case; this has the effect of band-limiting the signals, which in turn rounds off the sharp edges in the time domain response.

The aims of this study of matched filters, and the goals of this thesis (Section 1.2), namely to equip a *robotic agent* with sensing ability, allowing it to detect objects and traverse and map a scene, are best achieved with a system that is capable of sensing the largest volume with the best precision. Of the signal types considered for use with matched filters, this is best done by the Barker code type which allows greater signal energy in a system limited by peak transmitter power, while maintaining precision.

Appendix C

Optimisation of AR Spectral Estimators

This appendix describes methods for improving the efficiency of calculating the a_k parameters in the [YW](#) and [LSMYW](#) spectral estimators. Firstly, the properties of reversals are presented. They are used in the derivation of the Durbin algorithm, Section [C.2](#), which is used to optimise the [YW](#) spectral estimator. The optimisation of the [LSMYW](#) spectral estimator is described in Section [C.3](#). Finally, methods of checking the two optimised methods are described in Section [C.4](#).

C.1 Properties of Reversals

The following properties of reversals, denoted \sim , come from Table 2.5 of Therrien [[131](#)]. The meaning of vector reversal (where vectors are understood to be column vectors unless otherwise stated) is to take the elements,

$$\mathbf{x} = \begin{bmatrix} x_1 & x_2 & \dots & x_N \end{bmatrix}^T,$$

where T denotes the transpose, and to reverse their order,

$$\tilde{\mathbf{x}} = \begin{bmatrix} x_N & x_{N-1} & \dots & x_1 \end{bmatrix}^T.$$

Matrix reversal consists of reversing the row order and reversing the column order separately, in either order,

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}, \quad \tilde{\mathbf{A}} = \begin{bmatrix} a_{33} & a_{32} & a_{31} \\ a_{23} & a_{22} & a_{21} \\ a_{13} & a_{12} & a_{11} \end{bmatrix}. \quad (\text{C.1})$$

Some properties of reversals, for matrices and, where appropriate, for vectors as well, are

$$\widetilde{\mathbf{AB}} = \tilde{\mathbf{A}}\tilde{\mathbf{B}} \quad \widetilde{(\mathbf{A}^{-1})} = (\tilde{\mathbf{A}})^{-1} \quad \widetilde{(\mathbf{A}^T)} = (\tilde{\mathbf{A}})^T \quad (\text{C.2})$$

These properties are used in the derivation of the Durbin algorithm in Section [C.2](#).

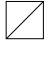
C.2 Durbin Algorithm

The normal equations ([11.22](#)) may be solved efficiently by using the Durbin algorithm, with complexity $O(p^2)$. The Durbin algorithm is also commonly known in the literature as the Levinson-Durbin recursion. The algorithm is only briefly discussed here. Full details are given by Therrien

[131, Section 8.5], Golub [45, Section 5.7], and Kay and Marple [75]. The code was written to algorithm 5.7-1 from Golub.

The matrix of autocorrelations, $R_x[i]$, on the left hand side of (11.22), repeated here for convenience,

$$\mathbf{T}_p = \begin{bmatrix} R_x[0] & R_x[1] & R_x[2] & \cdots & R_x[p-1] \\ R_x[-1] & R_x[0] & R_x[1] & \cdots & R_x[p-2] \\ R_x[-2] & R_x[-1] & R_x[0] & \cdots & R_x[p-3] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ R_x[1-p] & R_x[2-p] & R_x[3-p] & \cdots & R_x[0] \end{bmatrix}, \quad (\text{C.3})$$

has identical elements along each diagonal. The p in \mathbf{T}_p denotes the order of the square matrix. Matrices with constant diagonals are called Toeplitz matrices. Toeplitz matrices are also persymmetric, that is, they are symmetrical about the bottom-left to top-right diagonal  (Golub [45, Section 5.7]). Persymmetric matrices obey the relationship

$$\mathbf{T}_p^T = \widetilde{\mathbf{T}_p}, \quad (\text{C.4})$$

where \sim denotes reversal (Section C.1). The fact that \mathbf{T}_p is persymmetric can be clearly seen in (11.21).

An important property of persymmetric matrices is that their inverse is also persymmetric (Golub [45]), i.e.

$$(\mathbf{T}_p^{-1})^T = \widetilde{(\mathbf{T}_p^{-1})}. \quad (\text{C.5})$$

This property is used by Durbin's algorithm.

We may use (11.27) to re-write the negative lags of (C.3) in terms of the positive lags. Without loss of generality, it is assumed that $R_x[0] = 1$. If this is not the case, then we may construct the normalised autocorrelation

$$R'_x[i] = \frac{R_x[i]}{R_x[0]}. \quad (\text{C.6})$$

The denominator of (C.6) will not be zero unless the signal is zero. Furthermore, Bracewell [14] shows in an appendix to Chapter 3 that the autocorrelation always has its maximum value at 0. The division by $R_x[0]$ may be viewed as dividing by $R_x[0]$ on both sides of the normal equations (11.22). Therefore the solution, \mathbf{a}_p , remains unchanged by this procedure.

Using the normalised autocorrelation, $R'_x[i]$, and applying the autocorrelation symmetry property (11.27) yields the normalised autocorrelation matrix,

$$\mathbf{T}'_p = \begin{bmatrix} 1 & R'_x[1] & \cdots & R'_x[p-1] \\ R'_x[1] & 1 & \cdots & R'_x[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ R'_x[p-1] & R'_x[p-2] & \cdots & 1 \end{bmatrix}. \quad (\text{C.7})$$

This matrix is both symmetrical and persymmetric. Thus (C.4) becomes

$$\mathbf{T}'_p = \widetilde{\mathbf{T}'_p}. \quad (\text{C.8})$$

Because \mathbf{T}'_p is symmetric ($\mathbf{T}'_p = \mathbf{T}'_p{}^T$), its inverse is also symmetric ($\mathbf{T}'_p{}^{-1} = (\mathbf{T}'_p{}^{-1})^T$). Therefore (C.5) becomes

$$\mathbf{T}'_p{}^{-1} = \widetilde{(\mathbf{T}'_p{}^{-1})}. \quad (\text{C.9})$$

Durbin's algorithm solves the unit Yule-Walker equations, a modified version of (11.22),

$$\begin{bmatrix} 1 & R'_x[1] & \dots & R'_x[p-1] \\ R'_x[1] & 1 & \dots & R'_x[p-2] \\ \vdots & \vdots & \ddots & \vdots \\ R'_x[p-1] & R'_x[p-2] & \dots & 1 \end{bmatrix} \cdot \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = - \begin{bmatrix} R'_x[1] \\ R'_x[2] \\ \vdots \\ R'_x[p] \end{bmatrix},$$

or in vector notation

$$\mathbf{T}'_p \mathbf{a}_p = -\mathbf{r}'_p. \quad (\text{C.10})$$

Note that the p subscripts here denote the length of vectors and the square size of matrices. The matrix \mathbf{T}'_p is required to be positive definite (Anton [6]) for the Durbin algorithm to operate correctly.

The Durbin algorithm operates iteratively, building up the solution by including one more equation at each step. The solution for $k+1$ equations depends upon the solution of k equations, forming a recurrence. The recurrence is initialised with the solution for a single equation, $k=1$.

We turn firstly to the recurrence component. Assume that we have a solution \mathbf{A}_k , where $k < p$, to the partial system

$$\mathbf{T}'_k \begin{bmatrix} a_1 & a_2 & \dots & a_k \end{bmatrix}^T = - \begin{bmatrix} R'_x[1] & R'_x[2] & \dots & R'_x[k] \end{bmatrix}^T = -\mathbf{r}'_k,$$

where \mathbf{r}'_k is the column vector of autocorrelations, $R'_x[i]$, up to lag k . Each \mathbf{a}_k is unique to the stage that it is used in, and in general, the elements of \mathbf{a}_k change from one recurrence to the next. The matrix \mathbf{T}'_k and vector \mathbf{r}'_k do not have this property, as the k -th order system is a sub-matrix of the $k+1$ order system.

The $k+1$ order system is

$$\mathbf{T}'_{k+1} \begin{bmatrix} a_1 & a_2 & \dots & a_k & a_{k+1} \end{bmatrix}^T = - \begin{bmatrix} R'_x[1] & R'_x[2] & \dots & R'_x[k] & R'_x[k+1] \end{bmatrix}^T, \quad (\text{C.11})$$

which may also be written in terms of the solution of order k as the partitioned matrix system

$$\begin{bmatrix} \mathbf{T}'_k & \widetilde{\mathbf{r}'_k} \\ (\widetilde{\mathbf{r}'_k})^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{z}_k \\ \alpha_k \end{bmatrix} = - \begin{bmatrix} \mathbf{r}'_k \\ R'_x[k+1] \end{bmatrix}. \quad (\text{C.12})$$

Thus the desired solution, \mathbf{a}_{k+1} , is partitioned into a column vector \mathbf{z}_k and one additional scalar element, α_k . The column vector \mathbf{z}_k consists of the elements a_1 to a_k , which are different from the elements of \mathbf{a}_k from one iteration to the next. Expanding (C.12), the column vector \mathbf{z}_k is

$$\mathbf{z}_k = \mathbf{T}'_k{}^{-1} \left(-\mathbf{r}'_k - \alpha_k \widetilde{\mathbf{r}'_k} \right), \quad (\text{C.13})$$

along with the additional element

$$\alpha_k = - \left(\widetilde{\mathbf{r}'_k} \right)^T \cdot \mathbf{z}_k - R'_x[k+1]. \quad (\text{C.14})$$

Equation (C.13) may be simplified using (C.10) to replace $-\mathbf{T}'_k{}^{-1}\mathbf{r}'_k$, yielding

$$\mathbf{z}_k = \mathbf{a}_k - \alpha_k \mathbf{T}'_k{}^{-1} \widetilde{\mathbf{r}}'_k. \quad (\text{C.15})$$

We now seek a simplification of $\mathbf{T}'_k{}^{-1} \widetilde{\mathbf{r}}'_k$. By rearranging (C.10) and applying a reversal to both sides as described in Section C.1, we find

$$-\widetilde{(\mathbf{T}'_k{}^{-1})} \widetilde{\mathbf{r}}'_k = \widetilde{\mathbf{a}}_k.$$

Remembering that \mathbf{T}'_k is symmetric and persymmetric, and that $\mathbf{T}'_k{}^{-1}$ is also persymmetric, we may apply (C.9) to yield

$$-\mathbf{T}'_k{}^{-1} \widetilde{\mathbf{r}}'_k = \widetilde{\mathbf{a}}_k.$$

Substituting into (C.15), we may now write \mathbf{z}_k as

$$\mathbf{z}_k = \mathbf{a}_k + \alpha_k \widetilde{\mathbf{a}}_k. \quad (\text{C.16})$$

Substituting (C.16) into (C.14) yields

$$\alpha_k = -\left(\widetilde{\mathbf{r}}'_k\right)^T \cdot (\mathbf{a}_k + \alpha_k \widetilde{\mathbf{a}}_k) - R'_x[k+1].$$

Grouping the factors of α_k gives

$$\alpha_k \cdot \left(1 + \widetilde{\mathbf{r}}'^T_k \mathbf{a}_k\right) = -\left(\widetilde{\mathbf{r}}'_k\right)^T \mathbf{a}_k - R'_x[k+1]. \quad (\text{C.17})$$

The reversal of the term $\widetilde{\mathbf{r}}'^T_k \mathbf{a}_k$ may be neglected as the result, $\mathbf{r}'^T_k \mathbf{a}_k$, is a scalar. The final expression, allowing us to compute α_k in terms of the known k -th order solution, is

$$\alpha_k = \frac{-\left(\widetilde{\mathbf{r}}'_k\right)^T \mathbf{a}_k - R'_x[k+1]}{1 + \mathbf{r}'^T_k \mathbf{a}_k}. \quad (\text{C.18})$$

We now prove that the denominator is positive and hence non-zero. Remembering that \mathbf{T}'_{k+1} is positive definite, we use the definition of a positive definite matrix (Weisstein [137]),

$$\mathbf{X}^T \mathbf{T}'_{k+1} \mathbf{X} > \mathbf{0}_{(k+1) \times (k+1)}, \quad (\text{C.19})$$

where \mathbf{X} is an arbitrary $(k+1) \times (k+1)$ matrix. We now construct a particular matrix, \mathbf{X} , which is fortuitously known to provide the required result,

$$\mathbf{X} = \begin{bmatrix} \mathbf{I}_{k \times k} & \widetilde{\mathbf{a}}_k \\ \mathbf{0}_k^T & 1 \end{bmatrix}, \quad (\text{C.20})$$

where $\mathbf{I}_{k \times k}$ is the identity matrix, and $\mathbf{0}_k$ is the column zero vector of length k . The matrix \mathbf{T}'_{k+1} is partitioned as in (C.12). Multiplying out the first stage of (C.19) yields

$$\mathbf{X}^T \mathbf{T}'_{k+1} = \begin{bmatrix} \mathbf{T}'_k & \widetilde{\mathbf{r}}'_k \\ (\widetilde{\mathbf{a}}_k)^T \mathbf{T}'_k + (\widetilde{\mathbf{r}}'_k)^T & 1 + \widetilde{\mathbf{a}}_k^T \widetilde{\mathbf{r}}'_k \end{bmatrix}. \quad (\text{C.21})$$

Algorithm C.1 The Durbin algorithm with order $O(\frac{3}{2}p^2)$ flops.

Require: $R'_x[\cdot]$ contains the normalised autocorrelation sequence

Ensure: \mathbf{a}_p will contain the Yule-Walker coefficients

$\mathbf{a}_1[1] \leftarrow -R'_x[1]$ {Initialise the recurrence, (C.23).}

for $k = 1$ to $p - 1$ **do**

$\beta_k \leftarrow 1 + \mathbf{r}_k^T \mathbf{a}_k$ {Denominator of (C.18).}

$\alpha_k \leftarrow \frac{-(R'_x[k+1] + (\tilde{\mathbf{r}}_k^T)^T \mathbf{a}_k)}{\beta}$ {Equation (C.18).}

$\mathbf{z}_k \leftarrow \mathbf{a}_k + \alpha_k \tilde{\mathbf{a}}_k$ {Equation (C.16).}

$\mathbf{a}_{k+1} \leftarrow \begin{bmatrix} \mathbf{z}_k \\ \alpha_k \end{bmatrix}$ {Construct next order \mathbf{a}_{k+1} , eqns (C.11) and (C.12).}

end for

The lower left element may be shown to be the zero row vector, $\mathbf{0}_k^T$, using (C.8); a transposed and reversed form of (C.10), some properties of reversals (C.2) and the symmetric matrix property. The lower right element, $1 + \tilde{\mathbf{a}}_k^T \mathbf{r}_k^T$, is scalar, and hence the reversal may be ignored and the vector term is equal to its transpose.

We are now able to form the complete left hand side from the positive definite definition (C.19), which is

$$\mathbf{X}^T \mathbf{T}'_{k+1} \mathbf{X} = \begin{bmatrix} \mathbf{T}'_k & \mathbf{0}_k \\ \mathbf{0}_k^T & 1 + \mathbf{r}_k^T \mathbf{a}_k \end{bmatrix}. \quad (\text{C.22})$$

By the right hand side of equation (C.19), we may state that the right hand side of (C.22), and hence the lower right element in particular, is greater than zero. Thus (C.18) is well formed, and will not suffer from division by zero errors.

The recurrence relationship of the Durbin algorithm is initialised by the first order system of Yule-Walker equations,

$$\mathbf{T}'_1 \mathbf{a}_1 = 1 \cdot \mathbf{a}_1 = -R'_x[1]. \quad (\text{C.23})$$

The recurrence is then used to calculate \mathbf{a}_{k+1} for $k = 1, 2, \dots, p - 1$, from the previous solution \mathbf{a}_k , with total computational complexity $O(\frac{3}{2}p^2)$. The algorithm proceeds as shown in Algorithm C.1.

It is possible to improve upon the performance of Algorithm C.1 by reducing the computation required to obtain β_k , which is the denominator of (C.18). This is done by expanding

$$\beta_k = 1 + \mathbf{r}_k^T \mathbf{a}_k \quad (\text{C.24})$$

in terms of the lower order forms, using partitioned vectors and (C.16),

$$\beta_k = 1 + \begin{bmatrix} \mathbf{r}_{k-1}^T & R'_x[k] \end{bmatrix} \begin{bmatrix} \mathbf{a}_{k-1} + \alpha_{k-1} \tilde{\mathbf{a}}_{k-1} \\ \alpha_{k-1} \end{bmatrix}.$$

Multiplying out the matrices yields

$$\beta_k = 1 + \mathbf{r}_{k-1}^T \mathbf{a}_{k-1} + \alpha_{k-1} (\mathbf{r}_{k-1}^T \tilde{\mathbf{a}}_{k-1} + R'_x[k]).$$

Algorithm C.2 The Durbin algorithm with order $O(p^2)$ flops. The formulae written out here have dropped the use of the k subscript notation, and use only one array to store \mathbf{a} . The array \mathbf{z} may be eliminated as discussed in the text.

Require: $R'_x[\cdot]$ contains the normalised autocorrelation sequence

Ensure: \mathbf{a} will contain the order p Yule-Walker coefficients

```

{Initialise the recurrence, (C.23).}
 $\mathbf{a}[1] \leftarrow -R'_x[1]$ 
{ $\alpha$  and  $\beta$  now calculated before and at the end of each loop.}
 $\beta \leftarrow 1$ 
 $\alpha \leftarrow -R'_x[1]$ 
{Implement the recurrence}
for  $k = 1$  to  $p - 1$  do
   $\beta \leftarrow \beta \cdot (1 - \alpha^2)$  {Equation (C.25).}
   $\alpha \leftarrow \frac{-(R'_x[k+1] + \sum_{l=1}^k R'_x[k+1-l] \mathbf{a}[l])}{\beta}$  {Equation (C.18).}

  for  $l = 1$  to  $k$  do
     $\mathbf{z}[l] \leftarrow \mathbf{a}[l] + \alpha \cdot \mathbf{a}[k+1-l]$  {Equation (C.16).}
  end for
  {Fill in  $\mathbf{a}_{k+1}$ .}
  for  $l = 1$  to  $k$  do
     $\mathbf{a}[l] \leftarrow \mathbf{z}[l]$ 
  end for
   $\mathbf{a}[k+1] \leftarrow \alpha$ 
end for
```

Recognising and substituting the form $\beta_{k-1} = 1 + \mathbf{r}_{k-1}^T \mathbf{a}_{k-1}$, from (C.24), and the form in brackets as being the numerator of (C.18), we may simplify to

$$\beta_k = \beta_{k-1} + \alpha_{k-1} (-\beta_{k-1} \alpha_{k-1}) = \beta_{k-1} (1 - \alpha_{k-1}^2). \quad (\text{C.25})$$

This reduces the computation of β_k to a scalar recurrence relationship, instead of having to compute a vector product. The scalar values α_{k-1} and β_{k-1} are available from previous iterations. This simplification is incorporated into Algorithm C.2. The vector \mathbf{z} is included for clarity, but the implementation may be modified to remove it. This is done by computing $\mathbf{z}[l]$ and $\mathbf{z}[k+1-l]$, as a pair, and storing them in temporary variables. The values of $\mathbf{a}[l]$ and $\mathbf{a}[k+1-l]$ are no longer required once the temporary variables have been calculated, and may be overwritten by the new values of \mathbf{z} . The chirp application program implements Algorithm C.2 with the modification to eliminate \mathbf{z} .

