



MacKerron, Graham Henry (2011) *The detection of unknown waveforms in ESM receivers: FFT-based real-time solutions*. EngD thesis.

<http://theses.gla.ac.uk/3007/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Glasgow Theses Service

<http://theses.gla.ac.uk/>

theses@gla.ac.uk

THE DETECTION OF UNKNOWN WAVEFORMS IN ESM RECEIVERS: FFT-BASED REAL-TIME SOLUTIONS

Graham Henry MacKerron

MPhys (Hons)

Submitted in fulfilment of the requirements for the Degree of
Doctor of Engineering in System Level Integration (EngD)

School of Engineering

College of Science and Engineering

University of Glasgow

October, 2011



Abstract

Radars and airborne electronic support measures (ESMs) systems are locked in a tactical battle to detect each other whilst remaining undetected. Traditionally, the ESM system has a range advantage. Low probability of intercept (LPI) waveform designers are, however, more heavily exploiting the matched filter radar advantage and hence degrading the range advantage.

There have been literature and internal, SELEX Galileo proposals to regain some ESM processing gain of low probability of intercept (LPI) waveforms. This study, however, has sought digital signal processing (DSP) solutions which are: (1) computationally simple; (2) backward-compatible with existing SELEX Galileo digital receivers (DRxs) and (3) have low resource requirements. The two contributions are complementary and result in a detector which is suitable for detection of most radar waveforms.

The first contribution is the application of spatially variant apodization (SVA) in a detection role. Compared to conventional window functions, SVA was found to be beneficial for the detection of sinusoidal radar waveforms as it surpassed the fixed window function detectors in all scenarios tested.

The second contribution shows by simulation that simple spectral smoothing techniques improved DRx LPI detection capability to a level similar to more complicated non-parametric spectral estimators and far in excess of the conventional (modified) periodogram.

The DSP algorithms were implemented using model-based design (MBD). The implication is that a detector with improved conventional and LPI waveform detection capability can be created from the intellectual property (IP). Estimates of the improvement in SELEX Galileo DRx system detection range are provided in the conclusion.

Contents

Contributions	xiii
Preface	xiv
Acknowledgements	xvi
Declaration	xvii
Definitions/Abbreviations	xviii
Abbreviations	xviii
List of Symbols	xxii
1 Introduction & Background	25
1.1 Introduction	25
1.2 The Range Advantage	25
1.3 History of ESM System Development	28
1.4 Windowing/Apodisation	31
1.5 Super-resolution Techniques	34
1.6 Envelope and Energy Detection	35
1.7 Ideal Detection	38
1.8 Radar Designer's Concerns	39
1.9 LPI Waveforms	41
1.10 LPI Waveform Detection Methods	44
1.11 Summary	51
2 SVA Theory	52
2.1 Introduction to SVA	52
2.2 Raised-cosine Windows	52
2.3 DFT of Raised-cosine Windows	54
2.4 Spatially Variant Apodization (SVA)	55
2.5 5-point and 7-point SVA	59
2.6 SVA using the Kaiser Window	64
2.7 Noninteger SVA Algorithm	65
2.8 Complex-parametered SVA	66
3 Floating-point SVA	74
3.1 Detectors Summary	75
3.2 Full DCW	75

3.3	Randomly-filled DCW	85
3.4	Interference in Full DCW	89
3.5	Interference in Randomly-filled DCW	107
3.6	LPI signal in DCW	115
4	Non-Parametric Spectral Estimation	122
4.1	Detection	122
4.2	Multitaper Theory	126
4.3	Smoothed Periodogram	142
4.4	Conclusions	149
5	Floating-Point MTM & Periodogram Smoothing	150
5.1	Background	150
5.2	Full DCW	172
5.3	Randomly-filled DCW	179
5.4	Interference in Full DCW	183
5.5	Interference in Randomly-filled DCW	201
5.6	Conclusions	209
5.7	Summary Table	210
6	Thesis Conclusions & Future Work	211
6.1	Conclusions	211
6.2	Future Work	214
	Appendices	215
A	Helicopter Trials	217
A.1	General Information	217
A.2	Constraints	217
A.3	Aims	218
A.4	Experiment	219
A.5	Results	220
A.6	Summary of Results	233
B	Complex-Parametered SVA	236
B.1	Unconstrained Complex SVA	236
B.2	Lagrangian Multiplier complex SVA	240
B.3	Constrained complex SVA	241
C	Model-Based Design	243
C.1	Introduction	243
C.2	MBD Principles	243
C.3	SELEX Galileo MBD Methodology	244
C.4	System Generator [®]	246

C.5	Xilinx ISE Suite	249
D	Fixed-Point SVA	253
D.1	Implementation Decisions	253
D.2	Coding Methodology	255
D.3	Coding Phase 1	257
D.4	Coding Phase 2	262
D.5	Coding Explanation	267
D.6	Results	275
D.7	Conclusions	281
E	SVA Optimisations	282
E.1	Optimisations	282
F	SVA-on-FPGA	288
F.1	Implementation Decisions	288
F.2	Coding Methodology	289
F.3	Second Type of Test	308
F.4	Reports	312
F.5	Suggested Design Improvements	313
G	Navajo Trials	317
G.1	General Information	317
G.2	Constraints	317
G.3	Aims	317
G.4	Experiment	318
G.5	Results	323
G.6	Summary of Results	329
H	Fixed-Point Periodogram Smoothing	330
H.1	Implementation Decisions	330
H.2	Coding Methodology	330
H.3	Coding Explanation	337
H.4	Results	339
H.5	Conclusions	342
I	TSP-on-FPGA	343
I.1	Implementation Decisions	343
I.2	Coding Methodology	343
I.3	Monte Carlo Simulation Tests	350
I.4	Suggested Design Improvements	352
J	Patent App. & Paper	353

Glossary	375
Bibliography	376

List of Tables

1.4.1	Window Function Metrics	33
3.1.1	Conventional, SVA Detectors	75
3.2.1	Typical ESM system DRx Parameters	77
3.2.2	Example Contingency Table	80
3.6.1	Typical FMCW and ESM system DRx Parameters	115
3.6.2	Typical Barker-13 Pulse and ESM system DRx Parameters	117
3.6.3	Results Summary	121
4.1.1	Non-parametric Spectral Estimators	125
4.2.1	MTM Test Vector Parameters	141
4.3.1	Algorithm Complexity	148
5.1.1	Conventional, MTM & Periodogram Smoothing Detectors	171
5.2.1	LPI Waveform Parameters	172
5.2.2	DRx Parameters	174
5.7.1	Results Summary	210
A.4.1	Watchman A Radars	219
E.1.1	α Dividend Calculation	284
E.1.2	SVA FIR Operation	286
F.2.1	Phase 1 Synthesis Report	296
F.2.2	Phase 2 Place & Route Report	297
F.4.1	SVA-on-FPGA Static Timing Report	313
I.2.1	TSP-on-FPGA Phase 1 Place & Route Report	347
I.2.2	Extract from TSP-on-FPGA Static Timing Report	347
I.2.3	TSP-on-FPGA Phase 2 Place & Route Report	350
I.2.4	Extract from Phase 2 TSP-on-FPGA Static Timing Report	350

List of Figures

1.2.1	ESM Range Advantage	26
1.2.2	Matched Filter	27
1.3.1	Typical ESM System	28
1.3.2	Clustering	29
1.4.1	Window Parameters	34
1.4.2	Windowing	35
1.8.1	Ambiguity Functions	41
1.9.1	Chirps and PN	42
1.10.1	Noise Distribution	45
1.10.2	Cyclostationary Analysis Example	49
1.10.3	Self-mixer	50
2.2.1	Raised-cosine Windows	53
2.2.2	Window Function Choice	54
2.4.1	Example Dual-apodization	56
2.4.2	Joint- and Separate-SVA Comparison	59
2.5.1	3-, 5- and 7-point SVAs	63
2.5.2	Noisy 3-, 5- and 7-point SVAs	64
2.8.1	Variation of the Shape and Delay Parameters	66
2.8.2	Cost Function	68
2.8.3	Unconstrained Complex SVA	69
2.8.4	Lagrangian Minimisation SVA	70
2.8.5	Constrained Complex SVA Example 1	71
2.8.6	""Example 2	72
2.8.7	""Failing 1	73
2.8.8	""Failing 2	73
3.2.1	Full DCW Concept	76
3.2.2	Definition of Random Frequency	77
3.2.3	Detector A1	78
3.2.4	Detector B1	79
3.2.5	Full DCW ROC curves P_{fa} vs P_d	81
3.2.6	Full DCW ROC curves SNR_{out} vs P_d	82
3.2.7	Full DCW Scenario with 5-point SVA and 7-point SVA	84
3.3.1	Pulse Random Arrival Time	86
3.3.2	Randomly-filled DCW ROC curves	87
3.4.1	Detector C1	90

3.4.2	Detector D1	91
3.4.3	Estimating Spectral Leakage	92
3.4.4	Quadratic Peak Interpolation	94
3.4.5	Interference DCW, $\text{SNR}_{\text{out},B}=8$ dB, sep=2	95
3.4.6	"", $\text{SNR}_{\text{out},B}=8$ dB, sep=5	96
3.4.7	"", $\text{SNR}_{\text{out},B}=15$ dB, sep=5	97
3.4.8	"", $\text{SNR}_{\text{out},B}=15$ dB, sep=2	98
3.4.9	"", $\text{SNR}_{\text{out},B}=25$ dB, sep=2	99
3.4.10	"", $\text{SNR}_{\text{out},B}=25$ dB, sep=5	100
3.4.11	Detector E1	101
3.4.12	SVA with/without Variable Threshold	102
3.4.13	Interference DCW, 5-point and 7-point SVA	104
3.4.14	Complex SVA ROC Curves	105
3.4.15	Detector F1	106
3.5.1	Incorrect Spectral Leakage Envelope	108
3.5.2	Partial Interference DCW, $\text{SNR}_{\text{out},B}=8$ dB, sep=2	109
3.5.3	""", $\text{SNR}_{\text{out},B}=8$ dB, sep=5	110
3.5.4	""", $\text{SNR}_{\text{out},B}=15$ dB, sep=2	111
3.5.5	""", $\text{SNR}_{\text{out},B}=15$ dB, sep=5	112
3.5.6	""", $\text{SNR}_{\text{out},B}=25$ dB, sep=2	113
3.5.7	""", $\text{SNR}_{\text{out},B}=25$ dB, sep=5	114
3.6.1	FMCW ROC Curves	116
3.6.2	Random Portion of LPI Pulse	118
3.6.3	Barker-13 ROC Curves	120
4.2.1	WOSA and MTM	127
4.2.2	MTM Filter Response	128
4.2.3	χ^2 PDF	131
4.2.4	DPSS Eigenvalues, 0–7	133
4.2.5	DPSS Tapers, 0–3	135
4.2.6	DPSS Tapers, 4–7	135
4.2.7	Sine Tapers	138
4.2.8	Example MTM Spectral Estimates	140
4.3.1	Filter Responses	145
4.3.2	Smoothed Periodogram Methods	146
4.3.3	Frequency Bandwidth Clustering	149
5.1.1	Hypothesis Tests	153
5.1.2	Detector G1	155
5.1.3	Detector H1	157
5.1.4	Detector I1	158
5.1.5	Detector J1	160

5.1.6	Detector K1	161
5.1.7	Detector L1	164
5.1.8	Detector M1	165
5.1.9	Detector N1	167
5.1.10	Detector O1	168
5.1.11	Detector P1	170
5.2.1	Monitored Bandwidth	173
5.2.2	Full DCW, B-13 Pulse	175
5.2.3	"", m-sequence CW	176
5.2.4	"", B-13 Pulse, Conventional Methods	177
5.2.5	"", m-sequence CW, Conventional Methods	178
5.3.1	Randomly-filled DCW	179
5.3.2	"-", B-13 Pulse	180
5.3.3	"-", B-13 Pulse, Conventional Methods	181
5.4.1	Interference and Monitored Bandwidth	184
5.4.2	Interference Scenario Detector Input	184
5.4.3	Detector Q1	187
5.4.4	Detector R1	188
5.4.5	Detector S1	189
5.4.6	Interference DCW, B-13 Pulse, $\text{SNR}_{\text{out},B}=8$ dB	191
5.4.7	"", m-sequence CW, $\text{SNR}_{\text{out},B}=8$ dB	192
5.4.8	"", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=8$ dB	192
5.4.9	"", m-sequence CW, Conventional methods, $\text{SNR}_{\text{out},B}=8$ dB	193
5.4.10	"", B-13 Pulse, $\text{SNR}_{\text{out},B}=15$ dB	195
5.4.11	"", m-sequence CW, $\text{SNR}_{\text{out},B}=15$ dB	195
5.4.12	"", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=15$ dB	196
5.4.13	"", m-sequence CW, Conventional methods, $\text{SNR}_{\text{out},B}=15$ dB	197
5.4.14	"", B-13 Pulse, $\text{SNR}_{\text{out},B}=25$ dB	198
5.4.15	"", m-sequence CW, $\text{SNR}_{\text{out},B}=25$ dB	199
5.4.16	"", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=25$ dB	200
5.4.17	"", m-sequence CW, Conventional methods, $\text{SNR}_{\text{out},B}=25$ dB	200
5.5.1	Partial Interference DCW, B-13 Pulse, $\text{SNR}_{\text{out},B}=8$ dB	202
5.5.2	""", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=8$ dB	203
5.5.3	""", B-13 Pulse, $\text{SNR}_{\text{out},B}=15$ dB	205
5.5.4	""", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=15$ dB	206
5.5.5	""", B-13 Pulse, $\text{SNR}_{\text{out},B}=25$ dB	207
5.5.6	""", B-13 Pulse, Conventional methods, $\text{SNR}_{\text{out},B}=25$ dB	208
6.1.1	Extracted SVA Results 1	211
6.1.2	Extracted SVA Results 2	212
6.1.3	Extracted MTM Results 1	213
6.1.4	Extracted MTM Results 2	213

6.2.1	Integrated ESM System	215
A.4.1	Nyquist bands of ADC	220
A.5.1	Sample File 0 Rectangular-windowed Spectrogram	222
A.5.2	""Hann-windowed Spectrogram	223
A.5.3	""SVA Spectrogram	224
A.5.4	Sample File 8 Rectangular-windowed Spectrogram	226
A.5.5	""Hann-windowed Spectrogram	227
A.5.6	""SVA-windowed Spectrogram	228
A.5.7	Sample File 11 Rectangular-windowed Spectrogram	230
A.5.8	""Hann-windowed Spectrogram	231
A.5.9	""SVA Spectrogram	232
A.6.1	Heli Trials Example Periodogram 1	234
A.6.2	Heli Trials Example 2	235
C.1.1	Waterfall Methodology	243
C.2.1	MBD Methodology	245
C.3.1	SELEX Galileo MBD Methodology	245
C.4.1	System Generator [®] Design Flow	248
C.4.2	Example System Generator [®] Window	248
C.5.1	Xilinx [®] ISE Project Navigator	250
C.5.2	Example RTL diagram	250
C.5.3	FPGA Editor	252
D.1.1	Implementation Decisions	254
D.2.1	SVA Module Output	256
D.3.1	Detector A2	258
D.3.2	Detector B2	259
D.3.3	Detector C2	260
D.3.4	Detector E2	261
D.4.1	Detector A3	263
D.4.2	Detector B3	264
D.4.3	Detector C3	265
D.4.4	Detector E3	266
D.6.1	Full DCW, Fixed-point	278
D.6.2	"" , SVA ROC curves	278
D.6.3	Interference Full DCW, SNR _{out,B} =15 dB, sep=2, Fixed-point	280
D.6.4	"" , SNR _{out,B} =15 dB, sep=2, SVA ROC curves	281
E.1.1	α Optimisations	287
F.2.1	SVA-on-FPGA Pseudo Algorithm	291
F.2.2	""-"" , Phase 1 Model	294

F.2.3	"-","-","-", Phase 1 Xilinx [®] Blocks	295
F.2.4	SVA hardware co-simulation	298
F.2.5	SVA-on-FPGA, Phase 2 Model	299
F.2.6	"-","-","-", Phase 2 Xilinx [®] Blocks	300
F.2.7	"-","-","-", Sinusoid	304
F.2.8	"-","-","-", Sinusoid Plus AWGN	306
F.2.9	"-","-","-", Two Sinusoids Plus AWGN	308
F.3.1	"-","-","-", Full DCW	311
F.3.2	"-","-","-", Interference Full DCW	311
F.5.1	Proposed SVA-on-FPGA Architecture	316
G.4.1	UK TV Signals	319
G.4.2	Spreading Factor	321
G.4.3	RF/IF Chain	322
G.5.1	Navajo Trials Example Spectrogram	324
G.5.2	Isolation Filter	325
G.5.3	Navajo DPSS MTM Detections	327
G.5.4	Navajo Sine Taper MTM Detections	327
G.5.5	Navajo SMP Detections	328
G.5.6	Navajo TSMP MTM Detections	328
H.2.1	Tree-like Smoothing Modules	331
H.2.2	Tree Smoothing Fixed-point Error	332
H.2.3	Detector M2	333
H.2.4	Fx-pt TSSP Detector	335
H.2.5	Fx-pt Tree Smoothing Filter	336
H.4.1	Full DCW, Fixed-point, B-13 Pulse	340
H.4.2	Interference Full DCW, Fixed-point, SNR _{out,B} =15 dB, B-13 Pulse . .	342
I.2.1	TSP-on-FGPA Phase 1 Model	345
I.2.2	TSP-on-FGPA Phase 2 Model	349
I.3.1	Full DCW TSP-on-FPGA ROC Curves	351
I.3.2	Interference Full DCW TSP-on-FPGA ROC Curves	352

Contributions

This thesis contains two main areas of original contribution: (1) the use of SVA as a detector and (2) the application of the multitaper methods (MTMs) and periodogram smoothing techniques in pursuit of the detection of LPI radar waveforms. Both of these contributions are set in the context of achievable, real-time application in an ESM DRx. The supporting chapters (2 and 4) and appendices (A–I) greatly expand on these two themes.

The use of SVA in a detector has been filed as patent application number 1104107.6 with reference 001510GB, entitled “*Radio Frequency Digital Receiver System and Method*”.

Subsequently the use of SVA as a detector was included as an aspect in a paper for *Defense Applications of Signal Processing* (DASP) 2011, entitled “*Spatially Variant Apodization for Conventional and Sparse Spectral Sensing Systems*”.

Both of these documents have been inserted exactly as filed/submitted in appendix J.

Preface

Stephen Clark and Robert Cooper at SELEX Galileo originally posed the problem for this EngD project in terms of the detection of LPI radar waveforms. The problem statement was

Classically the best way of detecting a particular signal in noise is to use a matched filter. In fact, if the noise is white and Gaussian, it is the optimal procedure. Many communications and radar systems use the matched filter as the first line or corner stone of all detection and decoding.

Electronic warfare (EW) receivers must detect radar pulses and use the measured characteristics of the transmission to identify whether they are associated with potentially hostile systems. An important consideration is that the Electronic Warfare receiver should be able to make this observation before the radar is able to form a track on the host platform this is known as the range advantage. The disadvantage suffered by the EW receiver is in it not being matched to the transmitted pulse, and the worse the match is the more the range advantage is degraded. Modern EW systems employ digital receiver technology similar to that used in software defined radios. Digital EW receivers offer the ability to improve detection of radar pulses over earlier technology based on video detection, but current techniques have limited efficacy.

The matched filter technique employed by all modern radars is very powerful, inspiring the use of transmitted signals that are observed with degraded signal/noise by the opportunistic listener but which are readily detected by an appropriate listener who has knowledge of the transmitted signal. Faced with an environment of transmitted signals, what can we do as an opportunistic listener who does not know the transmitted waveform and thus cannot access the signal/noise improvements offered by the matched filter?

This is the challenge of the current project. One way forward, for example, is to exploit the restrictions imposed on the engineers who design the transmitted waveforms; a further possibility is to pursue recently developed concepts to do with soft detection theory and time-frequency analysis. The problem is multi-faceted and there are many dimensions to the solution space, but all solutions have to be mindful of the constraints of available processor speed, architecture and power.

Later the problem was expanded to improve the detection of conventional radar waveforms as it became clear that the use of the signal processing technique SVA would help with this. There were many signal processing techniques explored during the course of the research for this EngD problem, such as exploiting cyclostationarity and goodness-of-fit techniques to name a couple. However only the interesting and significant findings have made it into this thesis as main chapters, with internal reports being produced for SELEX Galileo on the other topics. During the course of the research the interesting signal processing techniques were applied in a MBD methodology as we were keen for the EngD to have a tangible output for SELEX Galileo.

The thesis is organised such that after the first background chapter, the next two chapters contain the SVA findings and the following two chapters contain the MTM and periodogram smoothing findings. These chapters are believed to be of most interest to examiners.

The appendices contain details of the signal processing techniques chosen for implementation in a MBD methodology, i.e. the practical side to the EngD.

— *Graham Henry MacKerron*
Edinburgh, 2011

Acknowledgements

Foremost I am extremely grateful to my academic and industrial supervisors: Professor Bernie Mulgrew, Chief Technologist Stephen Clark and Project Manager Systems Team Leader Robert Cooper. I could not have asked for a better academic supervisor than Bernie who is knowledgeable, encouraging, patient and approachable. I also enjoyed the time spent on my EngD project at SELEX Galileo, Capability Green under the supervision of Steve and Bob who did everything to facilitate my research and ensure my research was implementable.

In addition to Steve and Bob, there were many others whom I would like to thank for gladly aiding me with my research at Capability Green: Neil Cade, Bob Probyn, Graham Brown, Oliver Sims, John Saunders, Theo Theodossiou, Peter Seilly, Tony Cross, Paul Newham and Matthew Westwood.

I would also like to extend my thanks to the rest of the staff at Institute for Digital Communications, Edinburgh University and floor A1N3 at Capability Green.

In my personal life I thank my parents for supporting me throughout all my studies at University, my late grandma “Mama” for our trips out in the car around Auchterader and my Gran and late Grandpa for their words of encouragement. Finally, I am grateful to my girlfriend, Jenny, for putting up with my late evenings home from work whilst writing up my thesis.

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Engineering Doctorate Scheme and SELEX Galileo.

MATLAB, Simulink, Fixed-Point Toolbox, Embedded MATLAB and Simulink HDL Coder are registered trademarks of The MathWorks, Inc. Xilinx, the Xilinx logo and System Generator are registered trademarks of Xilinx, Inc. Their use is covered by *fair use* for the purpose of identification of products and companies and the characterisation and features of products.

Declaration

I declare that, except where explicit reference is made to the contribution of others, that this thesis entitled “*The detection of unknown waveforms in ESM receivers: FFT-based real-time solutions*” is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Signature

Name

Date

Graham H MacKerron

October 31, 2011

Definitions/Abbreviations

Abbreviations

AC	alternating current
ADC	analog-to-digital converter
AF	ambiguity function
ANSI	American National Standards Institute
AOA	angle-of-arrival
API	application programming interface
AR	autoregressive
AWGN	additive white Gaussian noise
BMWD	binary moving window detector
BPSK	binary phase-shift keying
CDF	cumulative distribution function
CORDIC	coordinate rotation digital computer
CPS	cyclic power spectrum
CSC	cyclic spectral coherence
CW	continuous wave
DC	direct current
DCW	data capture window
DFT	discrete Fourier transform
DPSS	discrete prolate spheroidal sequence
DRx	digital receiver
DSP	digital signal processing
DTFT	discrete-time Fourier transform
DWT	discrete wavelet transform
ELB	extended low band
ELINT	electronic signals intelligence
EM	electromagnetic
ERP	effective radiated power
ESM	electronic support measures
ESPRIT	estimation of signal parameters via rotational invariance techniques

EW	electronic warfare
FFT	fast Fourier transform
FH	frequency hopping
FIFO	first-in-first-out
FIR	finite impulse response
FM	frequency modulation
FMCW	frequency modulated continuous wave
FPGA	field-programmable gate array
FS	full scale
FSK	frequency-shift keying
FSM	finite state machine
FT	Fourier transform
GIGO	garbage-in-garbage-out
GPS	global positioning system
HDL	hardware description language
i.i.d.	independent identically distributed
IF	intermediate frequency
IFM	instantaneous frequency measurement
IO	input/output
IP	intellectual property
IPR	impulse response
ISAR	inverse synthetic aperture radar
JTAG	Joint Test Action Group
KKT	Karush-Kuhn-Tucker
LFM	linear frequency modulation
LO	local oscillator
LPI	low probability of intercept
LSB	least significant bit
LUT	look-up table
MACC	multiply-accumulate
MBD	model-based design
MOP	modulation-on-pulse
MOS	minimum operational sensitivity

MSB	most significant bit
MSE	mean square error
MTM	multitaper method
MUSIC	multiple signal classification
MVSE	minimum variance spectral estimation
NEB	noise equivalent bandwidth
OR	logical OR gate
PAF	periodic ambiguity function
PDE	partial differential equation
PDF	probability density function
PDW	pulse descriptor word
PN	pseudo-noise
PRF	pulse repetition frequency
PRI	pulse repetition interval
PSD	power spectral density
PSK	phase-shift keying
QIFFT	quadratic interpolated FFT
RAM	random access memory
RF	radio frequency
rms	root mean square
ROC	receiver operating characteristic
RSSR	rapidly-swept superheterodyne receiver
RTL	register transfer level
RV	random variable
RWR	radar warning receiver
SAR	synthetic aperture radar
SCF	spectral correlation function
SDFT	sliding discrete Fourier transform
SIR	signal-to-interference ratio
SMP	smoothed modified periodogram
SNR	signal-to-noise ratio
SP	smoothed periodogram
SSP	smoothed SVA periodogram
STFT	short-time Fourier transform
SVA	spatially variant apodization

TFR	time-frequency representation
TOA	time-of-arrival
TP	total power
TSMP	tree smoothed modified periodogram
TSP	tree smoothed periodogram
TSSP	tree smoothed SVA periodogram
TV	television
UHF	ultra high frequency
WOSA	Welch's overlapped segment averaging

List of Symbols

α	SVA parameter
α_c	Cycle frequency
α_K	Kaiser window parameter
B_{ESM}	Bandwidth of ESM system
β	Complex SVA parameter
β_I	Imaginary part of complex SVA parameter
β_R	Real part of complex SVA parameter
B_N	Bandwidth of DCW frequency bin filter
$\hat{S}_{xx}^{(\text{BT})}$	Blackman-Tukey spectral estimate
C	Cost function
$\chi_{\text{AF}}(\tau, f)$	Ambiguity function
d	Degrees of freedom
δ	Kronecker delta
δ_{det}	Ratio of power required to detect at ESM receiver to power required at radar receiver
δ_n	CORDIC rotation direction
E_x	Waveform energy
f	Frequency
f_c	Carrier frequency
f_{chip}	Chip-rate frequency
f_s	Sampling rate
FS_{dBm}	Full scale signal power
f_T	Target frequency
g	High pass filter impulse response
G_{AESM}	Gain of ESM antenna in direction of radar
$\Gamma(\dots)$	Gamma function
γ_e	Non-coherent integration efficiency factor
G_{ESM}	Gain of ESM antenna and receiver system in direction of radar
G_R	Gain of radar receive antenna in direction of target
G_{ESMr}	Gain of ESM receiver
G_T	Gain of radar transmit antenna in direction of target
h	Filter impulse response
H_0	Null hypothesis
H_1	Alternative hypothesis
hyp	Bandwidth hypothesis index
I_0	Modified zero-order Bessel function

K	CORDIC algorithm gain
k	Frequency bin index
k_B	Boltzmann's constant
L	Filter span
λ	Wavelength
L_{ESM}	Loss in ESM receiver
L_R	Loss in radar receiver
L_{sc}	Scalloping loss
$\hat{S}_{xx}^{(\text{MTM})}$	MTM spectral estimate
$\hat{S}_{p,xx}^{(\text{MTM})}$	p^{th} MTM eigenspectrum
N	Number of samples used in DCW
n	Time-domain sample
N_b	Number of bits
N_{DFT}	Number of samples in DFT
N_F	Noise figure
N_f	Noise factor
N_{FFT}	Number of samples in FFT
N_L	Number of overlapped samples
N_p	Noise power
$O(\dots)$	Computation complexity order
P	Number of tapers used
P_d	Probability of detection
$\hat{S}_{xx}^{(\text{P})}$	Periodogram spectral estimate
P_{fa}	Probability of false alarm
ϕ	Phase
$Q \sim \chi^2(d)$	Chi-square distribution with d degrees of freedom
ϕ_0	Initial phase
$\Phi_{\text{CDF}}(\dots)$	Standardised normal distribution CDF
$\phi_{\text{PDF}}(\dots)$	Standardised normal distribution PDF
Pow_r	Power received by a receiver
Pow_t	Transmitted power
S_{xx}	Power spectrum density function
R_{ESM}	ESM maximum detection range
ρ	Reference resistance
R_{Rad}	Radar maximum detection range
R_{xx}	Autocorrelation function
σ	Normal distribution standard deviation
σ_c	Radar cross-section of the target
σ_x^2	Variance of waveform
SNR_{out}	SNR after FFT

$\hat{S}_{xx}^{(\text{SP})}$	Smoothed periodogram spectral estimate
$\lambda(N, W_h)$	Spectral concentration
$\lambda_p(N, W_h)$	Spectral concentration of p^{th} taper
$\mathcal{S}_{xx}(f, \alpha_c)$	Cyclic spectrum density function
$\mathcal{SC}_{xx}(f, \alpha_c)$	Spectral correlation function
s_{pf}	Power spread factor
$s(t)$	Complex baseband pulse
T_ν	Nuttal's power law detector
T	DCW period
t	Time
T_0	Absolute standard temperature (290 Kelvin)
τ	Delay
T_C	Coherent integration time of the radar
T_{dwell}	Radar dwell time on a target
T_{ESM}	Non-coherent integration time of the ESM system
θ	Angle
$W_B(k)$	Bartlett window function in frequency-domain
$W(f)$	Continuous frequency-domain window function
W_h	Normalised half-bandwidth parameter
$W(k)$	Discrete frequency-domain window function
w_K	Order-K time-domain window function
W_N	Twiddle factor
$w(n)$	Time-domain window function
$W_p(k)$	p^{th} MTM taper in frequency-domain
$w_p(n)$	p^{th} MTM taper in time-domain
WV	Wigner Ville distribution
$x(t)$	Complex waveform in time-domain
$X_a(k)$	Filtered DFT
$X(f)$	Complex waveform in frequency-domain
$X_I(k)$	Imaginary part of DFT
$X(k)$	DFT output
$\hat{X}_p^{(\text{MTM})}$	p^{th} DFT for MTM
$\hat{X}_p^{(\text{WOSA})}$	p^{th} DFT for WOSA
$X_R(k)$	Real part of DFT

CHAPTER 1

Introduction & Background

1.1 Introduction

This thesis is laid out such that the content of academic interest is located in the main body of chapters whilst the industrial, practical implementation side is provided by the appendices.

The background to: (1) low probability of intercept (LPI) radar signals; (2) conventional system detection and (3) grouped techniques for improved LPI detection by electronic support measures (ESM) is given later in this chapter.

Following on from this two chapters describe the theory behind spatially variant apodization (SVA) and its application in the detection of conventional sinusoidal radar signals.

Chapters 4 and 5 describe computationally-efficient spectral estimation techniques which are suitable for improved detection of broader-band LPI radar signals in ESM systems.

Conclusions drawn from the main body of chapters and suggested future work feature in chapter 6. Estimates are provided for the improvement in the SELEX Galileo digital receiver (DRx) system detection range in the conclusion.

The appendices are designed to show the work is aligned with the subject matter of **system level integration**. This is through the concept of model-based design (MBD), which is specifically described in appendix C. Its application to SVA and a broader-band LPI radar signal detection method (periodogram smoothing) are methodically described with an appendix (D, F, H, I) devoted to each stage in the MBD methodology.

Interspersed throughout the MBD flow appendices are details of experiments conducted to apply the SVA and broader-band detection techniques to field data (appendices A and G) and suggested computational savings for applied SVA (appendix E).

1.2 The Range Advantage

The tactical situation between a radar and an ESM system is something of a competition in which both attempt to detect the other first. A particularly colourful analogy of this is given in [1] between Greek legends Achilles and Hector where the central

message is that both sides wish “to see and not be seen.” However there are different fundamental limitations which are placed on both sides and these are well documented, for example in [1, 2]. The ESM system has the strong advantage that it is a one-way path from the radar to the platform on which the ESM system is installed whereas a radar requires a two-way path from radar to platform and back again to the radar. In the highly simplified and idealised form the power density which arrives at the platform decreases with the square of the distance between the radar and platform but the power density of the received returns at the radar decreases with the fourth power of the distance. This effect is depicted in Fig 1.2.1.

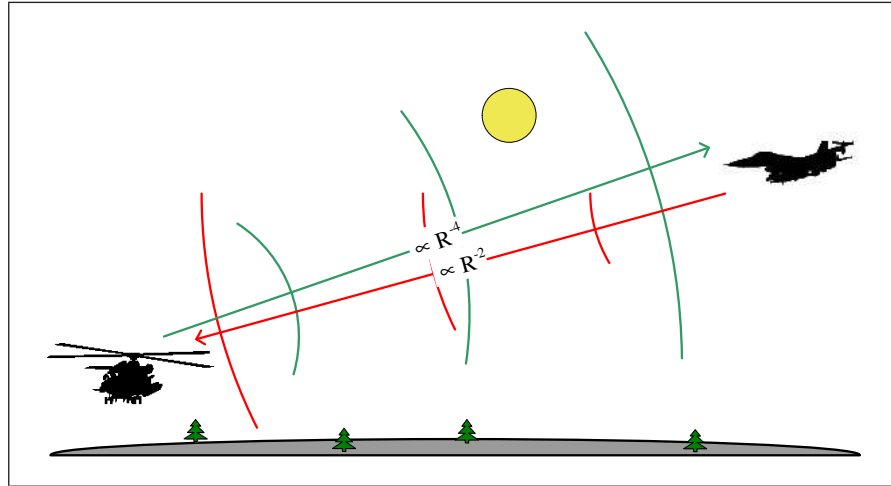


Figure 1.2.1 Range advantage of an ESM system on an airborne platform

A rigorous demonstration of the range advantage of the ESM system is given in [2]. When the simplifying assumption is made that the target the radar wishes to detect and the ESM platform are the same (i.e. detection of the radar's mainlobe with a radar warning receiver (RWR)), the detection range of the ESM system in terms of the detection range of the radar is given by

$$R_{\text{ESM}}^2 = R_{\text{Rad}}^4 \left[\frac{4\pi}{\delta_{\text{det}}} \frac{1}{\sigma_c} \frac{G_{\text{AESM}}}{G_R} \frac{L_{\text{ESM}}}{L_R} \right] \quad (1.2.1)$$

where:

- δ_{det} is the ratio of power required to detect at the ESM receiver to the power required at the radar receiver;
- G_R is the gain of the radar receive antenna in the direction of the ESM platform;
- G_{AESM} is the gain of the ESM's antenna in the direction of the radar (close to unity for omnidirectional antennas, i.e. the tactical situation);
- L_{ESM} and L_R are losses suffered in the ESM and radar receivers respectively;
- σ_c is the radar cross-section of the ESM platform.

The range advantage is evident from Eq (1.2.1) with the square and fourth powers due to the one-way and two-way paths.

The strong advantage of the radar is it knows to a high degree of accuracy the waveform that was transmitted. This concept is exploited by some LPI radar waveform designers through utilisation of more and more complicated coded waveforms which allow the radar to maximise its coherent integration time and thus its sensitivity. This concept would result in increases in the value of δ_{det} in Eq (1.2.1). In fact in [2] δ_{det} is defined as

$$\delta_{\text{det}} = (B_{\text{ESM}}T_C)^{1-\gamma_e} (T_{\text{dwell}}/T_{\text{ESM}})^{\gamma_e}$$

where:

- B_{ESM} is the bandwidth of the ESM system (which is almost certainly much larger than the radar bandwidth);
- γ_e is a non-coherent integration efficiency factor;
- T_C is the coherent integration time of the radar;
- T_{dwell} is the radar dwell time on a target;
- T_{ESM} is the non-coherent integration time of the ESM system.

Coherent integration time in a radar is maximised through the use of a matched filter, which is the conjugate time-reversal of the filter used to code the transmitted waveform. The matched filter concept is illustrated for digital filters in Fig 1.2.2 where panel (a) shows the transmitter side and panel (b) shows the receiver side.

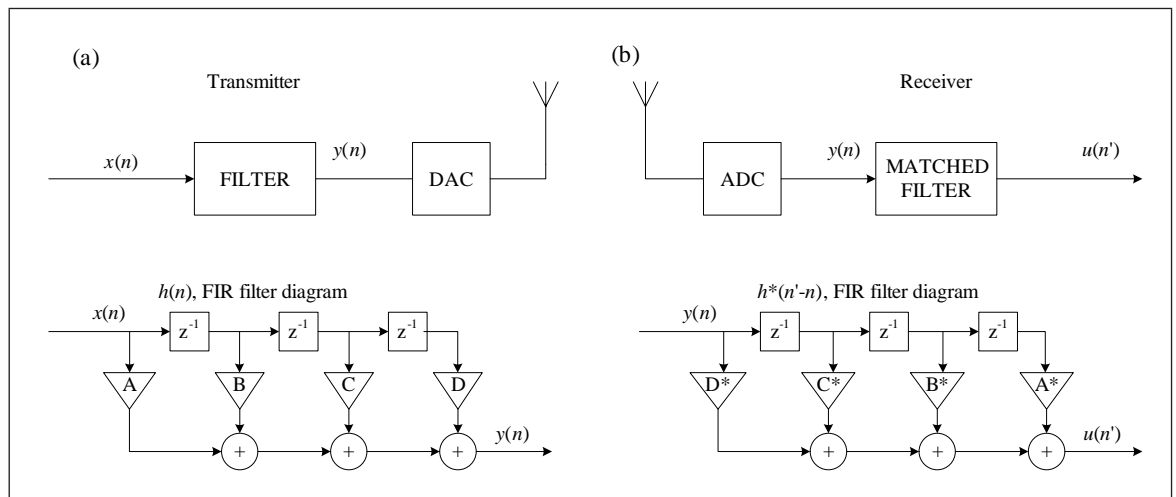


Figure 1.2.2 Matched filter concept with digital filters

The advent of LPI waveforms and the associated increase in the value of δ_{det} means radars which exploit this tactic can not be simply countered to the range advantage that was once offered when only pulsed sinusoid radars had to be considered.

The theme for this thesis is to attempt regain some of the lost ESM range advantage due to coded LPI waveforms in the tactical situation.

1.3 History of ESM System Development

This section gives a very brief summary of the main types of receivers upon which ESM systems have been typically based. The types of receivers are detailed in a loosely historical order.

The general trend for ESM systems over the years has been:

- to intercept a greater proportion of the electromagnetic (EM) spectrum;
- to incorporate greater degrees of autonomy;
- to increase azimuth and elevation coverage.

For airborne platforms, the bandwidth from 500 MHz to 20 GHz is considered essential for an ESM system. Likewise, antenna coverage of 360° azimuth and significant elevation is considered essential. The most advanced airborne ESM systems cover up to 40 GHz and have an almost spherical coverage.

The simplest view of the role of an ESM system is shown in Fig 1.3.1. The most important stage in Fig 1.3.1 as an initial building block for improved detection of LPI waveforms is the detection stage and this is considered in this thesis.

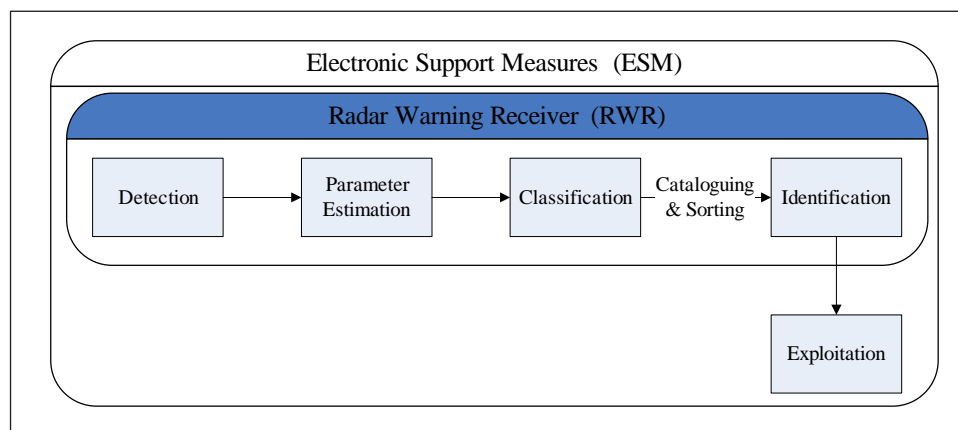


Figure 1.3.1 A top-level view of the functionality of a RWR and its relationship within an ESM system

Early ESM systems were based on wide-band crystal video receivers. These are analogue receivers which essentially accept a radio frequency (RF) input, remove the RF carrier and form an envelope signal. The envelope signal is passed through the functional blocks shown in Fig 1.3.1, where the pulses are detected on a pulse-by-pulse basis and attempts are made to reconstruct the original radar pulse train. The detection component is closely related to the concepts in a radiometer. Early ESM systems were used to counter simple radar signals in a relatively “quiet” operational environment.

As technology improved instantaneous frequency measurements (IFMs) were added to crystal video detectors which afforded the ESM system extra information about the frequency of a pulse. Unfortunately IFM measurements are only reliable against one signal at a time. This became a greater concern when complicated operational environments emerged where many pulses can be temporally coincident at the ESM receiver. This is indicative of the general trend in electronic warfare (EW) where a greater proportion of the EM spectrum is utilised as time passes.

Therefore chronologically the next feature to be added to ESM systems was frequency selectivity, generally in the form of rapidly-swept superheterodyne receivers (RSSRs), which are tunable, channelised systems. The RSSR sweeps across an enormous bandwidth in narrower-band “ESM dwells”. The ESM dwell times must be very short in comparison to the radar dwell time if the ESM system is to be able to capture the radar pulses.

As such, these receivers process only a small RF bandwidth or ESM dwell at a time, which leads to greater receiver sensitivity and ESM detection range (R_{ESM}) but has the disadvantage that the receiver is instantaneously “blind” to the other ESM dwells. The superheterodyne component of a RSSR downconverts the input RF signal to an intermediate frequency (IF), which is passed to a crystal video receiver which acts as the detector.

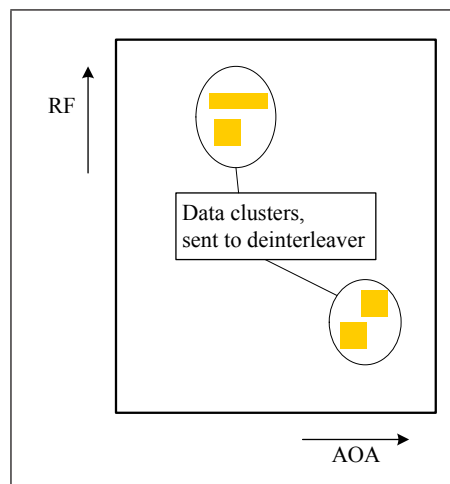


Figure 1.3.2 Clustering in frequency and AOA

By itself the frequency selectivity offered by an RSSR is typically insufficient and many more signal processing stages exist between the detection and classification/identification stages shown in Fig 1.3.1. These extra stages are required to handle the sheer number of detected pulses per second before a deinterleaver attempts to reconstruct the radar pulse trains in the classification and identification stages. The pulse trains are reconstructed through the use of clustering. This is demonstrated in Fig 1.3.2 which shows detected pulses by two characteristics: RF frequency and AOA (frequency and spatial selectivity). Pulses from the same emitters tend to cluster and these data clusters are

sent to the deinterleaver.

Clustering to vastly reduce the number of pulses passed to the deinterleaver is used for two reasons:

1. Only a dozen or so pulses from multiple “looks” at the same ESM dwell are required for identification (several pulses are required to identify stable emitters but tens of pulses are required for complicated emitters).
2. Technological limitations mean that signal analysis in software cannot cope with $\sim 50\,000$ pps.

The next milestone in the development of ESM systems was to replace the crystal video receiver with a DRx and create a hybrid system. Besides permitting digital processing techniques, use of a DRx: (1) removes the requirement for as many expensive RF, IF and video signal cables and (2) is not as expensive to maintain or replace as analogue equivalents.

The generic operation of the rapidly-swept superheterodyne system and DRx combination is to:

1. obtain an RF input;
2. downconvert to IF;
3. digitise through an analog-to-digital converter (ADC);
4. apply a variant of the Fourier transform (FT) to the samples in data capture windows (DCWs).

The FT is usually implemented with the use of a fast Fourier transform (FFT) algorithm or less often as a sliding discrete Fourier transform (SDFT). An excellent resource on the operation of DRxs is given by Tsui in [3].

It could be argued that the “instantaneous blindness” of the superheterodyne and DRx combinations could be overcome with a fully channelised system which consists of a number of DRxs each assigned to constantly monitor contiguous RF channels. However the cost to build such a system would be prohibitive. This is especially true because most ESM systems operate on data from a number of antennas to form maximal angular coverage and hence this multiplies up the number of channels.

The FFT algorithm in DRxs provides frequency-domain data which is passed to the detection stage in Fig 1.3.1. The FFT serves to further channelise the ESM dwells and provides greater frequency selectivity.

The next logical step which involves the FFT algorithm for DRxs is to link together consecutive DCWs to form a short-time Fourier transform (STFT). This forms a time-frequency history of the operational environment and can be subjected to a greater

number of digital signal processing (DSP) techniques. This trend is highlighted by research such as in [4].

1.4 Windowing/Apodisation

Windows are snapshots of data over a finite interval. Window functions themselves are defined to be zero outside the interval and a particular shape inside the interval such that the product of data multiplied by the window function is zero outside the interval and tapered by the window function shape inside the interval. Almost all practical window functions have a shape which resembles a bell-shaped curve, although their exact mathematical descriptions vary.

Window functions play a major role in spectral analysis which takes place in a DRx as part of an ESM system. The ESM system must periodically make decisions on whether it believes threats are present in the operational environment (decisions made in the order of hundreds of nanoseconds). In order to do so it is forced to take windows (given the name DCWs) of the constant stream of data and perform some spectral analysis by means of a discrete Fourier transform (DFT) to assess if there are any spectral peaks caused by potential threats. The discrete window functions are applied through multiplication of the window function sample by sample with the DCW.

The action of windowing causes some side-effects in the spectral estimate and these side-effects are directly linked to the shape of the window function. Harris's key paper [5] discussed shapes for windows, or in other words, window functions and the metrics by how they might be compared. A diagram which summarises the metrics by way of an example (a rectangular window function) is shown in Fig 1.4.1. Table 1.4.1 displays the values of the metrics for the window functions mentioned throughout this thesis. The table is based on the window functions table in Harris's paper. The metrics described in Harris's paper are described below.

1.4.1 NEB

The **noise equivalent bandwidth (NEB)** or equivalent noise bandwidth of a window function is a measure of the increase in the noise floor caused by the use of a window function ($w(n)$) relative to a rectangular window function. The increase in the noise floor is equivalent to an increase in the bandwidth of a hypothetical rectangular window which acts as a filter. Therefore it is given by the normalised ratio of the noise power to the peak signal power

$$\text{NEB} = \frac{\sum_{n=0}^{N-1} w(n)^2}{\left[\sum_{n=0}^{N-1} w(n)\right]^2} \quad (1.4.1)$$

where the non-rectangular window function ($w(n)$) consists of N coefficients from $n = 0, \dots, N - 1$. Processing gain, also defined in [5], is the reciprocal of NEB.

The NEB parameter is not apparent in Fig 1.4.1. To see this, additive white Gaussian noise (AWGN) would need to be processed (alongside the constant envelope signal) with a few different windows and the results would be required to be normalised to the rectangular-windowed noise floor. The NEB of other non-rectangular window functions would show as a loss in the peak response.

1.4.2 Overlap Correlation

Successive DFTs which overlap in the time-domain not only assist in mitigation against the unknown times-of-arrival (TOAs) of signals (as described in section 1.7) but they also help recover some signal information which is lost due to attenuation at the edges of bell-shaped window functions. This is useful to obtain as much signal information as possible for non-coherent integration or averaging of successive DFTs. However the overlapped samples contain a small amount of noise which is correlated and this reduces the variance reduction during non-coherent integration. The **overlap correlation** parameter quantifies this feature.

1.4.3 Cusping Loss

The DFT is a sampled version of the discrete-time Fourier transform (DTFT) in the frequency-domain. However the frequency samples are unlikely to match up exactly to the target signal frequency and therefore there is a processing loss (or cusping loss) which increases up to a maximum scalloping loss when the target frequency is exactly midway between DFT frequency samples. Therefore the **scalloping loss** represents the worst case and is given by

$$\text{Scalloping Loss} = \frac{W(1/2)}{W(0)} \quad (1.4.2)$$

The scalloping loss and NEB move in opposite directions for different window functions. Therefore the two can be combined in a third metric, a worst case net processing loss, which adds the NEB in decibels to the scalloping loss in decibels.

The cusping loss is demonstrated in Fig 1.4.1 as the peak sample of the DFT which does not occur at the same frequency as the peak response of the DTFT. The scalloping loss is also marked on the diagram.

1.4.4 Sidelobe Level

The spectral leakage pattern due to the window function is given by the magnitude-squared of the transform of the window function. The size and location of the spectral leakage sidelobes affects the ability of a detector to find weak signals which are close in frequency to strong signals. The level of the first sidelobe is usually specified as an important metric.

The relative level of the first sidelobe is shown in Fig 1.4.1.

1.4.5 Minimum Resolution Bandwidth

This metric also describes the spectral leakage pattern. The width of the mainlobe of the magnitude-squared of the transform of the window function is measured at its 3.0-dB and 6.0-dB points. The width is specified in terms of DFT frequency bins. This signifies the ability of the window function to discriminate between closely spaced signals of similar power. The conventional **3.0-dB resolution** is the width at the half-power points, however it is the **6.0-dB resolution** which actually determines the resolution of the window.

Both the 3.0-dB and 6.0-dB resolutions are shown in Fig 1.4.1.

1.4.6 Summary

Window	First Sidelobe Level (dB)	NEB (dB)	Scalloping Loss (dB)	Net Proc. Loss (dB)	3.0-dB Res. (bins)	6.0-dB Res. (bins)
Rectangular	−13.3	0	3.92	3.92	0.89	1.21
Hann	−31.5	1.76	1.42	3.18	1.44	2.00
Hamming	−42.7	1.34	1.78	3.12	1.30	1.81
Dolph- Chebychev	−60.0	1.79	1.44	3.23	1.44	2.01
Bartlett	−26.5	1.24	2.60	3.84	1.28	1.78

Table 1.4.1 The metrics of selected window functions mentioned in this thesis

1.4.7 Relevance to ESM

The previous section alluded to the two main weaknesses of a DFT-based detector: (1) a window of a waveform results in spectral leakage and (2) the unknown frequency of the target signal f_T usually results in a cusping loss up to a maximum scalloping loss. It is important to note that the choice of shape of the window in front of the DFT greatly affects both of these features. As an example, the rectangular (i.e. uniform) window function, which is constant within the DCW interval (if the constant is equal to one then the rectangular window function is the same as no window function), creates a spectral leakage pattern where the signal mainlobe is narrow but the sidelobes are quite high. Therefore to avoid erroneous detections from the sidelobes of strong signals, DRxs in ESM systems conventionally use bell-shaped window functions because they cause the sidelobe spectral leakage to be relatively low. The price to pay is a diminished frequency resolution because such bell-shaped window functions broaden the signal mainlobe to

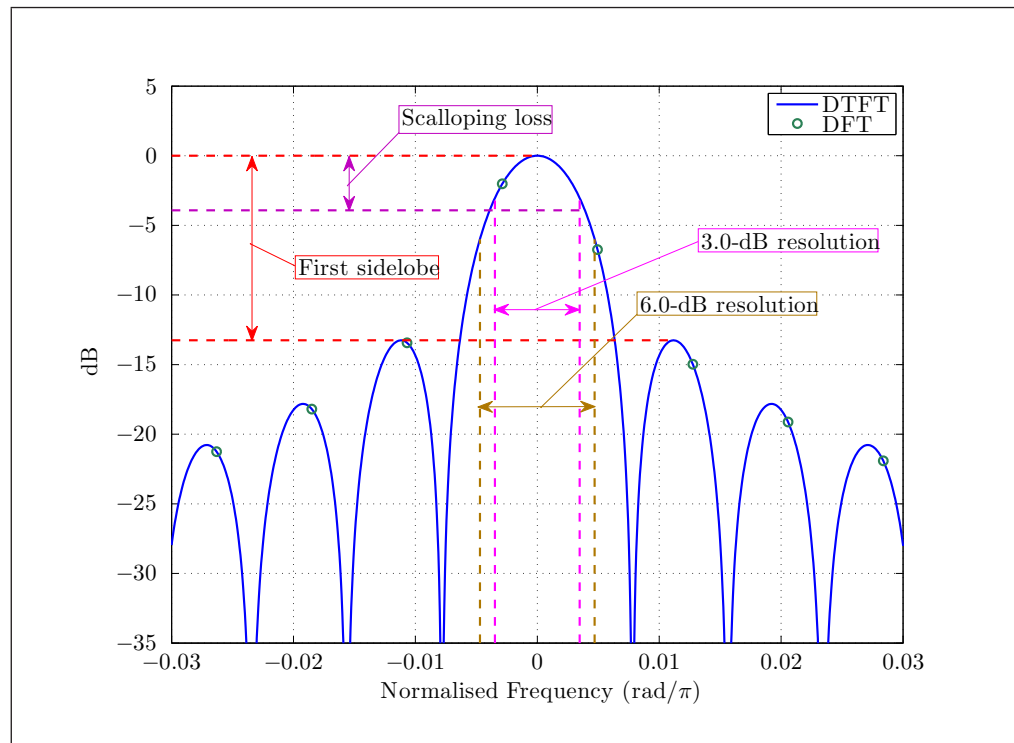


Figure 1.4.1 Window function parameters

some degree and an increase in NEB. Fig 1.4.2 describes the window process in a DRx as a graphic which shows a DCW of part of a sinusoid which is subsequently processed. The process could be undertaken with a rectangular window which would result in a spectral leakage pattern with a narrow mainlobe but high sidelobes or a bell-shaped curve which would result in a spectral leakage pattern with a wide mainlobe and low sidelobes.

1.4.8 Decision Rules

To prevent spurious detections caused by spectral leakage the detector in an ESM system usually has a crude knowledge of the approximate shape of the envelope of the spectral leakage pattern caused by the window function in use. For example if the window function used was a Hann window, then detections around the peak of a strong signal would be prohibited unless they were at least $1.0 \times f_s/N$ away from the peak. Records of the spectral leakage pattern of windows that permit or prohibit further detections are termed **decision rules**.

1.5 Super-resolution Techniques

There are so-called super-resolution techniques which can improve the frequency estimation of a target signal and therefore mitigate against one of the weaknesses of the FFT-based detector. These spectrum estimators are parametric and include some assumptions about the signal in their formulation. Examples of super-resolution techniques include:

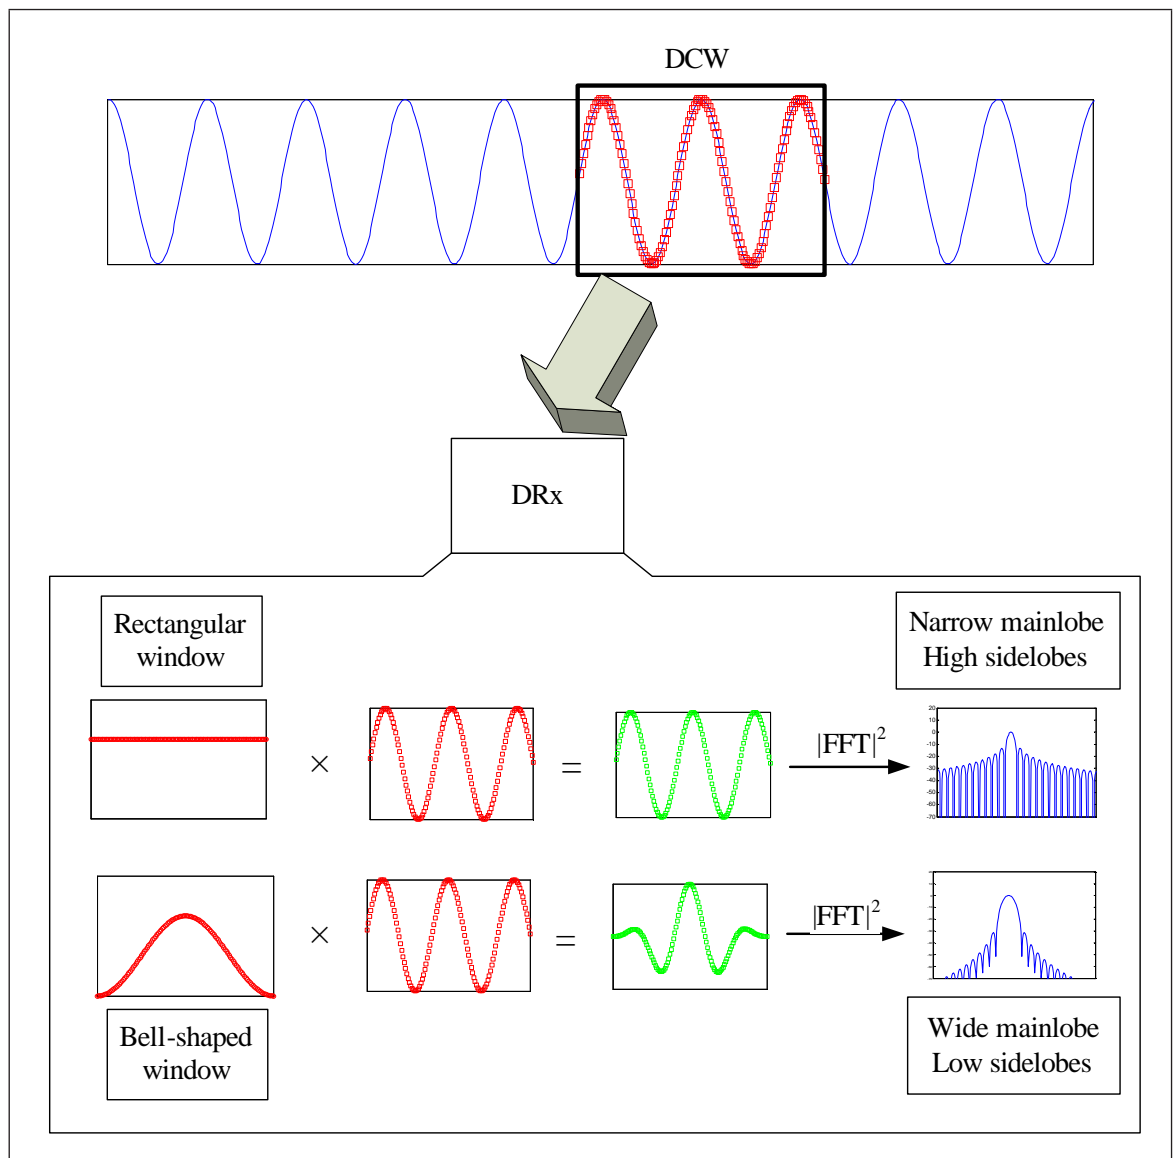


Figure 1.4.2 Introduction to the window process in a DRx

- the autoregressive (AR) method;
- multiple signal classification (MUSIC);
- estimation of signal parameters via rotational invariance techniques (ESPRIT);
- Prony's method;
- minimum norm method.

The route of the super-resolution techniques was not followed in this thesis primarily because the problem statement in section 1.2 does not call for a superresolution requirement.

1.6 Envelope and Energy Detection

Most traditional detection methods are based on signal energy detection or signal envelope detection. Both of these methods produce a statistic which is compared

against a threshold value to make the decision whether a signal is present (the alternate hypothesis H_1) or not (the null hypothesis H_0). Envelope detection examines the magnitude of a waveform, $x(t)$, which may be complex or real-only and energy detection examines the square of the envelope over a chosen window period of time, T . The physical energy of a waveform which might exist in a continuous-time crystal video detector over time T is given by

$$E_x = \frac{1}{\rho} \int_0^T |x(t)|^2 dt \quad (1.6.1)$$

where ρ is the resistance driven by $x(t)$. Where appropriate, a value of $\rho = 1 \Omega$ was taken in simulations in this thesis. In a discrete-time detector such as that in a DRx, Eq (1.6.1) is estimated by

$$E_x = \frac{1}{\rho} \sum_{n=0}^{N-1} |x(n)|^2 \Delta t \quad \text{where } T = N\Delta t$$

The simplest energy detector is a radiometer, which performs the operation in Eq (1.6.1) on a waveform as a function of time over a window of duration T and particular bandwidth B_{ESM} . This creates a value for E_x which is compared against a threshold value. Such a process is termed integrate-and-dump.

The energy detector is a type of detector from the group of Nuttall's power law detectors (T_ν), so-called because they were described by Nuttall in [6]. The Nuttall power law detector in discrete-time is described by

$$T_\nu = \sum_{n=0}^{N-1} |x(n)|^\nu \quad (1.6.2)$$

and if $\nu = 2$ then the energy is estimated.

Other important definitions are (assuming $\rho = 1 \Omega$):

envelope:	$ x(n) $
instantaneous power:	$ x(n) ^2$
average power:	$\frac{E_x}{N\Delta t}$
root mean square (rms):	$\sqrt{\frac{1}{N\Delta t} \sum_{n=0}^{N-1} x(n) ^2}$

Power itself is actually composed of two contributions from the constant direct current (DC) offset and the non-constant alternating current (AC) components.

Power is conventionally taken to be energy per unit time but a waveform can also be described in terms of its frequency content and it becomes meaningful to use the description (spectral) energy density. According to Parseval's theorem the energy in a window will be the same whether measured in the time-domain or frequency-domain.

The theory behind envelope detection and energy detection of square-integrable (or square-summable) waveforms is well documented, for example in [7, 8, 9].

The theory of envelope and energy detection includes many techniques to improve upon the basic threshold scheme through the use of non-coherent integration techniques such as the L -out-of- M detection method (also called a binary moving window detector (BMWD))[10] or double threshold detection method). L -out-of- M detection utilises a second, larger time window or larger frequency bandwidth which spans M detection trials and within this L smaller window detections are required to trigger a detection.

It can be instructive to look at the input ($x(t)$ or $X(f)$) over different time periods and different frequency bands and perform envelope or energy detection in different time-frequency cells. True channelisation involves subdivision of the entire bandwidth into smaller frequency bands and the simultaneous application of detectors to each smaller bandwidth channel. The superheterodyne component of the superhet and DRx combination means that these detectors do not quite fulfill this channelisation criterion but the DFT component in a DRx does. The DFT can be viewed as a filter bank which is composed of N filters with an integration time of the sampling rate (f_s) multiplied

by the N points in a DCW. The filter bandwidths are approximately f_s/N and are immediately followed by a downsample factor of N . Detection can be performed on the output of these filters simultaneously.

The most desirable outcome for an ESM system is when the time-frequency cell size is exactly matched to the unknown radar signal. In other words, when the signal time hops and frequency hops match the cell size. In this case, the radar signal would be entirely localised within the cell and the detector would be more likely to find it against the background noise without the need for further non-coherent integration. Thus the superhet and DRx combination is well-suited against pulsed sinusoidal signals because the FT is the matched filter for sines and cosines, however there are still some weaknesses in this type of detector which are worsened by the use of LPI waveforms.

1.7 Ideal Detection

The previous section describes how it would be highly desirable for a detector in an ESM system to concentrate radar signals into a single time-frequency cell. For a DRx which utilises an FFT, this would happen if the FFT was synchronised with a sinusoidal radar signal at the correct frequencies and intercepted pulses at exactly the right times. In this case the detector would act in a similar fashion to a matched filter, however this synchronisation is unlikely and there are weaknesses in the FFT-based DRx:

1. The arrival time and duration of a radar pulse is unknown to the DRx. These problems are alleviated if the FFTs are overlapped in time. However this adds considerably to the requirements of the ESM system.
2. An inherent difficulty with FFTs is that they operate on a window or snapshot of the waveform and hence the spectral estimate which results is the convolution of the transform of the window shape and the transform of the windowed waveform. The usual practice is to use a bell-shaped window function due to their more preferential transform shape, however this results in a loss in the ability of the DRx to resolve between closely-spaced signals in the frequency dimension and a loss in processing gain.
3. The frequency of the radar signal (f_T) is unknown to the DRx. Therefore the radar frequencies are unlikely to line up exactly with the FFT frequency bins and thus there is a cusping loss up to a maximum scalloping loss.

Radars which utilise LPI waveforms make ideal detection even harder for the detector in an ESM system because they spread their signal energy out over broader bandwidths which makes localisation more difficult.

The impact of the second of these weaknesses is minimised in the first contribution of this thesis.

Despite the difficulties, DRxs in ESM systems are almost certain to include an FFT signal processing element somewhere in the signal processing path due to its strong advantages of:

- the ability to separate radar signals at different frequencies which overlap in time (part of the frequency selectivity feature);
- provision of a frequency estimate for a pulse descriptor word (PDW);
- a processing gain which approaches that of a matched filter against pulsed sinusoids;
- its extremely efficient computational method.

1.8 Radar Designer's Concerns

The situation from a radar designer's point of view has been considered in the course of the analysis of this thesis problem. A radar designer has many factors such as cost, space, power consumption and bandwidth which limit what can be designed and built. Two important limitations which have been incorporated into all simulations in this project are:

1. The vast majority of radar pulses are approximately constant amplitude. This is a consequence of the use of saturated power amplifiers (more efficient than linear amplifiers at high radiated powers).
2. Radars generally have to operate within certain bandwidths to prevent out-of-band-interference. The spectral "skirt" of any radar transmissions are usually required to sit below a predetermined threshold. In the case of phase-coded waveforms this usually translates to a finite period of time for phase transitions.

The radar waveform which is used is very much linked to the philosophy of the radar platform. The waveform chosen depends on the purpose of the radar, as described in [11] and demonstrated by the waveforms in [12]. For example the LPI philosophy is to detect and take action against unsuspecting targets whereas the opposite philosophy is to be bold and allow the radar to be detected as a show of force.

It is believed in sources such as [13, 14, 15] that continuous wave (CW) waveforms will constitute the majority of LPI waveforms. By their nature CW waveforms present a few issues to a radar designer. For example, some method to transmit and receive at the same time must be employed. This can prove costly if the decision to go with two separate antennas is used. If two antennas is not feasible then a reflected power canceller must be used, however the radar would be limited to low radiated powers and hence short ranges.

Another consequence of long coded LPI waveforms is that it is difficult for the radar to have rapid revisit times in its scan modes. Therefore it is likely that LPI waveforms

are more useful for short ranges without the need for scan modes.

An advantage of CW waveforms is that large coherent integration gains can be obtained from long codes. The use of certain long codes also allows perfect cancellation of sidelobes in the periodic autocorrelation function (zero-Doppler cut of the periodic ambiguity function (PAF), see below for a description of the aperiodic version, the ambiguity function (AF)). These codes are termed perfect codes and examples which are categorised in section 1.9 include: the Frank codes, the Zadoff-Chu codes, the P4 code and Golomb codes.

It is stressed in [14] that the main tool used in the design of a radar waveform for finite duration signals is the AF ($\chi_{\text{AF}}(\tau, f)$). This is formed from the application of a matched filter to a delayed and Doppler-shifted version of the transmitted waveform signal as thus

$$\chi_{\text{AF}}(\tau, f) = \int_{-\infty}^{\infty} s(t)s^*(t - \tau)e^{-i2\pi ft} dt$$

where $s(t)$ is the complex baseband pulse (similar notation has been used as from [14]). The ideal AF surface has the shape of a “thumbtack” or Dirac delta function at the origin¹ such that the autocorrelation function at the zero-Doppler cut and the zero-delay cut would both be impulses.

Some example AF magnitude surfaces are shown in Fig 1.8.1². Panel (a) was produced from an uncompressed i.e. uncoded sinusoidal pulse of length $N = 64$ samples. The large ambiguity at zero-Doppler in delay is clearly visible as a triangular-shaped autocorrelation. The magnitude surface in panel (b) was created from an linear frequency modulation (LFM) pulse which was also composed from $N = 64$ samples but it also included a compression ratio (i.e. time-bandwidth product) of 13. In panel (b) the sidelobes were suppressed at zero Doppler because a Hamming window function was used to weight-on-receive. Therefore panel (b) actually depicts a cross-ambiguity function. The range-Doppler coupling (a phenomenon of chirps) is just visible in the plot.

Panel (c) was produced from a Barker-13 coded pulse of length $N = 65$ samples and the compression ratio was 13. The Doppler tolerance of the Barker-13 code is not as appealing as that of a chirp with a similar time-bandwidth product.

Panel (d) shows the AF magnitude surface created from a polyphase, Frank-coded pulse with a compression ratio of 16 and pulse length of $N = 64$ samples. The Frank code is a stepped-phase approximation to a linear chirp and the AF bears some similarities to (b), albeit with degradation.

¹With some form of Doppler processing executed elsewhere

²Negative Doppler range not shown in the plots because $|\chi_{\text{AF}}(\tau, f)| = |\chi_{\text{AF}}(-\tau, -f)|$

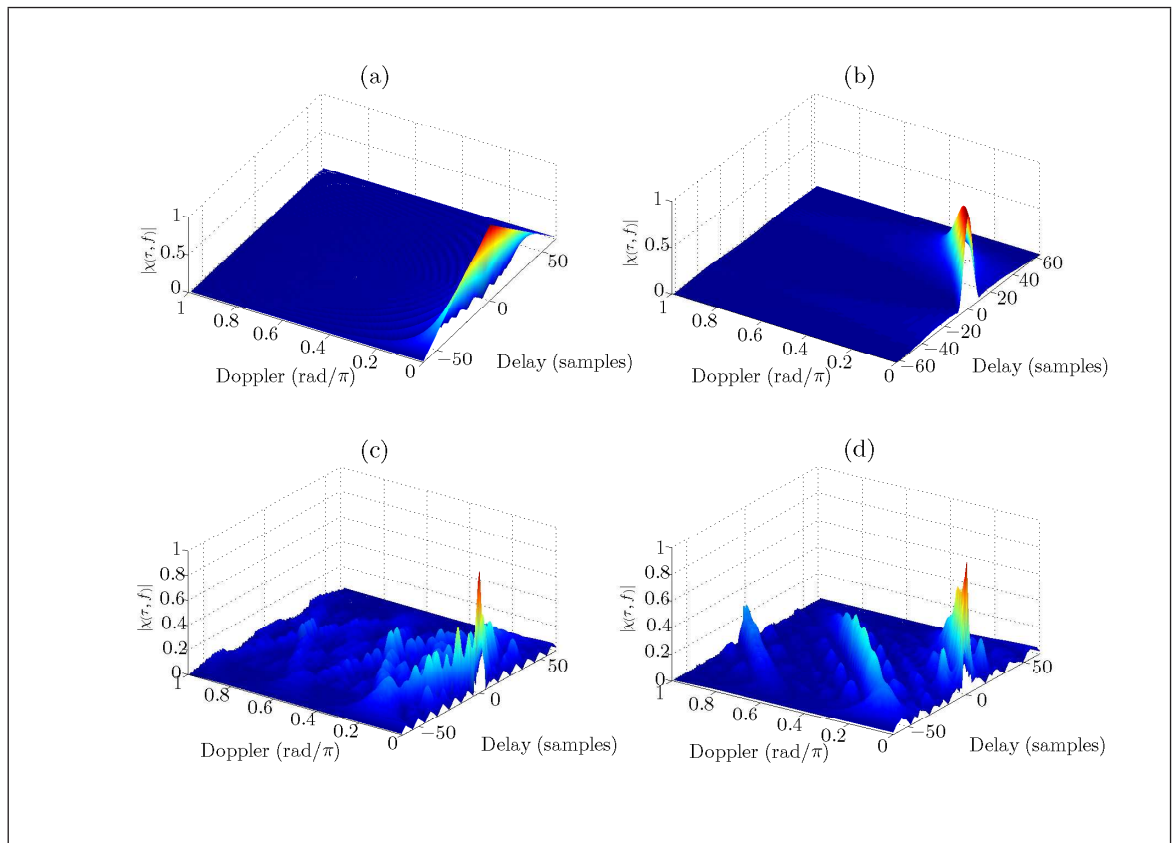


Figure 1.8.1 Ambiguity function magnitude surfaces for: (a) sinusoidal pulse; (b) Hamming-weighted chirp; (c) Barker-13 coded pulse; (d) Frank-coded pulse

1.9 LPI Waveforms

There are several techniques which radars may use to operate with a low probability of intercept. These include:

- bistatic techniques
- minimisation of transmit times
- use of special LPI waveforms.

The third technique of LPI waveforms is the subject of this thesis.

As succinctly described by Stove in [16] and which is the motivation behind the paper in [17], the notion behind the reduction of the probability of detection of a waveform is to spread its energy over as wide a set of dimensions as possible, be it time, frequency or space. To do so in time and frequency tends to result in a high duty cycle modulated waveforms or CW broadband waveforms. Doing so in space results in wider transmitter beams. In chapter 5 the LPI waveforms were modelled such that they were of the type which spread their energy out over the frequency dimension.

To achieve the spreading over many dimensions LPI waveforms may employ:

1. frequency hopping (FH)

2. power management (transmission of the minimum power necessary)
3. pulse compression
4. broader bandwidths for phase-coded pulses or CW waveforms.

Chapter 5 concentrates more on the detection of LPI waveforms which utilise the third and fourth method of generating a LPI waveform from the above list. For the purposes of this thesis, radar waveforms are grouped into three categories, two of which are LPI and the other a conventional:

- chirps
- pseudo-noise (PN) waveforms (phase-coded waveforms)
- sinusoids.

The main difference between the chirps and PN waveforms is that the chirps tend to appear more deterministic in a time-frequency representation (TFR) with a definite structure which changes with time, whereas PN waveforms appear less well defined and are characterised through the occupation of a certain bandwidth. A pictorial example of this is given in Fig 1.9.1 which shows two spectrograms for a chirp category signal and a PN category signal. The spectrograms have a linear rather than logarithmic scale and were produced from parameters which are typical for an ESM system with a superhet and DRx and CW radar waveforms. The spectrograms shown are actually zoom-ins of regions of interest from the whole spectrograms. Panel (a) shows the deterministic nature of an up-ramp for an frequency modulated continuous wave (FMCW) waveform and panel (b) shows the band of energy characteristic of a length $(2^{12} - 1)$ binary m-sequence waveform.

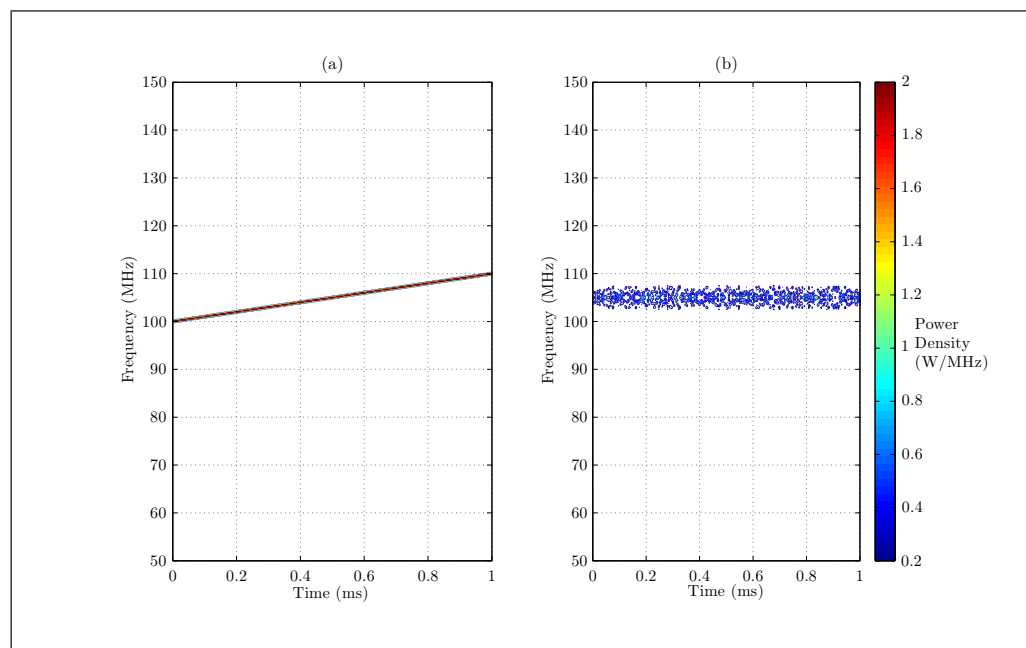


Figure 1.9.1 Spectrograms of a waveform from: 1. chirp category; 2. PN category

There are types of waveforms which do not easily fit into any of the above categories, for example polytime-coded waveforms [18] or hybrid-coded waveforms which combine FH and phase-shift keying (PSK) characteristics [19]. However the most common LPI code schemes are encapsulated by the above groups. The code structure of these categories is achieved through:

Frequency-coding In such waveforms the carrier frequency (f_c) of the waveform is directly changed according to a code scheme. Examples of waveform code schemes which are considered to be members of this category are:

- Pulsed LFM chirps and FMCW. Pulsed LFM chirps tend to be classified by chirp rate (slow, medium and fast). FMCW tactics are described in [15].
- Pulsed non-linear frequency modulation (FM) chirps[20, 21].
- Costas[22] or modified Costas[23] FH waveforms (a member of frequency-shift keying (FSK)).

Phase-coding In such waveforms the phase of the carrier is switched in accordance with a code scheme. These phase changes result in phase transitions in the instantaneous phase of the transmitted waveform. For CW waveforms the individual phase codes form the radar range cells. Examples of waveforms considered to be members of this category are:

- binary phase-shift keying (BPSK)-coded waveforms. For example with the use of Barker codes[24], binary m-sequences, Golomb biphasic codes[25] or Ipatov codes³. As BPSK-coded waveforms involve π phase shifts, the codes are easily applied through sign changes and thus implementation with an inverter in a radar.
- Polyphase-coded waveforms. For example with the use of polyphase Barker codes[26, 27, 28, 29, 30] or Huffman codes[31].
- Polyphase-coded waveforms which are phase-stepped approximations to LFM chirps. These include codes such as Frank[32, 33], P1, P2[34, 35], P3, P4[36], palindromic P4[37], Px[38] (P2 - palindromic P4 examined in [39]), Golomb[40], Zadoff[41] and Zadoff-Chu[42].
- Polyphase-coded waveforms which are phase-stepped approximations to non-linear FM chirps[43].
- Polytime-coded waveforms with the use of coding schemes such as the T1, T2, T3 or T4[18].

³Original paper in Russian but description given in [14]

1.10 LPI Waveform Detection Methods

Detection of LPI waveforms with the use of a DRx which employs the conventional methodology of the application of a threshold to the magnitude-squared of the FFT frequency bins is more difficult than detection of narrowband, sinusoidal waveforms. This is due to the lower peak power and broader bandwidth of LPI waveforms.

There are many proposed techniques to improve the ability of ESM systems to detect LPI radar waveforms. Included in these suggestions are readily understood ideas, such as from [44, 45, 16]. For example Stove suggests the use of correlation techniques or an ESM antenna with a higher directional gain. The link between parameters such as antenna pattern and LPI detection is described in [46]. However antennas with a higher directional gain are of less interest to the tactical problem posed in section 1.2 but they are of more interest in the strategic situation.

In contrast to this there are also algorithmic suggestions from other references. A summary of the algorithmic suggestions, loosely based on [47, 48] includes:

- Detection with the use of adaptive thresholds[49].
- Detection based on the statistics of sampled waveforms.
- Detection with the use of linear transforms.
- Detection with the use of quadratic transforms.
- Detection exploiting the cyclostationarity of signals.

Camuso *et al.* recently investigated selected techniques from the above list and compared their detection performances in [50]. These include wavelet and cyclostationary techniques.

Schemes such as that in [51] are attractive to improve the ability of ESM systems to detect LPI waveforms. In this scheme a simple, fast discriminator was used to discriminate between possible LPI and conventional waveform detections which were sent to a slower algorithm for further analysis and classification. This highlights that ESM systems need to be multi-functional and sensitive to both conventional radar waveforms and LPI radar waveforms. This also implies the requirement of fast algorithms with a circular buffer which occasionally send data sequences to more complicated algorithms for further analysis. The further analysis may involve attempts to recognise the class of waveform modulation with the use of methods such as in [52, 53].

A short description of some of the different detection techniques follows.

1.10.1 Stove's Suggestions

Stove's suggestions of correlation techniques or antennas with higher directional gains from [16] are more conventional solutions to the problem.

Correlation techniques could involve correlation of an input waveform with guesses of possible LPI waveforms from a library. However this would necessitate thousands of cross-correlations to be processed, which is highly impractical.

Correlation techniques could also involve the autocorrelation of an incoming waveform. However this technique has no advantage over energy detection (zero delay in an autocorrelation) for initial detection except when the ESM system has prior knowledge of the cycle frequency of the LPI waveform.

The use of an antenna with a higher directional gain would increase the value of G_{AESM} from Eq (1.2.1) and improve the probability of detection of an LPI waveform in an ESM system. However this is less applicable for the tactical situation posed in section 1.2 because this would be at the expense of lower sensitivity at other off-boresight angles, which is undesirable for an ESM system designed to give good angular coverage from a single antenna.

1.10.2 Detection using Statistics

For a superhet and DRx and for a noise-only input, captured windows of samples tend to be normally-distributed. This is because the noise in a superhet and DRx is usually thermal noise dominated. A typical example of recorded noise from a DRx is shown in Fig 1.10.1.

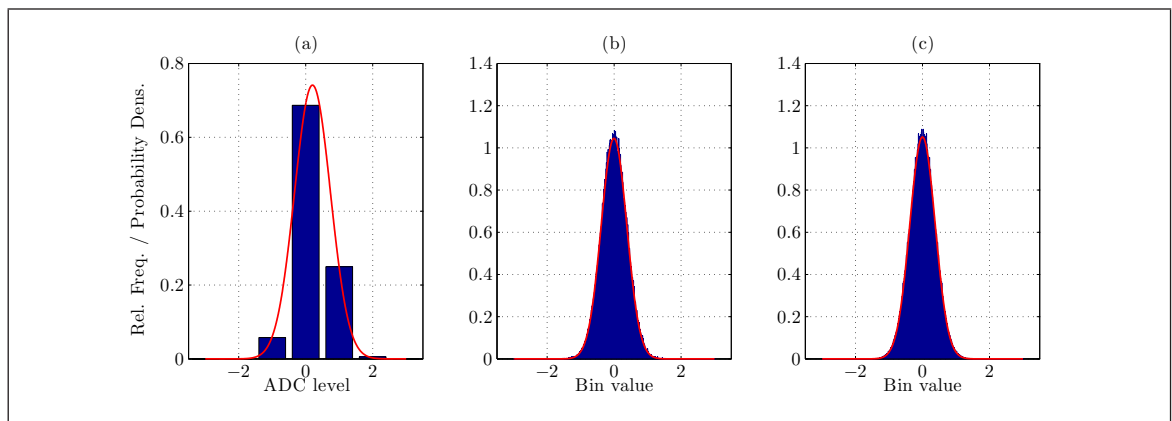


Figure 1.10.1 Typical example distribution of noise from a DRx: (a) time-domain, real channel; (b) real part in frequency-domain; (c) imaginary part in frequency-domain

The noise in Fig 1.10.1 was recorded on trials at Baddow on 03/07/2008 with the use of the SELEX Galileo DRx⁴. The 130938 real-only samples were recorded and panel (a) shows a normalised histogram of the ADC levels and an overlaid normal distribution

⁴The relevant file is `Baddow_03_07_08.Noise-FTDR1-0077.txt`

fit. It is not particularly clear from panel (a) that the noise-only input samples are normally-distributed because the ADC levels are quantised as integers. However upon transformation of the real data to the frequency-domain in panels (b) and (c), it is clear that the real and imaginary parts are approximately normally-distributed (with half the variance in each). The data in (b) and (c) was suitably normalised upon transformation to the frequency-domain and normalised further to produce normalised histograms.

If a radar signal is present in the samples sent to the DRx detection stages, the envelope of the samples tends to be drawn from a Rice distribution⁵. Accordingly the concept of deviation away from normality or “Gaussianity” is exploited in some detection techniques.

Higher-order cumulants and moments are useful to provide some measure of the deviation away from Gaussianity and can be used as decision statistics in detectors, for example [54, 55]. In addition there exists other goodness-of-fit techniques to measure the deviation by comparison to the normal probability density function (PDF) fit to the data[56, 57, 58].

Despite the success of techniques such as these in speech detection, in for example [59], the usefulness of these techniques does not translate well to the detection of LPI waveforms against AWGN. This is because in the case of signals hidden in AWGN, higher-order statistics are less-optimal than the use of a square-law detector, i.e. Eq (1.6.2) with $\nu = 2$ as a decision statistic. This is essentially stated in [60] for Gaussian-distributed signals but the criterion is also stated to be asymptotically true for non-Gaussian signals. Therefore detectors which utilise higher-order statistics are confined to the identification of LPI waveforms[61, 62] and detection amongst coloured noise such as in [63].

1.10.3 Linear Transforms

These transforms are formed from the projection of an input waveform onto a set of “mother” or reference functions. The DFT, which is the staple of most DRxs in ESM systems, falls into this category because it projects onto a set of complex exponentials. The windowed version of this is described in Eq (1.10.1).

$$X(k) = \sum_{n=0}^{N-1} w(n)x(n)W_N^{-kn}, \quad \text{where } k = 0, 1, 2, \dots, N-1 \quad (1.10.1)$$

where $w(n)$ is a window function of length N and $W_N = e^{-\frac{2\pi i}{N}}$.

Linear transforms are related to the square-law detectors because they attempt to

⁵Provided that the statistics of the samples have not been altered too greatly by filters before this point

concentrate signal power into as few coefficients as possible and simple detection is based on the square of these coefficients. This type of detection is described in [64, 65, 66].

Another common transform under this heading is the wavelet transform which projects onto a scaled and translated mother wavelet. However the one-dimensional form of this transform is more commonly found in speech and music analysis, for example [67, 68, 69], rather than the DRxs in ESM systems. This is because the wavelet transform characteristic of a longer duration DCW at higher frequencies and a shorter duration DCW at lower frequencies is less meaningful for radar signals. The discrete wavelet transform (DWT) can be understood as a series of two filters: (1) a low-pass filter (common notation: h , see Eq (1.10.2a)) and (2) a high-pass filter (common notation: g , see Eq (1.10.2b)). These two filters are quadrature mirror filters and are cascaded off the low-pass component at each layer. A downsample occurs at the output of each filter. Thus the high frequency components are analysed with a shorter duration DCW and broader frequency resolution, but the low frequency components are analysed with a longer duration DCW and narrower frequency resolution.

$$y_l(n') = \sum_{n=-\infty}^{\infty} x(n)h(2n' - n) \quad (1.10.2a)$$

$$y_h(n') = \sum_{n=-\infty}^{\infty} x(n)g(2n' - n) \quad (1.10.2b)$$

Nevertheless the DWT is often mentioned in the same field as LPI radar waveforms such as in [70, 71].