This completes the discussion of the efficient implementation of the YW equations, as used in Section 11.4.3.

C.3 LSMYW Optimisation

The Least Squares Modified Yule-Walker (LSMYW) spectral estimator requires the efficient solution of equation (11.32). The square matrix $\mathbf{R}_1^T \mathbf{R}_1$ from (11.32) was calculated directly from an

array of autocorrelation values, R_x , by using a formula similar to (6.10) from [20]:

$$(\mathbf{R}_1^T \mathbf{R}_1)[k, l] = \sum_{m=1}^t R_x[m-k] R_x[m-l], 1 \leq k, l \leq p \quad (\text{C.26})$$

This formula is different from Cadzow's because it computes the matrix on the left hand side of (11.31), whereas Cadzow computes the matrix on the left hand side of (11.30) and includes provision for a MA estimator as well.

As the autocorrelation data is real, $\mathbf{R}_1^T \mathbf{R}_1$ is symmetrical. This can be seen from (C.26), where interchanging the indices k and l results in the same formula being obtained.

The Durbin algorithm outlined in Section C.2 does not apply to this problem, as $\mathbf{R}_1^T \mathbf{R}_1$ is not Toeplitz. This may be seen by multiplying out $\mathbf{R}_1^T \mathbf{R}_1$ for $p = 3$ and $t = 4$, yielding :

$$\begin{bmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

with

$$\begin{aligned} a_{11} &= R_x[0]^2 + R_x[1]^2 + R_x[2]^2 + R_x[3]^2 \\ a_{12} = a_{21} &= R_x[-1] R_x[0] + R_x[0] R_x[1] + R_x[1] R_x[2] + R_x[2] R_x[3] \\ a_{13} = a_{31} &= R_x[-2] R_x[0] + R_x[-1] R_x[1] + R_x[0] R_x[2] + R_x[1] R_x[3] \\ a_{22} &= R_x[-1]^2 + R_x[0]^2 + R_x[1]^2 + R_x[2]^2 \\ a_{23} = a_{32} &= R_x[-2] R_x[-1] + R_x[-1] R_x[0] + R_x[0] R_x[1] + R_x[1] R_x[2] \\ a_{33} &= R_x[-2]^2 + R_x[-1]^2 + R_x[0]^2 + R_x[1]^2 \end{aligned}$$

The diagonals can be seen to be sums down the autocorrelation vector with offsets between the two terms related the diagonal number. However, each entry progressing down to the right operates as a rolling window, with one term dropping out and a new one being introduced. We can see that it is symmetrical.

The code solves (11.32) for \mathbf{a} using the symmetrical matrix routine `DSYSV` provided with LAPACK [4], having computational complexity $O(n^3/6)$ (LINPACK [33]). Since `DSYSV` only uses the upper right triangular area of the input matrix, (C.26) is only evaluated for $1 \leq k \leq l \leq p$.

C.4 Checking the Code

The implementation of the optimised methods has been checked in one of two ways depending upon the circumstances.

Two different algorithms: In this method, the result is computed twice, once by the optimised method and again by the direct method. The results are then automatically compared. When both give the same results over a wide variety of test inputs, the optimised version is deemed to be correct. This works as long as the implementation of the direct method is correct. As the code for the direct method is usually much simpler than for the optimised method, it is easier to implement correctly. Any special cases must also be tested explicitly for this method to be valid. This method was suggested by Maguire [89, P. 33].

Residuals: Matrix problems of the form

$$\mathbf{Ax} = \mathbf{y},$$

where \mathbf{x} is unknown, may be checked by taking the solution \mathbf{x} and then evaluating

$$\mathbf{Ax} - \mathbf{y} = \boldsymbol{\varepsilon},$$

where $\boldsymbol{\varepsilon}$ is the residuals vector. If the residuals are of the magnitude of machine precision, then a good solution has been found. However, some poorly conditioned problems may produce residuals much larger than machine precision. If this is the case, the programmer needs to be aware of the problem and perhaps find an alternate formulation which produces smaller residuals (Press [114]).

The verification was only applied in a debug build of the software. The checking was automatically removed in the optimised build using macros, as suggested by Maguire [89].

Appendix D

Geometrical Sonar Simulation

D.1 Review

Previous sonar simulation environments have been reported, with varying levels of complexity. The simplest is the Multiple Autonomous Mobile Robot Simulator (MAMROS) (Naumovski [99]), which simulates the environment and operation of one or more Yamabico mobile robots. MAMROS includes a 2D sonar simulator, which performs a first bounce geometrical simulation of the sonar operation. M^cKerrow [94] describes a 2D geometrical sonar visualisation program. Kuc and Siegel [85] describe a sonar simulation at the signal level, using an impulse response model for each of the components of the signal path. Zhu [147] describes a 2D Transmission Line Matrix (TLM) model which numerically solves the wave equation, thereby simulating the propagation of the sound wave through the air, frame by frame, at great computational expense.

D.2 This Work

For the purposes of this work, a 2D geometrical simulation was constructed in Mathematica [139], which allows a user to construct an environment consisting of planes and point reflectors. Corner reflections are processed by performing multiple plane reflections. By this generalisation, the model may also be used to investigate what are commonly called multi-path echoes, where the signal reflects from many objects in turn (M^cKerrow [97]).

A simplified model of the ultrasonic transducer is used, assuming that there is a fixed beam angle beyond which no signal can be detected. Additionally, no frequency domain effects are taken into account. The model exists purely to find geometric paths from the transmitter to the receiver, and to compute their exact lengths. Conceptually it operates similarly to ray tracing.

D.3 Operation

The simulation model accepts a two dimensional model consisting of lines and points. It also accepts a transmit transducer, and a list of receive transducers. Each transducer is defined by a Cartesian coordinate, $\{x, y\}$, an orientation, θ , and a beam width, ϕ . When a transducer is drawn (see, for example, Figure D.1), the position of the transducer is marked with a red dot, the principal direction is marked with a black line, and the extent of the beam is drawn with a cyan fan shape.

To simplify working with the complicated sensor head discussed in Chapter 7, an object was defined to allow a sensor head and its associated transducers to be manipulated as a single unit.

The model operation is illustrated in Figure D.1. The process extends the virtual source principle which was demonstrated in Figure 14.5. The signal path starts at the transmitter, T, which insonifies the world model in accordance with the beam width. Objects which are insonified may wholly or partially intersect with the beam. Obstructed targets are detected by the search and are not considered to be insonified. When the beam encounters a line, such as W2 in Figure D.1, a copy of the transmitter, the virtual source VT2, is made. The virtual source has its position and orientation reflected by the line W2. Furthermore, the beam angle of the virtual source VT2 is restricted, by different amounts on the left and right sides, so that the beam emerging from VT2 will intersect with W2 and appear to be the reflection from the transmitter, T. If the beam from the transmitter, T, does not insonify the whole line object, then only the insonified part acts as a reflector. This is implemented by restricting the beam angle of the virtual source VT2. The restriction may be asymmetrical, as shown in Figure D.1.

The beam from a virtual source is not considered to perform insonification in between the virtual source, VT2, and the line object, W2, from which it was reflected. Only the region on the front side of W2 is considered to be insonified. With this restriction, the model now continues recursively, using the virtual source VT2 as the new transmitter. In this way, any number of reflections may be modelled. For practical reasons of implementation, however, the recursion is terminated by either of two conditions. The first is an arbitrarily selected maximum reflection count of 5. The second is a maximum path distance, which is user specified. Either of these conditions alone will guarantee that the recursion will terminate.

At each iteration the model checks whether any of the virtual sources insonifies any of the receivers. If a path exists from a source to a receiver, the path is compared to the beam angle of the receiver. If this test succeeds, the model produces the list of virtual sources making up the path, and is capable of drawing the path (with or without the virtual sources) on a map of the environment in which the path was found.

Each line obstacle is deemed to have a front and a back. The front is the only surface which is acoustically active in the model. Signals from virtual sources may pass through the surface from the back without being obstructed. The front of a line is defined to be the face 90° clockwise from the line's vector. Thus a sequence of lines drawn clockwise to form a closed shape will have all their front sides on the inside of the shape. A wall must be constructed by placing two lines back to back with their 'fronts' facing outward. This construction does not obstruct the operation of virtual sources as each virtual source is associated with the particular line from which it is reflected. Other objects between the virtual source and the line which produces it cannot obstruct the beam from the virtual source.

This model is geometric in its focus, and does not attempt to model any of the frequency characteristics of the transducers or signal path, any form of absorption, or the finer points of transducer beam patterns. These features may all be added to this model, as discussed in Chapter 4, if required. Indeed, some extensions were made to model a Gaussian beam pattern, as described in Section 15.5.1, according to the models derived in Appendix H.




Figure D.1: Demonstration of how a path with multiple reflection points is constructed through the use of multiple intermediate virtual sources. The upper panel shows the path, while the lower panel shows the distance of flight. Only the path from the transmitter to the left receiver is shown. The transmitted signal insonifies the lower wall segment, W2. The reflection from this segment is modelled by creating the virtual source VT2, using the index of the reflecting wall segment. The beam of VT2 is narrowed so that it is not wider than the wall segment W2. The beam from this virtual transmitter subsequently insonifies the upper wall segment, W1. The reflection from the upper segment is modelled by creating the virtual source VT21 (and again adding the index of the reflecting wall). A line drawn directly between this virtual source and the receivers falls within the beam fan of both the transmitter and receiver, so a path is deemed to exist.

D.4 Example

An example of output from the model is shown in Figure D.2. This model is useful for automatically exploring possible paths of a multi-path echo.

D.5 Effect of Non-90° Corners

There are some interesting effects related to acute and obtuse corners which have not been previously published. The explanation of these effects also explains further the operation of the model.

D.5.1 Acute Corners

An environment with an acute corner is shown in Figure D.3. The beam from the transmitter insonifies both the upper and lower wall segments W1 and W2. Treating the upper wall segment first, the transmitter is reflected to form VT1. Note that the beam edges of VT1 are limited by the extent of the wall segment. Thus the virtual transmitter has a narrower beam width than the original transmitter.

Turning our attention to the beam from VT1, we see that both W1 and W2 lie within the beam. In order to produce the correct reflections in this case, each virtual source is paired with its reflecting element, and no reflection occurs between the virtual transmitter and its paired reflector. Thus the beam from VT1 insonifies W2 without occlusion by W1, as it must in order to model the reflected wave from the original transmitter.

As the beam from VT1 is considered to insonify W2 without occlusion, the model creates a new virtual source, VT12, by reflecting VT1 in the plane W2. The beam angle is again limited by the extent of the reflector. After considering again the pairing between VT12 and its reflecting object W2, and also the beam angle of VT12, we see that VT12 insonifies both of the receivers. Thus two paths are formed. An internal data structure describing all of the intermediate virtual sources and reflectors is created. The paths are drawn using the information recorded in this data structure. The path distance may also be calculated exactly using this information.

The paths involving VT21 may be constructed by following a similar process involving first W2 and then W1. The paths so formed exhibit the same properties over a wide variety of corner angles, although there is only a limited range of angles of the 80° corner orientation, with respect to the sonar head, in which the corner reflection is observed by both of the receivers. An important feature is that the two virtual sources, which fuse together as one virtual source when the corner is 90° (Figure 14.5), split into two distinct sources for the acute corner. The amount of splitting is controlled by the angle of the corner, and the distance of the sensor head from the corner. If the corner reflector is of finite extent, and the sensor head is at great distance from the acute corner, the dimensions of the corner reflector may be smaller than the required reflecting area, rendering the corner reflector invisible to the sensor.

D.5.2 Obtuse Corners

The paths for a 100° obtuse corner are shown in Figure D.4. The virtual sources are constructed

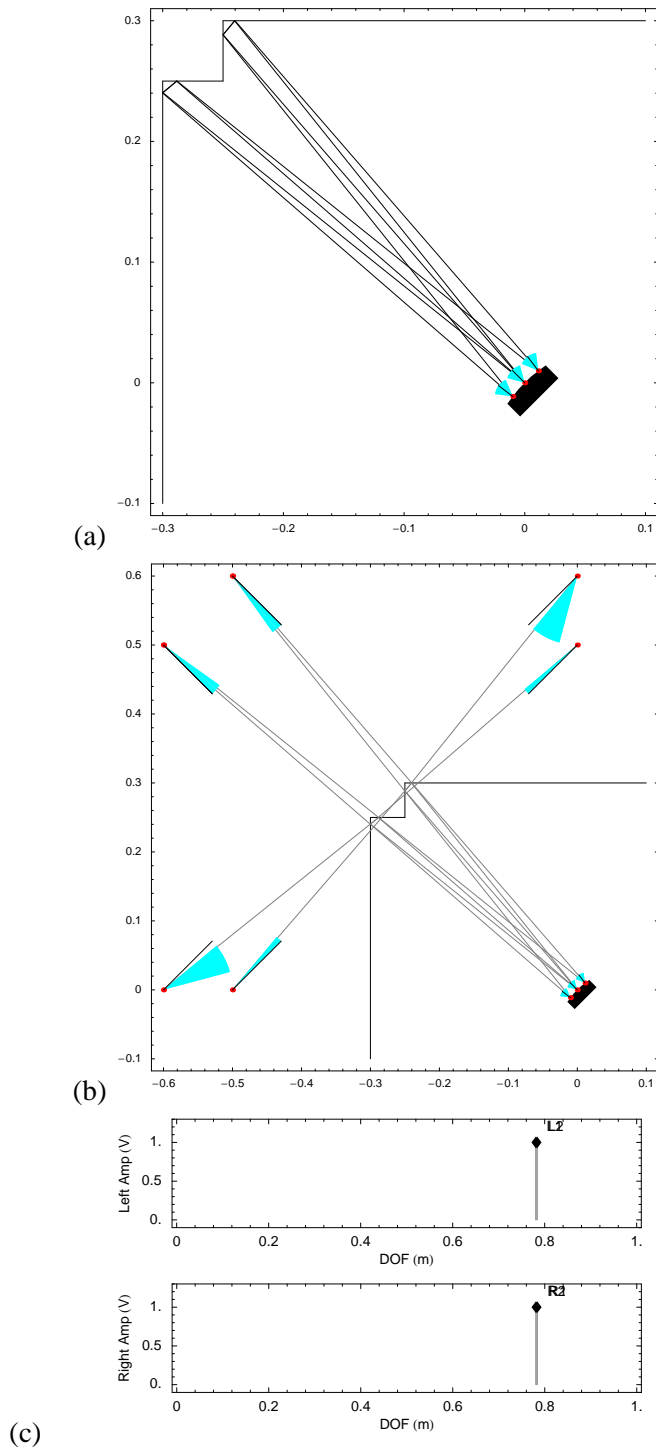


Figure D.2: Sample output from the 2D geometric ultrasonic model for a corner with trimming. This forms two corner reflectors. (a) Shows the sonar paths which strike the receivers. The x and y axes have dimensions of metres. (b) Shows the virtual sources, along with their principal directions and restricted beams. (c) Shows the sonogram of the echoes detected by the left and right receivers. The two echoes detected on each channel have path lengths of 0.7808 m and 0.7835 m. As the situation is symmetrical, the echoes on the left and right channels are identical. The amplitude is not modelled, so all echoes have unit amplitude.

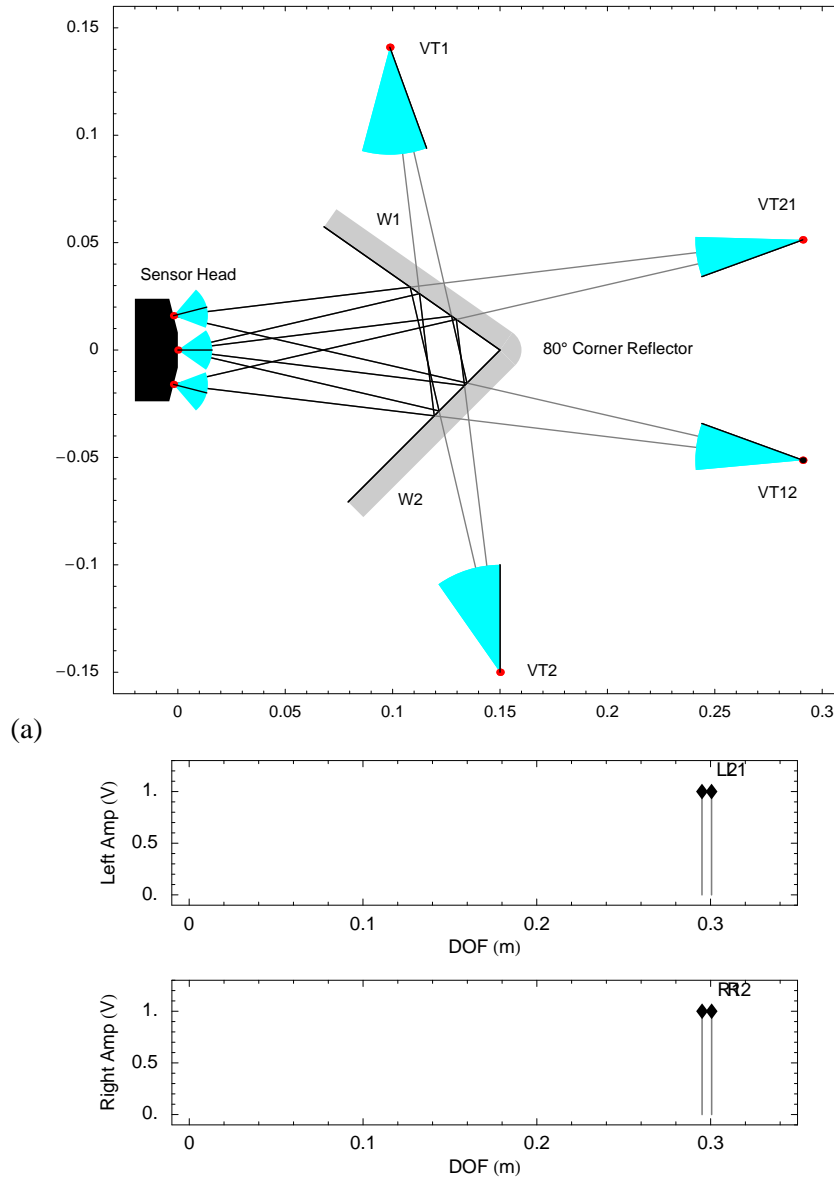


Figure D.3: An 80° acute corner. (a) Shows the paths from the transmitter to the receivers. The numbering of the virtual sources indicates the order of the walls from which they were reflected. There are two separate virtual sources, VT12 and VT21, which insonify the receivers. The asymmetry of the virtual sources, VT1 and VT2, is due to the corner being asymmetrically oriented with respect to the sensor head. (b) A sonogram of the echoes, which shows that there are two distinct echoes detected by each receiver.




Figure D.4: A 100° obtuse corner, showing the virtual sources. The numbering of the virtual sources indicates the order of the walls from which they were reflected. The beams from the virtual sources, VT12 and VT21, do not intersect with the receivers, and thus do not form a path. The only echo which will be detected from such a corner is the diffraction from the corner (not shown). No sonogram is shown as there are no echoes detected by this model.

by the same method used for the acute corner, but the reflection geometry turns out to be quite different. The positions of the virtual sources VT12 and VT21 are reversed when compared to those of the acute corner shown in Figure D.3. The virtual source beams of the obtuse corner do not intersect with either of the receivers. The only echo which will be detected from this reflector is diffraction from the corner itself. If the power of the diffracted signal is below the detection threshold of the sensor, as discussed in Section 4.2.2, then the obtuse corner will not be detected at all.

The results of the previous section, combined with the results of this section, demonstrate that corner reflectors which are to be used as sonar beacons must be precise 90° corners to obtain the desired results, as other angles will lead to the corner reflector being invisible to the sonar sensor.

D.6 Conclusion

The 2D geometrical model was used to check the operation of many portions of the work, and to provide reference path angles and distances of flight for the computer based experiments in Chapter 15. While simple in nature, useful results have been found. It is important to have a sonar model so that the detected echoes may be compared with the echoes which were expected, and unknown features identified and tracked. Future extension of models of this type may be made easier by exploiting the mature knowledge of the field of ray tracing, which solves a remarkably similar problem.

Appendix E

Atmospheric Effects

The formulae for the speed of sound in air and the absorption of sound by the air depend upon the thermodynamical properties of the air's constituents. As these properties are common to both the speed and the absorption formulae, a unified set of formulae is summarised from the literature and presented here for convenience. Due to the varying notations prevalent in the literature a unified notation has been adopted here. As many of the formulae appear in various sources, the most recently published form has been used. The source used for each formula is identified.

The first section lists the variables used by the formulae. Subsequent sections describe the formulae, in calculation order with the fundamental quantities first, leading finally to the formulae for the absorption, the speed of sound, and the compensation of the speed of sound for dispersion (frequency effects).

E.1 Variables

The following variables, which are listed alphabetically, are used in the formulae of this appendix, along with their constant value, status as a parameter, or an equation for calculation in terms of other known variables.

- | | |
|------------------------------|--|
| α | is the absorption coefficient of the pressure wave, in Nepers per metre. The Neper is defined in ANSI S1.26-1978 [2], and is abbreviated Np. The absorption is calculated using (E.9). |
| α' | is the absorption coefficient of the pressure wave, in decibels per metre. This is related to α by equation (E.10). |
| $\alpha_{v,N}, \alpha_{v,O}$ | are the molecular vibration relaxation coefficients of absorption (Nepers/m) calculated using (E.5) and (E.6). |
| c | is the speed of sound, in $\text{m} \cdot \text{s}^{-1}$, as used by the molecular vibration formulae in Section E.3. This may be computed approximately by (E.7), or precisely by (E.11) . |
| c_0 | is the speed of sound ($\text{m} \cdot \text{s}^{-1}$) at zero frequency, calculated using (E.11). |
| c_ϕ | is the speed of sound ($\text{m} \cdot \text{s}^{-1}$) including the effects of dispersion, calculated using (E.14). |
| d | is the distance of flight, in m. |

f_e	is the enhancement factor (dimensionless), as a function of static pressure and temperature, calculated using (E.13).
$f_{r,N}, f_{r,O}$	are the molecular vibration relaxation frequencies (Hertz) of N_2 and O_2 , calculated using (E.2) and (E.3).
h	is the absolute humidity (molar concentration of water vapour in) %, may be measured directly or calculated using (E.4).
h_r	is the relative humidity in %. This is a parameter.
f	is the frequency of the pure tone, in Hz. This is a parameter.
p_d	is the dynamic sound pressure after travelling a distance d .
p_0	is the dynamic sound pressure at the source, in Pascals.
p_s	is the static air pressure, in kilo Pascals. This is a parameter.
p_{s0}	is the reference atmospheric pressure at sea level, $p_{s0} = 101.325 \text{ kPa}$. This is the pressure referred to by Standard Temperature and Pressure (STP) .
p_{sat}	is the saturation vapour pressure of water in air, in kPa, calculated using (E.1).
T_a	is the absolute thermodynamic temperature in K. This is a parameter.
T_0	is the reference ambient temperature, $T_0 = 293.15 \text{ K}$. This is the temperature referred to by standard temperature and pressure .
T_3	is the triple point of water, $T_3 = 273.16 \text{ K}$.
x_w	mole fraction of water vapour, calculated using (E.12).

E.2 Saturation Pressure

The saturation pressure of water vapour, p_{sat} , may be calculated from (see Bass [9] eqn (2))

$$\log_{10} \left(\frac{p_{sat}}{p_{s0}} \right) = -6.8346 \left(\frac{T_3}{T_a} \right)^{1.261} + 4.6151, \quad (\text{E.1})$$

in kPa. Equation (E.1) is in good agreement with an older equation (see ANSI S1.26-1978 [2], eqn (D11)) within the range -20°C and 50°C . If temperatures outside this range are required, then the older formulation from ANSI S1.26-1978 should be used instead.

E.3 Molecular Vibrations

The molecular vibration relaxation frequencies of nitrogen $f_{r,N}$ and oxygen $f_{r,O}$ can be calculated using (see Bass [9] eqns (5) and (4))

$$f_{r,N} = \frac{p_s}{p_{s0}} \sqrt{\frac{T_0}{T_a}} \left(9 + 280h \exp \left\{ -4.17 \left[\left(\frac{T_0}{T_a} \right)^{\frac{1}{3}} - 1 \right] \right\} \right) \quad (\text{E.2})$$

and

$$f_{r,O} = \frac{p_s}{p_{s0}} \left(24 + 4.04 \times 10^4 h \frac{0.02 + h}{0.391 + h} \right), \quad (\text{E.3})$$

in Hertz. The absolute humidity h can be found from the relative humidity h_r (in %) by (see Bass [9] eqn (6))

$$h = h_r \left(\frac{p_{sat}}{p_s} \right) \quad (\text{E.4})$$

in %. The molecular vibration relaxation coefficients of nitrogen $\alpha_{v,N}$ and oxygen $\alpha_{v,O}$ can be found using (see ANSI S1.26-1978 [2], eqns (5) and (7))

$$\alpha_{v,N} = \left(\frac{2\pi \cdot 0.781}{35} \right) \left(\frac{3352.0}{T_a} \right)^2 \exp \left(\frac{-3352.0}{T_a} \right) \left(\frac{f}{c} \right) \left\{ 2 \frac{f}{f_{r,N}} \left[1 + \left(\frac{f}{f_{r,N}} \right)^2 \right]^{-1} \right\} \quad (\text{E.5})$$

and

$$\alpha_{v,O} = \left(\frac{2\pi \cdot 0.209}{35} \right) \left(\frac{2239.1}{T_a} \right)^2 \exp \left(\frac{-2239.1}{T_a} \right) \left(\frac{f}{c} \right) \left\{ 2 \frac{f}{f_{r,O}} \left[1 + \left(\frac{f}{f_{r,O}} \right)^2 \right]^{-1} \right\} \quad (\text{E.6})$$

in $\text{Np} \cdot \text{m}^{-1}$. The unit Np denotes a Neper, which is dimensionless. The formula used by ANSI S1.26-1978 for the speed of sound is (see ANSI S1.26-1978 eqn (6))

$$c = 343.23 \sqrt{\frac{T_a}{T_0}} \quad (\text{E.7})$$

in $\text{m} \cdot \text{s}^{-1}$. This is essentially the same as (12.6), but the lack of precision implied in this formula is not significant for most applications of the absorption formulae. The zero frequency speed of sound c_0 , which will be defined in Section E.5 as (E.11), may be used instead if it is available.

E.4 Absorption

ANSI S1.26-1978 [2] eqn (1) defines a standard method for calculating the absorption to be

$$p_d = p_0 e^{-\alpha d}, \quad (\text{E.8})$$

where the variables are listed below.

The absorption coefficient α is calculated as a function of the frequency, static pressure, temperature and humidity, and includes contributions from the two dominant gases, nitrogen and oxygen. The formula is (see ANSI S1.26-1978 [2] eqns (3) and (4))

$$\alpha = 1.84 \times 10^{-11} \left(\frac{T_a}{T_0} \right)^{\frac{1}{2}} f^2 \left(\frac{p_{s0}}{p_s} \right) + \alpha_{v,N} + \alpha_{v,O} \quad (\text{E.9})$$

in $\text{Np} \cdot \text{m}^{-1}$, where $\alpha_{v,N}$ and $\alpha_{v,O}$ are given by (E.5) and (E.6) respectively.

The value given by (E.9) may be converted to absorption in decibels per metre α' by the relationship

$$\alpha' = \frac{20}{\ln(10)} \alpha = 8.69 \alpha \quad (\text{E.10})$$

in $\text{dB} \cdot \text{m}^{-1}$. Isopleth plots of the variation in the absorption of sound α' are shown in Figure E.1. The plots show that the pressure has little effect upon the absorption. High frequencies are attenuated progressively more severely, imposing a practical limit on frequencies which may be used for air sonar systems. Higher temperatures in general increase the absorption. Changes in humidity affect the molecular vibrational relaxation frequencies, providing the most complex variations. For a full discussion of the physical mechanisms causing the absorption, refer to ANSI S1.26-1978 [2] Appendix B.

E.5 Speed Of Sound

A formula for the speed of sound in air is provided by Cramer [26]. While an overview was presented in Chapter 12, the details are described here. Cramer's formulation was found by using the best available values (from various sources) for all of the gas constituents of the standard atmosphere (ISO 2533-1975 [65]) to calculate values for γ , M and B in (12.3). He then fitted a polynomial approximation, valid over a small region of standard atmospheric parameters. This provides the speed of sound c_0 at zero frequency (see Cramer [26] eqn (15))

$$\begin{aligned} c_0(T_c, p_s, x_w, x_c) = & a_0 + a_1 T_c + a_2 T_c^2 + (a_3 + a_4 T_c + a_5 T_c^2) x_w \\ & + (a_6 + a_7 T_c + a_8 T_c^2) p_s + (a_9 + a_{10} T_c + a_{11} T_c^2) x_c \\ & + a_{12} x_w^2 + a_{13} p_s^2 + a_{14} x_c^2 + a_{15} x_w p_s x_c, \end{aligned} \quad (\text{E.11})$$

in $\text{m} \cdot \text{s}^{-1}$, where the coefficients a_0 - a_{15} are provided in Table E.1 and the parameters and their ranges of validity are listed below.

T_c	is the air temperature in degrees Celsius, valid range 0°C to 30°C .
p_s	is the static air pressure, valid range 75 000 Pa to 102 000 Pa. This formula specifically requires p_s to be expressed in Pascals.
x_w	is the mole fraction of water vapour, valid range 0 to 0.06 (dimensionless), given by (E.12) below.
x_c	is the mole fraction of CO_2 , valid range 0 to 0.01 (dimensionless), typical value is 0.000 314 (from Cramer [26]).




Figure E.1: Isopleth plots of the variation in the absorption of sound in air α' , in $\text{dB} \cdot \text{m}^{-1}$, as a function of frequency (horizontal axis) and humidity (vertical axis). The pressure increases across the table from left to right, as labelled on the horizontal axis of each plot. The temperature increases up the table. The isopleths are placed at intervals of $0.5 \text{ dB} \cdot \text{m}^{-1}$, and the areas are shaded according to the legend shown.

Table E.1: Coefficients from Cramer [26] for use with (E.11).

Coefficients	Values
a_0	331.5024
a_1	0.603 055
a_2	−0.000 528
a_3	51.471 935
a_4	0.149 5874
a_5	−0.000 782
a_6	$−1.82 \times 10^{-7}$
a_7	3.73×10^{-8}
a_8	$−2.93 \times 10^{-10}$
a_9	−85.209 31
a_{10}	−0.228 525
a_{11}	5.91×10^{-5}
a_{12}	−2.835 149
a_{13}	$−2.15 \times 10^{-13}$
a_{14}	29.179 762
a_{15}	0.000 486

The mole fraction of water vapour may be calculated from the relative humidity using (see Cramer [26] eqn (A1))

$$x_w = \frac{h_r}{100} f_e \frac{p_{sat}}{p_s}. \quad (\text{E.12})$$

The quantity x_w is similar to h given by (E.4), except that x_w is expressed as a fraction and includes the enhancement factor f_e . The enhancement factor f_e is calculated by (see Cramer [26] eqn (A2))

$$f_e = 1.00062 + 3.14 \times 10^{-8} p_s + 5.6 \times 10^{-7} T_c^2. \quad (\text{E.13})$$

E.5.1 Uncertainties

Cramer claims that (E.11) is accurate to ± 300 ppm or about $\pm 0.10 \text{ m} \cdot \text{s}^{-1}$ for dry air at STP. Wong [141], however, claims that Cramer neglected the error in some of the constants used to derive (E.11), and that the error bound is instead ± 545 ppm or $\pm 0.18 \text{ m} \cdot \text{s}^{-1}$ for dry air at STP.

E.6 Correction for Dispersion

Cramer notes that the correction for dispersion, as a function of frequency, is

$$\frac{1}{c_0} - \frac{1}{c_\phi} = \frac{\alpha_{v,N}}{2\pi f_{r,N}} + \frac{\alpha_{v,O}}{2\pi f_{r,O}}. \quad (\text{E.14})$$

Plots of the variation of the speed of sound as a function of frequency and the four environmental parameters are shown in Figure E.2. The effects are discussed in Section 12.2.

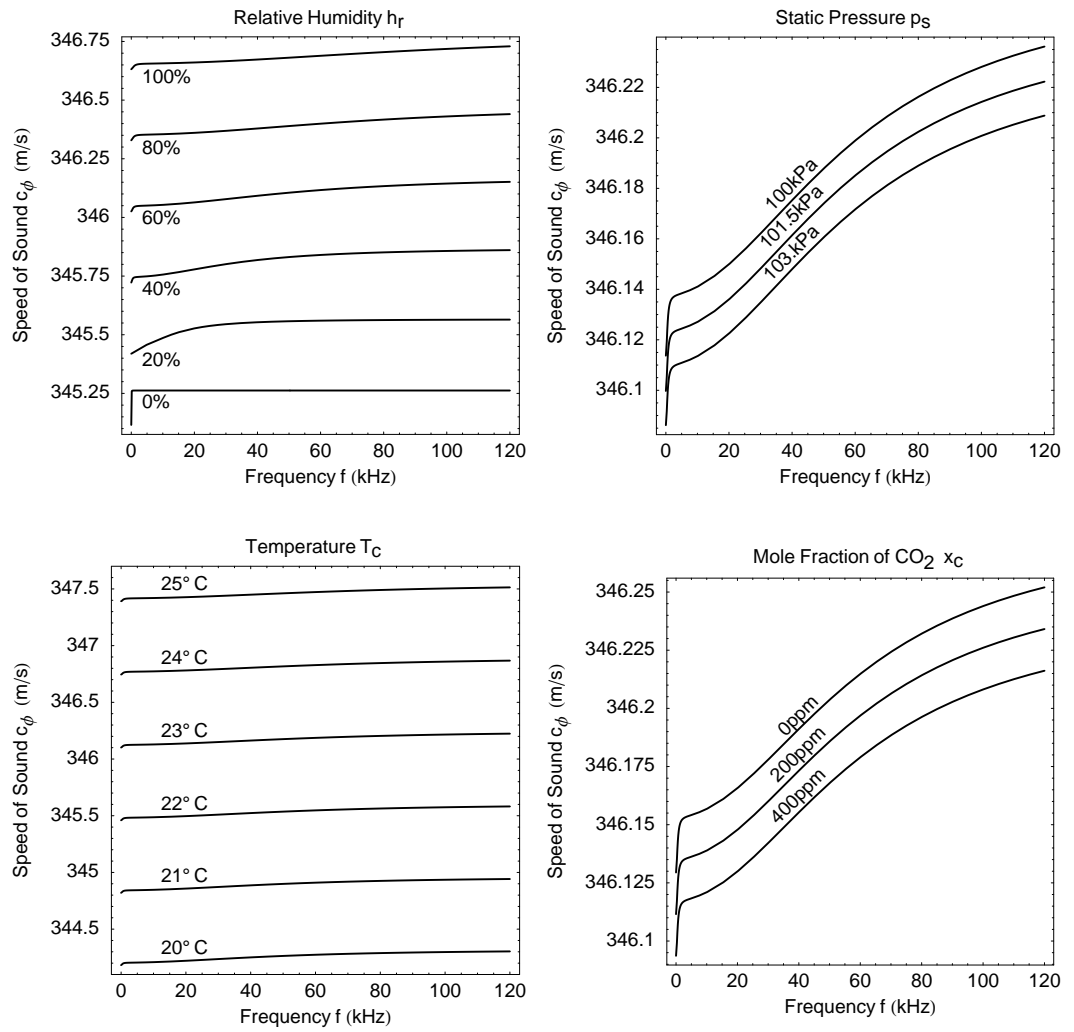


Figure E.2: Variation in the speed of sound c_ϕ as a function of frequency, calculated by (E.14). The four panels demonstrate the effects of varying each of the environmental parameters T_c , h_r , p_s and x_c in turn. The remaining parameters are held constant at $T_c = 23^\circ\text{C}$, $h_r = 65\%$, $p_s = p_{s0} = 101.325\text{ kPa}$ and $x_c = 314\text{ ppm}$, which are approximate average values of the experimental data.