A closely-related concept to the DWT is wavelet packet decomposition, which involves the cascade of the quadrature mirror filters off the high frequency components in addition to the low frequency components. In [72, 73], detection and feature extraction of LPI waveforms is based on examination of several layers of successive decomposition.

One transform which is well suited to the chirp category of radar signals is the chirplet transform. An example of its use against FMCW signals is given in [74] where an estimate of the chirp rate is used before projection of the input waveform onto many different guesses of the exact chirp rate.

1.10.4 Quadratic Transforms

These transforms are essentially formed from the comparison of an input waveform with itself. In a general comparison with linear transforms, the quadratic transforms create TFRs with better spectral resolution but they are degraded by cross-term interference between the components at the input to the waveform.

Cohen's class of functions falls under the header of quadratic transforms and perhaps

the two most well-known implementations are the AF and the Wigner Ville distribution (WV , continuous-time formulation shown in Eq (1.10.3)). The former is related to the latter by a two-dimensional Fourier transform (forwards in the direction of the delay dimension and backwards in the direction of the Doppler dimension).

$$WV_{xx}(t, f) = \int_{-\infty}^{\infty} x(t + \tau/2) * x(t - \tau/2) e^{-2\pi i f \tau} d\tau \quad (1.10.3)$$

where $*$ denotes convolution.

A windowed version of Eq (1.10.3) must be used in any implementation in a DRx as part of an ESM system to make it compatible with DCWs created in the DRx. The windowed version of the Wigner Ville distribution is called the pseudo-Wigner Ville distribution.

The Wigner Ville distribution has been explored as an analysis technique for the detection and/or classification of LPI radar waveforms or equivalent in, for example, [75, 76, 77, 78, 79, 80, 81]. However the cross-terms are a very large weakness of the technique and severely limit its usefulness in busy operational EM environments. In addition to this the Wigner Ville distribution, in common with all quadratic transforms, is extremely computationally complex which hampers real-time application in a DRx. With many TFRs it is quite common to require a second transform on the TFR output to successfully detect signals and this would further add to the complexity.

Time and frequency window functions (which have a smoothing filter interpretation) can be used to ease the impact of the inherent cross-term interference, however this comes at the expense of a reduction in time and frequency resolution. For example the Choi Williams distribution is produced from the application of a filter to the cross-terms which exist off the delay-Doppler axes in the AF interpretation. An attempt to quantify the best smoothing filters to use is given in [82]. Alternatively, in [83], the Wigner Ville distribution is combined with the STFT in an attempt to achieve the best of both. Ultimately, however with window functions or combined transforms, the cross-terms can never be removed to the level of the linear transforms.

The use of the Choi Williams distribution as an LPI waveform detection and classification mechanism was explored in [81, 84, 85], however in [84], as with other quadratic transforms, the conclusion was drawn that current technology would not permit it to be calculated in real-time in a feasible DRx.

1.10.5 Cyclic Feature Detectors

Broader-band LPI waveforms, and in general all signals which carry information, contain some modulation scheme. This results in signal adherence to a cyclostationary model, where the statistical parameters vary periodically with time. The square-law

detectors which are essentially used in most ESM DRxs assume a stationary model for both signals and noise. This may approximate in most cases because the modulation period is much longer than a DCW duration, but in certain circumstances detection based on the square-law may become confused[86]. For example the background noise level may change in the course of a DCW or LPI waveforms may be hidden amongst other signals in the same band.

Gardner brought the use of cyclic feature detectors to prominence in [87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97] and more recently there have been attempts to research and implement comprehensive cyclic feature detectors in systems such as in [98, 99]. In particular there are efforts which are aimed towards the integration into ESM systems such as [100, 101, 102, 103, 13] or similar systems [104, 105].

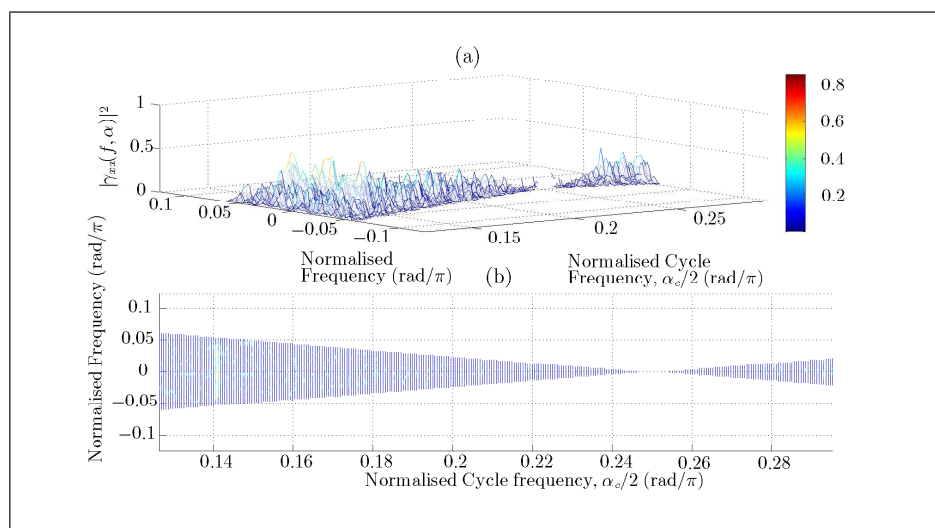


Figure 1.10.2 Cyclic-feature extraction with the aid of a CSC estimate: (a) from the side; (b) from above

Cyclic-feature detectors can be used to detect through the use of the hidden periodicity of LPI waveforms or extract cyclic-features, such as the chip-rate (and its harmonics), to provide information for classification. An investigation was undertaken during this project on cyclic-feature detectors and it was found that for typical, stationary thermal noise-dominated DRxs, cyclic-feature detectors offer no initial detection advantage over conventional (square-law) based techniques. A similar conclusion was found in [106]. However they could prove useful in the classification stage as found in [107]. An example of cyclostationary analysis is shown in Fig 1.10.2. A random BPSK signal with a normalised chip-rate of $f_{\text{chip}} = 0.141 \text{ rad}/\pi$ against a background of AWGN at signal-to-noise ratio $(\text{SNR})_{\text{in}} = 5 \text{ dB}$ was analysed⁶. The analysis involved:

1. initial detection with the use of conventional spectral thresholds;
2. estimation of the target frequency;
3. shift to approximately baseband;

⁶ SNR_{in} refers to the SNR after the ADC in a DRx

4. filter and sub-sampling;
5. cyclostationary analysis.

The cyclostationary technique used was a CSC function estimate which is a type of normalised cyclostationary function. The chip-rate frequency peak is clearly visible in both panels (a) and (b) which are simply different views of the same data.

The simplest cyclic-feature detector is the delay-and-multiply or self-mixer, as shown in Fig 1.10.3. The delay-and-multiply detector is unsuitable for all-purpose detection because it must be designed with delay parameter (τ) and bandpass filter centre frequency (α_c) to match the approximate characteristics of the LPI waveform which the detector wishes to detect. For example if the delay-and-multiply detector was loaded with parameters similar to the periodicity of class A chirps it would be unsuitable for the detection of class B and class C chirps (see Glossary for chirp class description). Therefore the delay-and-multiply detector must be loaded with parameters with help from *a priori* knowledge or the parameters must be scanned methodically.

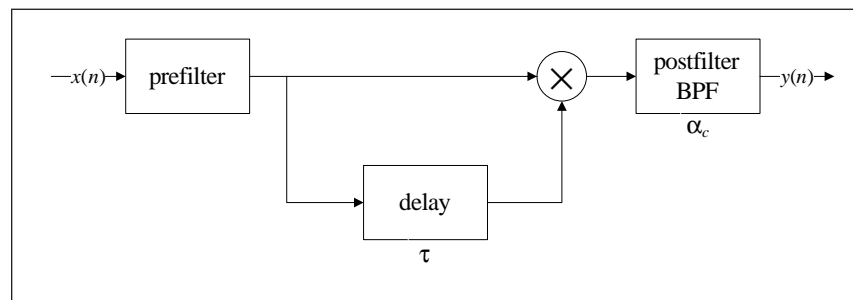


Figure 1.10.3 Self-mixer detector

A more general example is a cyclic spectrum analyser or spectral correlation analyser. The spectral correlation analyser creates a surface over all possible cycle frequencies α_c called a spectral correlation function (SCF). This is equivalent to the two-dimensional Fourier transform of the instantaneous autocorrelation function ($R_{xx}(t, \tau)$). With the use of similar notation to [108], this is given by

$$\mathcal{SC}_{xx}(f, \alpha_c) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_{xx}(t, \tau) e^{-i2\pi\alpha_c t} e^{-i2\pi f \tau} dt d\tau$$

and forms a surface-wise density with units of W/Hz². The cyclic spectrum analyser creates a series of evenly spaced lines (spaced apart by $\Delta\alpha_c$) because it evaluates the cyclic power spectrum (CPS) at a finite number of cycle frequencies. It is the equivalent of the Fourier transform of the cyclic autocorrelation function

$$\mathcal{S}_{xx}(f, \alpha_c) = \int_{-\infty}^{\infty} R_{xx}^{\alpha_c}(\tau) e^{-i2\pi f \tau} d\tau$$

which forms a line-wise power density with units W/Hz.

The cyclic spectrum analyser is applicable to DRxs and can be estimated with extensions to familiar non-parametric spectral estimators. A prime example is the cyclic periodogram and smoothed or averaged versions thereof. A large disadvantage of cyclic spectrum analysers, which they share with the other quadratic transforms is their computational complexity. However there are some techniques which enable computational complexity to be reduced at the expense of resolution. For example in [109] a coarse search for cycle frequencies is initiated before a fine cycle frequency estimation is undertaken on cycle frequencies of interest. Also in [110], two computationally efficient techniques for estimation of the CPS are presented which utilise the FFT algorithm and the topic is discussed further in [111, 112].

The SCF is related to the other quadratic transforms, such as the Wigner Ville distribution and AF, through various forwards and backwards combinations of the two-dimensional Fourier transform. The argument put forward by Gardner in [88] about the SCF and CPS seem true for ESM. That is, cycle frequency parameters about a signal extracted from an SCF or CPS are more useful to an ESM system for exploitation than information from a Wigner Ville distribution. However, like other quadratic transforms, the SCF and CPS suffer from cross-term interference when multiple signals are present. This interference manifests as beat cycle frequencies.

1.11 Summary

Whilst there are many proposed techniques to improve the detection performance of ESM systems against LPI waveforms, many are too computationally complex or expensive for real-time implementation across the entire ESM bandwidth (~ 500 MHz–20 GHz).

For a proposed initial detection algorithm to be easily implemented in a real-time ESM system, it should desirably be FFT-based, offer better performance than a simple energy detector and not compromise detection of conventional waveforms.

Chapters 2–5 discuss FFT-based algorithms which are shown to offer better detection performance than conventional FFT-based DRxs as part of ESM systems.

CHAPTER 2

SVA Theory

2.1 Introduction to SVA

This chapter aims to show that there is a way to avoid the seemingly inevitable window-feature trade-off between NEB, sidelobe level and frequency resolution described in section 1.4. Thereby the performance of DRxs which use a DFT for spectral analysis can be optimised. This is achieved through the use of a technique from the domain of synthetic aperture radar (SAR) and inverse synthetic aperture radar (ISAR) imagery called SVA and its application in the domain of EW.

SVA is effectively an adaptive window technique. Some important properties of SVA are: its simplicity; compatibility with a DFT-based DRx and the fact it requires no *a priori* information about the input waveform.

One interpretation of the SVA algorithm is as a special case of the minimum variance spectral estimation (MVSE) algorithm, first introduced by Capon. However SVA lacks the matrix inversion and therefore the associated computational complexity. SVA is constrained to window functions from the raised-cosine family[113].

As shown in Fig 1.4.2, it is conventional to apply a window function (data taper) to the DCW in the time-domain. The key concept behind SVA is transplanting the time-domain weight process to the frequency-domain and restricting the window functions to those from the raised-cosine family.

The next section describes the shapes and properties of raised-cosine window functions and the following sections describe how SVA applies them in the frequency-domain.

2.2 Raised-cosine Windows

Members of the raised-cosine window family derive from the generalised Hamming family of windows. The following is a commonly-quoted definition of the generalised Hamming family of windows of length N

$$w(n) = a_1 - 2a_2 \cos\left(\frac{2\pi n}{N}\right)$$

With the a_1 parameter set to the constant 1 and the a_2 parameter as a variable, a

definition for the conventional set of (first-order) raised-cosine windows of length N is obtained

$$w(n) = 1 - 2\alpha \cos\left(\frac{2\pi n}{N}\right) \quad (2.2.1)$$

where α is restricted to values $\{\alpha \in \mathbb{R} \mid 0 \leq \alpha \leq 1/2\}$ in order to keep the window always positive-valued in the time-domain. A selection of raised-cosine windows in the time-domain with $N = 128$ data points with varying α parameters are plotted in panel (a) of Fig 2.2.1.

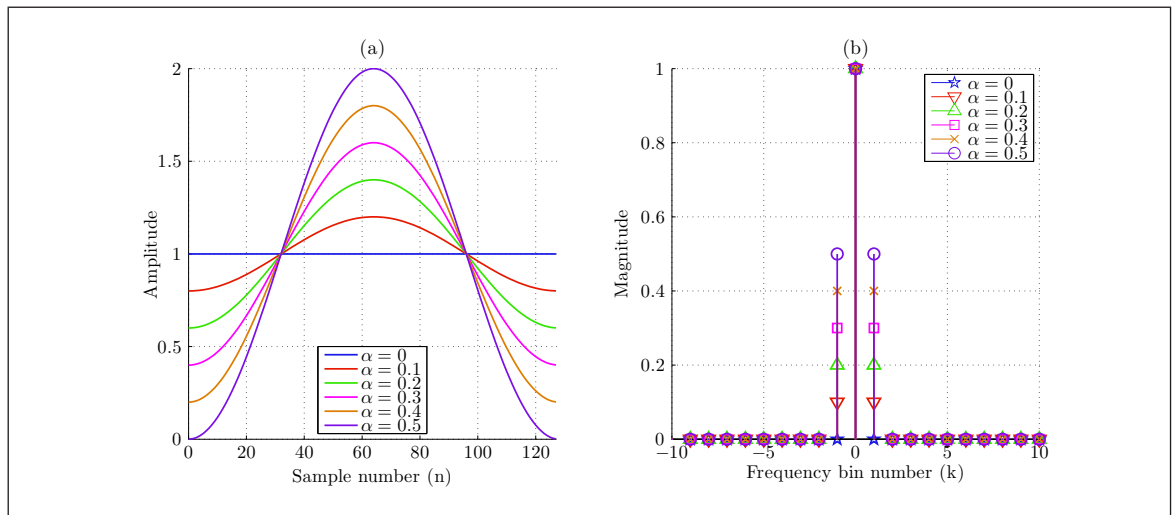


Figure 2.2.1 Some members from the raised-cosine window family with different α parameters: (a) in the time-domain; (b) in the frequency-domain (negative bin numbers same as positive bin numbers above Nyquist frequency)

At the $\alpha = 0$ limit, a uniform, rectangular window function is obtained, whilst at the $\alpha = 1/2$ limit a Hann window function is obtained.

Window functions such as those in panel (a) of Fig 2.2.1 are applied through multiplication of the N time-domain data samples in a DCW with the corresponding window weight in the time-domain. This is followed by a transformation of the data with a DFT

$$X(k) = \sum_{n=0}^{N-1} w(n)x(n)W_N^{-nk} \quad k = 0, 1, \dots, N-1 \quad (2.2.2)$$

The trade-offs mentioned in section 1.4 involved in the choice of window function in conventional spectral analysis are well-known and shown qualitatively in Fig 2.2.2. Fig 2.2.2 explains why in EW applications a moderate resolution and moderate dynamic range window function such as a Hann, Hamming or Dolph-Chebyshev window function is used: they are a safe option when it is uncertain what signals are in the operational environment.

Inevitably in conventional spectral analysis some resolution and sensitivity will be lost when, say, a Hann window function is chosen over the rectangular window, but it will have some positive benefits in dynamic range over, say, over a Blackman-Harris

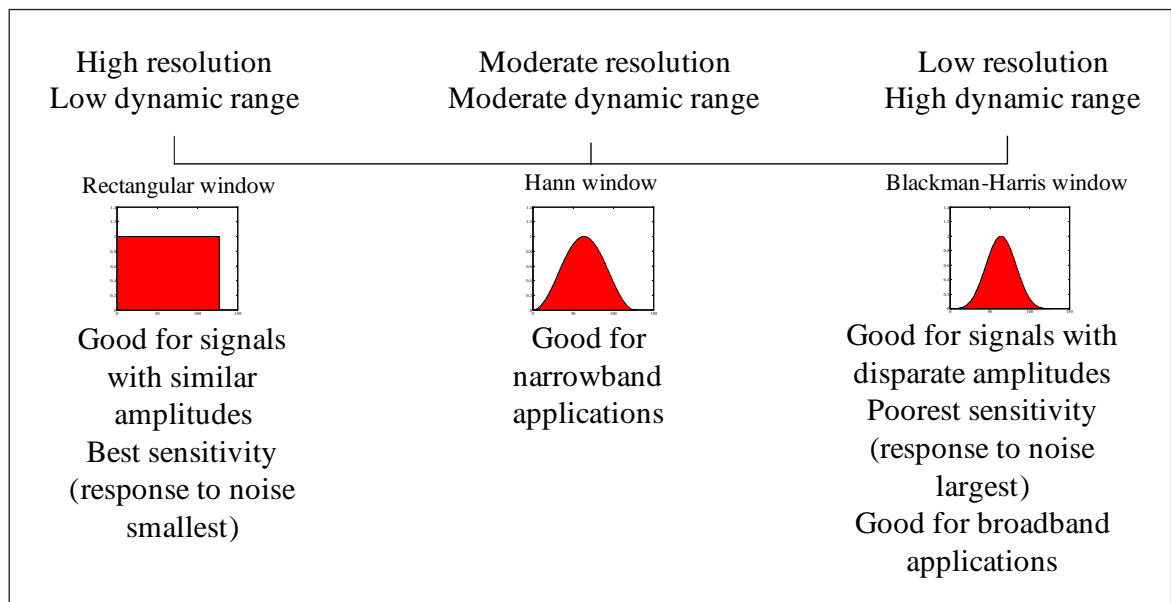


Figure 2.2.2 Window characteristics are all fixed once the choice of window function is made

window. If greater resolution or sensitivity is desired the current best solution is to choose a different window function for the next DCW.

However in the next section it will be shown that the raised-cosine window functions have an interesting frequency-domain response which will allow the window function trade-off to be mitigated somewhat.

2.3 DFT of Raised-cosine Windows

According to the convolution theorem it is entirely equivalent to perform the window function multiplication in the time-domain as a convolution in the frequency-domain. Specifically, instead of the multiplication of N window function weights with the data samples in the time-domain, the DFT of the window function can be convolved with the DFT of the data samples. Superficially this seems a retrograde idea because extra computational work is required as convolution is generally more computationally intensive than multiplication. However the $N_{\text{FFT}} = N$ DFT of the generic raised-cosine window function from Eq (2.2.1) is a simple 3-point frequency-domain response

$$W(k) = \delta(k) - \alpha\delta(k+1) - \alpha\delta(k-1) \quad (2.3.1)$$

Variation of α from Eq (2.3.1) between its limits (the shape of the time-domain window function) increases the magnitude of the Kronecker deltas one Nyquist-sampled frequency bin to the right and left of the central unit impulse Kronecker delta. This effect is illustrated in panel (b) of Fig 2.2.1.

The 3-point frequency-domain response when convolved with the unwindowed DFT of

a DCW is

$$X_a(k) = X(k) - \alpha X(k+1) - \alpha X(k-1) \quad (2.3.2)$$

There are two significant advantages from the application of a window function in the frequency-domain via convolution:

1. The unwindowed DFT data $X(k)$ can be made available for other processing techniques. This is useful for other processes which require an orthogonal transform of the data.
2. The window function applied at one frequency bin can be different from adjacent frequency bins, in effect they can each have their own corresponding time-domain window function. In the case of raised-cosine window functions, this idea can be simplified down to the adaptive selection of α parameters at each frequency bin. This concept is central to the SVA algorithm.

2.4 Spatially Variant Apodization (SVA)

The second point at the end of the last section is the basic SVA adaptive window concept. One of the simplest interpretations from [114] of the mechanism for selection of α parameters is achieved through consideration of **dual apodization**. This is a crude form of SVA, whereby two different α values (e.g. $\alpha = 0$ and $\alpha = 0.5$) are used in Eq (2.2.1) to create two different window functions which are then applied in Eq (2.2.2). The minimum magnitude-squared result at each frequency bin is selected from the two spectral estimates. This process is highlighted in Fig 2.4.1.

Naturally, the logical extension of dual-apodization is to select the minimum magnitude-squared value for each frequency bin from multiple spectral estimates in a process called **multiapodization** in [114, 115].

The calculation of many spectral estimates only to subsequently discard most of the results has a significantly negative impact on computational efficiency. A more rigorous process is to search for the optimal window function at each frequency bin. This search is undertaken in SVA through optimisation of the α parameter from Eq (2.3.2).

There are two established methodologies in SVA for selection of an optimal α parameter for the in-phase (i.e. real part $X_R(k)$) and quadrature (i.e. imaginary part $X_I(k)$) components of the complex output of a DFT. These are described in [114]:

1. search by the **joint** optimisation of $X_R(k)$ and $X_I(k)$;
2. search by the **separate** optimisation of $X_R(k)$ and $X_I(k)$.

In the first procedure, the cost function which is minimised in order to choose the optimal value for α is $|X_a(k)|^2$. This quantity is recognisable from section 1.6 as the instantaneous (spectral) energy density at the k^{th} frequency bin. The minimisation

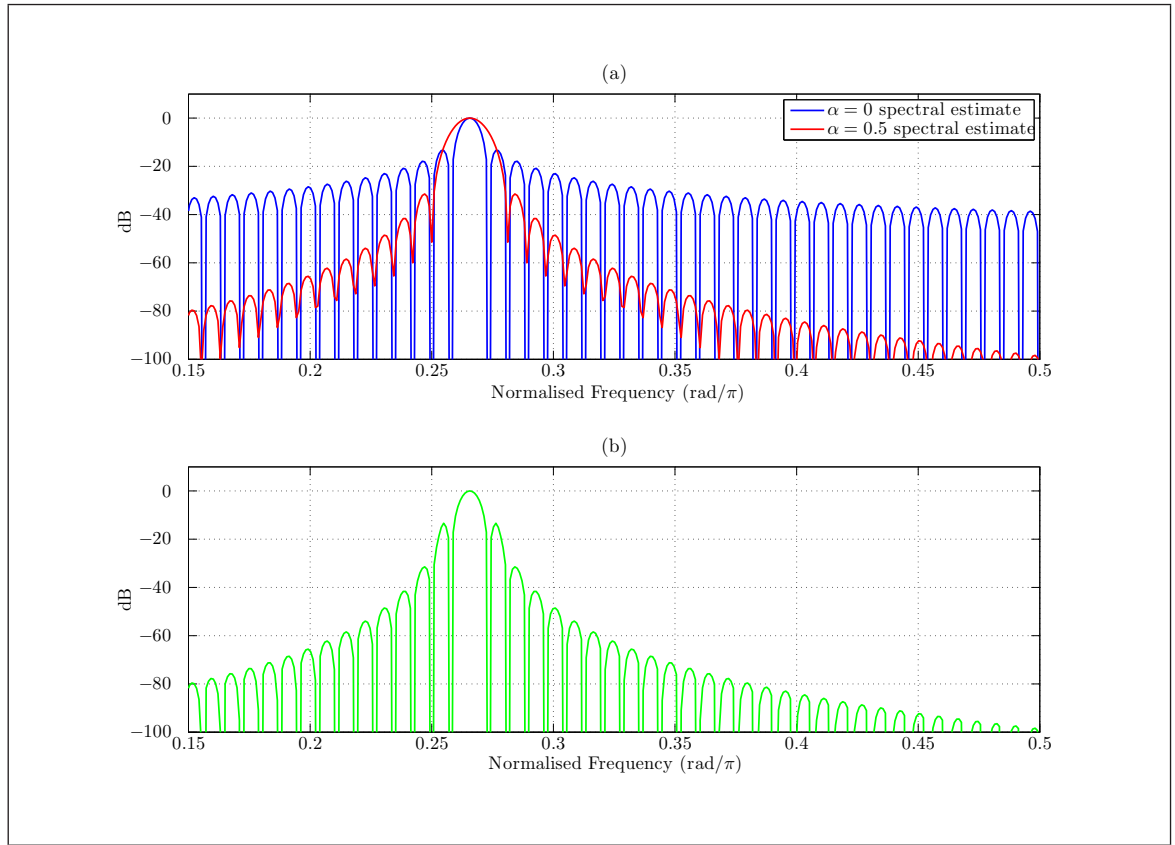


Figure 2.4.1 Example of dual-apodization: (a) two spectral estimates produced from same DCW but different window functions; (b) dual-apodized spectral estimate result

is solved by the partial derivative of $|X_a(k)|^2$ with respect to α equal to zero and rearranged for α to yield the solution

$$\alpha = \Re \left\{ \frac{X(k)}{X(k+1) + X(k-1)} \right\} \quad (2.4.1)$$

Equation (2.4.1) is subject to the constraint $\{\alpha \in \mathbb{R} \mid 0 \leq \alpha \leq 1/2\}$. Therefore the combination of Eq (2.3.2) and Eq (2.4.1) yields

$$X_a(k) = \begin{cases} X(k), & \text{if } \{\alpha < 0\}; \\ X(k) - \alpha X(k+1) - \alpha X(k-1), & \text{if } \{0 \leq \alpha \leq 1/2\}; \\ X(k) - 1/2 X(k+1) - 1/2 X(k-1), & \text{if } \{\alpha > 1/2\}. \end{cases} \quad (2.4.2)$$

The $X_a(k)$ vector is the filtered DFT after joint-SVA has been applied and thus a spectral estimate can be obtained from $|X_a(k)|^2$. In effect, the smallest magnitude-squared response in the frequency-domain is chosen for the k^{th} frequency bin through minimisation of the cost function $|X_a(k)|^2$. This has the effect of simultaneous minimisation of the sidelobe levels and mainlobe width of any signals present in the spectral estimate. A similar result would be obtained by multiapodization. Multiapodization, however would involve the application of many possible window functions from the set

described by Eq (2.2.1) to the same DCW in the time-domain; transformation with as many DFT operations and selection of the minimum magnitude-squared result at each frequency bin. The SVA approach is far less computationally expensive.

The other SVA technique described in [114] is the separate- $I&Q$ SVA technique. It follows a similar mathematical derivation, however, the cost function is separated into its real ($X_{R,a}(k)$) and imaginary ($X_{I,a}(k)$) parts before the magnitude-squared of each is minimised with respect to an α parameter. The unconstrained solutions for the α values are

$$\alpha_R = \frac{X_R(k)}{X_R(k+1) + X_R(k-1)} \quad (2.4.3)$$

$$\alpha_I = \frac{X_I(k)}{X_I(k+1) + X_I(k-1)} \quad (2.4.4)$$

The similarity with Eq (2.4.1) is clear. These α parameters are also constrained to $\{\alpha_R, \alpha_I \in \mathbb{R} \mid 0 \leq (\alpha_R, \alpha_I) \leq 1/2\}$. The real and imaginary parts of $X_a(k)$ show filtered DFTs according to the following rules

$$X_{R,a}(k) = \begin{cases} X_R(k), & \text{if } \{\alpha_R < 0\}; \\ 0, & \text{if } \{0 \leq \alpha_R \leq 1/2\}; \\ X_R(k) - 1/2X_R(k+1) - 1/2X_R(k-1), & \text{if } \{\alpha_R > 1/2\}. \end{cases} \quad (2.4.5)$$

$$X_{I,a}(k) = \begin{cases} X_I(k), & \text{if } \{\alpha_I < 0\}; \\ 0, & \text{if } \{0 \leq \alpha_I \leq 1/2\}; \\ X_I(k) - 1/2X_I(k+1) - 1/2X_I(k-1), & \text{if } \{\alpha_I > 1/2\}. \end{cases} \quad (2.4.6)$$

In [115], the separate- $I&Q$ SVA technique is rigorously examined. The thorough treatment includes a derivation of the PDF of the apodized $I&Q$ channels and the statistical moments of generic order. The PDF of either the $X_R(k)$ or $X_I(k)$ channel is derived as

$$\begin{aligned} f(X; \sigma) = & \frac{1}{2\sigma} \frac{1}{\sqrt{2\pi}} e^{-\frac{X^2}{2\sigma^2}} \\ & + \frac{1}{2\sigma} \frac{1}{\sqrt{3\pi}} e^{-\frac{X^2}{3\sigma^2}} \\ & \times \left[1 - \operatorname{erf} \left(\frac{|X|}{\sqrt{6}\sigma} \right) \right] \\ & + \frac{1}{\pi} \left(\frac{\pi}{2} - \arctan(\sqrt{2}) \right) \delta(X) \end{aligned}$$

It is noted in [115] that the separate-SVA technique is more “aggressive” in terms of the minimisation of sidelobes of possible signals as it assigns values drawn from a normal population to the value zero approximately 20% of the time, whereas the joint-SVA

technique is more likely to better preserve the statistics of the data. Therefore separate-SVA deforms the data PDF badly but represents the best sidelobe level reduction and joint-SVA is a compromise between PDF deformation and sidelobe level.

Examples of the use of the separate-SVA technique, joint-SVA technique, Hann window function and rectangular window function when a single tone is present throughout the entire DCW in the absence of noise are shown in Fig 2.4.2. Panel (a) shows the behaviour of the α parameter of joint-SVA for the spectral estimate in panel (b). Panel (b) shows the case without noise present whilst panel (c) shows the case with noise present. In panel (b) the mainlobe of the rectangular-windowed, joint-SVA and separate-SVA sit on top of each other. The “setting values to zero” effect of separate-SVA is apparent in the both plots as absent data points since the logarithm of zero is undefined. In both plots the separate-SVA sidelobe level oscillates between zero and the sidelobe level of the joint-SVA technique. Therefore for applications in DRxs in ESM systems the use of separate-SVA instead of joint-SVA would offer no advantage because the frequency bin location of the zeros is unknown *a priori* and so the sidelobe level of the joint-SVA must be assumed. Overall the use of the separate-SVA technique in a DRx would be to deform the data statistics badly for no extra gain. The simple nature of separate-SVA (even simpler than joint-SVA) makes it attractive as a fast technique to use in real-time on SAR data, but Chapter 3 studies only joint-SVA because it has better properties and conceivably should be possible to run in real-time in an ESM DRx.

The alternative interpretation of SVA to the multiapodization explanation was briefly mentioned in section 2.1. SVA can be shown to be a special case of MVSE. Indeed, MVSE is an adaptive window technique that utilises a window function which changes shape in accordance with the frequency of interest and the second-order statistics of the signal. Therefore MVSE provides some control over sidelobe levels. To obtain MVSE estimates, however, requires an inversion of the covariance matrix of the observations. This can be computationally intensive for large numbers of data samples. The SVA technique strips down MVSE complexity through its restriction of the window function to the raised-cosine window family from Eq (2.2.1) and assumes a simplistic estimate for the covariance matrix of the data samples. The SVA technique is simple enough that it could be implemented as a series of N adaptive 3-tap finite impulse response (FIR) filters in the frequency-domain. The central Kronecker delta in Eq (2.3.1) keeps the gain of each filter at unity for the frequency of interest, which reflects the distortionless requirement in MVSE.

The SVA techniques outlined so far are for the one-dimensional case. They can be applied to a single DFT output at a time and can be used along the frequency columns of an STFT as demonstrated in [116, 117].

As the SVA technique was primarily invented for use in SAR and ISAR imagery it

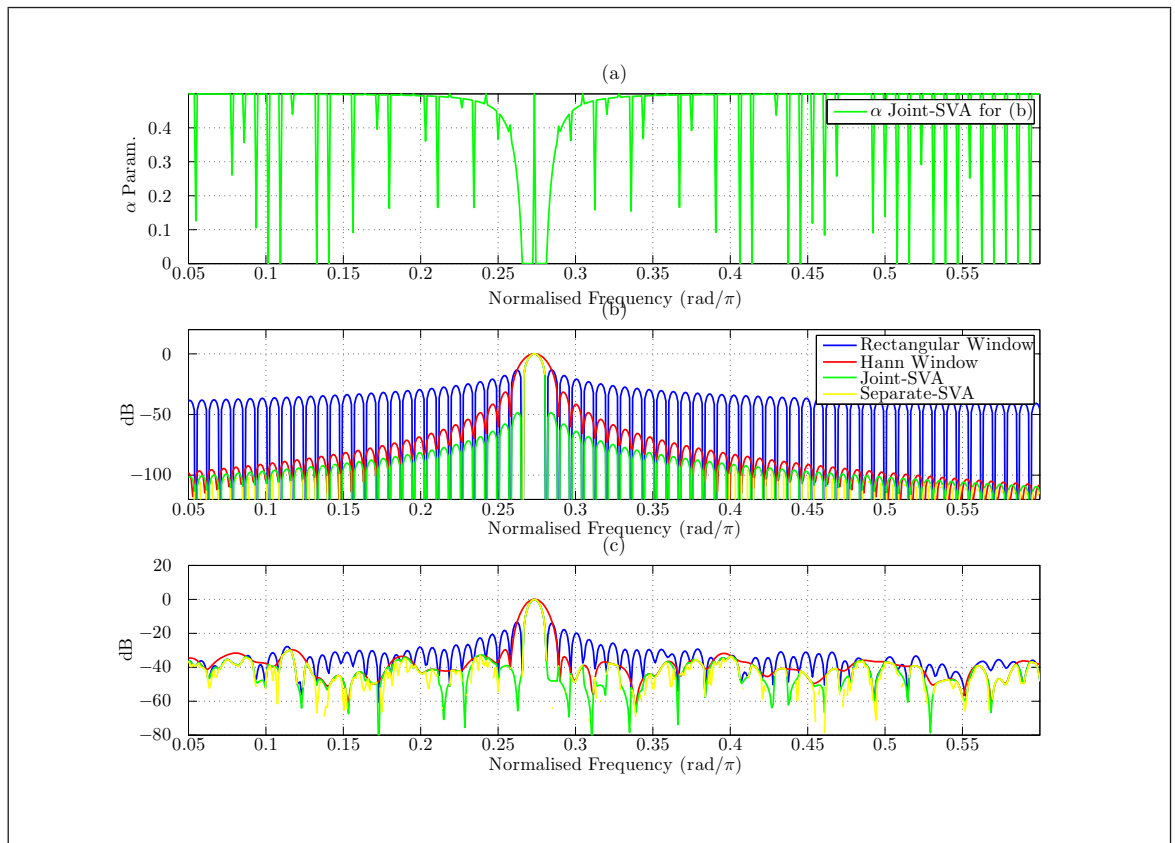


Figure 2.4.2 Comparisons of separate-SVA, joint-SVA, Hann-windowed DFT and rectangular-windowed DFT of: (a) a single tone; (b) a single tone plus AWGN at $\text{SNR}_{\text{in}} = 20$ dB

was originally extended for use in two dimensions. The mathematical steps for this are explained in [114], however in the context of a DRx in an ESM system based around an FFT processor, the two-dimensional formulation of SVA is not applicable.

As a point of interest the separate-SVA technique from Eqs (2.4.5) and (2.4.6) is patented for use in imaging systems in [118].

There are variations on the basic SVA concept, which are detailed in the following sections although they are disregarded for further study.

2.5 5-point and 7-point SVA

In Eq (2.3.1) there are three Kronecker delta points in the convolution formula, which as previously mentioned, are important since they create a very simple circular convolution in the frequency-domain. The relationship with the set of raised-cosine windows in the time-domain from Eq (2.2.1) however means that the maximum sidelobe roll-off of a signal peak in the frequency-domain is approximately limited to the sidelobe roll-off derived from the transform of the limiting case of a Hann window function at $\alpha = 0.5$.

A greater sidelobe roll-off rate can be achieved by extending the convolution formula to

include more Kronecker delta functions. This is done in [119] and it is described here. Using five points in the convolution formula is the equivalent to the use of window functions from the set of order-2 raised-cosine windows. Similarly, seven points in the convolution formula corresponds to the use of window functions from the set of order-3 raised-cosine windows. The general equation for the set of (unnormalized) raised-cosine windows of any order is

$$w_K(n) = 1 + \sum_{m=1}^K a_m (-1)^m \cos\left(\frac{2\pi mn}{N}\right) \quad (2.5.1)$$

To remove discontinuities at the DCW boundaries of the raised-cosine window functions, the derivatives of Eq (2.5.1) are desired to be zero. This is because the discontinuities determine the sidelobe roll-off rate: the fewer discontinuities, the greater the roll-off rate. Therefore for higher order raised-cosine windows, a greater roll-off rate is possible. All the odd-order derivatives of Eq (2.5.1) are zero at the boundaries because $\sin(m2\pi) = 0$ for all integer values of m . The even-order derivatives of Eq (2.5.1) with respect to n are the source of continuity equations which are described later.

In [119] the 5-point and 7-point convolution formulas are explicitly calculated. Firstly, the 5-point SVA technique is outlined. The derivation starts by substitution of $K = 2$ into Eq (2.5.1) to obtain an order-2 raised-cosine window function of the form

$$w_2(n) = 1 - \alpha_1 \cos\left(\frac{2\pi n}{N}\right) + \alpha_2 \cos\left(\frac{4\pi n}{N}\right) \quad (2.5.2)$$

Equation (2.5.2) can be specified in terms of one unknown α parameter upon examination of the zeroth derivative of Eq (2.5.1) with $K = 2$ (the order of the raised-cosine window) and the requirement that it sum to zero at the edges (where $\{n = 0, N\}$). This gives the continuity equation

$$1 - \alpha_1 + \alpha_2 = 0$$

which after dropping the subscript on the α parameter leads to

$$w_2(n) = 1 - \alpha \cos\left(\frac{2\pi n}{N}\right) + (\alpha - 1) \cos\left(\frac{4\pi n}{N}\right) \quad (2.5.3)$$

To find the convolution formula, the DFT of Eq (2.5.3) is taken to obtain the result

$$\begin{aligned} W_2(k) &= \delta_k - \alpha/2\delta(k-1) - \alpha/2\delta(k+1) + (\alpha-1)/2\delta(k-2) \\ &\quad + (\alpha-1)/2\delta(k+2) \\ &= (\delta(k) - \delta_{22}(k)/2) - \alpha/2(\delta_{11}(k) - \delta_{22}(k)) \end{aligned} \quad (2.5.4)$$

where Eq (2.5.4) has been simplified by grouping terms together to create the variables

$\delta_{11}(k) = \delta(k-1) + \delta(k+1)$ and $\delta_{22}(k) = \delta(k-2) + \delta(k+2)$. The convolution of Eq (2.5.4) with a DFT output in the frequency-domain, gives

$$X_a(k) = (X(k) - X_{22}(k)/2) - \alpha/2 (X_{11}(k) - X_{22}(k)), \quad (2.5.5)$$

where $X_{11}(k) = X(k-1) + X(k+1)$
and $X_{22}(k) = X(k-2) + X(k+2)$

In the same situation as the 3-point SVA techniques, there is one α parameter in Eq (2.5.5) for which an optimal value is chosen at each frequency bin in the frequency-domain. An analogous procedure is followed to the 3-point joint-SVA procedure where the cost function to be minimised is $|X_a(k)|^2$. The partial derivative of $|X_a(k)|^2$ with respect to α is calculated and the result set equal to zero. This yields the optimal values for α

$$\alpha = \Re \left\{ \frac{2X(k) - X_{22}(k)}{X_{11}(k) - X_{22}(k)} \right\} \quad (2.5.6)$$

The limits to be placed on α , in order to keep the time-domain order-2 raised-cosine windows always positive, are $\{\alpha \in \mathbb{R} \mid 0 \leq \alpha \leq 4/3\}$. Therefore the DFT in Eq (2.5.5) is

$$X_a(k) = \begin{cases} X(k) - X_{22}(k)/2, & \text{if } \{\alpha < 0\}; \\ X(k) - \alpha/2 X_{11}(k) + (\alpha-1)/2 X_{22}(k), & \text{if } \{0 \leq \alpha \leq 4/3\}; \\ X(k) - 2/3 X_{11}(k) + 1/6 X_{22}(k), & \text{if } \{\alpha > 4/3\}. \end{cases}$$

The start of the derivation for 7-point SVA technique is obtained upon substitution of $K = 3$ into Eq (2.5.1). This gives the order-3 raised-cosine window function formula

$$w_3(n) = 1 - \alpha_1 \cos\left(\frac{2\pi n}{N}\right) + \alpha_2 \cos\left(\frac{4\pi n}{N}\right) - \alpha_3 \cos\left(\frac{6\pi n}{N}\right) \quad (2.5.7)$$

Equation (2.5.7) can be expressed in terms of one unknown α parameter. This can be done with the help of two continuity equations. The first, Eq (2.5.8a), is obtained through the requirement of the zeroth derivative of Eq 2.5.1 with respect to n and with $K = 3$ (the order of the raised-cosine window) to sum to zero. The second, Eq (2.5.8b), is obtained through the requirement of the second derivative of (2.5.1) with respect to n and with $K = 3$ to sum to zero

$$1 - \alpha_1 + \alpha_2 - \alpha_3 = 0 \quad (2.5.8a)$$

$$-\alpha_1 + 4\alpha_2 - 9\alpha_3 = 0 \quad (2.5.8b)$$

When the unnecessary subscript on α is dropped and Eqs (2.5.7), (2.5.8a) and (2.5.8b)

are combined, the order-3 raised-cosine window function formula becomes

$$w_3(n) = \delta(k) - \alpha \cos\left(\frac{2\pi n}{N}\right) + (1.6\alpha - 1.8) \cos\left(\frac{4\pi n}{N}\right) - (0.6\alpha - 0.8) \cos\left(\frac{6\pi n}{N}\right) \quad (2.5.9)$$

Upon application of the DFT to Eq (2.5.9) and grouping terms (with the use of $\delta_{33}(k) = \delta(k-3) + \delta(k+3)$ and the same $\delta_{11}(k)$ and $\delta_{22}(k)$ formulas as in 5-point SVA) produces the 7-point convolution formula

$$W_3(k) = \delta(k) - \alpha/2\delta_{11}(k) + (0.8\alpha - 0.9)\delta_{22}(k) - (0.3\alpha - 0.4)\delta_{33}(k) \quad (2.5.10)$$

Convolution of Eq (2.5.10) with the DFT of a DCW creates a filtered DFT with the formula

$$X_a(k) = X(k) - \alpha/2X_{11}(k) + (0.8\alpha - 0.9)X_{22}(k) - (0.3\alpha - 0.4)X_{33}(k) \quad (2.5.11)$$

where $X_{33}(k) = X(k-3) + X(k+3)$. The same cost-function-minimisation procedure to find the optimal value for α at each frequency bin gives

$$\alpha = \Re \left\{ \frac{X(k) - 0.9X_{22}(k) + 0.4X_{33}(k)}{0.5X_{11}(k) - 0.8X_{22}(k) + 0.3X_{33}(k)} \right\}$$

The restrictions on α , in order to keep the time-domain order-3 raised-cosine window functions always positive, are $\{\alpha \in \mathbb{R} \mid 0.5 \leq \alpha \leq 1.5\}$. Therefore the overall filtered DFT from Eq (2.5.11) is

$$X_a(k) = \begin{cases} X(k) - 0.25X_{11}(k) - 0.5X_{22}(k) - 0.25X_{33}(k), & \text{if } \{\alpha < 0.5\}; \\ X(k) - 0.5\alpha X_{11}(k) + (0.8\alpha - 0.9)X_{22}(k) - (0.3\alpha - 0.4)X_{33}(k), & \text{if } \{0.5 \leq \alpha \leq 1.5\}; \\ X(k) - 0.75X_{11}(k) + 0.3X_{22}(k) - 0.05X_{33}(k), & \text{if } \{\alpha > 1.5\}. \end{cases}$$

The next section explores whether it is worthwhile to attempt to use the 5-point or 7-point SVA techniques over the joint 3-point SVA technique and conventional windowing as an algorithm in a DRx as part of an ESM system.

2.5.1 Evaluation of 5-point and 7-point SVA

To begin to draw comparisons between the different windowing techniques, some examples of spectral estimates are shown in Figs 2.5.1–2.5.2. These were produced by the:

- rectangular window
- Hann window
- (joint) 3-point SVA
- 5-point SVA
- 7-point SVA.

To produce Fig 2.5.1, a single tone, present throughout the DCW, was processed with the different window techniques. No noise was present in Fig 2.5.1. The spectral estimates are all normalised such that their peak frequency-domain responses rest on 0 dB. This was done to allow for an easy comparison of the spectral leakage patterns.

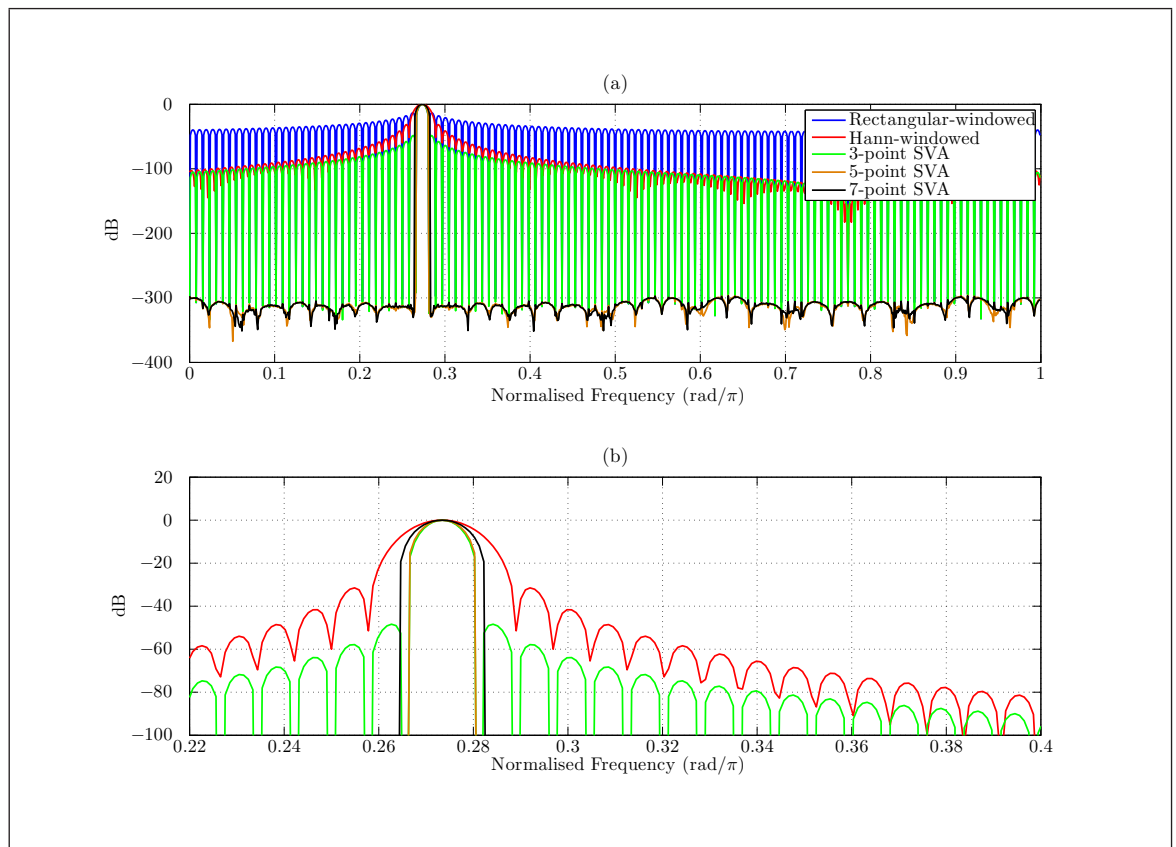


Figure 2.5.1 (Joint) 3-point SVA, 5-point SVA, 7-point SVA, Hann-windowed and rectangular windowed spectral estimates: (a) full view; (b) zoom-in on mainlobe

Panel (a) shows the 5-point and 7-point SVA window techniques are excellent at reducing sidelobe levels (the decibel level is close to machine numerical precision). Panel (b) is a zoom-in on the mainlobe in panel (a) without the rectangular window plot for clarity (its mainlobe sits on top of the (joint) 3-point SVA mainlobe). Panel (b) shows the expected result that the mainlobes are broader for the 5-point and 7-point SVA spectral estimates than for the 3-point SVA spectral estimates (although only marginally so in the case of the 5-point SVA). This is because the two (in the case of 5-point SVA) and three parameters (in the case of 7-point SVA) were combined into a single α parameter in the formulations and therefore there are additional constraints

on the window function shape. In other words the rectangular window function is no longer an option for the 5-point and 7-point SVA. What Fig 2.5.1 does not show because it has been normalised is that 5-point and 7-point will also not recover the window loss due to the lack of rectangular window function choice. This is a strong indication that it would be detrimental to use 5-point and 7-point SVA over 3-point SVA in an ESM system DRx.

Fig 2.5.2 examines the spectral estimates produced when the same signal is processed with AWGN also present in the DCW at $\text{SNR}_{\text{in}} = 20 \text{ dB}$. This scenario with noise present was more realistic and the results reveal the noise floor of 5-point and 7-point SVA is exactly the same as for the 3-point SVA and Hann-windowed DFT (the rectangular-windowed FFT also has the same noise floor but the rectangular-windowed spectral estimate is dominated by sidelobes close to the mainlobe). Again, the mainlobe of the rectangular-windowed and 3-point SVA sit on top of each other.

The SVA techniques are sidelobe-reducing algorithms not noise floor reducing algorithms, as implied in [117], therefore in noise-dominated regimes there would be no advantage in the use of 5-point or 7-point SVA over the 3-point SVA. This is especially true as the 5-point and 7-point SVA involve extra computation.

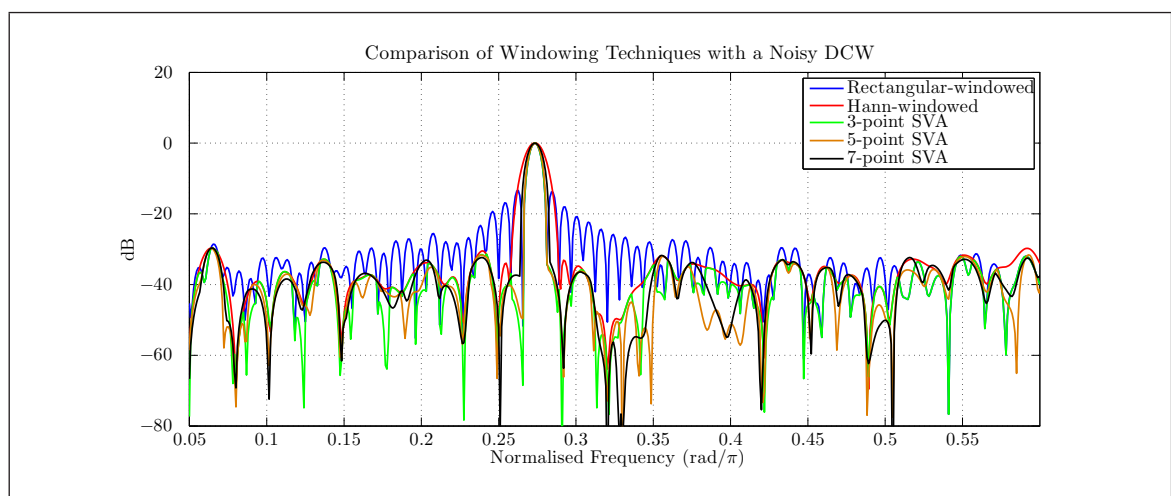


Figure 2.5.2 (Joint) 3-point SVA, 5-point SVA, 7-point SVA, Hann-windowed and rectangular windowed spectral estimates for signal plus AWGN

Extra arguments against the use of 5-point and 7-point SVA over 3-point SVA in an ESM system DRx are provided in Figs 3.2.7 and 3.4.13 and explained in sections 3.2 and 3.4.

2.6 SVA using the Kaiser Window

The previous section detailed a variation of the basic 3-point SVA with the use of higher order raised-cosine window functions. There are also varieties of SVA which do not rely solely on the use of window functions drawn from the family of raised-cosines. Again the idea of SVA as a special case of the MVSE method is used. When

other weighted window functions are used for the estimation of the sample covariance function, different results can be obtained.

According to [120, 121], where this formulation is derived, the use of the Kaiser window in SVA provides better spectral estimate results in comparison to the set of raised-cosine window functions. However the enormous disadvantage in departing from raised-cosine windows is the loss of the very simple 3-point convolution formula of Eq (2.3.1). This contrasts with the $N_{\text{FFT}} = N$ frequency-domain response of the Kaiser window function

$$W(k) = \frac{N}{I_0(\alpha_K \pi)} \frac{\sin \left(\sqrt{\alpha_K^2 \pi^2 - \left(\frac{Nk}{2} \right)^2} \right)}{\sqrt{\alpha_K^2 \pi^2 - \left(\frac{Nk}{2} \right)^2}}, \quad (2.6.1)$$

for $k = -\pi, \dots, -2\pi/N, 0, 2\pi/N, \dots, \pi$

where I_0 denotes the modified zero-order Bessel function.

To make it easier to find the optimal α_K parameter for Eq (2.6.1), assumptions need to be made. These include:

- empirical relationships between the value of α_K and the rate of sidelobe roll-off;
- how many k values contribute a significant amount to the convolution with the DFT of a DCW

The computation of an optimal α_K requires substantially more computational effort as was described in [120]. Therefore it would be unrealistic to expect it to run in real-time for use in a DRx in an ESM system. Therefore SVA using the Kaiser window function was not studied further.

2.7 Noninteger SVA Algorithm

The original formulation of SVA detailed in section 2.4 require data to be sampled at integer multiples of the Nyquist grid in the frequency-domain. If the data is sampled at a non-integer multiple of the Nyquist grid, upsampling and interpolation must be implemented. This method is described in [122, 123] under the name Super-SVA. However the intended application for this method is two-dimensional upsampled SAR images and it would be difficult to envisage a situation in a DRx in an ESM system where Nyquist-sampled data or integer multiples thereof would not suffice.

The computational complexity of upsampling and interpolation SVA algorithms makes them impractical for a DRx as part of an ESM system and the methods were not studied further in this work.

2.8 Complex-parametered SVA

This section details a novel idea for SVA, which is not detailed in the established literature. The approach for this SVA technique not only adaptively alters the window function shape but also the time shift of the window. These two degrees of freedom are shown in the panels of Fig 2.8.1. Panel (a) shows the effect of variation of the shape parameter whilst keeping the shift parameter fixed and panel (b) shows the effect of variation of delay whilst shape is constant. The motivation behind this technique was to improve the reduction of sidelobes further whilst maintaining mainlobe resolution.

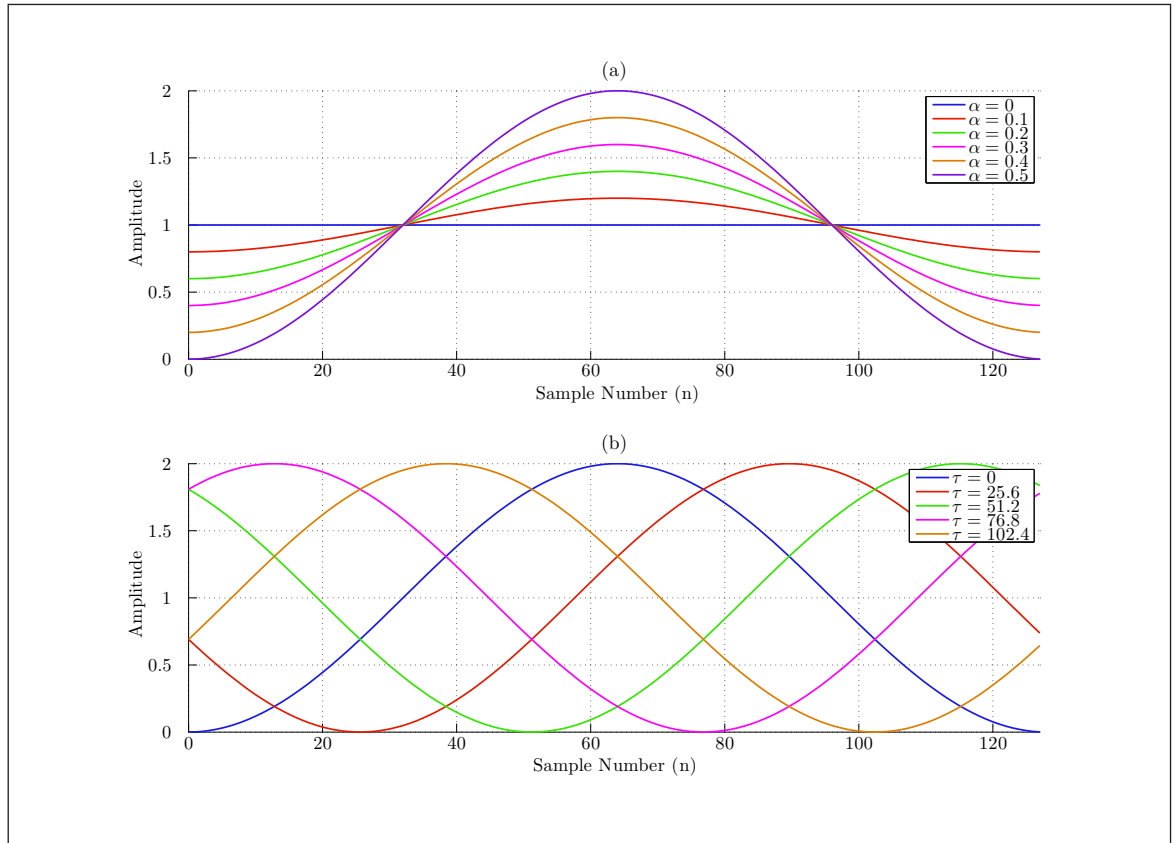


Figure 2.8.1 Variation in: (a) shape parameter whilst delay constant; (b) delay parameter whilst shape constant

2.8.1 Mathematical Basis

The equation which governs the shape of raised-cosine window functions, i.e. Eq (2.2.1), is modified such that it is circularly shifted (i.e. delayed) by τ samples (not necessarily integer samples) to become

$$w(n) = 1 - 2\alpha \cos\left(\frac{2\pi(n - \tau)}{N}\right) \quad (2.8.1)$$

As with 3-point joint-SVA (hereafter referred to simply as SVA), to keep the time-domain window function values positive, the α parameter is restricted to

$\{\alpha \in \mathbb{R} \mid 0 \leq \alpha \leq 1/2\}$. Upon transformation of Eq (2.8.1) by a length $N_{\text{FFT}} = N$ DFT

$$\begin{aligned} W(k) &= \delta(k) - \alpha e^{-\frac{i2\pi\tau}{N}} \delta(k+1) - \alpha e^{\frac{i2\pi\tau}{N}} \delta(k-1) \\ &= \delta(k) - \beta \delta(k+1) - \beta^* \delta(k-1) \end{aligned} \quad (2.8.2)$$

The time delay (τ) in Eq (2.8.1) in the time-domain is equivalent to a phase shift ($e^{i\phi}$) in Eq (2.8.2) in the frequency-domain. The characteristic parameter from SVA, α , which defines the shape of the raised-cosine window function in Eq (2.2.1) is generalised to a complex parameter β , which defines the shape and time delay (i.e. cyclic position). Explicitly, β as a complex number in polar co-ordinates, is

$$\beta = \alpha \exp\left(-\frac{i2\pi\tau}{N}\right) \quad (2.8.3)$$

In the special case of no time delay in Eq (2.8.3) the original SVA α parameter is recovered. Therefore the magnitude of β (that is, α) is restricted to $\{\alpha \in \mathbb{R} \mid 0 \leq \alpha \leq 1/2\}$ to keep the time-domain window functions always positive.

Equation (2.8.2) convolved with the DFT of a DCW yields

$$X_a(k) = X(k) - \beta X(k+1) - \beta^* X(k-1) \quad (2.8.4)$$

Equation (2.8.4) shows that the process of the application of a complex-parameterized version of SVA requires only a simple calculation to be done at each frequency bin in the frequency-domain.

The choice of the complex value to assign to the β parameter at each frequency bin in Eq (2.8.4) is more difficult. It can be chosen via cost function (C) minimisation. A suitable cost function may be $C = |X_a(k)|^2$, as in the SVA formulation. However the cost function is minimized with respect to a complex parameter. If the cost function is minimized in cartesian co-ordinates then the form of β that is used is $\beta = \beta_R + i\beta_I$ and there are two resulting partial differential equations (PDEs)

$$\begin{aligned} \frac{\partial C}{\partial \beta_R} &= 2A\beta_R - 4B\beta_I - 2E \\ \frac{\partial C}{\partial \beta_I} &= 2D\beta_I - 4B\beta_R - 2F \end{aligned}$$

where A, B, D, E and F are constants and are defined in appendix B. Full derivation of the PDEs is given in appendix B.

The following stationary point is obtained after combining the PDEs

$$(\beta_R, \beta_I) = \left(\frac{2BF + DE}{AD - 4B^2}, \frac{2BC + AF}{AD - 4B^2} \right) \quad (2.8.5)$$

This stationary point is shown to be a minimum or in the worst case part of a valley in appendix B.

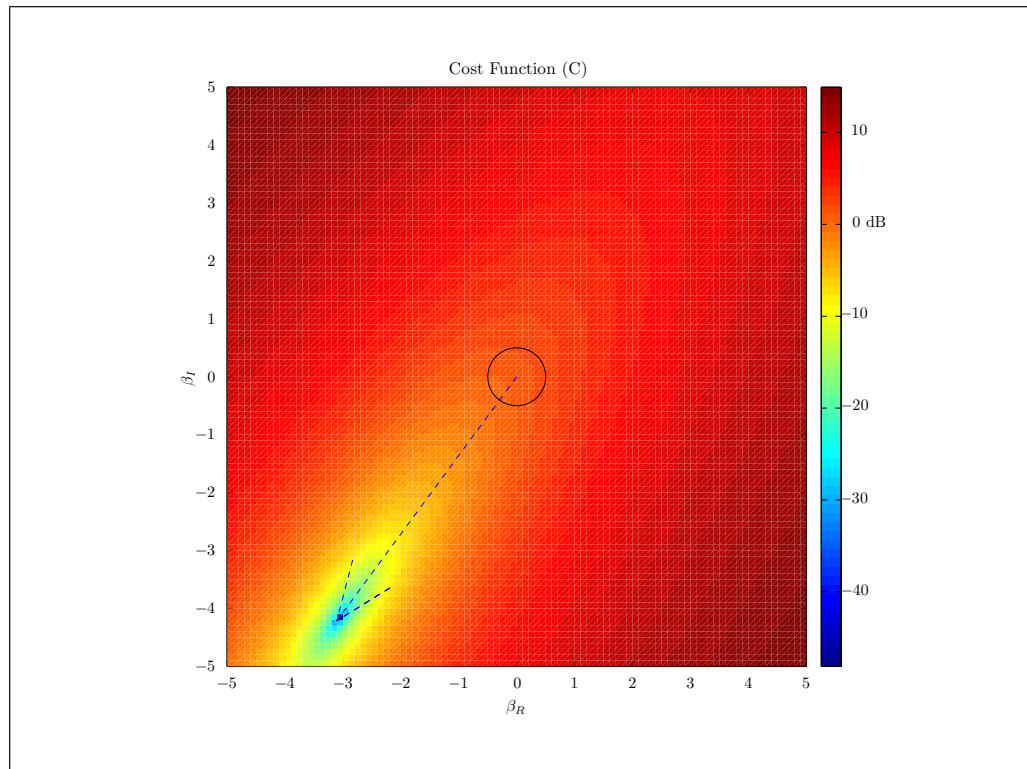


Figure 2.8.2 Example cost function surface for β variation

Unfortunately this minimum may not correspond to a time-shifted raised-cosine window function that is always positive-valued in the time-domain because the aforementioned minimisation of C was unconstrained. SVA requires that the α parameter of the raised-cosine window functions be limited in the interval $[0, 1/2]$. For complex-parametered SVA this requires the magnitude of the β parameter to be less than 0.5. The problem is illustrated in Fig 2.8.2 where the magnitude-squared cost function has been plotted for three adjacent frequency bins with variation in β . The black circle represents the maximum magnitude of 0.5 that the β parameter can assume. However the minimum point in the cost function surface occurs outside of the maximum-allowable magnitude of the β parameter. This is indicated by the blue arrow. It would be incorrect to simply use the phase at the minimum point of the cost function and limit the magnitude of the β parameter to 0.5. This is because the minimum point within the black circle may not necessarily be at that particular phase with a magnitude of 0.5.

If the solution to the unconstrained minimisation is used to produce a spectral estimate, it will not be useful because it is not distortionless. Therefore unconstrained complex SVA is not used in this work. An example of this failure is shown in Fig 2.8.3 for a single tone of $\text{SNR}_{\text{in}} = 20$ dB which is processed with various window techniques. The resultant spectral estimates are normalised.

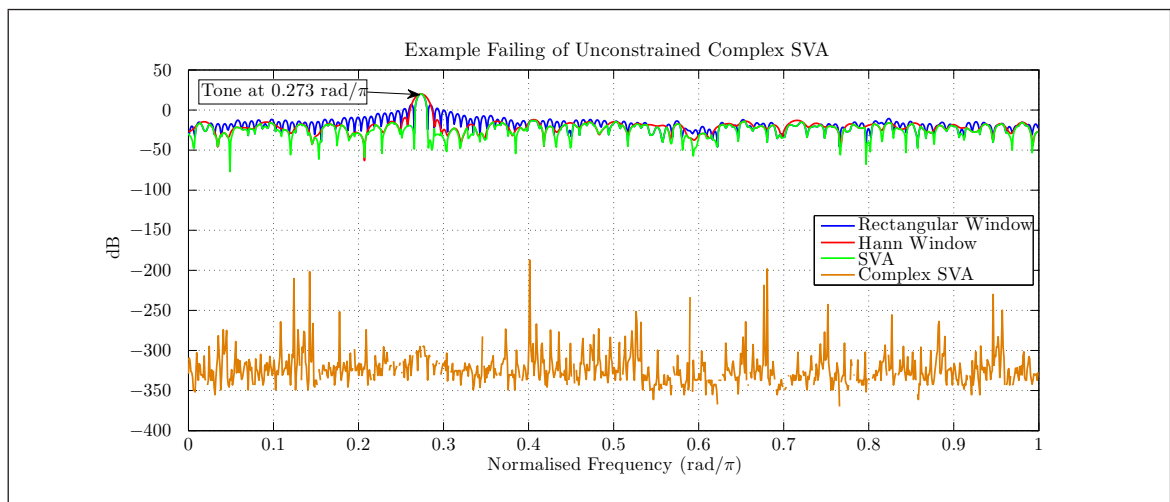


Figure 2.8.3 Example spectral estimates for different windowing techniques

2.8.2 Minimisation using a Lagrangian Multiplier

As a simple extension, which can be generalised later, a constrained minimisation problem can be solved with the use of a Lagrangian multiplier. The constraint could be taken to be $\alpha = 0.5$. This is not entirely correct but the Lagrangian multiplier technique operates with equality constraints. This would limit the solutions to those on the black circle in Fig 2.8.2. The solution satisfies the following equation

$$G \sin(2\phi) - B \cos(2\phi) + E \sin(\phi) - F \cos(\phi) = 0 \quad (2.8.6)$$

where $G = -(X_R(k+1)X_R(k-1) + X_I(k+1)X_I(k-1))$ and the complex parameter β is in polar coordinate form ($\beta_R = \alpha \cos(\phi)$ and $\beta_I = \alpha \sin(\phi)$). An example of the use of constrained minimisation with the use of a Lagrangian multiplier is shown in Fig 2.8.4. In both panels the constrained minimisation with the use of a Lagrangian technique is compared against simpler window methods. The example graphs demonstrate there would be no advantage in the use of this constrained minimisation version of SVA (i.e. Lagrangian SVA) over SVA. The Lagrangian SVA spectral estimate has a broader mainlobe and does not have better sidelobe reduction (in fact it virtually sits on top of the Hann-windowed estimate in panel (a) and its mainlobe sits on top of the Hann-windowed estimate in panel (b)). This is because the Lagrangian SVA is constrained to the use of Hann window functions albeit time-shifted.

2.8.3 Constrained Minimisation

Since Eq (2.8.6) has no simple analytical solution it is natural to extend the constraint $\alpha = 0.5$ to the desired $\alpha \leq 0.5$ and use the Karush-Kuhn-Tucker (KKT) conditions. This also has no analytical solution but can be solved numerically in programs such as MATLAB®.

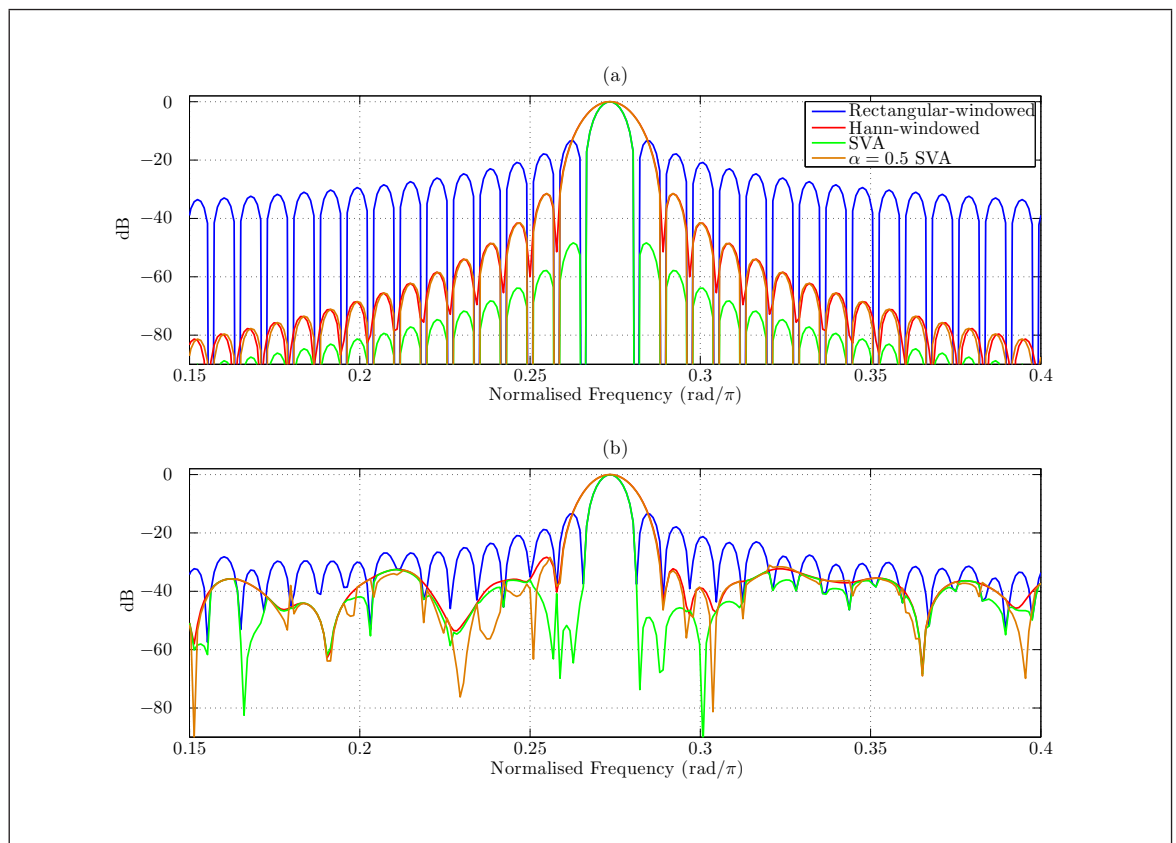


Figure 2.8.4 Example rectangular-windowed, Hann-windowed, SVA and Lagrangian SVA spectral estimates for: (a) a single tone present in all samples; (b) single tone with noise present at $\text{SNR}_{\text{in}} = 20$ dB

2.8.4 Evaluation of Constrained Complex SVA

The first example of how the constrained minimisation version of complex SVA compares to other windowing techniques is given in Fig 2.8.5. In this, a single tone of normalised frequency $0.109 \text{ rad}/\pi$ was present throughout the DCW and processed by the rectangular window, the Hann window, SVA and constrained complex SVA techniques. Panel (a) shows the result when no noise is present and panel (b) shows the result when the input DCW SNR_{in} was 20 dB. In panel (a) the constrained complex SVA technique appears to have sidelobes greatly reduced over the SVA technique and the complex SVA mainlobe appears even narrower than in the SVA technique (as usual, the rectangular-windowed and SVA mainlobes sit on top of each other).

In panel (b) the addition of AWGN brought the noise floor of the complex SVA technique up to the same level as the SVA technique and Hann-windowed technique. This was to be expected because, like the SVA technique, the complex SVA technique is a sidelobe-reduction technique not a noise-reduction technique.

Another example of the constrained complex SVA is given in Fig 2.8.6. In this example two tones which are closely spaced in frequency and of equal amplitude were processed by the different window techniques. The frequency spacing between the tones was actually $1.5f_s/N$. Panel (a) shows the case with no noise present and panel (b) shows

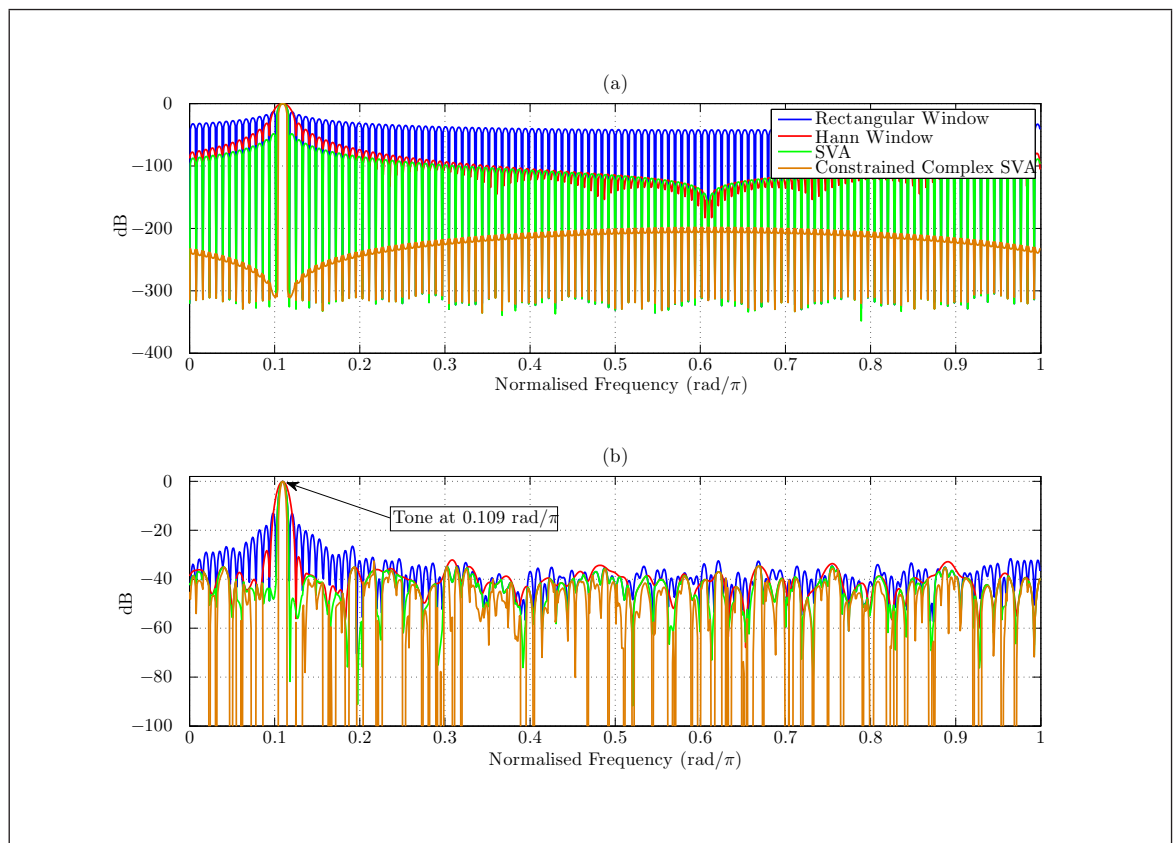


Figure 2.8.5 Example rectangular-windowed, Hann-windowed, SVA and constrained complex SVA spectral estimates for: (a) a single tone present in all samples; (b) single tone with noise present at $\text{SNR}_{\text{in}} = 20$ dB

the case at $\text{SNR}_{\text{in}} = 20$ dB. The results from panel (a) are encouraging because they show the ability of the constrained complex SVA technique to distinguish between two tones was better than the other window techniques. This ability was degraded to the noise floor by the addition of AWGN in panel (b).

However there are two major failings of the constrained complex SVA: (1) it struggles to produce accurate spectra when a weaker signal is below the SVA sidelobe level of another, stronger signal and (2) it is not distortionless in the case of interception of a partial pulse in a DCW.

The first of these disadvantages is demonstrated in Fig 2.8.7. In panel (a) two tones are clearly visible at frequencies $0.109 \text{ rad}/\pi$ and $0.137 \text{ rad}/\pi$. The relative power difference between the two tones was 20 dB and between the stronger tone and the noise floor the SNR_{in} was 100 dB. However, in panel (b), when the power of the second tone was reduced beneath the sidelobes of the SVA technique so that the power difference between strong and weak tones was 80 dB, the constrained complex SVA struggled to produce a spectral estimate which could detect the presence of the weaker tone. This is a large failing because it implies that the sidelobe reduction of constrained complex SVA is meaningless as it is unable to uncover any additional information below the sidelobe level of SVA.

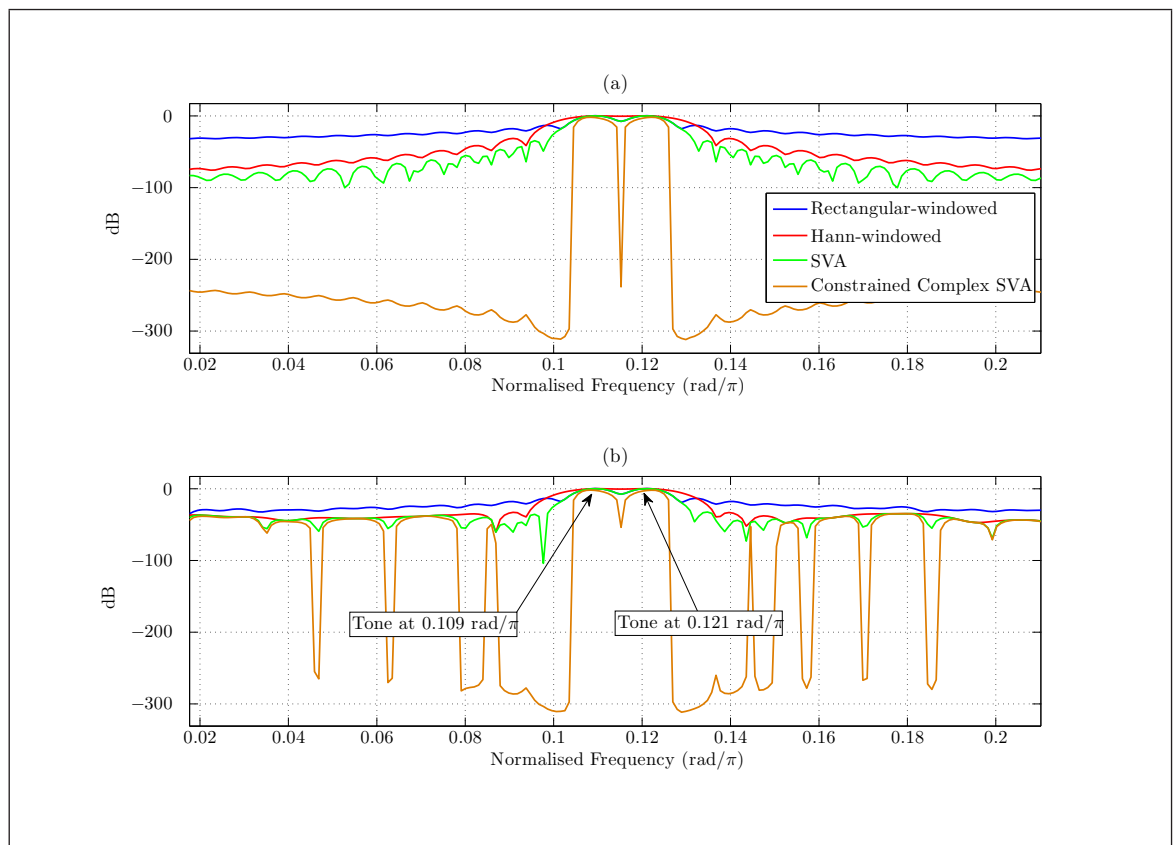


Figure 2.8.6 Example rectangular-windowed, Hann-windowed, SVA and constrained complex SVA spectral estimates for: (a) two tones present in all samples; (b) the same tones with noise present at $\text{SNR}_{\text{in}} = 20$ dB

The second of these failings is displayed in Fig 2.8.8. In this example, a tone was captured after appearing halfway through the DCW. Therefore only part of the available signal power was captured. This situation is pictorially explained in panel (b) by the blue captured signal. Panel (a) shows that whilst the SVA technique performed well for only half of the signal in time (it sits virtually on top of the rectangular-windowed), the constrained complex SVA did not. This is because constrained complex SVA minimised at the tone frequency by using a Hann window function which was shifted in the time-domain such that the zero or null was over the centre of the captured signal portion as demonstrated by the red window function in panel (b).

The two failings are major disadvantages of constrained complex SVA. Additional arguments against the usage of constrained complex SVA instead of SVA are given in section 3.4 where the performance of constrained complex SVA is compared against other windowing techniques in a Monte Carlo simulation.

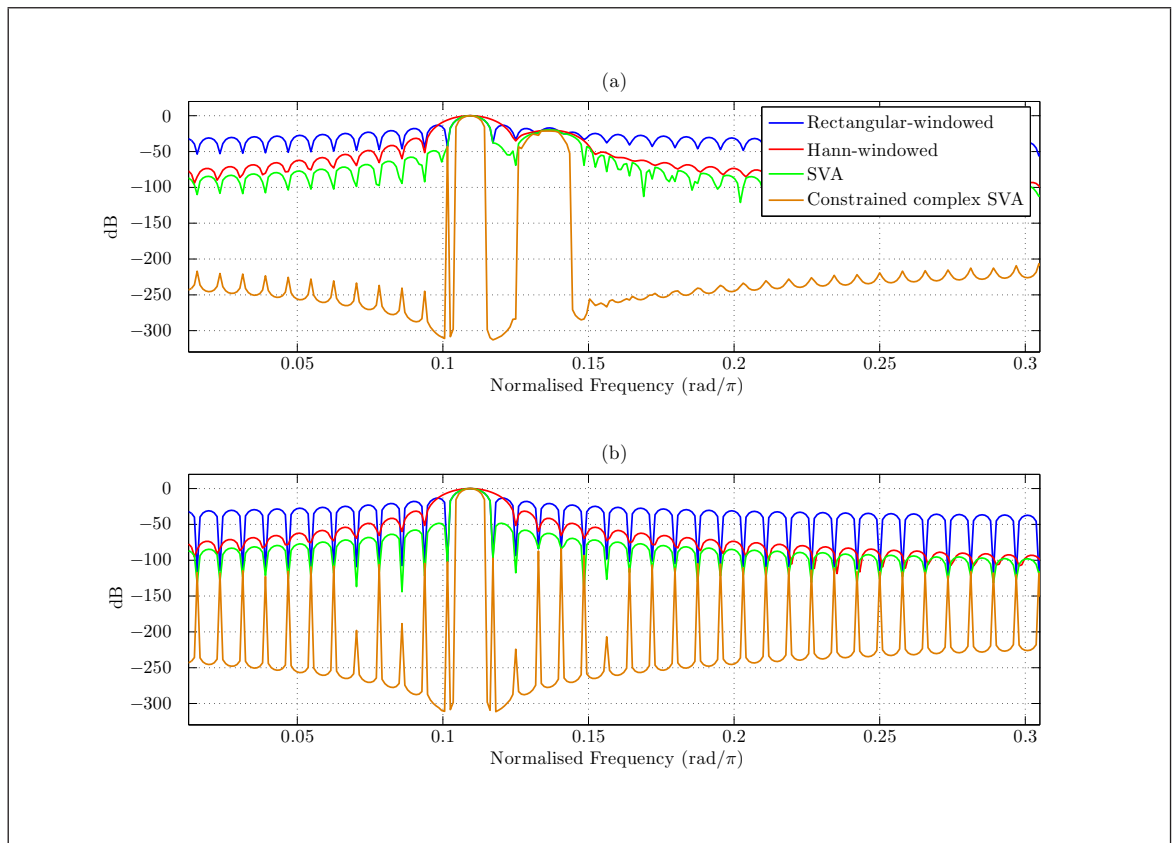


Figure 2.8.7 Example rectangular-windowed, Hann-windowed, SVA and constrained complex SVA spectral estimates for a strong tone and weak tone at: (a) 20 dB power difference; (b) 80 dB power difference

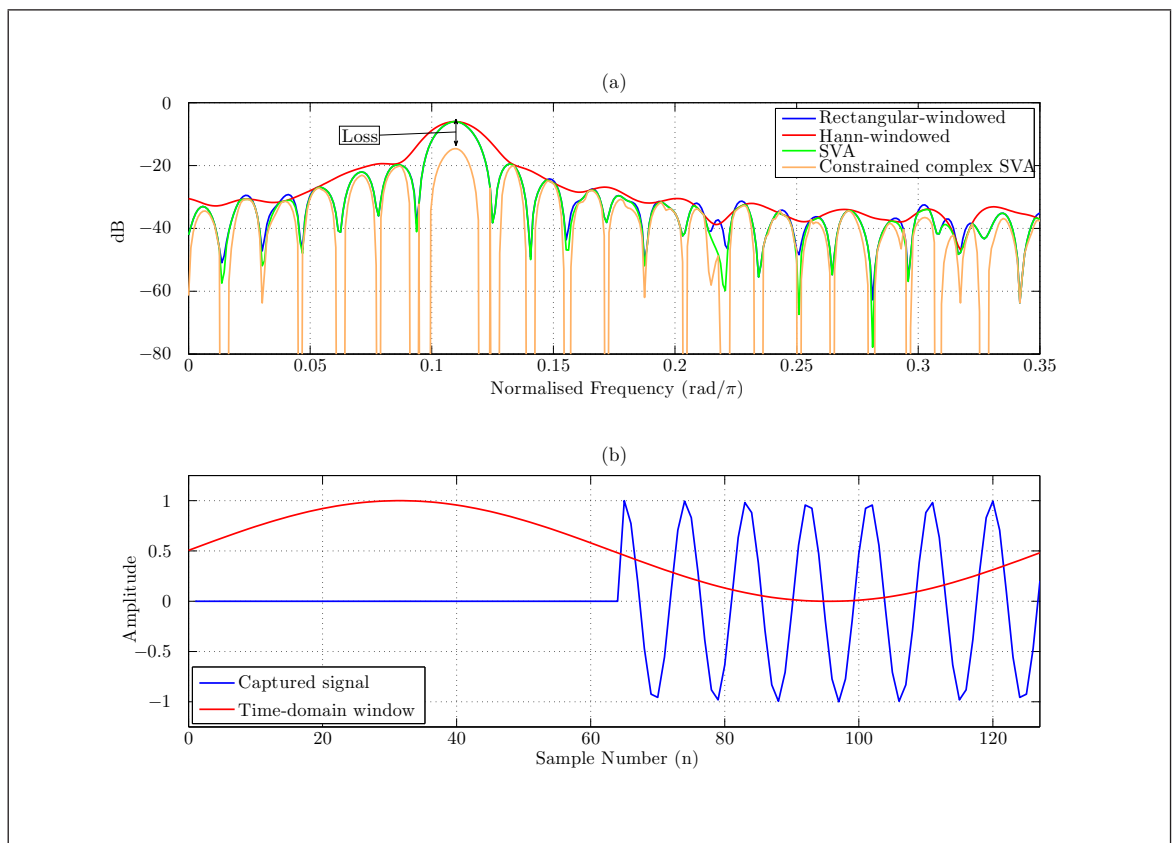


Figure 2.8.8 Example for partially-filled DCW: (a) example spectra when processed with different window techniques; (b) pictorial representation of such a DCW

CHAPTER 3

Floating-Point SVA

Section 2.4 advocated the use of SVA in a DRx as part of an ESM system. However section 2.4 did not offer rigorous simulation or experimental results to support the claim for the use of SVA. This floating-point SVA chapter aims to provide evidence for the advantages which SVA would offer when applied in a DRx. Section 2.4 showed that the joint 3-point SVA algorithm was the most likely candidate for eventual implementation in an EW DRx and all references to SVA hereafter refer to this particular flavour of SVA. Also some additional information is given in this section to show the merits of the joint 3-point SVA algorithm over other SVA varieties.