Appendix F

Calibration Device

The measurement of the environmental parameters was validated by comparing the modelled value of the speed of sound to an independently measured speed of sound. The modelled speed of sound was computed using the environmental parameters, as described in Section 12.4. The independent value for the speed of sound was measured directly by the sonar system using a calibration device, which is described in this appendix.

F.1 Theory

To validate the measurement of the environmental parameters an independent estimate of the speed of sound, c , is obtained from the sonar. This process also estimates the total delay, t_{off} , present in the system. These quantities, c and t_{off} , may be derived from the relationship between DOF, d , and TOF, τ ,

$$\tau = \frac{1}{c}d + t_{off}. \quad (F.1)$$

An experiment utilising this relationship is constructed using four fixed targets, each having DOF d_i . The sonar is used to measure the TOF, τ_i , of each target, using a single sonar observation. A line fitted to the points $\{d_i, \tau_i\}$ will have a slope $1/c$ and an offset t_{off} .

The ability to calculate the speed of sound, c , directly from sonar measurements provides a convenient method of validating that the environmental parameters used to calculate the speed of sound derived in Section 12.2 are measured correctly.

F.2 Design

The device shown in Figure F.1 provides four separate targets. The device can only be used with the left and right receivers, due to the vertical displacement of the reflection geometry. The top and bottom receivers may be used, if necessary, by using a rotation bracket which allows the sensor head to be rotated by 90° around the principal direction of the transmitter.

F.2.1 Manufacture

Timber was used to build the calibration device. Man-made board (MDF) was selected for the base to minimise warp and changes in size due to humidity changes. The risers were cut from solid




Figure F.1: The design of the calibration device. The poles are slanted to provide reflection points at tangents to the curves drawn, without obscuring targets at greater distances. Dashed lines show the sonar path, which was designed with the stated range r_i .




Figure F.2: The complete calibration device, with the sonar head mounted at one end. Note the strips of MDF along the sides of the spacers between the risers.

pine as this was available in a convenient size. The spacers between the risers were cut from the same stock, and attached to the MDF base using glue and dowels. Strips of MDF were attached to the sides of the spacers (see Figure F.2) to prevent warping along the length. The dimensions of the device are shown in Figure F.3.

The angles on the risers and spacers were cut using a mitre circular saw, and verified to be accurate using an angle measuring gauge. The length of the spacers between the risers was compensated for the exact thickness of the riser it was to mate with to maintain the correct horizontal spacing shown in Figure F.3.

The precision positioner described in Section 5.2 has key holes at the mounting points to ensure repeatable mounting of the sonar equipment. These key holes were reproduced in the base of the calibration device so that the sonar could be accurately positioned.




Figure F.3: Calibration device dimensions, in millimetres.

F.3 Results

F.3.1 Measuring Target DOFs

The completed calibration device was found to change size slightly with the amount of humidity absorbed by the timber. Therefore the target ranges, x_i , must be re-measured on a daily basis before the device is used. (While the air humidity may vary quickly, the humidity in the timber changes more slowly.) The measurement was carried out using a 1 m steel rule, which was held along the sides of the poles. The zero mark of the rule was aligned with the transmitter, and the distance along the normal to the target recorded. One person is required to sight each measurement point of the ruler. The normal may be found by rotating the ruler around the transmitter point and obtaining the shortest range measurement. The rule was braced by a piece of angle aluminium to avoid flexing while the measurement was carried out.

The geometrical DOF, d_i , corresponding to the range, x_i , of each target is computed using

$$d_i = \sqrt{(2x_i + d_T)^2 + d_R^2}, \quad (\text{F.2})$$

(which is the same as (13.1)) where

d_T, d_R are the distances of the transmitter and the receiver from the head centre point P_{HC} as shown in Figure 7.12, and

x_i is the range from the transmitter to the target, measured along the dashed lines shown in Figure F.1.

Some typical measurements obtained using the method described above are shown in the measured range column of Table F.1. The DOFs computed according to (F.2) are also shown in the table.

F.3.2 Measuring Time Of Flight

The CTFM sonar system was used, with a Yule-Walker spectral estimator, to measure the times of flight to each target. The output of the sonar system corresponding to one sonar measurement is shown in Figure F.4.

Table F.1: The measured ranges, x_i , to each of the calibration device’s four targets, and their geometrical DOFs, d_i , computed using (F.2). The head dimensions used to calculate the d_i were $d_T = 2\text{ mm}$ and $d_R = 15\text{ mm}$.

Designed Range	Measured Range x_i	Geometrical DOF d_i
mm	mm	mm
400.0	397.0	796.14
600.0	595.1	1192.3
800.0	796.3	1594.67
1000.0	996.0	1994.06




Figure F.4: Calibration device TOFs measured by the CTFM sonar system. Data from the left and right channels are presented. The black line represents the spectrum estimated by the Yule-Walker spectral estimator. The horizontal scale has been converted from frequency to the corresponding TOF. The red lines indicate the peaks which were found in the spectrum and their TOFs. The peaks near 0s are crosstalk. The next 4 large peaks are echoes from the calibration device. The reflection from the corner at the base of the first pole is visible just after the first pole reflection. There are several peaks after the four peaks due to the risers, which are multi-bounce echoes, between the sonar head and a riser.

F.3.3 Slope Results

The experiment was performed under computer control and 200 sonar measurements of the four poles were recorded. The environment sensor was used to record values of the environment parameters simultaneously with the sonar measurements. The speed of sound c was obtained from both the straight line fit and from the environment sensor using (E.14). The results obtained when the temperature was measured by an LM335 transducer attached to the aluminium case are shown in Figure 12.5a. This clearly indicates an error in a parameter of the speed of sound modelled by (E.14). As the temperature is the most important parameter in the speed of sound model, the measurement of temperature was checked. These investigations led to the sequence of improvements in the measurement of temperature which were described in Section 12.4.1.

The results obtained after the LM335 transducer was removed from the aluminium case, and instead mounted in the draught of a fan, are shown in Figure 12.5b. The improvement in the match between the two values of the speed of sound indicates that the lag in the measured temperature decreases when the fan is used.

The improvements made to the measurement of the temperature were made possible by the independent measurement of the speed of sound provided by the calibration device.

F.4 Problems

During a period of higher than normal humidity it was found that there was a significant difference in the speed of sound measured by the calibration device and by the environment sensor, where there had been agreement two days before. Inspection of the calibration device revealed a gap between its base and the table it was sitting on. The base had developed an arch.

The arching may be explained by examining the construction of the device. The pine used as spacers between the risers (Figure F.2) expands more when it absorbs moisture than the MDF base does. The pine spacers and MDF base then bend in the same way as a bi-metallic strip. The spacing of the risers is spread, causing a scaling error in the geometric DOFs. The scaling error minimally affects the computed offset time t_{off} , but directly affects the computed speed of sound.

The warping problem may be resolved by requiring that the measurement procedure described in Section F.3.1 be carried out on each day that the device will be used. The errors in the geometric DOFs will be eliminated by doing this. Thus the warping is not seen as a serious problem.

The warping problem could be eliminated entirely in future designs by utilising a single material which has a low thermal coefficient of expansion. Non-porous materials such as metals and plastics are also impervious to changes in humidity. The selection of a material meeting these criteria would allow the calibration device to be used without daily ruler measurements.

Appendix G

Improved Head Design

Section 14.1.2 discussed the errors due to treating all target types as plane reflectors. It is possible to design the sonar head so that these errors are minimised. Some exploration of the sonar head design parameter space using the error analysis technique of Section 14.1.2 shows that the errors are made much smaller when the transmitter is placed in-line with the receivers ($d_T = 0$, see Figure G.1). The splay angle of the receivers is maintained to support the IPD bearing technique which is described in Section 14.2. The receiver distance $d_R = 15$ mm also remains constant.

The formulae for target ranges and bearings may be simplified considerably when $d_T = 0$. The simplified formulae for each target type are described in the next two sections.

G.1 Planes and Ninety Degree Corners

The solution to the forward problem of computing the DOFs for plane targets and 90° corners with a position $\{r, \theta\}$ are

$$d_1 = \sqrt{d_R^2 + 4r^2 - 4d_R r \sin(\theta)} \quad (\text{G.1})$$

and

$$d_2 = \sqrt{d_R^2 + 4r^2 + 4d_R r \sin(\theta)}. \quad (\text{G.2})$$

The solution to the reverse problem, that of computing $\{r, \theta\}$ from measured DOFs d_1 and d_2 , is simplified as P_{HC} and T are now coincident. Thus $r = r'$ and $\theta = \theta'$, and the solutions are given directly by (14.5) and (14.6).




Figure G.1: The improved head design with $d_T = 0$ which reduces the error caused by not classifying targets. Note that the transmitter is located between the centres of the two receivers. Compare with Figure 7.12.

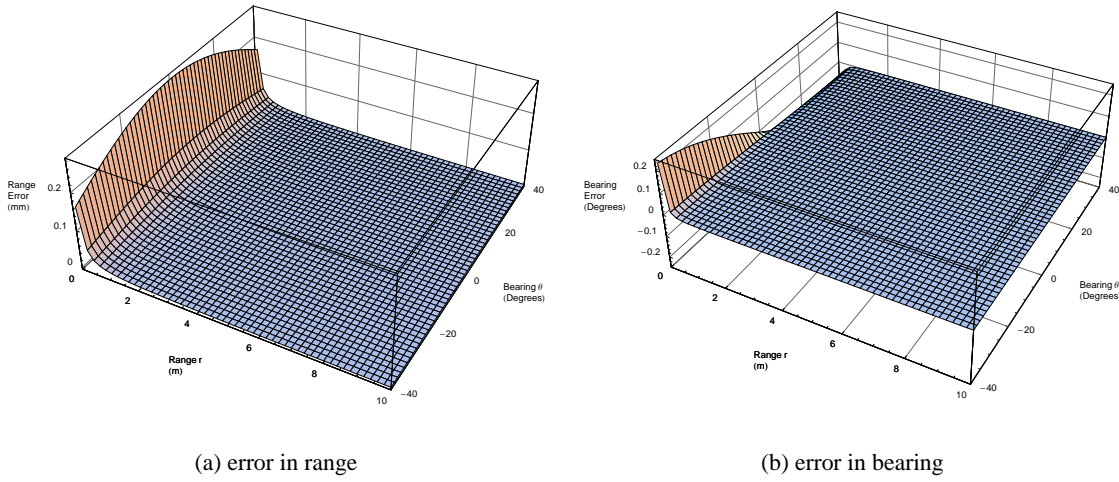


Figure G.2: The errors in range and bearing due to considering a point reflector as a plane reflector, for a sensor head with all transducers in-line. Compared with Figure 14.7, the bearing error is reduced for close range targets, and the systematic range error vanishes.

G.2 Edge Reflectors

The solution to the forward problem for an edge or point reflector is given by

$$d_1 = r + \sqrt{d_R^2 + r^2 - 2d_R r \sin(\theta)} \quad (\text{G.3})$$

and

$$d_2 = r + \sqrt{d_R^2 + r^2 + 2d_R r \sin(\theta)}. \quad (\text{G.4})$$

The position of the reflecting edge, $\{r, \theta\}$, may be found from the DOFs by

$$r = \frac{d_1^2 + d_2^2 - 2d_R^2}{2(d_1 + d_2)} \quad (\text{G.5})$$

and

$$\theta = -\sin^{-1} \left[\frac{(d_1 - d_2)(d_R^2 + d_1 d_2)}{d_R(d_1^2 + d_2^2 - 2d_R^2)} \right]. \quad (\text{G.6})$$

These are much simpler than (14.13) and (14.14), which calculate the same quantities but for a head with non-zero d_T .

G.3 Error Results

The bearing and range errors due to treating the edge reflector as a plane for the modified system are plotted in Figure G.2.

The bearing error has shrunk to 20% of its former value at its worst. The range error is now mostly zero, having a maximum value of 0.2mm for very close targets. Only for close targets,

with ranges of less than 0.1 m, are the range and bearing errors significant. The errors obtained using the in-line transducers is significantly less than those obtained by the unmodified design.

There are now two reasons to design sonar heads with the transmitter in line with the receivers. The first is that the formulae for computing the range and bearing to targets are vastly simplified compared to the forms (14.13) and (14.14). The second is that the error consequences of ignoring the target type when computing the bearing to a target are minimised.

While it was not possible to re-manufacture the sonar head for this work, future work should take into consideration the modifications suggested in this appendix.

Appendix H

Modelling of Transducer Path Angles

In Section 14.2.2, a Gaussian function was used to model the received echo power from a target as its bearing was changed. The method of estimating the Gaussian fit parameters is described in Section H.1.

The Gaussian function which is fitted to the received echo power data combines the effects of the two transducers traversed by the sonar path. The effects upon the signal of each individual transducer are derived in Section H.2, which assumes that the transducer parameters are known. Unlike the model used in Section 14.2.2, the model derived in Section H.2 incorporates the target range in addition to the target bearing. The two models converge to the same result for targets at sufficiently large range.

An important property of Gaussian functions, namely the product of two Gaussian functions, is derived in Section H.3. This allows the two separate transducer models derived in Section H.2 to be combined to model the whole effect upon the signal.

Section H.4 constructs formulae relating the measured IPD curves from the left and right channels to the Gaussian functions of each of the three transducers, incorporating the transducer incidence angles to a zero order approximation. This model can predict the IPD curve using the beam width, beam power and pointing angle of each of the three transducers.

Section H.5 demonstrates a method of inverting this relationship, allowing us to solve for the Gaussian beam pattern of each of the transducers individually, given only the received power curves from the left and right transducers, as shown in Figure 14.18. The results were used to construct transducer and IPD models suitable for testing the correspondence solver in Chapter 15.

H.1 Fitting for the Gaussian Parameters

This section discusses the problem of fitting the parameters a , b and θ_0 to the experimental echo amplitude data points from one channel. The parameters are described in Section 14.2.2. The curve fitting problem is non-linear in the amplitude space (Figure 14.17), but is linearised in the decibel power space (Figure 14.18). The Gaussian function may be converted to decibels by

$$20\log_{10}(f_{\text{gaus}}(\theta, a, \theta_0, b)) = 20\log_{10}\left(a \exp\left\{-\left(\frac{\theta - \theta_0}{b}\right)^2\right\}\right).$$

Expanding this and re-writing as a polynomial in θ yields

$$20\log_{10}(f_{\text{gaus}}(\theta, a, \theta_0, b)) = \frac{20}{\ln 10} \left(\ln a - \frac{\theta_0^2}{b^2} \right) + \frac{20}{\ln 10} \cdot \frac{2\theta_0}{b^2} \theta - \frac{20}{\ln 10} \cdot \frac{1}{b^2} \theta^2, \quad (\text{H.1})$$

which has linearised the fitting problem. Therefore a polynomial of the form

$$P(\theta) = \beta_0 + \beta_1 \theta + \beta_2 \theta^2 \quad (\text{H.2})$$

is fitted to the power data points

$$P_i = 20 \log_{10} A_i.$$

The fit is performed using the methods of Appendix I. The points used in the fit must not include data points past the point of separation (Section 7.2.2). The Gaussian parameters may be obtained from the fit parameters, β_i , by equating coefficients of equal powers of θ in (H.1) and (H.2), obtaining

$$b = \sqrt{\frac{20}{\ln 10} \cdot \frac{-1}{\beta_2}},$$

$$\theta_0 = \frac{b^2 \ln 10}{2} \frac{\beta_1}{20}$$

and

$$a = \exp\left(\frac{\ln 10}{20} \beta_0 + \frac{\theta_0^2}{b^2}\right).$$

Thus the parameters of the Gaussian function may be found using a linear fit to the echo power data.

H.2 Transducer Path Angles

The sonar paths for a plane target, shown in Figure H.1 and Figure H.2, subtend different angles to the normal of each transducer. We seek to derive these angles as a function of the target coordinates, $\{r, \theta\}$, the dimensions of the sonar head d_T and d_R , and the receiver splay angles, θ_{0R1} and θ_{0R2} .

The points appearing in Figures H.1 and H.2 are listed below.

- T is the transmitter.
- T_V is the image of the transmitter, T , reflected in the plane target.
- R_1, R_2 are the left and right receivers (channels 1 and 2).
- P_T, P_{R1}, P_{R2} are the intersection point of normals to the plane target drawn through the transducers T, R_1 and R_2 .
- B_1, B_2 are the points at which the left and right sonar paths are reflected from the plane target, obeying the law of specular reflection.
- P_{HC} is the point mid-way between the two receivers, R_1 and R_2 .
- A_1 is a point completing the rectangle formed by $TP_T P_{R1} A_1$. A similar point A_2 may be defined to complete the rectangle $TP_T P_{R2} A_2$.

The dimensions are




Figure H.1: Diagram showing the sonar paths from the transmitter to the receivers. The outline of the real sonar head is shown on the left, while the virtual sonar head is shown on the right. The plane target is shown as a vertical line in the middle. Lines parallel to TT_V are drawn through the receivers, R_1 and R_2 , to assist with the derivation. Further details are shown in Figure H.2.




Figure H.2: Closeup of the sonar path TB_1R_1 , showing the angles subtended to the normals of the transducers, and all of the details necessary to derive the relationship between them. The outline of the sonar head is not shown on this diagram. See text for more information.

- d_T, d_R , the parameters of the sonar head as defined in Section 7.3,
- d_{TR} , the distance between the transmitter and the receivers, defined by $d_{TR} = \sqrt{d_T^2 + d_R^2}$, and
- r , the range to the target, being the distance TP_T .

The angles are listed below.

- θ is the angle between the normal of the transmitter and the normal of the target passing through the transmitter, TP_T . This is the observed angle to the target.
- ϕ is the angle $P_{HC}R_1T$, calculated by $\phi = \tan^{-1}(d_T/d_R)$.
- β_1 is the angle P_TB_1 , which is geometrically equivalent to $B_1R_1P_{R1}$ and TT_VR_1 . A similar angle β_2 is defined for the right channel. The angle $2\beta_1$ is the specular reflection angle.
- α_1 is the angle A_1R_1T , which is defined by $\alpha_1 = \phi - \theta$. A similar quantity $\alpha_2 = \phi + \theta$ is defined for the right channel.
- θ_{0R1} is the splay angle between the transmitter normal angle and the receiver R_1 normal angle. A similar angle θ_{0R2} is defined for the right channel.
- θ_{R1} is the angle between the left channel echo path and the receiver R_1 normal angle. As positive angle is anti-clockwise relative to the transducer normal, the quantity labelled here has a negative sign. A similar quantity θ_{R2} is defined for the path to the receiver R_2 .
- θ_{T1} is the angle between the left channel echo path $TP_{B1}R_1$ and the transmitter T normal angle. A similar quantity θ_{T2} is defined for the path to the receiver R_2 .

The triangle A_1R_1T can be used to compute the distances A_1R_1 and A_1T . The angle β_1 is then found by applying the tangent function, \tan , to the right triangle $A_1T_VR_1$, yielding

$$\tan \beta_1 = \frac{d_{TR} \cos(\phi - \theta)}{2r + d_{TR} \sin(\phi - \theta)}. \quad (\text{H.3})$$

The desired angles θ_{T1} and θ_{T2} are

$$\theta_{T1} = \theta + \beta_1 \quad (\text{H.4})$$

and

$$\theta_{R1} = \theta - \beta_1 - \theta_{0R1}. \quad (\text{H.5})$$

Similar working for the right channel path $TP_{B1}R_1$ yields

$$\tan \beta_2 = \frac{d_{TR} \cos(\phi + \theta)}{2r + d_{TR} \sin(\phi + \theta)}, \quad (\text{H.6})$$

$$\theta_{T2} = \theta - \beta_2 \quad (\text{H.7})$$




Figure H.3: Plots showing the behaviour of the specular reflection angles β_1 and β_2 as functions of target range r and bearing θ .

and

$$\theta_{R2} = \theta + \beta_2 - \theta_{0R2}. \quad (\text{H.8})$$

This provides the exact path angles to the transducer normals as functions of the head geometry and the target position.

Plots of the specular reflection angles β_1 and β_2 for various target ranges are shown in Figure H.3. The dependence of β_1 and β_2 upon range indicates that the IPD model may require compensation for target range in order to remain accurate at ranges other than the calibration range. However, the variation with θ is much smaller and may potentially be ignored. The impact of β_1 and β_2 upon the accuracy of the IPD model have not been investigated.

H.3 Gaussian Products

This section demonstrates that the product of two Gaussian functions is itself a Gaussian function. More importantly, it provides the algebraic relationship between the parameters of the resulting Gaussian and the parameters of the two factors of which it is composed.

The Gaussian function used is equation (14.26), repeated here for convenience,

$$f_{\text{gaus}}(\theta, a, \theta_0, b) = a \exp \left\{ - \left(\frac{\theta - \theta_0}{b} \right)^2 \right\}, \quad (\text{H.9})$$

where a controls the height of the peak, b controls the peak width (the points of inflexion occur at $\theta_0 \pm \sqrt{2}b$), and θ_0 controls the offset of the peak from $\theta = 0$.

We wish to derive the product of two Gaussian functions. To begin, take

$$f_{\text{prod}} = f_{\text{gaus}}(\theta, a_1, \theta_{01}, b_1) \cdot f_{\text{gaus}}(\theta, a_2, \theta_{02}, b_2). \quad (\text{H.10})$$

Substituting in the definition of f_{gaus} from (H.9) yields

$$f_{\text{prod}} = a_1 \exp \left\{ - \left(\frac{\theta - \theta_{01}}{b_1} \right)^2 \right\} \cdot a_2 \exp \left\{ - \left(\frac{\theta - \theta_{02}}{b_2} \right)^2 \right\}.$$

Simplifying using the rules of exponentials yields

$$f_{prod} = a_1 a_2 \exp \left\{ - \left(\frac{\theta - \theta_{01}}{b_1} \right)^2 - \left(\frac{\theta - \theta_{02}}{b_2} \right)^2 \right\}.$$

Bringing the fractions over a common denominator gives

$$f_{prod} = a_1 a_2 \exp \left\{ - \frac{b_2^2 (\theta^2 - 2\theta\theta_{01} + \theta_{01}^2) + b_1^2 (\theta^2 - 2\theta\theta_{02} + \theta_{02}^2)}{b_1^2 b_2^2} \right\}.$$

Grouping terms in powers of θ yields

$$f_{prod} = a_1 a_2 \exp \left\{ - \frac{(b_2^2 + b_1^2) \theta^2 - 2(b_2^2 \theta_{01} + b_1^2 \theta_{02}) \theta + b_2^2 \theta_{01}^2 + b_1^2 \theta_{02}^2}{b_1^2 b_2^2} \right\}.$$

Dividing through the exponential term by $b_2^2 + b_1^2$ gives

$$f_{prod} = a_1 a_2 \exp \left\{ - \frac{\theta^2 - 2 \frac{b_2^2 \theta_{01} + b_1^2 \theta_{02}}{b_2^2 + b_1^2} \theta + \frac{b_2^2 \theta_{01}^2 + b_1^2 \theta_{02}^2}{b_2^2 + b_1^2}}{\frac{b_1^2 b_2^2}{b_2^2 + b_1^2}} \right\}. \quad (H.11)$$

At this point it is best to change variables. Fortuitous knowledge of the result suggests substituting

$$\theta_{0c} = \frac{b_2^2 \theta_{01} + b_1^2 \theta_{02}}{b_2^2 + b_1^2} \quad (H.12)$$

and

$$b_c^2 = \frac{b_1^2 b_2^2}{b_2^2 + b_1^2} \quad (H.13)$$

and

$$c = \frac{b_2^2 \theta_{01}^2 + b_1^2 \theta_{02}^2}{b_2^2 + b_1^2}. \quad (H.14)$$

Performing these substitutions, equation (444) is simplified to

$$f_{prod} = a_1 a_2 \exp \left\{ - \frac{\theta^2 - 2\theta_{0c}\theta + c}{b_c^2} \right\}.$$

Completing the square yields

$$f_{prod} = a_1 a_2 \exp \left\{ - \frac{(\theta - \theta_{0c})^2 + c - \theta_{0c}^2}{b_c^2} \right\}.$$

Extracting the term which is independent of θ from the exponential yields

$$f_{prod} = a_1 a_2 \exp \left\{ \frac{\theta_{0c}^2 - c}{b_c^2} \right\} \cdot \exp \left\{ - \frac{(\theta - \theta_{0c})^2}{b_c^2} \right\}.$$

Letting the factor which is independent of θ be a_c gives the final form

$$f_{prod} = f_{guas}(\theta, a_c, \theta_{0c}, b_c) = a_c \exp \left\{ - \left(\frac{\theta - \theta_{0c}}{b_c} \right)^2 \right\}, \quad (H.15)$$

which is in the same form as the Gaussian function (H.9).

The amplitude coefficient a_c may be found by substituting θ_c , b_c and c from (H.12), (H.13) and (H.14) and simplifying to yield

$$a_c = a_1 a_2 \exp \left\{ -\frac{(\theta_{01} - \theta_{02})^2}{b_1^2 + b_2^2} \right\}. \quad (\text{H.16})$$

Therefore the product of two Gaussian functions (H.9) is given by (H.15), with the parameters a_c , b_c and θ_{0c} defined by (H.16), (H.13) and (H.12) respectively.

H.4 Modelling the Combined Beam Pattern

This section relates the Gaussian parameters of the left and right sonar paths to the Gaussian parameters of each of the three transducers. The relationship is built incorporating a zero order approximation of the angles of incidence, β_1 and β_2 , derived in Section H.2.

Each of the three transducers, R_1 , R_2 and T , is assumed to have a Gaussian beam pattern $f_{\text{gaus}}(\theta, a, \theta_0, b)$, as defined in Section 14.2.2. The Gaussian has its centre at the transducer normal angle, defined by θ_0 . The angle at which the sonar beam enters or leaves the transducer relative to the transducer normal, derived in Section H.2, determines the response of the transducer. Assuming that the attenuation caused by the spreading loss and air absorption is common to both left and right sonar paths, the difference between the transducer responses is the dominant factor controlling the ratio of the echo power between the two channels. Thus, by determining the angles of the sonar path, the relative powers of the left and right echoes can be determined.

Referring to Figure H.2 and following the path $TP_{B1}R_1$, the received echo amplitude will be

$$f_{\text{gaus}}(\theta, a_L, \theta_{0L}, b_L) = f_{\text{gaus}}(\theta_{T1}, a_T, 0, b_T) \cdot f_{\text{gaus}}(\theta_{R1}, a_{R1}, 0, b_{R1}), \quad (\text{H.17})$$

where θ_{T1} , and θ_{R1} are defined by (H.4) and (H.5) respectively.

These formulae use θ in a nonlinear way. If (H.17) utilised θ with just an additive offset, then it would be possible to use the Gaussian combination relationships derived in Section H.3 to simplify (H.17). Seeking an approximation to β_1 which is independent of θ , we define a zero order approximation β_{10} as

$$\beta_{10} = \frac{d_{TR} \cos(\phi)}{2r + d_{TR} \sin(\phi)}. \quad (\text{H.18})$$

Using this approximation allows (H.17) to be re-written as

$$f_{\text{gaus}}(\theta, a_L, \theta_{0L}, b_L) = f_{\text{gaus}}(\theta + \beta_{10}, a_T, 0, b_T) \cdot f_{\text{gaus}}(\theta - \beta_{10} - \theta_{0R1}, a_{R1}, 0, b_{R1}). \quad (\text{H.19})$$

Referring to the definition of f_{gaus} for simplification allows the additions to θ to be moved to θ_0 , yielding

$$f_{\text{gaus}}(\theta, a_L, \theta_{0L}, b_L) = f_{\text{gaus}}(\theta, a_T, -\beta_{10}, b_T) \cdot f_{\text{gaus}}(\theta, a_{R1}, \theta_{0R1} + \beta_{10}, b_{R1}). \quad (\text{H.20})$$

A similar expression for the path to the right receiver is found to be

$$f_{\text{gaus}}(\theta, a_R, \theta_{0R}, b_R) = f_{\text{gaus}}(\theta, a_T, \beta_{20}, b_T) \cdot f_{\text{gaus}}(\theta, a_{R2}, \theta_{0R2} - \beta_{20}, b_{R2}), \quad (\text{H.21})$$

where $\beta_{20} = \beta_{10}$ identically. These formulae equate a Gaussian to the product of two Gaussian functions. Section H.3 demonstrated that these forms are equivalent, and provided a relationship between the parameters of the three Gaussian functions. For the left sonar path the combined effect is

$$a_L = a_T a_{R1} \exp \left\{ -\frac{(-2\beta_{10} - \theta_{0R1})^2}{b_T^2 + b_{R1}^2} \right\}, \quad (\text{H.22})$$

$$\theta_{0L} = \frac{b_T^2 (\theta_{0R1} + \beta_{10}) - b_{R1}^2 \beta_{10}}{b_T^2 + b_{R1}^2} \quad (\text{H.23})$$

and

$$b_L^2 = \frac{b_T^2 b_{R1}^2}{b_T^2 + b_{R1}^2}. \quad (\text{H.24})$$

An equivalent set of equations can be found for the right channel by substituting R for L subscripts and $R1$ for $R2$ subscripts and negating all terms involving β_{10} .

Thus the parameters of the Gaussian function representing the amplitude of the left and right paths may be determined from knowledge of the Gaussian functions representing the beams of the three transducers.

H.5 Solving For Transducer Gaussian Parameters

This section provides a method of solving for the Gaussian parameters of the transmitter and the two receivers from the Gaussian parameters fitted to the left and right channel paths in Section H.1. The solution is performed by inverting equations (H.22), (H.23) and (H.24) derived in the previous section. This procedure requires some additional information or assumptions to obtain a unique solution.

The experimental data shown in Figure 14.17 provides values for $\{a_L, \theta_{0L}, b_L\}$ and $\{a_R, \theta_{0R}, b_R\}$. To model the Gaussian functions of the transmitter and the two receivers for arbitrary targets, it is necessary to calculate values of $\{a_T, \theta_{0T}, b_T\}$, $\{a_{R1}, \theta_{0R1}, b_{R1}\}$ and $\{a_{R2}, \theta_{0R2}, b_{R2}\}$ from the values measured from the left and right paths. There are six known parameters and nine unknown variables, so there are three degrees of freedom. Three assumptions are made:

1. The transmitter pointing angle is assumed to be 0° , i.e. $\theta_{0T} = 0$.
2. The transmitter T and receiver R_1 beam widths are equal, i.e. $b_T = b_{R1}$. The left receiver is chosen arbitrarily.
3. The transmitter power is equal to that of the left IPD curve, i.e. $a_T = a_{R1}$.

These assumptions reduce the number of degrees of freedom to 0 and uniquely determine the solution. Other assumptions may lead to a different solution. The assumptions which are made are not important as long as the IPD curve which is produced by the separately modelled transducers matches the one which is determined experimentally.

Combining equations (H.22), (H.23) and (H.24) with the assumptions described above, and solving for the transmitter T beam parameters in terms of the known variables yields

$$a_T = \sqrt{a_L} \exp \left\{ \frac{(\beta_{10} + \theta_{0L})^2}{2b_L^2} \right\}, \theta_{0T} = 0, b_T = \sqrt{2}b_L. \quad (\text{H.25})$$

Solving the equations for the left receiver R_1 parameters yields

$$a_{R1} = \sqrt{a_L} \exp \left\{ \frac{(\beta_{10} + \theta_{0L})^2}{2b_L^2} \right\}, \theta_{0R1} = 2\theta_{0L}, b_{R1} = \sqrt{2}b_L, \quad (\text{H.26})$$

while solving for the right receiver R_2 parameters obtains

$$a_{R2} = \frac{a_R}{\sqrt{a_L}} \exp \left\{ \frac{(\beta_{10} - \theta_{0R})^2}{2b_L^2 - b_R^2} - \frac{(\beta_{10} + \theta_{0L})^2}{2b_L^2} \right\}, \quad (\text{H.27})$$

$$\theta_{0R2} = \frac{2\beta_{10}(b_L^2 - b_R^2) + 2b_L^2\theta_{0R}}{2b_L^2 - b_R^2} \quad (\text{H.28})$$

and

$$b_{R2}^2 = \frac{2b_L^2 b_R^2}{2b_L^2 - b_R^2}.$$

These results allow suitable beam functions to be estimated for the transmitter and both of the receivers using only the observed experimental curves obtained during the **IPD** calibration procedure in Section 14.2.2. These Gaussian models may be utilised to model the relative echo power between the two channels which will be detected from a target.

H.6 Conclusion

A more complete echo power model could be built by incorporating the effects of air absorption and spreading loss. A full echo power model, combined with the geometrical model of Appendix D, would allow echoes to be verified against known targets, and echoes from unknown targets identified. However, this model was developed expressly to test the correspondence solver in Chapter 15, and hence only the echo power variation with target bearing was required to be modelled.

Appendix I

Curve Fitting with Error Analysis

This appendix initially, in Section I.1, describes the nature of experimental errors, and shows how they may be represented. Section I.2 describes the calculation of the error in the output of a function whose inputs are themselves measured parameters with associated errors. Section I.3 describes how curves may be linearised and fitted to experimental data points, including the use of unequal weights where some points are known more precisely than others. This section also provides the **Variance CoVariance (VCV)** matrix of the fit parameters, which describes the precision with which the fit is known. Section I.4 describes a method of interpreting the **VCV** matrix. Section I.5 describes how to calculate the error in y , incorporating the error from both the stated error, Δx , in x , and the fitted curve. Section I.6 describes how the reverse procedure may be performed, calculating Δx from the curve fit parameters and Δy . Section I.7 describes how to evaluate the fit which has been achieved by the fitted curve, through examination of the residuals. Finally, Section I.8 describes how to perform a weighted average between two measurements whose mean and standard deviation are known.

I.1 The Nature of Errors

Experiments with real quantities involve errors from various sources, that is, instead of measuring the true value X , the measurement consists of

$$X + \varepsilon_x, \tag{I.1}$$

where ε_x is the error in a single measurement. These errors may be classified as either systematic or random in nature. Systematic errors are constant from one measurement to the next, e.g. using a rule with an inaccurate zero. We attempt to eliminate systematic errors from our experiments by careful design of the experiment and by using methods of measurement which are designed to eliminate the possibility of systematic error. On the other hand, random errors may be less than or greater than 0, and do not have a constant offset.

Consider a set of independent measurements of our quantity x . We obtain measurements $x_1, x_2, x_3, \dots, x_n$. Each measurement may be considered to be the true value x with some unknown additional error

$$x_i = X + \varepsilon_{x_i}.$$

We can use the standard methods of statistics to calculate the mean $\overline{\varepsilon_x}$ and standard deviation σ_{ε_x} of ε_x . Since we have eliminated any systematic error which would lead to $\overline{\varepsilon_x} \neq 0$, we shall assume

that $\overline{\varepsilon_x} = 0$. We shall assume that a measure of σ_{ε_x} is available. Furthermore, we shall assume that ε_x is a random variable, and that it is normally distributed $N(0, \sigma^2)$. If we have only a single measurement of x , then we must use the precision of our measuring equipment as an estimate of σ_{ε_x} . If we take a set of measurements and average them to find our estimate, \hat{x} , of the true value, X , then we should use the square of the standard error in the mean, σ_m , as our estimate of σ_{ε_x} ,

$$\sigma_{\varepsilon_x}^2 = \sigma_m^2 = \sigma_{mm} = \frac{\sigma^2}{n}, \quad (\text{I.2})$$

in as much as this does not exceed the accuracy of the measurement procedure [124, Section 3.4]. The notation of a double subscript shall be used henceforth to denote a variance. Where the subscript consists of two different variables, this denotes a covariance between the two random variables.