During algorithmic development it is not always clear that a candidate algorithm would be appropriate for implementation until modelling, simulation and testing are completed. In terms of MBD, the process may begin but then be abandoned because it becomes clear that the simulation results show that the algorithm, on balance, is inappropriate for the application. Therefore this chapter contains simulation results, which in addition to providing evidence for the advantages of the use of SVA in a DRx, also represent the start of the MBD process. With reference to appendix C, the start of the MBD process is the golden reference **floating-point model**.

This chapter compares simulation results from detectors which utilised SVA against conventional detector architectures. The simulated SVA detectors and conventional simulated detectors were essentially the same. However the former used the SVA technique whilst the latter used fixed window functions. There are many different scenarios in which the detectors could be tested. This is because in the field of EW, very few assumptions can be made about unknown signal waveforms that may be present in an operational environment. Nevertheless some scenarios were chosen and are detailed in this chapter. They were chosen because they were judged to be the most pertinent for an ESM system. The simulation results were obtained with the use of floating-point arithmetic. The chosen scenarios, which will be detailed later in this chapter were:

- full DCW
- randomly-filled DCW
- interference present in full DCW

- randomly-filled DCW with interference present
- LPI signals in DCW.

3.1 Detectors Summary

A summary of the detectors used in this chapter is shown in Table 3.1.1.

Detector	Algorithm
A1	Rectangular-windowed periodogram (single sig)
	Hann-windowed periodogram (single sig)
	Chebyshev-windowed periodogram (single sig)
B1	SVA periodogram (single sig)
C1	Rectangular-windowed periodogram
	Hann-windowed periodogram
	Chebyshev-windowed periodogram
D1	SVA periodogram (with decision rules)
E1	SVA periodogram (no decision rules)
F1	Complex-parametered SVA periodogram

Table 3.1.1 Conventional and SVA detectors

3.2 Full DCW

3.2.1 Scenario and Simulation Description

In this scenario a sinusoidal signal of unknown frequency (f_T) and unknown initial phase (ϕ_0) within the monitored bandwidth was present amongst background noise. The sinusoidal signal used was actually a complex exponential and the background noise was complex, with both real and imaginary components drawn from a normal distribution. The real and imaginary components of the noise were independent identically distributed (i.i.d.). The signal, by definition, was therefore extremely narrow band and would fall under the sinusoid radar signal category from chapter 1. The situation for this scenario was that **each** input sample was composed of both instantaneous power from the simulated radar signal and the background noise. In other words the DCW was full of signal power and hence the title for this section.

This scenario is preferable for detection algorithms in a DRx because the DCW happens to have intercepted as much of the signal as possible as permitted by the length of the DCW. The full DCW approach is demonstrated in Fig 3.2.1 where the eventual DCW

is formed from the summation of a complex signal and complex noise. It is clear that some of the signal is present in every sample in the DCW.

The monitored bandwidth of the detectors was the equivalent of one frequency bin width (f_s/N_{FFT}) but the exact frequency of the sinusoid was randomly chosen at the start of each DCW trial the frequency bin interval. This concept is illustrated in Fig 3.2.2 where it can be seen that most of the time f_T will not line up with the frequency samples of the DFT in the frequency-domain. Similarly, the initial phase of the sinusoid was randomly chosen at the start of each DCW trial.

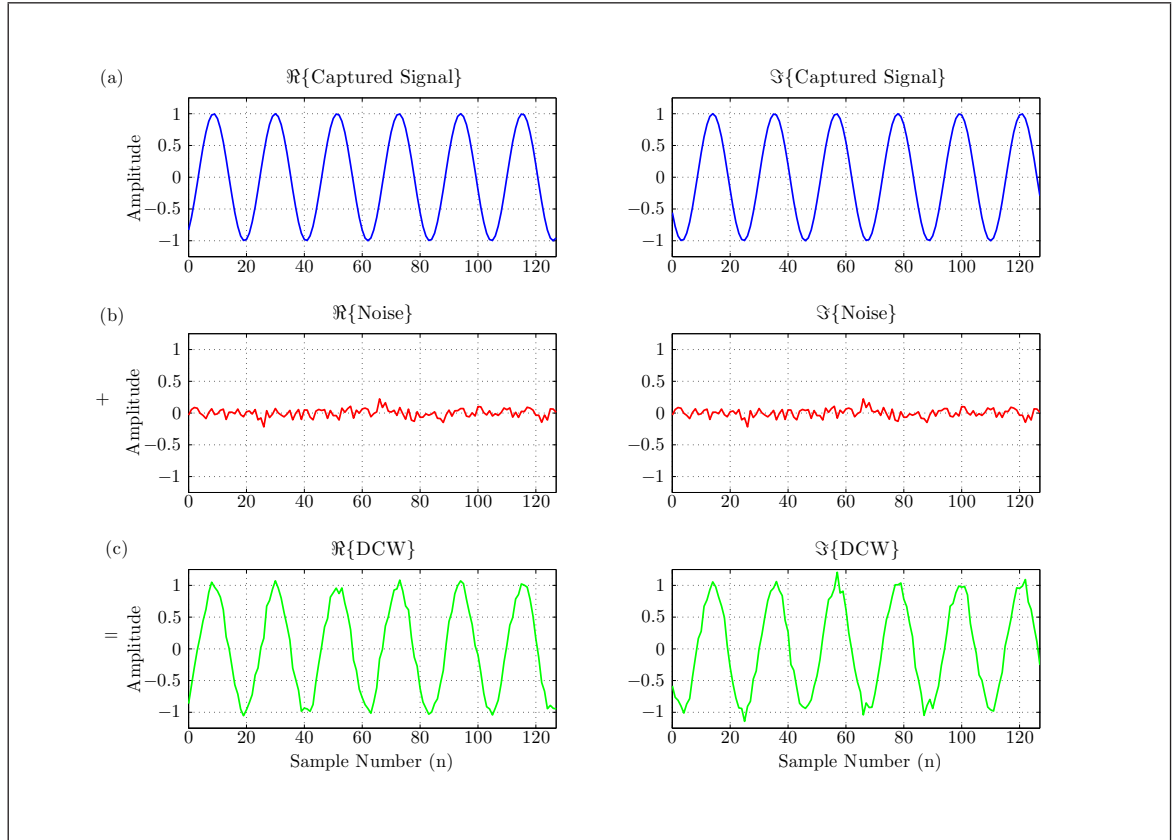


Figure 3.2.1 (a) complex signal; (b) complex noise at every sample; (c) the full DCW formed from the summation

There were no simulated interference radar signals in this simulation. Two different FFT lengths were tested for the same DCW length of $N = 128$ samples such that there was a non-zeropadded case of $N_{\text{FFT}} = 128$ data points and a zeropadded case consisting of $N_{\text{FFT}} = 256$ data points. Zeropadding the FFT up to twice the DCW's original length was done to allow for a performance comparison with non-zeropadded case. This is because calculating the odd frequency bins allows for more accurate peak detection and peak frequency location.

The conventional simulated detector applied one window function to the DCW out of a choice of: (a) a rectangular window function; (b) a Hann window function or (c) a (Dolph-)Chebychev window function. The system level diagram of the conventional simulated detector is illustrated in Fig 3.2.3. In this work, this detector architecture is

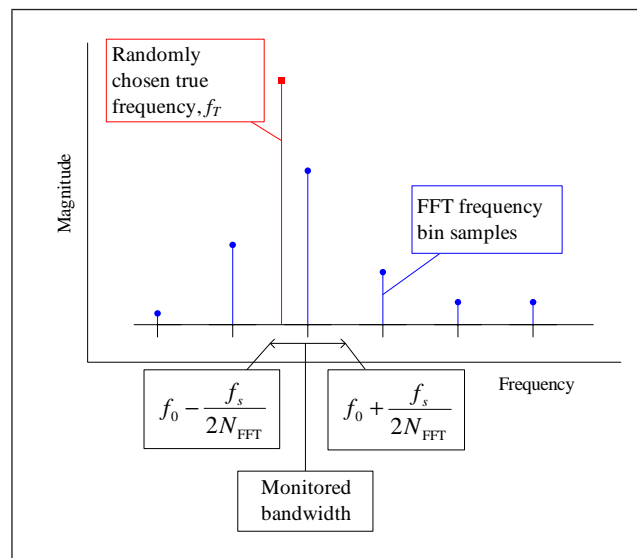


Figure 3.2.2 The signal frequency is randomly selected from the monitored bandwidth at the start of each DCW

referred to as **detector A1**. The simulated SVA detector applied the SVA technique to the FFT output. A flowchart of the data flow for the SVA detector used in this scenario is shown in Fig 3.2.4. In this work, it is referred to as **detector B1**. A comparison of the two detectors allowed the effect of each windowing technique to be seen. There are a large number of possible window functions which could be used in detector A1, any of which could have been chosen for the comparison. However the reason for the choice of the rectangular window function and Hann window function is that they are the two limiting cases from which SVA can select its equivalent time-domain window functions. Meanwhile the reason for the inclusion of the Chebychev window function is that it is commonly used within the aerospace industry as it provides a constant sidelobe envelope which is helpful for designing aspects such as threshold levels. A constant sidelobe level of -60 dBc was used in all of the simulations that involved the Chebychev window function.

Typical ESM system DRx parameters were used to populate the parameters of detectors A1 and B1 in this scenario and they are listed in Table 3.2.1¹.

DRx:	sampling rate	[1.0, 1.75] GHz
	sampling period	[571, 1000] ps
	DCW length	128 samples
	FFT mode	complex
	DCW overlap	0 samples
	non-zero padded frequency bin width	[7.8, 13.7] MHz
	DCW period	[73, 128] ns

Table 3.2.1 Typical parameters for a DRx as part of an ESM system

¹The exact parameters used are not shown

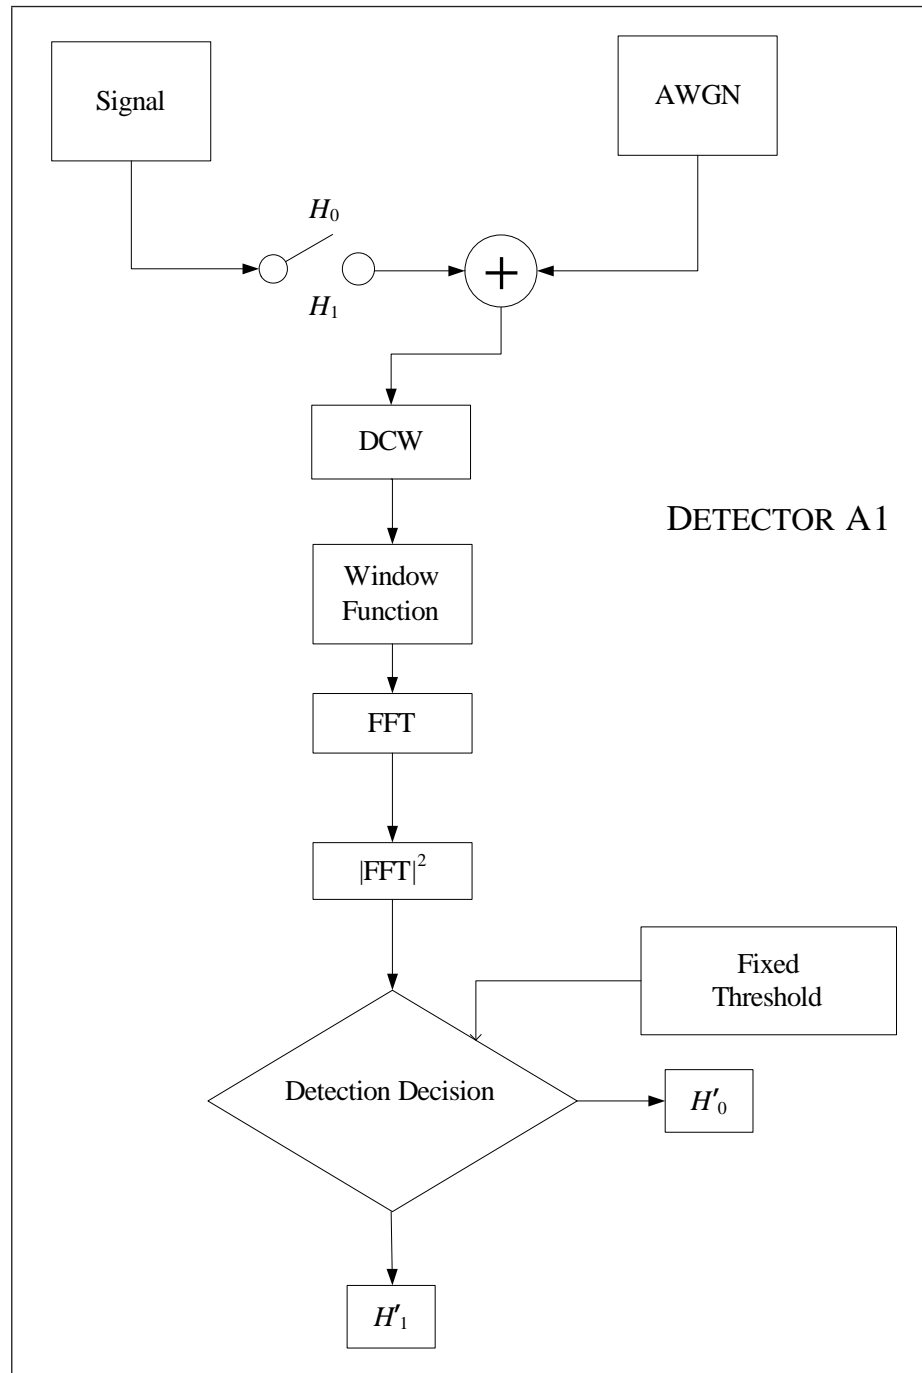


Figure 3.2.3 Detector A1 data flowchart used for the full DCW scenario

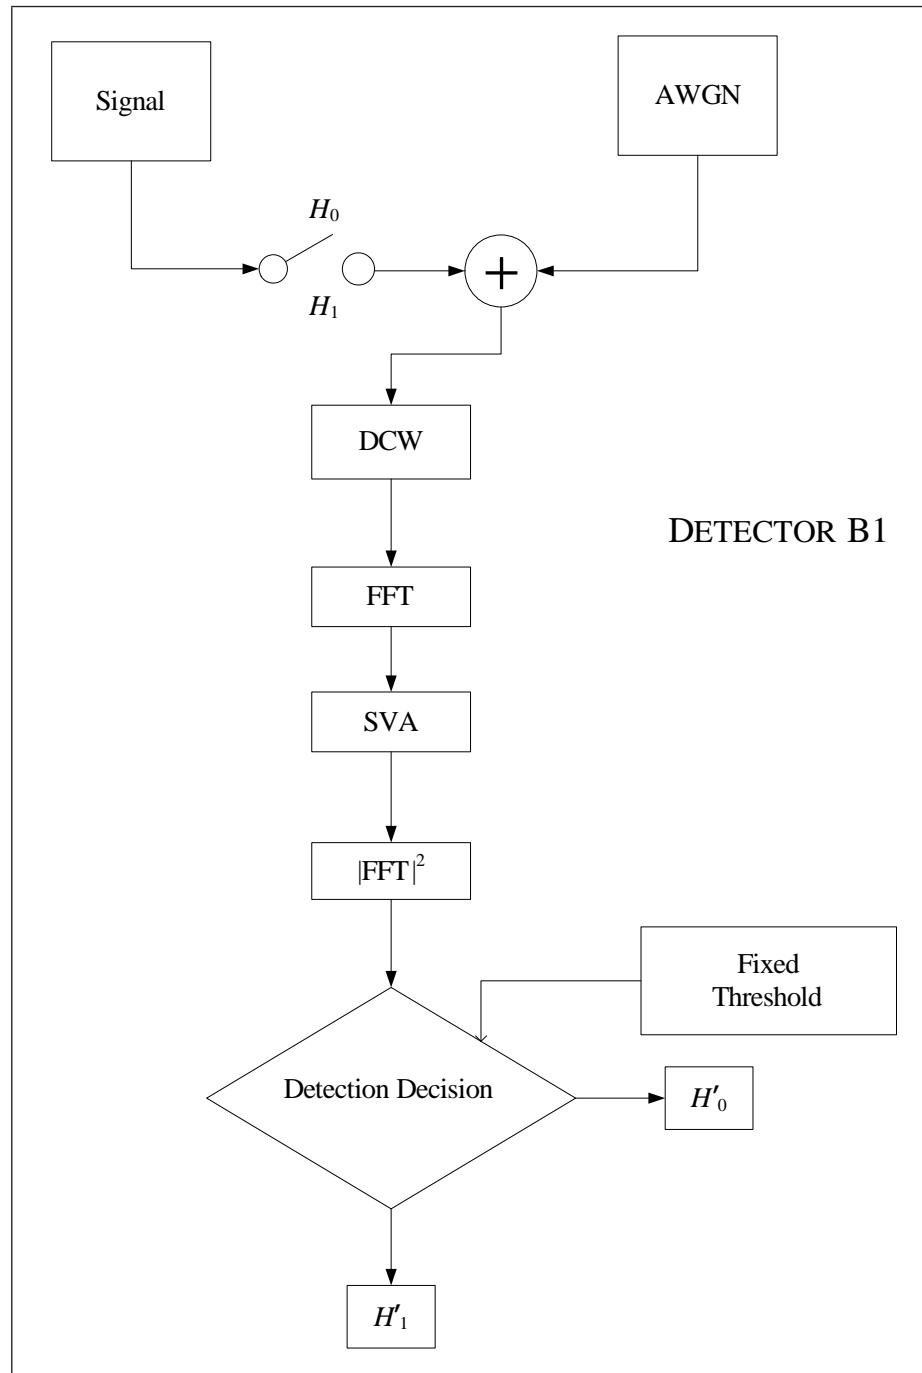


Figure 3.2.4 Detector B1 data flowchart used for the full DCW scenario

The simulated detectors placed a threshold on the monitored frequency bin and made decisions on whether a signal was present or not after each DCW. These decisions were tallied in a 2×2 contingency table or confusion matrix, an example of which is given in Table 3.2.2.

P_d was calculated such that it could be plotted in a receiver operating characteristic curve against either P_{fa} or SNR_{out} dependent on the parameter that was kept fixed during the simulation (SNR_{out} or P_{fa} respectively). SNR_{out} is defined as the SNR at the output of the FFT block

$$\text{SNR}_{\text{out}} = 10 \log_{10} \left(\frac{N \text{Amp}^2}{N_p} \right)$$

where Amp is the constant envelope amplitude of the simulated radar pulse signal and N_p is the noise power. This allowed the results to be the same regardless of DCW length because the coherent processing gain introduced by the FFT was addressed. The number of trials for each data point in the plots was 1×10^5 .

For the P_{fa} against P_d plots, the empirical P_{fa} was calculated from the summation of the total number of false alarms divided by the total number of DCWs when there was no signal present (X). In a comparable way, the empirical P_d was calculated from the summation of the total number of correct detections and divided by the total number of DCWs when a signal was present (Y). For the SNR_{out} against P_d plots the thresholds necessary to achieve a particular P_{fa} were required all detectors. Therefore simulations to obtain P_{fa} against P_d preceeded simulations to obtain SNR_{out} against P_d in order to calculate the thresholds (except in the minority of cases where thresholds could be calculated analytically). For simulation purposes and project-time constraints, a relatively high P_{fa} of 0.01 has been used in most scenarios in this thesis.

		Actual		Total
		H_1	H_0	
Decision	H'_1	Correct Detection	False Alarm	Y'
	H'_0	Miss	Correct Rejection	X'
Total		Y	X	

Table 3.2.2 Contingency table used in computing P_{fa} and P_d

3.2.2 Results and Discussion

The receiver operating characteristic (ROC) curves for P_{fa} against P_d are shown in Fig 3.2.5. Panel (a) shows the results for the non-zeropadded case and panel (b) for the zeropadded case. The nearer the curves to the top-left corner of the plot, the better the performance of the detector which produced the plot. From the results, it is clear that the rectangular-windowed and SVA detectors had better detector performances than

the Chebychev-windowed and Hann-windowed detectors. The drop-off in performance of the Hann-windowed and Chebychev-windowed detectors was caused by the poorer NEB of these bell-shaped window functions. At the edges of the bell-shaped window, the input samples were attenuated and since all of the samples in the DCW in this full DCW scenario contained signal power, this resulted in a loss of some signal power. This effect for bell-shaped window functions is directly related to the NEB of the window function and is termed the **window loss** throughout this work.

In Fig 3.2.5, there is also a jump in detector performance from the non-zero padded case to the zero padded case which also widens the performance gap between the two groups of detectors: (1) the rectangular-windowed detector and the SVA-detector (i.e. detector B1) and (2) the Hann-windowed detector and Chebychev-windowed detector. The reason for the non-zero padded and zero padded performance difference is that twice zero padding the DCWs allowed the detectors to more accurately estimate spectral peaks.

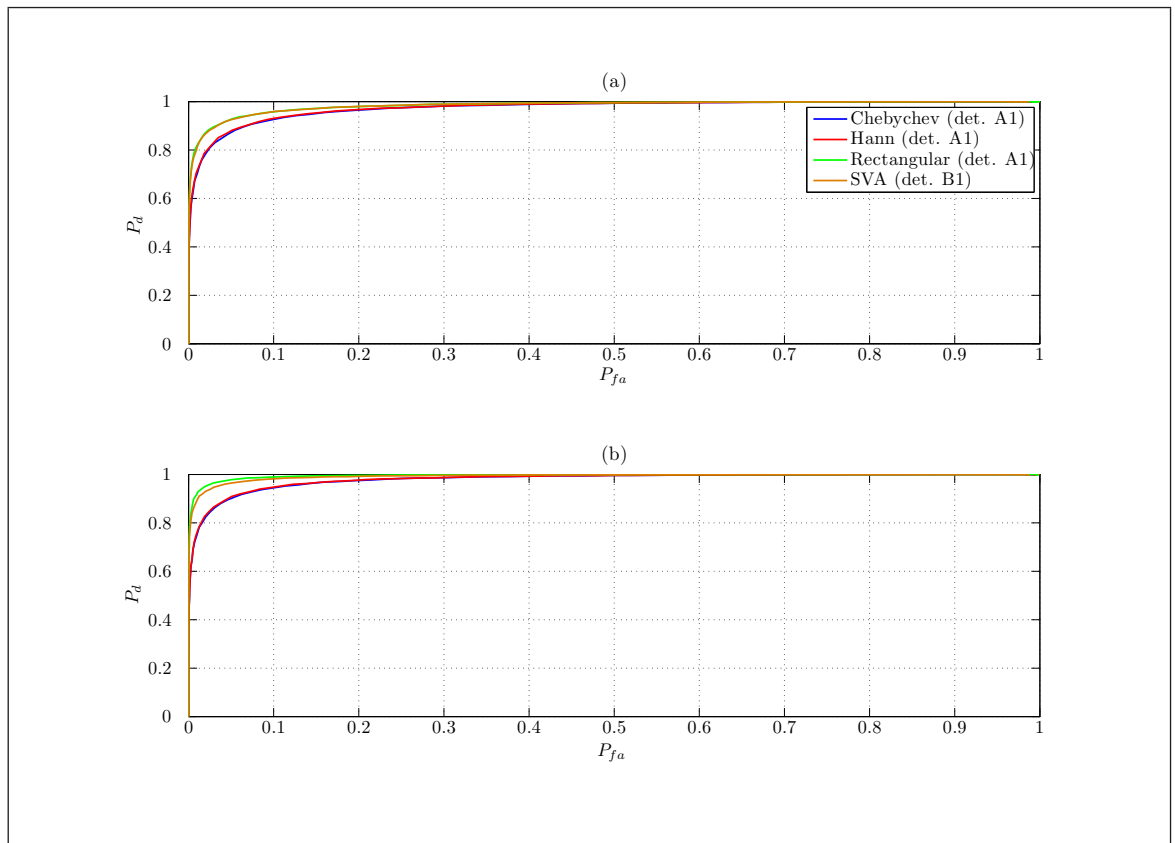


Figure 3.2.5 The simulated detector ROC curves of P_{fa} against P_d for: (a) the non-zero padded FFT; (b) the zero padded FFT

Typically this type of P_{fa} against P_d ROC plot is less useful for systems engineers wishing to get a feel of the processing gain of an algorithm and so they are shown in this full DCW scenario but later omitted from the other scenarios which show only SNR_{out} against P_d . One method to compare detectors is to specify the algorithmic processing gain at $P_d = 0.5$ which has been done throughout this thesis.

The ROC curves of SNR_{out} against P_d are shown in Fig 3.2.6. Panel (a) depicts the non-zeropadded case whilst panel (b) depicts the zeropadded case. For these results a threshold was set such that $P_{fa} = 0.01$. The results again show that the SVA detector (i.e. detector B1) and simulated rectangular-windowed detector have the strongest detector performances. In panel (a) the performance advantage associated with the use of SVA or a rectangular window is ~ 1 dB and in panel (b) it is ~ 1.5 dB. As in the other type of ROC curve plot, this performance difference is attributed to the window loss which the Hann and Chebychev window functions suffer. Although they have been grouped together in this analysis, there is a very slight difference in performance between the rectangular-windowed detector and SVA detector in both panels (a) and (b). This slight drop-off in performance of the SVA detector was caused by the algorithmic noise introduced during estimation of α parameters.

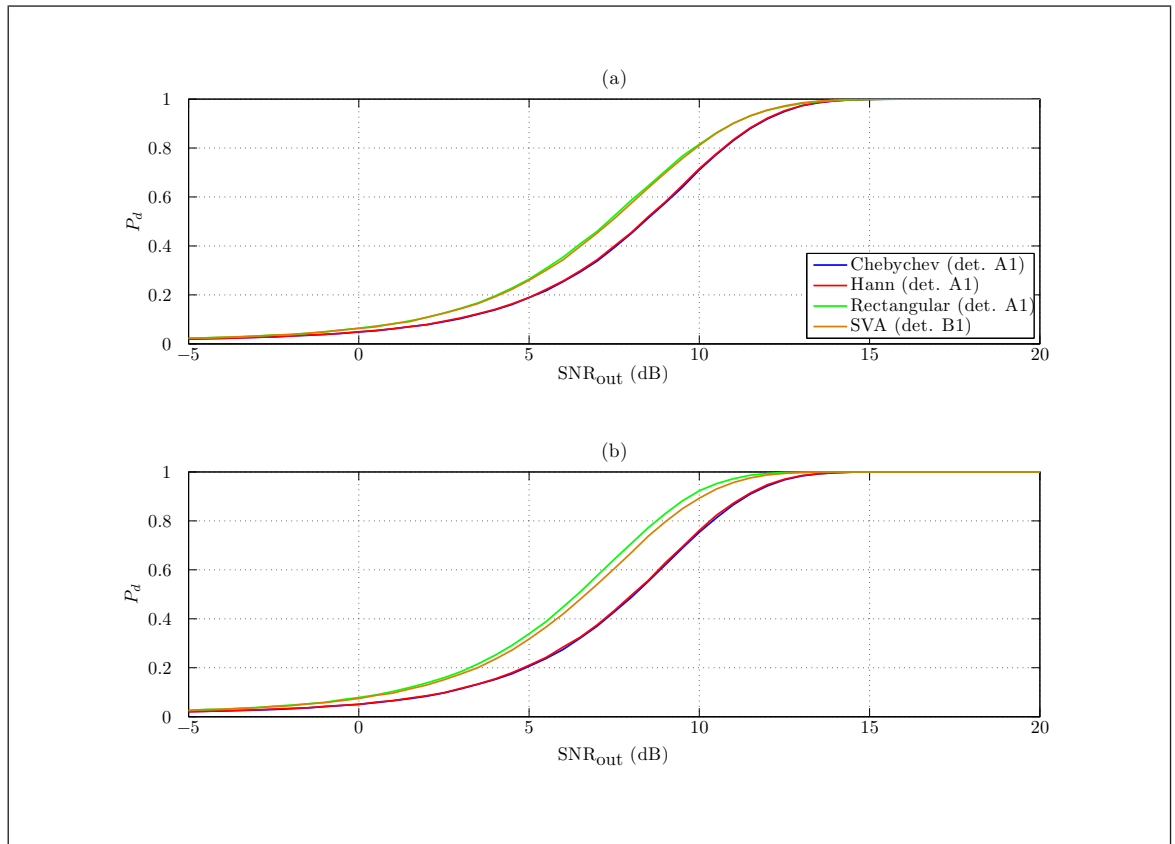


Figure 3.2.6 The simulated detector ROC curves of SNR_{out} against P_d for: (a) the non-zeropadded FFT; (b) the zeropadded FFT

In section 2.1 it was explained that in EW, signals in the environment and their parameters are inevitably unknown and to be cautious a moderate dynamic range window function such as a Hann or Chebychev window is almost always used. The important point therefore from these results is that the SVA technique retains the cautious approach of a moderate dynamic range window **and** is able to mostly recover the window loss of a moderate dynamic range window.

A full DCW situation is common for a DRx in an ESM system when a signal is present because for most real-life signals the pulse duration is much greater than the DCW

duration and therefore some of the DCWs will be full of signal power. There are however other scenarios to be considered. The scenario in section 3.3 removes one of the ideal conditions (the full DCW) by randomising the TOA of the signal pulse.

3.2.3 5-point and 7-point SVA Comparison

It was stated at the end of section 2.5 that further arguments would be provided against the use of 5-point SVA and 7-point SVA over the simpler (joint 3-point) SVA. Fig 3.2.7 presents one of these arguments. It depicts the use of 5-point SVA and 7-point SVA in the full DCW scenario described above. The data from the simulated (joint 3-point) SVA detector, Hann-windowed detector, Chebychev-windowed detector and rectangular-windowed detector from Fig 3.2.6 are simply redisplayed in Fig 3.2.7. The new 5-point SVA and 7-point SVA data originated from more simulations of detector B1 with the obvious change that the SVA block of detector B1 referred to the 5-point and 7-point SVA.

The results in Fig 3.2.7 show that the 5-point and 7-point SVA did not recover the window loss as well as the 3-point SVA. In fact in this scenario, the trend for the performance of the simulated SVA-based detectors was for them to fall in performance for each higher order of SVA. The performance of the 7-point SVA simulated detector was actually worse than that of the Hann-windowed and Chebychev-windowed detectors. Therefore Fig 3.2.7 provides strong evidence for sticking with joint 3-point SVA and not utilising 5-point SVA or 7-point SVA in a detector.

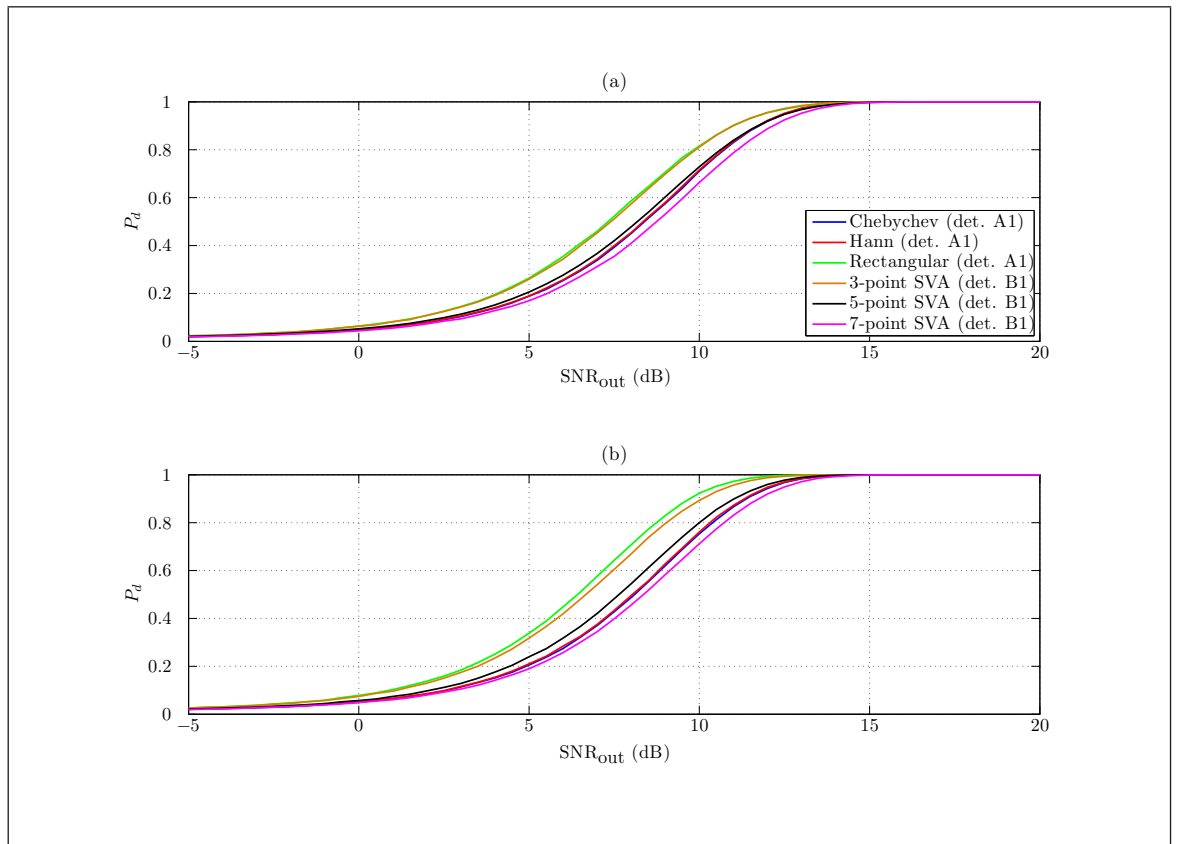


Figure 3.2.7 The results from Fig 3.2.6 alongside results using 5-point SVA and 7-point SVA in simulated detectors for: (a) the non-zero padded FFT; (b) the zero padded FFT

3.3 Randomly-filled DCW

3.3.1 Scenario and Simulation Description

This scenario introduced the random TOA aspect of a signal pulse. As a results of this the DCWs tested contained a fraction of the signal pulse. This was an informative scenario because a real-life DRx is unlikely to be time-synchronised to unknown radar signal pulses. A number of important assumptions were made about real-life DRxs in ESM system in order to create this scenario environment:

1. The DRx designer would select the DCW length in time to be matched to the shortest duration radar pulse that the designer would like to detect.
2. There would be no sample overlap between consecutive DCWs.
3. Where a radar pulse spans two DCWs, the DCW which contains the largest proportion of signal would be retained to test.
4. There are several DCWs per ESM dwell and they are considered separately.

The first point is a common design assumption, see for example [3]. The second assumption was included because DRxs as part of ESM systems operate in real-time and would struggle with the extra computational overhead associated with time-overlapped DCWs. Whenever a rectangular window (and therefore by extension SVA) is used it is unnecessary to overlap DCWs for the purposes of mitigating against window loss. The only reason to attempt to overlap when a rectangular window function (or SVA) is used would be to make it more likely that one of the DCWs would be better matched in time to the signal pulse arrival time. On balance, the extra computational effort required to overlap was deemed too great for the benefit it would provide because the typical signal pulse width is much greater than the typical DCW duration and therefore some DCWs would be full of samples which contain signal power.

With the third and fourth assumption, it should be stressed that a real-life DRx in an ESM system would constantly monitor its operational environment by capturing data in consecutive DCWs. The panels in Fig 3.3.1 illustrate this point. Urgent decisions over whether threats are present are made after each DCW. A signal pulse may arrive partway through a DCW and be partially present in the next DCW, but detection decisions are made after each DCW.

The extreme case, from point 1 above, was used in this scenario whereby the signal pulse duration was **as short as** the duration of the DCW. The detection process was simulated through examination of the processed DCW that contained the most signal after each test iteration. Therefore, with this pulse duration, the worst case occurred when the unknown signal appeared or finished in the middle of a DCW. This concept is illustrated in panel (b) of Fig 3.3.1, where, in this picture the DCW used in the test

iteration would be half-filled with signal power.

For a pulse as short as the DCW duration, the best case occurred when it arrived exactly at the start of a DCW because this created a DCW maximally filled with signal power. This is exactly the same DCW case dealt with in section 3.2. This case is illustrated in panel (a) of Fig 3.3.1.

All other partially-filled DCW cases occurred with equal probability between the best and worst cases. An example of this is shown in panel (c) of Fig 3.3.1, where in this picture, DCW 3 would be retained for testing. Aside from the random arrival time of

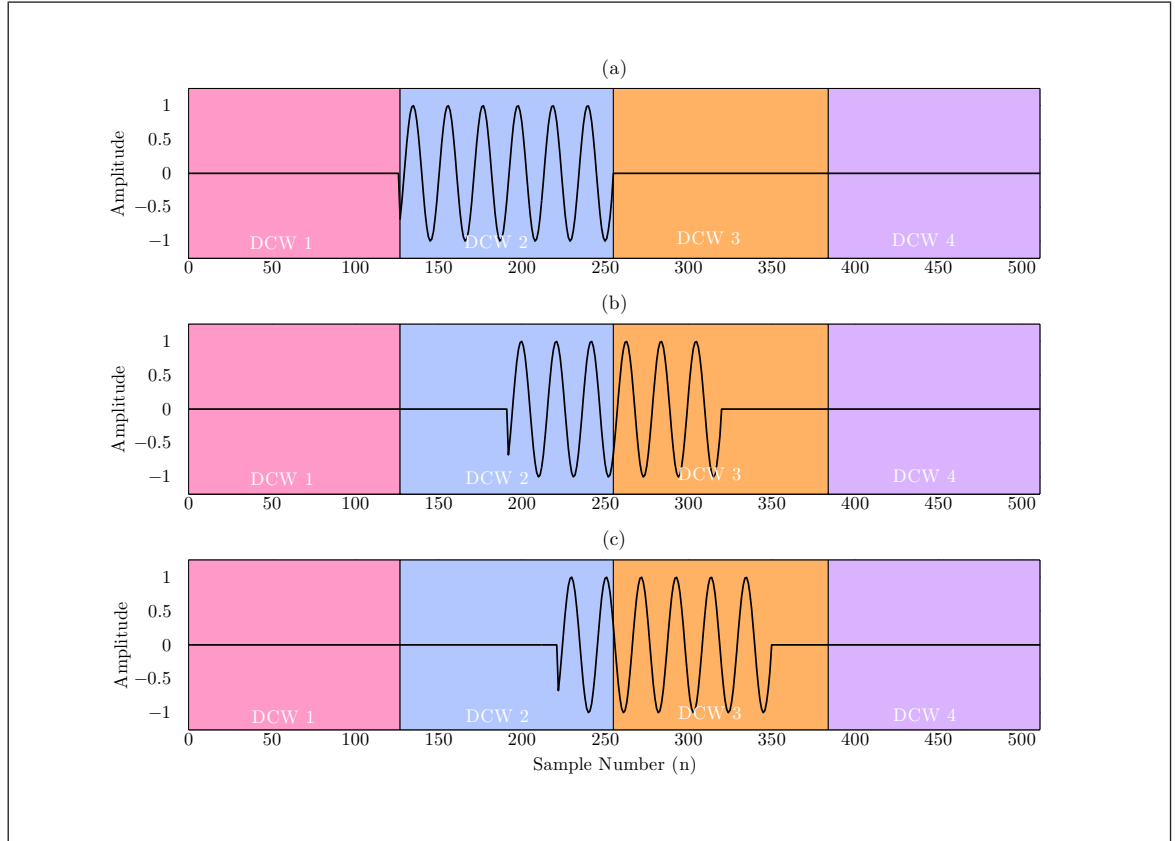


Figure 3.3.1 Random arrival time of pulses: (a) pulse completely fills DCW 2; (b) pulse is split equally across DCW 2 and DCW 3; (c) pulse partially fills DCW 2 and DCW 3, detection test performed on DCW 3

the signal pulses, this scenario used the same situational environment and typical DRx parameters as the full DCW scenario in section 3.2. The simulated detectors operated were detector A1 from Fig 3.2.3 and detector B1 from Fig 3.2.4.

3.3.2 Results and Discussion

The P_d against SNR_{out} results of Monte Carlo simulations are shown in Fig 3.3.2 for constant $P_{fa} = 0.01$. Each data point was produced from 1×10^5 detection decisions. Panel (a) in Fig 3.3.2 displays the results from the non-zero padded case, whilst panel (b) shows the results from the zero padded case. The results show that there was a smaller performance difference between the different simulated detectors when compared to

Fig 3.2.6 for the full DCW scenario in section 3.2. In panel (a) the performance gap is only ~ 0.4 dB and in panel (b) the performance gap is only ~ 0.6 dB. This closure of the gap is because the window loss effect was much reduced. As with the full DCW scenario results, there is also a slight discrepancy between the rectangular-windowed detector and the SVA detector due to SVA algorithmic noise.

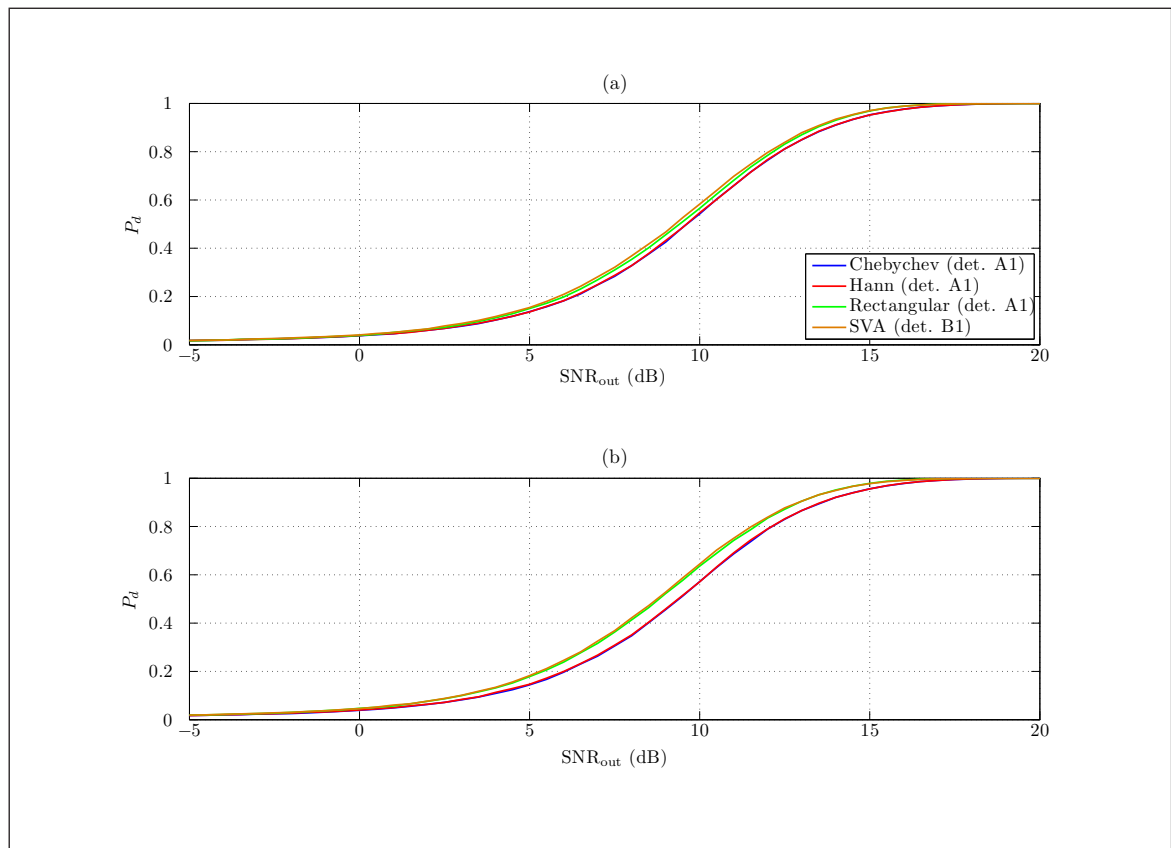


Figure 3.3.2 The simulated detector ROC curves of SNR_{out} against P_d for: (a) the non-zero padded FFT; (b) the zero padded FFT

As an explanation of why the performance gap between the different detectors was smaller in this scenario when compared against the full DCW scenario results, it is necessary to consider the many instances when the signal pulse was only partially present in a DCW. In this instance the rectangular-windowed and SVA detectors included only noise power at one end of the DCW during calculation of spectral estimates. Meanwhile the Hann and Chebychev window functions attenuated both ends of the DCW and hence effectively discarded noise samples at one end of the DCW. This contrasts with the full DCW scenario in section 3.2 where the rectangular-window and SVA had the advantage of the inclusion of signal power from both ends of the DCW. However in this scenario, in most trials, they only included signal power from one end. The overall effect was that the Hann-windowed and Chebychev-windowed detectors were able to close the performance gap in comparison to the full DCW results in section 3.2.

It should be noted that this randomly-filled DCW scenario is an **extreme case** because in this scenario the shortest possible signal pulse duration was used for which the detectors were designed. This scenario tested detectors A1 and B1 to their design

limits. In real-life, radar signal pulses tend to be longer in duration than the duration of a DCW in a DRx in an ESM system.

The scenario in section 3.2 favoured the use of the rectangular-windowed detector and also the SVA detector over the Hann-windowed and Chebychev-windowed detectors because the rectangular window function and SVA maximised the amount of possible signal power that could be sent to the FFT block and then onto the detector block. This was also the case, although only marginally so, for the scenario in this section.

This scenario and the scenario in section 3.2 monitored a frequency bin of interest and did not consider the output of other frequency bins. Therefore the effect of high sidelobes caused by the use of the rectangular window function was irrelevant. This is not an effect which can be ignored in real-life DRx signal processing and is the main reason why the rectangular window function is seldom chosen for use in detectors. The effect of sidelobes is considered in the following two sections.

In the partially-filled DCW and full DCW scenarios, the SVA detector performed almost as well as the rectangular-windowed detector because the SVA technique causes a distortionless response at the frequency bin of interest. The slight detector performance discrepancy between the SVA detector (i.e. detector B1) and rectangular-windowed detector was caused by the algorithmic noise from the estimation of the α parameter from the data.

The situations described in the next two scenarios deal with more complicated situations which highlight the advantages of the use of SVA over the rectangular window function by no longer ignoring the effects of high sidelobes.

3.4 Interference in Full DCW

3.4.1 Scenario and Simulation Description

A third scenario was tested where more than one radar signal was present against a background of AWGN. In this scenario the monitored bandwidth was extended from the previous scenarios in section 3.2 and section 3.3. This multiple signal situation was simulated by two sinusoids of unknown carrier frequencies ($f_{T,A}$ and $f_{T,B}$) present throughout the entire DCW to represent two radar signals from the sinusoid category from section 1.8. In this scenario both signals were full DCW. In all interference scenarios the H_0 case was AWGN plus interference present and the H_1 case was signal plus AWGN plus interference present.

The sinusoids were complex exponentials and whilst their exact frequencies were chosen at random from a frequency bin width interval (f_s/N_{FFT}), the separation between them in frequency was, on average, a particular multiple of f_s/N . The signals were tested at different powers and as such the power ratio between them was defined as the signal-to-interference ratio (SIR). SIR_{out} is used to keep consistent notation, although $\text{SIR}_{\text{out}} = \text{SIR}_{\text{in}}$. The simulated detectors were set up to detect one of the sinusoids (the desired, S_A) in the presence of the other sinusoid (the interference, S_B). Where $\text{SNR}_{\text{out},B}$ is mentioned in this section it refers to the ratio of powers between S_B and the noise floor both of which were held constant throughout the simulations. For clarity, it was the power of S_A which was varied to produce the results.

The simulated detector architectures used were those shown in Figs 3.4.1 and 3.4.2. In this work these are referred to as **detectors C1** and **D1**. Detectors C1 and D1 more closely resemble the process stages in a real-life DRx in an ESM system than detectors A1 and B1. The additional stages over detectors A1 and B1 are related to the estimation of the effect of sidelobe spectral leakage and mainlobe width caused by a strong signal which impacts on other frequency bins. A real-life DRx might not even be as complicated as the architectures shown for detectors C1 and D1 and may only include a few rough decision rules which attempt to prevent sidelobe spectral leakage or broad mainlobes causing erroneous detections in adjacent frequency bins. For example a real-life detector might contain a rule which states that if a rectangular window function is used no further detections are allowed in neighbouring bins until the sidelobe level drops 13 dB.

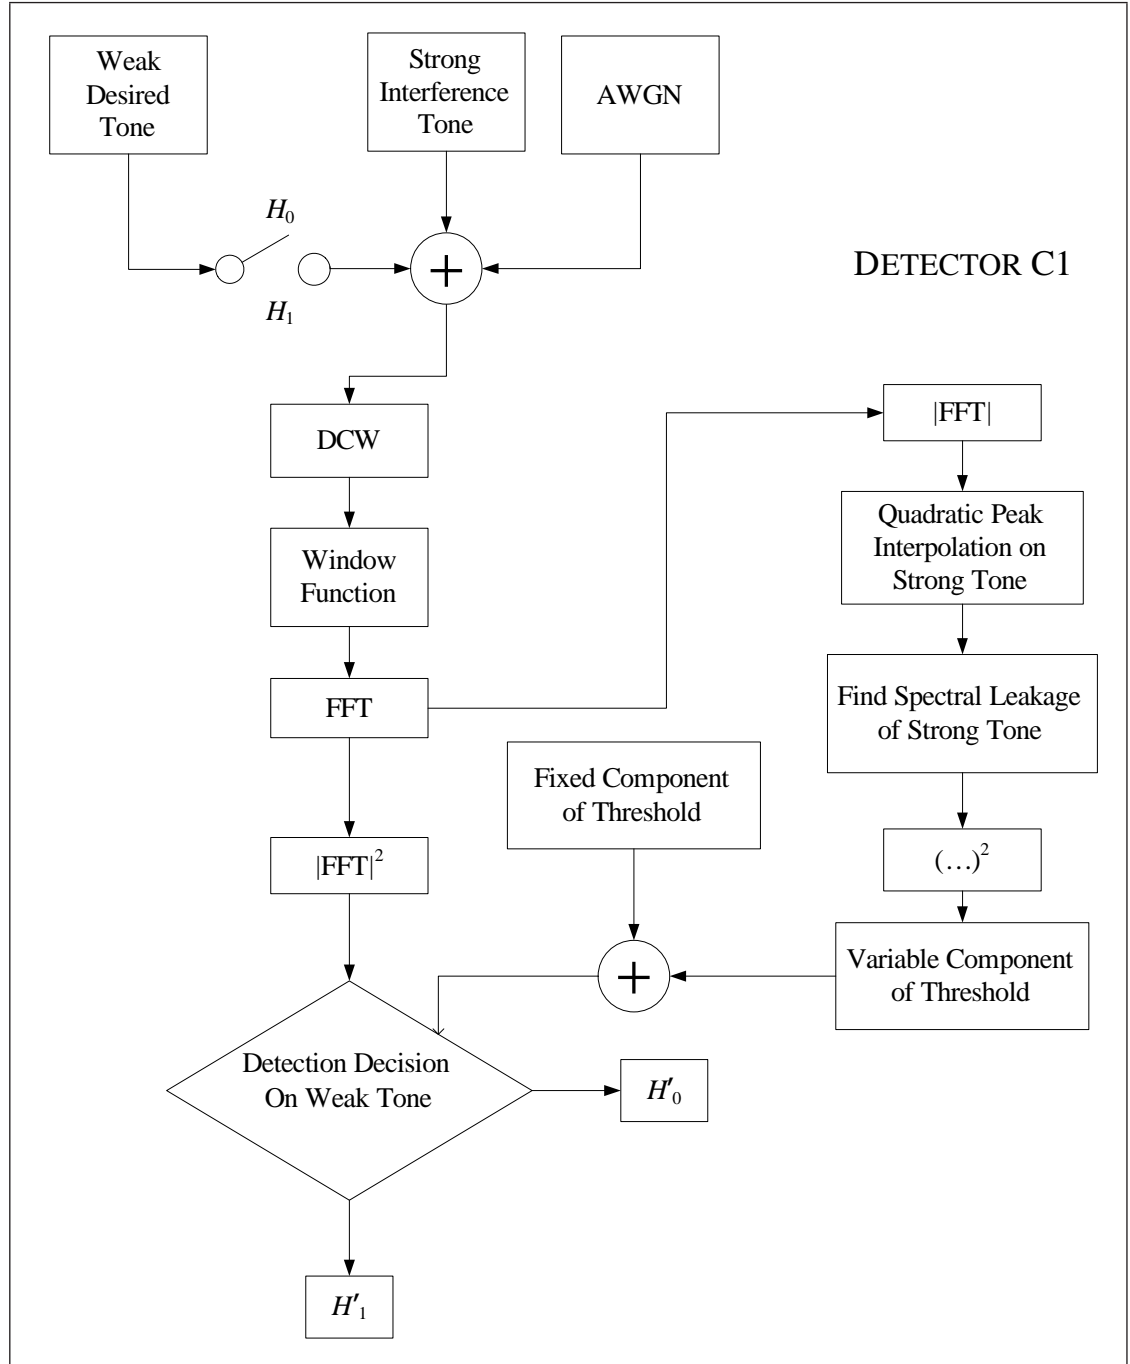


Figure 3.4.1 Detector C1 flowchart describing the structure of the rectangular-windowed, Hann-windowed and Chebbychev-windowed detectors

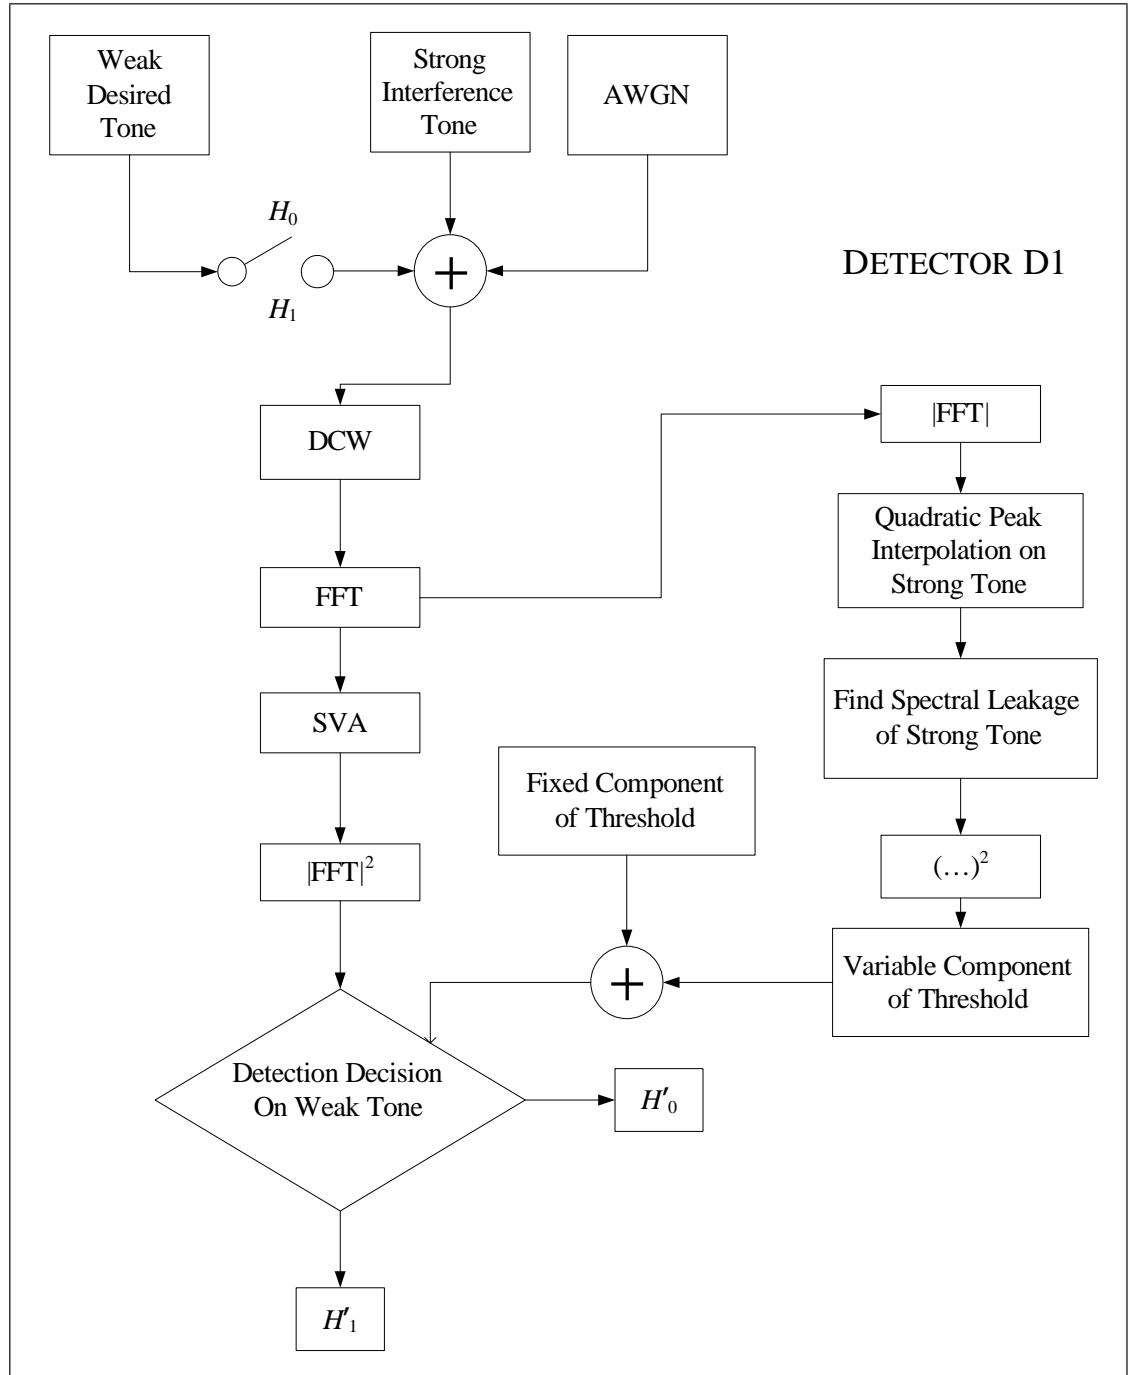


Figure 3.4.2 Detector D1 flowchart describing the structure of the SVA detector

In this scenario, the different stages of how the detectors estimated the spectral leakage from an interfering signal are illustrated in Fig 3.4.3. These stages are described here. The first stage in the calculation of the spectral leakage was to identify the interference signal S_B , which is easily done in Fig 3.4.3 as it is the largest spectral peak. The peak calculated could suffer from a cusping loss because the true frequency ($f_{T,B}$) was unknown. Therefore the simulated detectors attempted to account for this by interpolation for the true peak of S_B where possible. This concept is demonstrated with the red peak in Fig 3.4.3. The estimated true peak was used to create a spectral leakage envelope around the peak. An envelope as opposed to the actual spectral leakage pattern was used because the estimated true peak was only an estimate and therefore there was no certainty of the exact location of the sidelobe peaks and troughs. A demonstration of the spectral leakage envelope for a Hann window is shown in Fig 3.4.3 as a green envelope. Fig 3.4.3 shows how this envelope was used to form a variable component of threshold, which was used in the detectors to attempt to detect the desired signal S_A . The other part of the threshold was derived to ensure a certain false alarm rate above the noise floor and was thus a fixed component of threshold. In the example in Fig 3.4.3, the desired signal is at the normalised frequency of $0.3 \text{ rad}/\pi$ and it would be detected because it is greater than the sum of the variable component of threshold and the fixed component of threshold. Theoretical formulas

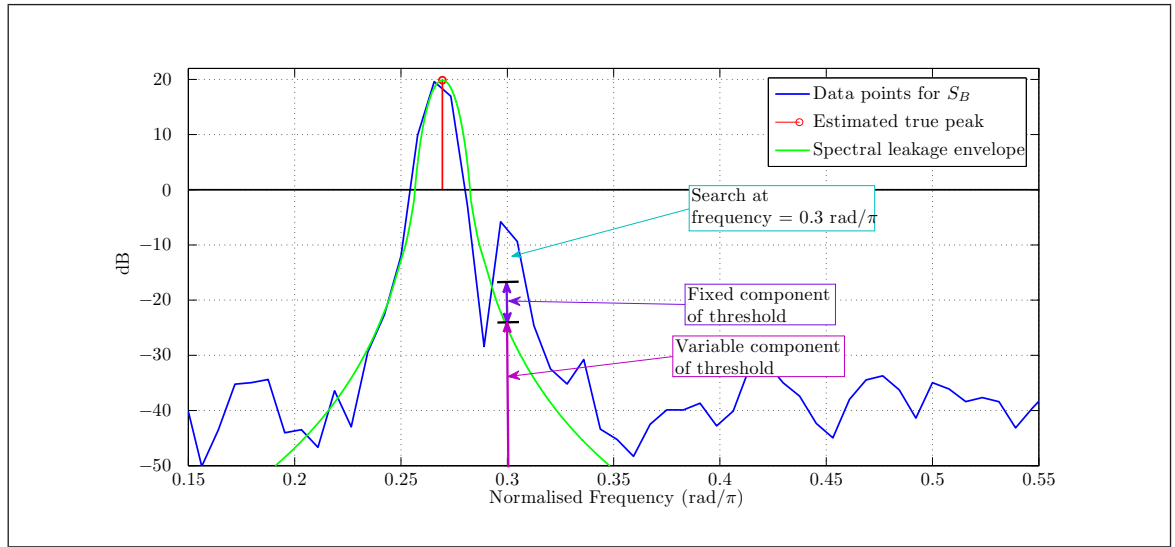


Figure 3.4.3 How spectral leakage is estimated in the simulated detectors

were used initially to calculate the spectral leakage envelopes for the different detector C1 window functions. The spectral leakage envelopes were oversampled ($N_{ov} \gg N$) in the frequency-domain and stored in detector C1 such that it was unnecessary to calculate the envelopes for each trial. The basis for the rectangular window was the Dirchlet (aliased-sinc) function

$$W(k) = \left| \frac{\sin\left(\frac{\pi f_{T,B} k}{f_s}\right)}{N_{ov} \sin\left(\frac{\pi f_{T,B} k}{f_s N_{ov}}\right)} \right|, \quad k = 0, 1, \dots, N_{ov} - 1 \quad (3.4.1)$$

The basis for the Hann window spectral envelope was the sum of three Dirchlet functions

$$W(k) = \left| \frac{\sin\left(\frac{\pi f_{T,B}k}{f_s}\right)}{2N_{ov} \sin\left(\frac{\pi f_{T,B}k}{f_s N_{ov}}\right)} + \frac{\sin\left(\frac{\pi f_{T,B}(k-1)}{f_s}\right)}{4N_{ov} \sin\left(\frac{\pi f_{T,B}(k-1)}{f_s N_{ov}}\right)} + \frac{\sin\left(\frac{\pi f_{T,B}(k+1)}{f_s}\right)}{4N_{ov} \sin\left(\frac{\pi f_{T,B}(k+1)}{f_s N_{ov}}\right)} \right|, \quad (3.4.2)$$

$$k = 0, 1, \dots, N_{ov} - 1$$

The basis for the Dolph-Chebyshev spectral envelope was

$$W(k) = \left| \frac{\cos\left\{N_{ov} \arccos\left[b \cos\left(\frac{\pi k}{N_{ov}}\right)\right]\right\}}{\cosh[N_{ov} \arccos(b)]} \right|, \quad (3.4.3)$$

where $b = \cosh\left[\frac{1}{N_{ov}} \operatorname{arccosh}(10^a)\right]$
and $k = 0, 1, \dots, N_{ov} - 1$

and where $a = 3$ was selected to create sidelobe levels of -60 dBc.

A quick comment is required on the interpolation method that was used in this scenario. It was quadratic peak interpolation and the reason for the choice was one of simplicity. As always, a real-time DRx would struggle with more complicated interpolation methods. To zeropad is the most accurate method of interpolation and this has been used in the previous scenarios in section 3.2 and 3.3. The simulations in this scenario attempted quadratic peak interpolation with and without zeropadding. The zeropadding served to make the quadratic peak interpolation more accurate. Quadratic peak interpolation in a zeropadded FFT, or quadratic interpolated FFT (QIFFT), is popular in audio applications due to its simplicity and accuracy.

The use of quadratic interpolation is described here. It required the frequency bin and magnitude (f_q and q) of the maximum magnitude peak of the interference signal in the FFT output data. The frequencies ($f_q - \frac{f_s}{N_{FFT}}$ and $f_q + \frac{f_s}{N_{FFT}}$) and magnitudes (p and r) of the two surrounding nearest neighbour frequency bins were also required. The condition that, of these magnitude samples, the main central frequency bin magnitude peak was the largest, was also necessary ($q > p$ and $q > r$). In this case, the interpolated peak was located in the frequency-domain at the frequency given by

$$f_{\text{true}} = f_q + \Delta f$$

where Δf is in the interval $[-0.5, 0.5] \times f_s/N_{FFT}$ and was estimated using the following equation

$$\Delta f = \frac{p - r}{2(p - 2q + r)} \quad (3.4.4)$$

The estimate for the magnitude of the true peak (Q) at f_{true} is given by

$$Q = q - \frac{\Delta f}{4} (p - r) \quad (3.4.5)$$

In [124], quadratic peak interpolation is described as the most successful technique,

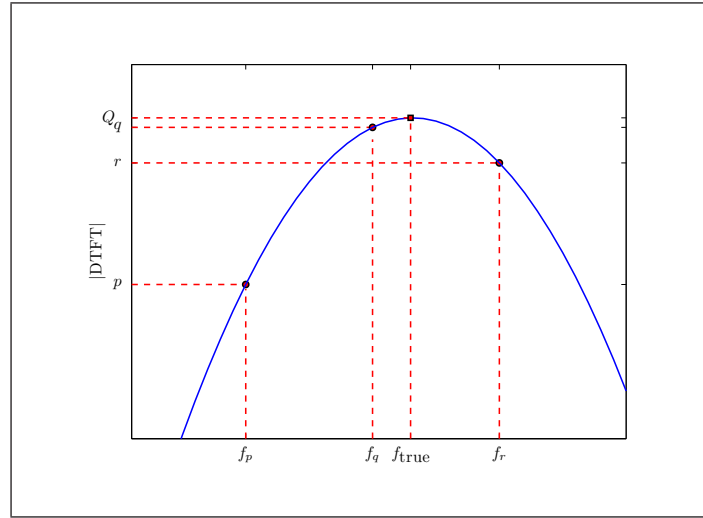


Figure 3.4.4 Quadratic peak interpolation

notwithstanding zeropadding, in low noise environments and in situations where multiple signal carrier frequencies are present and are not closely spaced in the frequency dimension. It is also described as most accurate when used over shorter distances between the three frequency-domain samples (f_p , f_q and f_r).

Obviously the calculation of the spectral leakage from one signal onto another used here was not entirely accurate as it ignored the interaction between the phases of the signals. It was an acceptable approximation however because it would only ever overestimate the variable component of threshold. This consideration is another example of trading-off accuracy so that rapid detection decisions can be made in real-time DRxs.

The other features of the simulations were otherwise the same as in section 3.2 and section 3.3.

3.4.2 Results and Discussion

ROC curve results of SIR_{out} against P_d at a constant $\text{SNR}_{\text{out},B} = 8 \text{ dB}$ for an average frequency separation of $2f_s/N$ are shown in Fig 3.4.5 and for an average frequency separation of $5f_s/N$ are shown in Fig 3.4.6.

The $\text{SNR}_{\text{out},B}$ was increased to 15 dB whilst the separations were kept the same to produce the ROC curves in Figs 3.4.8 and 3.4.7. The $\text{SNR}_{\text{out},B}$ was further increased to 25 dB to create the results in Figs 3.4.9 and 3.4.10. Every point in the ROC curves was calculated from 1×10^5 detection decisions within the simulated detectors and the probability of false alarm was set at $P_{fa} = 0.01$. Panel (a) in all figures shows the

results when a non-zero-padded FFT was used and panel (b) in all figures shows the results when a twice-zero-padded FFT was used in the detectors.

The reason for undertaking simulations at three different $\text{SNR}_{\text{out},B}$ levels was to understand detector performances in low-, medium- and high- $\text{SNR}_{\text{out},B}$ regimes. The $\text{SNR}_{\text{out},B}$ used to produce the results in Figs 3.4.5 and 3.4.6, was low at 8 dB. Therefore the noise floor was the dominant factor which limited the performance of the detectors and the sidelobe spectral leakage and mainlobe width from S_B were less significant.

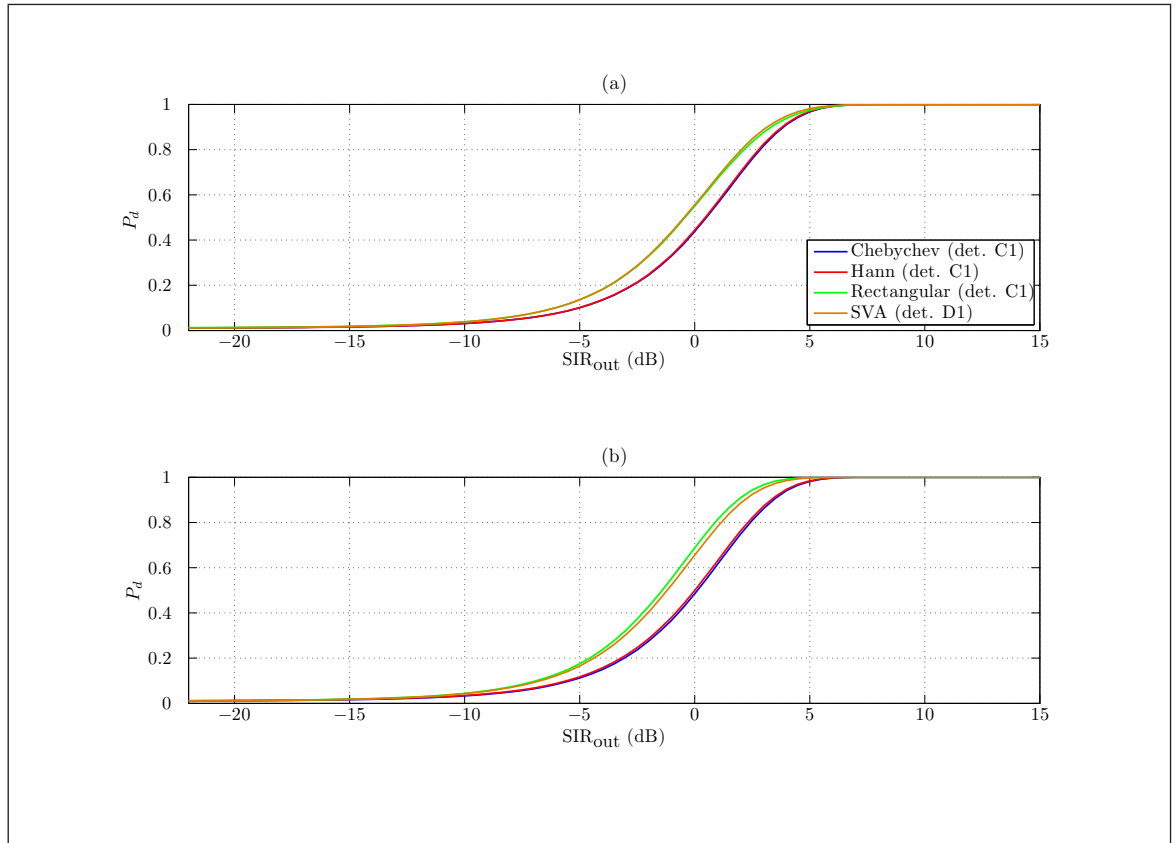


Figure 3.4.5 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 8$ dB with an average separation of $2f_s/N$ for: (a) the non-zero-padded case; (b) the zero-padded case

At $\text{SNR}_{\text{out},B} = 8$ dB, the noise floor was dominant both close to the mainlobe of S_B , which equates to the average separation of $2f_s/N$ test, and further away from the mainlobe of S_B , which equates to the average separation of $5f_s/N$ test. This explains why the results at $\text{SNR}_{\text{out},B} = 8$ dB in Figs 3.4.5 and 3.4.6 hardly differ in the two different average frequency separation cases. As the noise floor was the dominant factor, the window loss was the principal effect in determining the performance of the different detectors. From panel (a) of both Fig 3.4.5 and Fig 3.4.6 the rectangular-windowed detector and SVA detector (i.e. detector D1) are shown to have a performance advantage of ~ 1 dB over the Hann-windowed and Chebychev-windowed detectors. In panel (b) of Figs 3.4.5 and 3.4.6 this advantage grew to ~ 1.5 dB. All the detector performances improved when they utilised a twice-zero-padded FFT because this improved their spectral peak estimation. This is evident in the changes in detector performances

from panel (a) to panel (b) in both Figs 3.4.5 and 3.4.6. The results at both separations were almost as if S_A was the only signal present and the results resemble those from the full DCW scenario in Fig 3.2.6. In this low- $\text{SNR}_{\text{out},B}$ regime, the results imply, either the rectangular-windowed detector or the SVA detector should be used in a real-life DRx as they do not suffer from the window loss. However the rectangular-windowed detector would need to be used with caution because there would be no guarantee that for an all-purpose DRx the regime would always be low- $\text{SNR}_{\text{out},B}$.

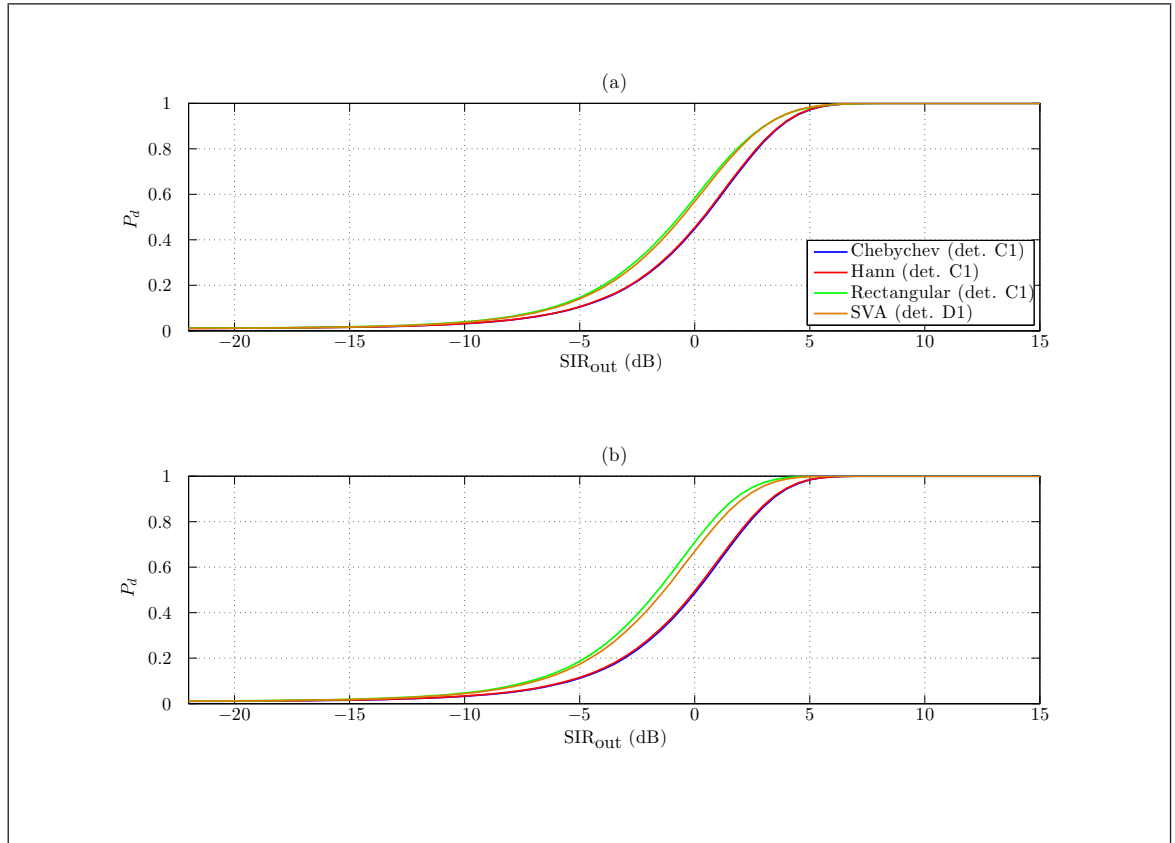


Figure 3.4.6 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 8$ dB with an average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

When the $\text{SNR}_{\text{out},B}$ was raised to 15 dB, the noise floor continued to be dominant in further apart signals. The results to produce Fig 3.4.7 were from signals which were separated by an average of $5f_s/N$. At this relatively large separation and at $\text{SNR}_{\text{out},B} = 15$ dB, the spectral leakage from S_B into S_A was negligible compared to the noise level. Therefore, in this case, the principal effect which limited the detector performances was the window loss. This explains why the rectangular-windowed detector and SVA detector had performances ~ 1 dB better than the Hann-windowed and Chebychev-windowed detectors in panel (a) of Fig 3.4.7 and ~ 1.5 dB better in panel (b) of Fig 3.4.7.

However the results in Fig 3.4.8 were produced from signals separated by an average of $2f_s/N$ and show the level of spectral leakage became important at closely-spaced frequencies. The results displayed in panel (a) show that the SVA detector had a

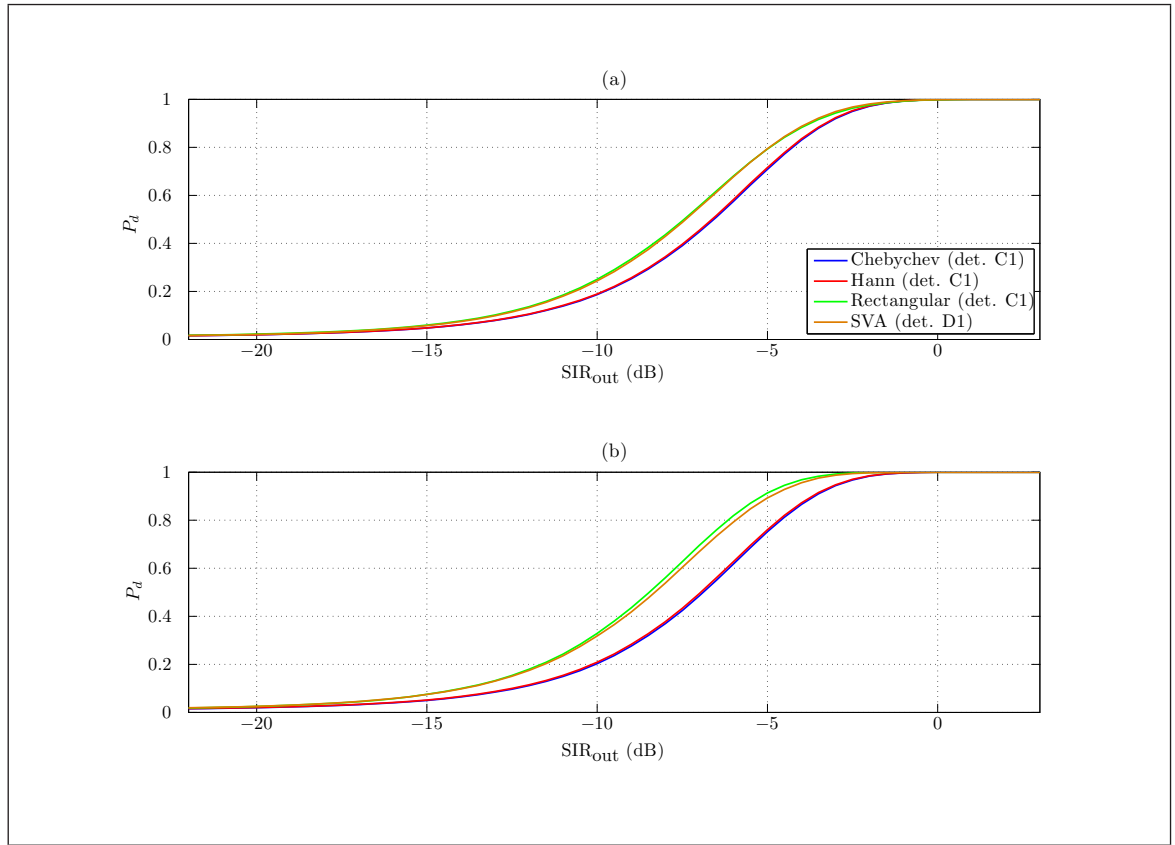


Figure 3.4.7 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 15 \text{ dB}$ with an average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

better performance than the other detectors by $\sim 0.8 \text{ dB}$. This was due to the sidelobe spectral leakage and mainlobe width minimisation properties of SVA and therefore the performance of the SVA detector was only limited by the noise floor. At the same time the rectangular-windowed detector suffered greatly from sidelobe spectral leakage which explains the fall in its performance from $\text{SNR}_{\text{out},B} = 8 \text{ dB}$ (panel (a) of Fig 3.4.5) to $\text{SNR}_{\text{out},B} = 15 \text{ dB}$ (panel (b) of Fig 3.4.8). However because detector C1 utilises a QIFFT, the rectangular-windowed detector was able to recover its performance when the FFT was twice zeropadded. This was because the sidelobe spectral leakage level was fairly accurately estimated so that the rectangular-windowed detector became noise floor dominated again. This is shown by panel (b) of Fig 3.4.8, which overall shows that the rectangular-windowed detector and the SVA detector had a performance $\sim 1.3 \text{ dB}$ better than the Hann-windowed and Chebychev-windowed detectors.

In the medium- $\text{SNR}_{\text{out},B}$ regime, the SVA detector was clearly the best option for a non-zeropadded FFT detector.

Figs 3.4.9 and 3.4.10 display the results when the $\text{SNR}_{\text{out},B}$ was increased to 25 dB . From the closely-spaced in frequency plots the merits of the use of SVA should become clear. Fig 3.4.9 displays these results. In these simulations the signals were closely spaced in frequency and there was a significant amount of spectral leakage from the interference tone (S_B) into the desired tone (S_A). The sidelobe spectral leakage and

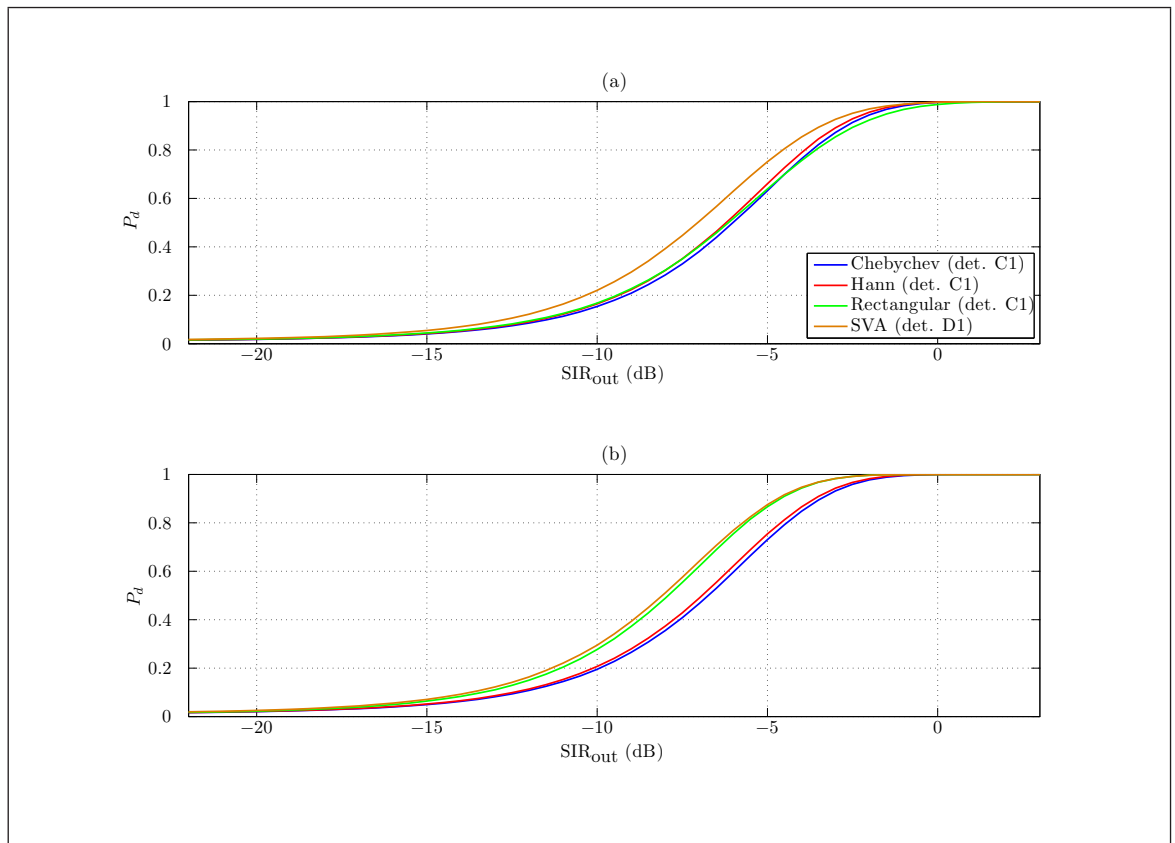


Figure 3.4.8 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 15\text{ dB}$ with an average separation of $2f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

mainlobe width were no longer masked by the noise floor as occurred in the lower $\text{SNR}_{\text{out},B}$ regimes. Spectral leakage was the dominant factor in this high- $\text{SNR}_{\text{out},B}$ and closely-spaced-frequency regime. Therefore SVA, which minimises sidelobe spectral leakage and mainlobe width, is shown as the best detector in both panel (a) and panel (b). Interestingly the results also show that the Hann-windowed detector had a better performance than the Chebychev-windowed detector because the Hann window frequency domain response has a sidelobe roll-off whereas the Chebychev window frequency domain response has constant sidelobes. Therefore there was scope for the Hann-windowed detector to minimise spectral leakage more so than the Chebychev-windowed detector. The QIFFT-based detector C1 enabled the rectangular-windowed, Hann-windowed and Chebychev-windowed detectors to close the performance gap when the DCWs were zeropadded but even the QIFFT-based detector C1 could not completely mitigate such large levels of spectral leakage.

This is evident as a closure of the gap from panel (a) to panel (b) of Fig 3.4.9. For the simulations which produced the results in Fig 3.4.10 the frequency separation between the signals was greater and hence the spectral leakage from S_B into S_A was smaller. However the results from panel (a) to panel (b) show that the rectangular-windowed detector, which was obviously suffering due to large amounts of spectral leakage, was able to recover somewhat when a zeropadded FFT was used. This was because detector

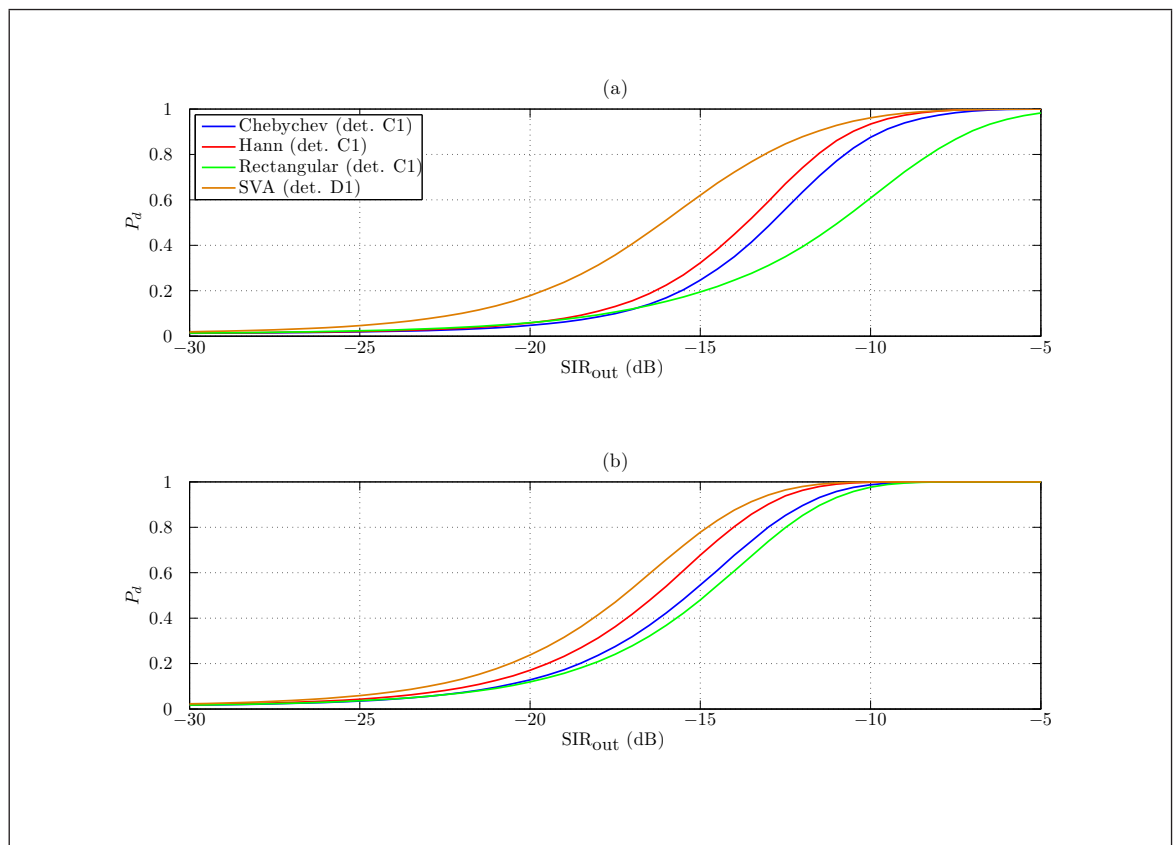


Figure 3.4.9 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 25\text{ dB}$ with an average separation of $2f_s/N$ for: (a) the non-zero padded case; (b) the zero padded case

C1 was able to account for the spectral leakage better when it was zero padded. As a result panel (b) starts to resemble the familiar full DCW scenario results, where only the window loss is the significant effect. In panel (a) of Fig 3.4.10 the Hann-windowed and Chebychev-windowed detectors appear to match the performance of the SVA detector. This was because the high- $\text{SNR}_{\text{out},B}$ regime was spectral-leakage-dominated even at larger separations and with reference to panel (a) of Fig 2.4.2 the sidelobe reduction levels of SVA and other bell-shaped windows are quite similar. This meant that the performances of the Hann-windowed, Chebychev-windowed and SVA detectors were similar.

In the high- $\text{SNR}_{\text{out},B}$ regime, the SVA detector was clearly the best option to use in an all-round detector.

Since the spectral leakage caused by the SVA technique is in actual fact extremely low, a different, simpler detector architecture was attempted for the SVA technique to be compared against the results from the more complicated detector D1. This new detector was called **detector E1** and is shown in Fig 3.4.11. The core difference over detector D1 is the removal of the variable threshold and associated calculations.

A comparison of the results from detector D1 and the simpler detector E1 are shown in Fig 3.4.12. The left column shows the results for the nominal separation of $2f_s/N$

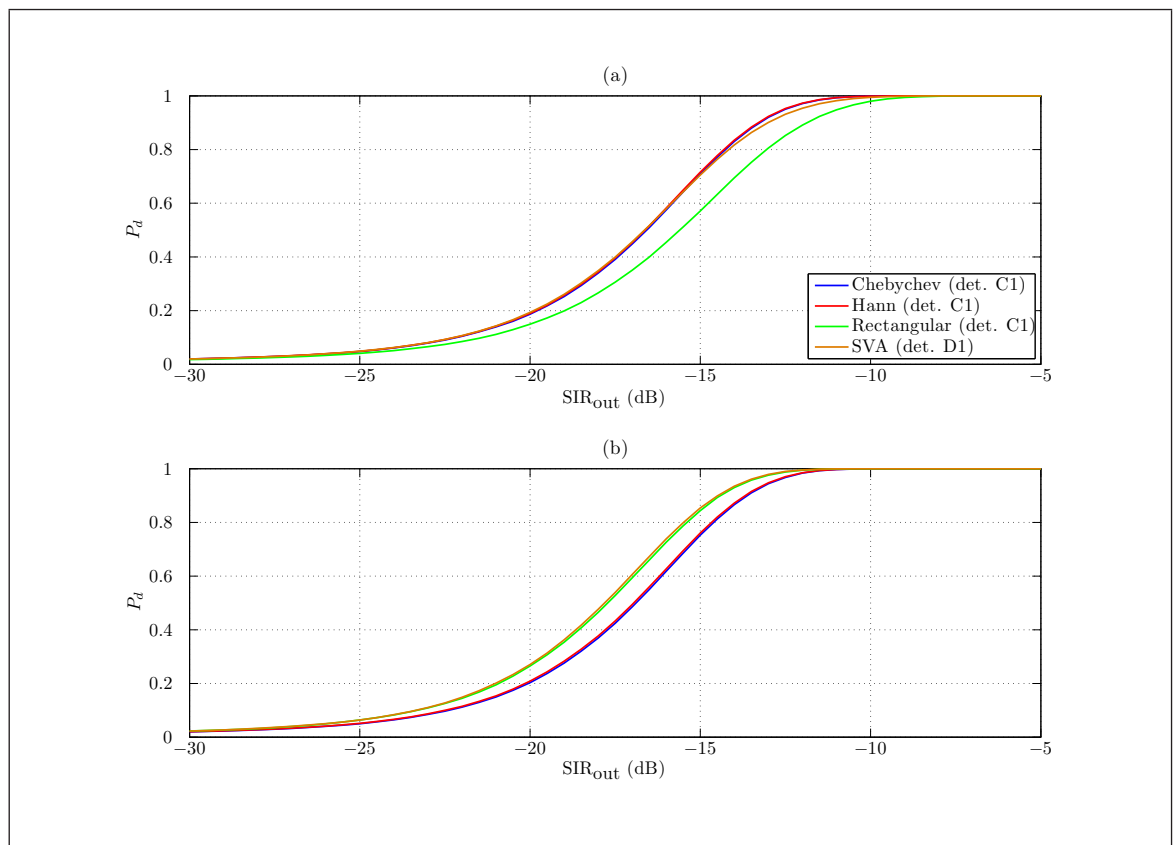


Figure 3.4.10 Simulated detector ROC curves for the full DCW interference scenario at $\text{SNR}_{\text{out},B} = 25\text{ dB}$ with an average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

and the right column shows the results for the nominal separation of $5f_s/N$.

The results in Fig 3.4.12 indicate that there was little to be gained from the variable threshold calculations in the SVA detector. This was due to the negligible amount of spectral leakage caused by the SVA technique. There would be no point in the calculation of the variable component in a real-life DRx which used SVA as the variable component is a small number and likely to be below the quantisation noise in most cases. Therefore a decision was made that all subsequent models which utilised SVA in the MBD methodology would omit the calculation of the variable component of threshold for SVA and be based on the detector E1 architecture. The results in Fig 3.4.12 also have another important implication for a real-life DRx. A real-life DRx often does not have the luxury of the ability to calculate an accurate variable component of threshold from spectral leakage patterns in the time with which they are provided. This is the reason why many simply use decision rules which are approximations to spectral leakage patterns. However as the spectral leakage caused by SVA is negligible, decision rules are no longer necessary and therefore computational savings can be obtained from their omission in a real-life DRx. In fact, this point was hinted at in [117], which states that “the SVA [technique] reduces the sidelobes until it reaches the noise floor”, which is equivalent to “the spectral leakage caused by SVA is negligible.”

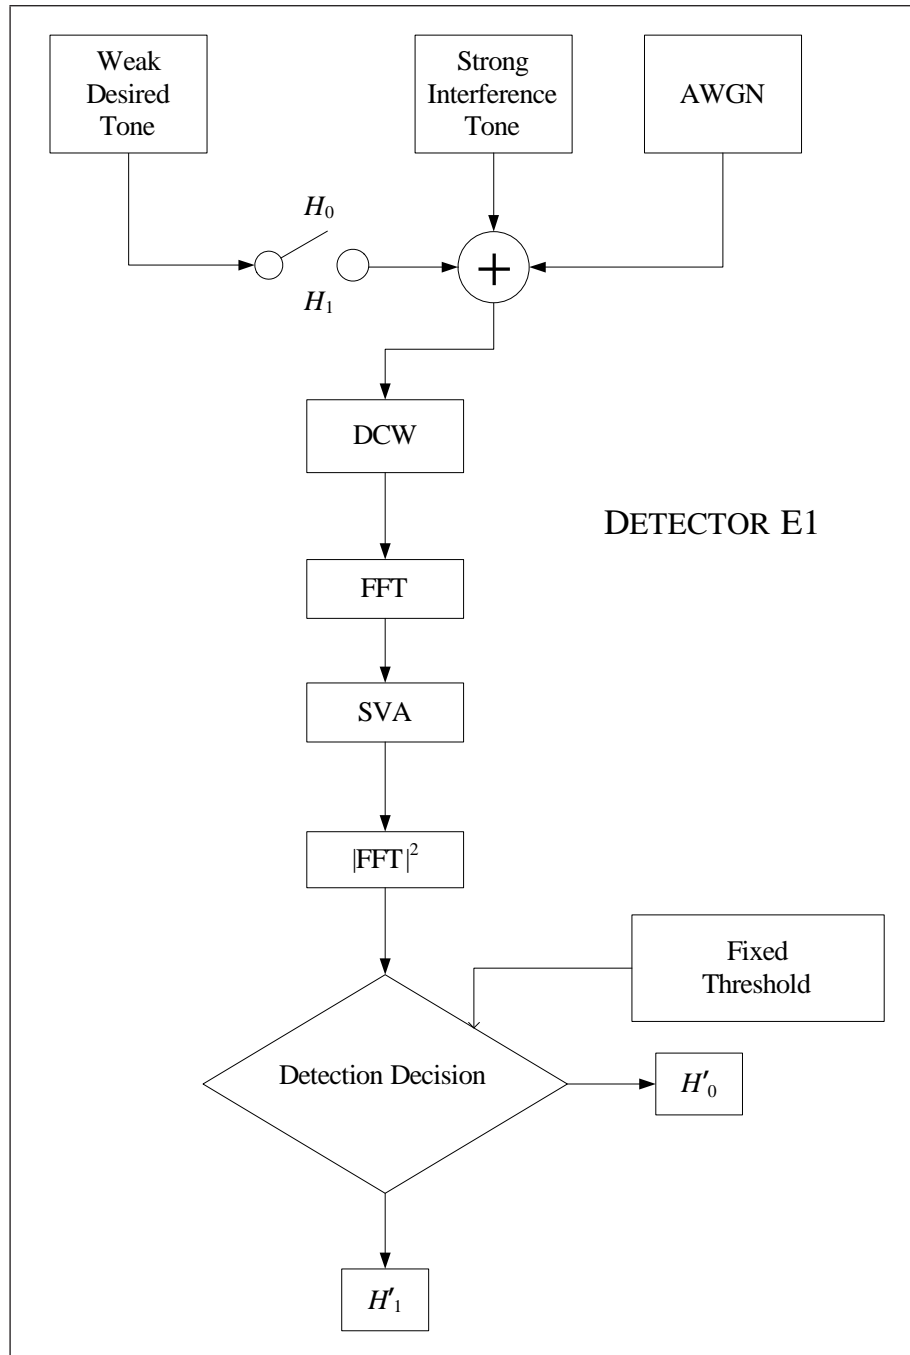


Figure 3.4.11 Detector E1 is a simple detector which assumes no spectral leakage from the use of SVA

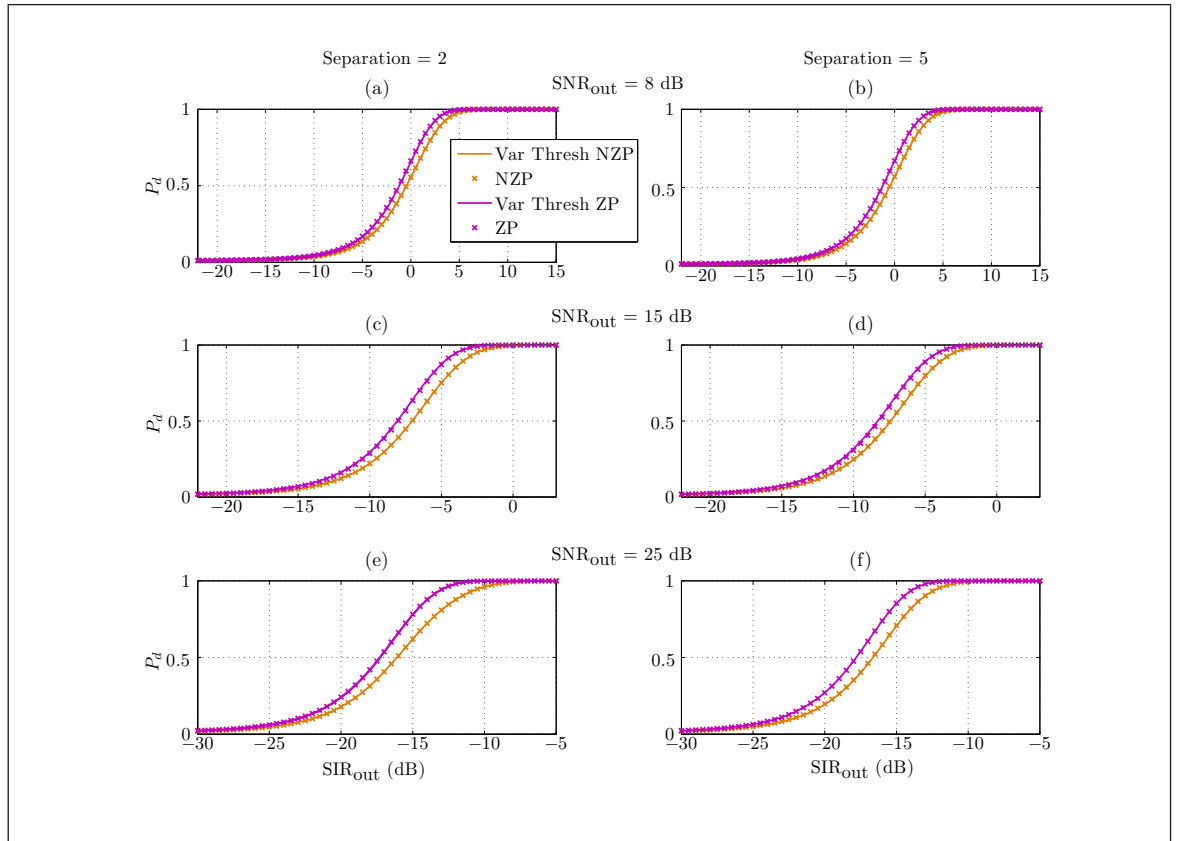


Figure 3.4.12 Comparison of the SVA results from detectors D1 (Var Thresh) and E1. Both non-zeropadded (NZP) and zeropadded (ZP) results are shown on the same graphs for: (a) $\text{SNR}_{\text{out},B} = 8 \text{ dB}$, separation = $2f_s/N$; (b) $\text{SNR}_{\text{out},B} = 8 \text{ dB}$, separation = $5f_s/N$; (c) $\text{SNR}_{\text{out},B} = 15 \text{ dB}$, separation = $2f_s/N$; (d) $\text{SNR}_{\text{out},B} = 15 \text{ dB}$, separation = $5f_s/N$; (e) $\text{SNR}_{\text{out},B} = 25 \text{ dB}$, separation = $2f_s/N$; (f) $\text{SNR}_{\text{out},B} = 25 \text{ dB}$, separation = $5f_s/N$

Overall the SVA technique is expected to provide a performance improvement when used in a DRx as part of an ESM system in situations where there are multiple sinusoidal signals present in the DCW because SVA minimises spectral leakage from signals into neighbouring frequency bins. In particular SVA is expected to:

- have an advantage over moderate dynamic range window functions such as the Hann or Chebychev windows when signals are closely spaced in frequency and are of a similar amplitude;
- have an advantage over low dynamic range window functions such as the rectangular window when signals are of disparate amplitude.

The interference scenario detailed in this section suited the problem that SVA was designed to tackle. However this scenario, that the DCW was always full of signal power, may not always be the case as random TOAs can produce partially-filled DCWs. To consider this scenario, section 3.5 tested the detectors against a scenario where a random proportion of both the signals was captured in a DCW.

3.4.3 5-point and 7-point SVA Comparison

The end of section 2.5 declared that further lines of reason would be provided against the use of 5-point and 7-point SVA over the simpler (joint 3-point) SVA in a detection role. One of these arguments was provided at the end of section 3.2 which showed that the 5-point and 7-point SVA were not as successful as the 3-point SVA in recovering the window loss. The results in Fig 3.4.13 provide another strong case against the use of 5-point and 7-point SVA. It shows the performance of simulated detectors which utilised the 5-point and 7-point SVA techniques in the full DCW interference scenario. Furthermore, the scenario parameters were copied from the case where 3-point SVA was the best performer so that the performance of 5-point SVA and 7-point SVA could be compared against this. These parameters were $\text{SNR}_{\text{out},B} = 25 \text{ dB}$ and a nominal frequency separation of $2f_s/N$ between S_A and S_B , where $\text{SNR}_{\text{out},B}$, S_A and S_B are as defined earlier in the full DCW interference section.

The data from the simulated 3-point SVA detector, Hann-windowed detector, Chebychev-windowed detector and rectangular-windowed detector from Fig 3.4.9 were replicated into Fig 3.4.13. The new 5-point SVA and 7-point SVA ROC data curves derived from the simulation of detector D1 with 5-point or 7-point SVA techniques.

The results demonstrate that the 5-point and 7-point SVA simulated detectors performed badly when presented with closely-spaced in frequency signals in a scenario where spectral leakage was the dominant factor over the noise floor. As before, in the full DCW scenario, the 7-point SVA detector suffered a performance drop relative to the 5-point SVA detector due to extra algorithmic noise. The use of twice-zero-padded

FFT blocks enabled an improvement in performance of all detectors. The main conclusion to draw, however, is that there is no advantage in the implementation of 5-point or 7-point SVA over 3-point SVA in a detection role. This is because 5-point and 7-point SVA do not offer further optimisation over 3-point SVA in spectral leakage-dominated scenarios where it might be expected that they would.

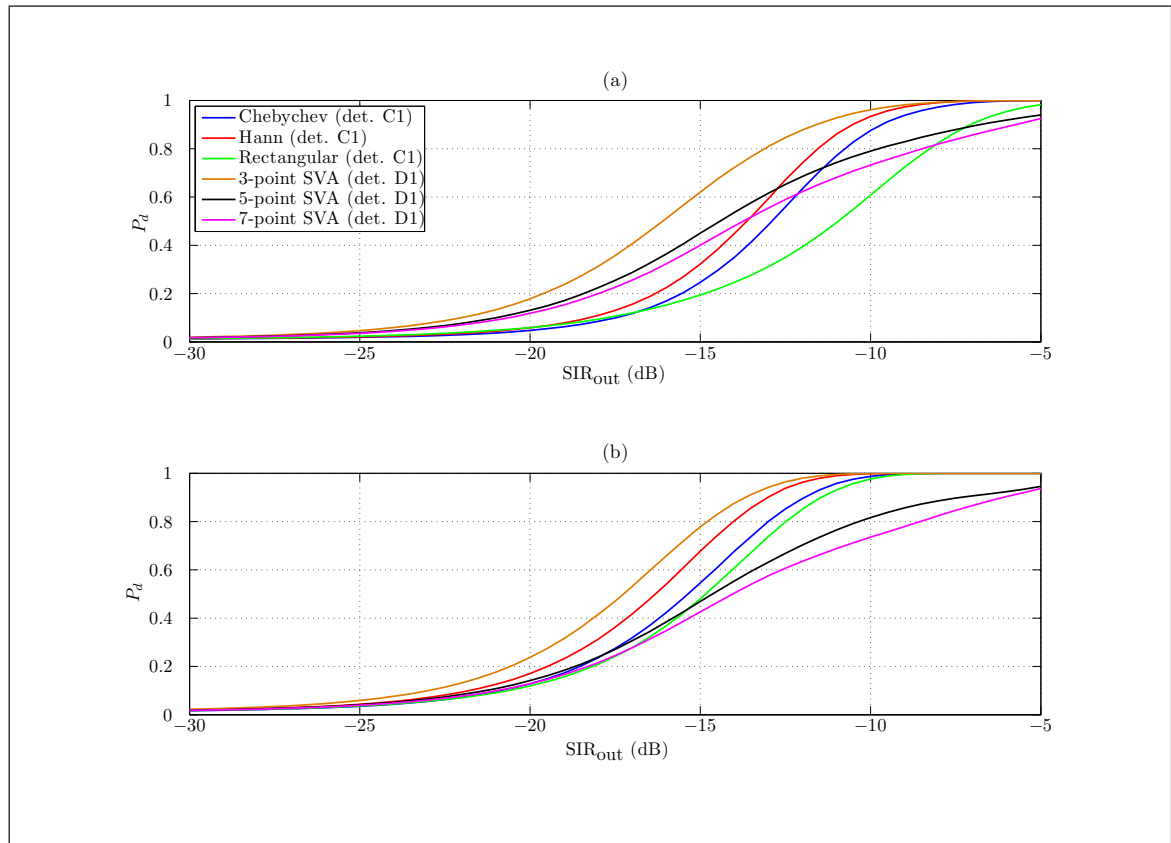


Figure 3.4.13 The results from Fig 3.4.9 alongside results from the use of 5-point and 7-point SVA detectors for: (a) the non-zero padded case; (b) the zero padded case

3.4.4 Constrained Complex SVA Comparison

The interference scenario was also used to compare the constrained complex SVA technique against rectangular-windowed, Hann-windowed, Chebychev-windowed and (joint, 3-point) SVA detectors. The reason for the use of the interference scenario test was to uncover whether it was worth extensively exploring constrained complex SVA over the 3-point SVA algorithm.

The simulation conditions here were mostly the same as those described earlier, specifically:

- The signals were sinusoids in the form of complex exponentials.
- During the simulations both strong and weak signals were sinusoids and were present for the duration of the DCW.

- The actual frequencies ($f_{T,A}$ and $f_{T,B}$) of the signals were random over the interval of f_s/N_{FFT} .
- The initial phase of the signals ($\phi_{0,A}$ and $\phi_{0,B}$) was random.
- The signals were nominally separated by the equivalent of $3.5f_s/N$.

However in this scenario the weak signal was very much buried in the sidelobes of the strong, interference signal. The data flowchart for the rectangular-windowed, Hann-windowed and Chebychev-windowed detectors was detector C1, depicted in Fig 3.4.1. The data flowchart for the SVA detector was detector D1 depicted in Fig 3.4.2.

Detectors C1 and D1 used the variable component and a fixed component to produce a threshold. As in simulations in the interference section, the variable component derived from the envelope of the spectral leakage of the identified strong, interference signal and the fixed component derived from a threshold set to achieve a desired false alarm rate. For an illustration of how the variable component was calculated see Fig 3.4.3 and refer back to the discussion at the start of section 3.4.

The data flowchart for the constrained complex SVA detector is shown in Fig 3.4.15 as **detector F1**. There is no well-defined spectral leakage pattern for the constrained complex SVA technique. Therefore a variable component of threshold was not used. Without a variable component, only a fixed threshold was used which was calculated to produce a constant false alarm rate with the interference signal and AWGN present. The constant false alarm rate was set such that $P_{fa} = 0.01$.

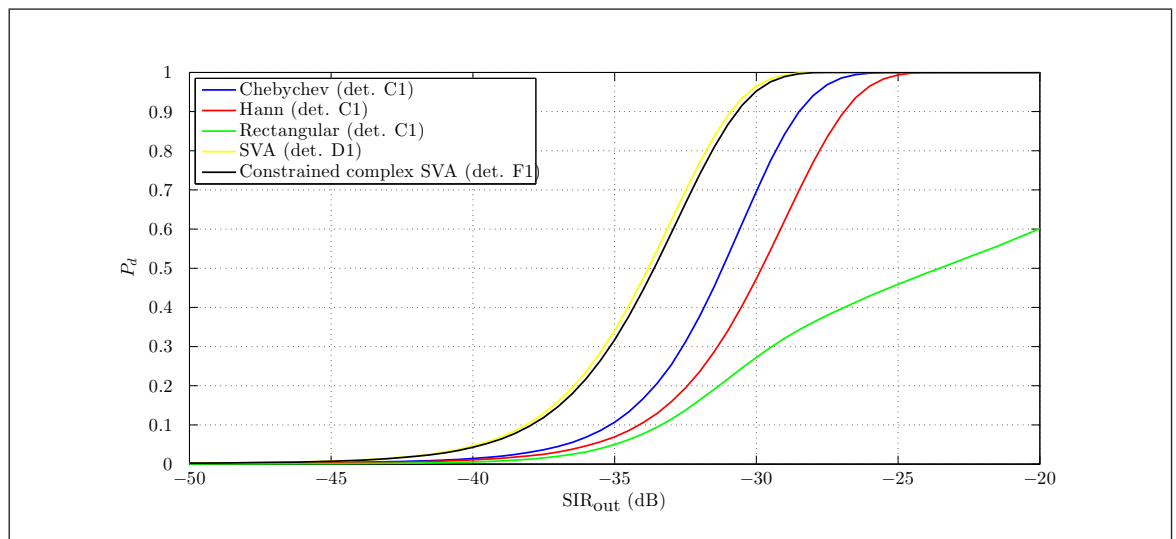


Figure 3.4.14 A comparison of the performance of detectors which use a rectangular window, Hann window, Chebychev window, SVA and constrained complex SVA in a difficult detection scenario

The ROC curve results are shown in Fig 3.4.14. The results show that SVA and constrained complex SVA simulated detectors had the best performance in this difficult scenario. The rectangular-windowed detector performed particularly poorly as a result of large sidelobes through this window function choice.

Ultimately the main conclusion to draw from this set of results is that there would be no advantage to be gained with the use of constrained complex SVA over the computationally faster SVA. Therefore the constrained complex SVA was not studied further as a viable algorithm for possible application in an DRx in an ESM system.

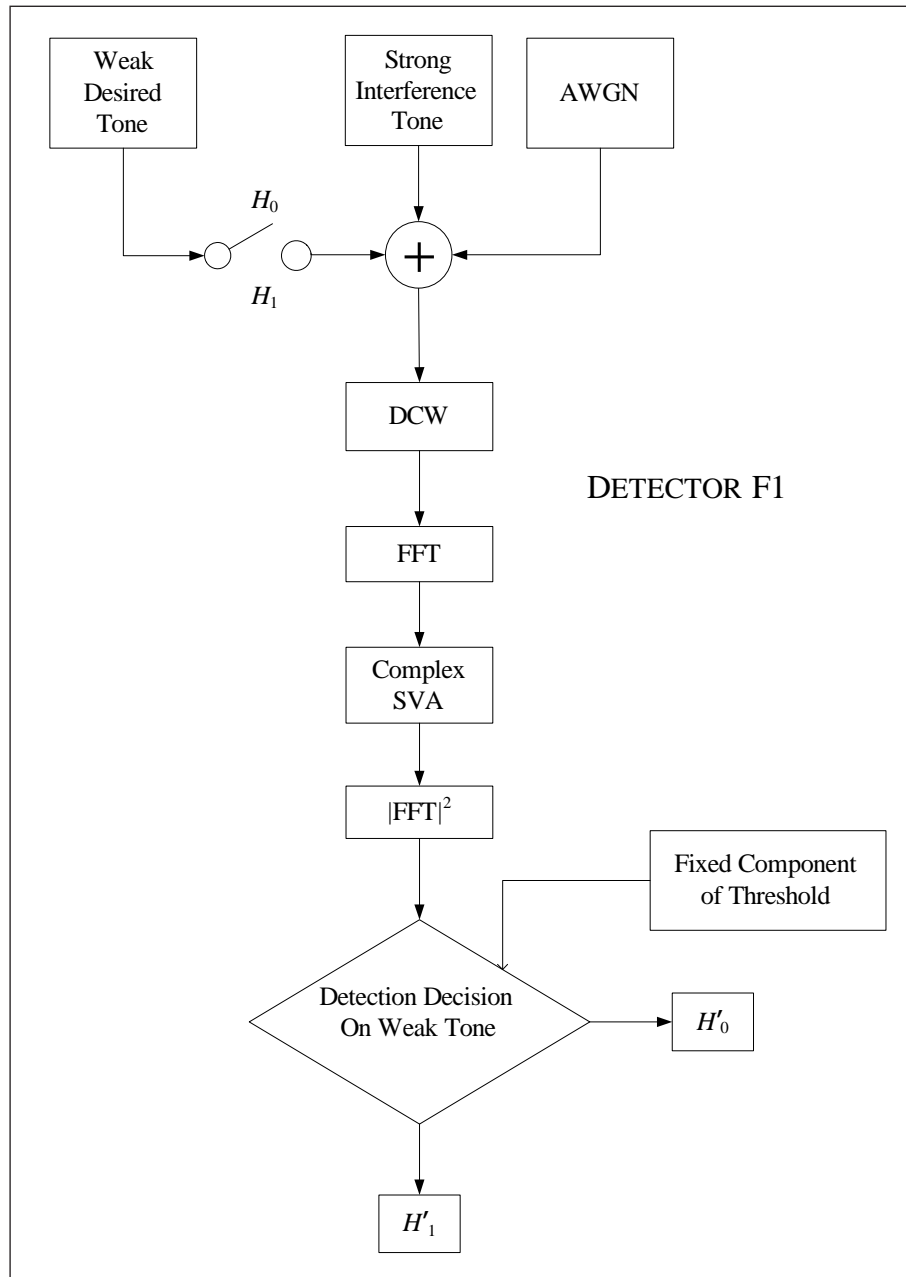


Figure 3.4.15 Detector F1 flowchart describing the structure of the constrained complex SVA detector

3.5 Interference in Randomly-filled DCW

3.5.1 Scenario and Simulation Description

The simulations from section 3.4 indicated that SVA affords a performance improvement in the case of interception of the maximal amount of signal power from two sinusoidal radar signals close in frequency. The next step was to test whether the SVA technique could still perform reasonably well in comparison to conventional window detectors in the case where the DCW contained only part of either of the signal pulses. This conceptual step was exactly the same as the one that was taken from the full DCW scenario in section 3.2 to the randomly-filled DCW scenario in section 3.3. As such, the randomly-filled part of this scenario was arranged in an analogous way to the randomly-filled DCWs in section 3.3 and used the same set of assumptions. The most important aspect of this scenario was that it tested detectors C1 and E1 to their limits through the extreme case that the signal pulses of S_A and S_B were as short in duration as the detector DCW. As mentioned in section 3.3, in the worst case this translated to a half-filled DCW with samples that had a signal contribution. This had important consequences for detector C1 which utilised conventional windows. The accuracy of detector C1 was most affected by the missed S_B signal power from the samples as when the spectral leakage envelope was calculated the mask was no longer entirely accurate and hence the variable component of threshold was no reliable. This happened because absent signal samples in a DCW have a similar effect to zeropadding, in that in the frequency-domain the spectral peak and sidelobe pattern are generally broadened due to a loss of information.

This effect is illustrated for a Hann-windowed detector in Fig 3.5.1. In panel (a) there is a DCW full of signal power from a single signal. When the Hann window function is applied and an FFT taken, the frequency-domain picture shown in panel (b) results. The sidelobe envelope pattern can be predicted with the use of the Hann window function frequency-domain response. This was exactly the mechanism that was exploited in detector C1 in section 3.4. In panel (c) the worst case is shown where the DCW is only half-filled with signal power. After a Hann window function is applied and an FFT taken, the predicted spectral leakage envelope is no longer accurate in panel (d).

The number of missing signal samples in a DCW in this scenario would be unknown and therefore the extent of this effect would also be unknown. This means it would be very difficult to compensate for this effect. Therefore this scenario was simply to test the existing detector C1 architecture through application in a case where it was not entirely accurate.

The interference signal in this scenario was arranged in an analogous way to the interference signal in section 3.4, except the interference was also partially captured in

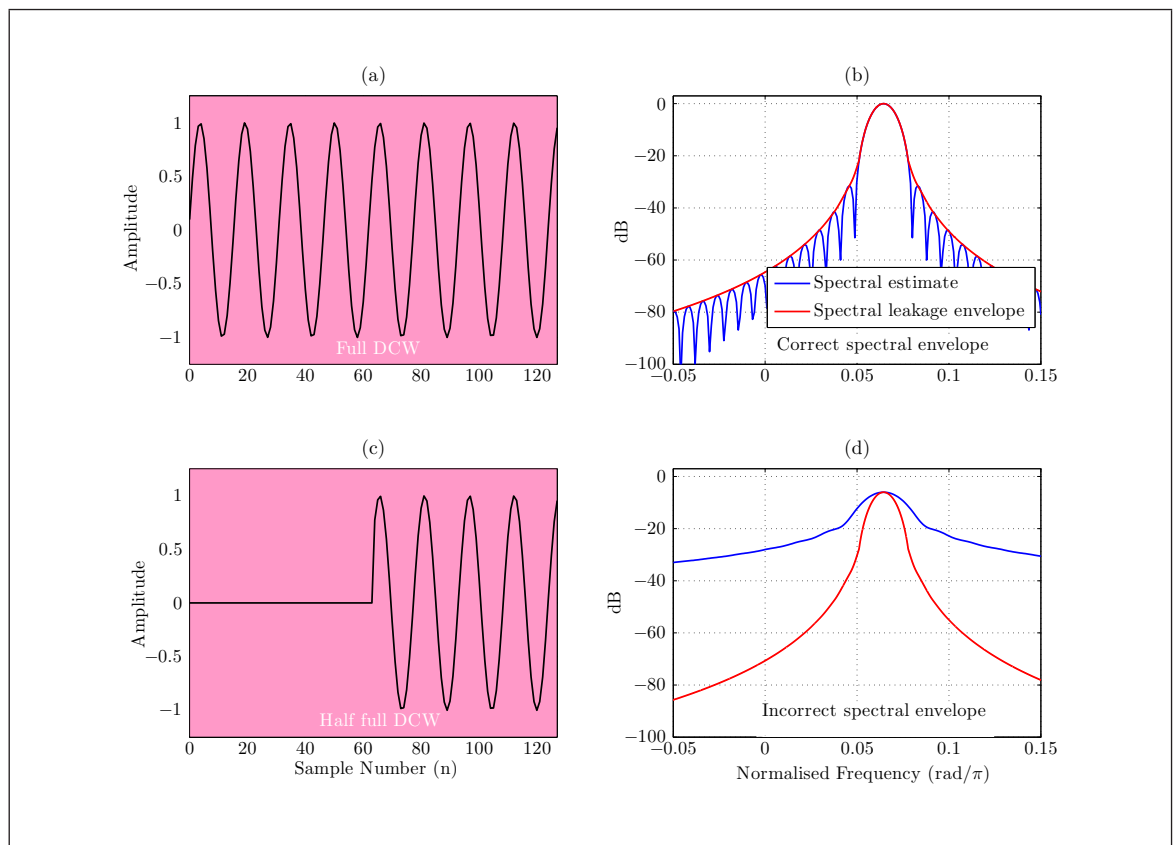


Figure 3.5.1 The effect missed signal samples has on the variable threshold calculation blocks: (a) a full DCW; (b) the Hann window spectral pattern and Hann spectral leakage envelope; (c) a half-full DCW; (d) spectral leakage pattern and incorrect Hann spectral leakage envelope

the DCWs. For a pictorial reminder of how the randomly-filled DCWs would have appeared an example is shown in Fig 3.3.1.

3.5.2 Results and Discussion

The ROC results for the detectors are shown in Figs 3.5.2–3.5.7. The results set was created from a similar variety of parameters to the interference scenario in section 3.4. Namely, a relatively low- $\text{SNR}_{\text{out},B}$, medium- $\text{SNR}_{\text{out},B}$ and high- $\text{SNR}_{\text{out},B}$ were tested. These $\text{SNR}_{\text{out},B}$ regimes were further broken down into closely-spaced frequency and spaced-apart frequency tests.

Fig 3.5.2 displays the results for an average frequency separation of $2f_s/N$ and Fig 3.5.3 shows the results for an average of $5f_s/N$. Both of these figures show the situation at $\text{SNR}_{\text{out},B} = 8 \text{ dB}$.

The $\text{SNR}_{\text{out},B}$ was increased to 15 dB and the results are displayed in Fig 3.5.4 and Fig 3.5.5. It was further increased to 25 dB and the results obtained feature in Fig 3.5.6 and Fig 3.5.7. The same two different average frequency separations were tested at all the $\text{SNR}_{\text{out},B}$ regimes. Panel (a) in each figure displays the results obtained when non-zero-padded DCWs were used and panel (b) in each figure shows the results when

twice-zeroadded DCWs were used. Every data point was calculated from 1×10^5 trials and the probability of false alarm was set at $P_{fa} = 0.01$.

One instantly noticeable aspect about all of the randomly-filled interference results in comparison to the full DCW interference results is that where there is a performance gap between detectors in the interference full DCW scenario, it is smaller in the interference randomly-filled DCW scenario. Another obvious change in the results from full interference DCW to randomly-filled interference DCW is that the general performance of all the detectors is worse. This feature is easily explained by the reduced S_A signal power that the detectors received as a result of the randomly-filled DCWs.

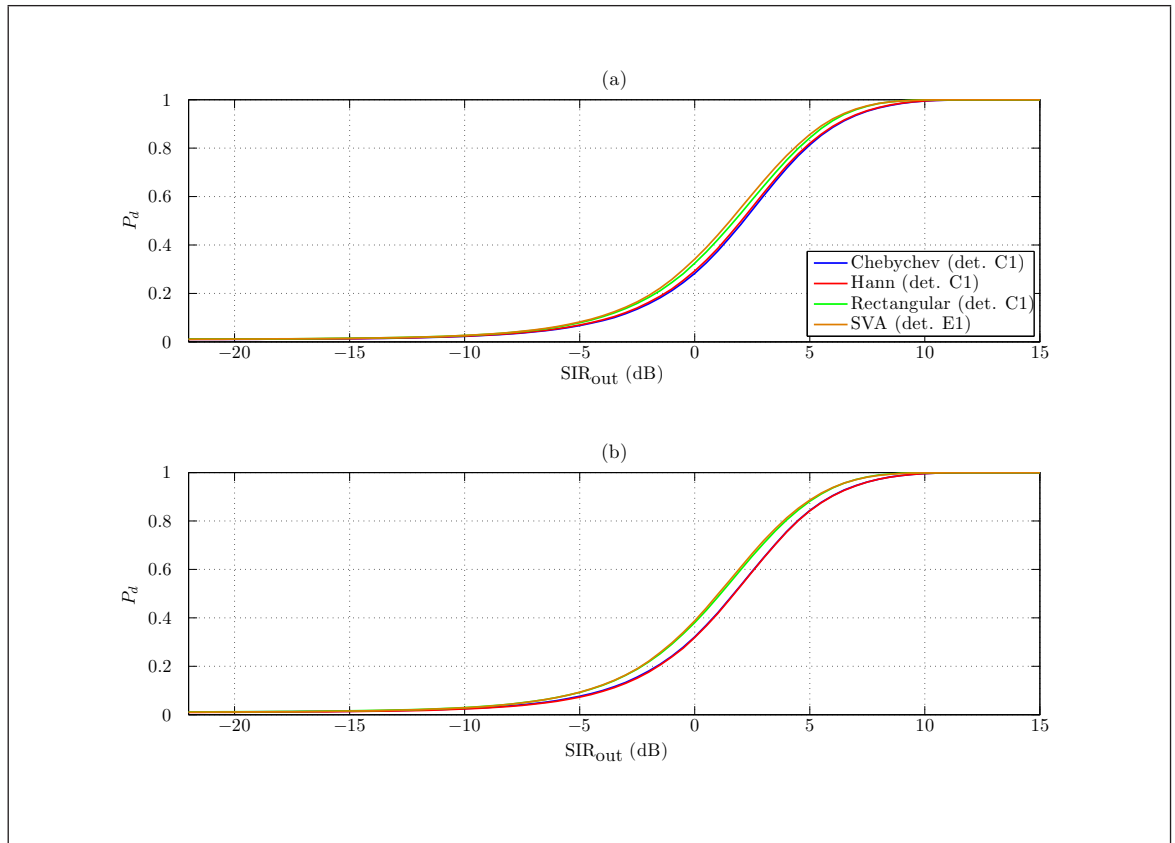


Figure 3.5.2 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 8$ dB with average separation of $2f_s/N$ for: (a) the non-zeroadded case; (b) the zeroadded case

Upon examination of the results in Fig 3.5.2, which were produced from the low- $\text{SNR}_{\text{out},B}$ regime, the SVA detector was the best performer as shown by both panels (a) and (b). This time in the interference randomly-filled DCW scenario the SVA detector was marginally better than the rectangular-windowed detector, whereas in the interference full DCW scenario the rectangular-windowed detector was marginally better than the SVA detector due to the additional algorithmic noise associated with the SVA detector. The switch in performance was because the rectangular-windowed detector was affected by the randomly-filled DCW aspect more significantly than the SVA detector due to the variable threshold component calculation in detector C1, which was not entirely accurate. The results in Fig 3.5.3, which were produced from

the same $\text{SNR}_{\text{out},B}$ but from an average separation of $5f_s/N$, are similar in nature in that the SVA detector is seen to be the best. The reason behind this is that $\text{SNR}_{\text{out},B} = 8 \text{ dB}$ is a low $\text{SNR}_{\text{out},B}$ regime and therefore the noise floor was the dominant factor which constrained detection performance rather than spectral leakage level.

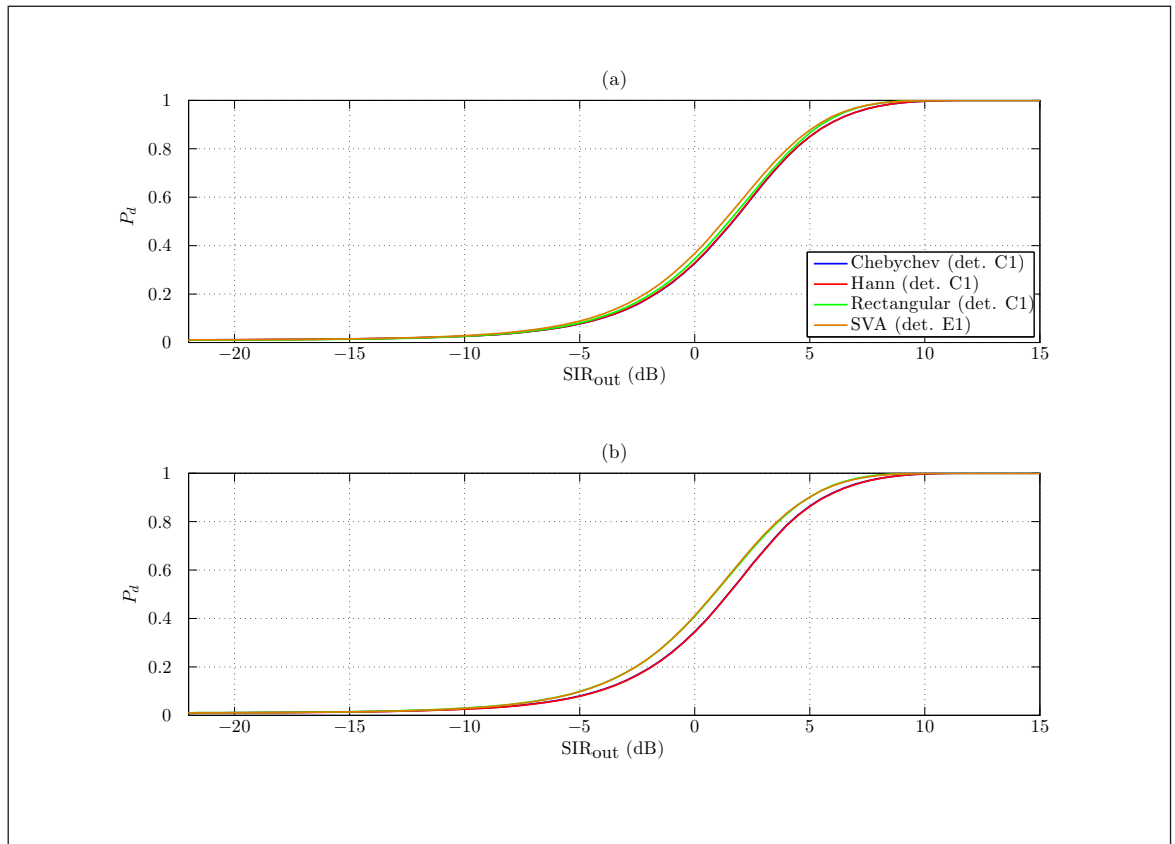


Figure 3.5.3 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 8 \text{ dB}$ with average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

Consideration of the results from both separations, revealed that the window loss was the dominant effect at both closely spaced frequencies, shown in Fig 3.5.2, and further apart frequencies, shown in Fig 3.5.3. When zeropadded the detection performance of all the detectors improved because the use of zeropadded DCWs allowed for more accurate spectral peak estimation. This is shown by the performance improvement visible from panel (a) to panel (b) in both Figs 3.5.2 and 3.5.3.

In the medium- $\text{SNR}_{\text{out},B}$ regime, for which results are displayed in Figs 3.5.4 and 3.5.5, the SVA detector continued to be the best detector because its detector architecture was affected less severely by the randomly-filled DCW. At the average separation of $5f_s/N$, the noise floor was still the dominant factor which limited detector performance and therefore the window loss was the chief effect which determined the relative detector performances in Fig 3.5.5. This explains why Fig 3.5.5 closely resembles Fig 3.5.2 and Fig 3.5.3, as they were produced from noise-floor-dominated regimes. However at an average separation of $2f_s/N$ and at the medium- $\text{SNR}_{\text{out},B}$ regime, the spectral leakage could not be ignored. In fact spectral leakage was the dominant factor over the noise

floor and hence the reason why the SVA detector was the best detector by a wider margin in Fig 3.5.4 than in Fig 3.5.5.

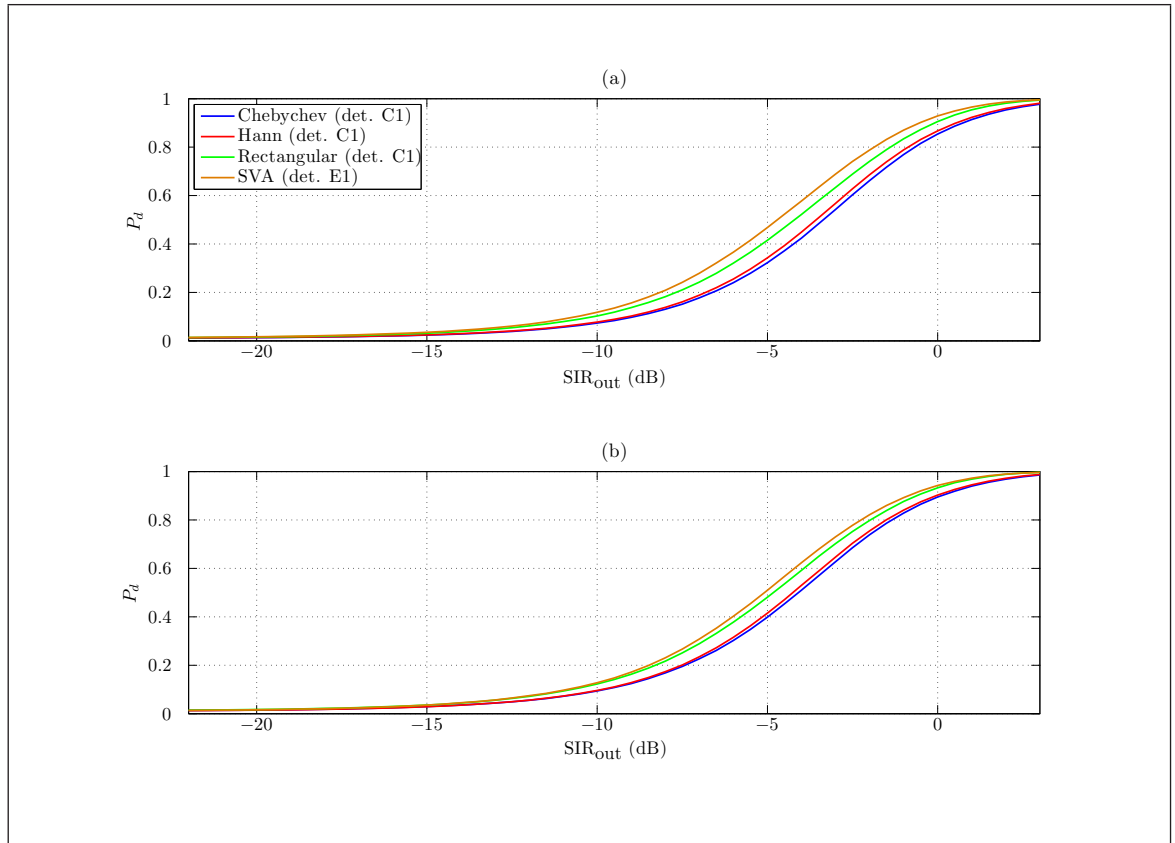


Figure 3.5.4 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 15$ dB with average separation of $2f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

As in the interference full DCW scenario, the use of a high $\text{SNR}_{\text{out},B}$ clearly portrayed the benefits of the use of SVA in a detector. This is especially true of closely spaced frequencies and can be seen from the results in Fig 3.5.6 where the spectral leakage of S_B was large. Fig 3.5.7 displays the results when the average frequency spacing was $5f_s/N$ and it shows the SVA detector was still the best performer in both panels (a) and (b) and definitely preferable to the rectangular-windowed detector but perhaps only marginally better than the Hann-windowed and Chebychev-windowed detectors. Like the outcome from the interference full DCW scenario the SVA, Hann-windowed and Chebychev-windowed results can be explained by noting that at greater frequency spacings the sidelobe reduction of the SVA technique and other non-rectangular windowing techniques becomes similar (for example panel (a) of Fig 2.5.1).

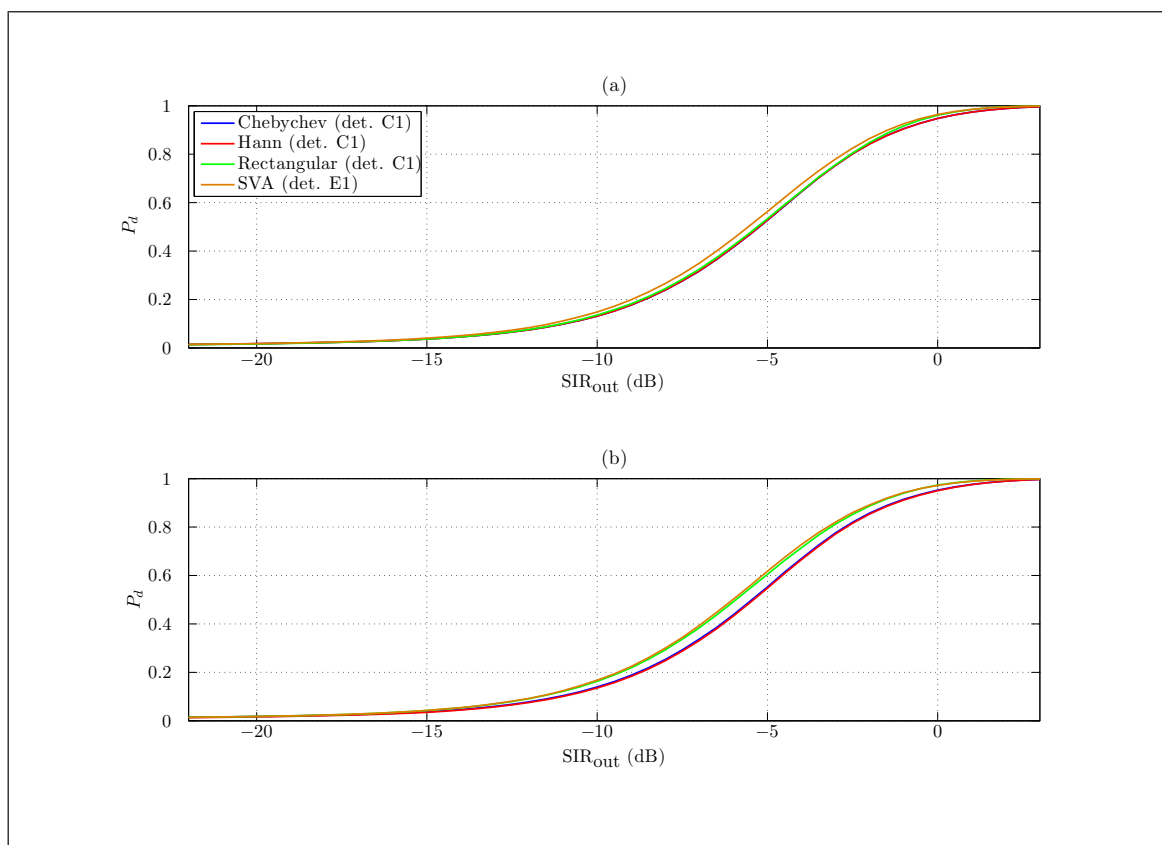


Figure 3.5.5 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 15 \text{ dB}$ with average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

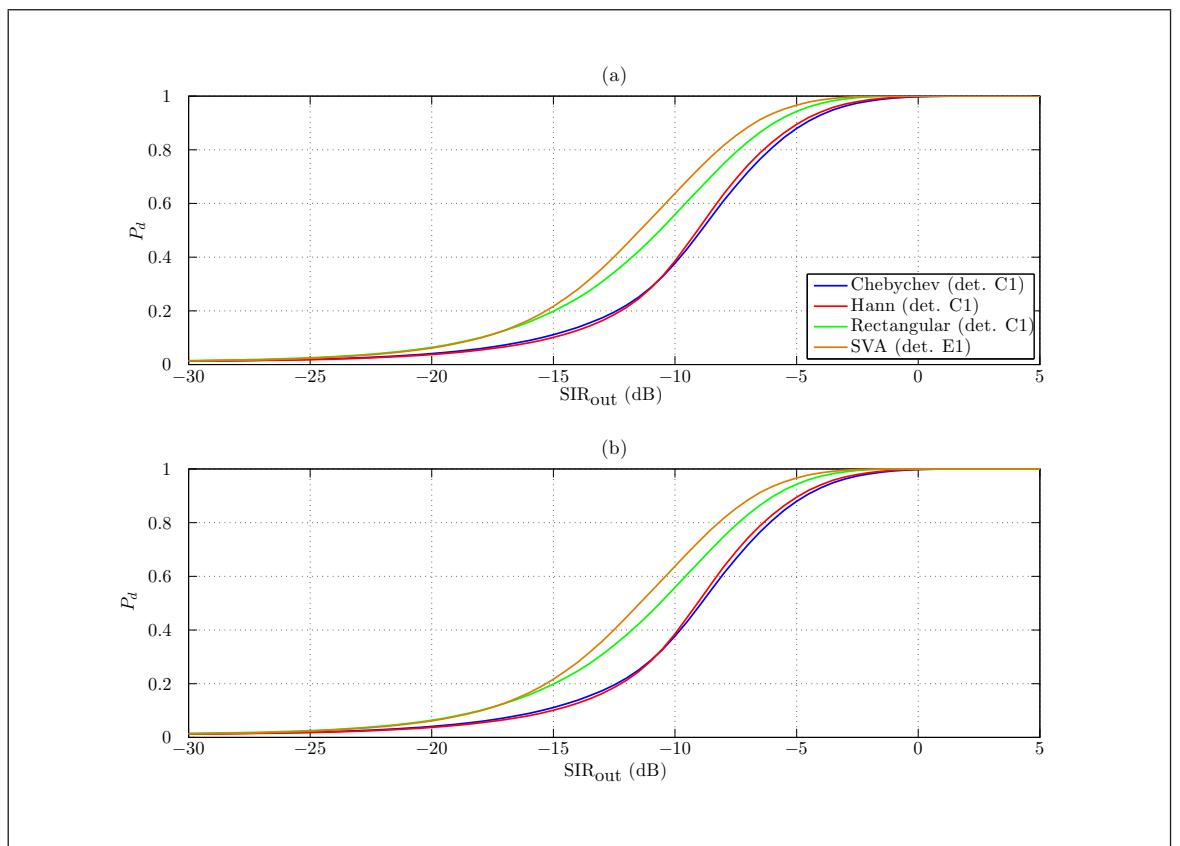


Figure 3.5.6 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 25$ dB with average separation of $2f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

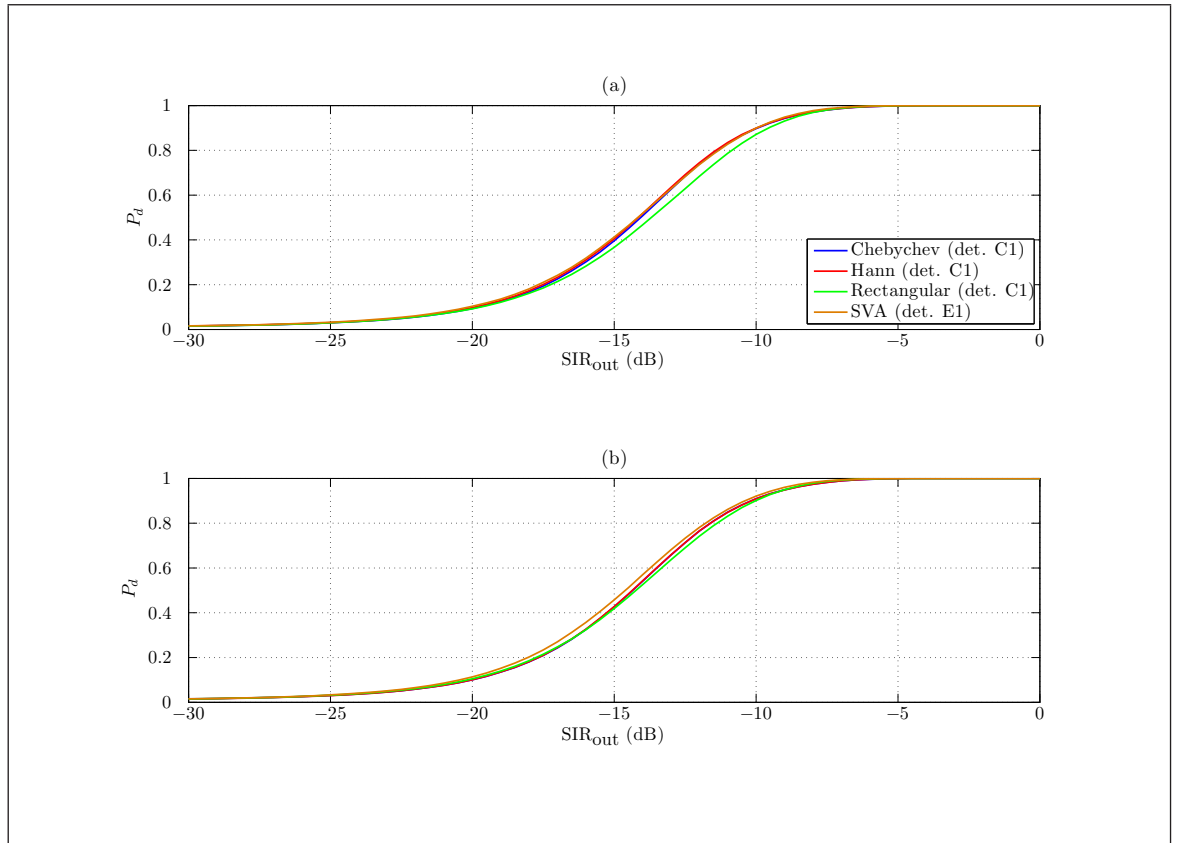


Figure 3.5.7 ROC curves for randomly-filled DCW interference scenario at $\text{SNR}_{\text{out},B} = 25$ dB with average separation of $5f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

3.6 LPI signal in DCW

From the previous sections it can be concluded that the SVA detector is successful as an all-round technique for many different scenarios in which a sinusoidal radar signal could be captured. The SVA technique seems able to assist with the detection of multiple sinusoids, which are necessarily narrowband. However SVA is not specifically designed to help with the detection of LPI radar signals that can be broader-band. In this section, results are obtained for an SVA detector tested against LPI waveforms.

3.6.1 Frequency-coded LPI Waveforms

FMCW Waveform:	bandwidth sweep	10 MHz
	sweep period	1 ms
	sweep rate	10 GHz/s
DRx:	sampling rate	[1.0, 1.75] GHz
	sampling period	[572, 1000] ps
	DCW length	1024 samples
	FFT mode	complex
	DCW overlap	0 samples
	non-zero padded frequency bin width	[0.97, 1.71] MHz
	DCW period	[585, 1024] ns

Table 3.6.1 Typical parameters for an FMCW waveform and DRx as part of an ESM system

One type of frequency-coded LPI waveform known to be in use in the PILOT radar system is linear FMCW[13, 15]. Essentially a single period of an FMCW waveform is composed of an up chirp ramp immediately followed by a down chirp ramp. The units are repeated indefinitely until the configuration is altered in some way, hence the CW part of the abbreviation. FMCW signals come under the chirps category of radar signals from section 1.8.

Typical parameters taken from [15] for an FMCW period are 10 MHz swept in 1 ms. This, along with typical parameters for a DRx as part of an ESM system are shown in Table 3.6.1. A longer FFT size and hence DCW duration than was used in the previous scenario is shown in Table 3.6.1. These parameters were used to form the simulations in this section. Notably, the FMCW sweep is greater than the bandwidth of an individual frequency bin but the sweep is not rapid enough to pass through a typical frequency bin width over the course of a typical DCW.

3.6.2 Scenario and Simulation Description

The scenario was created with detectors A1 and B1 attempting to detect a captured portion of the FMCW waveform through detection on the central frequency bin of the signal which was spread out in frequency. The parameters of the FMCW pulse were

taken from Table 3.6.1. However the central frequency bin of the captured portion of the FMCW signal was chosen randomly over the interval of f_s/N_{FFT} and the initial phase of the captured portion was also chosen at random. This was done so that the FMCW signal did not neatly coincide with the detector parameters. No interference signals were present in this scenario. The probability of false alarm was set at $P_{fa} = 0.01$.

3.6.3 Results and Discussion

The simulation results are shown in Fig 3.6.1. The results are almost identical to those of Fig 3.2.6 for the full DCW scenario. This was because the channelisation caused by the FFT decomposing the digitised IF into frequency bins in the simulated detectors was large relative to the bandwidth swept by the FMCW signal in the time taken for one DCW. Even with a relatively large DCW size of $N = 1024$ samples the bandwidth of the individual frequency bins was still relatively large. Therefore the FMCW signal appeared to the simulated detectors to have a bandwidth narrower than that of a single frequency bin after each DCW detection event.

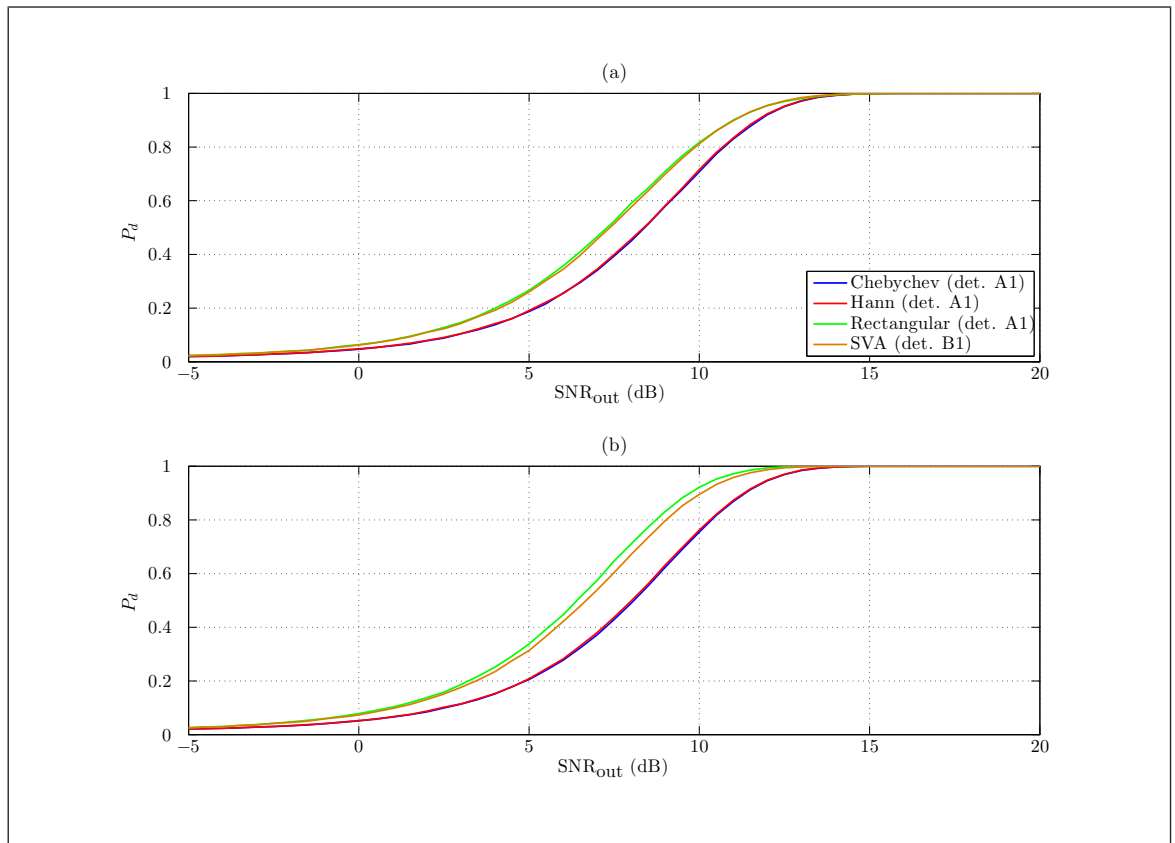


Figure 3.6.1 ROC curves for detection of an FMCW signal in detectors in two cases: (a) the non-zeropadded case; (b) the zeropadded case

The results in panel (a), which derived from the use of a non-zeropadded DCW in all the detectors, show the rectangular-windowed and SVA detectors had the best performance. The window loss associated with the Hann and Chebychev window functions caused their respective detectors to have a detector performance loss in comparison to

the rectangular-windowed detector. The interpretation of the SVA technique results is that it recovered the window loss. Panel (b) displays the results obtained through the use of a twice-zero padded DCW in all the detectors. Upon examination of the differences between the results from panel (a) to panel (b), it is seen that all the detectors improved upon their performances, however the rectangular-windowed detector and SVA detector (i.e. detector B1) increased their performances at a greater rate than the Hann-windowed and the Chebychev-windowed detectors. Zero padding improved detection because it enabled better estimation of the location of the spectral peaks.

The purpose of this scenario was to check whether the use of SVA would adversely affect the detection of chirp-like signals in a real-life, DFT-based DRx in an ESM system. This is because in DRxs the same mechanism for detection of sinusoidal signals is usually used for the detection of chirps, that is, through the detection of spectral peaks. The usual method for the identification of chirps is to track the movement of spectral peaks as they shift to adjacent frequency bins in subsequent DCWs. Therefore, the results show if SVA was added to a DRx it would not reduce the ability to detect the spectral peaks of part of an FMCW period and therefore the SVA technique would not interfere with the established procedure for identifying chirps.

3.6.4 Phase-coded LPI Waveforms

Another common type of LPI waveform known to be used in radar systems are Barker-coded, BPSK waveforms. In this scenario pulses were created which were encoded with a Barker-13 code. A technique sometimes called modulation-on-pulse (MOP). This signal falls under the category of a pseudo-noise radar signal from section 1.8. To create the simulations, the typical parameters from Table 3.6.2 were used. The pulse parameters were based on data captured by SELEX Galileo in the year 2006².

Barker-13 Pulse:	chip-rate	3.25 MHz
	pulse width	4 μ s
	PRF	60 kHz
	IF cycles/chip	~ 21
DRx:	sampling rate	[1.0, 1.75] GHz
	sampling period	[572, 1000] ps
	DCW length	1024 samples
	FFT mode	complex
	DCW overlap	0 samples
	non-zero padded frequency bin width	[0.97, 1.71] MHz
	DCW period	[585, 1024] ns

Table 3.6.2 Typical parameters for a Barker-13 pulse and DRx as part of an ESM system

²Based on intercepted pulses from the TTCP Trials 2006

3.6.5 Scenario and Simulation Description

The assumption was made that the simulated detectors would process data from full DCWs. This assumption was made because with the use of the parameters in Table 3.6.2, the detectors would record the $4\mu\text{s}$ long pulses and at least four consecutive DCWs would always be completely filled with samples which contain signal power (assuming several DCWs per ESM dwell).

In this scenario one test iteration involved the retention of one of the four DCWs which would contain a random portion of the pulse and an attempt to detect it. Another interpretation of this methodology is that the DCW could capture a random portion of the signal pulse but the random portion must be within the pulse temporal boundaries in order to satisfy the full DCW requirements. This concept is illustrated in Fig 3.6.2, where panel (a) shows a Barker-13 coded pulse and panel (b) shows the DCW can capture any part of this pulse within its limits. 1×10^5 trials were used per data point in the ROC curves.

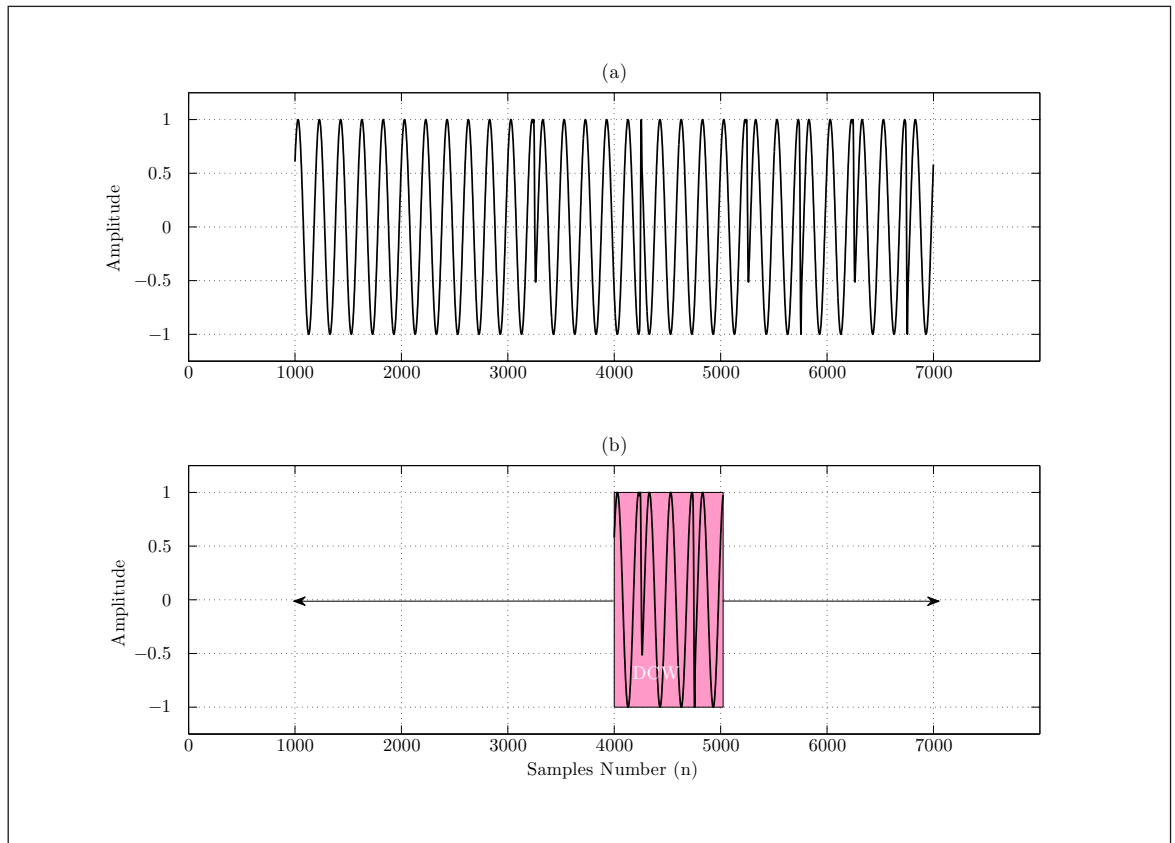


Figure 3.6.2 Capture of a random portion of LPI pulse: (a) a Barker-13-coded pulse; (b) DCW capturing a random portion within limits

As with all previous simulations, it was desirable to keep the parameters of the signal separate from the parameters of the simulation detectors. To that end, the frequency (f_T) of the Barker pulses had a nominal value but it was chosen at random at the start of each DCW test such that it was random over the equivalent of f_s/N_{FFT} around the nominal value. The initial phase (ϕ_0) was also selected at random at the start

of each DCW. The phase jumps associated with the actual Barker code were made over finite transitions (5% of chip interval) instead of theoretical discontinuous jumps which are often described in textbooks, for example in [13]. This was done to simulate the output from a real-life signal generator. The detector architectures were similar to those used for FMCW detection, however the detection logic was slightly different. In fact detector N1 from section 5.1.5 was used for the conventional window detection and detector O1 was used for SVA detection. More detail on these detectors is given in section 5.1.5. With a signal chip-rate of 3.25 MHz and with a frequency bin width of $[0.97, 1.71]$ MHz, the signal was expected to be spread across a few frequency bins. However since the exact target frequency was unknown and the number of frequency bins is always an integer number, 3 frequency bins were checked for peaks caused by the Barker-13 coded pulse. For the twice-zeroadded detectors the number of bins that were checked was 5.

Overall the detectors captured a random part of the Barker-13 coded pulse and attempted to detect on the frequency bins which could contain signal power from the Barker-13 pulse. Like the FMCW scenario, no interfering signals were present and the probability of false alarm was set at $P_{fa} = 0.01$.

3.6.6 Results and Discussion

The simulation results are shown in Fig 3.6.3. As with the FMCW ROC curves in Fig 3.6.1, the picture is pretty similar to that obtained from the full DCW scenario in Fig 3.2.6. That is, the SVA technique helped to recover the window loss from which a detector would suffer if it used a moderate dynamic range window function like the Hann or Chebychev window function. Also the zeroadding improved detection performance for all detectors because it gave a better estimate of the location and magnitude of any spectral peaks.

The incentive for this scenario was exactly the same as for the FMCW scenario. It was desirable to uncover whether the use of SVA would somehow penalise detection of LPI radar signals. In this scenario, the simulated detectors attempted to detect a pseudo-noise radar signal in the form of a Barker-13 pulse and the results show that the detection was not adversely affected through the use of SVA.

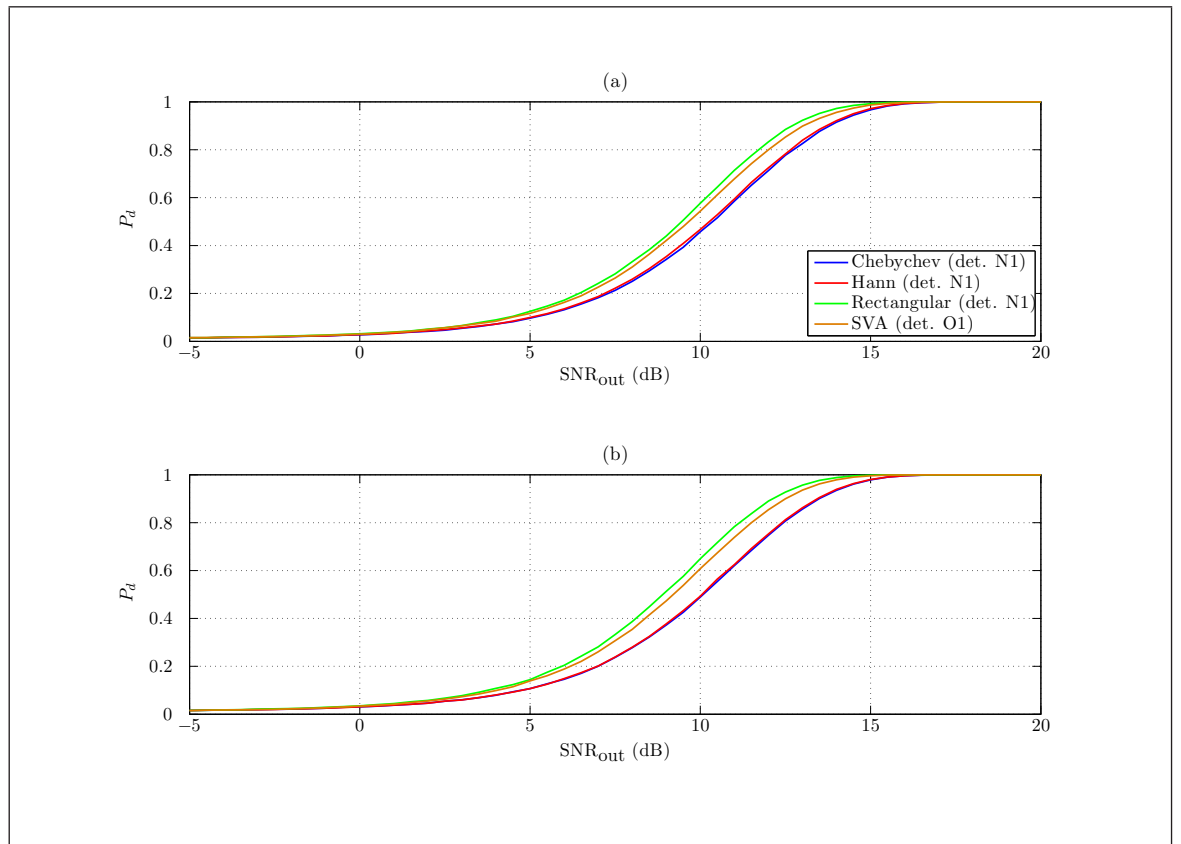


Figure 3.6.3 ROC curves for detection of a Barker-13 pulse in detectors in two cases: (a) the non-zeropadded case; (b) the zeropadded case

3.6.7 Summary Table

The results from this chapter are summarised in Table 3.6.3.

Scenario	Best Detector(s)	Description
Full DCW, no interference	Rectangular-windowed detector, SVA detector	The SVA detector overcame the window loss suffered by Chebychev-windowed detector and Hann-windowed detector
Randomly-filled DCW, no interference	Rectangular-windowed detector, SVA detector	The window loss effect was much reduced but SVA still recovered the loss to a similar level of performance as the rectangular-windowed detector
Full DCW, interference present	SVA detector	The SVA detector reduced the impact of spectral leakage from the interference and recovered the window loss
Randomly-filled DCW, interference present	SVA detector	The window loss and spectral leakage effects were smaller but SVA remained the best detector
LPI signal, no interference	Rectangular-windowed detector, SVA detector	The rectangular detector performed well because it allowed as much signal power to be collected as possible. The SVA detector performed almost as well as the rectangular-windowed detector

Table 3.6.3 Floating-point simulation results summary

CHAPTER 4

Non-Parametric Spectral Estimation

4.1 Detection

Detection and identification of radar threats is the primary goal of a DRx as part of an ESM system. As many radar signals are sparse in the frequency-domain, the common method to detect and identify threats is to estimate the true operational environment power spectral density (PSD) over different time intervals and detect peaks from this.

Forming a good spectral estimate from which initial detections are found can be the basic building block for a good ESM detector. The initial detections can be passed to more complicated classification algorithms for further filtering as per Fig 1.3.1.

The PSD is, for wide-sense stationary processes, the Fourier transform of the autocorrelation function ($\mathcal{F}\{R_{xx}(\tau)\}$), for which it would be possible to exactly specify if $x(t)$ was known for all t . Naturally however in an ESM DRx, $x(t)$ is sampled in DCWs and never fully known because it is corrupted by background noise and interference. Therefore many ESM DRxs employ a spectral estimation technique.

This is commonly achieved through the use of an FFT algorithm due to its simplicity. However the use of the magnitude-squared of an FFT is one of many spectral estimation techniques.

Spectral estimation can be neatly split into two groups:

Parametric Assume a data-generating process or model.

Non-parametric Use only the data provided in the DCWs.

In general, non-parametric methods are computationally simple but have limited accuracy. A DRx in an ESM system processes vast amounts of data in short periods of time and therefore computational simplicity is usually of greater importance than limited accuracy. In addition to this a DRx generally can not assume much about the characteristics of unknown signals initially. Therefore non-parametric methods are generally employed by DRxs in ESM systems.

There are a few properties of spectral estimators to consider:

Bias Whether the spectral estimate will asymptotically converge to the true PSD (S_{xx}) with greater N , i.e.

$$\lim_{N \rightarrow \infty} \mathbb{E} \left\{ \hat{S}_{xx}(f) \right\} = S_{xx}(f).$$

Variance Does the amount of deviation from the PSD change with a bigger or smaller N , i.e.

$$\lim_{N \rightarrow \infty} \text{Var} \left\{ \hat{S}_{xx}(f) \right\} = 0.$$

Resolution The extent of the blurring of the spectral estimate compared to the PSD.

The magnitude-squared of an FFT is a simple spectral estimator and is valid because the Fourier transform of the sampled autocorrelation function would be used which is given by

$$R_{xx}(l) = \lim_{N \rightarrow \infty} \left\{ \frac{1}{2N+1} \sum_{n=-N}^N x(n+l)x(n) \right\}, \quad l = 0, \pm 1, \pm 2, \dots, \pm N-1 \quad (4.1.1)$$

however a real-life DRx would have only a finite number of samples N , with which to work. Therefore it would approximate the ideal autocorrelation in Eq (4.1.1) with a finite sample autocorrelation

$$\hat{R}_{xx}(l) = \frac{1}{N} \sum_{n=0}^{N-1-|l|} x(n+l)x(n) \quad (4.1.2)$$

The implicit assumption in Eq (4.1.2) is that $x(n)$ is zero outside of the DCW and therefore simplifies to

$$\hat{R}_{xx}(l) = \frac{1}{N} x(l) * x(-l) \quad (4.1.3)$$

and upon the DFT of Eq (4.1.3), the method of “FFT and magnitude-squared” is recovered

$$\hat{S}_{xx}^{(P)}(k) = \frac{1}{N} |X(k)|^2, \quad \text{where } k = 0, 1, 2, \dots, N_{\text{FFT}} - 1 \quad (4.1.4)$$

This spectral estimation method, which was used in chapter 3, is called the periodogram. The bias of a periodogram is

$$\mathbb{E} \left\{ \hat{S}_{xx}^{(P)}(k) \right\} = W_B(k) * S(k)$$

where $W_B(k)$ is the Fourier transform of the Bartlett window (i.e. $W_B(k) = \frac{1}{N} \left[\frac{\sin \pi k N}{\sin \pi k} \right]$). When $N \rightarrow \infty$, $W_B(k) \rightarrow \delta(0)$ and therefore the periodogram is asymptotically unbiased. This is a beneficial feature of the periodogram, however the variance of the periodogram is

$$\text{Var} \left\{ \hat{S}_{xx}^{(P)}(k) \right\} \approx S_{xx}^2(k)$$

The variance is large and independent of N and therefore no matter how many samples are processed in a DCW the variance never decreases. This property of the variance

was evident in the jagged nature of many of the periodogram-type plots in chapter 2, for example Figs 2.5.2 and 2.8.5.

There are other non-parametric spectral estimators which seek to remedy the variance issue of the periodogram through a trade-off of variance for bias or resolution. A list of possible non-parametric spectral estimators is given in Table 4.1.1. The SVA work is included in the table and highlighted as orange as a modification of the periodogram method. The methods highlighted in red trade some spectral resolution of the periodogram for a reduction in variance. In terms of computational complexity, they are the next most computationally complex methods after the periodogram (and its relations: the modified periodogram and SVA periodogram). The candidate methods are examined in chapter 5 and some theory follows in this chapter.