The final result for a measurement is normally quoted as

$$x \pm \Delta x, \quad (\text{I.3})$$

where the error bound, Δx , is calculated from

$$\Delta x = c_{\text{int}} \cdot \sigma_x = c_{\text{int}} \cdot \sigma_{\varepsilon_x}, \quad (\text{I.4})$$

where c_{int} is the statistical confidence interval. Using a value of $c_{\text{int}} = 1$ gives a 66% probability of the errors being within the error bounds, $c_{\text{int}} = 2$ gives 96%, and so on according to the normal distribution model $n(0, 1)$. If the measurement is a single observation, half the quantisation of the measuring instrument is used.

I.2 Error Analysis of Multi-Parameter Functions

Now we analyse the error in the output of a multi-parameter function, g , due to the errors present in each of the parameters, μ_i , which are determined from experimental data. The function is

$$y = g(\mu_1, \mu_2, \mu_3, \dots). \quad (\text{I.5})$$

Since the μ_i are derived from measurements, each will be composed of the true value M_i and some measurement error:

$$\mu_i = M_i + \varepsilon_{\mu_i}.$$

The measurement error is a 0-mean random variable, as discussed previously. The standard deviation, $\sigma_{\varepsilon_{\mu_i}}$, of the error, ε_{μ_i} , is the same as σ_{μ_i} , and may be found from the quoted error bound, $\Delta\mu_i$, and confidence interval, c_{int} , for the measurement, as described in Section I.1. For some particular measurement using (I.5), we have:

$$y + \varepsilon_y = g(\mu_1 + \varepsilon_{\mu_1}, \mu_2 + \varepsilon_{\mu_2}, \mu_3 + \varepsilon_{\mu_3}, \dots).$$

As an intermediate step to finding ε_y , we use a Taylor series expansion about the point $(\mu_1, \mu_2, \mu_3, \dots)$:

$$\begin{aligned} y + \varepsilon_y &= g(\mu_1, \mu_2, \mu_3, \dots) \\ &+ \frac{\partial g}{\partial \mu_1} \varepsilon_{\mu_1} + \frac{\partial g}{\partial \mu_2} \varepsilon_{\mu_2} + \frac{\partial g}{\partial \mu_3} \varepsilon_{\mu_3} + \dots \\ &+ \frac{\partial^2 g}{\partial \mu_1^2} (\varepsilon_{\mu_1})^2 + \frac{\partial^2 g}{\partial \mu_2^2} (\varepsilon_{\mu_2})^2 + \frac{\partial^2 g}{\partial \mu_3^2} (\varepsilon_{\mu_3})^2 + \dots \\ &+ \dots \end{aligned} \quad (I.6)$$

For mathematical convenience, we discard the terms involving $\frac{\partial^2 g}{\partial \mu_i^2} (\varepsilon_{\mu_i})^2$ and higher orders. This is justified by assuming that $(\varepsilon_{\mu_i})^2 \ll \varepsilon_{\mu_i}$, such that the contributions of the higher order terms are negligible. This is usually the case. Subtracting (I.5) from (I.6), we arrive at the approximation

$$\varepsilon_y \approx \frac{\partial g}{\partial \mu_1} \varepsilon_{\mu_1} + \frac{\partial g}{\partial \mu_2} \varepsilon_{\mu_2} + \frac{\partial g}{\partial \mu_3} \varepsilon_{\mu_3} + \dots \quad (I.7)$$

We now consider a set of measurements, each with its own value for $\{\varepsilon_{\mu_1}, \varepsilon_{\mu_2}, \varepsilon_{\mu_3}, \dots, \varepsilon_y\}$. Each of these variables is a random variable, of the form $n(0, \sigma^2)$, having a mean and a standard deviation. As discussed before, we expect that $\overline{\varepsilon_{\mu_i}} = 0$. This implies, through (I.7), that $\overline{\varepsilon_y} = 0$ as well.

The variance, σ_{yy} , is defined to be (Mendenhall [91])

$$\sigma_y^2 = \sigma_{yy} = E \left\{ (\varepsilon_y - \overline{\varepsilon_y})^2 \right\} \quad (I.8)$$

where $E \{ \}$ is the expected value operator. As $\overline{\varepsilon_y} = 0$, (I.8) reduces to

$$\sigma_{yy} = E \left\{ \varepsilon_y^2 \right\}. \quad (I.9)$$

Substituting (I.7) into (I.9) gives

$$\begin{aligned} E \left\{ \varepsilon_y^2 \right\} &= E \left\{ \left(\frac{\partial g}{\partial \mu_1} \varepsilon_{\mu_1} + \frac{\partial g}{\partial \mu_2} \varepsilon_{\mu_2} + \frac{\partial g}{\partial \mu_3} \varepsilon_{\mu_3} + \dots \right)^2 \right\} \\ &= \frac{\partial g}{\partial \mu_1} \frac{\partial g}{\partial \mu_1} E \left\{ \varepsilon_{\mu_1} \varepsilon_{\mu_1} \right\} + \frac{\partial g}{\partial \mu_1} \frac{\partial g}{\partial \mu_2} E \left\{ \varepsilon_{\mu_1} \varepsilon_{\mu_2} \right\} + \frac{\partial g}{\partial \mu_1} \frac{\partial g}{\partial \mu_3} E \left\{ \varepsilon_{\mu_1} \varepsilon_{\mu_3} \right\} + \dots \\ &\quad + \frac{\partial g}{\partial \mu_2} \frac{\partial g}{\partial \mu_1} E \left\{ \varepsilon_{\mu_2} \varepsilon_{\mu_1} \right\} + \frac{\partial g}{\partial \mu_2} \frac{\partial g}{\partial \mu_2} E \left\{ \varepsilon_{\mu_2} \varepsilon_{\mu_2} \right\} + \frac{\partial g}{\partial \mu_2} \frac{\partial g}{\partial \mu_3} E \left\{ \varepsilon_{\mu_2} \varepsilon_{\mu_3} \right\} + \dots \\ &\quad + \frac{\partial g}{\partial \mu_3} \frac{\partial g}{\partial \mu_1} E \left\{ \varepsilon_{\mu_3} \varepsilon_{\mu_1} \right\} + \frac{\partial g}{\partial \mu_3} \frac{\partial g}{\partial \mu_2} E \left\{ \varepsilon_{\mu_3} \varepsilon_{\mu_2} \right\} + \frac{\partial g}{\partial \mu_3} \frac{\partial g}{\partial \mu_3} E \left\{ \varepsilon_{\mu_3} \varepsilon_{\mu_3} \right\} + \dots \\ &\quad + \dots \end{aligned}$$

Substituting $\sigma_{\mu_i \mu_j} = E \left\{ \varepsilon_{\mu_i} \varepsilon_{\mu_j} \right\}$ and re-writing in matrix form, we have

$$\sigma_{yy} = \begin{bmatrix} \frac{\partial g}{\partial \mu_1} & \frac{\partial g}{\partial \mu_2} & \frac{\partial g}{\partial \mu_3} & \dots \end{bmatrix} \cdot \begin{bmatrix} \sigma_{\mu_1 \mu_1} & \sigma_{\mu_1 \mu_2} & \sigma_{\mu_1 \mu_3} & \dots \\ \sigma_{\mu_2 \mu_1} & \sigma_{\mu_2 \mu_2} & \sigma_{\mu_2 \mu_3} & \dots \\ \sigma_{\mu_3 \mu_1} & \sigma_{\mu_3 \mu_2} & \sigma_{\mu_3 \mu_3} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \cdot \begin{bmatrix} \frac{\partial g}{\partial \mu_1} \\ \frac{\partial g}{\partial \mu_2} \\ \frac{\partial g}{\partial \mu_3} \\ \vdots \end{bmatrix} \quad (I.10)$$

This may be re-written as

$$\sigma_{yy} = \mathbf{G} \Sigma_{\mu} \mathbf{G}^T, \quad (I.11)$$

where the matrices are listed below:

Σ_μ is the **Variance CoVariance (VCV)** matrix of the function parameters, μ_i . The subscript is used to reflect the variable or variables the **VCV** refers to.

G is the design matrix corresponding to the function g of equation (I.5).

The contents of these matrices may be identified from (I.10). The terms, $\sigma_{\mu_i\mu_i}$, along the diagonal of Σ are variances, while the off-diagonal terms, $\sigma_{\mu_i\mu_j}$, $i \neq j$, are the covariances. In the case where the covariances are 0, i.e. random variables μ_i and μ_j are statistically independent, (I.10) reduces to

$$\sigma_{yy} = \left(\frac{\partial g}{\partial \mu_1} \right)^2 \sigma_{\mu_1}^2 + \left(\frac{\partial g}{\partial \mu_2} \right)^2 \sigma_{\mu_2}^2 + \left(\frac{\partial g}{\partial \mu_3} \right)^2 \sigma_{\mu_3}^2 + \dots, \quad (\text{I.12})$$

which appears in Squires [124, Section 4.1].

The stated error, Δy , may now be determined using the standard deviation, σ_y , and the confidence interval, c_{int} , as in (I.4).

I.3 Curve Fitting

The problem of curve fitting assumes that there are some parameters $\xi_0, \xi_1, \xi_2, \dots, \xi_n$ which we may control, and one dependent variable y . The problem is to find linear parameters, β_i , which satisfy the relationship

$$y = \beta_0 \xi_0 + \beta_1 \xi_1 + \beta_2 \xi_2 + \dots + \beta_n \xi_n, \quad (\text{I.13})$$

assuming that such a linear relationship exists. Note that a non-linear function, such as $y = x^2$, may be linearised in the parameters, ξ_i , using a polynomial basis,

$$\xi_0 = 1, \xi_1 = x, \xi_2 = x^2, \dots, \xi_n = x^n, \quad (\text{I.14})$$

up to whatever order n is necessary, giving

$$y = \beta_0 + \beta_1 x + \beta_2 x^2 + \dots + \beta_n x^n. \quad (\text{I.15})$$

Thus a polynomial fit is formed. It is further possible to have multiple input variables by adding extra polynomial terms in a different parameter. It is also possible to linearise the fit by substituting functions such as $\xi_0 = \log_{10} x$.

If we have a set of m measurements of parameters, x_i , which are assumed to have no significant measurement error, and measurements y_i and associated σ_{y_i} , we can write (I.15) in matrix form as

$$\mathbf{Y} = \mathbf{A}\boldsymbol{\beta} + \mathbf{E}. \quad (\text{I.16})$$

The measurements, (x_i, y_i) , are represented by the vectors and matrix

$$\mathbf{Y} = \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}, \mathbf{A} = \begin{bmatrix} 1 & x_1 & x_1^2 & \dots & x_1^n \\ 1 & x_2 & x_2^2 & \dots & x_2^n \\ 1 & x_3 & x_3^2 & \dots & x_3^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_m & x_m^2 & \dots & x_m^n \end{bmatrix}, \boldsymbol{\beta} = \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \vdots \\ \beta_n \end{bmatrix}, \mathbf{E} = \begin{bmatrix} e_1 \\ e_2 \\ e_3 \\ \vdots \\ e_m \end{bmatrix},$$

The matrices are listed below.




Figure I.1: A first order $n = 1$ polynomial fit showing the measured data points, the fitted curve, and the errors of each point from the fitted curve.

Y contains the y_i values of the measurements.

A is the design matrix. The rows are derived from the data points, x_i , by equation (I.14).

β is a vector of the fit parameters.

E represents the residual errors between each measured point and the fitted curve.

We can also see that the rows of the design matrix, **A** , are the derivative of (I.15) with respect to the fit parameters, β_i , i.e.

$$\left[\frac{\partial g}{\partial \beta_1} \quad \frac{\partial g}{\partial \beta_2} \quad \frac{\partial g}{\partial \beta_3} \quad \cdots \quad \frac{\partial g}{\partial \beta_n} \right]$$

This gives one row of **A** . When this is applied to a set of measurements, we obtain the full matrix **A** .

The fitted curve parameters, **β** , can be represented graphically for a first order polynomial fit (straight line) in Figure I.1.

I.3.1 Weighted Measurements

If we know the relative errors of each measurement in the y direction (not the fit errors), e.g.

$$\Delta Y = \begin{bmatrix} \Delta y_1 & \Delta y_2 & \Delta y_3 & \cdots & \Delta y_m \end{bmatrix}^T,$$

then we may take that information into account by using the following procedure. Use the stated error, Δy_i , for each measurement, along with the confidence interval, c_{int} , to form the measurement **variance covariance** matrix using (I.4) to calculate $\sigma_{y_i y_i}$:

$$\Sigma_y = \begin{bmatrix} \sigma_{y_1 y_1} & 0 & 0 & \cdots & 0 \\ 0 & \sigma_{y_2 y_2} & 0 & \cdots & 0 \\ 0 & 0 & \sigma_{y_3 y_3} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma_{y_m y_m} \end{bmatrix},$$

where the covariances between separate measurements are assumed to be 0, i.e. the measurements are independent. Then the observation \mathbf{VCV} , $\mathbf{\Sigma}_y$, is inverted to form a weights matrix, \mathbf{P} :

$$\mathbf{P} = \mathbf{\Sigma}_y^{-1}.$$

When the \mathbf{VCV} matrix is diagonal, it may be inverted simply by inverting its elements (Anton [6]). The weight matrix is applied to the fit equation (I.16) as

$$\mathbf{PY} = \mathbf{PA}\boldsymbol{\beta} + \mathbf{E}. \quad (\text{I.17})$$

When no estimate of the measurement uncertainty is provided, we may use (I.16) instead of (I.17).

I.3.2 Curve Fitting Solution

The process of curve-fitting is to find the parameters, $\boldsymbol{\beta}$, which minimise the sum of squared errors, $SSE = \mathbf{E}^T \cdot \mathbf{E}$, which when expanded gives

$$SSE = \begin{bmatrix} e_1 & e_2 & e_3 & \dots \end{bmatrix} \cdot \begin{bmatrix} e_1 \\ e_2 \\ e_3 \\ \vdots \end{bmatrix}.$$

The fit parameters vector, $\boldsymbol{\beta}$, which minimises the sum of squared errors, SSE , may be estimated in the least squares sense using the Moore-Penrose pseudo-inverse (Weissstein [137]) to solve (I.17). The estimate, $\hat{\boldsymbol{\beta}}$, of the fit parameters, $\boldsymbol{\beta}$, is

$$\hat{\boldsymbol{\beta}} = (\mathbf{A}^T \mathbf{PA})^{-1} \mathbf{A}^T \mathbf{PY}. \quad (\text{I.18})$$

It is necessary to have $m \geq n$ to obtain a solution. Setting $\mathbf{P} = \mathbf{I}$, where \mathbf{I} is the identity matrix, in (I.18) obtains the solution to (I.16), which is used when weighting is not desired.

Now we must measure our confidence in the fit parameters $\hat{\boldsymbol{\beta}}$. This is done by first estimating the residual errors $\hat{\mathbf{E}}$, obtained by solving (I.16) for \mathbf{E} :

$$\hat{\mathbf{E}} = \mathbf{Y} - \mathbf{A}\hat{\boldsymbol{\beta}}.$$

This allows us to calculate an estimate of the standard deviation of the measurements in the y direction, by weighting the errors


$$\hat{\sigma}^2 = \frac{\hat{\mathbf{E}}^T \mathbf{P} \hat{\mathbf{E}}}{m - n}.$$

The estimated fit parameters variance covariance matrix, $\hat{\mathbf{\Sigma}}_{\boldsymbol{\beta}}$, is calculated from

$$\hat{\mathbf{\Sigma}}_{\boldsymbol{\beta}} = \frac{\hat{\mathbf{E}}^T \mathbf{P} \hat{\mathbf{E}}}{m - n} (\mathbf{A}^T \mathbf{PA})^{-1}$$

and looks like

$$\hat{\mathbf{\Sigma}}_{\boldsymbol{\beta}} = \begin{bmatrix} \hat{\sigma}_{\beta_1\beta_1} & \hat{\sigma}_{\beta_1\beta_2} & \hat{\sigma}_{\beta_1\beta_3} & \dots & \hat{\sigma}_{\beta_1\beta_n} \\ \hat{\sigma}_{\beta_2\beta_1} & \hat{\sigma}_{\beta_2\beta_2} & \hat{\sigma}_{\beta_2\beta_3} & \dots & \hat{\sigma}_{\beta_2\beta_n} \\ \hat{\sigma}_{\beta_3\beta_1} & \hat{\sigma}_{\beta_3\beta_2} & \hat{\sigma}_{\beta_3\beta_3} & \dots & \hat{\sigma}_{\beta_3\beta_n} \\ \dots & \dots & \dots & \ddots & \vdots \\ \hat{\sigma}_{\beta_n\beta_1} & \hat{\sigma}_{\beta_n\beta_2} & \hat{\sigma}_{\beta_n\beta_3} & \dots & \hat{\sigma}_{\beta_n\beta_n} \end{bmatrix}.$$

Figure I.2: Ellipsoid of Standard Deviation for Σ_{β} .

Since the relationship between covariances is symmetrical, $\sigma_{\beta_i\beta_j} = \sigma_{\beta_j\beta_i}$. This in turn leads to the fit parameter **VCV** matrix $\widehat{\Sigma}_{\beta}$ being symmetrical.

I.4 Meaning of the Variance Covariance Matrix

An interpretation of the fit parameter **Variance CoVariance (VCV)** matrix $\widehat{\Sigma}_{\beta}$ may be found as follows (after Hamilton [53, Section 4.2]). The quadratic equation below is parameterised by the elements of β , and evaluates to a scalar:

$$S_0 = (\beta - \widehat{\beta})^T \widehat{\Sigma}_{\beta}^{-1} (\beta - \widehat{\beta}).$$

For convenience, we write $\Delta = (\beta - \widehat{\beta})$, where $\Delta = \begin{bmatrix} \Delta\beta_0 & \Delta\beta_1 & \Delta\beta_3 & \dots \end{bmatrix}^T$, thus

$$S_0 = \Delta^T \widehat{\Sigma}_{\beta}^{-1} \Delta$$

The locus of points Δ which satisfies this equation forms a hyper-ellipsoid centred at $\Delta = \mathbf{0}$. When using a straight line fit, $\Delta = \begin{bmatrix} \Delta\beta_0 & \Delta\beta_1 \end{bmatrix}^T$. A graph of the locus (obtained using a contour plot at $S_0 = 1$) for the **VCV** matrix $\widehat{\Sigma}_{\beta}$ corresponding to the data set in Figure I.1 is shown in Figure I.2. This particular locus, where $S_0 = 1$, is called the ellipsoid of standard deviation.