Method	Procedure	Comment
Periodogram	DFT of a time series and calculate the magnitude-squared	Large variance, inconsistent estimator
Modified Periodogram	Apply window function to time series prior to DFT and magnitude-squared	Same as above, some spectral resolution traded-off against spectral leakage
SVA Periodogram	Apply DFT to time series, execute SVA algorithm and take magnitude-squared	Same as periodogram, but spectral leakage much reduced
Blackman-Tukey	Fourier transform of truncated autocorrelation function after weighted with a lag window	Obsolete since advent of fast computers and FFT as Fourier coefficients directly estimated with FFT
Smoothed Periodogram	Estimate periodogram, apply a smoothing filter such as moving average filter	Spectral resolution traded for reduction in variance. Frequency-domain equivalent of Blackman-Tukey
WOSA	Divide the time series into smaller, overlapped, windowed segments. Find periodogram of each segment and average. When no window function or overlap used, known as Bartlett's method	Trades spectral resolution for variance reduction. Bias problems for short time series.
Multitaper Method	Orthogonal window functions applied to same time series, DFT and magnitude-squared. Results averaged	Spectral resolution traded for reduction in variance. Trade-off easily quantified
Singular Spectrum Analysis	Decompose time series into sum of components. Identify components due to noise and reconstruct time series without them. Form spectral estimates from de-noised time series	Largest computational load out of all the above

Table 4.1.1 Non-parametric Spectral Estimators

4.2 Multitaper Theory

4.2.1 Introduction

In addition to the information in Table 4.1.1 this section provides some additional reasons for the study of the multitaper method (MTM) as an algorithm of interest for a DRx as part of an ESM system. The MTM is often overlooked, however its usefulness is highlighted by its success in spectral holes sensing in cognitive radio in [125], where it is also combined with cyclostationarity in a Lo  ve transform to detect and classify signals in [126, 127].

Some common characteristics of the MTM and SVA are revealed in this section and in section 4.3 the smoothed periodogram is examined as it is similar to one of the MTMs.

Some common characteristics of the MTM and SVA:

- The MTM attempts to overcome the window loss when a window function is applied to a DCW prior to a DFT.
- The MTM can be applied in an automatic fashion as it is a non-parametric spectral analysis technique, which requires no *a priori* information.
- The MTM is fully compatible with spectral analysis based on an FFT.
- The MTM is a relatively, computationally simple technique.
- One particular version of the MTM involves simple convolutions in the frequency-domain.

As a technique, the MTM was formulated in [128] as an alternative to Welch’s overlapped segment averaging (WOSA), which is a popular technique for the reduction of variance in non-parametric spectral analysis. WOSA can have bias problems for shorter DCWs as it comprises the steps of: (1) split the DCW into shorter, overlapping segments and (2) average the spectral estimates from each segment. This is an important point in ESM as there may be no knowledge of the target signal pulse duration *a priori*. Therefore our DCW length may be a poor choice which could be worsened by WOSA. Additionally the variance reduction in WOSA is not proportional to the number of segments averaged due to correlation between noise samples in overlapped segments. WOSA was not considered further for study

The MTM, can however, use the full DCW through the steps of: (1) application of multiple orthogonal tapers (i.e. window functions) to the DCW (2) perform FFTs and (3) average the spectral estimates from each window function. Fig 4.2.1 pictorially contrasts the basic operations in WOSA and the MTM.

A further contribution came from [129] which introduced a different set of tapers to the original discrete prolate spheroidal sequence (DPSS) tapers from [128]. This different

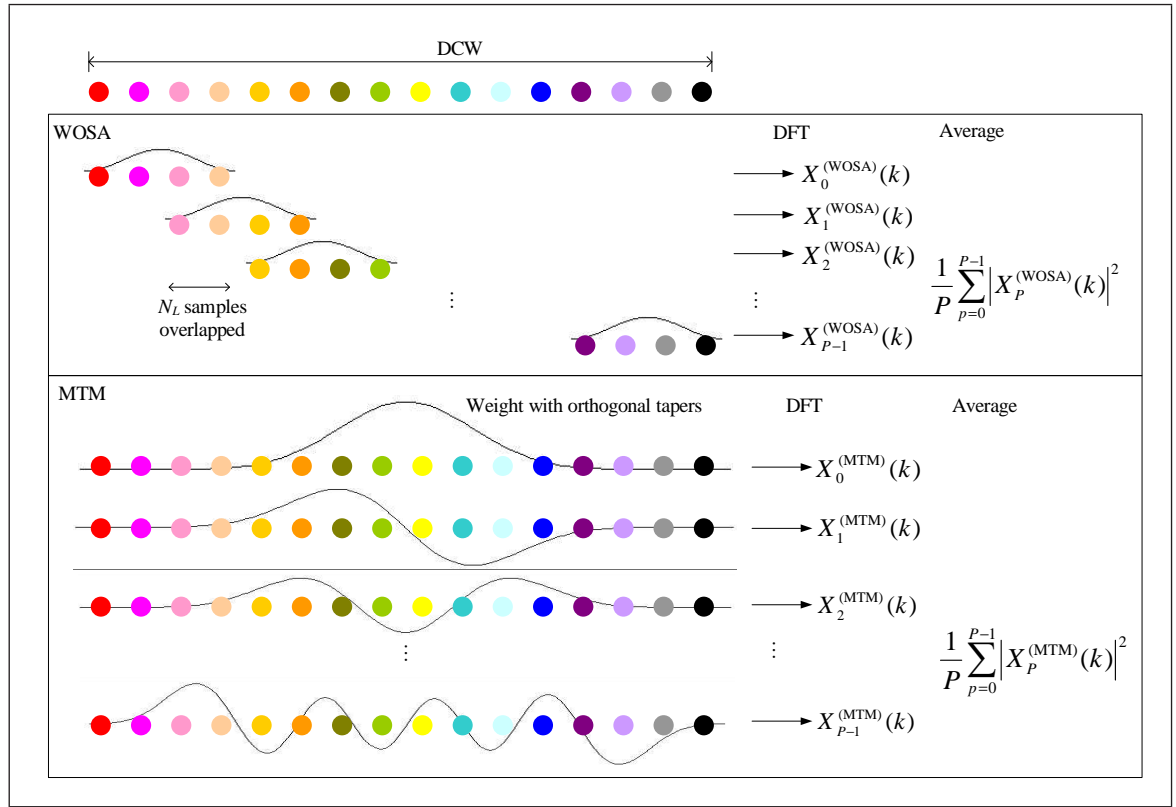


Figure 4.2.1 Basic operations in WOSA and MTM

set is called the sine tapers and possess the advantages of an analytical formulation and a simple frequency-domain representation.

In the following sections the rationale for the investigation of the MTM in terms of the detection of unknown LPI radar waveforms is explained, which leads onto a discussion of even simpler smoothed periodogram methods.

4.2.2 MTM

Here the motivation for investigation of the MTM is explained. LPI radar signals may attempt to be covert through a spread of their peak power across a frequency band by the use of, for example, phase-modulation (see section 1.9). If the intercepted signal is strong enough, such LPI radar signals appear as a band of energy in spectrogram-based or periodogram-based techniques (the periodogram is merely 1 DCW or slice through the spectrogram). A pictorial example of this is displayed in panel (b) of Fig 1.9.1.

When conventional spectrograms and periodograms are applied in this manner in a DRx it effectively **over-resolves** the LPI signal in the frequency dimension. This is because the filter shape of the individual FFT bins is too narrow for the signal. The magnitude filter response of each FFT bin is the same as the magnitude frequency-domain response of the window function that was used at the input of the FFT. A better match to the bandwidth of the signal bandwidth is obtained through the MTM.

In the MTM, as the number of tapers used increases, so too does the bandwidth of

each frequency bin. This effect is demonstrated in the series of plots in Fig 4.2.2. For $N_{\text{FFT}} = 1024$ and a design half-bandwidth (in normalised terms) of $W_h = 0.0195 \text{ rad}/\pi$, the average of the MTM taper magnitude-squared responses is shown on a logarithmic scale for an increasing number of tapers used up to the Shannon number. The Shannon number and parameter W_h are explained in section 4.2.5. The number of tapers used, from top to bottom, is $p = 10, 20, 30, 40$. The response becomes more brick-wall-like as more tapers are added up to the Shannon number. This shape more closely resembles the band of energy shown in panel (b) of Fig 1.9.1 and therefore the MTM could be used to collect as much signal power as possible into the bin upon which the signal band is centred, provided the correct half-bandwidth parameter W_h is chosen. Therefore this is a strong motivation for the investigation of the MTM implemented in an ESM DRx as a simple means to improve LPI radar signal detection.

As the average MTM taper response shape becomes wider and flatter, the frequency resolution of the MTM spectral estimate becomes coarser. The normalised frequency resolution approximates to $P/N \text{ rad}/\pi$, where P is the number of tapers used and N is the number of samples in the DCW. However this may not be a problem for a real-life DRx which could use SVA for the detection of narrower band signals like sinusoids and the MTM to detect broader-band, LPI, pseudo-noise signals.

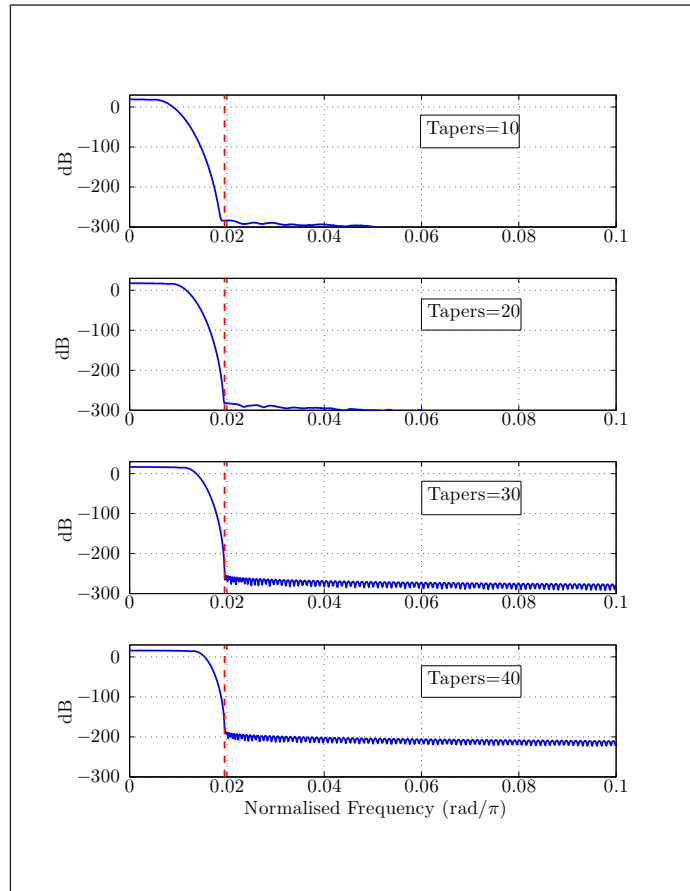


Figure 4.2.2 Average of the magnitude-squared filter responses of MTM tapers for increasing numbers of tapers designed for $W_h = 0.0195 \text{ rad}/\pi$

The simplest formulation of the MTM is an average of P spectral estimates ($\hat{S}_{p,xx}^{(\text{MTM})}$) formed from the same DCW as

$$\hat{S}_{xx}^{(\text{MTM})}(k) = \frac{1}{P} \sum_{p=0}^{P-1} \hat{S}_{p,xx}^{(\text{MTM})}(k), \quad \text{where } k = 0, 1, \dots, N-1 \quad (4.2.1)$$

$\hat{S}_{p,xx}^{(\text{MTM})}$ are simply modified periodograms formed from: (1) the application of normalised window functions ($w_p(n)$ with $n = 0, 1, 2, \dots, N-1$), which are pairwise orthogonal to each other to a DCW, (2) application of a DFT and (3) the calculation of the magnitude-squared of the Fourier coefficients.

The orthogonality condition is

$$\sum_{n=0}^{N-1} w_p(n) w_q(n) = \begin{cases} 1, & \text{if } p = q; \\ 0, & \text{otherwise.} \end{cases}$$

and application of a DFT and subsequent calculation of the magnitude-squared yields

$$\hat{S}_{p,xx}^{(\text{MTM})}(k) = \left| \sum_{n=0}^{N-1} w_p(n) x(n) W_N^{-nk} \right|^2$$

There are MTM spectral estimators other than that of Eq (4.2.1) which use a non-uniformly weighted average scheme but the discussion is limited to the uniformly weighted case for now.

The use of any non-rectangular window function on white noise amounts to removal of some information in a DCW as the spectrum is smoothed and has an increased variance as described in [130]. Throughout the SVA work this phenomenon was termed a window loss. As shown in [131, p. 347–349], this loss can be overcome through the use of the MTM as many orthogonal spectral estimates or eigenspectra are formed from the same DCW and averaged. If all possible tapers are used then the window loss is recovered. However, as is shown in Fig 4.2.6, higher-order tapers, past the Shannon number, are undesirable because they actually increase the “broadband bias” of the MTM spectral estimate ($\hat{S}_{xx}^{(\text{MTM})}(k)$) outside the normalised bandwidth of interest $[-W_h, W_h]$ rad/ π . Therefore they are not generally included in the averaging process. Fortunately, in [131] it is shown that the higher-order tapers, when used in the average process, contribute only a relatively modest amount to a decrease in variance of $\hat{S}_{xx}^{(\text{MTM})}(k)$. Therefore their omission from the MTM process has little impact on the variance reduction property.

When complex AWGN is input into the MTM, each eigenspectrum consisting of frequency bins ($\hat{X}_p^{(\text{MTM})}(k)$) is distributed as an exponential random variable (RV)

$$f(X; \lambda) = \begin{cases} \lambda \exp(-\lambda X) & \text{if } X \geq 0; \\ 0 & \text{if } X < 0. \end{cases}$$

When the AWGN has unit variance $\sigma^2 = 1$, this is the same as a chi-square RV with 2 degrees of freedom (typically denoted as $Q \sim \chi^2(2)$). The chi-square distribution arises from the sum of the squares of independent, standard normal RVs A_1, A_2, \dots, A_M as

$$Q = \sum_{m=1}^M A_m^2$$

The PDF for a chi-square distributed RV is given by

$$f(X; d) = \frac{1}{2^{d/2} \Gamma(d/2)} X^{d/2-1} \exp(-X/2) \text{Pos}_{\{X \geq 0\}}$$

where d denotes the degrees of freedom, $\Gamma(\dots)$ denotes the Gamma function and $\text{Pos}_{\{x \geq 0\}}$ shows that the chi-square distribution is only defined for non-negative x .

The summation of the squares of independent RVs resembles the process which occurs in the MTM when P eigenspectra are averaged. This is because the eigenspectra, which are theoretically pairwise orthogonal and independent, are, in the simplest MTM, summed and divided by the number of eigenspectra. Therefore, when complex AWGN is processed, the multitaper spectral estimate $\hat{S}_{xx}^{(\text{MTM})}(k)$ should approximately be distributed as a scaled version of the chi-square distribution with $2P$ degrees of freedom. A scaled chi-square distribution would be given by

$$f(X; d) = \frac{d^{d/2}}{2^{d/2} \sigma^d \Gamma(d/2)} X^{d/2-1} \exp\left(-\frac{X d}{2\sigma^2}\right) \text{Pos}_{\{X \geq 0\}} \quad (4.2.2)$$

where σ^2 is the noise power and the degrees of freedom (d) would be replaced by $2P$. An example of the agreement between the PDF given by Eq (4.2.2) and empirically obtained MTM data is shown in Fig 4.2.3. Here complex AWGN with variance $\sigma^2 = 1$ was processed by the MTM using DCWs and FFTs of length $N = N_{\text{FFT}} = 128$. The time-bandwidth product was arbitrarily chosen to be $NW_h = 5$. 1×10^3 DCWs were processed such that a total of 128,000 data points were grouped into bins in a normalised histogram. The number of degrees of freedom was $2 \times (2NW_h) = 20$ and this was used to create the theoretical PDF curve from Eq (4.2.2).

Such theoretical results are useful because through integration of Eq (4.2.2) between a chosen lower limit and an upper limit of $+\infty$ probabilities of false alarm can be obtained. The lower limit on integration would represent a threshold in a real-life

detector which utilised the MTM and therefore from the theoretical distributions it is possible to get an idea of the thresholds which would need to be set in order to achieve a particular P_{fa} . This concept was used where possible in the simulations in chapter 5.

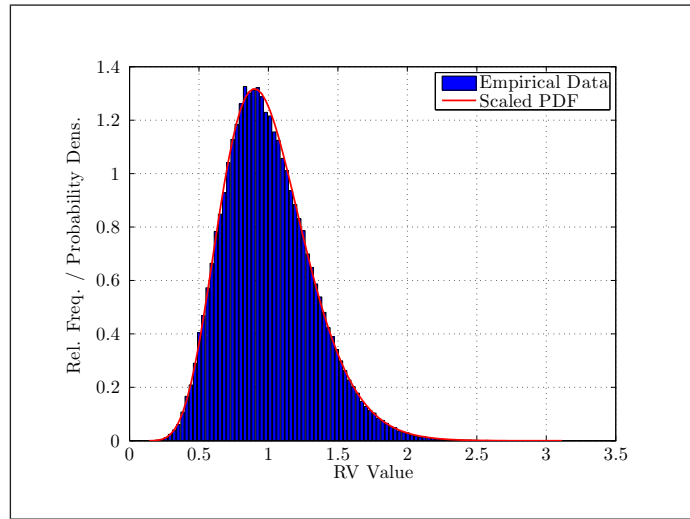


Figure 4.2.3 $Q \sim \chi^2(20)$ PDF overlaid on a normalised histogram of 128,000 MTM values

The theoretical PDF given in Eq (4.2.2) is valid for the simplest version of the MTM where the eigenspectra are uniformly weighted by 1, however there are variations on the weight scheme which are detailed in the next section for which the PDF in Eq (4.2.2) is only be a good fit for large P .

4.2.3 MTM Variants

The uniformly-weighted average MTM is the simplest MTM spectral estimator. However in [131] a derivation of an unbiased MTM spectral estimator is given which essentially weights each eigenspectrum by its eigenvalue

$$\hat{S}^{(\text{MTM})}(k) \equiv \frac{\sum_{p=0}^{P-1} \lambda_p \hat{S}_p^{(\text{MTM})}}{\sum_{p=0}^{P-1} \lambda_p} \quad (4.2.3)$$

The significance of the eigenvalues is explained in the next section, however, from Fig 4.2.4 it can be seen that as the order of the eigenvalue increases, the size of the eigenvalue decreases towards 0. Evidently from Eq (4.2.3) less emphasis is placed on the higher-order eigenspectra, which is intuitively better as higher-order eigenspectra become more polluted with spectral leakage. The difference between the uniformly-weighted MTM and the eigenvalue-weighted MTM is slight since the number of tapers (P) is usually chosen to be $P = 2NW_h, 2NW_h - 1$ or $2NW_h - 2$ and therefore most of the eigenvalues are close to the value of 1 anyway. The importance of the Shannon number $2NW_h$ is explained further in the next section.

Also in [131] the derivation of another estimator which has adaptive weights is shown as

$$\hat{S}_{xx}^{(\text{MTM})}(k) \equiv \frac{\sum_{p=0}^{P-1} b_p(k)^2 \lambda_p \hat{S}_{p,xx}^{(\text{MTM})}(k)}{\sum_{p=0}^{P-1} b_p(k)^2 \lambda_p} \quad (4.2.4)$$

where the value of $b_p(k)$ is estimated from

$$b_p(k) = \frac{\hat{S}_{xx}(k)}{\lambda_p \hat{S}_{xx}(k) + (1 - \lambda_p) \hat{\sigma}_x^2} \quad (4.2.5)$$

Initially with 1 eigenspectrum or 2 eigenspectra, a spectral estimate ($\hat{S}_{xx}^{(\text{MTM})}(k)$) is formed from Eq (4.2.4), which along with an estimated variance is substituted into Eq (4.2.5) to obtain estimates for all the weights $b_p(k)$. The weights allow a new, updated spectral estimate to be made with the use of Eq (4.2.4), which can subsequently be inserted back into Eq (4.2.5) to iterate the process. As stated in [131], a reasonable result is usually obtained after a few iterations.

4.2.4 Spectral Concentration Problem

Qualitatively, it is beneficial if the required P window functions have low sidelobe levels in the frequency-domain in order to minimise spectral leakage. Put differently, window functions which maximise the ratio of power in a given frequency interval $[-W_h, W_h]$ rad/ π to the total power in the normalised frequency band interval $[-1/2, 1/2]$ rad/ π are required. This power ratio is typically called the spectral concentration (common notation: $\lambda(N, W_h)$) of the particular window function. Explicitly the spectral concentration is given by

$$\lambda(N, W_h) = \frac{\int_{-W_h}^{W_h} |W(f)|^2 df}{\int_{-1/2}^{1/2} |W(f)|^2 df}$$

where the DTFT of the window functions ($W(f)$) is defined as

$$W(f) = \sum_{n=0}^{N-1} w(n) W_N^{-nf}$$

The spectral concentration problem is solved by maximising $\int_{-W_h}^{W_h} |W(f)|^2 df$ subject to the total power being fixed to unity, $\int_{-1/2}^{1/2} |W(f)|^2 df = 1$. This leads to the eigenvalue equation

$$\sum_{n'=0}^{N-1} \frac{\sin 2\pi W (n - n')}{\pi (n - n')} w(n') = \lambda w(n), \quad n = 0, 1, 2, \dots, N - 1 \quad (4.2.6)$$

The eigenvectors of the eigenvalue equation are the tapers which solve the spectral concentration problem. These eigenvectors are termed DPSS tapers or Slepian sequences. The tapers are normalised such that the sum of the squares of the eigenvectors is equal to unity

$$\sum_{n=0}^{N-1} w_p(n)^2 = 1$$

Contained within Eq (4.2.6) is an $N \times N$ positive-definite, real, symmetric matrix given by

$$V_{n,n'} = \frac{\sin 2\pi W (n - n')}{\pi (n - n')}$$

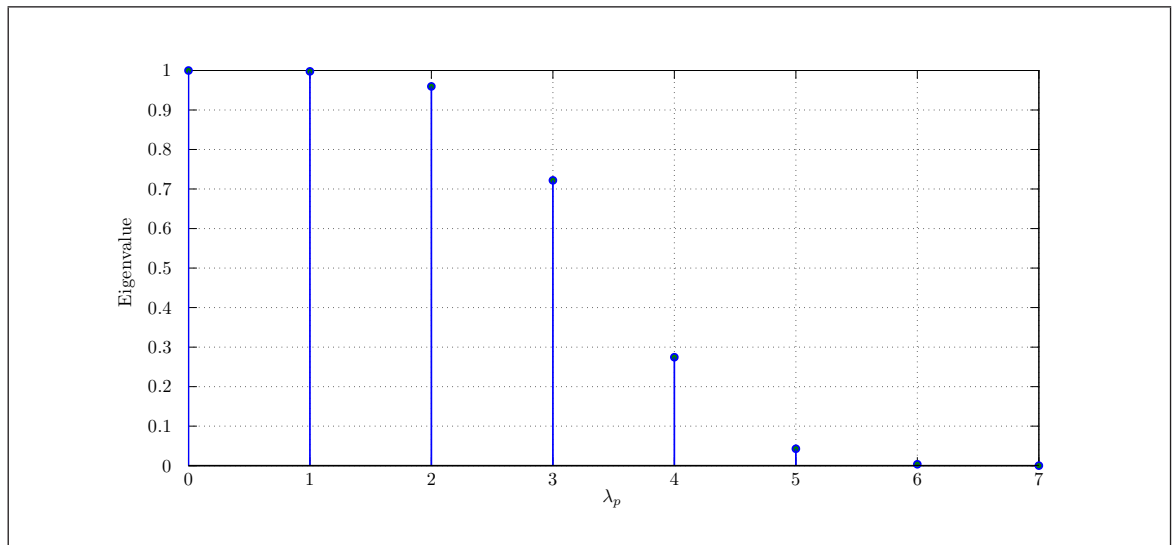


Figure 4.2.4 First eight eigenvalues for $N = 64$ and $NW_h = 2$

Since this matrix is positive-definite, the eigenvalues must be positive. Each eigenvalue has a corresponding eigenvector. For example the $\lambda_0(N, W_h)$ eigenvalue has the zeroth-order DPSS taper associated with it. When these eigenvalues are sorted into descending order ($\lambda_0(N, W_h) > \lambda_1(N, W_h) > \dots > \lambda_{P-1}(N, W_h)$), the first $2NW_h$ eigenvalues have significant values close to 1 and the rest are close to 0. $2NW_h$ is the Shannon number. Since the eigenvalues are also spectral concentrations, it follows that only the first $2NW_h$ DPSS tapers have good spectral leakage properties and concentrate power within the frequency interval $[-W_h, W_h]$ rad/ π . In fact in [128] to be cautious, it is recommended that only the first $2NW_h - 1$ or $2NW_h - 2$ DPSS tapers be selected for use.

As an example, with the use of the parameters $N = 64$ and $NW_h = 2$, the first eight eigenvalues are shown in Fig 4.2.4. The Shannon number is apparent as the first $2NW_h = 4$ eigenvalues are closer to 1 than 0 and thereafter the eigenvalues are closer to 0 than 1. The reason for the recommendation to use $2NW_h - 1$ or $2NW_h - 2$ tapers is also clear because the indexed eigenvalues after this number drop more rapidly towards 0. The DPSS tapers which correspond to the first eight eigenvalues are displayed in the left-hand side of Figs 4.2.5 and 4.2.6. The frequency-domain responses of the tapers

are displayed alongside each of them in the right-hand column of Figs 4.2.5 and 4.2.6. The dotted red lines indicate the design half-bandwidth $W_h = 0.0313 \text{ rad}/\pi$ in which it would be desirable if the frequency response were mostly concentrated.

The relationship between the spectral concentrations (eigenvalues) and the DPSS tapers is obvious upon examination of Figs 4.2.4–4.2.6. As the eigenvalues fall away from the value of 1, the corresponding DPSS tapers becomes more pronounced at the taper edges and the amount of spectral leakage outside the design half-bandwidth W_h becomes greater.

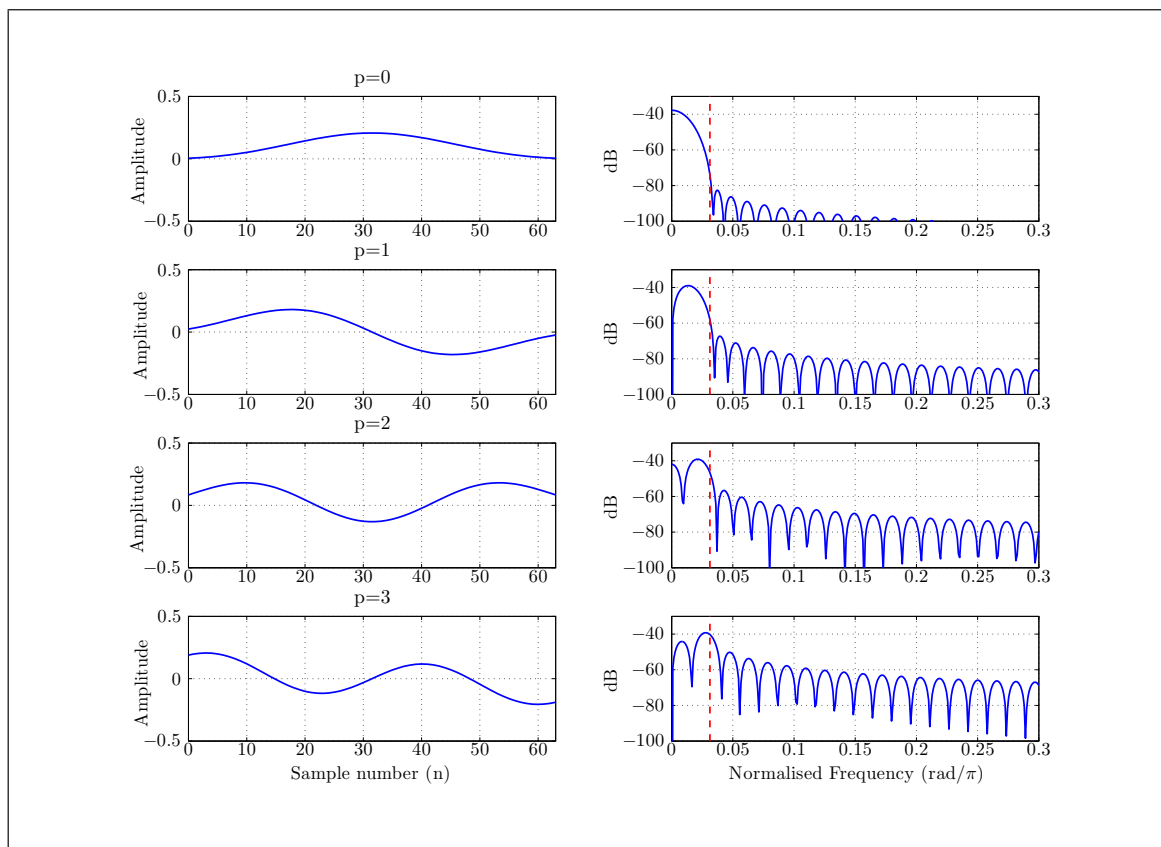


Figure 4.2.5 With $N = 64$ and $NW_h = 2$, DPSS tapers of orders 0, 1, 2, 3 from top to bottom on left-hand side, corresponding frequency-domain responses on right-hand side

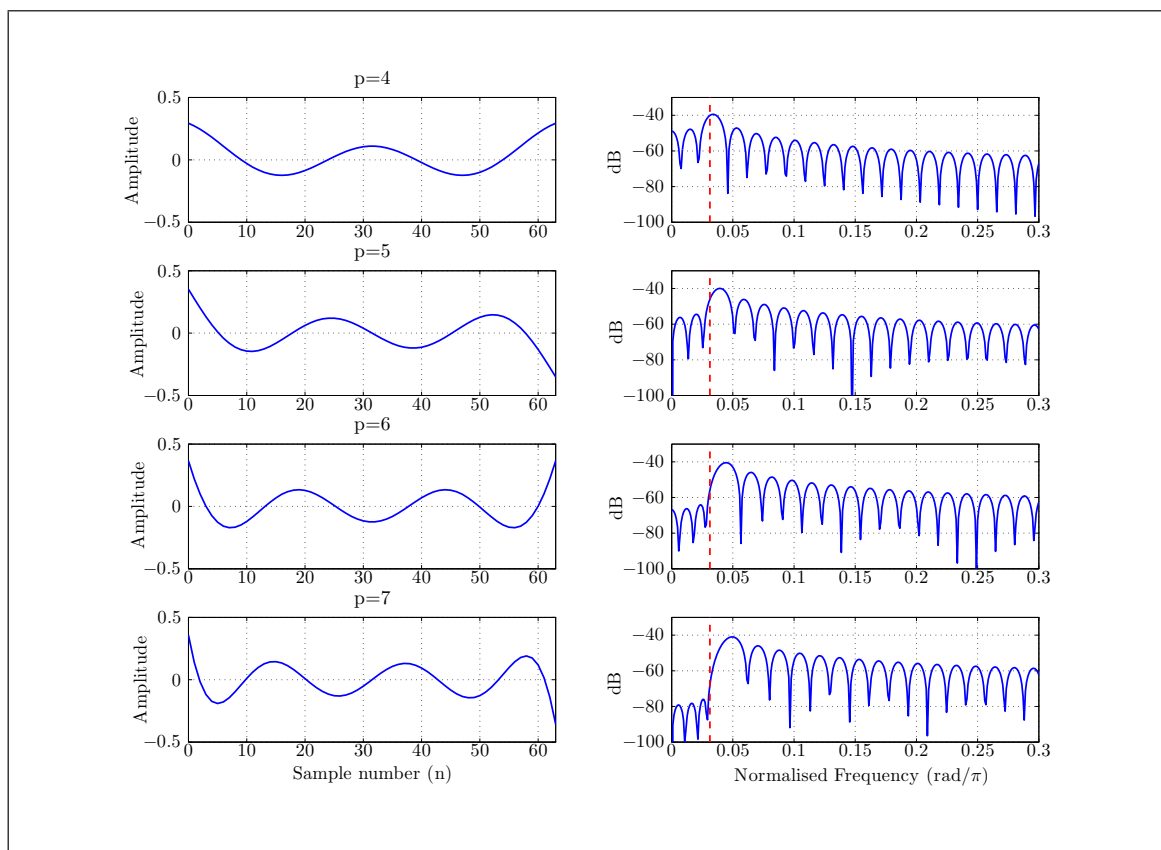


Figure 4.2.6 With $N = 64$ and $NW_h = 2$, DPSS tapers of orders 4, 5, 6, 7 from top to bottom on left-hand side, corresponding frequency-domain responses on right-hand side

4.2.5 DPSS Windows

The MTM allows for a reduction in variance over the periodogram through an average of P spectral estimates. A reduction in variance is important because it allows a lower detection threshold to be set for a given P_{fa} . In order for the variance to be reduced by an approximate factor of $1/P$, the spectral estimates need to be pairwise uncorrelated with each other. For this to be possible the tapers need to be orthogonal. The DPSS tapers fulfill this requirement. They are characterised by the half-bandwidth parameter W_h and the number of discrete points in the sequence N .

The DPSS tapers can be found from the direct solution of Eq (4.2.6) for eigenvectors $\mathbf{w}_0(N, W_h), \mathbf{w}_1(N, W_h), \dots, \mathbf{w}_{P-1}(N, W_h)$ subject to the constraints that:

- they are normalised: $\sum_{n=0}^{N-1} w_p(n)^2 = 1$ (along with the orthogonality condition this makes them orthonormal);
- the even-order tapers have a positive average: $\sum_{n=0}^{N-1} w_p(n) > 0$;
- the odd-order tapers start with a positive lobe: $\sum_{n=0}^{N-1} (N - 1 - 2n) w_p(n) > 0$.

However direct solution of Eq (4.2.6) is non-trivial. For example if a mathematical software tool is used numerical precision needs to be carefully monitored and the rate of convergence to a solution for the higher-order eigenvalues can be slow.

In [128], an alternative method for the calculation of the DPSS tapers is suggested which utilises numerical integration. However this method can also suffer from precision difficulties. A third method from [132] involves the use of a tridiagonal matrix from which it is possible to solve for only the required eigenvectors. This is generally the method that is used to find the DPSS taper coefficients in modern software tools.

Regardless of the method used to find the tapers it would be far too demanding for a real-life DRx in an ESM system to calculate the DPSS tapers to use for every DCW. A more feasible scheme would involve pre-computation of the DPSS tapers for a variety of N and W_h parameters and store the results in a look-up table (LUT). The storage requirements would be predictable because values for the N parameter would be chosen as the values which the FFT core in the DRx can assume, for example $N_{\text{FFT}} = 64, 128, 256, 512, 1024, 2048$. Likewise the W_h parameter would be chosen as a consequence of the operational environment in which the DRx is likely exist and what signals it may be required to intercept.

4.2.6 Sine Tapers

There are 2 practical disadvantages involved with the MTM when DPSS tapers are used:

1. The DPSS tapers are required to be calculated for each combination of the parameters N and W_h .
2. Each DPSS taper requires an FFT operation.

For large numbers of tapers, these two points create a significant computational burden. The first point can be mitigated in a real-life DRx through precalculation of the tapers for all possible combinations of N and W_h which the DRx may encounter.

There exists a lesser-known MTM which can overcome the problem stated in the second point. This MTM utilises **sine tapers** which have a sparse representation in the frequency-domain. Therefore in a similar fashion to SVA, the sine tapers can be applied in the frequency-domain via convolution for a relatively small computational cost. Therefore only one FFT operation per DCW is necessary. The sine tapers briefly feature in [133], but its unique advantages over the DPSS tapers are not explored. Additionally, the only reference to the use of the sine tapers in the field of radar and EW is [134]. Although useful, this analysis of radar signals with the use of the sine taper MTM does not utilise the sparse frequency-domain representation of the sine tapers nor mention the other important property of the sine taper MTM in that different frequency bins can have different numbers of tapers. These advantages are explored in chapter 5.

The sine tapers were originally formulated in [129] in order to produce tapers which create a minimum in “local bias” within the bandwidth of interest $[-W_h, W_h]$ rad/ π at the expense of slightly greater spectral leakage than the DPSS tapers. The fortunate property of the sine tapers is that they have an analytical expression and do not require calculation from eigenvalue decomposition unlike the DPSS tapers. The time-domain equation for the orthogonal set of P normalised sine tapers is given by¹

$$w_p(n) = \sqrt{\frac{2}{N}} \sin\left(\frac{\pi(p+1)n}{N}\right), \quad n = 0, 1, \dots, N-1 \quad (4.2.7)$$

The prefactor at the front of the right-hand side of Eq (4.2.7) normalises the sine tapers. In the frequency-domain the sine tapers have the response

$$W_p(k) = \delta\left(\frac{k+p}{2}\right) - \delta\left(\frac{k-p}{2}\right), \quad k = 0, 1, 2, \dots, N_{\text{FFT}} - 1 \quad (4.2.8)$$

Therefore if a DCW is twice-zero padded, rectangular-windowed and an FFT applied, it is possible to apply as many or as few sine tapers as desired at each frequency bin in the frequency-domain via a two point circular convolution. For example if the

¹Often $N+1$ and $n = 1, 2, \dots, N$ are used in the definition of the sine tapers to: (1) avoid window functions with zero valued samples (which throw away samples for little apparent gain), (2) make the odd-order windows symmetric and (3) give the even-order windows a 180° rotational symmetry. To recover these definitions replace N with $N+1$ and adjust the indices n and k in Eqs (4.2.7) and (4.2.8)

eigenspectrum from the zeroth-order sine taper was desired then a length $N_{\text{FFT}} = 2N$ FFT would be performed on a rectangular-windowed DCW and each frequency bin would be replaced by the complex difference between its neighbour to the left and its neighbour to the right. The magnitude-squared would be calculated to create the eigenspectrum. The first four orders of sine tapers characterised by $N = 128$ sample points are shown in Fig 4.2.7. The left-hand column, (a), shows the time-domain representation of the sine tapers and the right-hand column, (b), shows the imaginary part of the corresponding frequency-domain representation. The frequency-domain representation has been shifted so that the DC term is in the centre of the plot and the index k has been renumbered from $k = 0, 1, \dots, N_{\text{FFT}} - 1$ to $k = -N_{\text{FFT}}/2, -N_{\text{FFT}}/2 + 1, \dots, N_{\text{FFT}}/2 - 1$.

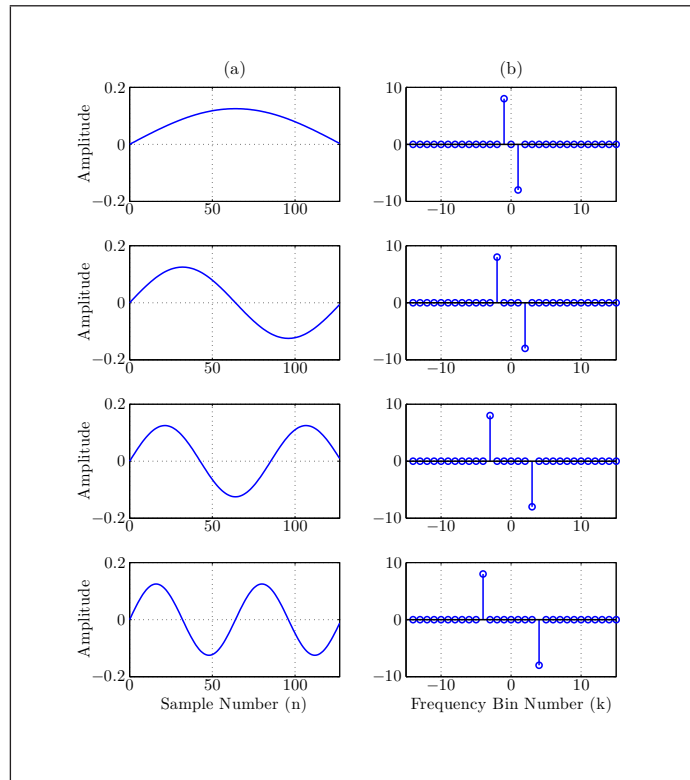


Figure 4.2.7 Sine tapers for first four orders in: (a) time-domain; (b) frequency-domain

As before in Eq (4.2.1), eigenspectra $\hat{S}_{p,xx}^{(\text{MTM})}(k)$ are averaged to get the full MTM spectral estimate:

$$\hat{S}_{xx}^{(\text{MTM})}(k) = \frac{1}{2P} \sum_{p=0}^{P-1} \left| X\left(\frac{k+p}{2}\right) - X\left(\frac{k-p}{2}\right) \right|^2 \quad (4.2.9)$$

The minimum bias property of the sine tapers is a little more obvious from Eq (4.2.9) as the frequency sidelobe from $X\left(\frac{k+p}{2}\right)$ cancels with the sidelobe of $X\left(\frac{k-p}{2}\right)$ as described in [129].

There are two important points to draw from the use of the sine tapers for practical applications in real-life detectors:

1. As the tapers are applied in the frequency-domain via convolution, only one FFT on a twice-zero-padded, rectangular-windowed DCW is required to be calculated.
2. It is optional to execute the convolution at every frequency bin. If desired, only the bins of interest can have as many or as few tapers applied to them.

These points provide opportunities for computational savings in a real-life DRx which utilises the sine taper MTM. The second point in particular highlights the flexibility of the sine taper MTM in that the DPSS MTM is for fixed bandwidth estimation only but different numbers of sine tapers can be used at different bins in the same spectral estimate. This is because the DPSS tapers are functions of the half-bandwidth parameter W_h whereas the sine tapers are not.

The possibility of using a variable number of tapers in the frequency-domain and hence a variable resolution and error across frequency was exploited in the original formulation of the sine tapers in [129]. In the reference, the optimal number of tapers, in the sense of minimisation of the local loss (within $[-W_h, W_h]$), was found to be

$$P_{opt} \sim \left[\frac{12S_{xx}(k)N^2}{S''_{xx}(k)} \right]^{2/5} \quad (4.2.10)$$

Equation (4.2.10) states that fewer tapers should be used when the true spectrum varies rapidly from bin to bin. In a real-life DRx, the true discrete PSD ($S_{xx}(k)$) and its second derivative ($S''_{xx}(k)$) would both be unknown and therefore direct use of Eq (4.2.10) would not be possible. However a spectral estimate with a variable number of tapers in the frequency-domain can be achieved by performing a mean square error (MSE) adaptive estimation. This entails: (1) a pilot spectral estimate with an appropriately low number of tapers $\hat{S}_{xx}^{(MTM)}(k)$, (2) estimation of $S''_{xx}(k)$ and (3) utilisation of a form of Eq (4.2.10). This would provide an estimate of the optimum number of tapers at each bin, which in turn provides an updated estimate of $S_{xx}(k)$, from which $S''_{xx}(k)$ can be estimated to insert back into Eq (4.2.10). The process is iterated several times to converge on a solution for the optimal number of tapers P_{opt} and hence a good estimation of the true discrete PSD ($S_{xx}(k)$).

Where the adaptive sine taper MTM has been used in this work, the estimation of $S''(k)$ has followed a similar procedure to that in [135]. This involves a quadratic fit to $\ln S_{xx}(k)$ and rather than the estimation of $|S_{xx}(k)/S''_{xx}(k)|$, the use of

$$\frac{1}{Y''(k) + Y'(k)^2}, \quad \text{where } Y(k) = \ln S_{xx}(k)$$

In this work, some approximation of the algorithm has been used to:

1. provide an upper limit on the number of tapers so that the time-bandwidth product does not exceed $(N/2 - 1)$;
2. overcome points where $S_{xx}''(k)$ vanishes (which is fairly common in most spectra);
3. manage the growth of the number of tapers so that frequency intervals are not engulfed by neighbour frequency intervals.

All of the same conclusions with respect to the MTM but not specific to the DPSS tapers hold. These include that for an AWGN input, the output from the uniformly-averaged sine taper MTM follows a chi-square distribution and eigenspectra within the MTM spectral estimate can be individually weighted. In [129] parabolic weights for the eigenspectra are proposed as the preferred weight scheme for the subject matter (nuclear physics) in order to decrease the weights smoothly down to zero.

An example comparison of the adaptive sine taper MTM against uniformly-weighted eigenspectra and parabolically-weighted eigenspectra is shown in Fig 4.2.8. The use of the parabolic weight scheme tends to results in more tapers at each frequency bin in the adaptive algorithm.

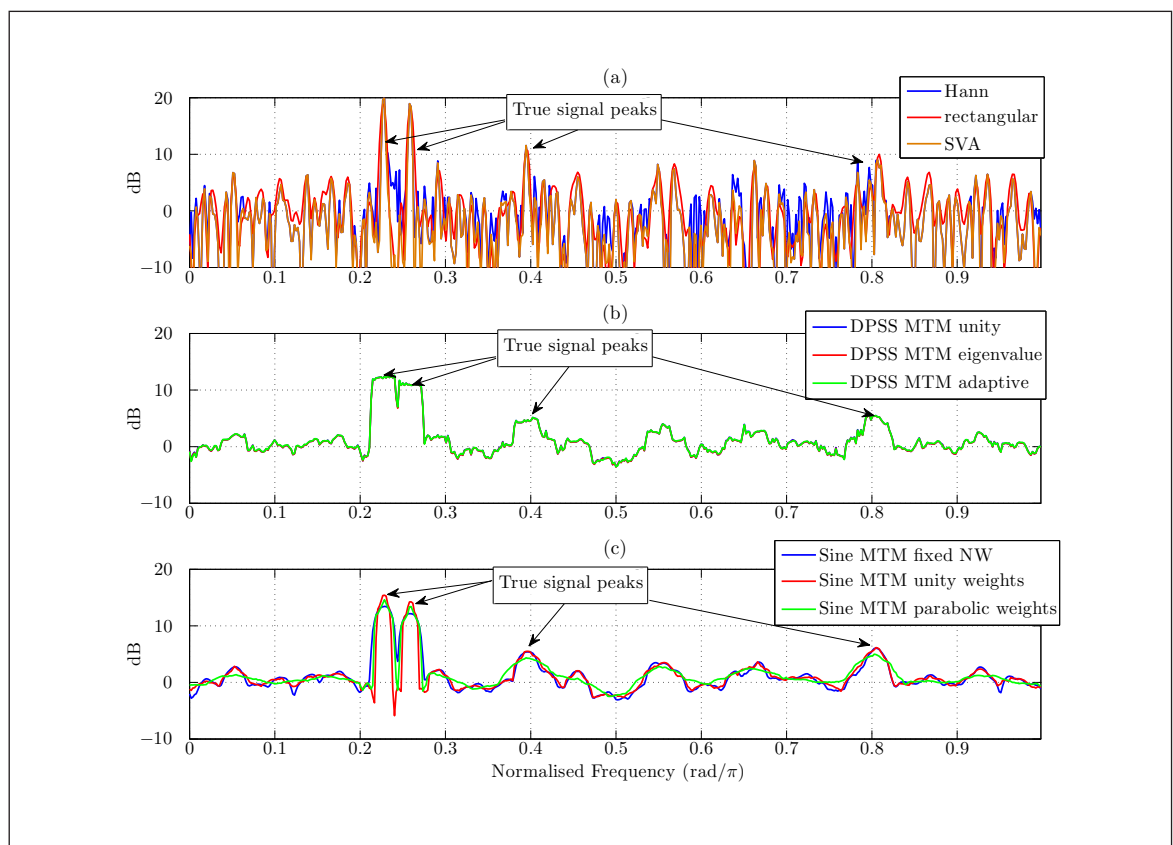


Figure 4.2.8 Spectral estimates of same DCW with use of: (a) periodogram, modified periodogram and SVA; (b) DPSS MTM variants; (c) sine taper MTM variants

The sine tapers are more convenient to calculate and manipulate than the DPSS tapers and according to [108], DPSS tapers are relevant only for nearly white signals. However in [136] the sine tapers are not recommended for use on processes which have highly

coloured background noise or involve a very large dynamic range due to the inferior spectral leakage properties of the sine tapers over the DPSS tapers.

An example of spectral estimates from different MTMs is given in Fig 4.2.8. Also included for comparison are two conventional spectral analysis techniques and SVA in panel (a). The conventional spectral analysis techniques are simply formed from the rectangular-windowed and Hann-windowed FFT, in other words, the periodogram and modified periodogram respectively.

To create the spectral estimates a particular test vector was used to form a length $N = 256$ DCW. This test vector consisted of four signals and an AWGN background. The signal parameters are listed in Table 4.2.1. All the signals were created so that they were present throughout the duration of the DCW. The signals were chosen such that S_{sig1} and S_{sig2} were relatively strong and close together in frequency, S_{sig3} was relatively weak and further apart from the S_{sig1} and S_{sig2} group and S_{sig4} was broader-band. The length of FFT used in all the estimates was $N_{\text{FFT}} = 512$.

Signal 1:	Type	Sinusoid (S_{sig1})
	SNR_{out}	20 dB
	Frequency	$0.2278 \text{ rad}/\pi$
	Initial phase	0.9950 rad
Signal 2:	Type	Sinusoid (S_{sig2})
	SNR_{out}	20 dB
	Frequency	$0.2587 \text{ rad}/\pi$
	Initial phase	5.2842 rad
Signal 3:	Type	Sinusoid (S_{sig3})
	SNR_{out}	12 dB
	Frequency	$0.3954 \text{ rad}/\pi$
	Initial phase	4.9574 rad
Signal 4:	Type	Pseudo-noise (S_{sig4})
	SNR_{out}	15 dB
	Target centre frequency	$0.7972 \text{ rad}/\pi$
	Initial phase	5.680 rad
	Chip-rate	$0.0287 \text{ rad}/\pi$

Table 4.2.1 MTM test vector parameters

In panel (a) the large variance of the periodogram, modified periodogram and SVA technique are clearly visible as the estimate is extremely jagged. The peaks of the sinusoidal signals (S_{sig1} , S_{sig2} and S_{sig3}) are apparent but S_{sig4} is over-resolved as it is not clear if it is broader band than the sinusoids. There are also many other peaks due to the AWGN background which would be erroneously marked as additional detections in a detector which relied on these estimation techniques. Also of note in panel (a)

is the spectral leakage between the peaks of S_{sig1} and S_{sig2} for the periodogram. The modified periodogram and SVA-periodogram do not suffer from this spectral leakage.

In panel (b), the DPSS MTMs were estimated with the use of tapers calculated for $NW_h = 4$, $W_h = 0.0156 \text{ rad}/\pi$, i.e. with the use of $2NW_h - 1 = 7$ tapers (the eighth taper was dropped). This was close to but not an exact match to the half bandwidth of S_{sig4} ($W_{h,\text{sig4}} = 0.0143 \text{ rad}/\pi$). In reality, the NW_h parameter must be estimated according to a heuristic or automatic approach. The heuristic approach is used for the MTMs in this thesis for computational efficiency reasons, however automatic approaches such as described in [137] could be used.

There is very little difference between the DPSS MTM spectral estimates. The four largest humps correspond to the four signals. The resolution of the DPSS MTM of $P/N \text{ rad}/\pi$ is also noticeable as the sinusoidal peaks appear flat-topped. There are additional peaks but these are far fewer in number than in panel (a) as the variance is reduced by approximately $\sim \sigma^2/7$, which causes the appearance of the noise floor to be much flatter.

With the use of $P = 7$ tapers in the fixed NW_h sine taper MTM spectral estimate, the differences between the DPSS taper MTMs and the sine taper MTM is evident. Use of the sine taper MTM results in spectral peaks which are more rounded (smaller bias to the actual spectral peak shape) and the spectral estimate appears smoother. As with the DPSS MTM, the four largest humps in the sine taper MTMs correspond to the actual four signals and the noise floor is fairly flat.

The variable NW_h sine taper MTMs result in sharper spectral peaks of S_{sig1} and S_{sig2} as fewer tapers are used around those points. Of the sine taper MTMs, the method which uniformly-weighted the eigenspectra appears to work best as it better preserves the peaks of S_{sig3} and S_{sig4} . The parabolic weight scheme appears to smooth the spectral estimate too much away from S_{sig1} and S_{sig2} .

4.3 Smoothed Periodogram

In section 4.2 the sine taper MTM is highlighted as a computationally efficient flavour of the MTM. It was declared in section 4 that the periodogram smoothing method is also a computationally simple technique to obtain a spectral estimate with a lower variance than that of the periodogram. Hence some consideration is given to it in this section so that its use in a detector can be compared against a detector which employs a MTM.

For completeness the WOSA method has been discounted as a spectral estimator for two reasons: (1) the actual target signal duration is assumed unknown and subdividing the DCW into smaller segments may worsen the spectral estimate and (2) the variance reduction is not proportional to the number of segments being averaged due to the

correlation between noise in overlapped samples.

Periodogram smoothing has its origins with the Blackman-Tukey method. In accordance with the Wiener-Khinchin theorem, the true PSD is given by the Fourier transform of the autocorrelation function. However in a real-life DRx the autocorrelation function is approximated by the sample (biased) autocorrelation $\hat{R}_{xx}(l)$ given in Eq (4.1.2). Unfortunately the number of terms used in Eq (4.1.2) at various lags (l) decreases as the value of $|l|$ gets larger. Ultimately at the ends of the sample autocorrelation only one term is used to obtain an estimated data point. This is the origin of the large variance of the periodogram. To deal with this, the Blackman-Tukey method ($\hat{S}_{xx}^{(BT)}$) deemphasises the ends of the sample autocorrelation through multiplication of $\hat{R}_{xx}(l)$ by a lag window as²

$$\hat{S}_{xx}^{(BT)}(k) = \sum_{l=-(N-1)}^{N-1} \hat{R}_{xx}(l)w(l)W_N^{-kn}, \quad k = 0, 1, 2, \dots, N_{\text{FFT}} - 1; \quad (4.3.1)$$

$$n = 0, 1, 2, \dots, N - 1.$$

By the convolution theorem, Eq (4.3.1) in the frequency-domain is

$$\hat{S}_{xx}^{(SP)}(k) = \mathcal{F} \left\{ \hat{R}_{xx}(k) \right\} * \mathcal{F} \left\{ w(l) \right\}$$

which upon use of the step from Eq (4.1.3) to Eq (4.1.4) yields

$$\hat{S}_{xx}^{(SP)}(k) = \frac{1}{N} |X(k)|^2 * W(k) \quad (4.3.2)$$

where $W(k)$ is the DFT of the lag window ($w(l)$).

The choice of coefficients for $w(l)$ has a direct effect on the convolution in the frequency-domain. If suitable weights are chosen for $w(l)$, then the convolution could be implemented with the use of a simple FIR filter.

Equation (4.3.2) shows a smoothed periodogram ($\hat{S}_{xx}^{(SP)}$) is achieved through application of filters of various spans to the periodogram or one of its variants such as the modified periodogram. The simplest such filter is a moving average filter. The degree of smoothness, frequency resolution and variance of the final spectral estimate are controlled by the properties of the filter. Longer filters provide smoother spectral estimates, whilst narrower filters provide rougher estimates. There is a balance to be struck, for, if too long a filter is used then signal peaks may become merged, if too short a filter is used then the detection threshold could not be lowered significantly.

Periodogram smoothing in this fashion closely resembles the sine taper MTM both

²The unbiased sample autocorrelation, $\hat{R}_u(l) = \frac{1}{N-|l|} \sum_{n=0}^{N-1-|l|} x(n)x(n+l)$, is not used as a means to deemphasise the end points because its use in the Wiener-Khinchin theorem results in a spectral estimate with an even larger variance

in terms of the mathematical equations and computational load (see Table 4.3.1). Therefore it is better to compare the MTMs against this method than it would be to compare the MTMs against the WOSA technique. In fact, as is shown in Table 4.3.1, spectral smoothing in this way is even more computationally efficient than the best MTM and is easier to implement. Therefore the comparison of detectors which use the MTMs and periodogram smoothing is of great interest. For example, if there is only a small advantage in the use of the MTMs over periodogram smoothing, then it is not worth the extra effort to implement the MTM. This is the question which is investigated in chapter 5.

One interpretation of periodogram smoothing is that it is like the power summation of neighbour frequency ordinates and concentrates the result in the ordinate under test. This is similar to the MTM although the actual filter shape does not resemble a brick-wall shape as strongly as in the MTMs. As was mentioned in section 4.2.2, the filter shape interpretation provides extra motivation for the study of the smoothed periodogram because it is like the use of a filter whose shape is a better match to broader-band LPI radar signals.

For a selected half bandwidth $W_h \approx 0.0352 \text{ rad}/\pi$, the equivalent average response for various spectral estimators are shown in Fig 4.3.1. The notion would be to choose the response to match the bandwidth of an LPI radar signal. The plot series shown in Fig 4.3.1 were created from: (1) the sine taper MTM with uniform weights, (2) the periodogram smoothing method with the use of a moving average filter on a twice-zeropadded FFT, (3) the DPSS MTM with uniform weights and (4) the periodogram smoothing method with the use of a moving average filter on a non-zeropadded FFT. The length of the time series was $N = 128$ samples and for the MTMs the number of tapers used was $2NW_h - 1 = 8$. The smoothed periodograms used moving average filters of span $L = 8$.

The advantages and disadvantages of each spectral estimator can be extracted from Fig 4.3.1. The cut-off frequency is related to the spectral leakage of the method, the passband ripple is related to the local bias and the stopband attenuation is related to the broad band bias. The sine taper MTM has a similar cut-off frequency to the smoothed periodogram methods and good broad band bias properties. The non-zeropadded smoothed periodogram has poor local bias but reasonable broad band bias. The zeropadded smoothed periodogram improves the passband ripple at the expense of the stopband attenuation. The DPSS MTM has the best transition rate from passband to stopband as would be expected because the DPSS tapers offer the best spectral leakage protection.

By smoothing a periodogram with a moving average filter, effectively a summation of independent values is performed. Therefore the discussion on chi-square distributed

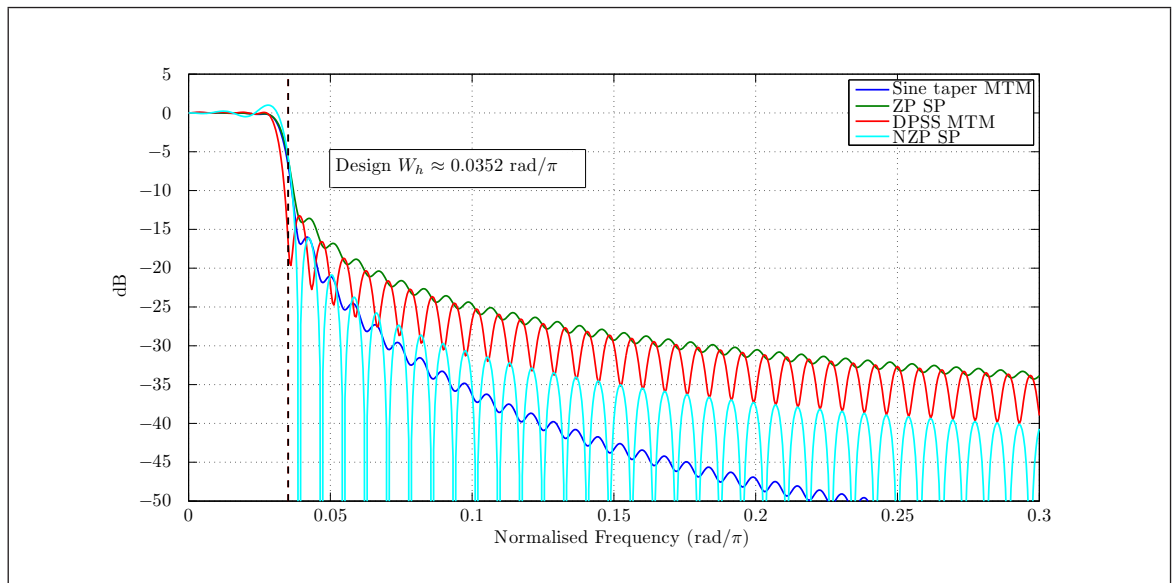


Figure 4.3.1 $W_h \approx 0.0352 \text{ rad}/\pi$ filter responses for different spectral estimators: (1) sine taper MTM; (2) twice-zeropadded smoothed periodogram; (3) DPSS MTM; (4) non-zeropadded smoothed periodogram

RVs in section 4.2.2 applies. For example, if the input to the periodogram is complex AWGN, the moving average filter outputs applied to the periodogram follow the scaled chi-square distribution from Eq (4.2.2) with $2L$ degrees of freedom. As with the uniformly-weighted MTMs, this theoretical distribution is useful for the determination of detection thresholds. However the output from the smoothing filters is not chi-square distributed if the modified periodogram or SVA periodogram are smoothed. This is because the ordinates are not independent. An additional condition to ensure the filter outputs follow a chi-square distribution is that the filter is required to have uniform coefficients (like the moving average filter).

4.3.1 Tree Smoothed Periodogram

This section describes a scheme for smoothing the periodogram with a primitive filter. The motivation to devise such a scheme is to reduce the computational load for a potential implementation of a smoothed periodogram. This is achieved through restrictions on the filter design and hence limitations on the different possible bandwidths which can be heuristically tested.

In this scheme, the filter span (L) is limited to two frequency bins, however this filter can be implemented in layers such that the output from one layer is used as the input to the next. Fewer possible bandwidths can be tested as a result in comparison to the moving average filter periodogram smoothing method because the possible bandwidths increase as the power of two.

With a filter span of two bins at each layer, the same filter is reused at successive layers. Not only does this cut down on the resources required in an implementation, but the previous layer results form the summation for the next layer. Additionally with a filter

span of two there are further computational savings because the normalisation of the summation by $1/2$ can be implemented with fixed-point barrel shifts.

The reuse of the results from the previous layer is the reason for the name **tree smoothed periodogram (TSP)**. For clarity, there is no downsampling by a factor of two after each layer. A comparison of the smoothing filter methodologies is shown in Fig 4.3.2. Panel (a) illustrates the concept behind periodogram smoothing with the use of a moving average filter. A filter is selected with a chosen span and a new value is calculated for each frequency bin as the filter is applied across the spectral estimate. Meanwhile panel (b) illustrates that the TSP can be composed of a multitude of filters which average two adjacent frequency bins in a series of filters. The filters are, like the moving average filter, applied across the periodogram to produce a new value for each frequency bin.

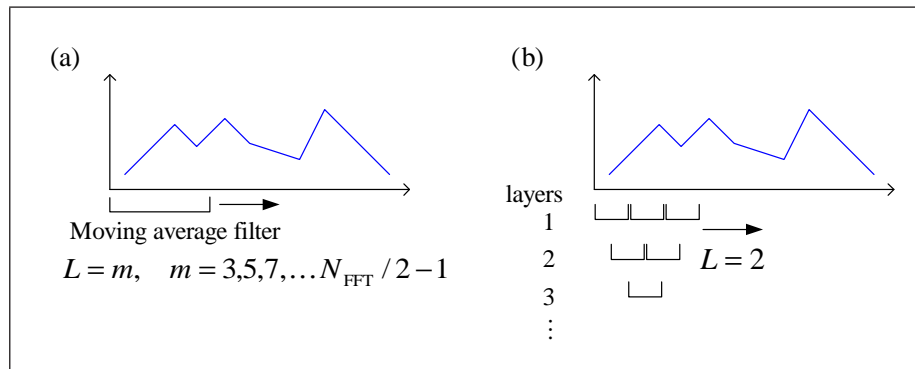


Figure 4.3.2 Methodologies of: (a) smoothed periodogram with the use of a moving average filter; (b) TSP

The loss in bandwidth resolution may not impact greatly on the performance of a real-life detector because for the tree periodogram smoothing method the bandwidth resolution is finer for small possible bandwidths and this is where LPI radar signals are likely to be found in a superhet-DRx combination.

One consequence of the average of an even number of ordinates (in this case two) is the filter output at each layer is offset by the equivalent of $0.5 \times f_s / N_{\text{FFT}}$ and this needs to be considered in the results.

The same conclusion with respect to the smoothed periodogram and chi-square distribution also holds for the TSP, since it is simply a different way of summing ordinates. Therefore if a periodogram of complex AWGN is calculated (as opposed to a modified periodogram or SVA periodogram) the ordinate values follow a chi-square distribution with 2^{hyp+1} degrees of freedom, where *hyp* is shorthand for a hypothesis layer number.

4.3.2 Complexity

In this section a brief analysis of the computational complexity of some of the simplest MTMs and smoothed periodogram methods is given. The methods analysed

are: (1) uniformly-weighted DPSS MTM, (2) uniformly-weighted sine taper MTM, (3) smoothed periodogram with the use of a moving average filter and (4) TSP.

Table 4.3.1 displays the number of operations required for each of these methods. N denotes the number of samples in a DCW and P denotes the number of eigenspectra required or the number of smoothed periodograms required, each with a different test half-bandwidth parameter W_h . Although the TSP method does not have the same bandwidth resolution as the other methods the number of operations is given for comparative purposes **as if** it was desired to have P eigenspectra.

The number of operations is defined as the number of real additions/subtractions and multiplies, but not trivial operations such as barrel shifts. There are amendments required to the details in Table 4.3.1 if other methods are used. For example smoothing a modified periodogram with a moving average filter adds N operations to the total in the periodogram smoothing column as the length N DCW would be required to be multiplied by a window function.

For the uniformly-weighted DPSS MTM each term in Table 4.3.1 corresponds to a required operation. These are to:

1. multiply the DCW by P tapers;
2. perform P FFTs on the windowed-DCWs;
3. take the magnitude-squared of each FFT output;
4. sum the eigenspectra;
5. normalise the spectral estimate.

Similarly, for the uniformly-weighted sine taper MTM, the operations are to:

1. perform one twice-zeroadded FFT;
2. apply P sine tapers by convolution in the frequency-domain;
3. form the eigenspectra through the magnitude-squared of each output;
4. sum the eigenspectra.

Also, for the periodogram smoothed with a moving average filter:

1. perform an FFT;
2. take the magnitude-squared of the output;
3. filter;
4. normalise the spectral estimate.

Lastly, for the TSP:

1. perform an FFT;

2. take the magnitude-squared of the output;
3. filter.

DPSS MTM	Sine Taper MTM	Periodogram Smoothing	Tree Periodogram Smoothing
PN $+5PN \log_2(N)$ $+3PN$ $+PN$ $+N$	$10N \log_2(2N)$ $+10PN$ $+2PN$ $+2N$	$5N \log_2(N)$ $+3N$ $+2PN$ $+N$	$5N \log_2(N)$ $+3N$ $+N \text{ round } \{\log_2(P)\}$

Table 4.3.1 Algorithmic complexity ($O(\dots)$) of MTMs and periodogram smoothing methods

Table 4.3.1, for realistic N and P , in order of increasing computational complexity is: (1) TSP; (2) smoothed periodogram (3) sine taper MTM; (4) DPSS and MTM. This complexity was a consideration which affected the decision of which algorithm to implement in fixed-point and on field-programmable gate array (FPGA) in appendices H and I.

4.3.3 Bandwidth Hypothesis Testing

The ability of the MTMs and periodogram smoothing techniques to detect broader-band LPI signals more successfully than basic periodogram techniques is demonstrated in Fig 4.3.3. The picture shows how the MTMs may be used at greater and greater hypothesis bandwidths ($2W_h$) or the smoothing filters may be used at greater and greater spans (L). Around the approximate frequency and bandwidth combination of the signal, detection is more likely.

Fig 4.3.3 was created from the accumulation of detections at all possible hypothesis bandwidths applied to one thousand length $N = 128$ full DCWs. The input LPI waveform was a CW m-sequence signal at $\text{SNR}_{\text{in}} = -10$ dB and the processing technique utilised was the twice-zero-padded unity-weighted DPSS MTM. The signal parameters were $f_T \approx 0.36 \text{ rad}/\pi$ and $B \approx 0.12 \text{ rad}/\pi$, but otherwise random. The ordinates at the output of DPSS MTM were thresholded at the hypothesis bandwidths ($W_h = [0, 0.992] \text{ rad}/\pi$, i.e. 1–127 tapers) such that the P_{fa} for an AWGN input would be 1×10^{-6} . The thresholds were calculated from the scaled chi-square distribution (see Eq (4.2.2)). The detections after each DCW trial were accumulated to produce the plot and are shown to indicate the correct bandwidth and frequency, much like the clustering technique (section 1.3). The accumulation of the detections is akin to the non-coherent BMWD techniques described in section 1.6.

The bandwidth hypothesis tests increase the probability of detection of LPI signals and indicate the approximate frequency and bandwidth of the signal. The left-hand edge of Fig 4.3.3 shows that in this example no detections would be made with the

modified periodogram which is the non-parametric technique generally used in ESM DRxs. Similar plots can be produced for the other MTMs and periodogram smoothing techniques. This heuristic bandwidth hypothesis testing is the premise used in the detectors in section 5.

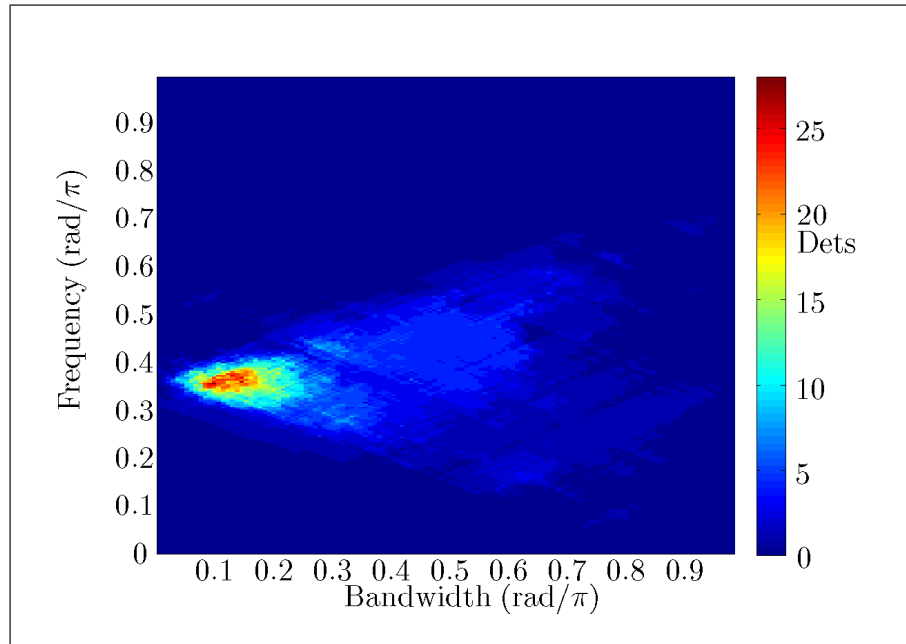


Figure 4.3.3 Clustering of detections with the use of frequency and bandwidth selectivity

4.4 Conclusions

As a result of its computational efficiency and lack of assumptions, non-parametric spectral estimation is essential for ESM systems to quickly detect threats. The MTMs and spectral smoothing are more appropriate for the detection of broader-band LPI radar signals than periodogram-based techniques because they can be better matched to the signal bandwidth.

The next chapter investigates the candidate MTMs and smoothing techniques against conventional ESM spectral detection methods.

CHAPTER 5

Floating-Point MTM & Periodogram Smoothing

5.1 Background

The previous chapter 4.2 claims that the MTMs and periodogram smoothing methods would be better suited to the detection of broader-band LPI radar signals, which would normally have their signal power spread across a few periodogram ordinates, because the equivalent average response shapes of these methods are better matched to broader-band LPI radar waveforms. The goals of this chapter are to provide an indication that these methods do indeed improve detector performance and to show their relative performance. Furthermore, as these methods vary in computational complexity it was essential to find the best candidate algorithm and balance its performance gains against the computational load it entails.

The intention was that the best candidate algorithm would be selected and taken further along the MBD methodology. Hence the selected simulated detector from this chapter represented the golden reference model. All models at this stage were created in the floating-point domain. The other candidate simulated detector models which were not chosen as the golden reference were dropped from the MBD methodology.

This section examines the performance of various simulated detectors for their ability to detect LPI waveforms. The detection of sinusoids was not investigated because the MTMs and periodogram smoothing methods are less well suited to this signal type. Rather, the conclusion from chapter 3 is that the SVA periodogram would be a favourable method to improve detection of sinusoidal radar signals. The LPI signal coding schemes inspected and reported in this section are:

Barker 13 Pulse This is a type of BPSK-coding which is commonly used due to the ease of implementation. This coding scheme was previously used in simulations in section 3.6. This falls into the category of a pseudo-noise signal.

Binary M-Sequence CW This is another type of BPSK-coding scheme which like the Barker-13 code is easily generated. As the code sequence generated increases, it provides an asymptotically ideal PAF with low sidelobes. As such, this coding scheme was used as a phase code in a CW waveform. This also falls into the

category of pseudo-noise signal.

This means one pulsed LPI waveform and one CW LPI waveform were tested. Whilst this is not enough to prove anything conclusively, it is a start. The detection scenarios deemed important for an ESM DRx which feature in this section are:

- full DCW
- randomly-filled DCW
- interference present in a full DCW
- randomly-filled DCW with interference present.

Suffice that the randomly-filled scenarios were not applicable for the CW waveform. Greater detail is given on each of these scenarios later in this chapter.

The simulated algorithms, which were of interest and were employed in each of the scenarios were:

- DPSS MTM
- sine taper MTM
- smoothed periodogram (SP)
- smoothed modified periodogram (SMP)
- smoothed SVA periodogram (SSP)
- tree smoothed periodogram (TSP)
- tree smoothed modified periodogram (TSMP)
- tree smoothed SVA periodogram (TSSP).

These candidate algorithms were expressly considered because it was conjectured they would provide better detection capability against the LPI signals than conventional detection techniques. Therefore some conventional algorithms were also used within simulated detectors for comparative purposes and these were:

- OR periodogram
- OR modified periodogram
- OR SVA periodogram
- periodogram
- modified periodogram
- SVA periodogram
- total power (TP).

The detectors (except those which used the the sine taper MTMs algorithms) were tested with both twice-zero padded and non-zero padded DCWs. Chiefly, this was done to examine if zero padding offered any advantage in the scenarios.

The sine taper MTMs were only tested with a twice-zero padded DCW because the formulation of the sine taper MTMs utilised was the computationally-efficient frequency-domain formulation which relied on convolution with a twice-zero padded DFT. Twice-zero padding the other detectors therefore also allowed for a fair comparison with the sine taper MTMs.

A short section on each detectors follows.

5.1.1 DPSS MTM Detector

The DPSS MTM is a function of two parameters: N and W_h . Whenever a new combination of N and W_h is required, calculation of a new set of tapers is required and the MTM must be performed with these tapers. In typical applications, the MTM is only calculated with the use of a small number of small-valued time-bandwidth products (NW_h) to enjoy its benefits. The examples of [131, p. 371] demonstrate this. An exhaustive search of all possible N and W_h combinations is seldom used. Thankfully this is actually appropriate for the detection of LPI signals. This is because LPI radar signals usually utilise a bandwidth of a few MHz in a large baseband bandwidth typically in the interval [1.0, 1.75] GHz in a DRx and therefore upon an FFT the LPI signal power is spread across only a few frequency bins.

Therefore the DPSS MTM detector searched for detections across multiple time-bandwidth hypotheses. The various time-bandwidth products tested were $NW_h = 0.5, 1.0, 1.5, \dots, 3.5$. This encompassed the actual NW_h s for the test LPI signals which were: $NW_h \approx 2.73$ for the Barker-13 pulse and $NW_h \approx 2.90$ for the m-sequence CW waveform. A diagram of how detection decisions were made is shown Fig 5.1.1. The diagram shows there were multiple hypotheses tested and an “OR” operation was performed on all of the logical decisions derived from each hypothesis NW_h . The multiple hypotheses are denoted with the shorthand *hyp*. The probability of false alarm was defined for the overall system rather than the probability of false alarm for the individual hypotheses. The probability of false alarm for the individual multiple hypotheses detectors were set to have equal individual probabilities of false alarm ($P_{fa, \text{individual}}$) subject to the overall system $P_{fa} = 0.01$. This was followed for all of the detectors which used multiple hypotheses.

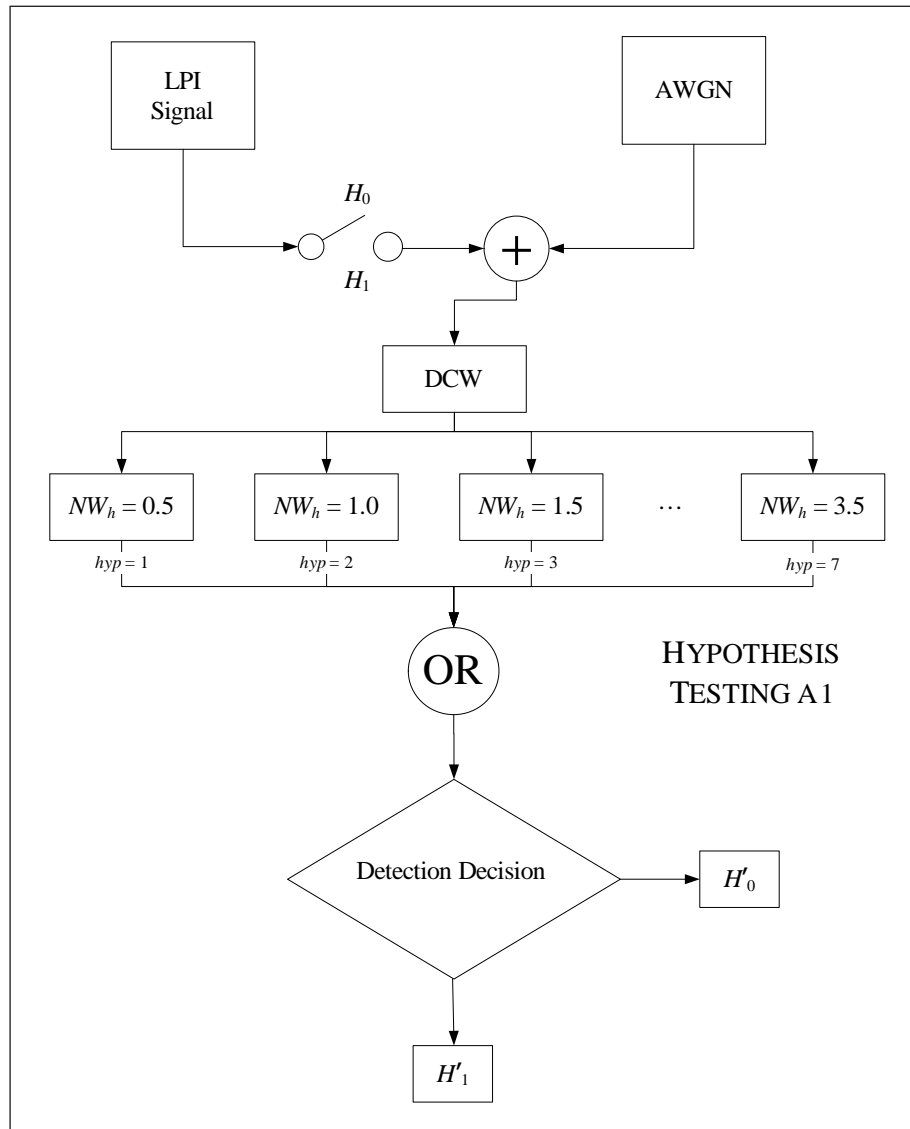


Figure 5.1.1 Tests at small-valued time-bandwidth products hypotheses

At each hypothesis NW_h , the simulated detector, represented as a system flowchart in Fig 5.1.2, was used. This detector architecture is called **detector G1**. The tapers required calculation at each hypothesis and there were $P = 2NW_h$ tapers required at each hypothesis. For example, at one extreme for the hypothesis of $NW_h = 0.5$, only one taper was used but at the other extreme of $NW_h = 3.5$, seven tapers were used. Only the results from the unity weight scheme are described in this work because, as demonstrated in Fig 4.2.8, the different DPSS MTM spectral estimates do not give noticeably different detector performances.

Although not a great concern in the floating-point simulations, each individual detector G1 requires to calculate P FFTs in a serial manner or it would require the resources to perform P FFTs in parallel.

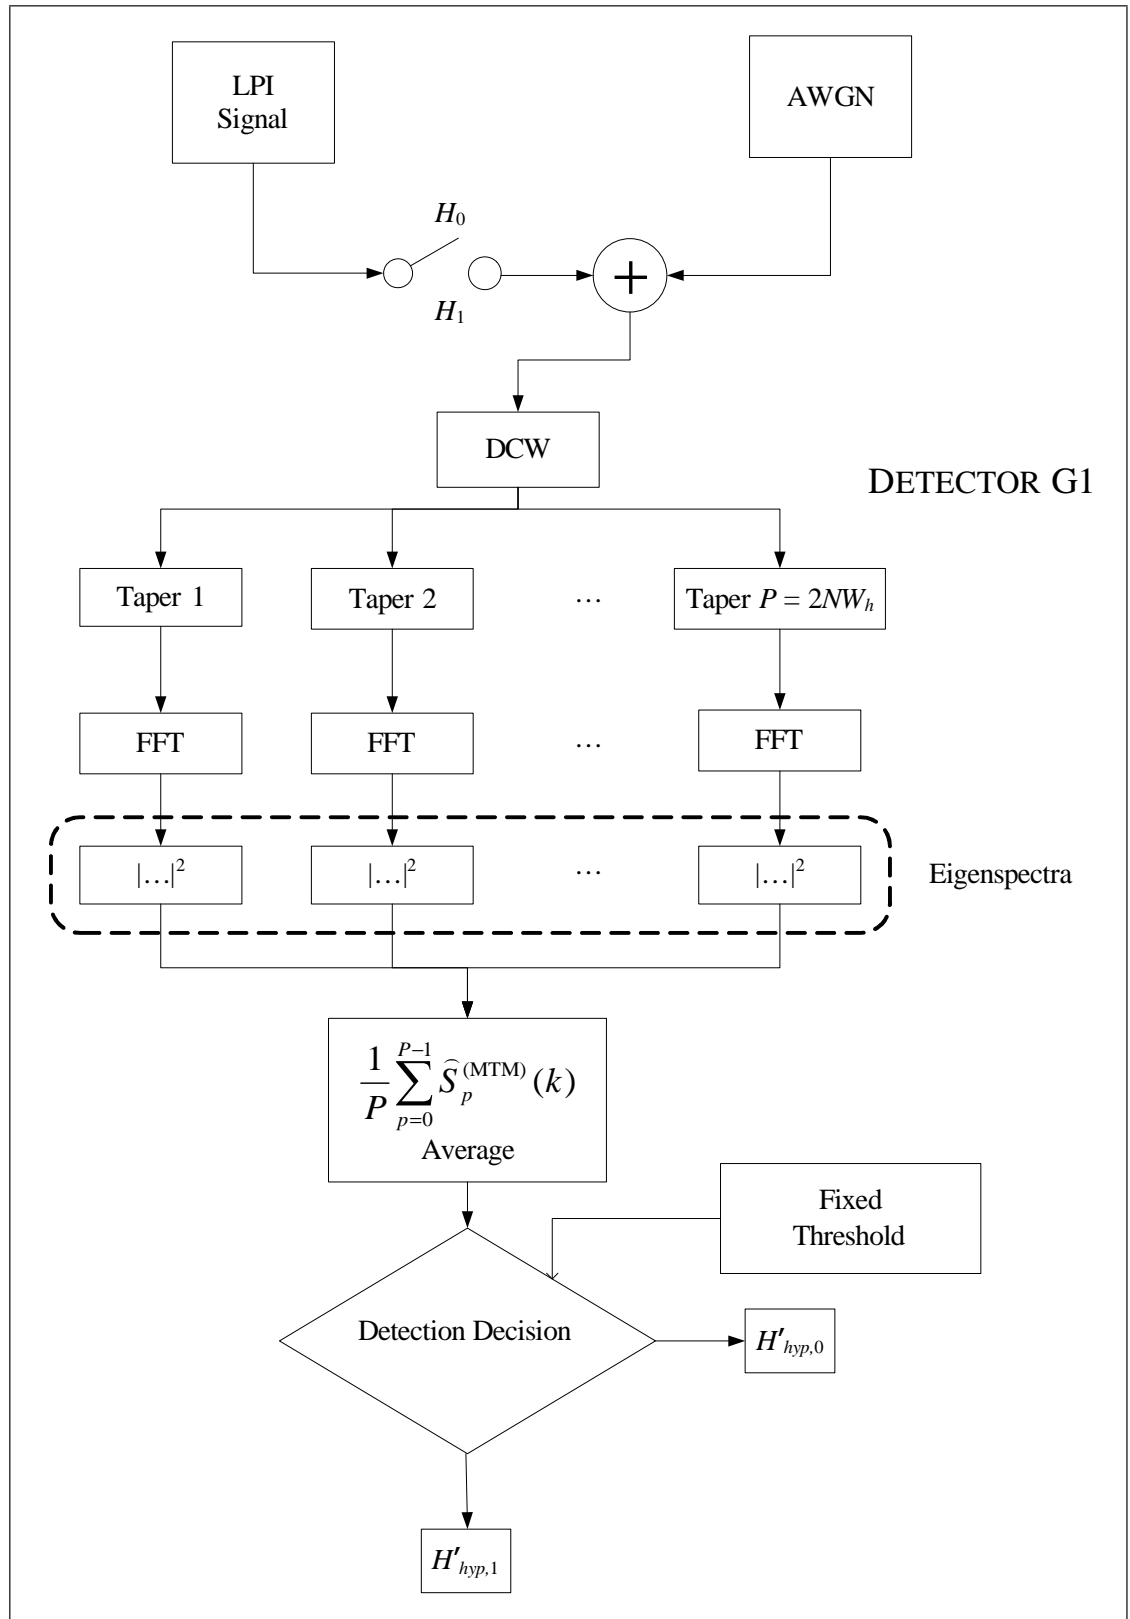


Figure 5.1.2 Detector G1 for single signal scenarios

5.1.2 Sine Taper MTM Detectors

Only two of the three different sine taper MTMs are portrayed in this work: (1) the unity-weighted sine taper MTM and (2) the adaptive sine taper MTM. The sine taper MTMs algorithms are described in section 4.2.6.

The unity-weighted sine taper MTM was used rather than the parabolically-weighted sine taper MTM because the parabolic weight scheme did not appear to be applicable to the type of spectrum which an ESM DRx might encounter as it caused the detector performance to be degraded over the use of the unity weight scheme. This is partially evident from inspection of the spectral estimates in panel (c) of Fig 4.2.8.

The same small-valued hypothesis time-bandwidth products were tested in the unity-weighted sine taper MTM detector as in the DPSS MTM detectors. This was $NW_h = 0.5, 1.0, 1.5, \dots, 3.5$ and these hypotheses were indexed by $hyp = 1, 2, 3, \dots, 7$. The same method of detection through the application of an OR operator to all of the logical outcomes of the individual hypotheses was used. This is illustrated in Fig 5.1.1. The probability of false alarm was defined for the output of the overall system.

The sine tapers have an analytical expression, which is given in Eq (4.2.7). Therefore there was no requirement to calculate different sine tapers for different time-bandwidth products. Furthermore the sparse frequency-domain representation of the sine tapers was utilised by convolution in the frequency-domain. This is clear from Fig 5.1.3 where **detector H1**, utilises the unity-weighted sine taper MTM. In the unity-weighted sine taper MTM, the closed-form expression for the sine tapers also meant that the eigenspectra from previous smaller valued NW_h detectors were re-used in larger valued NW_h detectors. The pattern was thus: the eigenspectra from $hyp = 1$ were reused in $hyp = 2$, the eigenspectra from $hyp = 2$ were reused in $hyp = 3$ and so on up to $hyp = 7$. The analytical sine taper expression, the use of a single FFT and the re-use of the eigenspectra all enabled computational savings to be made over the DPSS MTM (although this was less of a concern during floating-point simulation).

The sine taper definition in Eq (4.2.7) wastes one sample as it contains a zero-weight coefficient, but this was found to have a negligible impact on all results for the length of FFT tested in the simulations in this chapter. This was found through comparison with results where the DCW length included an extra sample. Therefore the sine taper definition used was Eq (4.2.7), albeit in the frequency-domain.

The unity-weighted sine taper MTM uses a fixed number of tapers at each frequency bin at each hypothesis (hyp). However the adaptive sine taper MTM detector can use as many or as few sine tapers as necessary at each frequency bin. Thus there was no hypothesis bandwidth involved in the adaptive sine taper MTM detector. The use of a variable number of tapers at each frequency bin has advantages as the adaptive sine taper MTM detector is not restricted to a maximum of P tapers (instead it is limited

to $N/2 - 1$). The disadvantage was that it is far more computationally intensive than the unity-weighted sine taper MTM. A system flow diagram of the adaptive sine taper MTM is shown in Fig 5.1.4. It is named **detector I1**.

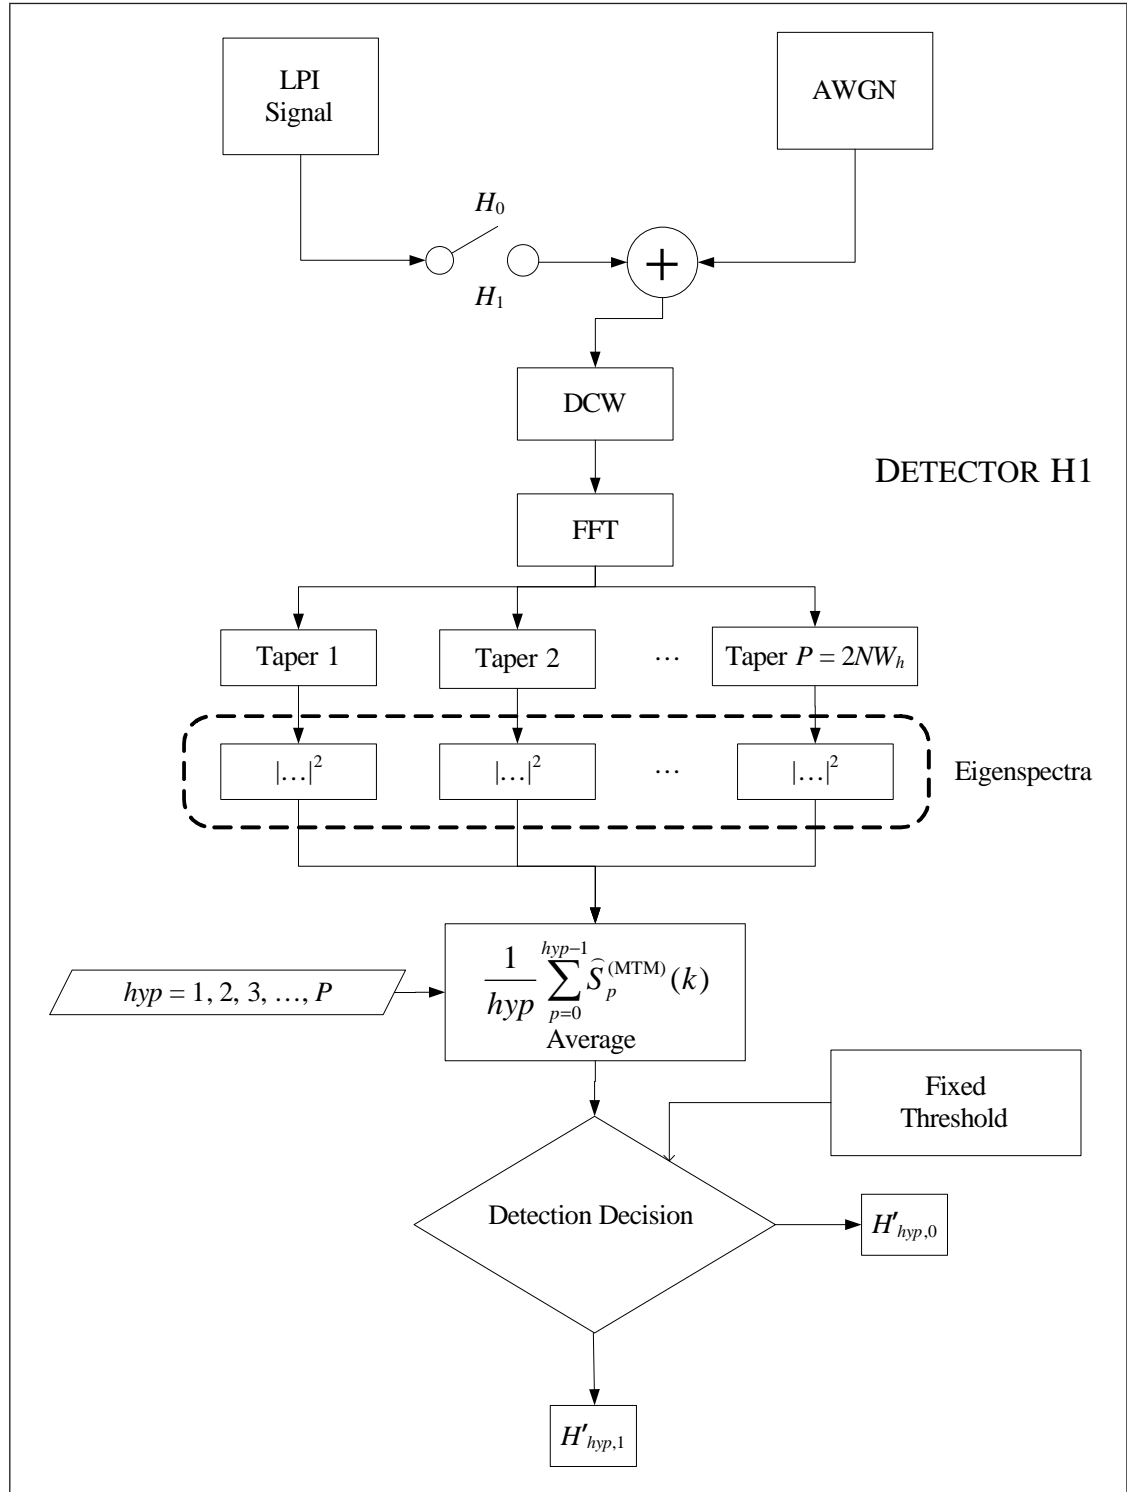


Figure 5.1.3 Detector H1 for single signal scenarios

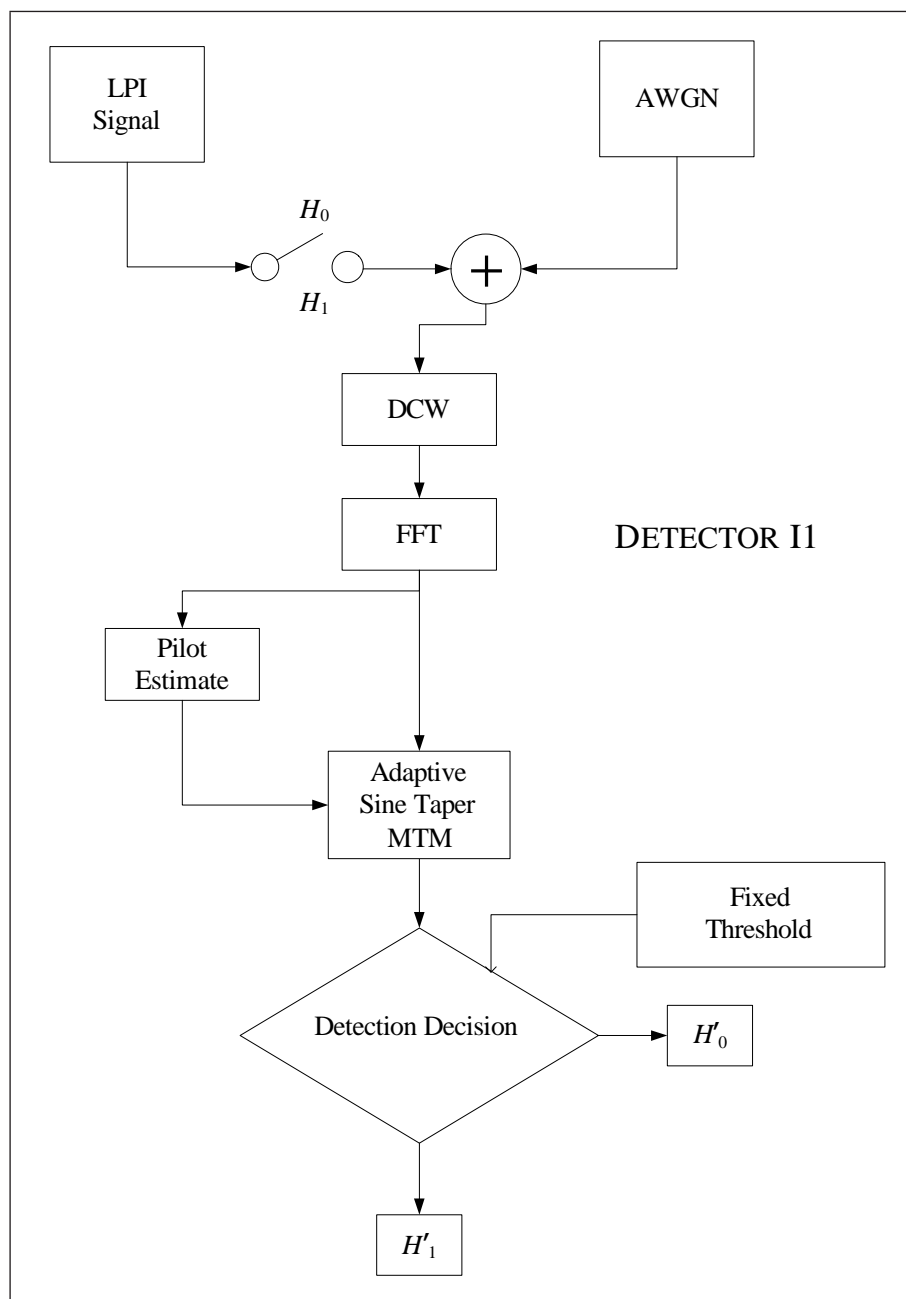


Figure 5.1.4 Detector I1 for single signal scenarios

5.1.3 Smoothed Periodogram Detectors

Although the title of this section is “smoothed periodogram detectors”, this section covers the SP detector, the SMP detector and the SSP detector because they are closely related. The basic algorithm involved in these detectors was the calculation of the periodogram/modified periodogram/SVA periodogram and application of odd-valued uniformly-weighted smoothing filters of increasing span (L) up to a length limit. Uniformly-weighted filters were used for simplicity reasons so that the operation was simply the summation of adjacent ordinates and division by the summation length.

The periodogram (rectangular window function) includes as much signal power as possible because it does not suffer from a window loss. The modified periodogram (bell-shaped window function) suffers window loss but attenuates sidelobe spectral leakage. Therefore it was expected that the SP detector would perform better than the SMP detector in scenarios where only one signal was considered, however the situation would be reversed in scenarios where sidelobe spectral leakage could not be ignored. This was the justification as to why both the SP detector and SMP detector were tested as they were expected to perform differently depending on the scenario. The SVA periodogram approaches the processing gain of the periodogram and has good spectral leakage minimisation like the modified periodogram. This was the rationale for the inclusion of a SSP detector in the comparison.

The system diagram for the SP and SMP detectors is shown in Fig 5.1.5 and labelled **detector J1**. In this diagram the window function block can be taken to mean either the rectangular window function in the case of the SP detector or the Hann window function in the case of the SMP detector. Fig 5.1.6 illustrates the system diagram for the SSP detector, labelled as **detector K1**. The only change between Fig 5.1.5 and Fig 5.1.6 is the location of the window from the time-domain to the frequency-domain.

To enable a fair comparison against the MTMs, the filter span (L) was limited to a maximum value which was the equivalent of $NW_h = 3.5$. Therefore for the non-zeropadded case, this involved spans of $L = 1, 3, 5, 7$ and for the zeropadded case this involved spans of $L = 1, 3, 5, 7, 9, 11, 13$. This corresponded to four hypotheses to test for the presence of a signal in the non-zeropadded case and seven hypotheses to test in the zeropadded case. The individual hypotheses test outcomes were combined to produce a signal-present or no-signal-present declaration in an analogous way to that shown in Fig 5.1.1, except that $hyp = 4$ was the maximum for the non-zeropadded case. As with the MTMs, the probability of false alarm was defined at the output of the overall system.

The multiple hypotheses involved the calculation of moving average filters of increasing, although odd-numbered, spans and therefore there was no need to recalculate the entire summation for each filter because an accumulated sum was stored and reused. For example, from an $L = 3$ filter to an $L = 5$ filter, the sets of three ordinates which

were summed in the $L = 3$ filter were stored as an accumulated total to which only sets of two further ordinates were added for the $L = 5$ filter. This reuse of previous results is symbolised in Figs 5.1.5 and 5.1.6. Although not a major concern in floating-point simulations, the reuse of previously calculated results would enable computational savings in a real-life DRx.

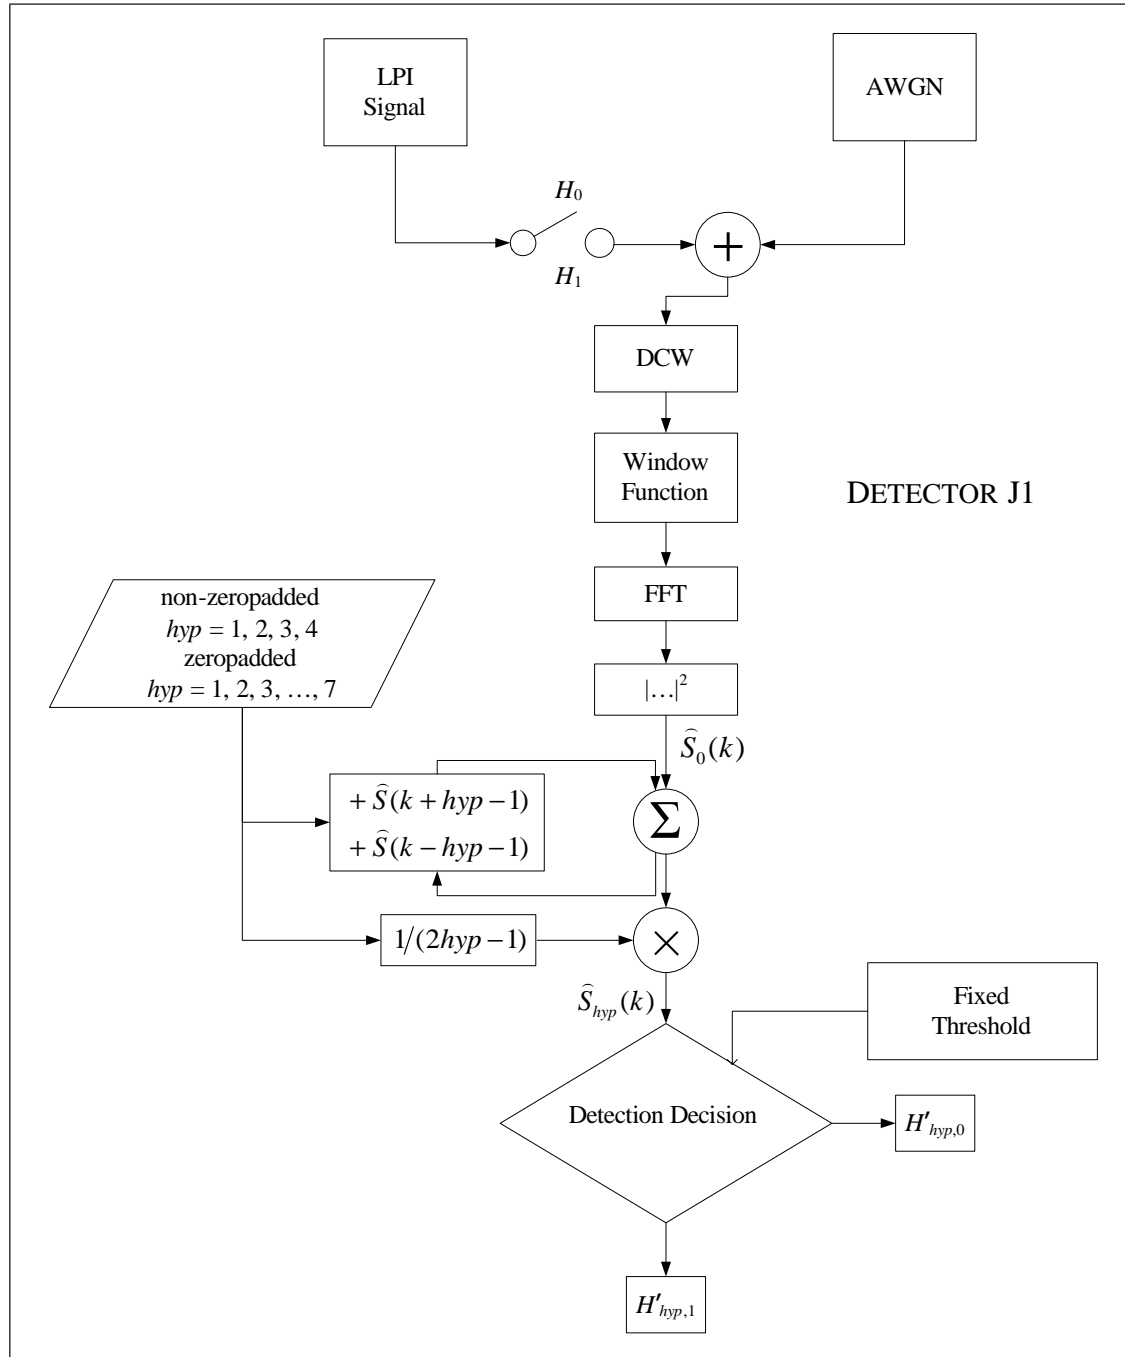


Figure 5.1.5 Detector J1 for single signal scenarios

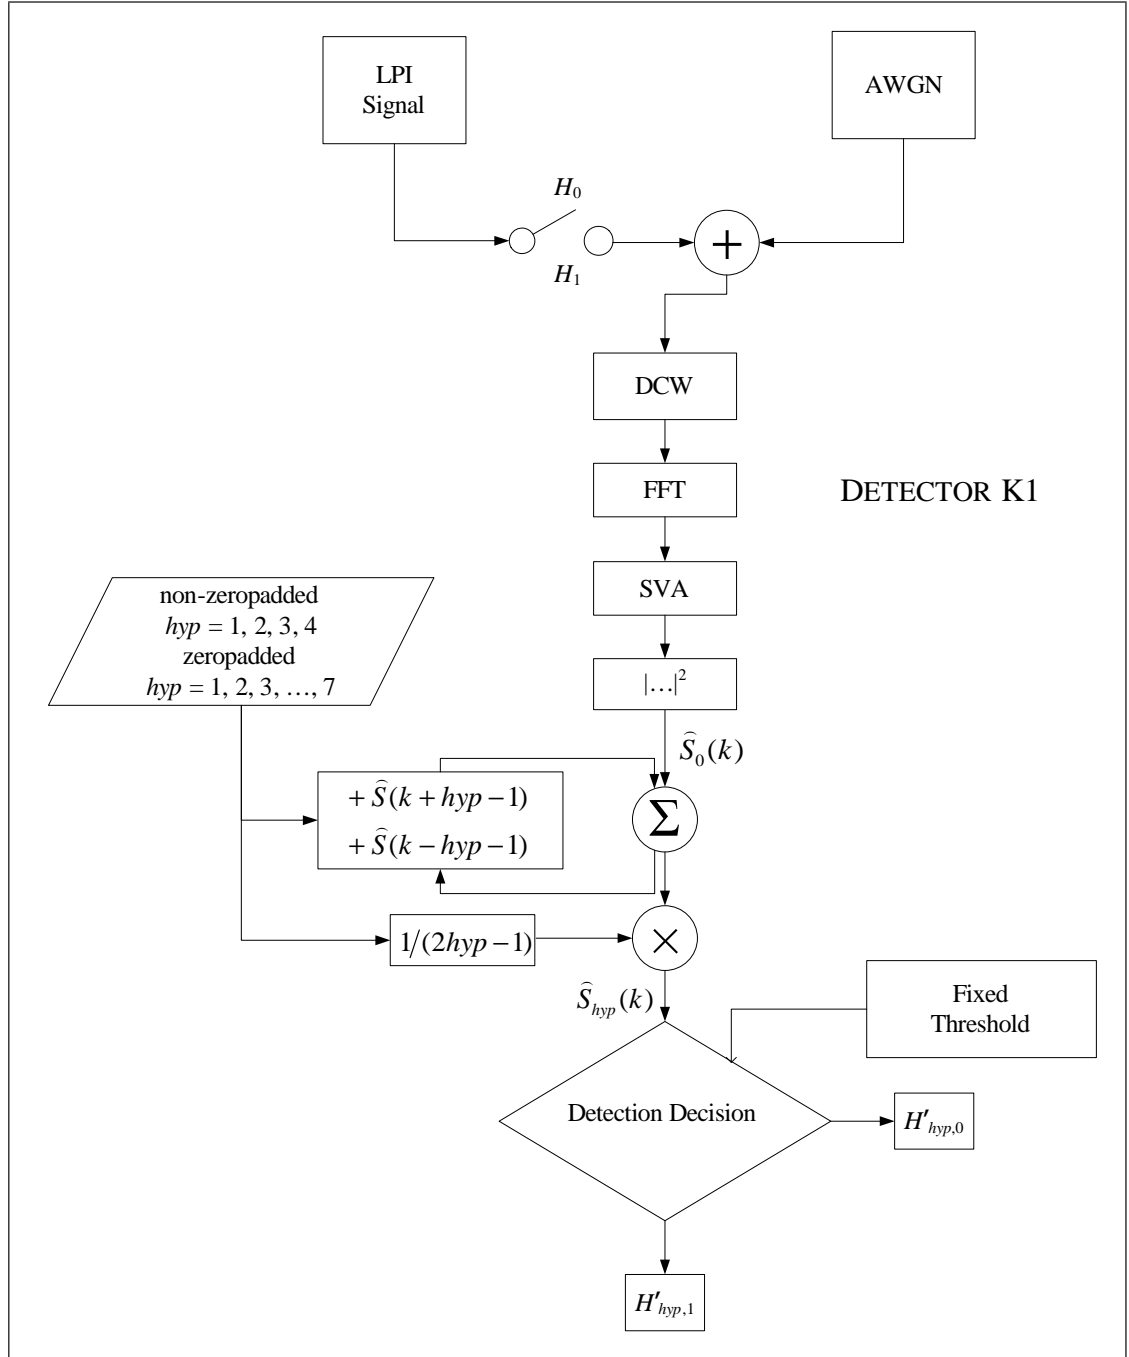


Figure 5.1.6 Detector K1 for single signal scenarios

5.1.4 Tree Smoothed Periodogram Detectors

The descriptions given in this section encompass not only the TSP detector but also the TSMP and TSSP detectors. These detectors are grouped together because the differences between them are trivial. The fundamental processes which took place in these detectors were to calculate either: (1) the periodogram or (2) the modified periodogram or (3) the SVA periodogram. The neighbour ordinates of these spectral estimates were averaged in increasing powers of two up to a limit. The tree smoothed detectors are thus a special case of the moving average filter detectors as the spans are restricted to powers of two. For clarity no special weight scheme was applied to the ordinates prior to the average to maintain simplicity.

The motivation for the calculation of averages in increasing powers of two was that the results from a previous “layer” could be reused in the subsequent “layer”. For example, if sets of two ordinates are averaged in the first layer, then the results from this could be used in the formulation of the sets of averages of four ordinates. This would provide great computational savings for a real-life DRx.

The restriction of the filter span to powers of two caused the number of hypothesis bandwidths which could be tested to be curtailed. For example, a hypothesis bandwidth which is the equivalent of three ordinates could not be tested. This loss in bandwidth test accuracy becomes coarser as the bandwidth increases. However for the DRx parameters given in Table 5.2.2 most LPI signals would span only a handful of frequency bins where the loss in bandwidth test accuracy is not so severe. Therefore the loss in detector performance was conjectured to be tolerable when traded against the computational savings.

For the same reasons given in section 5.1.3, the use of a rectangular window function, Hann window function and SVA was explored in the periodogram tree-like smoothing detectors. The names given to these detectors were: (1) the TSP detector; (2) the TSMP detector and (3) the TSSP detector respectively.

It was expected that the TSP detector would perform well in scenarios where spectral leakage is not a problem whilst the TSMP detector would perform well in scenarios where spectral leakage could not be ignored and the TSSP detector would perform well in both of these scenarios. The TSP and TSMP detectors are shown at the system level in Fig 5.1.7. The detector in Fig 5.1.7 is labelled **detector L1**. The choice of the window function used in Fig 5.1.7 was either a rectangular window function or Hann window function, depending on whether detector L1 operated as a TSP detector or TSMP detector. Likewise, the TSSP detector is described at a system level in Fig 5.1.8. The detector in Fig 5.1.8 is labelled as **detector M1**. The only change from detector L1 to detector M1 is the location of the window. In detector L1 it is in the time-domain and in detector M1 it is in the frequency-domain. In Figs 5.1.7 and 5.1.8 the simple periodogram-based spectral estimation is perceptible as the $\hat{S}(k)$

estimate formed after the magnitude-square operation. The smoothing in a tree-like architecture is represented in the diagram with the subsequent operations. The loop and the sigma symbol show the tree smoothed spectral estimate results in one layer are accumulated and reused to produce the next tree smoothed spectral estimate for the next layer. The \gg symbol represents a barrel shift ($hyp - 1$) places to the right (i.e. division by 2^{hyp-1}). The layers are labelled with the index hyp .

To compare the tree smoothed detectors with the other detectors, four different bandwidth hypotheses were tested in the non-zeropadded case $hyp = 1, 2, 3, 4$ and five hypotheses in the zeropadded case $hy = 1, 2, 3, 4, 5$. This was the equivalent of moving average filters of spans $L = 1, 2, 4, 8$ in the non-zeropadded case and $L = 1, 2, 4, 8, 16$ in the zeropadded case. In terms of the hypotheses time-bandwidths being tested, this corresponded to $NW_h = 0.5, 1.0, 2.0, 4.0$ in the non-zeropadded case and $NW_h = 0.5, 1.0, 2.0, 4.0, 8.0$ in the zeropadded case. The individual hypothesis bandwidths were tested for the presence of signals and the multiple hypotheses were combined into a single signal-present or no signal-present declaration in exactly the same manner as in Fig 5.1.1.

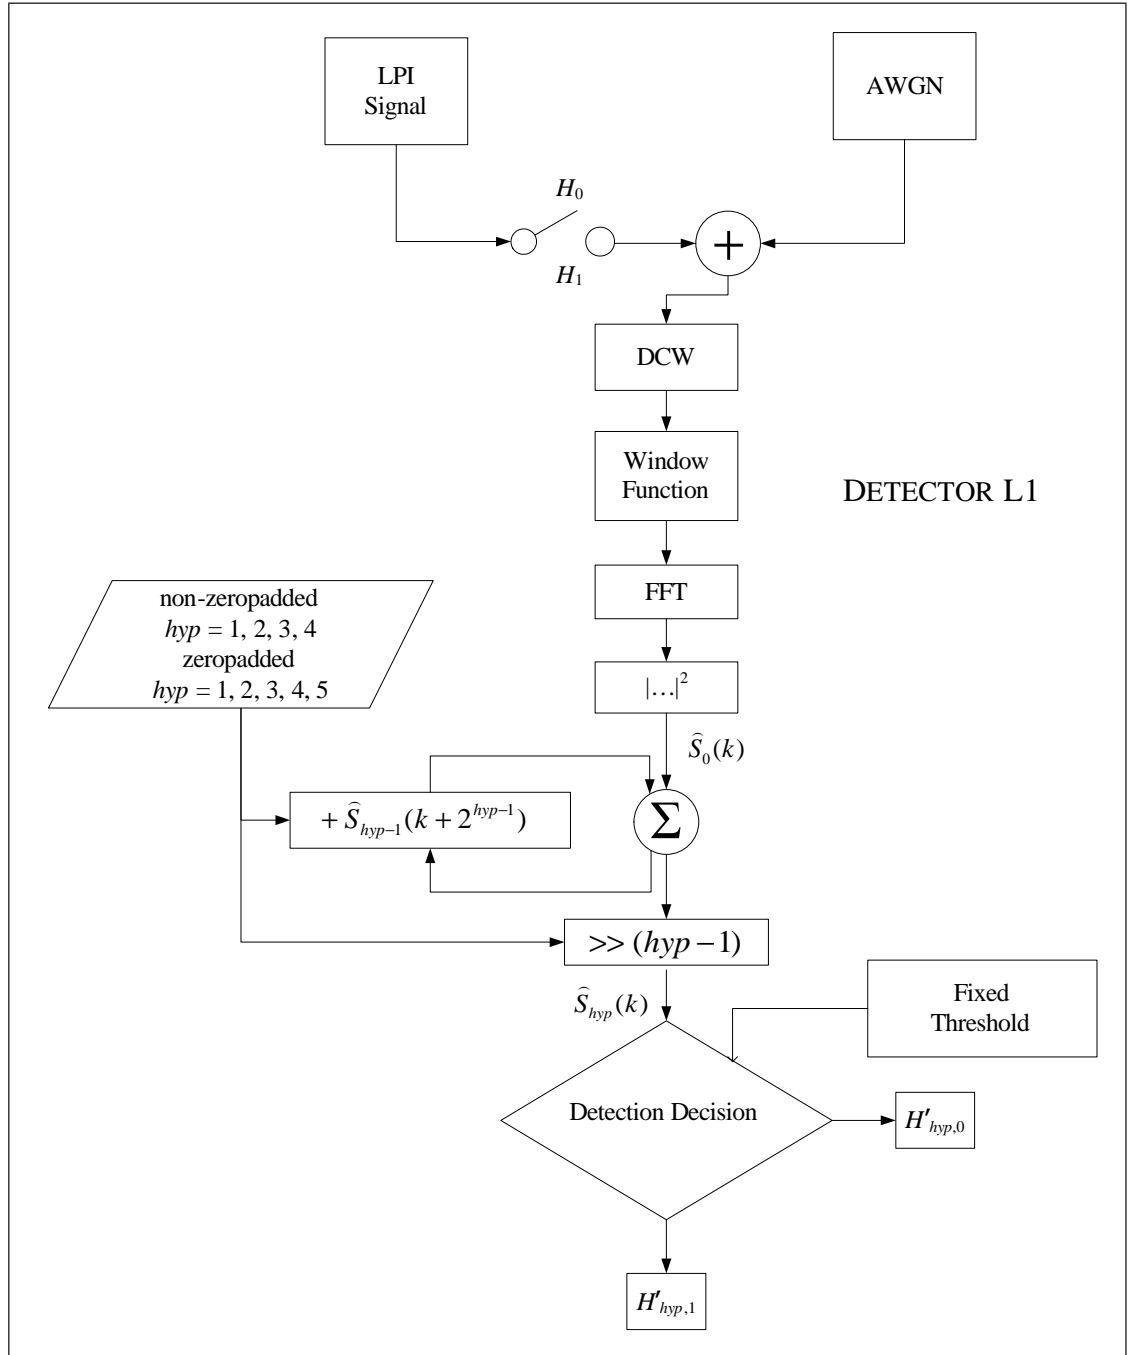


Figure 5.1.7 Detector L1 for single signal scenarios

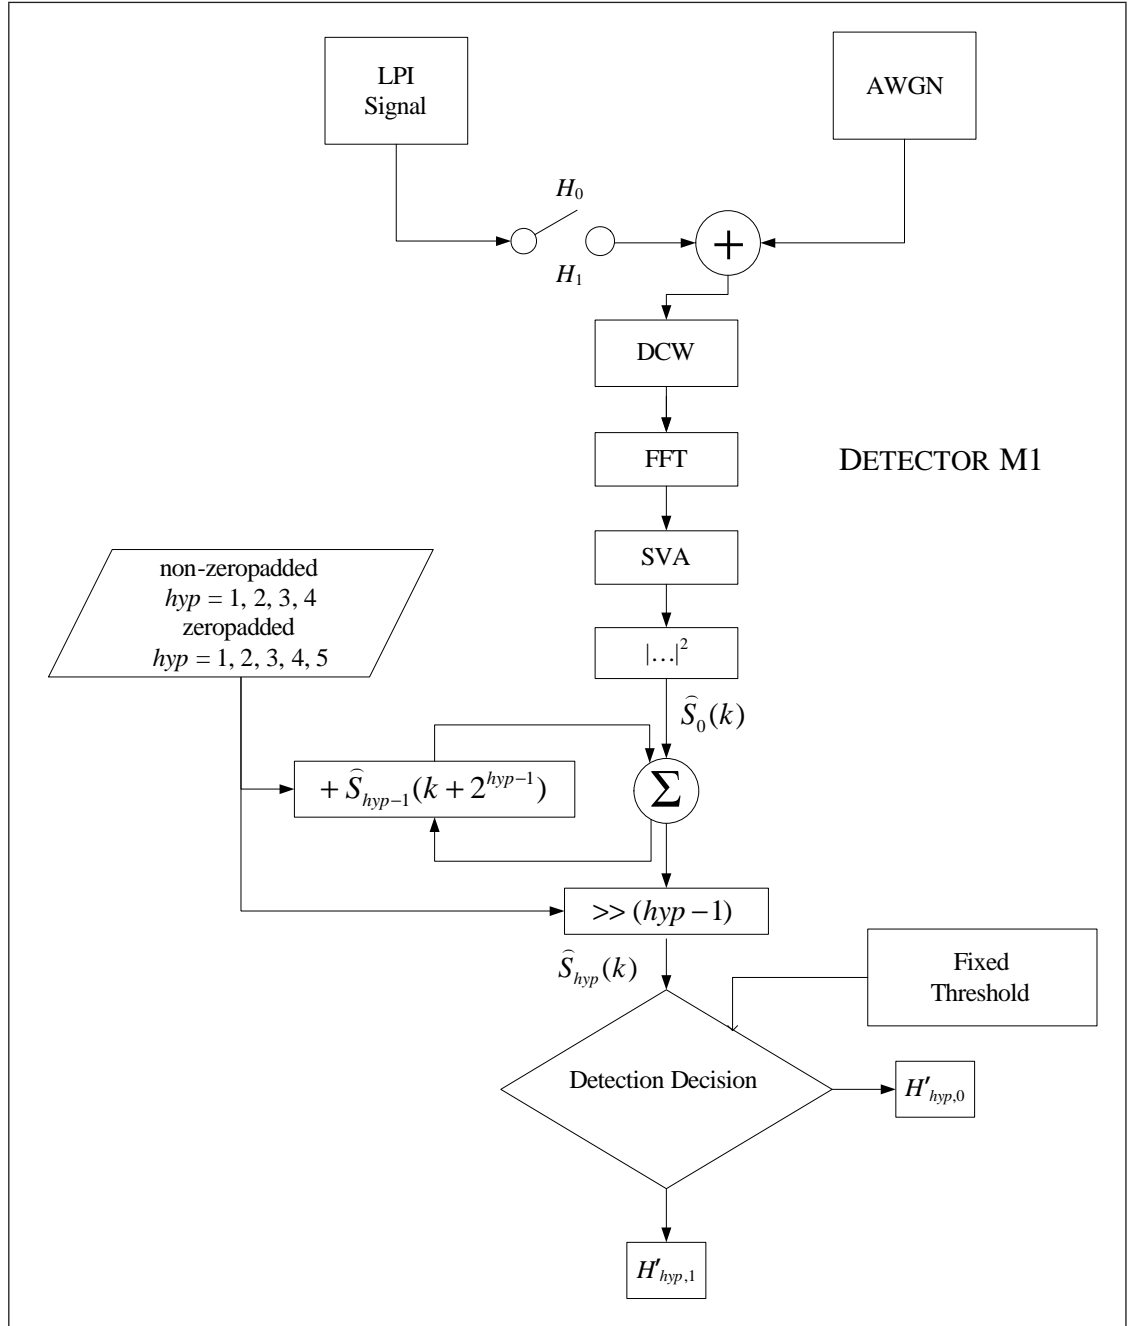


Figure 5.1.8 Detector M1 for single signal scenarios

5.1.5 OR Periodogram Detectors

This is the first of the descriptions about the conventional detection techniques which were used for comparative purposes. The title of this section is “OR periodogram detectors” but the explanations cover the OR periodogram detector, OR modified periodogram detector and OR SVA periodogram detector. As a conventional detection technique, the processes involved in these detectors were fairly basic. Essentially the periodogram, modified periodogram or SVA periodogram were calculated and the ordinates within the bandwidth of interest were checked against a threshold to determine if a signal was present. This detection mechanism was already used in the detection of the Barker 13 signal in the SVA chapter in section 3.6.

These detectors represent the best detection procedure which could be achieved in a conventional FFT-based DRx as part of an ESM system. In fact these detectors may be considered to be generous to conventional DRx processes because in a real-life situation the bandwidth of a target signal is initially unknown and if there were detections encountered in a few adjacent ordinates it would be difficult for a conventional DRx to decide if these were due to many closely-spaced in frequency, individual sinusoidal signals or a broader band signal. The purpose of the OR detectors was to provide a best possible case benchmark against which the MTM detectors and periodogram (tree) smoothing detectors could be compared.

In this scenario the number of ordinates that were checked for the presence of LPI signals was seven in the non-zeropadded DCW case and thirteen in the zeropadded DCW case. This was done so that these detectors could be fairly compared against the MTM detectors and the periodogram (tree) smoothing detectors which searched up to a maximum hypothesis time-bandwidth of $NW_h = 3.5$. The ordinates checked were centred around the target frequency of the LPI signal.

The same arguments apply from the previous two sections as to why the periodogram, modified periodogram and SVA periodogram were chosen as the fundamental operations in the detectors. A representation of the OR periodogram and OR modified periodogram detectors is given in Fig 5.1.9 and labelled as **detector N1**. In Fig 5.1.9 the window function block is a rectangular window function for the OR periodogram detector and a Hann window function for the OR modified periodogram detector. The block which shows the indices selected from the spectral estimate, $[\hat{S}_{xx}(k - P), \hat{S}_{xx}(k + P)]$, indicates that seven ordinates were checked in the non-zeropadded case and thirteen ordinates were checked in the zeropadded case. All of the ordinates were checked against the same fixed threshold to produce seven or thirteen Boolean values depending on whether the DCW was non-zeropadded or zeropadded. The Boolean values were then short-circuited by an OR operation into one signal-present or no signal-present decision.

The system level diagram of the OR SVA periodogram detector is given in Fig 5.1.10

and is labelled as **detector O1**. The difference which detector O1 exhibits in comparison to detector N1 is that the window is in the frequency-domain.

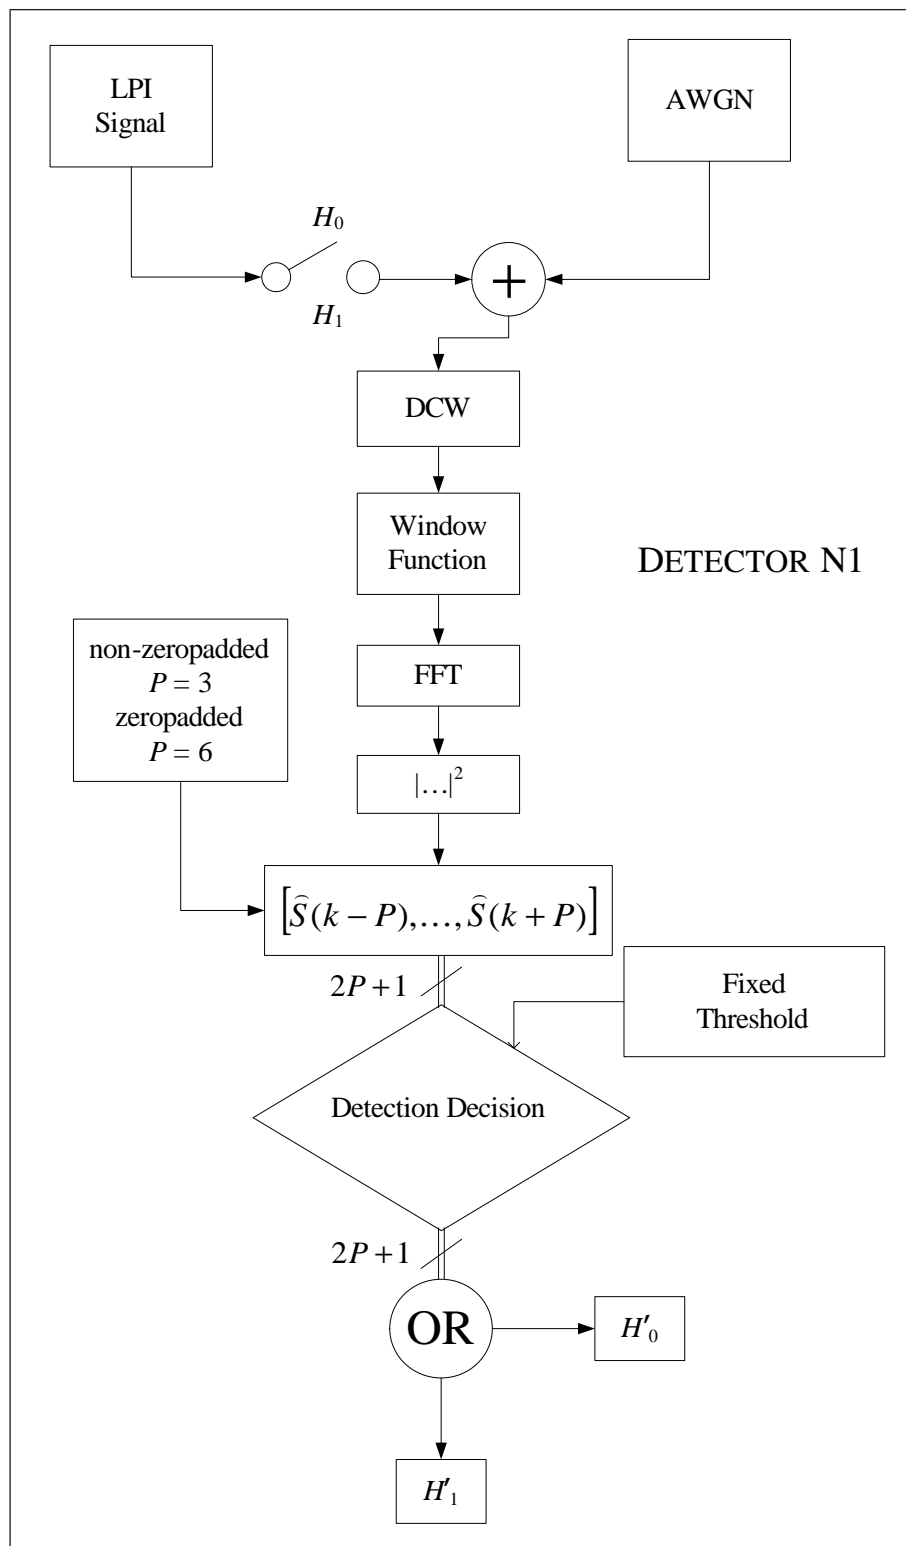


Figure 5.1.9 Detector N1 for single signal scenarios

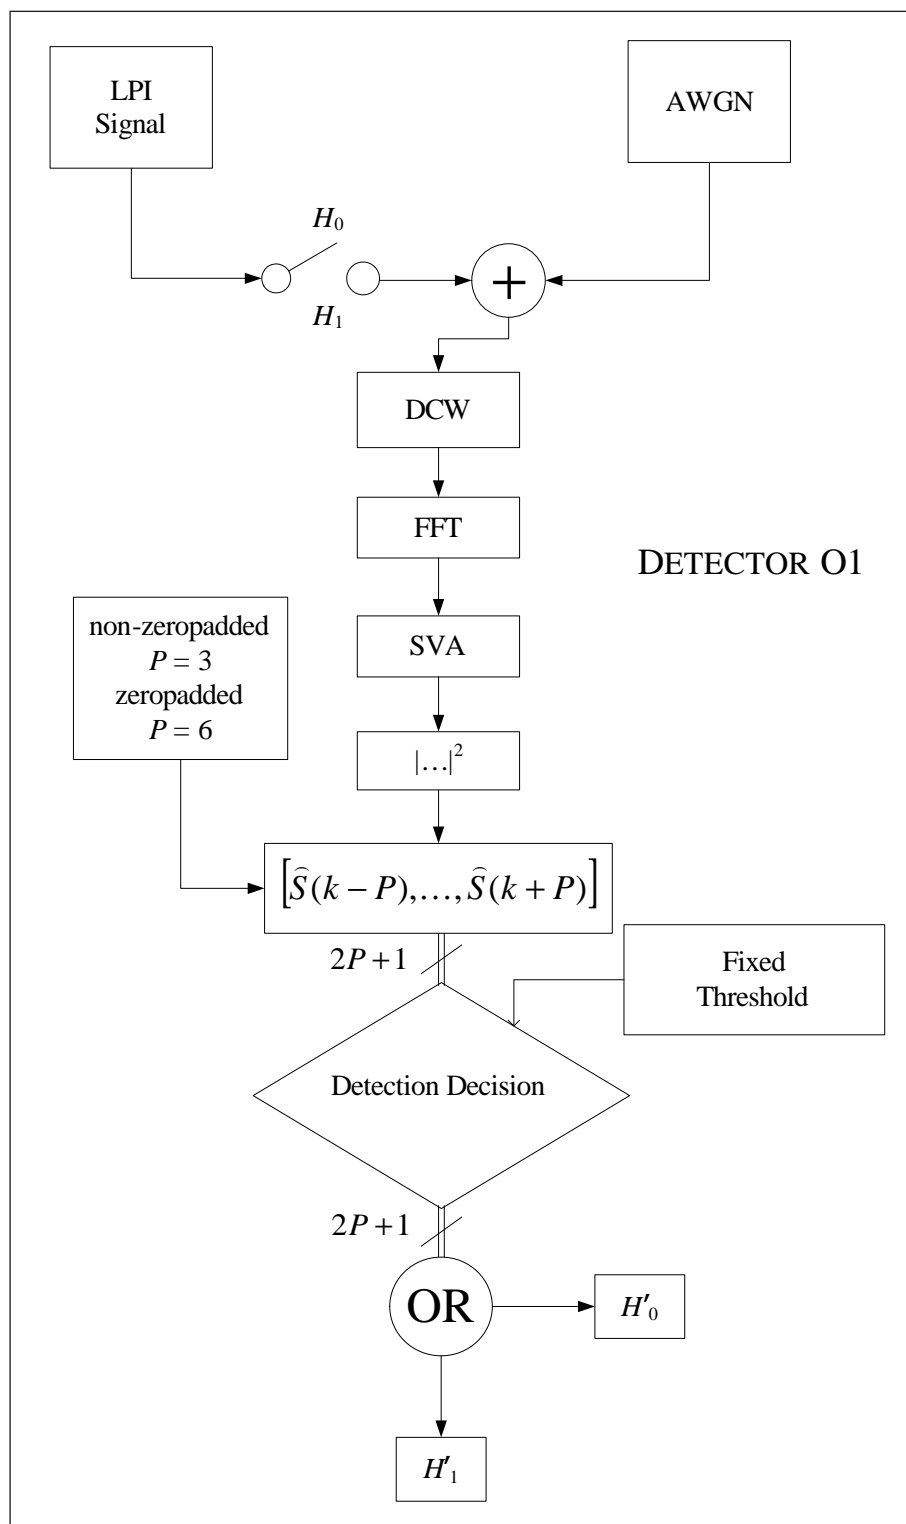


Figure 5.1.10 Detector O1 for single signal scenarios

5.1.6 Periodogram Detectors

Detectors A1, B1, C1 and E1 from the SVA chapter were also reused in the simulations in this chapter as additional benchmarks.

Both the periodogram-based detector and modified periodogram-based detector were tested as variations of detectors A1 and C1. To produce a periodogram detector, a rectangular window function was chosen for the window function block in Fig 3.2.3 and to produce a modified periodogram, a Hann window function was chosen for the window function block. The SVA periodogram was tested with the use of detector B1 and E1 models. These detectors are grouped under the category “periodogram detectors”. As with the periodogram smoothing detectors and the OR periodogram detectors, the Hann window function, rectangular window function and SVA were chosen as window techniques for simulation to link to the SVA chapter and its conclusions.

The periodogram detectors differed from the OR periodogram detectors in that only the ordinate of the LPI signal target frequency was checked for the presence of a signal. This type of architecture more closely resembles real-life DRx decision logic which traditionally searches for the presence of narrow band sinusoidal signals and was therefore a more realistic benchmark against which the MTM and periodogram smoothing methods were tested. As a consequence there were no bandwidth hypotheses involved in these detectors.

5.1.7 Total Power Detector

The final conventional detector model which was used to provide another benchmark was the total power detector. This was selected as another benchmark because the total power methodology is used in older conventional ESM systems (see chapter 1.3). In older systems the envelope or square of the envelope of an analogue DCW was thresholded, typically with an integrator circuit. However the digital equivalent of the square of the envelope was modelled in these total power detectors through the summation of the magnitude-squared of the complex samples from a DCW (i.e. square-law detection). The total power was calculated in the time-domain. This is indicated in Fig 5.1.11.

The summation of the magnitude-squared of the complex samples in a DCW produced a single decision statistic therefore there were no bandwidth hypotheses to be tested in the total power detectors.

The decision statistic obtained was tested against a fixed threshold to make the decision whether a signal was present or not. This fixed threshold was set at a level such that the probability of false alarm was $P_{fa} = 0.01$. To find such a P_{fa} in the full DCW scenario, the formula for the PDF of the sum of N independent, normally distributed

RVs was used. This is given by

$$f(x; N) = \frac{1}{2^N \sigma^{2N} \Gamma(N)} x^{N-1} \exp\left(-\frac{x}{2\sigma^2}\right) \text{Pos}_{\{x \geq 0\}} \quad (5.1.1)$$

The derivation of Eq (5.1.1) is described in [3, p. 294–295]. The PDF expression in Eq (5.1.1) is similar to that of Eq (4.2.2). The main difference between the two expressions is that the PDF in Eq (4.2.2) is for the summation of P independent, **standardised**, normal RVs whereas the PDF in Eq (5.1.1) is for the summation of N independent, normal RVs.

There are similarities between the total power detector and the SP detector or TSP detector because they are both summations of RVs. The main difference between the total power detector and the SP detector or TSP detector is that the SP or TSP truncate the summation at a much lower number of ordinates and test the decision statistic after each term is added to the accumulated sum.

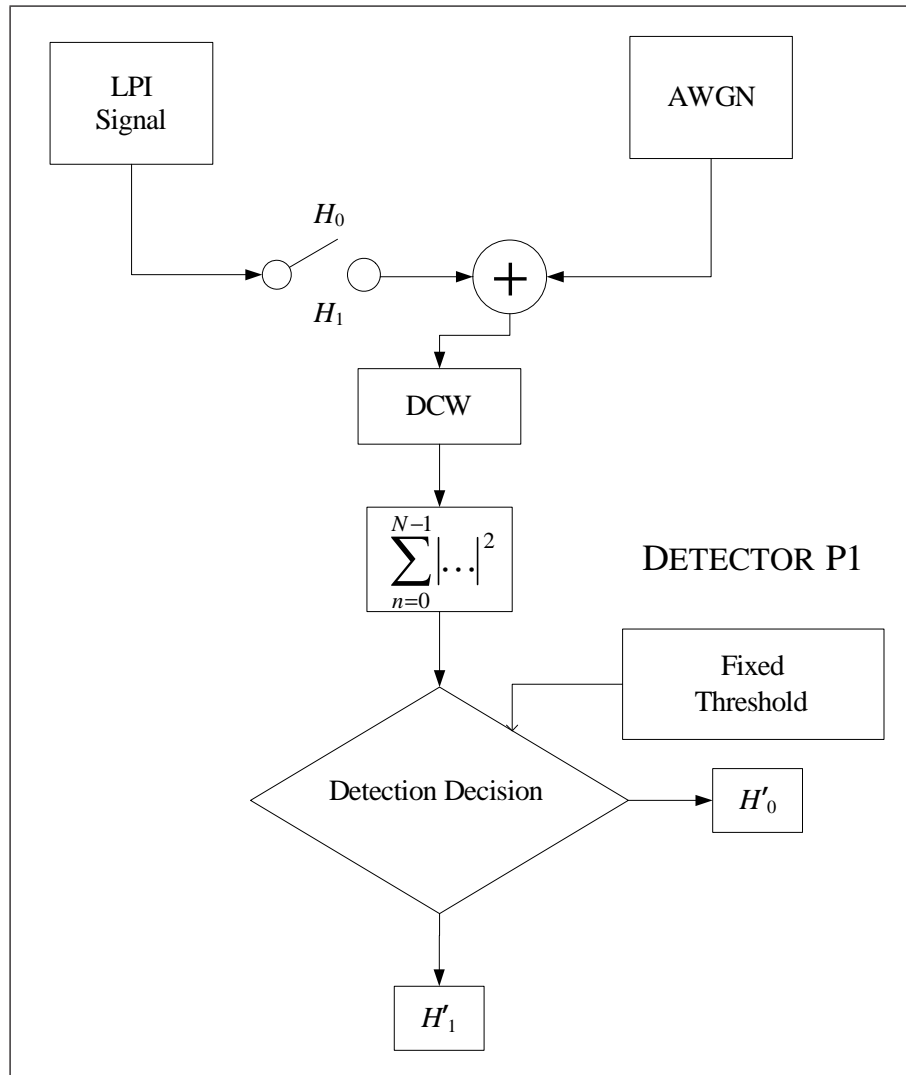


Figure 5.1.11 Detector P1 for single signal scenarios

5.1.8 Detectors Summary

A summary of the algorithms and corresponding detectors used in this chapter is shown in Table 5.1.1.

Detector	Algorithm	Notation
A1	Periodogram (single sig)	P
	Modified periodogram (single sig)	MP
B1	SVA periodogram (single sig)	SVA P
C1	Periodogram	P
	Modified periodogram	MP
E1	SVA periodogram	SVA P
G1	Unity-weighted DPSS MTM	DPSS MTM
H1	Unity-weighted sine taper MTM	Sine MTM
I1	Adaptive sine taper MTM	Sine MTM adapt
J1	Smoothed periodogram (single sig)	SP
	Smoothed modified periodogram (single sig)	SMP
K1	Smoothed SVA periodogram	SSP
L1	tree smoothed periodogram (single sig)	TSP
	tree smoothed modified periodogram (single sig)	TSMP
M1	tree smoothed SVA periodogram	TSSP
N1	OR periodogram (single sig)	OR P
	OR modified periodogram (single sig)	OR MP
O1	OR SVA periodogram	OR SVA P
P1	total power	TP
Q1	smoothed periodogram	SP
	smoothed modified periodogram	SMP
R1	tree smoothed periodogram	TSP
	tree smoothed modified periodogram	TSMP
S1	OR periodogram	OR P
	OR modified periodogram	OR MP

Table 5.1.1 Conventional, MTM & periodogram smoothing detectors

5.2 Full DCW

5.2.1 Scenario and Simulation Description

The detectors in this scenario searched for the presence of signals against a background of AWGN. As indicated above, the signals tested were LPI and used the coding schemes: Barker-13 and m-sequence. The coding was applied to a sinusoidal carrier using phase shifts. Where phase changes between 0 and π occurred they were implemented over a very short although finite time interval ($\sim 5\%$ of the chip interval).

The sinusoidal carrier was actually a complex exponential with an unknown target frequency (f_T) and unknown initial phase (ϕ_0). In a similar fashion to the simulations from chapter 3, the sinusoidal carrier frequency was chosen at random from a frequency bin interval (f_s/N_{FFT}) at the start of each test run. This concept is illustrated in Fig 5.2.1 for the Barker-13 pulse. Panels (b) and (c) illustrate that the whole approximate bandwidth shifted up or down in frequency along with the centre target frequency. A similar picture would occur for the m-sequence CW signal but it would show a slightly wider approximate bandwidth.

Barker-13 Pulse:	chip-rate	4 MHz
	pulse width	3.25 μs
	PRF	50 kHz
	IF cycles/chip	~ 79
M-Sequence CW:	chip-rate	4.25 MHz
	compression ratio	255 : 1
	IF cycles/chip	~ 75

Table 5.2.1 LPI Waveform parameters

In the case of the Barker-13 MOP, the signals were created as a series of pulses and each pulse was individually coded with the phase code scheme. In the case of the m-sequence code, the signal was created as a CW waveform, where periods of the waveform were coded with the phase code scheme. The parameters of the LPI waveforms used are given in Table 5.2.1.

In this scenario, the DCWs which were tested for the presence of a signal at each test iteration were full of signal instantaneous power, or in other words, each sample contained a power contribution from the signal. This is always true for CW signals. For the Barker-13 pulses, the full DCW scenario was important because typically an ESM DRx DCW duration is much less than the duration of a signal pulse. Therefore when typical parameters were applied in the Barker-13 simulations it was possible to always find a DCW which was full of signal power. The typical DRx parameters which were used in this scenario are listed in Table 5.2.2. Despite a relatively long DCW and

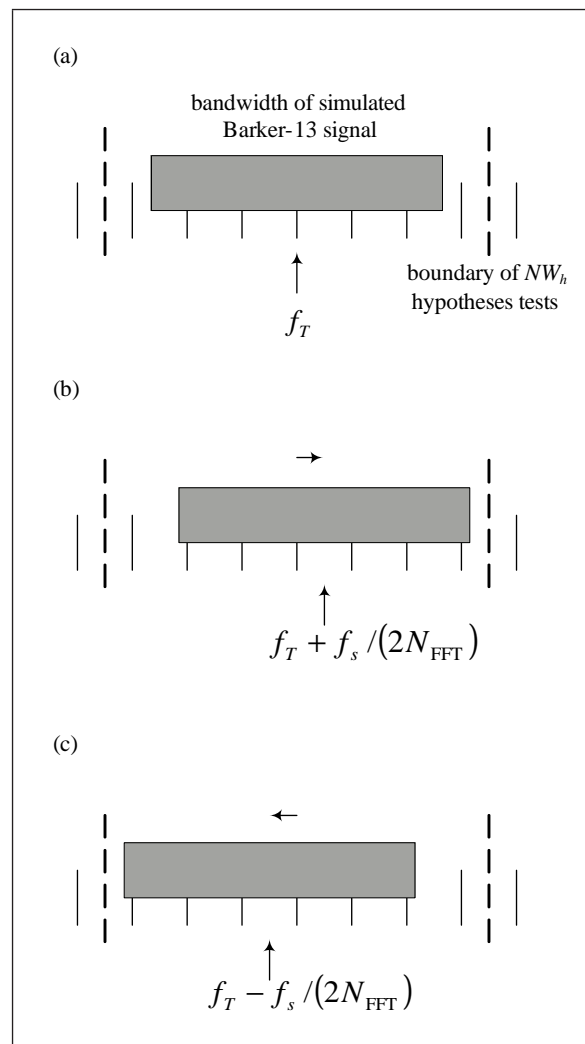


Figure 5.2.1 Monitored bandwidth where the Barker-13 target frequency: (a) coincides with central monitored bin; (b) is offset by $+f_s/N_{\text{FFT}}$; (c) is offset by $-f_s/N_{\text{FFT}}$

application in a non-overlapped fashion, the DCW duration was still much less than the Barker-13 pulse duration.

In the test procedure it was necessary to simulate that it would be unlikely in real-life for the DCW start times and Barker-13 pulses or m-sequence periods to be synchronised. Therefore a random portion of the pulse or CW was captured in the DCW. For the Barker-13 pulses this entailed the obvious caveat that the DCW was full of signal power. This was arranged such that the pulse did not start after the start time of the DCW nor did it finish before the end time of the DCW. A diagram of how this was achieved is shown in Fig 3.6.2 in section 3.6.

A similar definition of random target frequency to that in Fig 3.2.2 was used. However the difference between here and the simulations from chapter 3 was that the LPI signals were broader-band. Fig 5.2.1 shows how the random choice of f_T at the start of each test could shift the whole bandwidth up or down in frequency. Panel (a) shows the ideal case where f_T happened to be coincident on the frequency bin of interest. Panel

(b) and (c) show the two extreme cases where f_T is offset by the equivalent of plus or minus half a frequency bin. The broader band nature of the LPI signals required information from the frequency bin of interest and a number of its neighbours to be processed. The detectors themselves processed the frequency bins in different ways in order to make detections.

DRx: sampling rate	[1.0, 1.75] GHz
sampling period	[571, 1000] ps
DCW length	2048 samples
FFT mode	complex
DCW overlap	0 samples
non-zero padded frequency bin width	[488, 855] kHz
DCW period	[1.17, 2.05] μ s

Table 5.2.2 DRx parameters

Empirical probabilities of detection (P_d) were found with the same procedure detailed in chapter 3, in that 1×10^5 trials per ROC data point were analysed by the detectors. The thresholds for all detectors were selected from earlier P_{fa} against P_d simulations such that they gave $P_{fa} = 0.01$.

The variable which was altered to produce the ROC curve results was SNR_{out} . As before in chapter 3, SNR_{out} referred to the SNR after the FFT block and therefore was independent of the FFT length and its processing gain.

5.2.2 Results and Discussion

ROC curves were obtained for the algorithms of specific interest. These were results based on the detectors::

- DPSS MTM (detector G1);
- sine taper MTM (detector H1);
- adaptive sine taper MTM (detector I1);
- SP (detector J1 variant);
- SMP (detector J1 variant);
- SSP (detector K1);
- TSP (detector L1 variant);
- TSMP (detector L1 variant);
- TSSP (detector M1).

However for clarity only the detectors with the best performance and one of the worst performances are shown in the figures. Fig 5.2.2 displays the results for the Barker-13

pulse input whilst Fig 5.2.3 displays the results for the m-sequence CW input

Panels (a) of Fig 5.2.2 and 5.2.3 display the results from the non-zero padded case and panels (b) from each figure show the results from the twice-zero padded case (results for detectors that employed one of the sine taper MTM can only be included only in (b) panels because the frequency-domain interpretation of the sine taper MTM was used in detectors H1 and I1 and this by its nature involved a twice zero padded FFT).

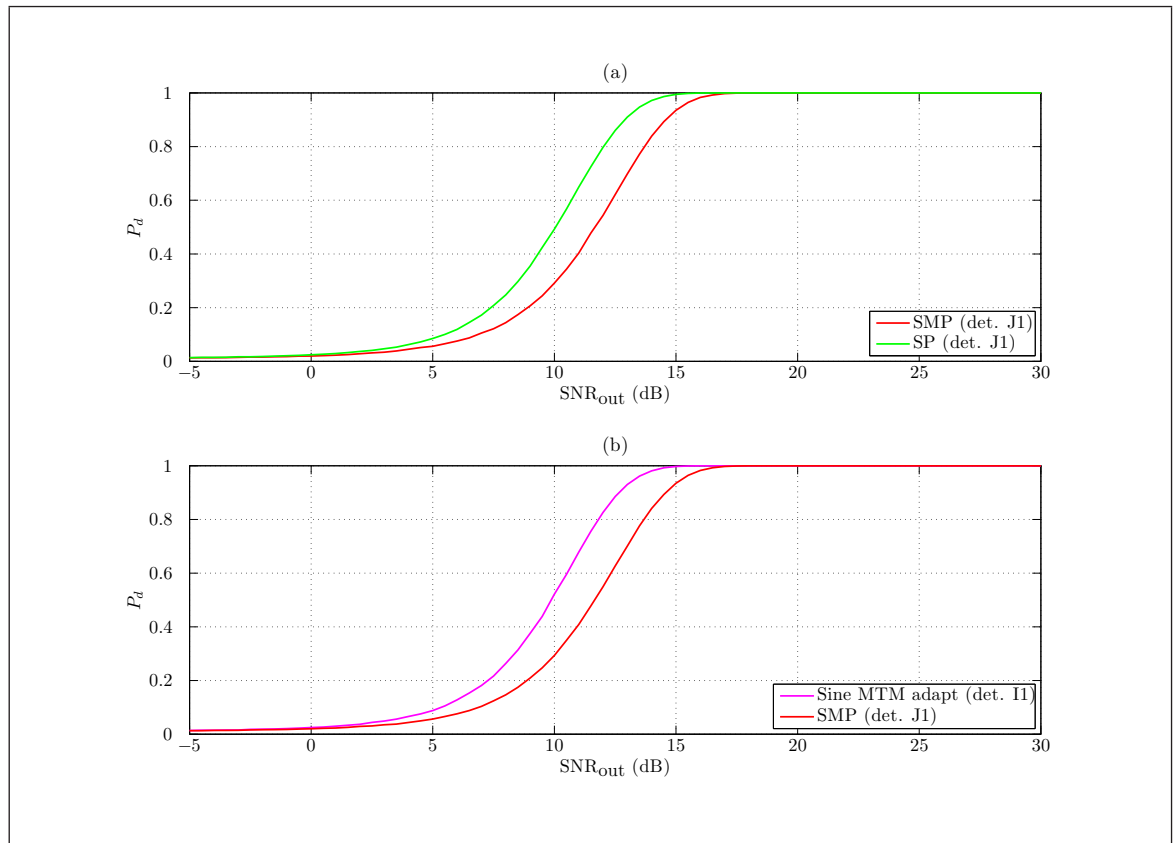


Figure 5.2.2 ROC curves of SNR_{out} against P_d for Barker-13 input and LPI detectors: (a) the non-zero padded case; (b) the zero padded case

In the results obtained and in the figures it is immediate from panel (a) to panel (b) that a twice-zero padded DCW provided a negligible increase in performance for the detectors where this was an option. This result is significant because if there is no major advantage to be gained from a twice-zero padded DCW in a DRx then the computational cost of its implementation would outweigh its benefits. From the full DCW scenario evidence in Figs 5.2.2 and 5.2.3 it would appear that to zero pad affords no improvement in detector ability to find LPI signals.

The results obtained showed two groups of detectors, separated by $\sim 1.5\text{dB}$, with similar performances. This is also evident across Figs 5.2.2 and 5.2.3. The group of detectors with the poorer performances were derived from the modified periodogram (detectors J1 and L1 variants). The reason the detectors based on the modified periodogram did not perform so well as the other detectors was due to the window loss. This mirrors the conclusions drawn from the results obtained in the floating-point SVA section 3.2 for the full DCW scenario.

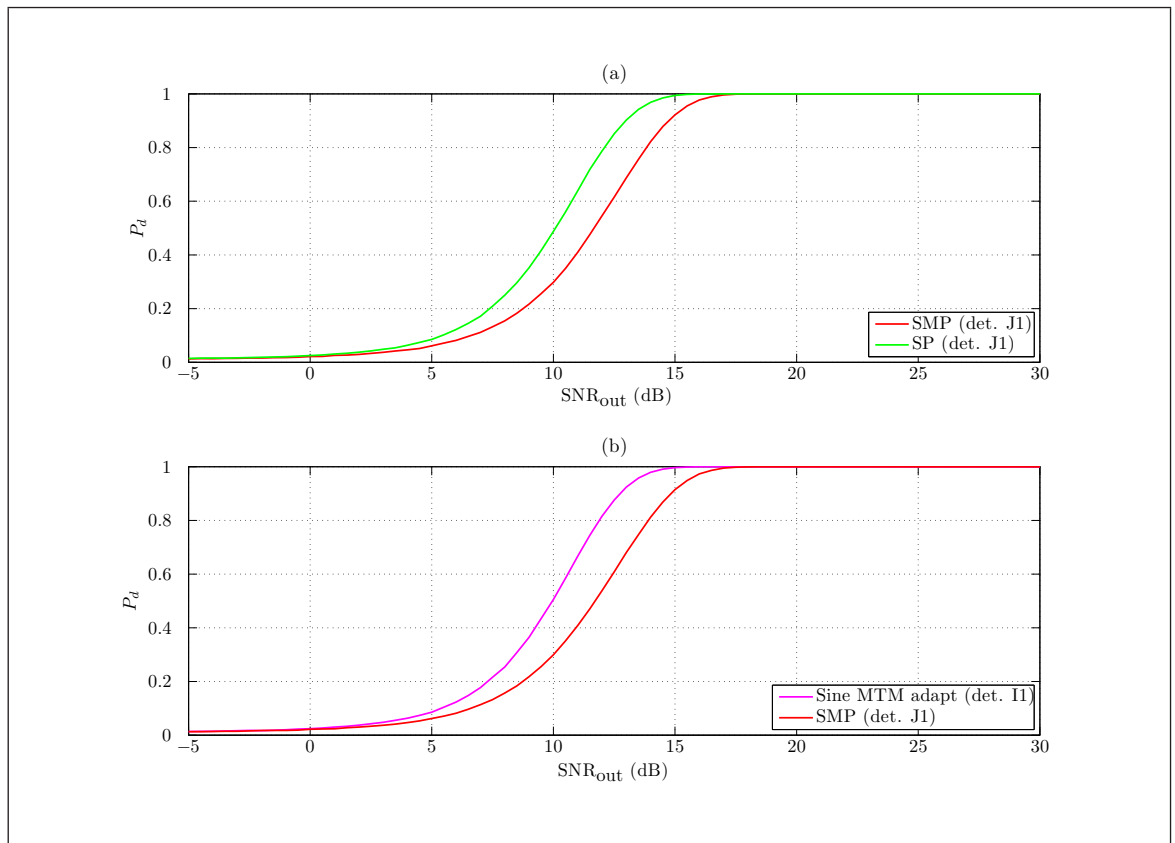


Figure 5.2.3 ROC curves of SNR_{out} against P_d for m-sequence input and LPI detectors: (a) the non-zeropadded case; (b) the zeropadded case

All the other detectors of interest belonged in the group with the better performance. In the (b) panels, the adaptive sine taper MTM (detector I1) and sine taper MTM (detector H1) had the best performances (ROC curves on top of each other, but only adaptive sine taper MTM shown in panels (b) Figs 5.2.2 and 5.2.3). In the (a) panels the smoothed periodogram (detector J1 variant) and tree smoothed periodogram (detector L1 variant) had the best performances (again, ROC curves on top of each other so only smoothed periodogram shown in Figs 5.2.2 and 5.2.3). However this was by a very small margin (~ 0.5 dB).

The explanation for the ~ 0.5 dB degradation in performance of detector G1 in comparison to the other MTM detectors (H1 and I1) is because the sine taper MTM is a minimum bias method whereas the DPSS MTM is concerned with the minimisation of spectral leakage. Since spectral leakage is of no concern in this single signal present scenario, the sine taper MTMs were better suited to the detection process.

The results across Figs 5.2.2 and 5.2.3 were similar because in this full DCW scenario they both appear as BPSK-coded waveforms with only slightly different chip-rates and hence slightly different NW_h products. This manifests itself as only slight changes to the ROC curve patterns between Figs 5.2.2 and 5.2.3.

The results obtained from the conventional detection algorithms demonstrated the dangers of simple reliance on the modified periodogram (detector A1 variant), periodogram

(detector A1 variant), total power (detector P1) or even the SVA periodogram (detector B1) for detection of an LPI signal from an FFT output where the equivalent filter bandwidths are less than the LPI signal bandwidth. These techniques, except the SVA periodogram, are all well-established techniques for use in DRxs. Therefore the results indicated that simple thresholds on individual ordinates, which is the norm in many detectors, is a poor method for the detection of LPI signals whose power is spread-out across frequency space. As extracted results, the periodogram (detector A1 variant) and total power (detector P1) are shown in Figs 5.2.4 and 5.2.5 alongside the best performance conventional detector (OR periodogram, detector N1).

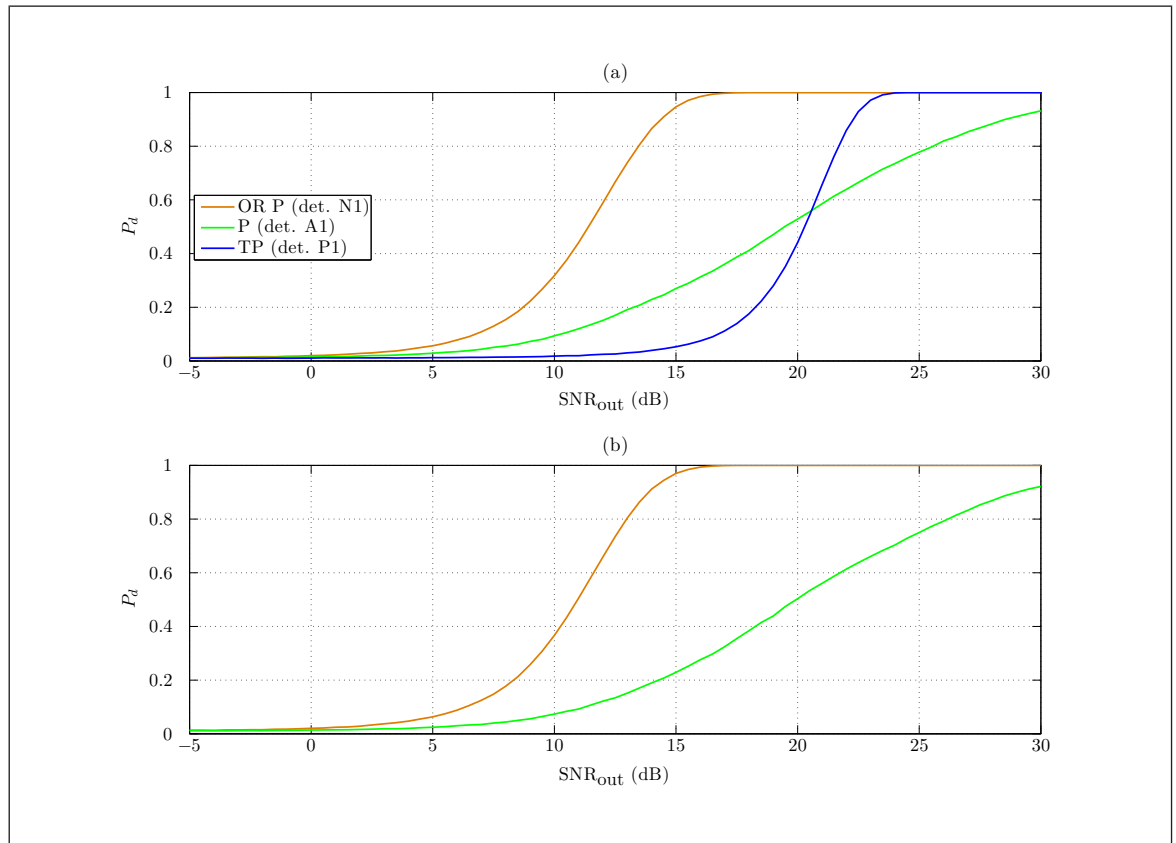


Figure 5.2.4 ROC curves of SNR_{out} against P_d for Barker-13 input and conventional detectors: (a) the non-zeropadded case; (b) the zeropadded case

The performance of detector P1 which was based on the total power contained within a DCW did not change across the two different signal types. This was because the total signal power in the DCWs did not change with modulation type in the full DCW scenario.

Another point was that the use of a twice-zeropadded DCW (similar to the findings from the algorithms of interest) did not make a significant difference to the performance of the conventional detectors against LPI signals.

The SVA periodogram (detector B1) performed particularly poorly against the Barker-13 pulse. In general, this is because the SVA technique is ideally suited to narrowband sinusoidal signal detection. However the detector performance of the SVA periodogram

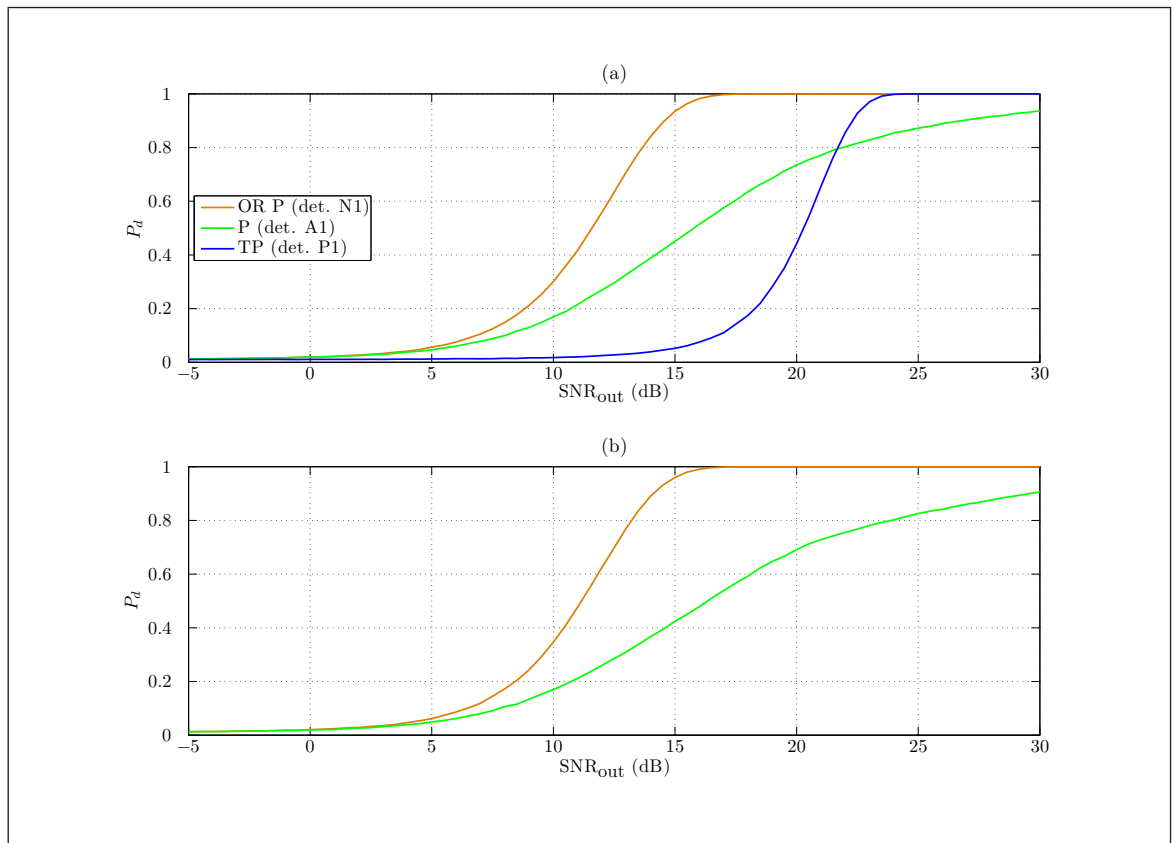


Figure 5.2.5 ROC curves of SNR_{out} against P_d for m-sequence input and conventional detectors: (a) the non-zeropadded case; (b) the zeropadded case

in addition to the periodogram and modified periodogram all improved once the detection logic included an OR operation between the monitored ordinates for the presence of the LPI signal. This is precisely how the Barker-13 pulse was detected in section 3.6. The improvement in performance with OR gates is illustrated in Figs 5.2.4 and 5.2.5 by the jump from the periodogram to the best detector from the conventional technique, i.e. the OR periodogram (detector N1).

Despite the use of an OR operation, detectors N1 and O1 did not perform as well as the detectors based on the MTMs or the best smoothing detectors. Their detector performance was found to be ~ 1.5 dB poorer. The conclusion from this result was that most of the algorithms of interest (except for the SMP and TSMP) were better than the best conventional detection techniques in the full DCW scenario.

5.3 Randomly-filled DCW

5.3.1 Scenario and Simulation Description

For this scenario, the simulations were designed to use the same Monte Carlo framework as the previous full DCW scenario. Within the framework, the signal characteristics, noise characteristics and detectors were the same. The only change was that the full DCW constraint was relaxed. In this scenario the effect on detector performances of the interception of a random part of the Barker-13 LPI pulse was studied.

Although the duration of the test Barker-13 pulse from Table 5.2.1 was longer than the design DCW from Table 5.2.2, there could conceivably be cases where only the leading or trailing part of the pulse would be captured by a superhet and DRx combination. Therefore in this scenario a random portion of the Barker-13 pulse was captured in the same detectors from the full DCW scenario in section 5.2. The constraint on this randomly-filled DCW scenario was that a DCW would contain at least one sample with a contribution from the signal power. This concept is illustrated in Fig 5.3.1.

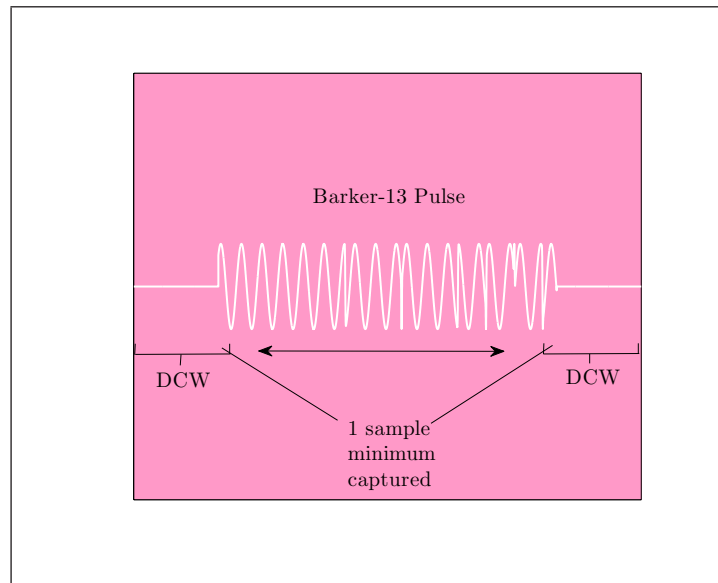


Figure 5.3.1 Barker-13 pulse captured in a DCW

This constraint was necessary in order to produce ROC curves of SNR_{out} against P_d . The constraint was also different to that used in the randomly-filled DCW scenario in section 3.3 which tested detectors A1 and B1 against sinusoidal pulses (i.e. here the assumption of several DCWs per ESM dwell was removed). In section 3.3 the detectors were tested to their design limits through searches for sinusoidal pulses whose duration was the same as the DCW duration. Clearly in this section the pulses have a much longer duration. In fact in many instances, in the randomly-filled DCW scenario in this section the DCWs would be completely filled with samples which contained a signal power contribution (as the Barker-13 pulse duration was over twice as long as the DCW

duration). Therefore to obtain at least some partially-filled DCWs the assumption of several DCWs per ESM dwell was removed.

The m-sequence CW LPI signal was not tested under this scenario because the results would be identical to those from the full DCW scenario in section 5.2.

5.3.2 Results and Discussion

For clarity, only the extracted best and one of the worst detectors which utilised the algorithms of interest are shown in Fig 5.3.2 and only a few conventional detectors are shown in Fig 5.3.3 to illustrate the important trends.

The ROC curve results were very distorted from the conventional ROC curve shape as a consequence of the capture of a random portion of the Barker-13 pulse (this effect is evident in Figs 5.3.2 and 5.3.3). The results showed that each individual detector performance was degraded in comparison to its performance in the full DCW scenario. This was expected because on average there was less signal energy in the randomly-filled DCWs.

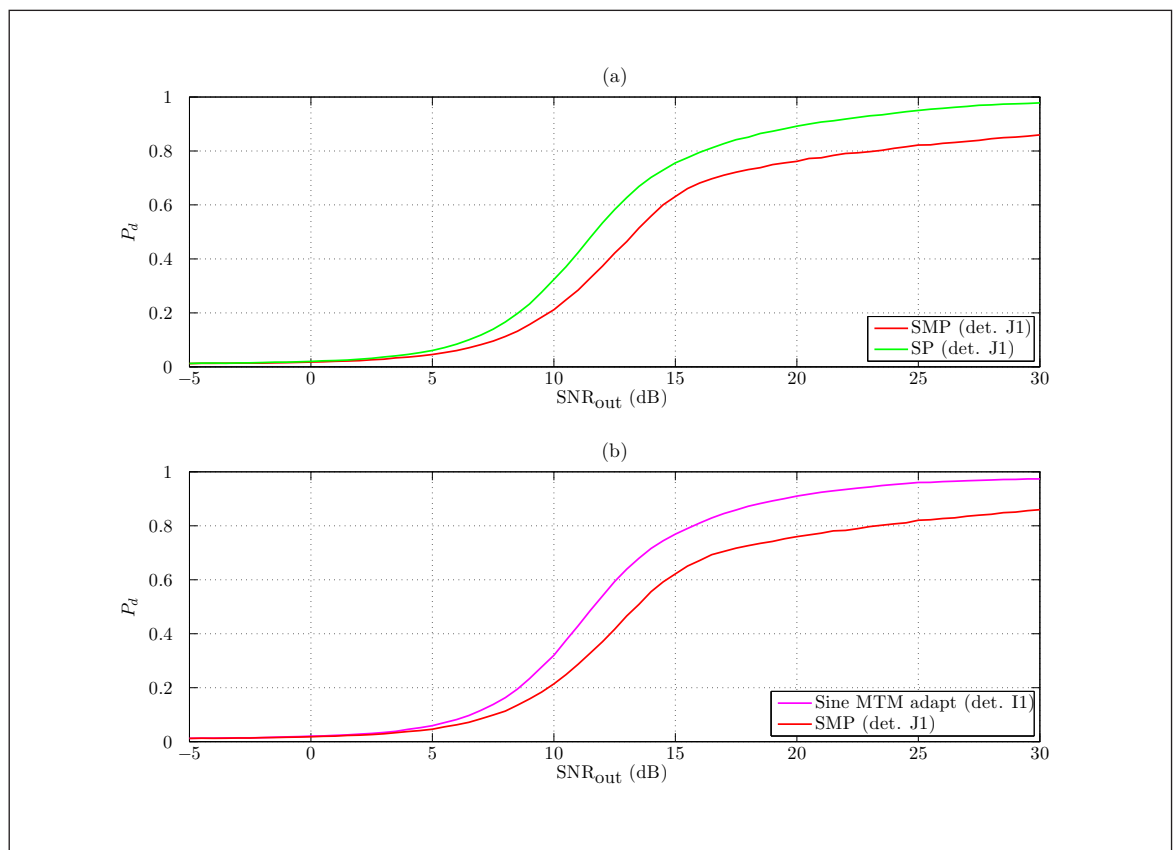


Figure 5.3.2 ROC curves of SNR_{out} against P_d for Barker-13 input and LPI detectors: (a) the non-zero padded case; (b) the zero padded case

For the algorithms of interest and low SNR_{out} values the relative detector performances remained the same as in the full DCW scenario. However as the SNR_{out} increased past ~ 13 dB, the gradient of the performances of the: (1) SMP (detector J1 variant);

(2) TSMP (detector L1 variant); (3) SSP (detector K1) and (4) TSSP (detector M1) diverged from the rest of the detectors. This was caused by the loss of signal power at the edges of the DCWs. The same effect was observed with the detectors which utilised the SVA algorithm as SVA minimised spectral leakage, which in this scenario had the adverse effect to select the equivalent of a Hann window function for some frequency bins. The effect was not quite as severe as the application of a fixed Hann window across all frequency bins, nevertheless it had a detrimental effect on the performance of the SSP and TSSP (detectors K1 and M1). This effect of the utilisation of SVA was not noticeable in the partially-filled DCW scenario in section 3.3 because in that scenario the condition had been set that at least one DCW would be at least half full of samples with a signal power contribution.