The estimated **VCV** for the fit to this particular data set was

$$\widehat{\Sigma}_{\beta} = \begin{bmatrix} 3.30941 & -2.42152 \\ -2.42152 & 2.30621 \end{bmatrix},$$

which gives

$$\sigma_{\beta_0} = \sqrt{\sigma_{\beta_0\beta_0}} = 1.82, \sigma_{\beta_1} = \sqrt{\sigma_{\beta_1\beta_1}} = 1.52.$$




Figure I.3: A family of ellipsoids of standard deviation demonstrating the effect of changing the correlation coefficient $c_{\beta_0\beta_1} = \sigma_{\beta_0\beta_1} / (\sigma_{\beta_0} \sigma_{\beta_1})$ (and hence covariance $\sigma_{\beta_0\beta_1}$) through the range $-1 \dots 1$. The standard deviations σ_{β_0} and σ_{β_1} are kept constant.

Figure I.2 shows that the ellipsoid of standard deviation is bounded by a rectangle whose side lengths are equal to two standard deviations along each side, according to the standard deviation of each axis. The effect of the covariance $\sigma_{\beta_0\beta_1}$ is shown in Figure I.3. When $\sigma_{\beta_0\beta_1} < 0$, the ellipsoid slopes to the left. This indicates that a positive error in β_0 is very likely to be associated with a negative error in β_1 . When $\sigma_{\beta_0\beta_1} = 0$, there is no correlation between the errors in β_0 and β_1 . When $\sigma_{\beta_0\beta_1} > 0$, the ellipse slopes to the right. This indicates that a positive error in β_0 is very likely to be associated with a positive error in β_1 .

It is convenient to define the correlation coefficient, $c_{\beta_0\beta_1}$, to be

$$c_{\beta_0\beta_1} = \frac{\sigma_{\beta_0\beta_1}}{\sigma_{\beta_0} \sigma_{\beta_1}},$$

as this quantity describes the skewness of the ellipsoid of standard deviation in a normalised manner. Legitimate values of the correlation coefficient, $c_{\beta_0\beta_1}$, fall in the range $-1 \dots 1$. A correlation coefficient of -1 or 1 will produce a singular **VCV** matrix.

I.5 Forward Error Calculation

In this section we consider how to use the fitted curve to calculate $y \pm \Delta y$ given $x \pm \Delta x$, the estimated curve fit parameters, $\hat{\beta}$, and the curve fit **VCV** matrix, $\hat{\Sigma}_{\beta}$. This procedure will also take into account the error associated with the curve fit. We shall only consider the case of a single parameter straight line fit (polynomial of order one). Higher orders can be similarly derived. The value for y is calculated from

$$y = \beta_0 + \beta_1 x. \quad (\text{I.19})$$




Figure I.4: Error bounds on the fitted line for $\Delta x = 0$, incorporating covariances between β_0 and β_1 .

The stated error in x , Δx , must be converted to a variance σ_{xx} using (I.4). The variance in y , σ_{yy} , is calculated according to (I.11), where

$$\mathbf{G} = \begin{bmatrix} \frac{\partial y}{\partial x} & \frac{\partial y}{\partial \beta_0} & \frac{\partial y}{\partial \beta_1} \end{bmatrix} = \begin{bmatrix} \beta_1 & 1 & x \end{bmatrix}$$

and

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_{xx} & 0 & 0 \\ 0 & \sigma_{\beta_0\beta_0} & \sigma_{\beta_0\beta_1} \\ 0 & \sigma_{\beta_1\beta_0} & \sigma_{\beta_1\beta_1} \end{bmatrix}$$

The variance in y , being σ_{yy} , may be converted to the stated error, Δy , using (I.4). Thus we have calculated $y \pm \Delta y$ incorporating errors from all sources.

Figure I.4 shows the error in the fitted curve from our example data shown in Figure I.1. The narrowest point of the error band is at \bar{x} , and σ_{β_0} is the width of the error band at $x = 0$. The effect of the covariances $\sigma_{\beta_0\beta_1}$ is to shift the narrowest point of the error band away from the origin.

Knuth [83] describes some efficient methods for computing polynomials, which are useful for evaluating (I.19), especially for large polynomial orders.

I.6 Reverse Error Calculation

A reverse error calculation allows us to calculate $x \pm \Delta x$ given $y \pm \Delta y$, the fit parameters, $\hat{\boldsymbol{\beta}}$, and the fit parameter VCV matrix, $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\beta}}$. For a single parameter straight line fit, the principal value, x , can be found in terms of y by solving (I.19):

$$x = \frac{y - \beta_0}{\beta_1}.$$

The stated error, Δx , must be converted to the variance σ_{xx} using (I.4). The information needed to calculate (I.11) is

$$\mathbf{G} = \begin{bmatrix} \frac{\partial x}{\partial y} & \frac{\partial x}{\partial \beta_0} & \frac{\partial x}{\partial \beta_1} \end{bmatrix} = \begin{bmatrix} \frac{1}{\beta_1} & \frac{-1}{\beta_1} & \frac{\beta_0 - y}{\beta_1^2} \end{bmatrix}$$




Figure I.5: The error bounds $x \pm \Delta x$ for a reverse estimation.

and

$$\Sigma = \begin{bmatrix} \sigma_{yy} & 0 & 0 \\ 0 & \sigma_{\beta_0\beta_0} & \sigma_{\beta_0\beta_1} \\ 0 & \sigma_{\beta_1\beta_0} & \sigma_{\beta_1\beta_1} \end{bmatrix}.$$

Subsequent application of (I.11) yields σ_{xx} , and Δx may be found using (I.4). These error bounds are plotted in Figure I.5. Whereas Figure I.4 shows the error lines displaced vertically from the line of best fit, Figure I.5 shows the error lines displaced horizontally.

Application of the technique to a reverse error calculation of a quadratic is a little trickier, because there are two possible roots in this case. There is also the possibility that one or both of the roots do not exist. The forward formula is

$$y = \beta_0 + \beta_1 x + \beta_2 x^2. \quad (\text{I.20})$$

The well known quadratic formula for obtaining the roots of this equation is

$$x = -\frac{\beta_1 \pm \sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}}{2\beta_2}, \quad (\text{I.21})$$

which may have between 0 and 2 real solutions. The design matrix \mathbf{G} is

$$\mathbf{G} = \begin{bmatrix} \frac{\partial x}{\partial y} & \frac{\partial x}{\partial \beta_0} & \frac{\partial x}{\partial \beta_1} & \frac{\partial x}{\partial \beta_2} \end{bmatrix}.$$

The two solutions, where they exist, are given by

$$\mathbf{G}_1 = \begin{bmatrix} \frac{1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} \\ \frac{-1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} \\ \frac{\beta_1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} - 1 \\ \frac{2\beta_2}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} \\ \frac{y - \beta_0}{\beta_2 \sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} - \frac{-\beta_1 + \sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}}{2\beta_2^2} \end{bmatrix}^T$$

and

$$\mathbf{G}_2 = \begin{bmatrix} \frac{-1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} \\ \frac{1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} \\ \frac{-\beta_1}{\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} - 1 \\ \frac{2\beta_2}{\beta_2\sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}} + \frac{\beta_1 + \sqrt{\beta_1^2 + 4(y - \beta_0)\beta_2}}{2\beta_2^2} \end{bmatrix}^T.$$

We shall use a **VCV** matrix which incorporates the variance of \mathbf{x} :

$$\Sigma_\beta = \begin{bmatrix} \sigma_{xx} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_0\beta_0} & \sigma_{\beta_0\beta_1} & \sigma_{\beta_0\beta_2} \\ 0 & \sigma_{\beta_1\beta_0} & \sigma_{\beta_1\beta_1} & \sigma_{\beta_1\beta_2} \\ 0 & \sigma_{\beta_2\beta_0} & \sigma_{\beta_2\beta_1} & \sigma_{\beta_2\beta_2} \end{bmatrix},$$

which remains the same for both solutions to (I.21).

Extension of the reverse error calculation technique to polynomials of order greater than 2 is not trivial, due to the number of solutions available to choose from ($0 \dots n$, where n is the order of the polynomial). Sometimes there may be no real solutions to choose from - the non-real solutions will be complex. Explicit formulae for the roots of a polynomial of orders 3 and 4 are available (Abramowitz [1, Section 3.8]), but the formulae become quite complicated. Roots of higher order polynomials are usually determined numerically, so a numerical approximation to the derivatives must also be obtained in such a case. It must be pointed out that there exist matrix techniques for efficiently finding the roots of large polynomials, by finding the eigenvalues of the companion matrix (Golub [45, Section 7.4]). This technique will produce n roots for an order n polynomial, although some (or all) roots may be complex.

Only polynomials of order 1 and 2 have been used for reverse error estimation in this thesis, so the formulae mentioned previously suffice for this work. The companion matrix root finding technique has not been applied in the context of reverse error calculation, but appears to be a non-trivial extension due to the requirement of calculating the derivatives which are not provided by the numerical technique.

I.7 Significance of the Regression

Once a curve has been fitted to the data, it is important to check whether the curve matches the data well, or whether a different curve should have been used. This can be assessed by plotting the fit errors \mathbf{E} on a graph, and determining if there is a pattern in the errors. If the errors appear to follow a pattern, then the wrong function was used for the fit, and another one should be selected. When the errors appear similar to random noise, a satisfactory fit has been obtained. This technique is illustrated well in the Using Matlab guide [130, Chapter 6].

A more formal mathematical description of the above technique is described by Draper [35, Section 1.3].

I.8 Weighted Averaging

Weighted averaging is used to combine two statistical measurements of a quantity, $\{\mu_1, \sigma_{11}\}$ and $\{\mu_2, \sigma_{22}\}$, where μ_i represents the mean of the sample and σ_{ii} represents the variance. These two measurements can be combined by using a weighted average

$$\mu_c = \frac{\sum_i w_i \mu_i}{\sum_i w_i} = \frac{w_1 \mu_1 + w_2 \mu_2}{w_1 + w_2}, \quad (\text{I.22})$$

where w_1 and w_2 are the weights. The more precise measurements, i.e. those with smaller σ_{ii} , may be given more importance by using the weighting function

$$w_i = \frac{1}{\sigma_{ii}}. \quad (\text{I.23})$$

Substituting (I.23) back into (I.22) yields an expression for μ_c in terms of known quantities:

$$\mu_c = \frac{\sigma_{22} \mu_1 + \sigma_{11} \mu_2}{\sigma_{11} + \sigma_{22}}. \quad (\text{I.24})$$

This is the best combined value. Now the variance of the combined measurement is calculated from (I.11), where the design matrix is

$$\mathbf{G} = \begin{bmatrix} \frac{\partial \mu_c}{\partial \mu_1} & \frac{\partial \mu_c}{\partial \mu_2} \end{bmatrix} = \begin{bmatrix} \frac{\sigma_{22}}{\sigma_{11} + \sigma_{22}} & \frac{\sigma_{11}}{\sigma_{11} + \sigma_{22}} \end{bmatrix} \quad (\text{I.25})$$

and the variance covariance matrix is

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_{11} & \sigma_{21} \\ \sigma_{12} & \sigma_{22} \end{bmatrix}. \quad (\text{I.26})$$

Multiplying out the matrices and simplifying yields

$$\sigma_c^2 = \frac{\sigma_{11} \sigma_{22} (\sigma_{11} + 2\sigma_{12} + \sigma_{22})}{(\sigma_{11} + \sigma_{22})^2} \quad (\text{I.27})$$

for the general case where the measurements are correlated. When the measurements are uncorrelated ($\sigma_{12} = 0$) this simplifies to

$$\sigma_c^2 = \frac{\sigma_{11} \sigma_{22}}{\sigma_{11} + \sigma_{22}}. \quad (\text{I.28})$$

Thus we may use (I.24) to find the best combined mean of two normal distributions. The variance is found by (I.27) for correlated distributions, and by (I.28) for uncorrelated distributions. The combined variance will be larger for correlated measurements than for uncorrelated measurements.

I.9 Conclusion

The techniques of curve fitting and error analysis were widely utilised throughout this thesis to analyse and interpret the results, by fitting the expected models to the experimental data. The technique of weighted averages was used in Section 15.4 to combine the **IDD** and **IPD** bearing measurements.

Bibliography

- [1] ABRAMOWITZ, M., AND STEGUN, I. A., Eds. *Handbook of Mathematical Functions*. Dover, New York, 1972.
- [2] ACOUSTICAL SOCIETY OF AMERICA. American national standard for the calculation of the absorption of sound by the atmosphere. Tech. Rep. ANSI S1.26-1978, American Institute of Physics for the Acoustical Society of America, New York, 1978.
- [3] AKBARALLY, H. A vision supplemented sonar sensor for mobile robotics. Master's thesis, Department of Electrical and Computer Systems Engineering, Monash University, Clayton, VIC., Australia, Aug. 1995.
- [4] ANDERSON, E., ET AL. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, 1992. <http://www.netlib.org/lapack>.
- [5] ANKE, D. Luftshallwandler nach dem sell-prinzip fur frequenzen von 50 kHz bis 100 kHz. *Acustica* 30 (1974), 30–39.
- [6] ANTON, H. *Elementary Linear Algebra*, 5th ed. John Wiley and Sons, New York, 1987.
- [7] BAILEY, T., ET AL. Behaviour-based docking using the DAMN arbiter. In *Australian Conference on Robotics and Automation* (1999), Australian Robotics and Automation Association, pp. 38–43.
- [8] BARTLETT, M. S. Smoothing periodograms from time series with continuous spectra. *Nature* 161 (May 1948), 686–687.
- [9] BASS, H. E., ET AL. Atmospheric absorption of sound: Further developments. *Journal of the Acoustical Society of America* 97, 1 (1995), 680–683.
- [10] BENJAMIN, R. *Modulation, Resolution and Signal Processing in Radar, Sonar and Related Systems*. Pergamon Press, London, 1966.
- [11] BERKOWITZ, R. *Modern Radar: Analysis, Evaluation, and System Design*. John Wiley and Sons, New York, 1965.
- [12] BIBER, C., ET AL. The Polaroid ultrasonic ranging system. In *67th Convention of the Audio Engineering Society* (New York, 1980).
- [13] BLACKMAN, R. B., AND TUKEY, J. W. *The Measurement of Power Spectra*. Dover, New York, 1958.

- [14] BRACEWELL, R. N. *The Fourier Transform and its Applications*, 2nd ed. McGraw-Hill, New York, 1986.
- [15] BUDENSKE, J., AND GINI, M. Why is it so difficult for a robot to pass through a doorway using ultrasonic sensors? In *ICRA* (1994), pp. 3124–3129.
- [16] BUDERI, R. *The Invention that Changed the World*. Simon and Schuster, New York, 1996.
- [17] BUREAU INTERNATIONAL DES POIDS ET MESURES. http://www.bipm.fr/enus/3_SI/base_units.html, Feb. 2001.
- [18] BUREAU INTERNATIONAL DES POIDS ET MESURES. http://www.bipm.fr/enus/5_Scientific/e_rad_phot/temperature.html, Mar. 2002.
- [19] BURRUS, C. S., ET AL. *Computer-Based Exercises for Signal Processing using Matlab*. Prentice Hall, Englewood Cliffs, N.J., 1994. Electronic supplement available at <ftp://eedsp.gatech.edu/>.
- [20] CADZOW, J. A. Spectral estimation: An overdetermined rational model equation approach. *Proceedings of the IEEE* 70, 9 (1982), 907–939.
- [21] CHONG, K. S. *Simultaneous Mapping and Localisation for a Mobile Robot using Sonar Sensing*. PhD thesis, Department of Electrical and Computer Systems Engineering, Monash University, Clayton, VIC., Australia, Aug. 1997.
- [22] CHONG, K. S., AND KLEEMAN, L. Accurate odometry and error modelling for a mobile robot. In *IROS* (Grenoble, France, Sept. 1997), IEEE.
- [23] CLARK, S., AND DURRANT-WHYTE, H. The design of a high performance mmw radar system for autonomous land vehicle navigation. In *International Conference on Field and Service Robotics* (1997), A. Zelinsky, Ed., Australian Robot Association Inc, pp. 292–299.
- [24] COOLEY, J. W., AND TUKEY, J. W. An algorithm for the machine computation of complex Fourier series. *Journal of Mathematical Computation* 19 (Apr. 1965), 297–301.
- [25] COOLEY, J. W., AND TUKEY, J. W. Historical notes on the Fast Fourier Transform. *IEEE Transactions on Audio and Electroacoustics AU-15* (June 1967), 76–79.
- [26] CRAMER, O. The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity and CO₂ concentration. *Journal of the Acoustical Society of America* 93, 5 (May 1993), 2510–2516.
- [27] CROW, E. L., DAVIS, F. A., AND MAXFIELD, M. W. *Statistics Manual*. Dover Publications, New York, 1960.
- [28] DAVIES, J. K. *3D CTFM Airborne Ultrasonic Imaging*. PhD thesis, Faculty of Technology, The Open University, Milton Keynes, UK, Aug. 1994.
- [29] DE ROOS, A., ET AL. The detection and classification of objects lying on the seafloor. *Journal of the Acoustical Society of America* 84, 4 (Oct. 1988), 1456–1477.