From the algorithms of interest, aside from the detectors based on the modified periodogram (detectors J1 and L1 variants) and SVA periodogram (detectors K1 and M1), the other detectors demonstrated similar relative performances to the full DCW scenario. Therefore the same conclusions that were drawn about those detectors in section 5.2 also hold in this scenario.

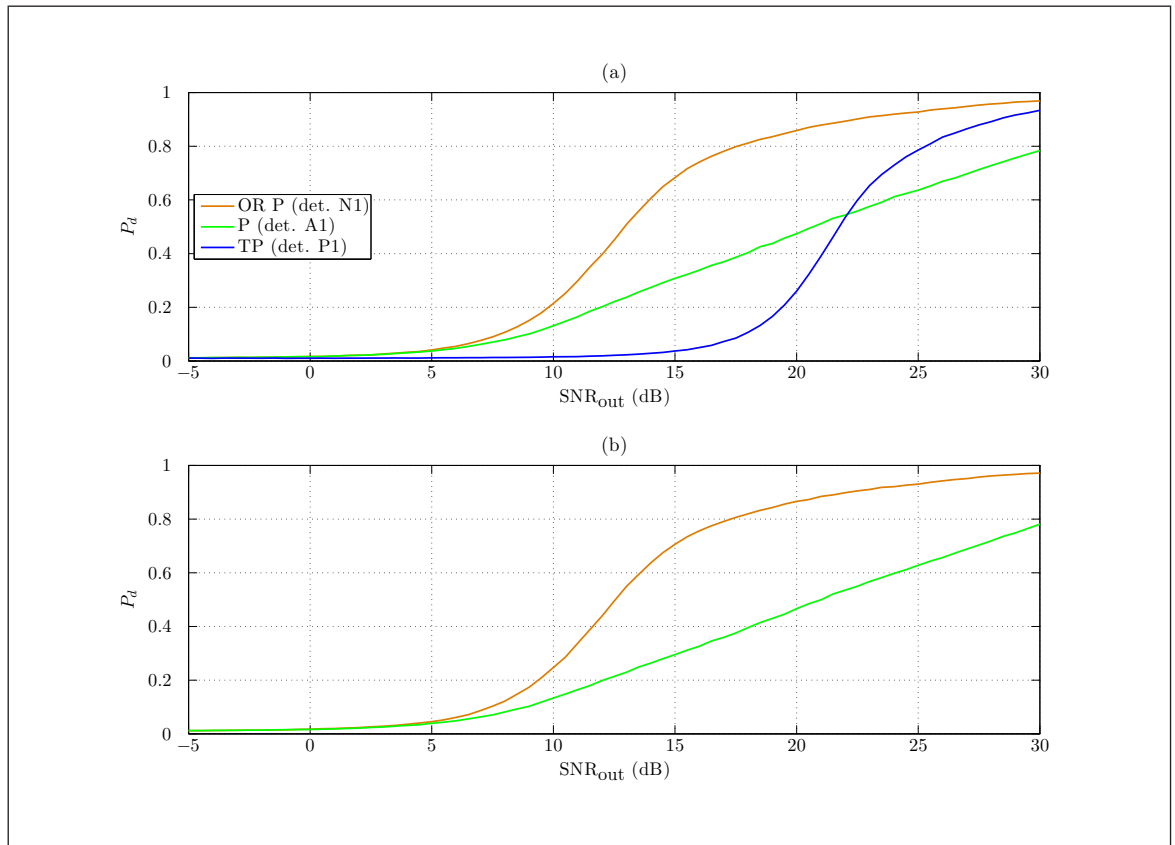


Figure 5.3.3 ROC curves of SNR_{out} against P_d for Barker-13 input and conventional detectors: (a) the non-zero padded case; (b) the zero padded case

From the conventional detectors, the results based on the modified periodogram (detectors A1 and N1 variants) and SVA periodogram (detectors B1 and O1) were similarly affected by the attenuation of samples at the edges of DCWs. However the overall

relative detector performances remained the same as the corresponding detector performance from the full DCW case.

With consideration of the results from the full DCW scenario and partially-filled DCW scenario, it is clear that the algorithms of interest outperformed the detectors which utilised conventional methods. The performance difference can be very large at ~ 10 dB if a comparison is made between the simplest conventional periodogram (detector A1) and the more complicated sine taper MTMs (detectors H1 and I1). However within the algorithms of interest the increase in computational cost does not lead directly to an increase in detector performance. For example sine taper MTM (detector H1) was only fractionally better than the SP (detector J1).

Upon this evidence it would be tempting to state that the SSP or TSSP (detectors K1 or M1) would be the best all-round detectors because the SVA algorithm would ensure good detector performance against narrowband, sinusoidal signals and the smoothing would ensure good performance against LPI signal. These detectors are certainly less computationally complicated than the detectors based on the MTMs so their implementation would make good engineering sense. However it was necessary to examine the scenarios where there was interference present as it was unknown how the algorithms of interest would perform under these conditions. The next two sections deal with this issue.

5.4 Interference in Full DCW

5.4.1 Scenario and Simulation Description

The task for the detectors became more difficult in the interference DCW scenario. The detectors searched for one of the LPI signals from Table 5.2.1 with an interference sinusoid (S_B) of frequency $f_{T,B}$ present adjacent to the bandwidth of the LPI signal. The desired LPI signal (S_A) and the interference sinusoidal signal (S_B) were both present for the duration of the DCW. This represented the scenario where a DRx would coincidentally capture as much signal power from S_A and S_B as temporally possible in a DCW.

The target frequencies of S_A and S_B were chosen at random from a frequency bin interval at the start of each test run. Therefore, although the exact frequency separation between S_A and S_B varied, there was, on average a nominal frequency separation. The relationship between the frequencies, in the case of the Barker-13 signal, is demonstrated in Fig 5.4.1. As examples of the randomly chosen frequencies, panel (a) shows the case where the frequencies of S_A and S_B align exactly on frequency bins, panel (b) shows the case where the frequencies are maximally separated and panel (c) displays the case where the frequencies are minimally separated.

A very similar relationship was true for the m-sequence signal and the interference sinusoid except the bandwidth of the m-sequence signal was slightly larger than for the Barker-13 pulse. The larger bandwidth led to a smaller possible frequency separation between S_B and the bandwidth edges of S_A .

The differences between the power levels of S_A and S_B and the power levels of S_B and the noise floor were expressed as SIR_{out} and SNR_{out} with the use of the previous definitions from section 3.4. To produce ROC curve results the SIR_{out} level was varied for $P_{fa} = 0.01$ using the methodology and definitions of H_0 and H_1 from section 3.4.

Essentially all aspects of the LPI signals from the full DCW scenario in section 5.2 were maintained and the introduction of an interference tone (S_B) was the major difference. Therefore the comments on the LPI signals (S_A) from section 5.2 are also valid for this section. For example remarks on the random initial phase and random code portions of the signal captured remain accurate. The introduction of an interference tone (S_B) was a simple extension to the signal generation simulation such that the same basic ideas were reused. The sinusoid (S_B) was generated with a random initial phase and it was present throughout the duration of the DCW. Therefore every time-domain sample in a DCW contained a power contribution from the LPI signal, the interference tone and the AWGN background.

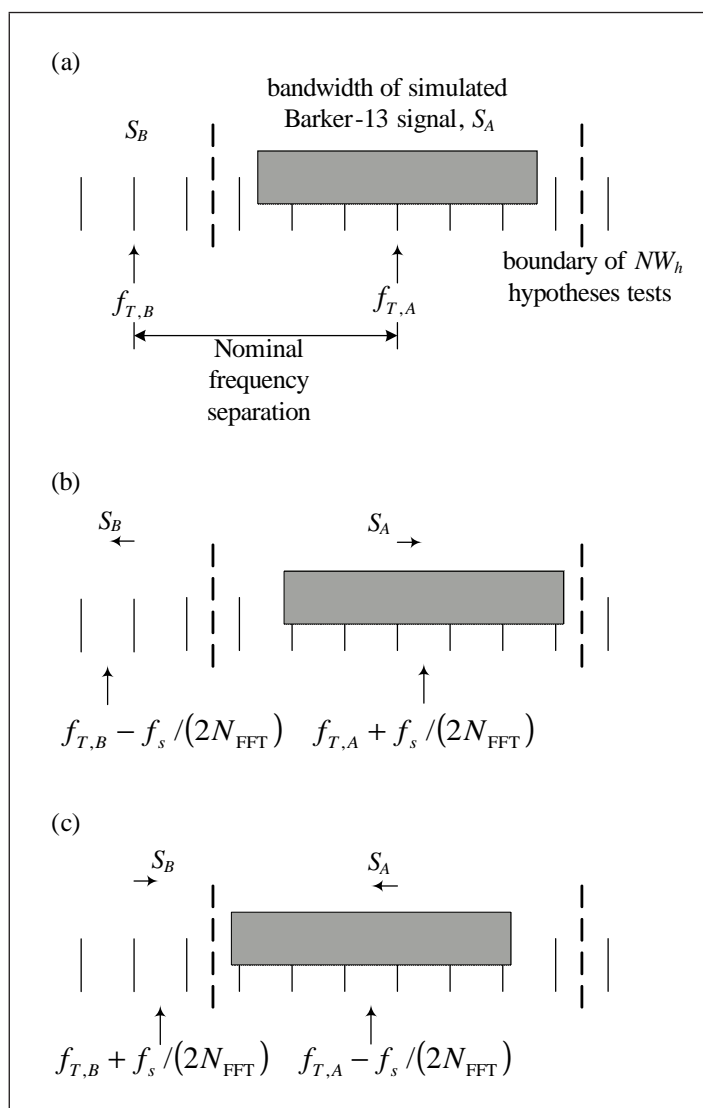


Figure 5.4.1 S_B adjacent to monitored bandwidth in: (a) S_A coincides with central monitored frequency; (b) S_A and S_B are maximally separated; (c) S_A and S_B are minimally separated

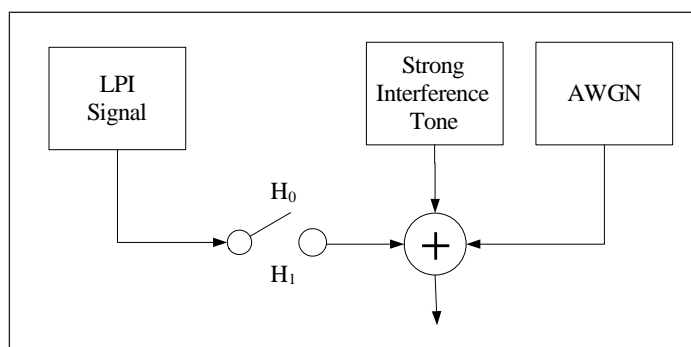


Figure 5.4.2 Interference in full DCW scenario detector input

Some of the detectors used in the full DCW and randomly-filled DCW scenarios were simply re-implemented with the trivial change of the addition of the interference tone (S_B) at the input in this scenario. Pictorially this small alteration is shown in Fig 5.4.2. This was because these detectors were designed with the use of techniques which do not require explicit steps to remove sidelobe spectral leakage or take account of mainlobe width. This is a positive feature of these detectors since they are naturally more versatile to different scenarios and do not require post-FFT steps to manually mitigate against sidelobes spectral leakage or broader mainlobes. From the algorithms of interest these were:

- DPSS MTM (detector G1)
- unity-weighted sine taper MTM (detector H1)
- adaptive sine taper MTM (detector I1)
- SSP (detector K1)
- TSSP (detector M1).

All that is required to obtain a system level description is to replace the input section of each detector diagram from Figs 5.1.2–5.1.11 with that of Fig 5.4.2.

The conventional detection techniques which did not require alterations for this scenario were the:

- OR SVA periodogram (detector O1)
- SVA periodogram (detector E1)
- total power (detector P1).

The detectors which utilised the SVA technique did not require post-FFT processing steps due to the conclusion drawn in section 3.4 from Fig 3.4.12, that is, the SVA technique reduced sidelobe spectral leakage and mainlobe width sufficiently that they could be ignored in realistic scenarios.

However the detectors which were fundamentally based on the procedure which involved the application of a weighted window function and an FFT **did** require steps to recognise the frequency-domain response of the window function. These necessary steps are analogous to the decision rules which exist in ESM DRxs to minimise erroneous detections as a result of spectral peaks in neighbour ordinates. The steps are outlined in section 3.4 and demonstrated visually in Fig 3.4.3 and therefore are not repeated here.

From the set of the algorithms of interest the following were modified from their corresponding form in the full DCW and partially-filled DCW scenarios:

- SP (detector J1)

- TSP (detector L1).

Also from the conventional techniques the following required modification:

- OR periodogram (detector N1)
- periodogram (detector A1).

The changes resulted in detectors with different labels:

detector J1 \rightarrow detector Q1

detector L1 \rightarrow detector R1

detector N1 \rightarrow detector S1

detector A1 \rightarrow detector C1

A diagram of detector C1 is shown in Fig 3.4.1 in section 3.4 and diagrams of detectors Q1, R1 and S1 are shown in Figs 5.4.3–5.4.5.

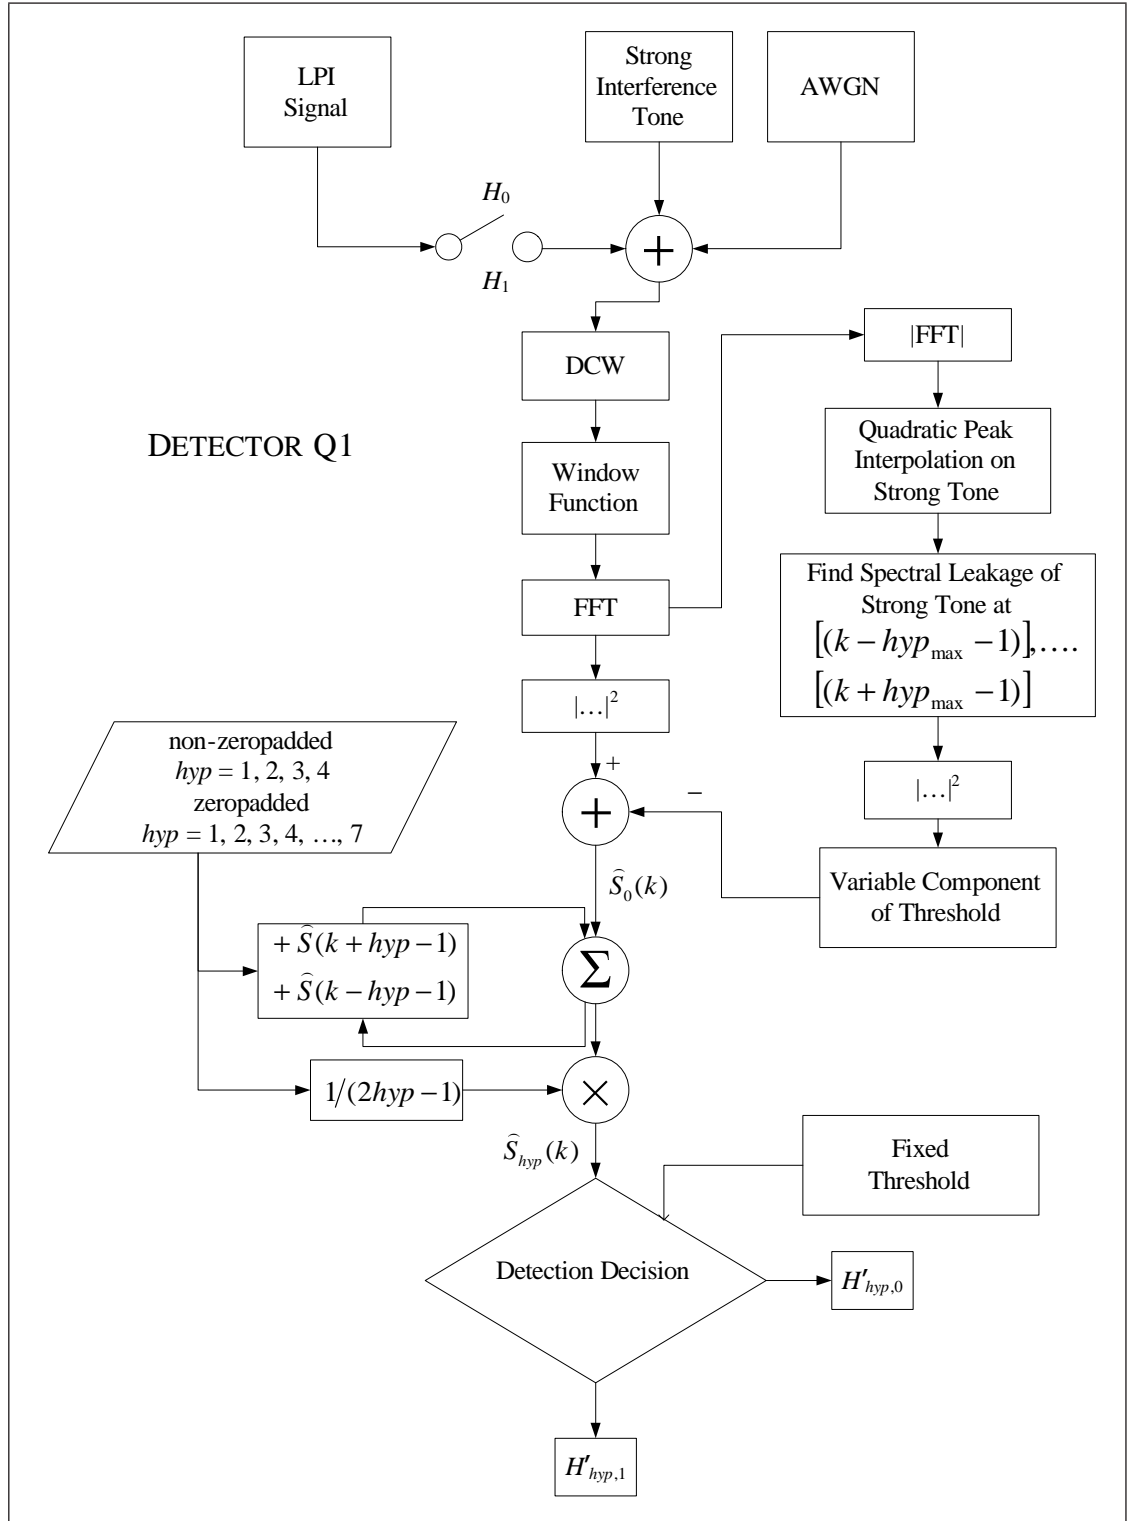
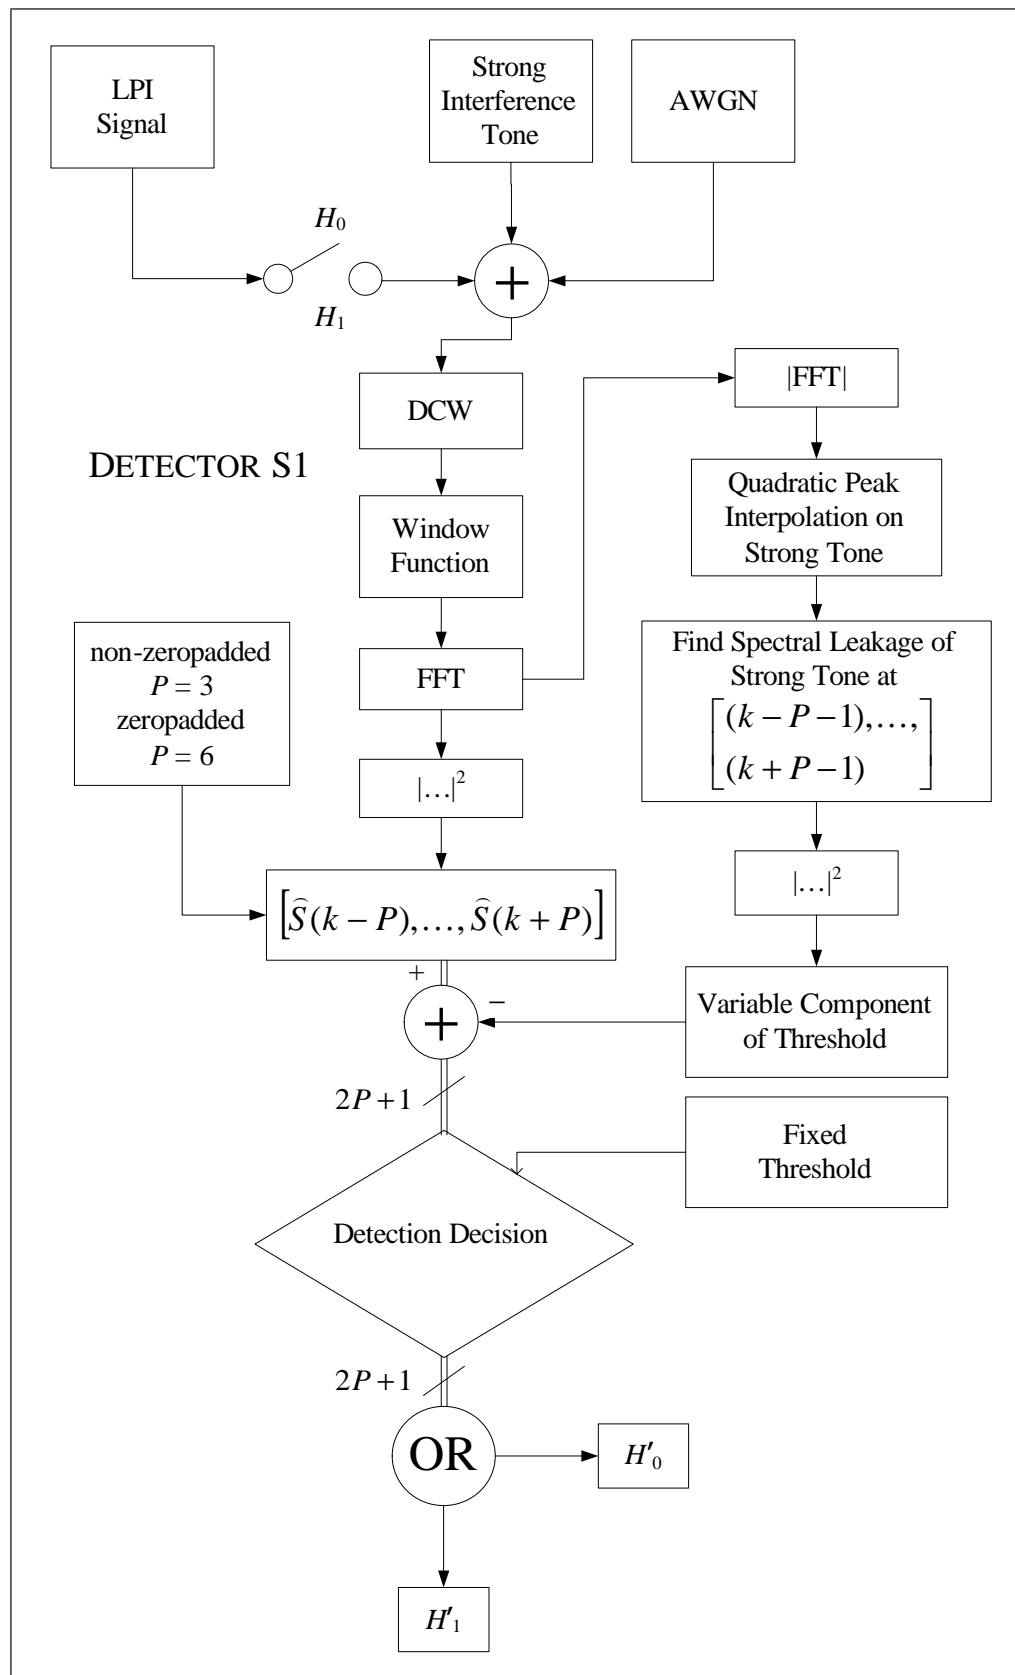


Figure 5.4.3 Detector Q1 for the interference DCW scenarios



5.4.2 Results and Discussion

Low $\text{SNR}_{\text{out},B}$

In the low $\text{SNR}_{\text{out},B}$ regime the power ratio between S_B and the noise floor was $\text{SNR}_{\text{out},B} = 8$ dB and the nominal frequency separation between test signal (S_A) and interference tone (S_B) was $5f_s/N$ (which was also maintained in the medium and high $\text{SNR}_{\text{out},B}$ regimes). For the algorithms of interest tested against a Barker-13 pulse the relative performance of the detectors in Fig 5.4.6 from best to worst in both zeropadded and non-zeropadded cases was:

1. DPSS MTM (detector G1)
2. unity-weighted sine taper MTM (detector H1), SP (detector Q1 variant) and TSP (detector R1 variant)
3. SSP (detector K1) and TSSP (detector M1)
4. adaptive sine taper MTM (detector I1)
5. SMP (detector Q1 variant) and TSMP (detector R1 variant).

For clarity, only the best and worst detectors from the algorithms of interest are shown in Fig 5.4.6. The smoothed and tree smoothed detectors for all the various periodogram types are grouped together in the above list however there was a small performance drop in the tree smoothed equivalent detectors. This was caused by the coarser bandwidth hypothesis reduction of the tree smoothed detectors at larger hypothesis bandwidths.

The SMP (detector Q1 variant) and TSMP (detector R1 variant) had the worst performance and were separated from the next best detector by ~ 1 dB. The other detectors were grouped together in a performance range of ~ 0.6 dB in the non-zeropadded case and ~ 0.8 dB in the zeropadded case.

Overall the relative results were as expected. In the low- $\text{SNR}_{\text{out},B} = 8$ dB regime, the noise rather than the spectral leakage was dominant and hence the results obtained to some extent resemble those obtained from the full DCW scenario. However there were slight differences in the results between this scenario and the full DCW scenario. Specifically, in the full DCW scenario the DPSS MTM (detector G1) had a performance which was almost the same as the SP (detector J1 variant) and also the sine taper MTMs (detectors H1 and I1) actually had the best detector performance. In this scenario, however, the best detector was the DPSS MTM (detector G1) and this was attributed to its slightly better spectral leakage protection properties than the sine taper MTM (detector H1) (although the impact of spectral leakage was very small in this low- $\text{SNR}_{\text{out},B}$ regime).

The relative detector performance for the algorithms of interest against an m-sequence input was almost identical to the results obtained against the Barker-13 pulse in the

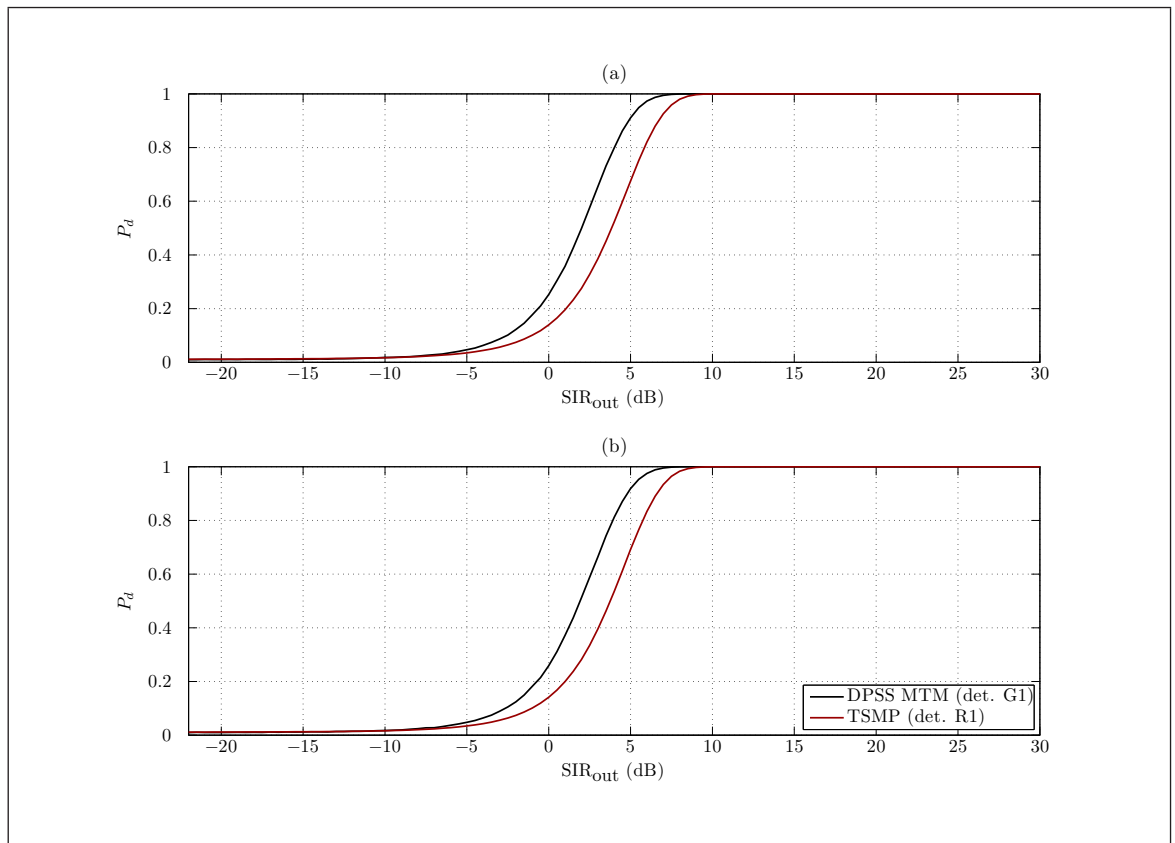


Figure 5.4.6 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $\text{SNR}_{\text{out},B} = 8$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

low- $\text{SNR}_{\text{out},B}$ regime. Extracted results for the best and worst detectors and shown in Fig 5.4.7.

In the same regime and for the conventional detectors against the Barker-13 signal input the relative detector performances in order from best to worst was found to be:

1. OR periodogram (detector S1 variant)
2. OR SVA periodogram (detector Q1 variant)
3. OR modified periodogram (detector S1 variant)
4. modified periodogram (detector C1 variant)
5. periodogram (detector C1 variant)
6. SVA periodogram (detector E1 variant).

For clarity only some results are shown in Fig 5.4.8.

In the results there were slight differences in the relative detector performances between the non-zeropadded case and zeropadded case. In the non-zeropadded FFT case there was a performance gap of ~ 0.12 dB between the OR periodogram (detector S1 variant) and the OR SVA periodogram (detector O1) but in the zeropadded case the performance gap widened to ~ 0.22 dB.

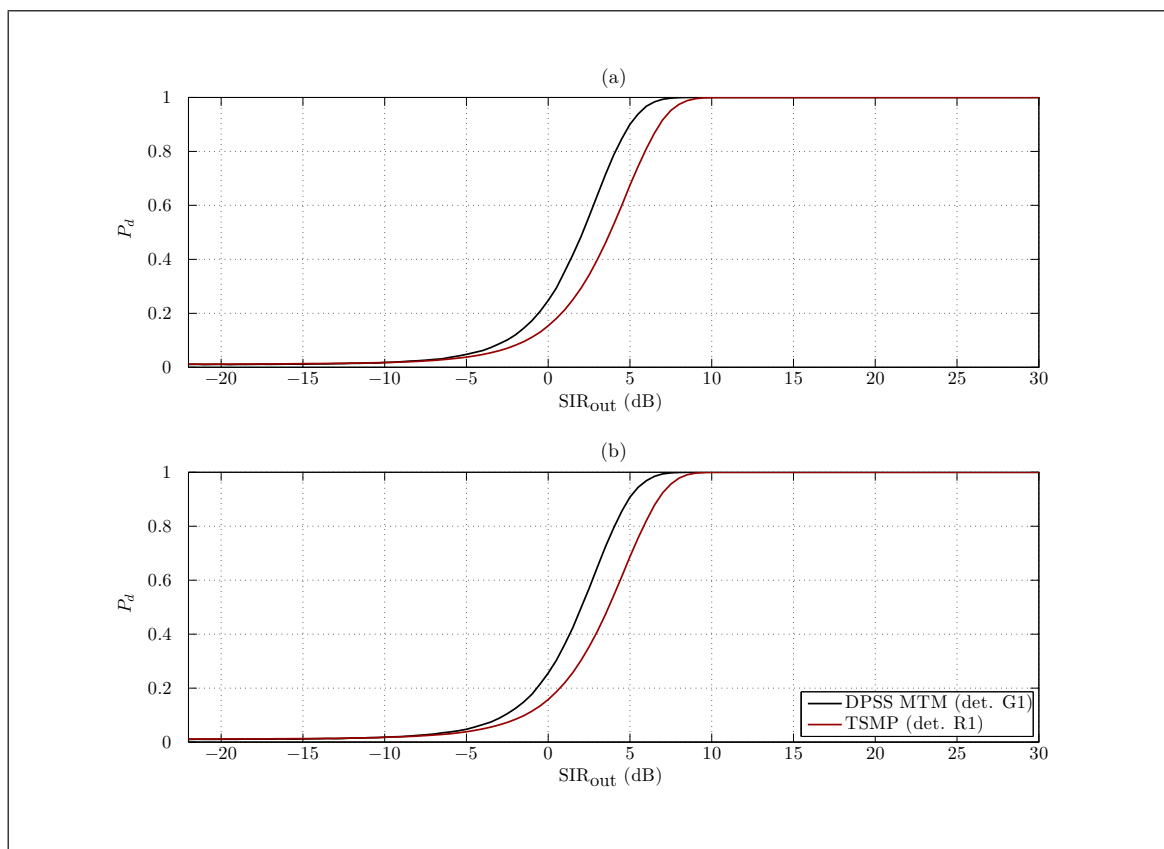


Figure 5.4.7 ROC curves of SIR_{out} against P_d for m-sequence input and LPI detectors at $\text{SNR}_{\text{out},B} = 8$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

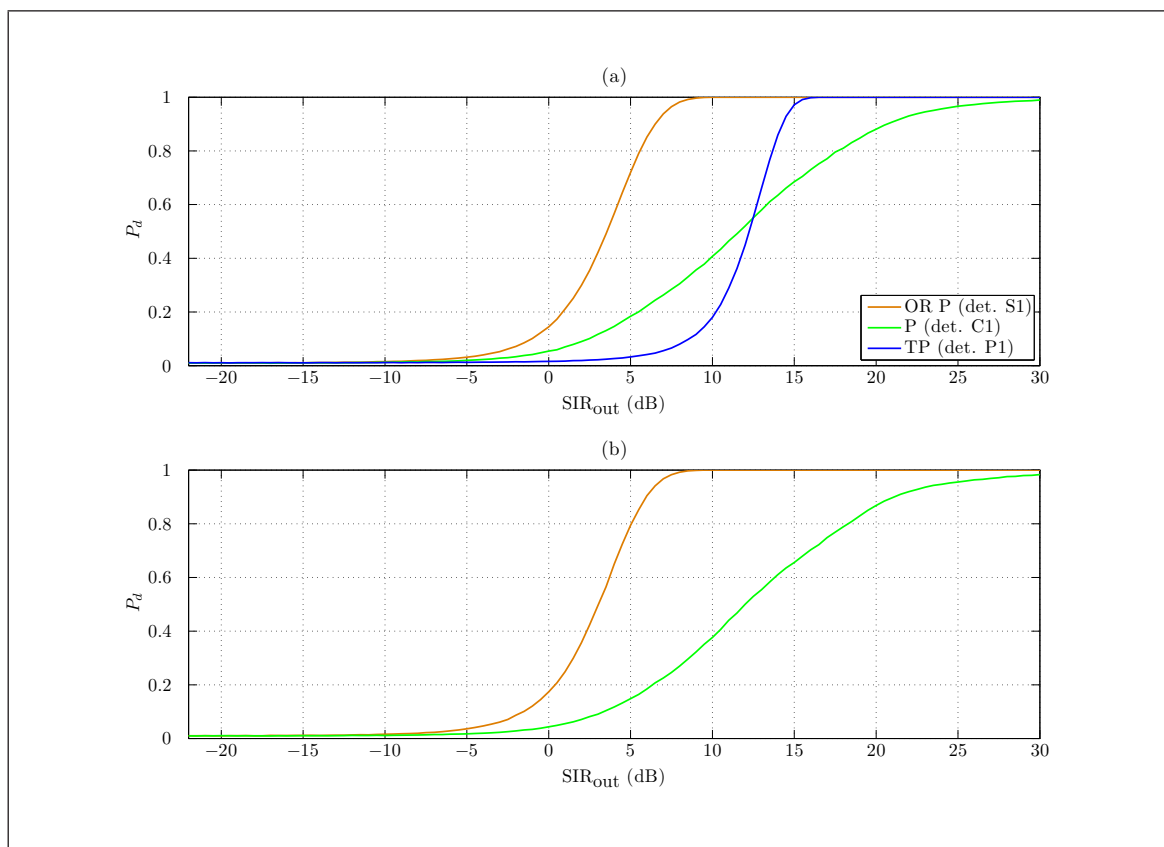


Figure 5.4.8 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 8$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

The results obtained from the conventional detectors against the m-sequence input had some noteworthy differences in comparison to the Barker-13 input. The main difference was the modified periodogram (detector C1 variant) and periodogram (detector C1 variant) switched places in the relative detector performances. Another difference was that the performance gap between the OR periodogram (detector S1 variant) and OR SVA periodogram (detector O1) actually widened in both zeropadded and non-zeropadded cases for the m-sequence test signal. Extracted results for the m-sequence test signal and conventional detectors are shown in Fig 5.4.9.

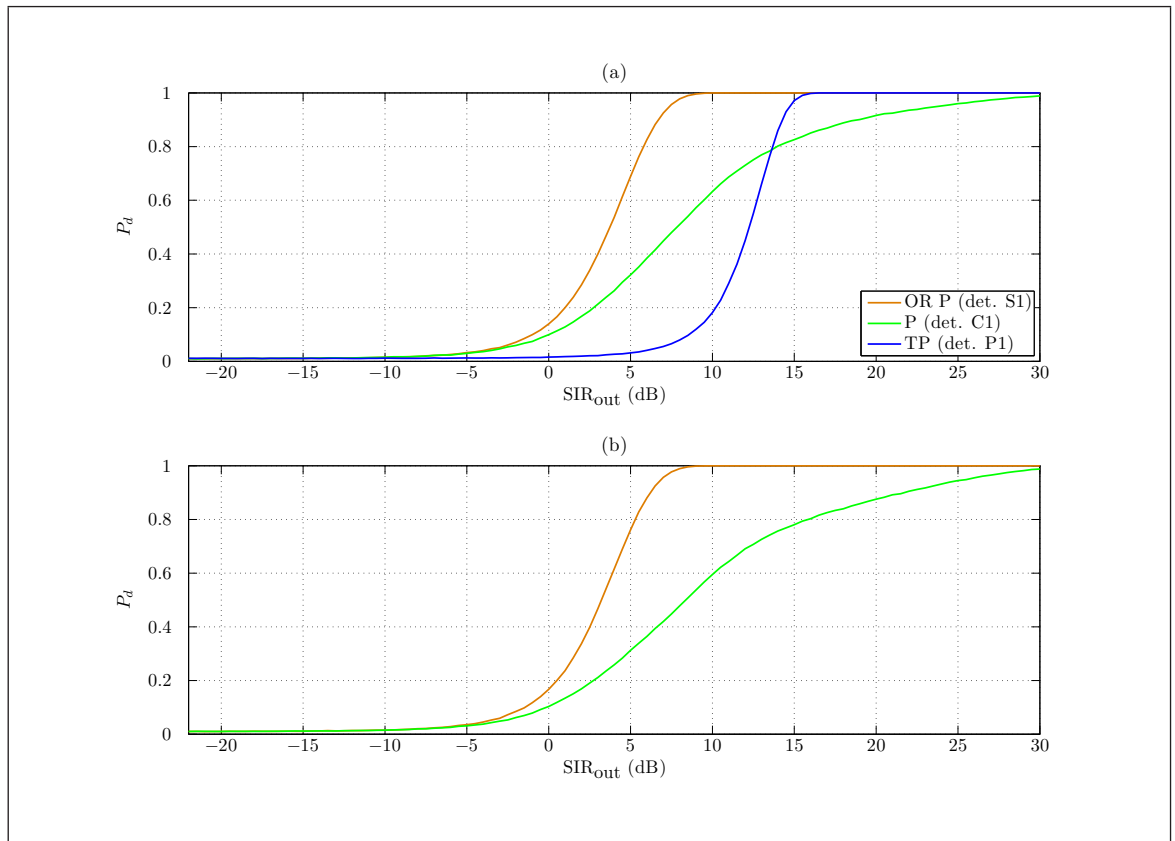


Figure 5.4.9 ROC curves of SIR_{out} against P_d for m-sequence input and conventional detectors at $\text{SNR}_{\text{out},B} = 8 \text{ dB}$ for: (a) the non-zeropadded case; (b) the zeropadded case

The clear message from the conventional detector results was that the periodogram, modified periodogram and SVA periodogram (detectors C1 and E1) from the SVA chapter can be greatly improved upon for the detection of LPI signals. The ROC curves for these detectors were heavily distorted by the presence of an interference tone and they are less suitable as LPI signal detectors in comparison to those based on the algorithms of interest and the conventional OR detectors.

The total power detector ROC curve shape was undistorted, however as it non-coherently sums power in samples it had a poorer processing gain than the OR detectors. The observed performance difference between the worst OR detector and the total power detector was $\sim 8 \text{ dB}$.

In a situation where only a conventional detector could be used, the OR periodogram (detector S1 variant) or OR SVA periodogram (detector O1) would clearly be the best options. However these two conventional detectors had a performance which was degraded by $\sim [1.1, 1.6]$ dB (depending on the zeropad case and signal input) against the best detector which used an algorithm of interest (the DPSS MTM (detector G1) in this low- $\text{SNR}_{\text{out},B}$ regime). All the detectors which used algorithms of interest had reasonably close performances except the SMP (detector Q1 variant) and TSMP (detector R1 variant). Therefore on the evidence from the low- $\text{SNR}_{\text{out},B}$ regime, the algorithms of interest (except the aforementioned variants of detectors Q1 and R1) are improvements over conventional DRx detection techniques in the search for LPI signals.

Moderate $\text{SNR}_{\text{out},B}$

In this regime the $\text{SNR}_{\text{out},B}$ was set at 15 dB. For both Barker-13 and m-sequence test signals the relative performance of the algorithms of interest was broadly consistent with the $\text{SNR}_{\text{out},B} = 8$ dB results except for one significant difference. The sine taper MTM (Detector I1) had the worst detector performance in the moderate- $\text{SNR}_{\text{out},B}$ regime from the algorithms of interest. The apparently erratic detection performance of the sine taper MTM (detector I1) from one of the best detectors in the low- $\text{SNR}_{\text{out},B}$ regime to the worst detector in the moderate- $\text{SNR}_{\text{out},B}$ regime can be ascribed to the particular set of parameters with which it was configured for all the simulations in this chapter. In other words, in some scenarios the parameters were a good match, whereas in others they were a poor match.

The best and worst detectors for the algorithms of interest are shown in Fig 5.4.10 for the Barker-13 test signal and in Fig 5.4.11 for the m-sequence test signal.

The slight degradation in the performance results between the smoothed and corresponding tree smoothed detectors was also present in the moderate- $\text{SNR}_{\text{out},B}$ regime results as it had been for the low- $\text{SNR}_{\text{out},B}$ regime results. This was for the same reasons that the tree smoothed detectors were coarser in terms of bandwidth hypothesis.

In comparison to the low- $\text{SNR}_{\text{out},B}$ regime results, the performance of the sine taper MTM (detector H1) improved slightly in the moderate- $\text{SNR}_{\text{out},B}$ regime such that there was no discernible difference between the sine taper MTM (detector H1) and DPSS MTM (detector G1). This was true for both Barker-13 pulse and m-sequence test signals in both non-zeropadded and zeropadded cases.

In comparison to the low- $\text{SNR}_{\text{out},B}$ regime, the moderate- $\text{SNR}_{\text{out},B}$ regime saw the influence of sidelobe spectral leakage grow and the dominance of the noise floor diminish. This was demonstrated by the smaller detector performance gap from the low- $\text{SNR}_{\text{out},B}$ regime to moderate- $\text{SNR}_{\text{out},B}$ regime between the grouping of the SP/TSP (detectors Q1 variant/R1 variant) and the grouping of the SSP/TSSP (detectors K1/M1).

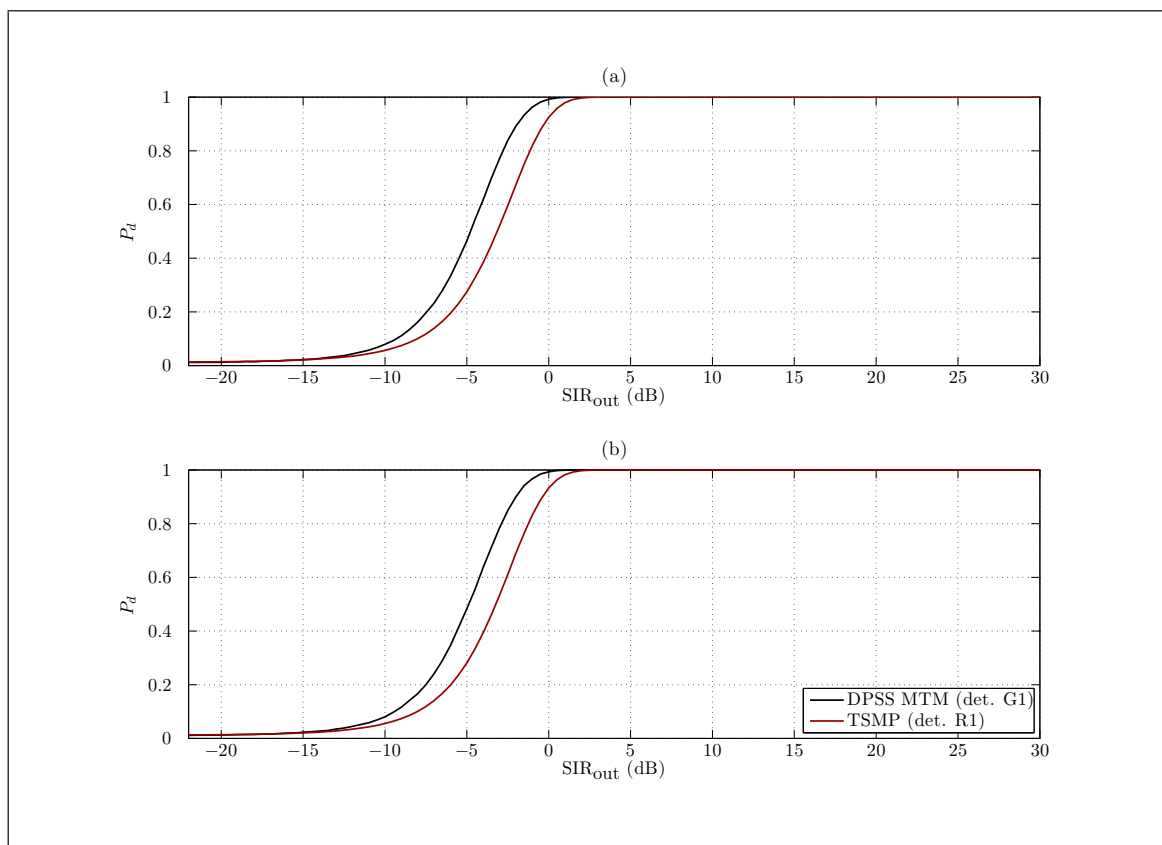


Figure 5.4.10 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $\text{SNR}_{\text{out},B} = 15$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

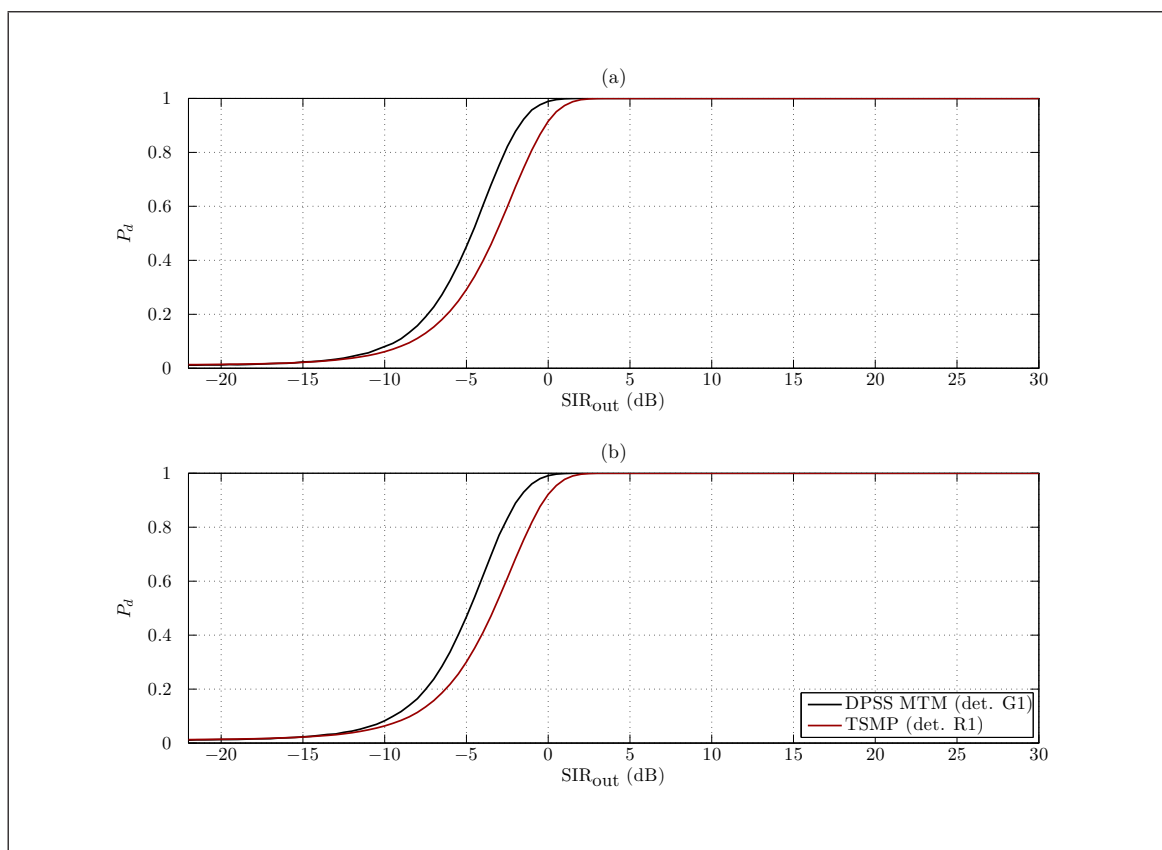


Figure 5.4.11 ROC curves of SIR_{out} against P_d for m-sequence input and LPI detectors at $\text{SNR}_{\text{out},B} = 15$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

The SVA algorithm offers better protection against spectral leakage with a small window loss. Therefore as sidelobe spectral leakage became a more significant effect, the SSP and TSSP (detectors K1 and M1) improved their performance relative to the SP (detector Q1 variant) and TSP (detector R1).

The results obtained from the conventional detectors for the moderate- $\text{SNR}_{\text{out},B}$ regime showed that for the Barker-13 pulse, the OR SVA periodogram (detector O1) performed the best in the non-zeropadded case but in the zeropadded case the OR periodogram (detector S1 variant) was able to match the performance of the OR SVA periodogram (detector O1). This was due to similar reasons given in section 3.4, that is, the QIFFT used in the zeropadded OR periodogram (detector S1 variant) was able to mitigate against sidelobe spectral leakage in this moderate-leakage regime. Extraced results for the Barker-13 test signal are shown in Fig 5.4.12.

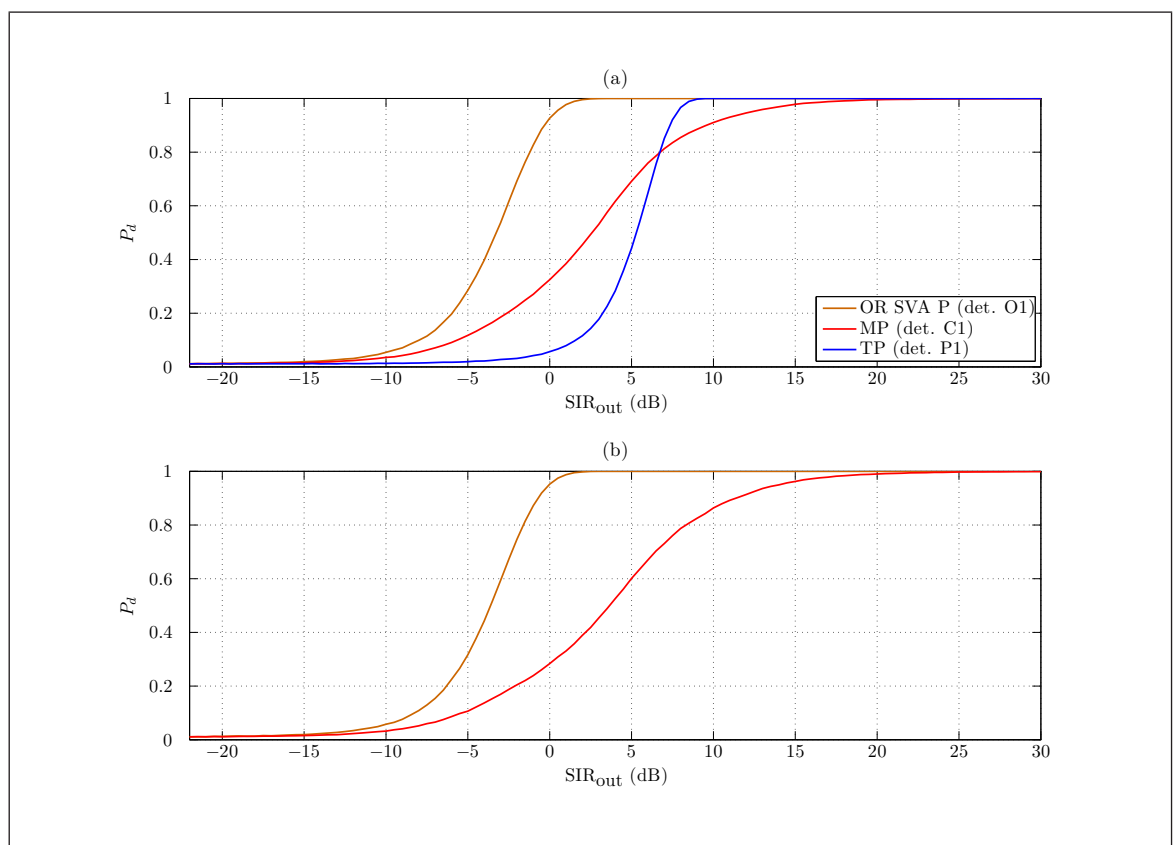


Figure 5.4.12 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 15$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

The main limitation of the OR modified periodogram (detector S1 variant) was the window loss which the QIFFT cannot improve. Therefore the performance of the OR modified periodogram (detector S1 variant) was relatively unchanged from the non-zeropadded to zeropadded case in this regime and for the Barker-13 input.

Fig 5.4.13 shows some extracted conventional detector results for the m-sequence input under this regime. As per the Barker-13 test signal input, the results for the m-sequence input showed the group of OR detectors to be the best out of the conventional

techniques for the detection of these LPI signals.

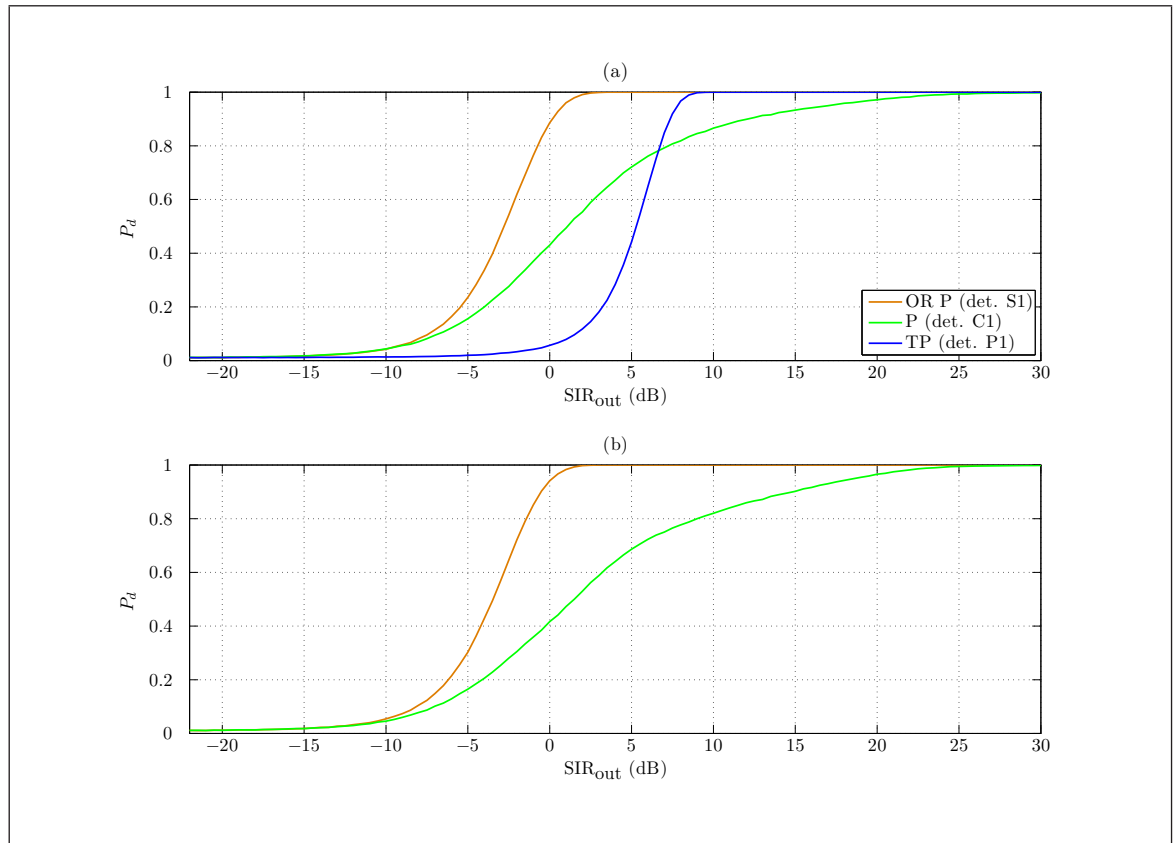


Figure 5.4.13 ROC curves of SIR_{out} against P_d for m-sequence input and conventional detectors at $SNR_{out,B} = 15$ dB for: (a) the non-zero padded case; (b) the zero padded case

The OR SVA periodogram (detector O1) improved its performance relative to the OR periodogram (detector S1 variant) from low- $SNR_{out,B}$ regime to moderate- $SNR_{out,B}$ regime. This was due to the greater importance of sidelobe spectral leakage.

Detection based on the modified periodogram, periodogram and SVA periodogram (detectors C1 and E1) was found to be poorly suited for LPI signal detection. In addition the notion that the best detector based on conventional techniques was poorer than the majority of detectors based on the algorithms of interest was reinforced.

In the moderate- $SNR_{out,B}$ regime the advantage associated with the use of a MTM or a smoothing technique over the best conventional technique ranged from ~ 0.9 dB to ~ 1.3 dB dependent on whether it was zero padded and the LPI signal type.

High $SNR_{out,B}$

For the high spectral leakage regime where the $SNR_{out,B}$ was 25 dB. Extracted results for the best and worst detectors from the algorithms of interest are shown in Fig 5.4.14 for the Barker-13 test signal and in Fig 5.4.15 for the m-sequence test signal. The sine taper MTM (detector H1) gave the best detector performance in the zero padded case (as it obviously did not exist in the non-zero padded case). However this was only by a negligible margin because all the MTMs (detectors G1, H1 and I1) were grouped

together with similar performance levels. Closer inspection of the results revealed the adaptive sine taper MTM (detector I1) had a degraded performance relative to the others in this group from Barker-13 to m-sequence input. This was because the adaptive sine taper MTM was less well suited in the m-sequence case.

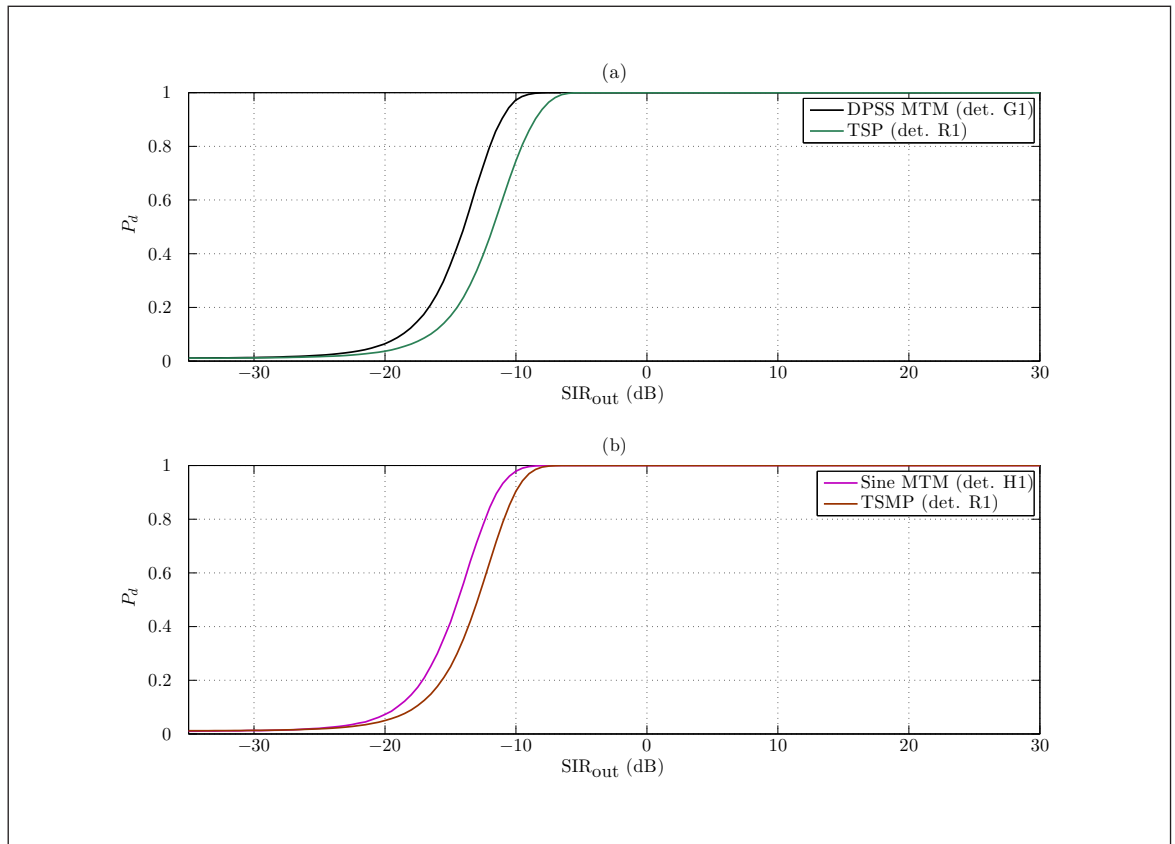


Figure 5.4.14 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $SNR_{out,B} = 25$ dB for: (a) the non-zero padded case; (b) the zero padded case

It was also clear that in the high- $SNR_{out,B}$ regime the SSP and TSSP (detectors K1 and M1) had better detector performances than the SP and TSP (detector Q1 and R1 variants). This was a consequence of the dominance of the sidelobe spectral leakage rather than the noise floor at the monitored ordinates. The SSP (detector K1) was only ~ 0.4 dB poorer in terms of detection performance than the sine taper MTM (detector H1) in the worst case analysed. The TSSP (detector R1) was only slightly degraded from this value. These performance drops could comfortably be tolerated when traded against other engineering advantages which the use of a smoothed SVA periodogram would provide.

The final set of results examined the conventional detection techniques in the high- $SNR_{out,B}$ regime (extracts of which to show the best conventional detector are in Figs 5.4.16 and 5.4.17). The OR SVA periodogram (detector O1) had the best detector performance. The ROC curves of the OR modified periodogram (detector S1 variant) were the next best, followed by the OR periodogram (detector S1 variant). However the OR SVA periodogram (detector O1) was between $\sim [1.5, 1.6]$ dB worse

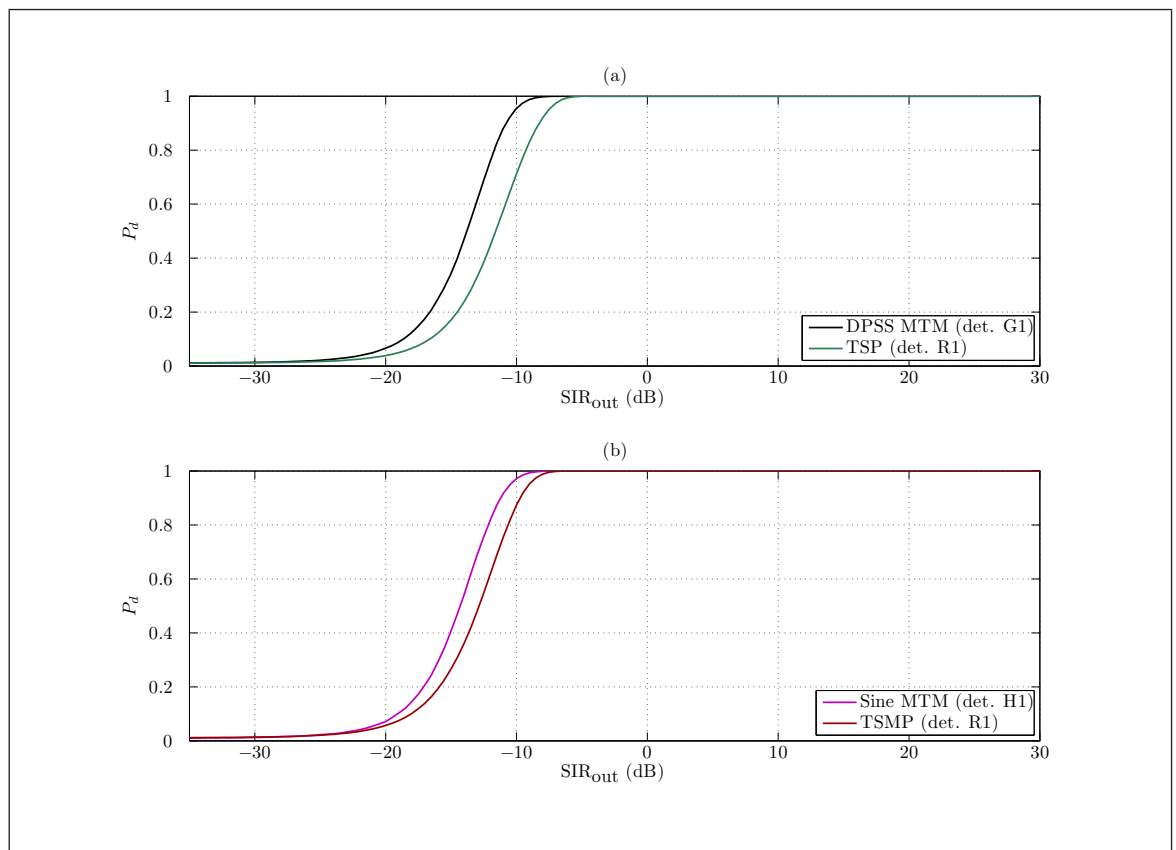


Figure 5.4.15 ROC curves of SIR_{out} against P_d for m-sequence input and LPI detectors at $\text{SNR}_{\text{out},B} = 25$ dB for: (a) the non-zero padded case; (b) the zero padded case

than the DPSS MTM (detector G1) for both the Barker-13 pulse input and the m-sequence input in both the non-zero padded and zero padded cases.

The overall trend from low- $\text{SNR}_{\text{out},B}$ regime to high- $\text{SNR}_{\text{out},B}$ regime was the detector performance of the SSP and TSSP (detectors K1 and M1) to improve at a faster rate than the SP and TSP (detectors Q1 and R1 variants). Solely from the analysis in this chapter it could be concluded that the DPSS MTM and sine taper MTM (detectors G1 and H1) would be the best algorithms to choose due to their superior detection performance. This would not give due consideration to the impact of their implementation. The SSP and TSSP (detectors K1 and M1) are a compromise solution. This is partly due to their reasonable performances in detectors K1 and M1 and partly due to their simpler computational demands than MTMs MTM (detectors G1 and H1).

The limitations of the periodogram, modified periodogram, SVA periodogram and total power (detectors C1, E1 and P1) against LPI signals was highlighted from the results in this scenario. The use of OR detectors offered an improvement over the basic periodogram, modified periodogram and SVA periodogram (detectors C1 and E1). However the OR detectors in general lagged behind the MTMs and smoothing techniques by > 1 dB. This performance gap was found to widen as greater spectral leakage from the interference tone was found in the monitored ordinates.

The next section deals with the scenario when the TOAs of S_A and S_B are random.

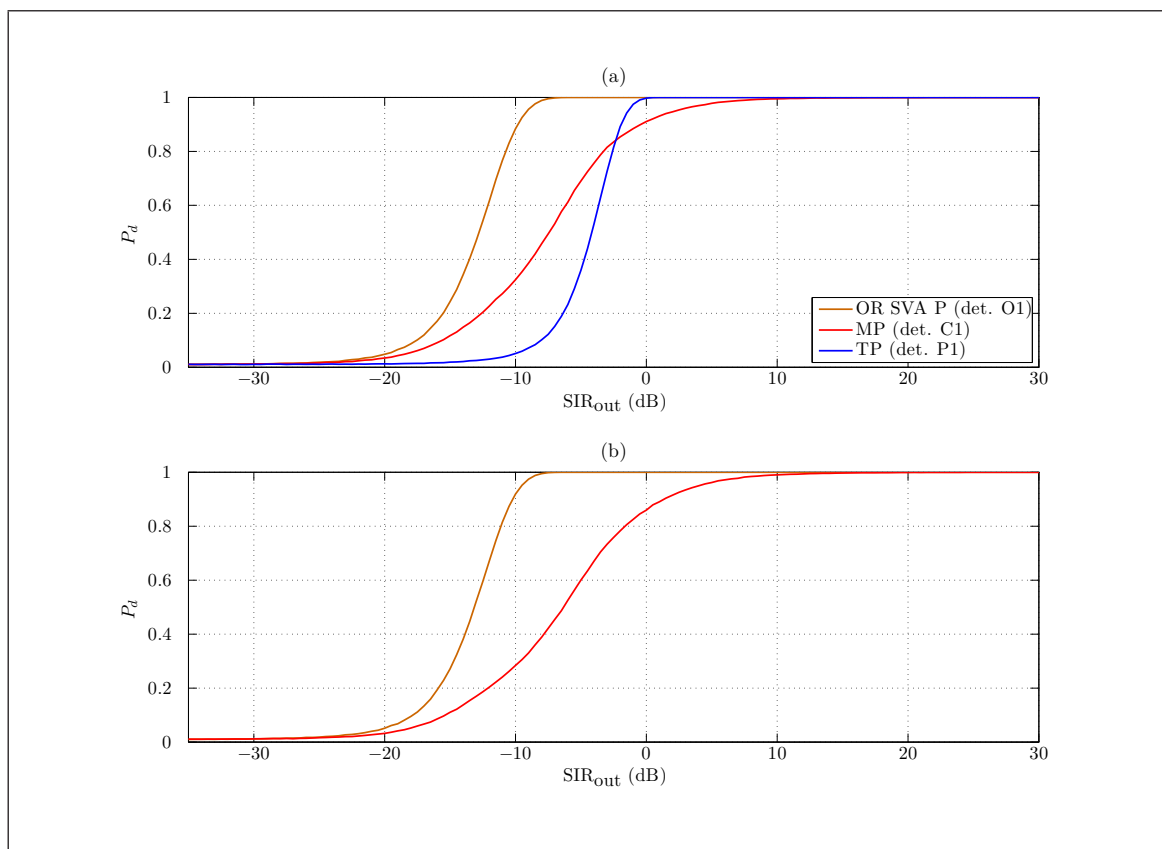


Figure 5.4.16 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 25$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

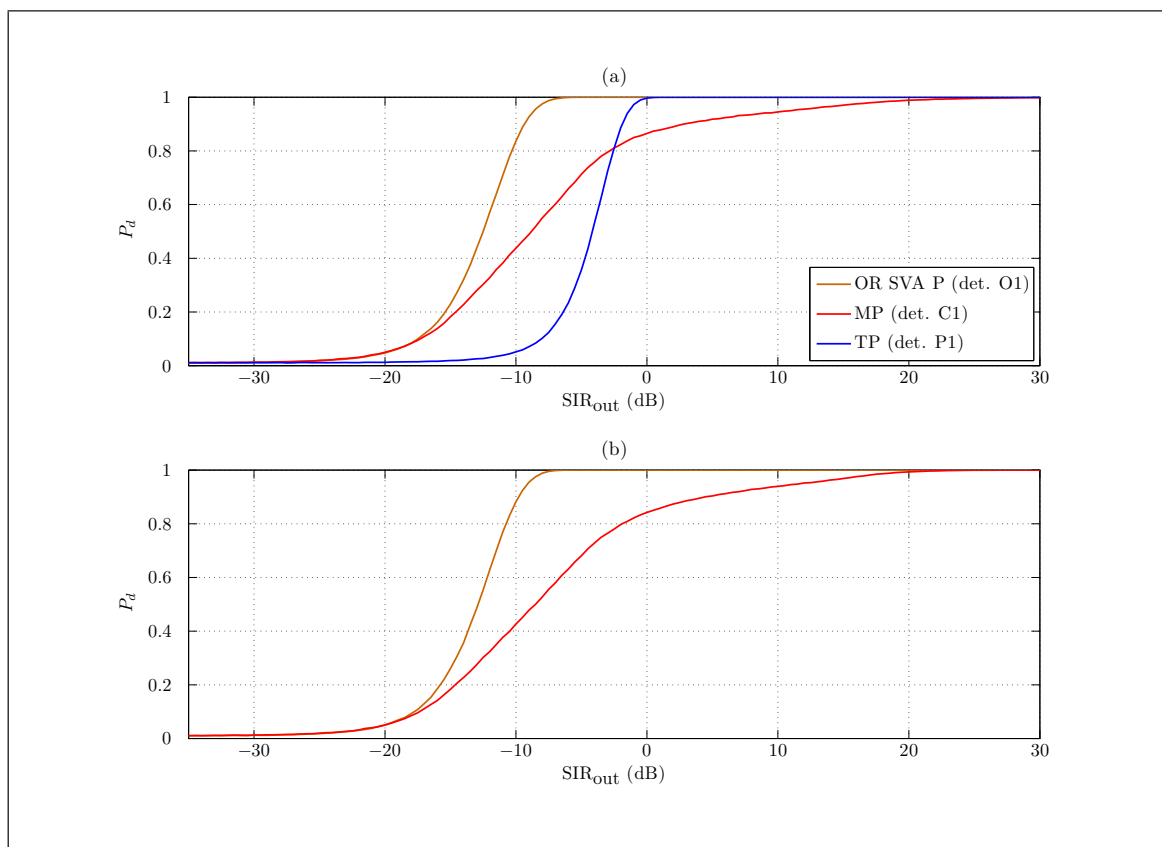


Figure 5.4.17 ROC curves of SIR_{out} against P_d for m-sequence input and conventional detectors at $\text{SNR}_{\text{out},B} = 25$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

5.5 Interference in Randomly-filled DCW

5.5.1 Scenario and Simulation Description

The simulations in section 5.4 revealed the relative detector performances for the algorithms of interest and the conventional detection techniques. The unity-weighted sine taper MTM (detector H1), the SSP (detector K1) and the TSSP (detector M1) performed particularly well over the range of conditions tested. The scenario in this section extended the interference full DCW scenario in the same manner that the full DCW scenario was extended to the partially-filled DCW scenario.

Since the CW m-sequence signal would always be present in a DCW, only the Barker-13 pulse signal was used as the desired LPI signal (S_A). The same mechanism to capture a random proportion of the desired Barker-13 pulse (S_A) in a DCW as in Fig 5.3.1 in section 5.3 was used in this scenario. Whereas the mechanism to capture a random proportion of the interference tone (S_B) was the same as was used in section 3.3 and illustrated in Fig 3.3.1.

In this scenario the detectors which incorporated steps to estimate sidelobe spectral leakage and mainlobe width were tested to their limits. The estimates which these detectors created became less and less accurate as the proportion of S_B in a DCW decreased. The explanation for this is given in section 3.5 and exhibited in Fig 3.5.1. Therefore this scenario was expected to reveal the limitations of detectors Q1, R1, S1 and C1.

All other aspects of the interference full DCW scenario were maintained in the interference in randomly-filled DCW scenario.

5.5.2 Results and Discussion

Low $\text{SNR}_{\text{out},B}$

The first set of results were created with:

- a Barker-13 pulse input;
- $\text{SNR}_{\text{out},B} = 8 \text{ dB}$;
- utilisation of the algorithms of interest;
- nominal target frequency separation of $5f_s/N$.

From the results, the detector performances in descending order were:

1. DPSS MTM (detector G1), unity-weighted sine taper MTM (detector H1), SP (detector Q1 variant) and TSP (detector R1 variant)

2. SSP (detector K1) and TSSP (detector M1)
3. adaptive sine taper MTM (detector I1)
4. SMP (detector Q1 variant) and TSMP (detector R1 variant).

The best and worst detectors from the algorithms of interest are shown in Fig 5.5.1.

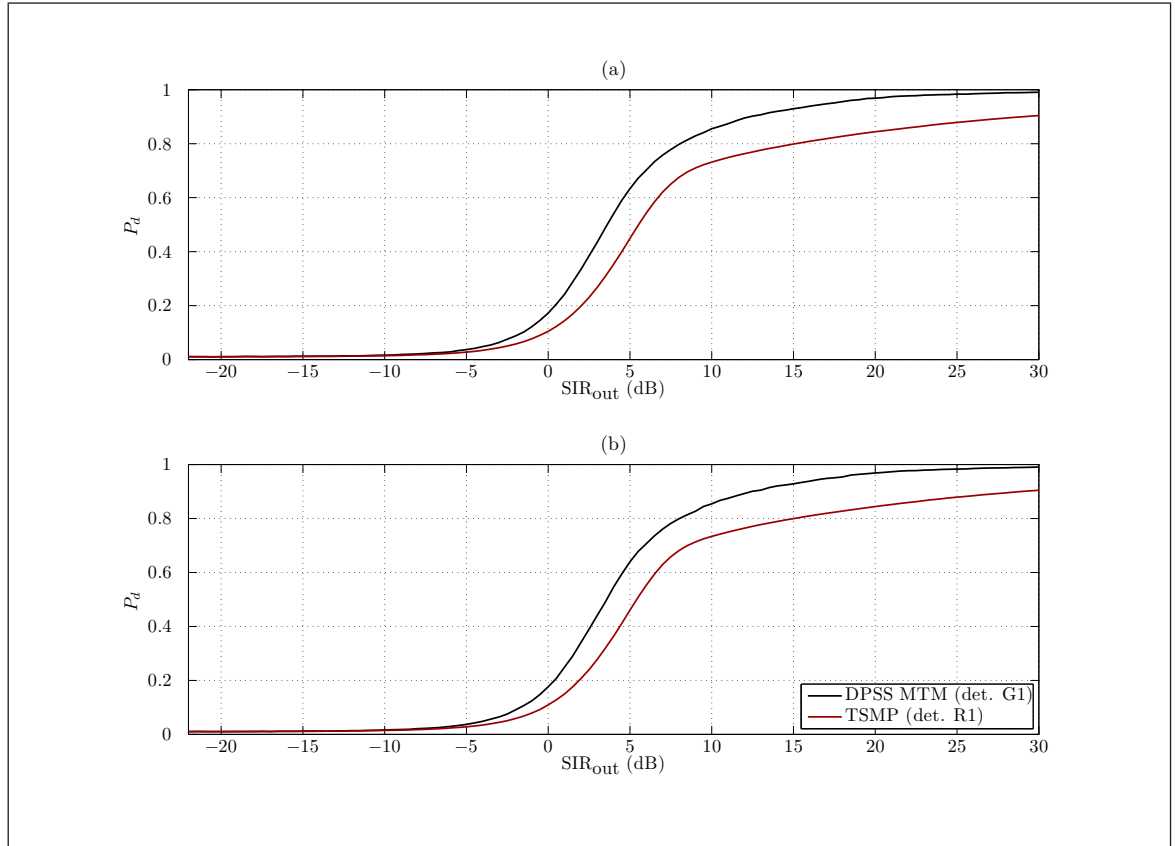


Figure 5.5.1 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $\text{SNR}_{\text{out},B} = 8$ dB for: (a) the non-zero padded case; (b) the zero padded case

The ROC curves of all the detectors were distorted. This was a consequence of the capture of S_A and S_B with random TOAs. The proportion of S_A power captured in a DCW could be as low as one sample whereas the proportion of S_B captured could go as low as half of the samples in a DCW. Therefore the ROC curves were distorted such that they approached $P_d = 1$ more slowly than in the interference full DCW scenario.

The adaptive sine taper MTM (detector I1) was poorly suited to this particular scenario in comparison to the interference full DCW scenario from section 5.4. However its performance surpassed the SVA-based detectors K1 and M1 at higher SIR_{out} ratios. This was because in cases of partially-filled DCWs, the SVA technique tended to select bell-shaped window functions to apply to the DCWs which thus attenuates the power of the desired signal (S_A). The same effect was seen through the use of the SMP because in this detector variant it applied a Hann window. Therefore the SVA-based detectors K1 and M1 and the SMP (detector Q1 variant) and TSMP (detector R1 variant) asymptotically approached P_d less rapidly than the other detectors.

The results demonstrated that the relative detector performances were the same regardless of whether the FFT was zeropadded. The results also closely resembled those obtained for the partially-filled DCW scenario, which suggested that spectral leakage was fairly insignificant in this scenario and $\text{SNR}_{\text{out},B}$ regime.

The results for the conventional technique detectors bore a resemblance to the results from the partially-filled DCW scenario. That is, the order of detector performance in from best to worst was:

1. OR periodogram (detector S1 variant) and OR SVA periodogram (detector O1)
2. OR modified periodogram (detector S1 variant)
3. total power (detector P1)
4. periodogram (detector C1 variant)
5. modified periodogram (detector C1 variant)
6. SVA periodogram (detector E1).

The relative detector performances were the same regardless of zeropadding. Extracted results to demonstrate the best conventional detector are shown in Fig 5.5.2.

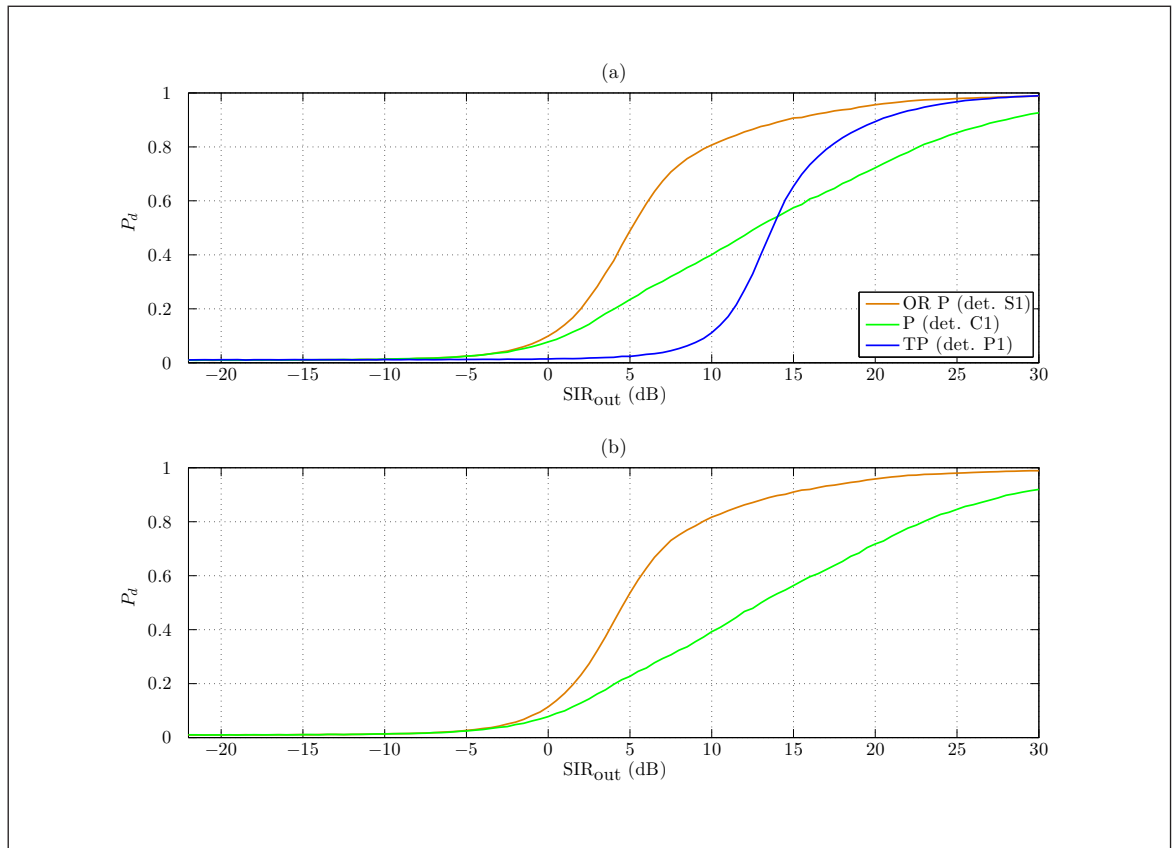


Figure 5.5.2 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 8$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

The OR periodogram (detector S1 variant) and OR SVA periodogram (detector O1) are grouped together in the above list because their detection performances were almost

identical. However at high SIR_{out} levels, the OR SVA periodogram (detector O1) asymptotically approached $P_d = 1$ at a slower rate. This was due to the previously described effect where SVA attenuates the signal power in a DCW where the signal is mostly concentrated at the DCW band edges.

The detector performance gap between the grouping of the OR periodogram/OR SVA periodogram (detectors S1 variant/O1) and the OR modified periodogram (detector S1 variant) was ~ 0.6 dB in the non-zeropadded case and ~ 0.75 dB in the zeropadded case. This was typical for the low- $\text{SNR}_{\text{out},B}$ regime where the modified periodogram-based detectors (which used a Hann window function) suffered a window loss. Additionally the OR modified periodogram (detector S1 variant) displayed the same pattern as the OR SVA periodogram (detector O1) whereby its performance approached $P_d = 1$ at a slow rate.