- [30] DE ROOS, D. *Spectral Analysis Classification Sonars*. PhD thesis, University of Canterbury, Christchurch, New Zealand, 1986.
- [31] DELPHI AUTOMOTIVE SYSTEMS. Forewarn adaptive cruise control (radar). <http://www.delphiauto.com>, 2000.
- [32] DiFRANCO, J. V., AND RUBIN, W. L. *Radar Detection*. Prentice-Hall, Englewood Cliffs, N.J., 1968.
- [33] DONGARRA, J. J., ET AL. *LINPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, 1979. <http://www.netlib.org/linpack>.
- [34] DONGARRA, J. J., ET AL. An extended set of FORTRAN basic linear algebra subprograms. *ACM Transactions on Mathematical Software* 14 (1988), 1–17. <http://www.netlib.org/blas>.
- [35] DRAPER, N. R., AND SMITH, H. *Applied Regression Analysis*, 2nd ed. John Wiley and Sons, New York, 1981.
- [36] EISENHART, C. Realistic evaluation of the precision and accuracy of instrument calibration systems. *Journal of Research of the National Bureau of Standards* 67C, 2 (Apr. – June 1963), 161–187.
- [37] ELACHI, C. *An Introduction to radar remote sensing*. IEEE, New York, 1987.
- [38] ET AL, J. The global structure of the annual and semiannual sea surface height variability from geostat altimeter data. *Journal of Geophysical Research* 97 (1992), 17813–17828.
- [39] EVANS, L. B., AND BASS, H. E. Tables of absorption and velocity of sound in still air at 68°F(20°C). Tech. Rep. WR 72-2, Wyle Laboratories, Huntsville, AL, Jan. 1972. National Technical Information Service report number AD-738576.
- [40] FLYNN, A. M. Combining sonar and infrared sensors for mobile robot navigation. *International Journal of Robotics Research* 7, 6 (1988), 5–14.
- [41] FRIGO, M., AND JOHNSON, S. G. The fastest Fourier transform in the west. Tech. Rep. MIT-LCS-TR-728, MIT, Sept. 1997. <http://www.fftw.org>.
- [42] FRIGO, M., AND JOHNSON, S. G. FFTW: An adaptive software architecture for the FFT. In *ICASSP Conference Proceedings* (1998), vol. 3, pp. 1381–1384. <http://www.fftw.org>.
- [43] GABOR, D. Theory of communication. *Journal of the Institution of Electrical Engineers* 93, 3 (1946), 429–457.
- [44] GIANCOLI, D. C. *Physics: principles with applications*, 2nd ed. Prentice-Hall, Englewood Cliffs, N.J., 1985.
- [45] GOLUB, G. H., AND VAN LOAN, C. F. *Matrix Computations*. John Hopkins University Press, Baltimore, 1983.

- [46] GOUGH, P. T., AND DE ROOS, A. Continuous transmission FM sonar with one octave bandwidth and no blind time. *IEE Proceedings* 131, 3 (June 1984), 270–274. Part F.
- [47] GRIFFITHS, D. J. *Introduction to Electrodynamics*. Prentice Hall, Englewood Cliffs, New Jersey, 1989.
- [48] HACKMANN, W. *Seek and Strike: Sonar, anti-submarine warfare and the Royal Navy, 1914-54*. London: Her Majesty's Stationery Office, 1984.
- [49] HACKMANN, W. Asdics at war. *IEEE Review* (May 2000). http://www.iee.org/Oncomms/pn/history/library_temp.cfm.
- [50] HAHN, S. L. *Hilbert Transforms in Signal Processing*. Artech House, Inc., Norwood, MA, 1996.
- [51] HALLIDAY, D., AND RESNICK, R. *Fundamentals of Physics*, third ed. John Wiley and Sons, 1988.
- [52] HAMBLEY, A. R. *Electrical Engineering: principles and applications*. Prentice Hall, Upper Saddle River, N.J., 1997.
- [53] HAMILTON, W. C. *Statistics in Physical Science*. The Ronald Press Company, New York, 1964.
- [54] HARPER, N. Control software for a precision positioner. Honours thesis, Department of Computer Science, University of Wollongong, Wollongong, N.S.W., Australia, 1993.
- [55] HARPER, N. *Classification of Plants by the Interpretation of CTFM Sonar Data*. PhD thesis, School of Information Technology and Computer Science, University of Wollongong, Wollongong, N.S.W., Australia, 1999.
- [56] HARPER, N. L., AND M^CKERROW, P. J. Classifying plants with ultrasonic sensing. In *Australian Conference on Robotics and Automation* (1999), Australian Robotics and Automation Association, pp. 166–171.
- [57] HARRIS, F. J. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE* 66, 1 (1978), 51–84.
- [58] HAYES, M. P. *A CTFM Synthetic Aperture Sonar*. PhD thesis, University of Canterbury, Christchurch, New Zealand, Sept. 1989.
- [59] HEALE, A. High speed sonar with interference rejection and classification. Master's thesis, Department of Electrical and Computer Systems Engineering, Monash University, Clayton, VIC., Australia, 2001.
- [60] HEALE, A., AND KLEEMAN, L. A real time DSP sonar echo processor. In *IROS* (2000).
- [61] HECHT, E. *Optics*, 2nd ed. Addison-Wesley, Reading, Mass., 1987.

- [62] HONG, M. L., AND KLEEMAN, L. A low sample rate 3D sonar sensor for mobile robots. In *IEEE International Conference on Robotics and Automation* (1995), pp. 3015–3020.
- [63] HUNT, F. V. *Electroacoustics*, 2nd ed. American Institute of Physics, 1982.
- [64] HUNT, R. A. *Calculus with analytic geometry*. Harper and Row, New York, 1988.
- [65] ISO 2533-1975. Standard atmosphere. Tech. rep., International Organisation for Standardization, Geneva, Switzerland, 1975.
- [66] JOSUTTIS, N. M. *The C++ Standard Library: A Tutorial and Reference*. Addison-Wesley, Reading, Mass., 1999.
- [67] KAY, L. The technology. Introduction to the Technology of Ultrasonics.
- [68] KAY, L. A comparison between pulse and frequency modulation echo-ranging systems. *Journal of the British Institute of Radio Engineers* 19, 2 (Feb. 1959), 105–113.
- [69] KAY, L. An experimental comparison between a pulse and a frequency-modulation echo-ranging system. *Journal of the British Institute of Radio Engineers* 20, 10 (Oct. 1960), 785–796.
- [70] KAY, L. An ultrasonic sensing probe as a mobility aid for the blind. *Ultrasonics* (Apr. – June 1964), 53–59.
- [71] KAY, L. A sonar aid to enhance spatial perception of the blind: engineering design and evaluation. *The Radio and Electronic Engineer* 44, 11 (Nov. 1974), 605–627.
- [72] KAY, L. Airborne ultrasonic imaging of a robot workspace. In *Robot Sensors: Tactile and Non-Vision (International Trends in Manufacturing Technology)*, A. Pugh, Ed., vol. 2. IFS, 1986, pp. 287–295.
- [73] KAY, L. A CTFM acoustic spatial sensing technology: its use by blind persons and robots. *Sensor Review* 19, 3 (1999), 195–201.
- [74] KAY, L. Auditory perception of objects by blind persons, using bioacoustic high resolution air sonar. *Journal of the Acoustical Society of America* 107, 6 (June 2000), 3266–3275.
- [75] KAY, S. M., AND MARPLE, JR., S. L. Spectrum analysis - a modern perspective. *Proceedings of the IEEE* 69, 11 (1981), 1380–1419.
- [76] KAYE, G. W. C., AND LABY, T. H. *Tables of Physical and Chemical Constants*, 15th ed. Longman Group Limited, London, 1986.
- [77] KIMBALL, J. W. Kimball's biology pages. <http://biology-pages.info/C/CarbonCycle.html>, 2002.
- [78] KINSLER, L. E., ET AL. *Fundamentals of Acoustics*, 3rd ed. John Wiley and Sons, New York, 1982.

- [79] KLEEMAN, L. Optimal estimation of position and heading for mobile robots using ultrasonic beacons and dead-reckoning. In *ICRA* (1992), pp. 2582–2587.
- [80] KLEEMAN, L. Scanned monocular sonar and the doorway problem. In *IROS* (1996), pp. 96–103.
- [81] KLEEMAN, L. Ultrasonic sensors. In *The Industrial Electronics Handbook*, J. D. Irwin, Ed. CRC Press, IEEE Press, Boca Raton, Florida, Apr. 1997, pp. 738–745.
- [82] KLEEMAN, L., AND KUC, R. Mobile robot sonar for target localization and classification. *International Journal of Robotics Research* 14, 4 (Aug. 1995), 295–318.
- [83] KNUTH, D. E. *The Art of Computer Programming*, 3rd ed. Addison Wesley, Reading, Mass., 1998.
- [84] KUC, R. Biologically motivated adaptive sonar system. *Journal of the Acoustical Society of America* 100, 3 (Sept. 1996), 1849–1854.
- [85] KUC, R., AND SIEGEL, M. W. Physically based simulation model for acoustic robot navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-9*, 6 (Nov. 1987).
- [86] KURIE, F. N. D., Ed. *FM Sonar Systems*. Sonar Devices Division, University of California Division of War Research, United States Navy Laboratory, San Diego, 1946.
- [87] LANGER, D. An integrated MMW radar system for outdoor navigation. Tech. rep., Carnegie Mellon University, Pittsburgh, Penn., USA, Jan. 1997.
- [88] LITTLE, W., FOWLER, H. W., AND COULSON, J. *The Shorter Oxford English Dictionary*, 3rd ed. Oxford, London, 1959.
- [89] MAGUIRE, S. *Writing Solid Code*. Microsoft Press, Redmond, Washington, 1993.
- [90] MARTIN, G. *Electronics and Transducers for an Ultrasonic Blind Mobility Aid*. PhD thesis, University of Canterbury, Christchurch, New Zealand, 1969.
- [91] MENDENHALL, W., WACKERLY, D. D., AND SCHEAFFER, R. L. *Mathematical Statistics with Applications*, 4th ed. Thomson Information Publishing Group, Boston, Mass., 1990.
- [92] METROWERKS. Power plant. <http://www.metrowerks.com/>, 1995.
- [93] M^CKERROW, P., AND HARPER, N. Plant acoustic density profile model of CTFM ultrasonic sensing. *IEEE Sensors Journal* 1, 4 (Dec. 2001), 245–255.
- [94] M^CKERROW, P. J. Simulation of sonar echolocation. In *Proceedings First Workshop on Domestic Robots and Second Workshop on Medical and Healthcare Robotics* (Newcastle Upon Tyne, UK, Sept. 1989), International Advanced Robotics Programme, pp. 127–136.
- [95] M^CKERROW, P. J. *Introduction To Robotics*. Addison-Wesley, Sydney, Australia, 1991.

- [96] M^CKERROW, P. J. Progress in ultrasonic sensing for service robots. In *Robotics for the Service Industries* (Sydney, Australia, 1995), Australian Robot Association Inc., pp. 44–52.
- [97] M^CKERROW, P. J., AND ZHU, S. Modelling multiple reflection paths in ultrasonic sensing. In *IROS* (Osaka, Japan, Nov. 1996), pp. 284–291.
- [98] MURATA. <http://www.murata.com/>.
- [99] NAUMOVSKI, J., ET AL. A multiple autonomous mobile robot simulator. In *International Conference on Mechatronics and Machine Vision in Practice* (Queensland, Australia, Sept. 1994).
- [100] NIELSEN, R. O. *Sonar Signal Processing*. Artech House, Boston, Mass., 1991.
- [101] NORTH, D. O. Analysis of the factors which determine signal/noise discrimination in radar. Tech. rep., RCA, June 1943. PTR-6C, later republished as [102].
- [102] NORTH, D. O. Analysis of the factors which determine signal/noise discrimination in radar. *Proc. IEEE* 51, 7 (July 1963), 1015–1027.
- [103] OHNO, T., OHYA, A., AND YUTA, S. An improved sensory circuit of an ultrasonic range finder for mobile robot's obstacle detection. In *Robots for Australian Industries* (Melbourne, July 1995), IEEE, pp. 178–185.
- [104] OPPENHEIM, A. V., ET AL. *Signals and Systems*. Prentice Hall, Englewood Cliffs, N.J., 1983.
- [105] OPPENHEIM, A. V., AND SCHAFER, R. W. *Discrete-Time Signal Processing*, second ed. Prentice Hall, Upper Saddle River, N.J., 1998.
- [106] OSBORNE, C. F. Ray and diffraction effects of acoustic waves when atmospheric conditions produce nonlinear acoustic velocity profiles. *Journal of the Acoustical Society of America* 91, 1 (Jan. 1992), 86–90.
- [107] PEREMANS, H. A maximum likelihood algorithm for solving the correspondence problem in tri-aural perception. In *IEEE MFI* (Las Vegas, NV, USA, Oct. 1994), pp. 485–492.
- [108] PEREMANS, H. *Tri-aural perception for mobile robots*. PhD thesis, Department of Electronics and Information Systems, University of Ghent, Ghent, Belgium, 1994.
- [109] PEREMANS, H., AUDENAERT, K., AND CAMPENHOUT, J. A high-resolution sensor based on tri-aural perception. *IEEE Transactions on Robotics and Automation* 9, 1 (1993), 36–48.
- [110] PEREMANS, H., AND VEELAERT, P. Improved tri-aural perception through robot motion. In *Robotic and Manufacturing Systems* (Montpellier, France, 1996), vol. 3, TSI Press, pp. 623–630.
- [111] POLITIS, Z., AND PROBERT, P. J. Modeling and classification of rough surfaces using CTFM sonar imaging. In *International Conference on Robotics and Automation* (Detroit, Michigan, USA, 1999), pp. 2988–2993.

- [112] POLITIS, Z., AND PROBERT, P. J. Target localization and identification using CTFM sonar imaging: the AURBIT method. In *Computational Intelligence in Robotics and Automation* (1999), IEEE, pp. 256–261.
- [113] POOLE, H. H. *Fundamentals of Robotics Engineering*. Van Nostrand Reinhold, New York, 1989.
- [114] PRESS, W. H., ET AL. *Numerical Recipes in C*, 2nd ed. Cambridge University Press, Cambridge, 1992.
- [115] PRESTON-THOMAS, H. The international temperature scale of 1990 (ITS-90). *Metrologica* 27 (1990), 3–10. <http://www.its-90.com>.
- [116] PROAKIS, J. G., AND MANOLAKIS, D. G. *Digital Signal Processing*, 3rd ed. Prentice Hall, Englewood Cliffs, N.J., 1996.
- [117] RATNER, D. *Landmark Navigation of the Titan 4WD Outdoor Mobile Robot Using Narrow Beam CTFM Sonar*. PhD thesis, School of Information Technology and Computer Science, University of Wollongong, Wollongong, N.S.W., Australia, Jan. 2002.
- [118] RIHACZEC, A. W. *Principles of High-Resolution radar*, 2nd ed. Artech House, Inc., Norwood, MA, 1996.
- [119] ROGOWITZ, B. E., AND TREINISH, L. A. Data visualization: the end of the rainbow. *IEEE Spectrum* (Dec. 1998).
- [120] ROWELL, D. *Auditory Display of Spatial Information*. PhD thesis, Electrical Engineering, University of Canterbury, Christchurch, New Zealand, 1970.
- [121] SICK AG. *LMS 210/220/290 Laser Measurement System (Outdoor Version): Technical Description*. Reute, Germany, 1998.
- [122] SKOLNIK, M. I. *Radar Handbook*. McGraw-Hill, Inc, New York, 1970.
- [123] SOSTMANN, H. Fundamentals of Thermometry. <http://www.its-90.com/onref.html>, 1999.
- [124] SQUIRES, G. L. *Practical Physics*, third ed. Cambridge University Press, Cambridge, 1985.
- [125] STORER, J. E., AND TURYN, R. Optimum finite code groups. *Correspondence, Proceedings of the IRE* 46 (1958), 1649.
- [126] STRAHAN, R., Ed. *Complete Book of Australian Mammals*. Angus & Robertson Publishers, Melbourne, Australia, 1983.
- [127] STROUSTRUP, B. *The C++ Programming Language*, 3rd ed. Addison-Wesley Longman, Murray Hill, N.J., 1997.

- [128] SUGA, N. Biosonar and neural computation in bats. *Scientific American* (June 1990), 34–41.
- [129] SWORDS, S. S. *Technical History of the Beginnings of RADAR*. Peter Peregrinus Ltd, London, 1986.
- [130] THE MATHWORKS. Matlab. <http://www.mathworks.com>, 1999.
- [131] THERRIEN, C. W. *Discrete Random Signals and Statistical Signal Processing*. Prentice Hall, Englewood Cliffs, N.J., 1992.
- [132] TREES, H. L. V. *Detection, estimation and modulation theory*, vol. 1. John Wiley and Sons, New York, 1968.
- [133] VAN VALKENBURG. *Analog Filter Design*. Oxford University Press, New York, 1982.
- [134] VON KROGE, H. *GEMA: Birthplace of German Radar and Sonar*. Institute of Physics Publishing, Bristol, England, 2000. English translation by Louis Brown.
- [135] WARD, K., ZELINSKY, A., AND M^CKERROW, P. Learning to avoid objects and dock with a mobile robot. In *Australian Conference on Robotics and Automation* (1999), Australian Robotics and Automation Association, pp. 132–133.
- [136] WATERLOO MAPLE INC. Maple. <http://www.maplesoft.com/>, 2002.
- [137] WEISSTEIN, E. Eric Weisstein’s World of Mathematics. <http://mathworld.wolfram.com/>, 2002.
- [138] WELCH, P. D. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio Electro-Acoustics AU-15* (June 1967), 70–73.
- [139] WOLFRAM, S. *The Mathematica Book*, 4th ed. Wolfram Media, Champaign, Ill., 1999. <http://www.wolfram.com>.
- [140] WONG, G. S. K. Speed of sound in standard air. *Journal of the Acoustical Society of America* 79, 5 (May 1986), 1359–1366.
- [141] WONG, G. S. K. Comments on “The variation of the specific heat ratio and the speed of sound in air with temperature, pressure, humidity and CO₂ concentration” [J. Acoust. Soc. Am. 93, 2510–2516 (1993)]. *Journal of the Acoustical Society of America* 97, 5 (May 1995), 1359–1366.
- [142] YATA, T. *Direction Measurable Ultrasonic Sensing Systems for Mobile Robots*. PhD thesis, Intelligent Robot Laboratory, University of Tsukuba, Tsukuba, Japan, Mar. 2000.
- [143] YATA, T., KLEEMAN, L., AND YUTA, S. Fast-bearing measurement with a single ultrasonic transducer. *The International Journal of Robotics Research* 17, 11 (Nov. 1998), 1202–1213.

- [144] YATA, T., KLEEMAN, L., AND YUTA, S. Wall following using angle information measured by a single ultrasonic transducer. In *IEEE International Conference on Robotics and Automation* (Leuven, Belgium, May 1998), pp. 1590–1596.
- [145] YATA, T., OHYA, A., AND YUTA, S. A fast and accurate sonar ring sensor for a mobile robot. In *IEEE International Conference on Robotics and Automation* (Detroit, Michigan, USA, May 1999), pp. 630–636.
- [146] YULE, J.-D., Ed. *Methuen Concise Encyclopedia of Science and Technology*. Methuen of Australia, Sydney, 1978.
- [147] ZHU, S.-M. *Visualisation of Echolocation*. PhD thesis, Department of Computer Science, University of Wollongong, Wollongong, N.S.W., Australia, 1996.