The total power (detector P1) was heavily handicapped by the inclusion in its summation of ordinates that contained noise power only. Also the periodogram, modified periodogram and SVA periodogram (detectors C1 and E1) were again found to be ill-suited to the detection of LPI signals.

In this low- $\text{SNR}_{\text{out},B}$ regime and with a random proportion of S_A in the DCW, the DPSS MTM (detector G1) had the best performance and was better than the OR periodogram (detector S1 variant), which had the best performance from the conventional techniques, by ~ 1.6 dB in the non-zeropadded case and ~ 1 dB in the zeropadded case. Therefore in this scenario and regime there are strong arguments for the use of a MTM or smoothing technique over the conventional techniques.

Moderate $\text{SNR}_{\text{out},B}$

In the moderate- $\text{SNR}_{\text{out},B} = 15$ dB regime the relative detector performances changed very little from the low- $\text{SNR}_{\text{out},B} = 8$ dB regime. The best and worst detectors for the algorithms of interest are shown in Fig 5.5.3.

There was a fairly large performance gap between the TSSP (detector M1) and the SMP (detector Q1 variant) of ~ 1.3 dB in the zeropadded case and ~ 1.4 dB in the non-zeropadded case. This was caused by the fixed Hann window applied to every DCW which attenuated captured signal power when it was concentrated at the DCW edges. Meanwhile the SSP and TSSP (detectors K1 and M1) results were better because they had the ability to be somewhat more flexible in their choice of window function and did not always attenuate DCW edges in all cases.

Again the adaptive sine taper MTM (detector I1) did not perform as well as the other MTMs because of the poor match between its adaptive parameters and this particular regime and scenario. However the attenuation of S_A power at the DCW edges caused the SVA and modified periodogram-based detectors to be the worst at most SIR_{out} points. The results for the SVA and modified periodogram-based detectors approached

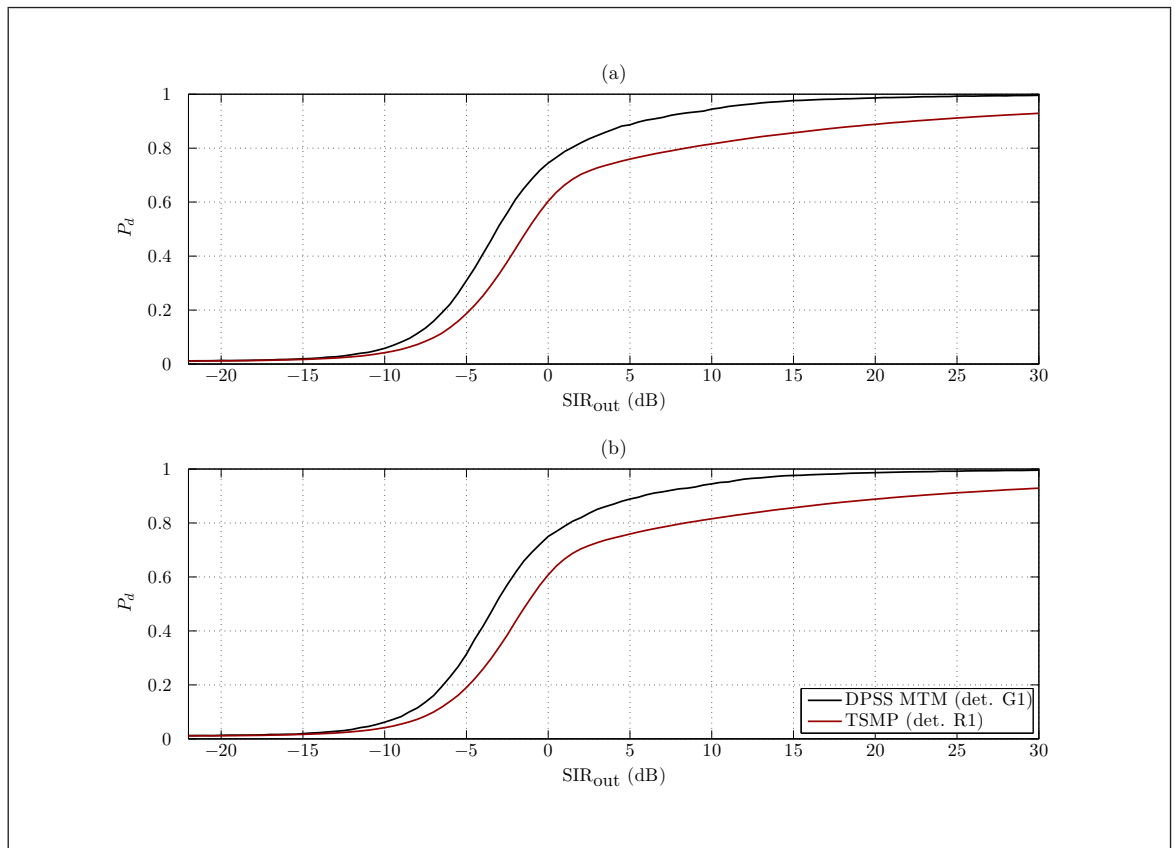


Figure 5.5.3 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $\text{SNR}_{\text{out},B} = 15$ dB for: (a) the non-zero padded case; (b) the zero padded case

$P_d = 1$ very slowly.

The relative performance results for the conventional detectors for the moderate- $\text{SNR}_{\text{out},B}$ regime were almost identical to those from the low- $\text{SNR}_{\text{out},B}$ regime. One minor difference was that the performance gap between the OR SVA periodogram (detector O1) and the SMP (detector S1 variant) expanded to ~ 0.75 dB in the non-zero padded case and ~ 0.8 dB in the zero padded case in the moderate- $\text{SNR}_{\text{out},B}$ regime. Extracted ROC curves for the conventional detectors are shown in Fig 5.5.4.

The point that the algorithms of interest are superior for the detection of LPI waveforms was reinforced by the DPSS MTM (detector G1) which exhibited a performance that bettered the OR periodogram (detector S1 variant) by ~ 1.5 dB in the non-zero padded case) and ~ 1.25 dB in the zero padded case.

Since the results from the low- $\text{SNR}_{\text{out},B}$ regime and moderate- $\text{SNR}_{\text{out},B}$ regime were almost identical in this scenario, it is reasonable to conclude that between these levels the loss of signal samples from the DCW outweighed any effects due to spectral leakage.

High $\text{SNR}_{\text{out},B}$

For the algorithms of interest and in the high- $\text{SNR}_{\text{out},B} = 25$ dB regime there were many similarities with the relative detector performances in the low and moderate regimes. There were however two details of interest in the results. One concerns the fact that

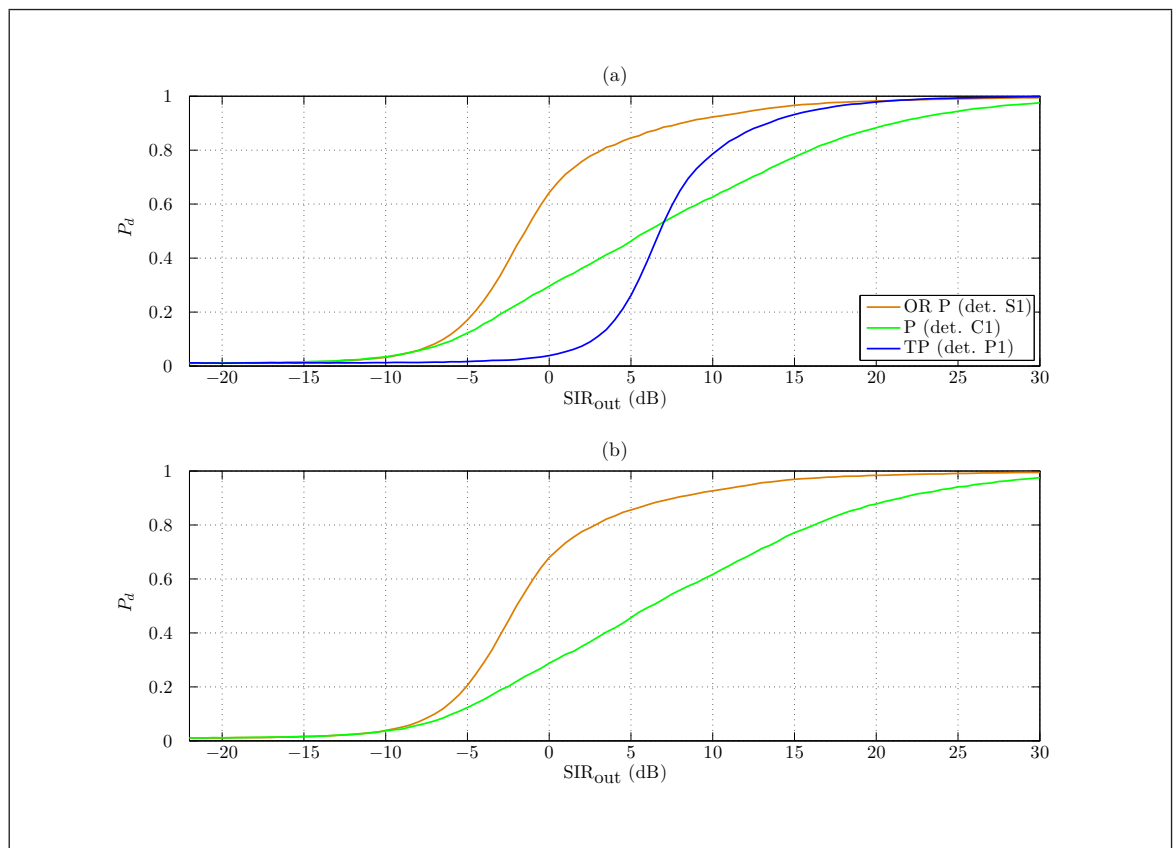


Figure 5.5.4 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 15$ dB for: (a) the non-zero padded case; (b) the zero padded case

adaptive sine taper MTM (detector I1) had a performance imbetween the sine taper MTM (detector H1) and the SP (detector Q1 variant). This differed from the other regimes and is explained by the better match between the set choice of parameters for the adaptive algorithm and this scenario and regime.

The second detail relates to the larger performance gap between the grouping of the SSP and TSSP (detectors K1 and M1) and the grouping of the SMP and TSMP (detector Q1 and R1 variants) than in the previous regimes. This was because the modified periodogram-based detectors could not catch up with the other detectors due to the constant attenuation of the DCW edges where S_A power is likely to be concentrated. The problem was exasperated in comparison to the results from the section 3.5 by the strict conditions in this scenario of the minimum proportion of samples which contained signal power that could be captured (this was as low as one sample). The best and worst detectors from the algorithms of interest are shown in Fig 5.5.5.

The conventional detectors were also tested in the high- $\text{SNR}_{\text{out},B}$ regime and some results are displayed in Fig 5.5.6. There were some differences in comparison to the two previous $\text{SNR}_{\text{out},B}$ regimes. In the non-zero padded case the OR SVA periodogram (detector O1) was actually the best detector with the caveat that at higher $\text{SIR}_{\text{out},B}$ levels its approach to $P_d = 1$ slowed relative to the OR periodogram (detector S1 variant). This indicated that in this high- $\text{SNR}_{\text{out},B}$ regime the effect of sidelobe spectral

leakage in monitored frequency bins was as important as the random TOAs of S_A and S_B .

Another difference was that the performance gap between the OR modified periodogram (detector S1 variant) and the other OR detectors expanded to a large difference of ~ 3.9 dB in the non-zeropadded case and ~ 2.8 dB in the zeropadded case in comparison to the lower $\text{SNR}_{\text{out},B}$ regimes. This highlights the potential detector processing gains to be achieved through the avoidance of the conservative approach of a modified periodogram as the basis of a DRx detection circuit against LPI waveforms.

In a similar situation to the other regimes, the high- $\text{SNR}_{\text{out},B}$ regime results revealed that the best algorithm of interest detector provided a better detector performance than the best conventional detector. This difference was ~ 2.6 dB in the non-zeropadded case and ~ 2.2 dB in the zeropadded case.

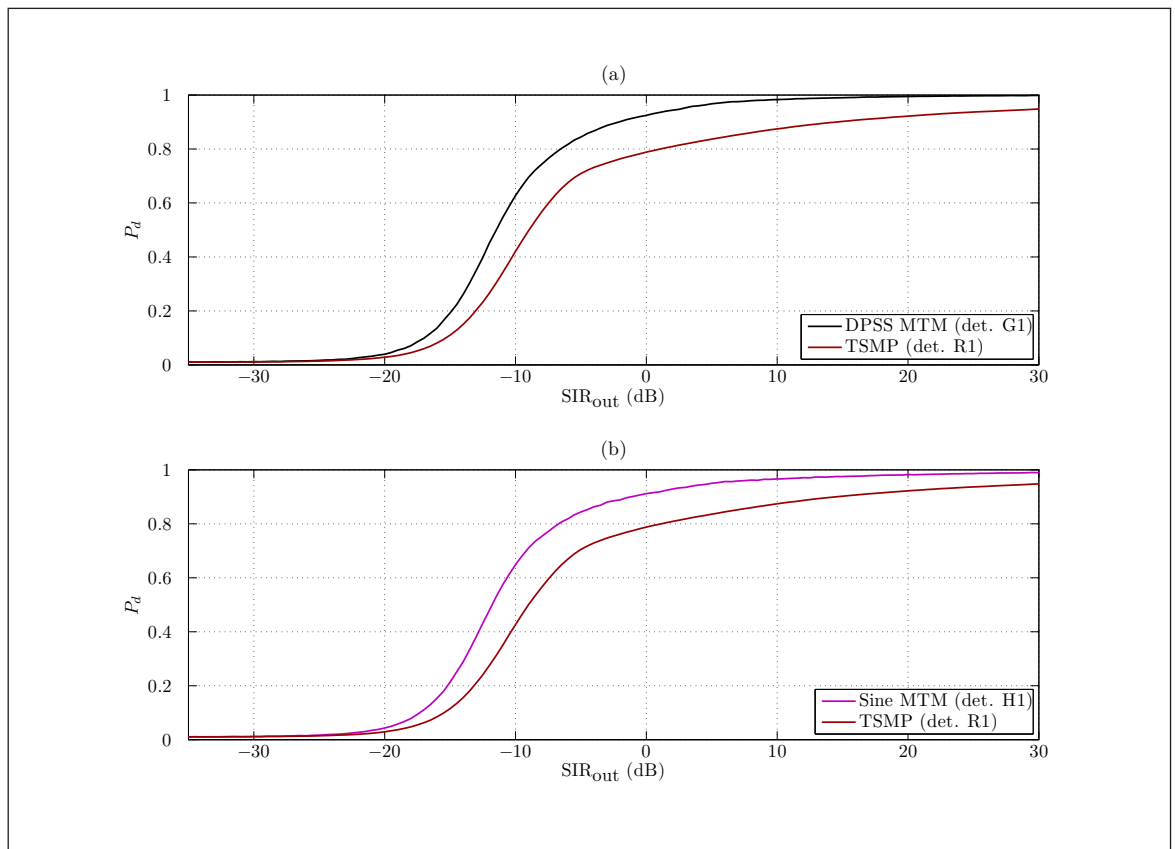


Figure 5.5.5 ROC curves of SIR_{out} against P_d for Barker-13 input and LPI detectors at $\text{SNR}_{\text{out},B} = 25$ dB for: (a) the non-zeropadded case; (b) the zeropadded case

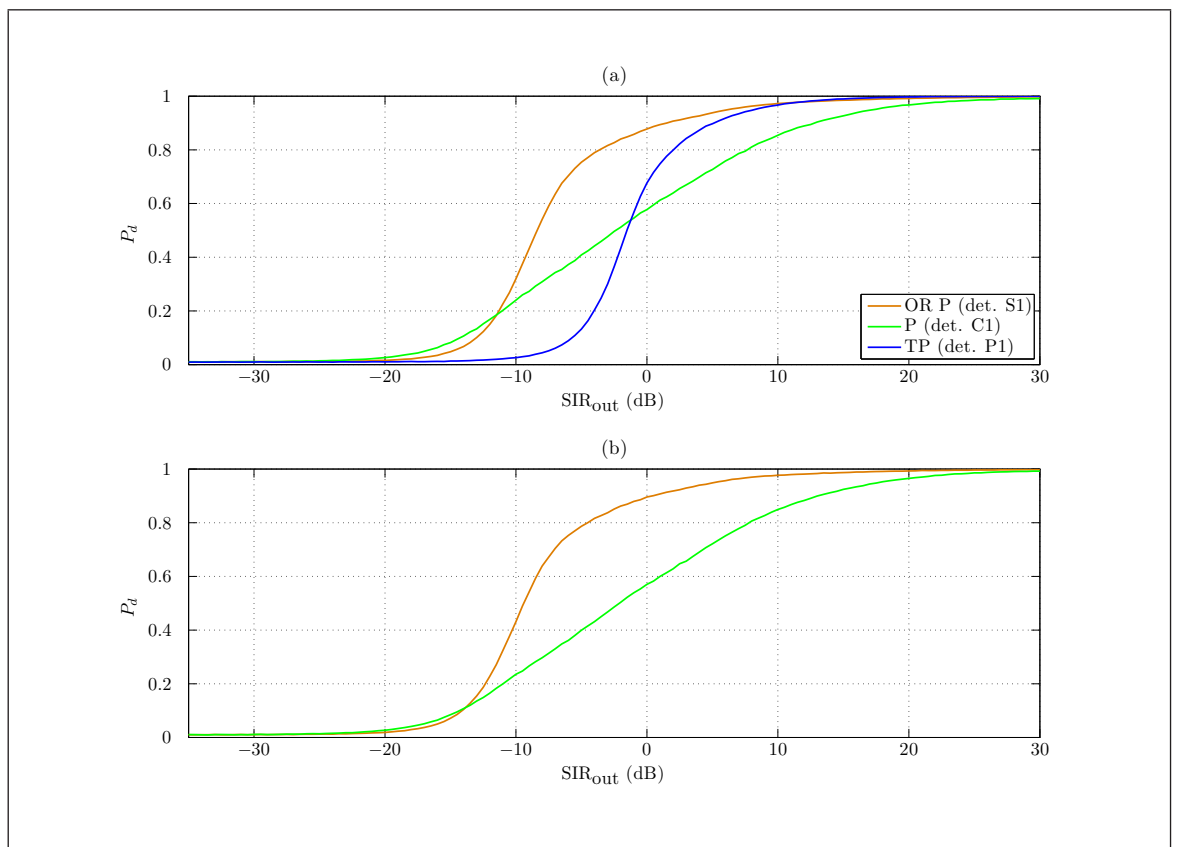


Figure 5.5.6 ROC curves of SIR_{out} against P_d for Barker-13 input and conventional detectors at $\text{SNR}_{\text{out},B} = 25 \text{ dB}$ for: (a) the non-zeropadded case; (b) the zeropadded case

5.6 Conclusions

The aim of this chapter was to investigate whether detection algorithms which tested for LPI signals spread across a broad frequency band would offer an improvement over conventional techniques.

One general finding from this chapter was that the use of the DPSS MTM (detector G1) and the unity-weighted sine taper MTM (detector H1) provided an improvement over all the conventional techniques tested in the examined scenarios. In addition, most of the detectors which utilised smoothing techniques also surpassed the performance of the conventional techniques. In that regard, this investigation was able to show the MTMs and smoothing techniques were improvements over conventional DRx techniques. Out of the smoothing techniques, the SSP and TSSP (detectors K1 and M1) were relatively consistent across all the scenarios tested.

Although a smoothed SVA detector, whether SSP or TSSP (detectors K1 or M1), would not necessarily offer the best detection performance in all scenarios, as an implementable algorithm in a real-life DRx it would certainly be more feasible than the MTMs because the MTM would require multiple FFT core instantiations. Detector G1 which utilised the DPSS MTM would provide the best performance in most scenarios but it would require a great deal of computational power as evidenced from Table 4.3.1. Detector H1 based on the unity-weighted sine taper MTM would be more feasible than the DPSS MTM (detector G1) because it would require only one FFT block. However from Table 4.3.1 it would still be more computationally complex than a smoothing technique. Another important consideration was that the SVA technique was shown in chapter 3 to be a successful method for the detection of sinusoidal signals. Therefore if the recommendation from chapter 3 was followed and an SVA periodogram was formed for the detection of pulsed sinusoidal radar signals in an ESM DRx, then it could be reused in a SSP or TSSP in detectors K1 or M1 for the detection of broader-band LPI radar signals.

Since the performances of the SSP and TSSP (detectors K1 and M1) had almost negligible differences between them in the scenarios tested, the computational savings provided by the TSSP (detector M1) would outweigh any slight degradation in performance. Whilst the MTMs are of more academic interest and possible future implementation, the use of the TSSP (detector M1) was chosen for further investigation due to engineering considerations. The appendices H–I detail the translation of the TSSP (detector M1) down the MBD methodology.

5.7 Summary Table

The results for the various scenarios examined are summarised in Table 5.7.1.

Scenario	Best Detector(s)	Description
Full DCW, no interference	Sine taper MTMs (detectors H1 and I1), SP (detector J1 variant), TSP (detector L1 variant), DPSS MTM (detector G1), SSP (detector K1), TSSP (detector M1)	Techniques which collect together the most signal power
Randomly-filled DCW, no interference	As above	Detector performances were degraded by random loss of signal power but the best techniques still gathered as much signal power as possible from bandwidth
Full DCW, interference present	DPSS MTM (detector G1), unity-weighted sine taper MTM (detector H1), SSP (detector K1), TSSP (detector M1)	The DPSS MTM, unity-weighted sine taper MTM and SSP techniques minimise impact of spectral leakage and result in best detector performances over a range of interference regimes
Randomly-filled DCW, interference present	DPSS MTM (detector G1), unity-weighted sine taper MTM (detector H1), SP (detector Q1 variant), TSP (detector R1 variant), OR periodogram (detector S1 variant), SSP (detector K1), TSSP (detector M1)	Detector performances degraded in comparison to the above. The loss more noticeable in the SSP techniques than the other “best” detectors

Table 5.7.1 Floating-point simulations results summary

CHAPTER 6

Thesis Conclusions & Future Work

6.1 Conclusions

For real-time applications of detector algorithms in ESM systems the importance of computational efficiency is paramount and algorithms tend to be based around the FFT.

The approach taken in this thesis was to enhance the ability of a typical FFT-based DRx as a constituent part of an ESM system to detect both sinusoidal radar signals and broader-band LPI radar signals. This was achieved through the use of SVA against sinusoidal, narrowband signals and the MTMs and spectral smoothing techniques against broader-band LPI signals.

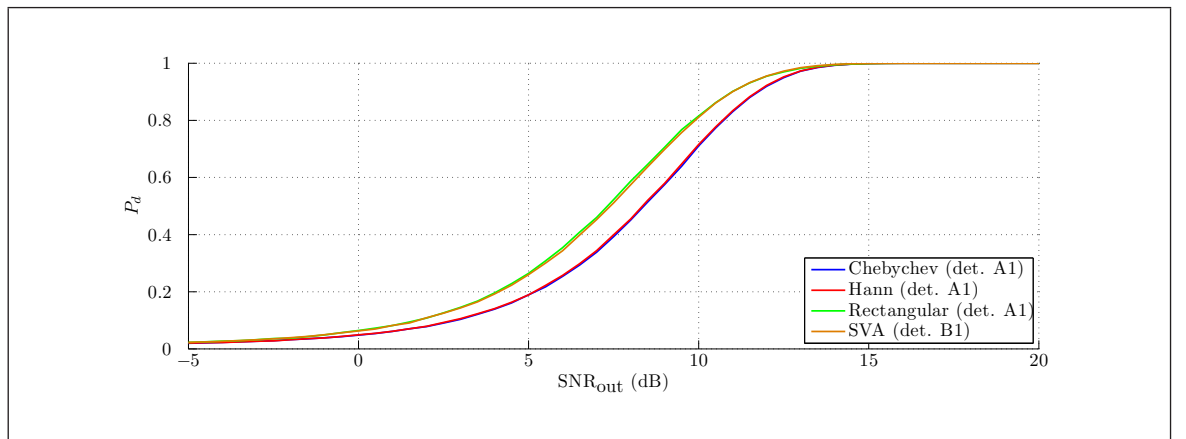


Figure 6.1.1 Advantages of the use of SVA for full DCW scenario

The SVA technique is a computationally efficient, FFT-based algorithm. In “quiet” scenarios, where only one signal is present in a bandwidth of interest, the SVA technique overcomes the window loss of conventional fixed window analysis suffered in modified periodograms. Although a similar performance can be achieved with the use of the rectangular window, this is seldom used in practice due to its consequential spectral leakage. SVA can minimise spectral leakage **and** overcome the window loss. Extracted non-zero padded floating-point results in Fig 6.1.1 demonstrate this effect where SVA recovered the ~ 1 dB window loss suffered by a detector using a bell-shaped window

function. This would translate to an estimated real-life detector range improvement of ~ 17.1 km for the SELEX Galileo DRx and extended low band (ELB) antenna system against a typical unmodulated pulsed radar

Against narrowband sinusoidal signals, SVA demonstrably affords a detection advantage over conventional fixed window techniques in scenarios with strong narrowband interferers which are close in frequency. Extracted non-zero padded FFT floating-point results (separation $2f_s/N_{FFT}$ and $\text{SNR}_{\text{out},B} = 15$ dB) in Fig 6.1.2 demonstrate this effect where SVA improves the detector performance by ~ 0.8 dB. This would lead to an estimated real-life detection range improvement of ~ 14.8 km for the SELEX Galileo DRx and ELB antenna system (which does not currently able to use a zero padded FFT) against a typical unmodulated pulsed radar.

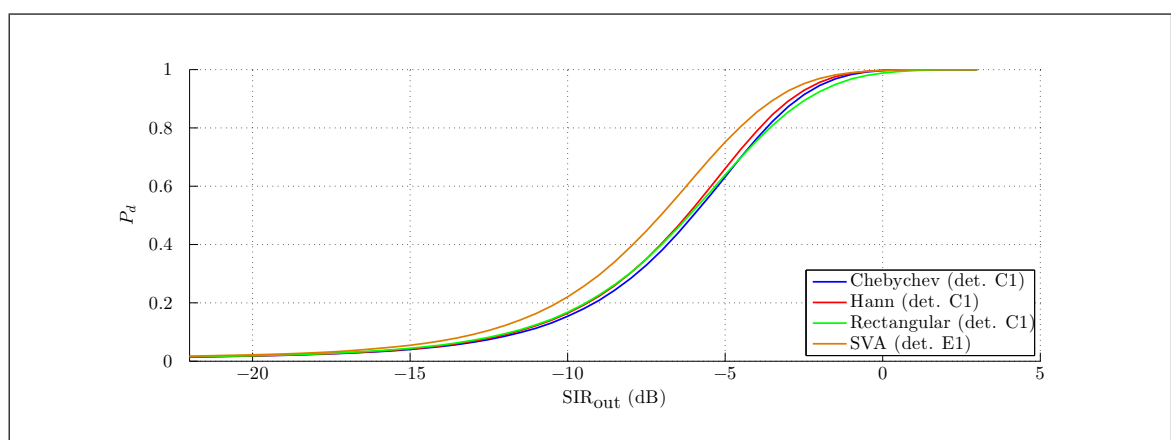


Figure 6.1.2 Advantages of the use of SVA for full interference DCW scenario

As members of non-parametric spectral estimation, the MTMs and spectral smoothing techniques are relatively computationally efficient. They are appropriate for the detection of broader-band LPI radar signals in FFT-based detectors where simple (modified) periodograms would over-resolve the signal. In contrast to this the MTMs and smoothing techniques can be used to test for the presence of broader-band signals at a few different hypothesis bandwidths.

The results showed the MTMs and all smoothing techniques were a great improvement in detector performance over conventional ESM DRx techniques.

For the scenario of detection of a single Barker-13 pulse in a full DCW, extracted floating-point results show the best MTM/smoothing technique (sine taper MTM, detector H1) outperformed the best conventional technique (OR periodogram, detector N1) by ~ 1.5 dB which would lead to an estimated real-life detection range improvement of ~ 12.5 km for for the SELEX Galileo DRx and ELB antenna system against a typical pulsed radar.

For the scenario of detection of a Barker-13 pulse in the presence of a close-in frequency interference tone, extracted results in Fig 6.1.4 ($\text{SNR}_{\text{out},B} = 15$ dB) show the

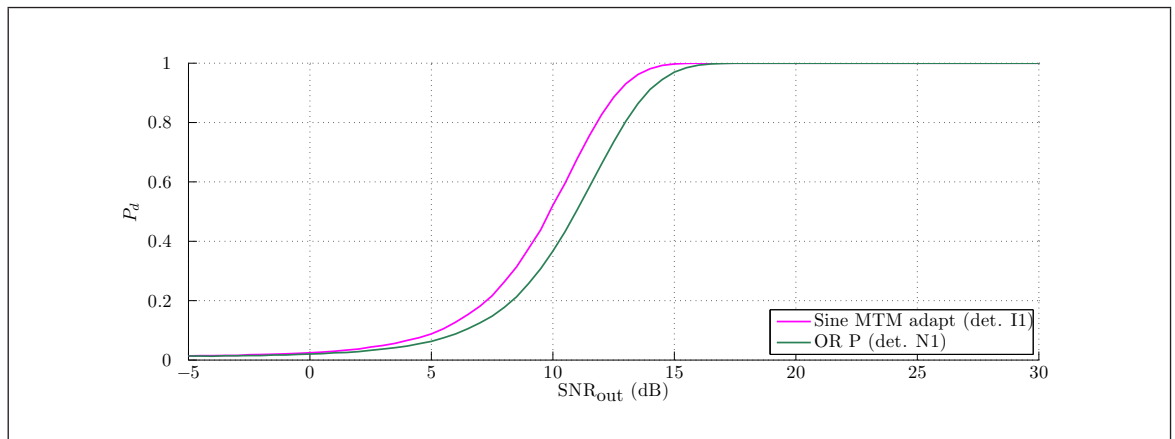


Figure 6.1.3 Advantages of the use of the MTMs/smoothing techniques for full DCW scenario

best MTM/smoothing technique (DPSS MTM, detector G1) outperformed the best conventional technique (OR SVA periodogram, detector O1) by ~ 1.3 dB which would lead to an estimated real-life detection range improvement of ~ 11.4 km for the SELEX Galileo DRx and ELB antenna system against a typical pulsed radar.

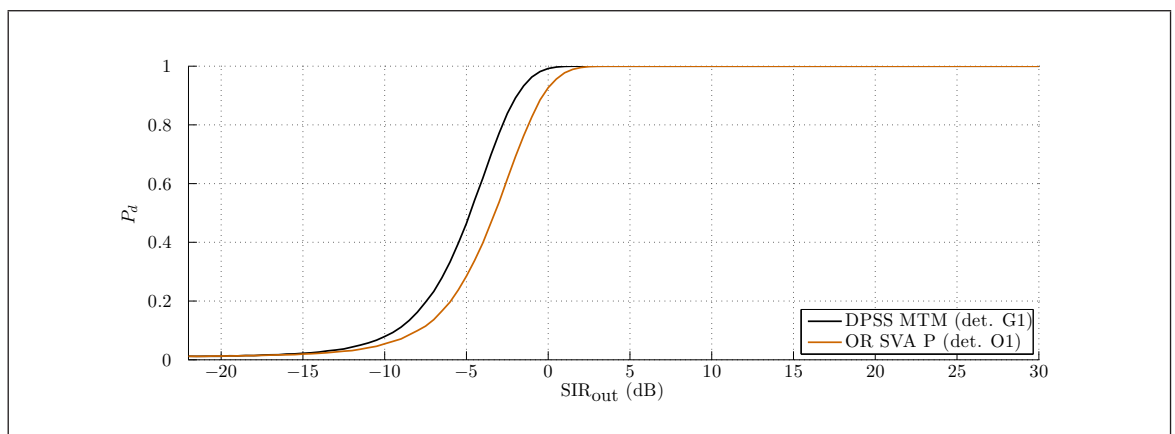


Figure 6.1.4 Advantages of the use of the MTMs/smoothing techniques for interference full DCW scenario

A similar pattern of results was achieved for the MTMs and smoothing techniques in pursuit of broader-band LPI radar signals as was observed for SVA and conventional window techniques in pursuit of narrowband, sinusoidal radar signals. That is, the MTMs (which as a by-product overcome the window loss) and smoothing techniques, which used the periodogram or SVA periodogram as their basis, excelled in the scenarios where only one broader-band signal in a bandwidth of interest had to be considered. This was due to the maximisation of the amount of signal power processed in these algorithms, i.e. there was no or virtually no window loss. Also as a strong interference tone was introduced to a frequency close to the desired signal, the MTMs and SVA periodogram smoothing algorithms surpassed the performance of the other smoothed fixed window detectors owing to their better spectral leakage protection.

The application of SVA should be employed in the pursuit of narrowband, sinusoidal

radar signals and a MTM or smoothing technique should be employed against broader-band LPI radar signals. The next section for future work describes how these algorithms may be combined as part of an all-round detector.

The appendices describe the progress which has already been made in taking the SVA and the TSSP (one of the smoothing techniques) through the MBD process to integrate these algorithms into an ESM system.

6.2 Future Work

The problem space for successful conventional and LPI radar waveform detection and classification is huge. The contributions offered in this thesis have simply dealt with the question of initial detection to gain or recover some range advantage. In doing so many algorithms (such as those based on signal cyclostationarity or signal statistical properties) were by-passed. Decisions were taken in the course of research to by-pass these algorithms because they are currently too computationally complex to run in real-time on data supplied from a superhet and DRx combination at a rate within [1.0, 1.75] GHz or they did not offer an initial detection performance improvement over energy detection in AWGN.

The algorithms investigated in this thesis concentrated on those based on the FFT which would be viable for real-time operation. However the problem of classification remains and here is where more computationally intensive algorithms may find a more suitable role as they may be sent small subsets of data from clusters. These clusters may derive from data flagged as possible detections by the algorithms studied in this thesis. Therefore future work could concentrate on the creation of a system with a rapid FFT-based core for initial detection and slower algorithms processing occasional interesting frames of data to assist with detection confirmation and classification. This system would require at the very least an extended PDW.

The SVA and broader-band LPI waveform detection algorithms described in this thesis are complementary. However some work is required to research a system which combines the two in one system. Fig 6.2.1 shows two options for an all-round detector: (a) with the use of the TSSP described in appendices H and I and (b) with the use of the sine taper MTM. Such a system would have the advantage of conventional radar waveform and LPI waveform detection but some work would be required for the control logic and decision circuitry to make it successful.

The immediate further work is to continue the research into the separate application of SVA and the TSSP algorithms into the SELEX Galileo DRx. This would entail: (1) detection tests at a lower P_{fa} ; (2) implementation of the suggested design improvements from sections F.5 and I.4 and (3) further computational improvements to the SVA implementation. The first point is important because the P_{fa} used in the simulations in this thesis was relatively high for a real application. It was selected at this

level so that the simulations would achieve results within project timescales and to prove the concept.

The second point is explained in the relevant appendices. The third point follows on from a technical review by other SELEX Galileo engineers of the SVA work covered in this thesis. The avenues of exploration for further computational improvements in the implementation of SVA are to observe the effect of a reduction of the input/output (IO) number of bits and the effect of replacing the divide operator with a set of discrete α parameters (more akin to the multiapodization interpretation).

Work packages have been drawn up to tackle the immediate separate application of SVA.

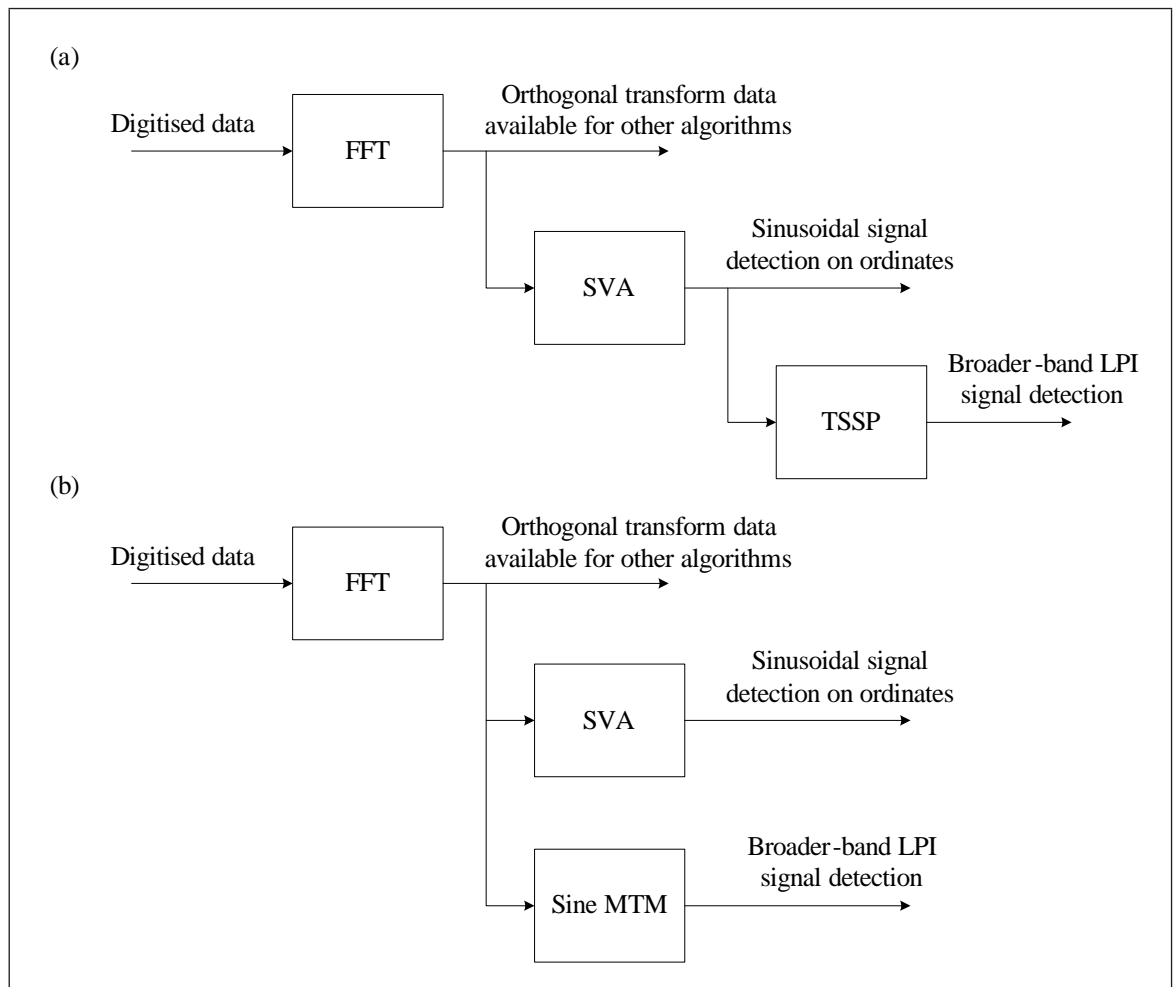


Figure 6.2.1 Proposed integrated ESM systems: (a) with the use of TSSP; (b) with the use of sine taper MTM

Appendices

APPENDIX A

Helicopter Trials

A.1 General Information

The floating-point simulation results obtained with the use of SVA in section 3 were promising. These Monte Carlo simulation results were rigorous in that they produced ROC curves for different simulated scenarios. However it should be noted that both the data input and the simulated detectors in these results were from high-level, floating-point software models. This is not a negative point because algorithms developed using a MBD methodology always begin at this level. However it was desirable to test these software models against some real-life signal data captured using a DRx in an ESM system.



The “helicopter trials” lasted from 18th August to 20th August 2009 and provided a unique opportunity to subject the SVA algorithm to further scrutiny with the use of real-life data captured with the use of a DRx from the field. From the SELEX Galileo point of view the opportunity also demonstrated what could be achieved with their DRx design.

Originally the plan was to attempt to detect a Captor radar in keeping with the “tactical situation” described in section 1.2. However this was not possible due to staff shortages during the trials. Therefore a backup plan was used to fly over the Firth of Forth. Whilst not exactly the desired tactical situation, this provided a wealth of shipping radar emitter activity against which the SVA algorithm could be tested.

A.2 Constraints

There were some engineering constraints which applied at the time of the experiment:

1. The DRx hardware could not switch between a “sample capture mode” and a threat display mode whilst in-flight.
2. The specific superhet frequency band to which the DRx was tuned could not be changed in-flight (i.e. the ESM dwell).
3. The global positioning system (GPS) location system in the DRx was offline.

Working within these constraints, the DRx was used in sample capture mode to obtain buffered ADC data which then provided an input to a laptop taken on-board the helicopter to execute MATLAB[®] functions. These functions performed spectrogram techniques on the data, with and without the application of SVA. Before each flight, the DRx was tuned to a dwell which covered the shipping navigation range of frequencies.

A.3 Aims

Predictions of which emitters would be active during the flight was too difficult, even with a partial list of known emitters, because the time of arrival of the signals at the helicopter's antennas would be unknown. In this kind of situation, when faced with the analysis of an unknown spectrum, the standard engineering practice is to apply a bell-shaped window function (such as a Hann or Chebychev window function) before the FFT to ensure a moderate dynamic range of signals can be detected. This cautious approach has been mentioned in section 1.4. The penalty to pay with the use of bell-shaped windows is that the DRx spectral estimate frequency resolution is effectively reduced because the mainlobes of all spectral peaks become wider. An explanation of this effect is because the window effectively diminishes the significance of samples (which contain signal power) at the ends of the DCW. Use of the SVA technique on the captured data however permits the frequency resolution of a rectangular window function, suppresses the sidelobe level and approximately maintains the processing gain of the rectangular window.

Therefore the aim of the experiment was to capture opportunistically real-life data, process it with the use of SVA and other conventional window functions and prove the advantages of SVA were evident on real-life data. It was expected that the SVA technique would afford advantages in the following possible scenarios which could have been encountered:

1. When there are emitters very closely spaced in frequency and at similar received power. A Hann window function or Chebychev window function would be unable to tell them apart.
2. When signal spectral peaks occur over a moderately-large dynamic range, in other words, when there are strong spectral peaks and weak spectral peaks. The weak spectral peaks would be masked by the sidelobes of the strong spectral peaks if a rectangular window function was used.
3. When a processing gain similar to the rectangular window function was required.

A.4 Experiment



The advantages of SVA are prominent in the detection of narrowband signals. This is described and shown in chapter 3. Marine radar beacons, or racons for short, are radar emitters with very simple sinusoidal pulses. Therefore the SVA algorithm would be expected to perform well in their detection. Racons usually operate in both the X-band and the S-band. The S-band emissions normally occur within the frequency band 2.920–3.100 GHz. To cover this, the ESM dwell that was selected pre-flight in the DRx was ESM dwell number 4¹.

It was also known that the Glasgow airport Watchman A radar and the Edinburgh airport Watchman A radar operated within the frequency range of this ESM dwell and therefore it was predicted that some of their pulses would also be captured at the same time. The Watchman radars had been recorded during previous trials with parameters displayed in Table A.4.1² and so it was known what possible parameters their pulses might exhibit in intercepted data.

Identity	RF (MHz)	Pulse Train Width (ms)	Radar Dwell (ms)
Watchman A radars	[2750, 2950]	[0.5, 20]	[40, 1050]
	Agile from pulse-train to pulse-train	Agile	Staggered

Table A.4.1 Watchman A Parameters

The ESM dwell input was mixed down by the local oscillator (LO) in the superhet to a lower IF frequency band. This bandwidth was sampled using real-only sampling with an ADC at a particular sampling rate f_s from the interval [1.0, 1.75] GHz. The superhet mixed the bandwidth down into the second Nyquist band of the ADC which promptly folded down into the first Nyquist band (conjugate symmetry about the Nyquist frequency, see panel (b) of Fig A.4.1 for a pictorial explanation of this). The Nyquist frequency was greater than the desired sampled bandwidth, i.e. the input was oversampled. The reason for this was because it allowed a less-than-perfect analogue filter with a gentle roll-off to be used after the LO. Consequentially, the degree of complexity of the analogue filters was traded-off against the number of useful digital

¹This number is an index for a frequency range not shown here

²Exact parameters are not shown

FFT frequency bins because some bins at the edges representing the analogue filter transition band were discarded.

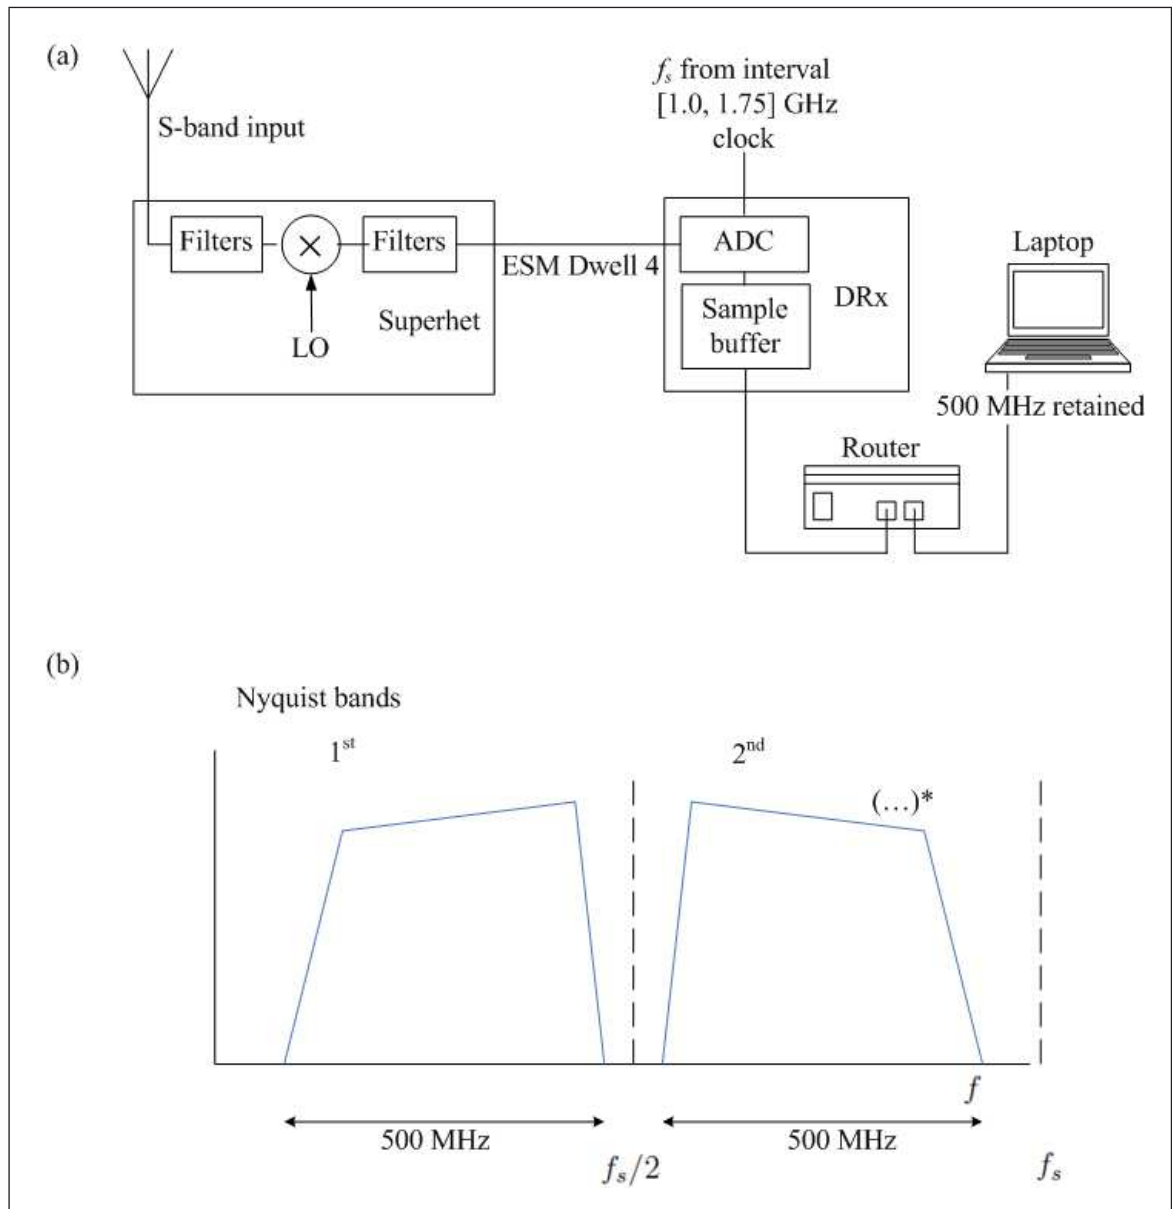


Figure A.4.1 Experimental set-up: (a) system design (detail omitted); (b) depiction of the ADC Nyquist bands

A.5 Results

The captured sample data were processed in MATLAB[®] with the use of an on-board laptop. The custom MATLAB[®] functions were designed to model a real-life DRx. For example the sample data was processed with non-overlapped and non-zero padded $N_{\text{FFT}} = 1024$ FFTs of which the magnitude-squared was taken. This was chosen as it is a reasonably sized FFT which a real-life DRx would be able to process within time constraints.

The actual experiment produced twenty-one 8 MB files which contained sample data

and a full report was produced on these results however only a selection of the most relevant data is shown in Figs A.5.1–A.5.9. The dynamic range of the z-axis scale is maintained between plots of the same sample file. The units of the z-axis are dB/Hz, i.e. power spectral density. To convert to physical units of dBm/Hz the following conversion is required:

$$S(f, T)_{\text{dBm/Hz}} \approx \text{FS}_{\text{dBm}} - \rho \frac{S(f, T)_{\text{dB/Hz}}}{20} - 3.01N_b + 33.01 \quad (\text{A.5.1})$$

where T is the duration of a DCW, ρ is the reference resistance, N_b is the number of ADC bits, FS_{dBm} is the full-scale level of the ADC and there is an addition of 30 involved to convert Watts to milliWatts. The conversion is approximate because full-scale is slightly asymmetric for positive and negative levels in a linear ADC.

Plots of rectangular-windowed and Hann-windowed spectrograms are compared against the same data processed with the use of the SVA algorithm. The SVA algorithm was applied along the frequency axis of the rectangular-windowed STFT. There is one plot for each type of window technique under the three example sample files selected. Only relevant frequency bins which contain a 500 MHz portion of the S-band mentioned in section A.4 are displayed in the spectrogram plots. Also only real data which was analysed from one antenna are displayed in the plots because the results obtained from analysis of other antennas were almost identical to the first.

A.5.1 Sample File 0

This sample file was created to test the system design was working correctly inside the hangar before the flight. However it was useful as it served as an indication of the noise floor and any interference signals present inside the helicopter. There was clearly a continuous interference tone present with frequency ~ 2.74 GHz. This is marked on the plots in Figs A.5.1–A.5.3.

Although there are no clear external signals of interest in this sample file, the broad-mainlobe effect which the Hann window function causes is conspicuous upon examination of the continuous interference in Figs A.5.1–A.5.3. This results in the Hann-windowed spectrogram having a coarser resolution in the frequency direction. The spectral line caused by the continuous interference in the Hann-windowed spectrogram is roughly double the width in the frequency dimension than that of the rectangular-windowed spectrogram and SVA-spectrogram.

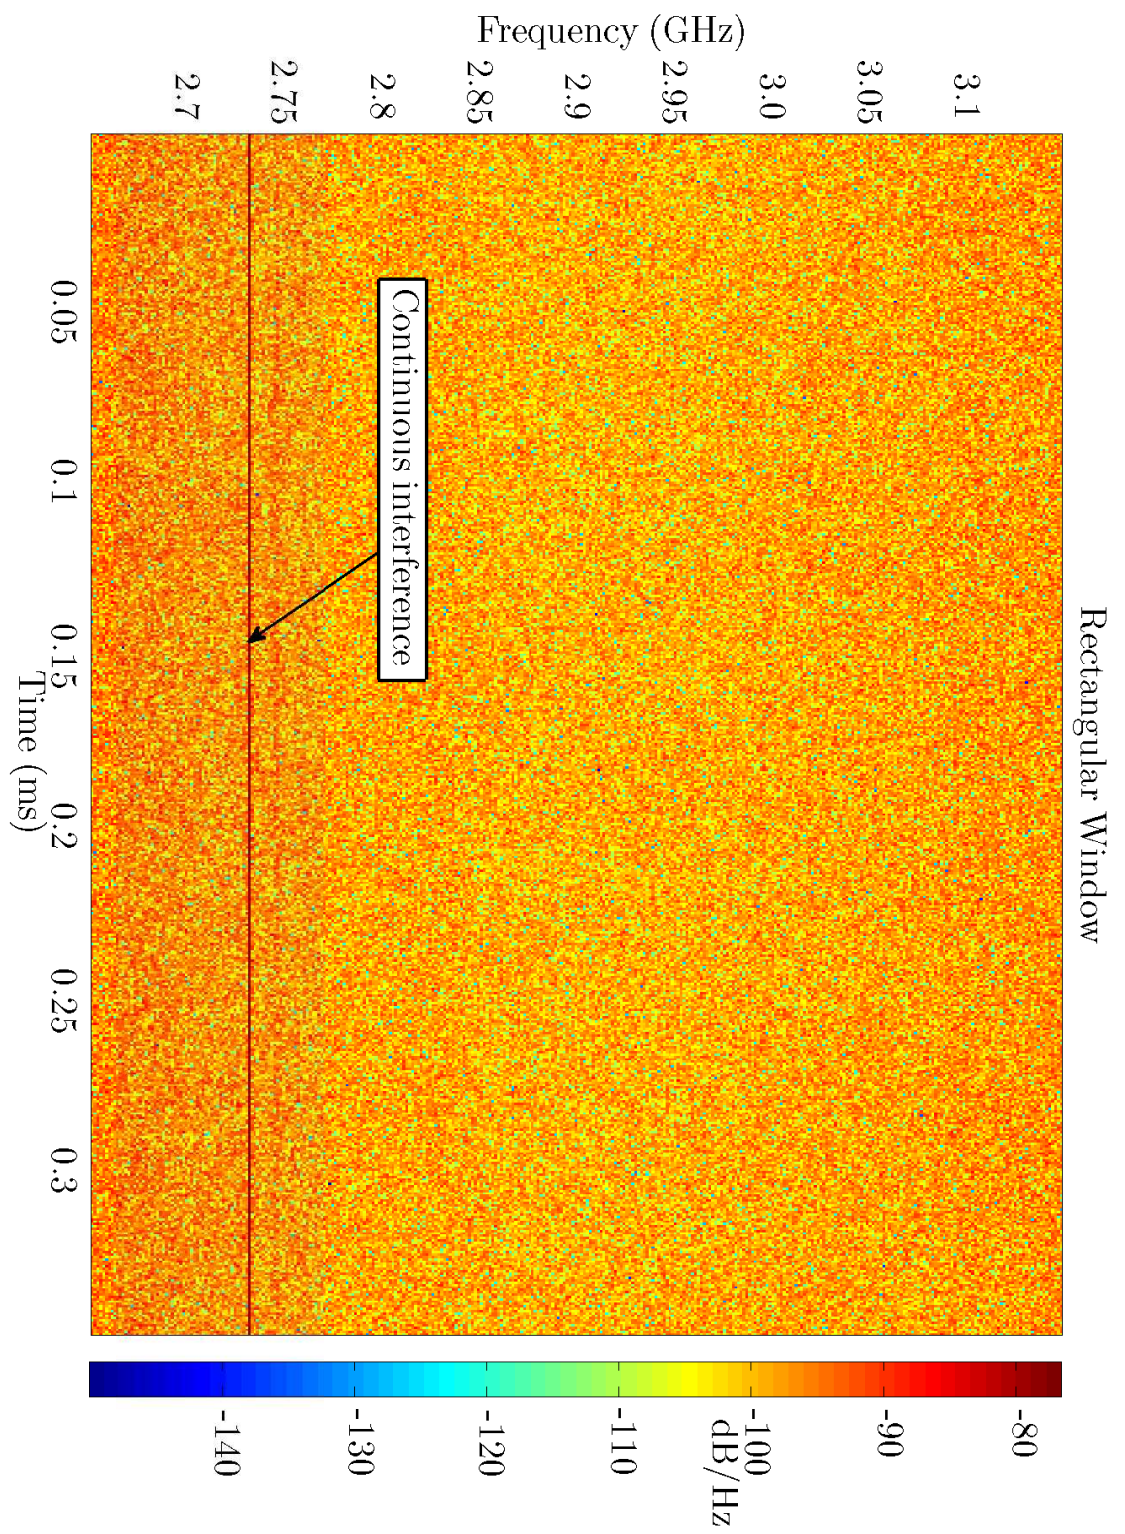


Figure A.5.1 Sample file 0 rectangular-windowed spectrogram

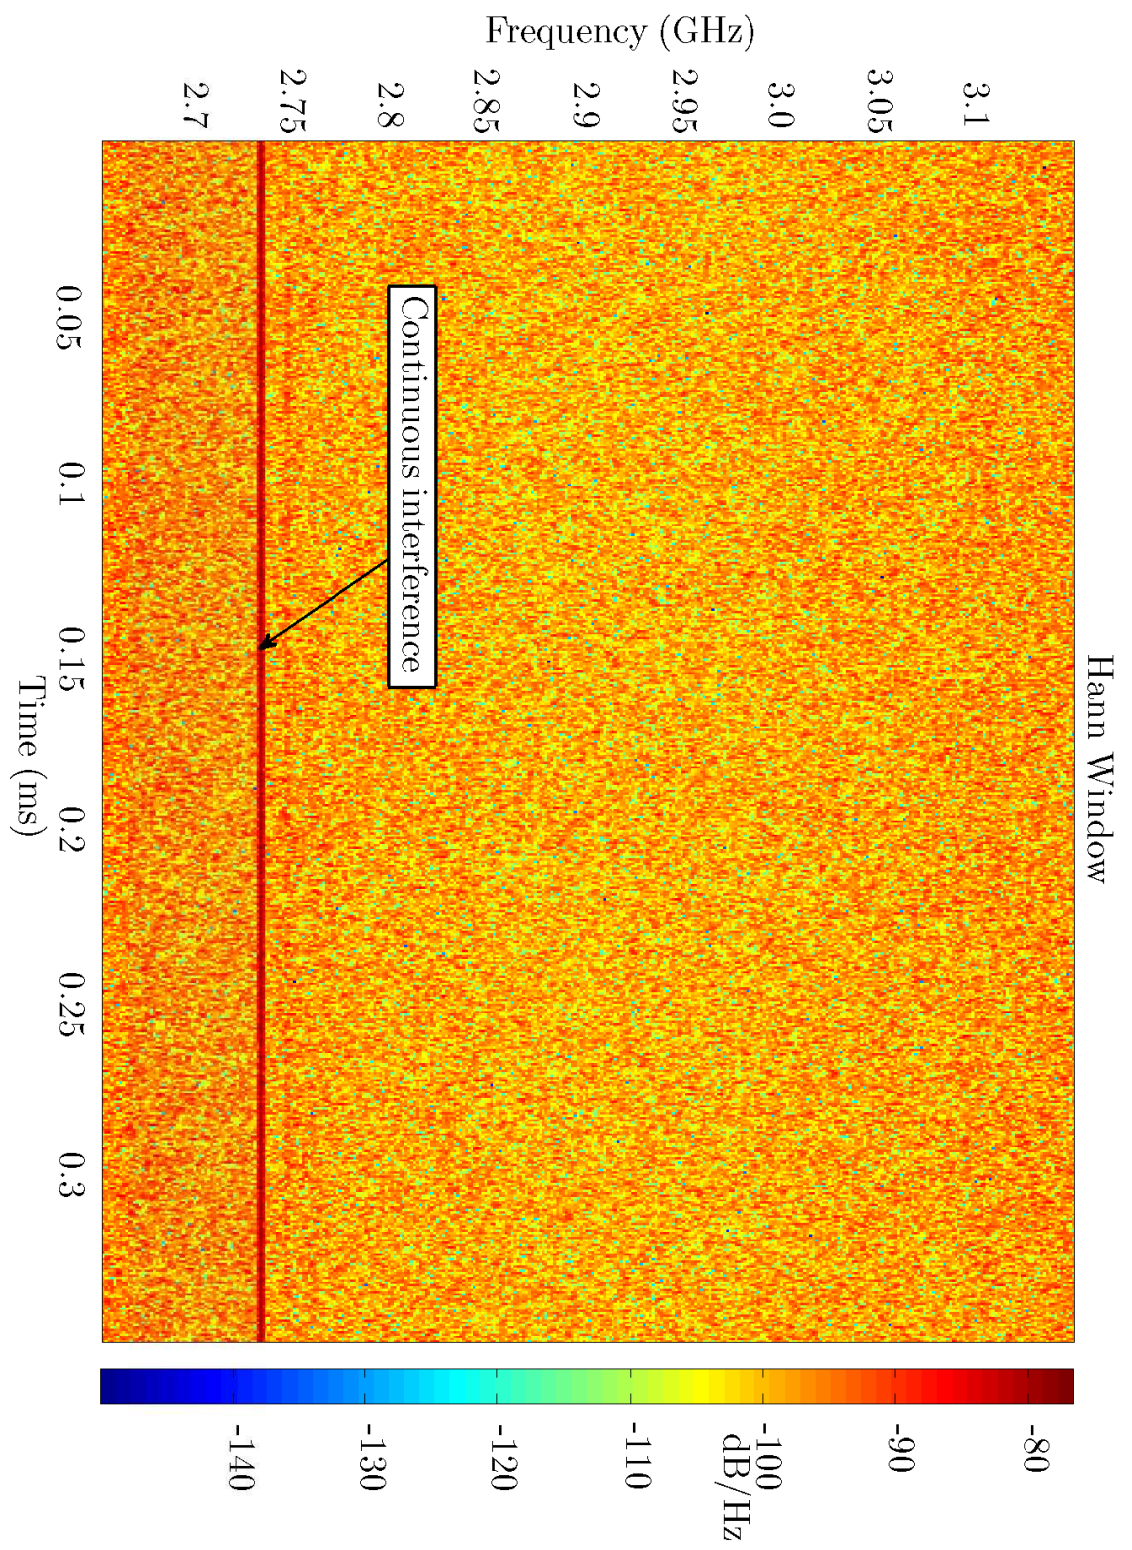


Figure A.5.2 Sample file 0 Hann-windowed spectrogram

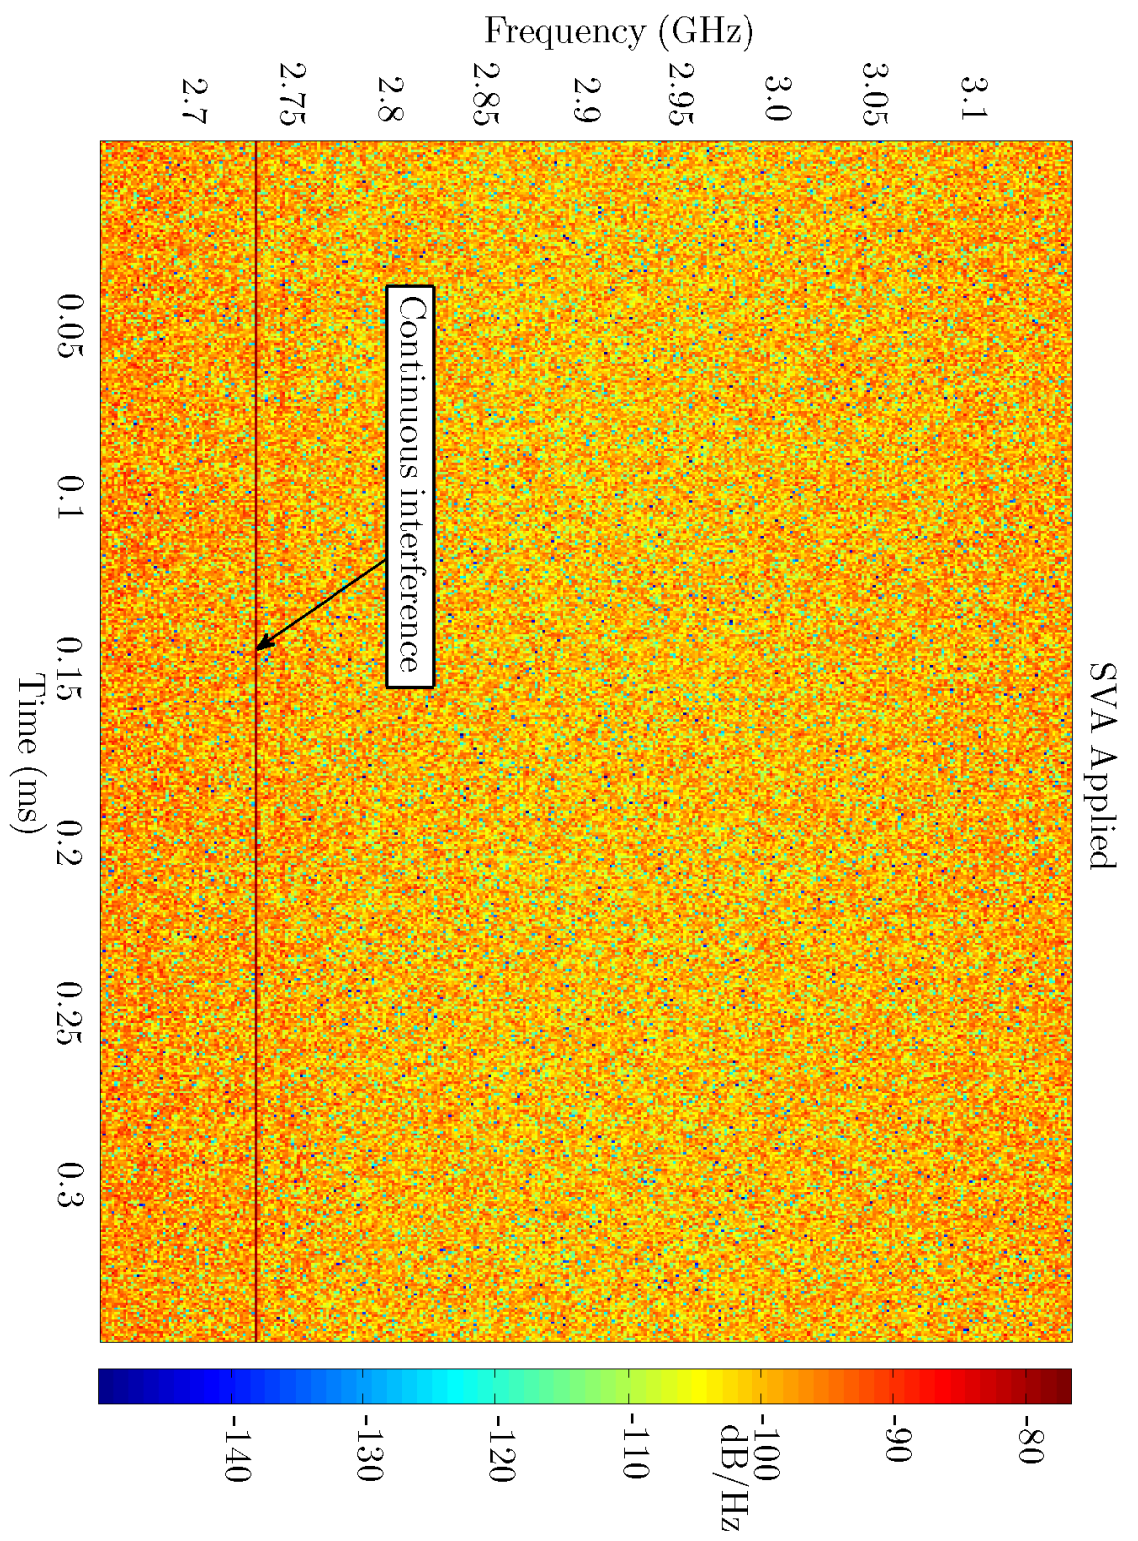


Figure A.5.3 Sample file 0 SVA spectrogram

A.5.2 Sample File 8

A racon signal was the strongest signal captured in this sample file. This is visible from the results shown in Figs A.5.4–A.5.6 which were produced with the use of a rectangular-windowed spectrogram, Hann-windowed spectrogram and SVA spectrogram respectively. The racon signal had a centre frequency of approximately 2.965 GHz and a pulse repetition interval (PRI) of approximately 0.7 ms. The PRI parameter was obtained from examination of later samples. Upon closer examination of the fine structure of the signal it appeared as a small chirp signal of span a few megahertz.

The strong signal is interesting because it corresponded to the second scenario listed from section A.3. The results from the strong signal neatly demonstrates one of the advantages of the use of SVA. In Fig A.5.4 the results show that the use of a rectangular-windowed spectrogram resulted in large amounts of spectral leakage from high sidelobes of the strong signal. This helps to support the point repeated throughout the study, that a DRx designer will usually avoid the rectangular window function because high sidelobes can cause erroneous detections. The use of the Hann-windowed spectrogram and SVA spectrogram resulted in much lower sidelobes for the strong signal in Figs A.5.5 and A.5.6 respectively.

A weaker signal from the Watchman radar at Edinburgh airport was also present in the sample file and is noticeable in all the plots. When analysed alongside later samples, this captured signal appears as pulse trains of three or four pulses which then hop to a different carrier frequency. The continuous interference tone was also still present in the sample file.

From inspection of the general features of the plots in Figs A.5.4–A.5.6, the other advantages of the use of SVA become clear. The use of SVA resulted in an improved frequency resolution than when the Hann window was used. This is similar to the first point from the list in section A.3. In addition to this, visually, the use of SVA in a spectrogram creates a lower noise floor than the use of the Hann window.

Overall the conclusion which can be drawn from the examination of the data from this sample file is that the use of the SVA algorithm did provide the expected advantages over the use of a rectangular window function and a bell-shaped window function in the guise of the Hann window.

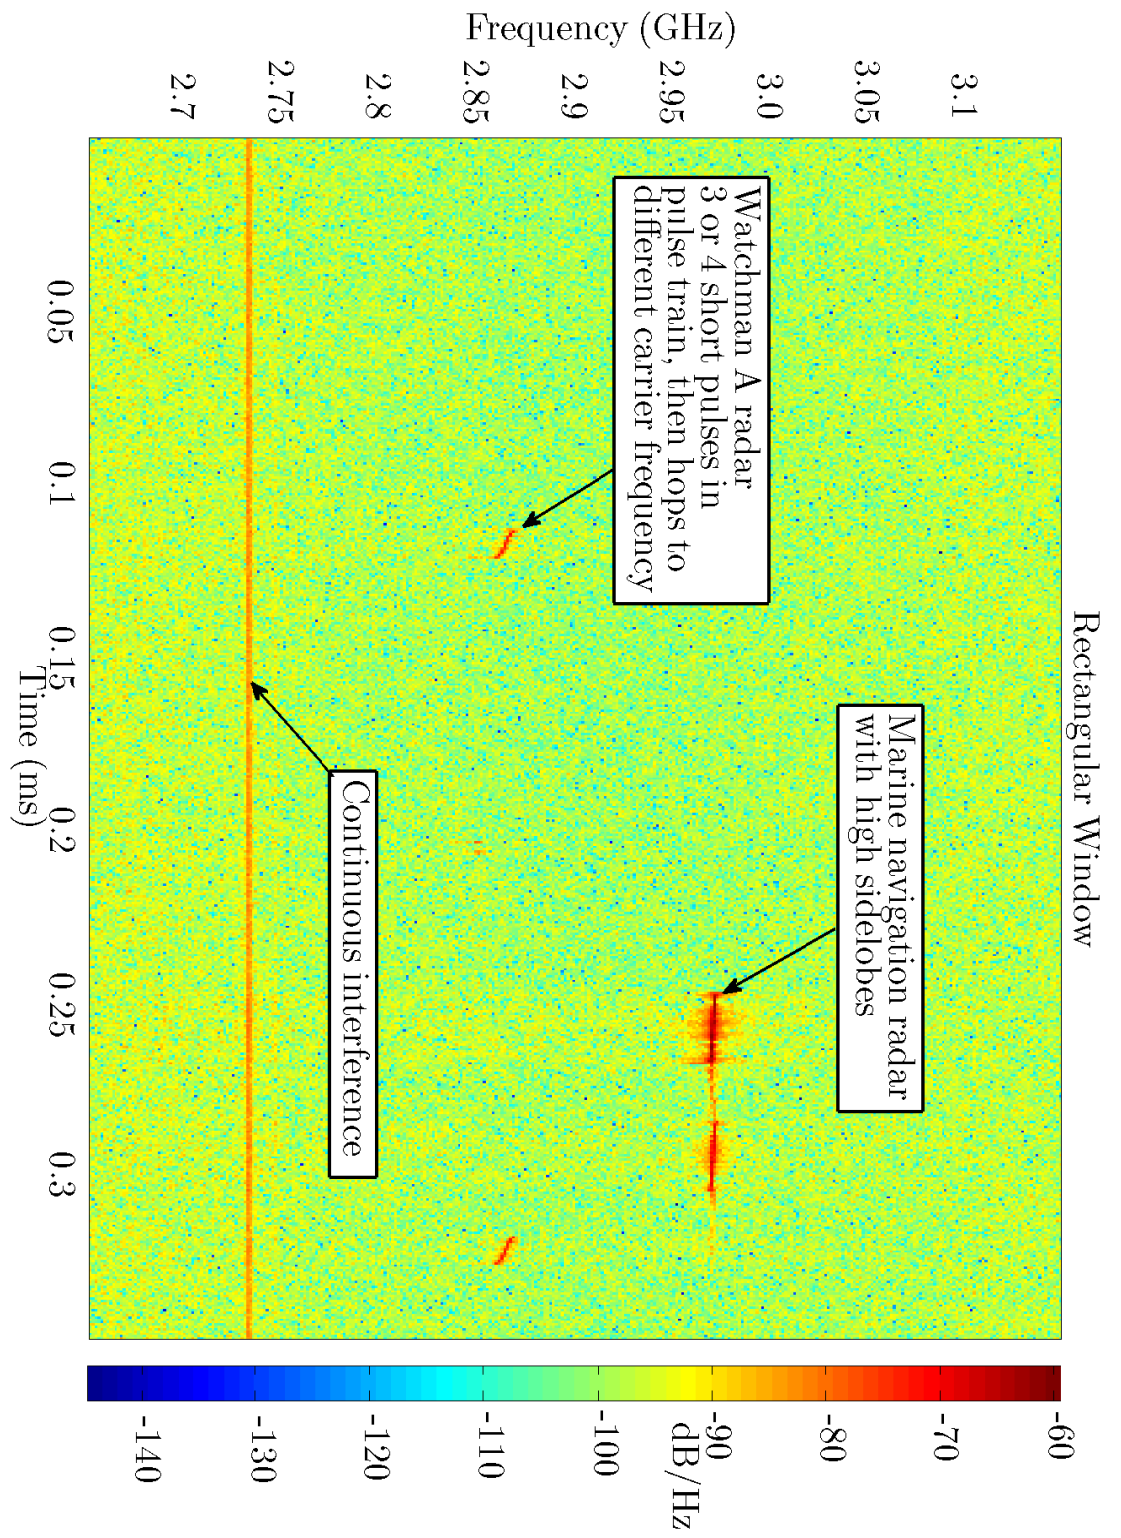


Figure A.5.4 Sample file 8 rectangular-windowed spectrogram

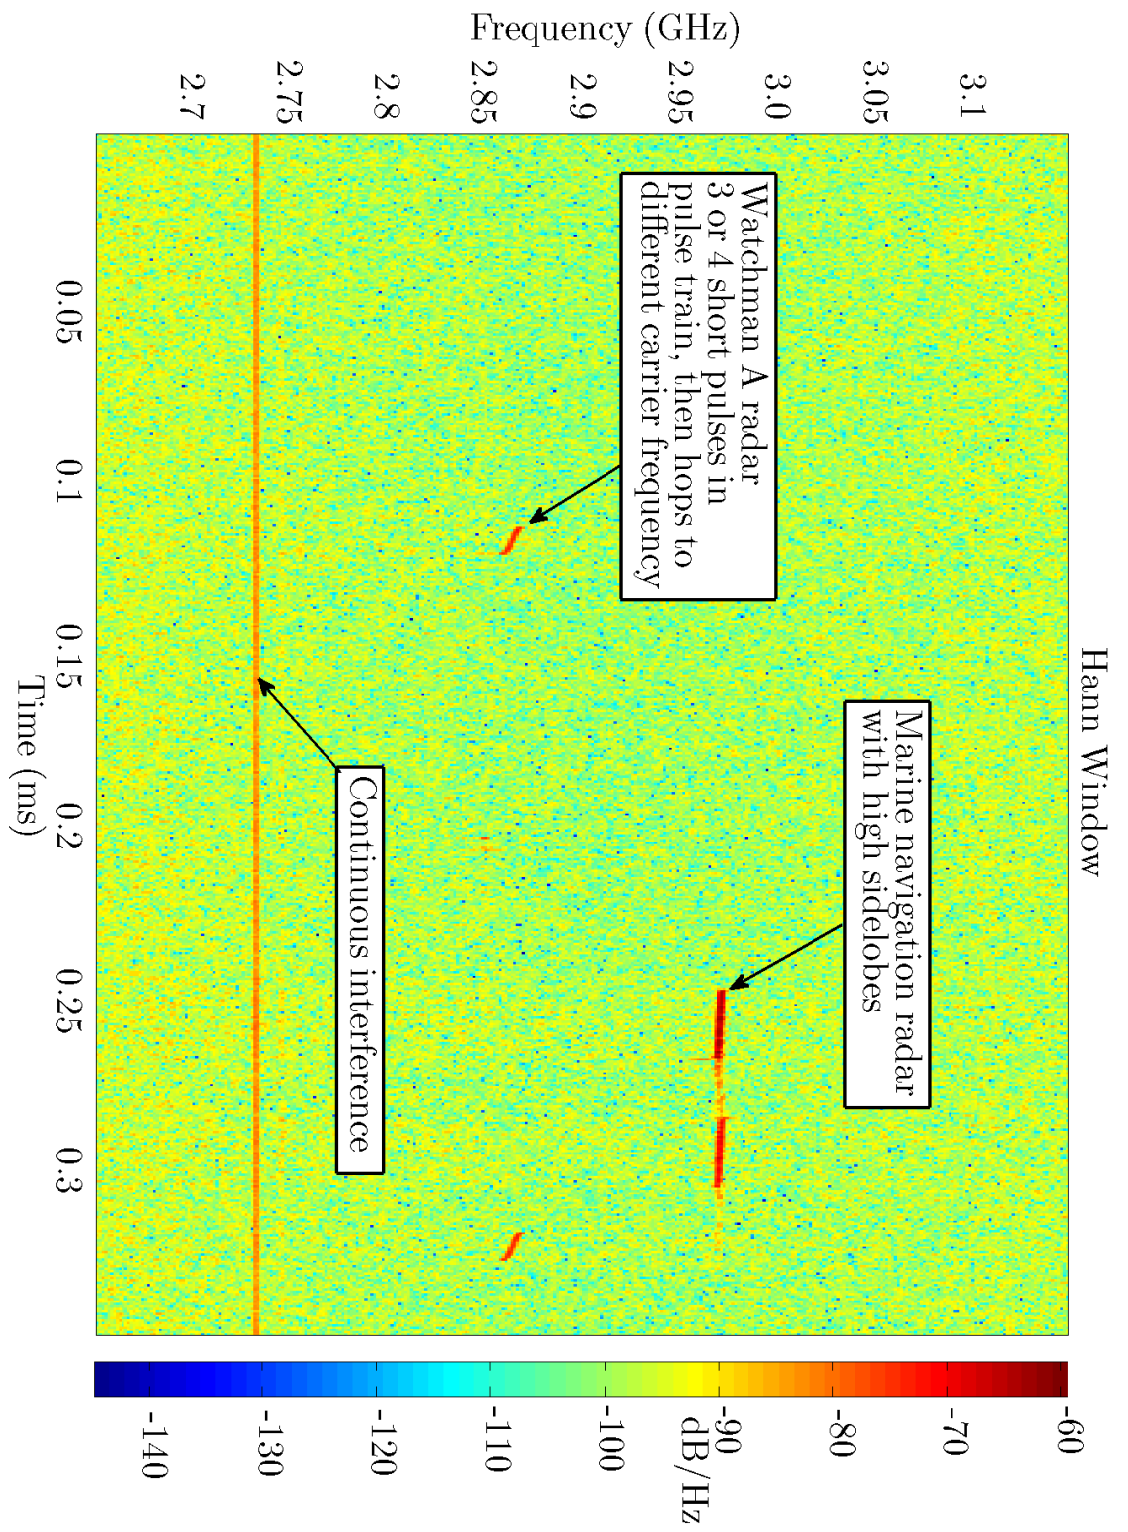


Figure A.5.5 Sample file 8 Hann-windowed spectrogram

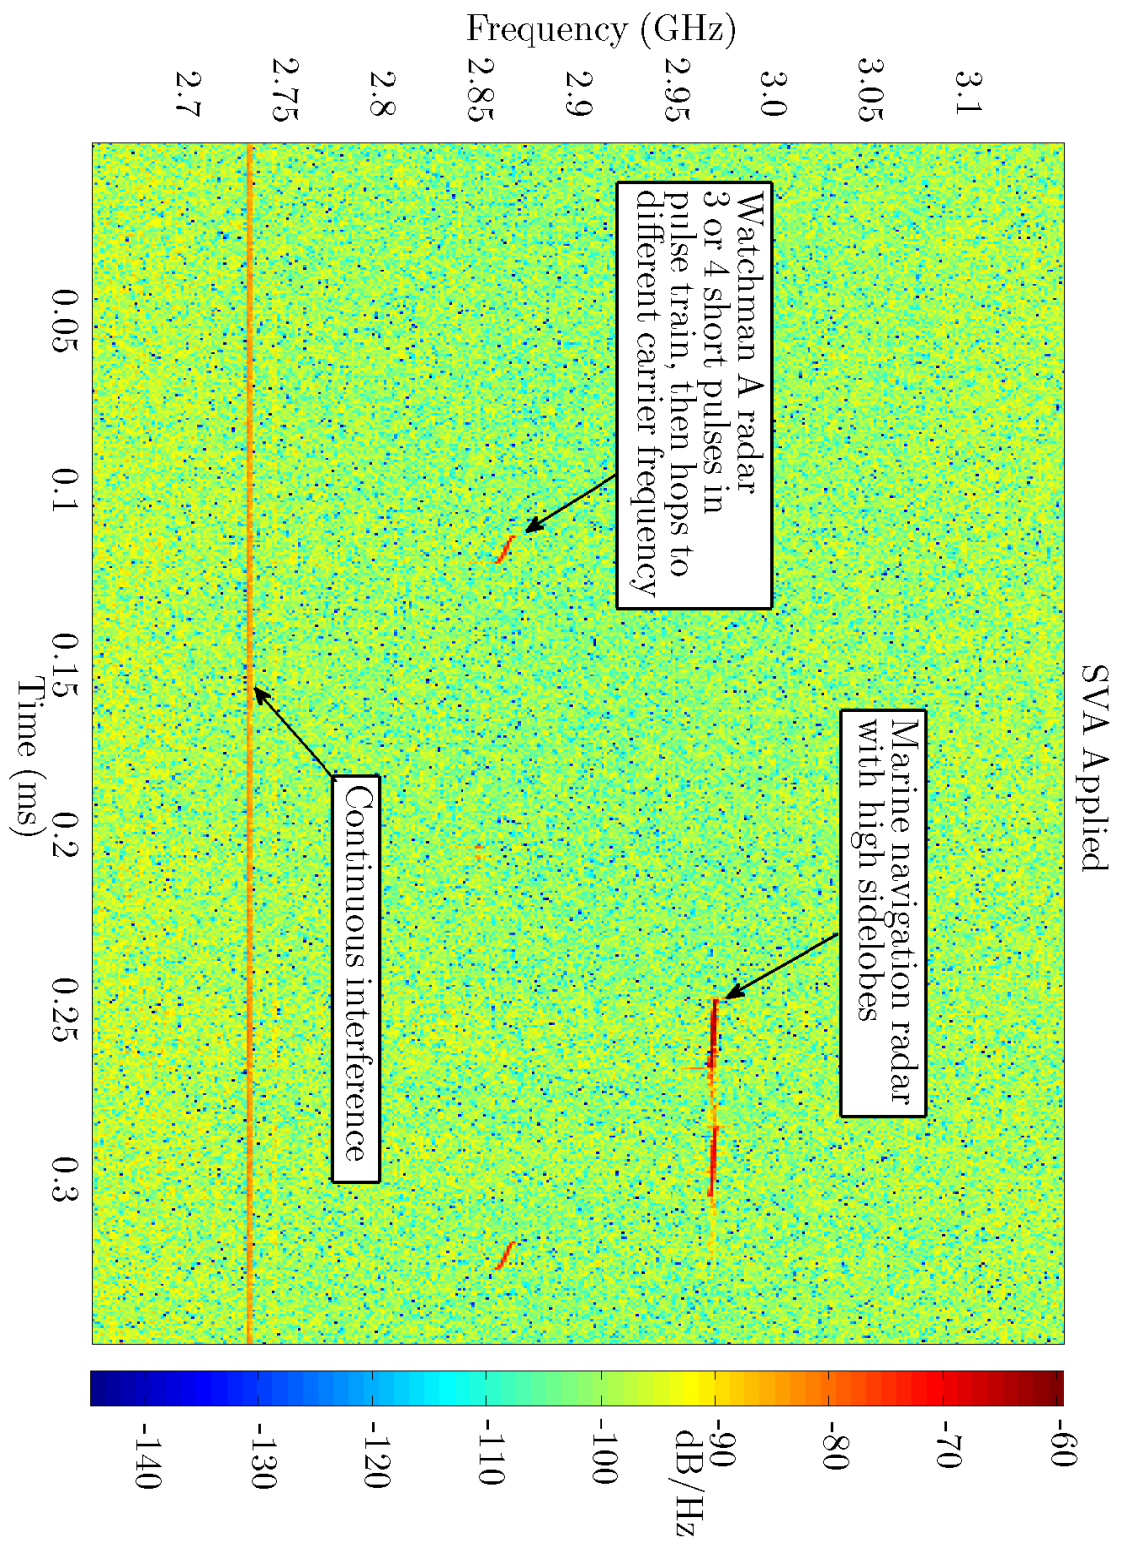


Figure A.5.6 Sample file 8 SVA-windowed spectrogram

A.5.3 Sample File 11

There were several signals contained within this sample file however only a few are visible in Figs A.5.7–A.5.9 because these display only an extract of the samples from the file. The spectral peaks due to the Watchman radar are again visible and when viewed alongside later samples they appear as pulse trains of three or four pulses between the approximate frequencies ~ 2.89 GHz and ~ 2.94 GHz. There were also two different racon signals collected in the sample file one of which is marked on the plots. As from the previous sample file, the same conclusions were drawn with regards to the differences between the spectrograms in Figs A.5.7–A.5.9 which were calculated with the use of the rectangular window, Hann window and SVA:

- The SVA spectrogram has better sidelobe attenuation than the rectangular-windowed spectrogram.
- The SVA spectrogram has a better frequency resolution than the Hann-windowed spectrogram.
- The SVA spectrogram has a better response to the noise floor than the Hann-windowed spectrogram.

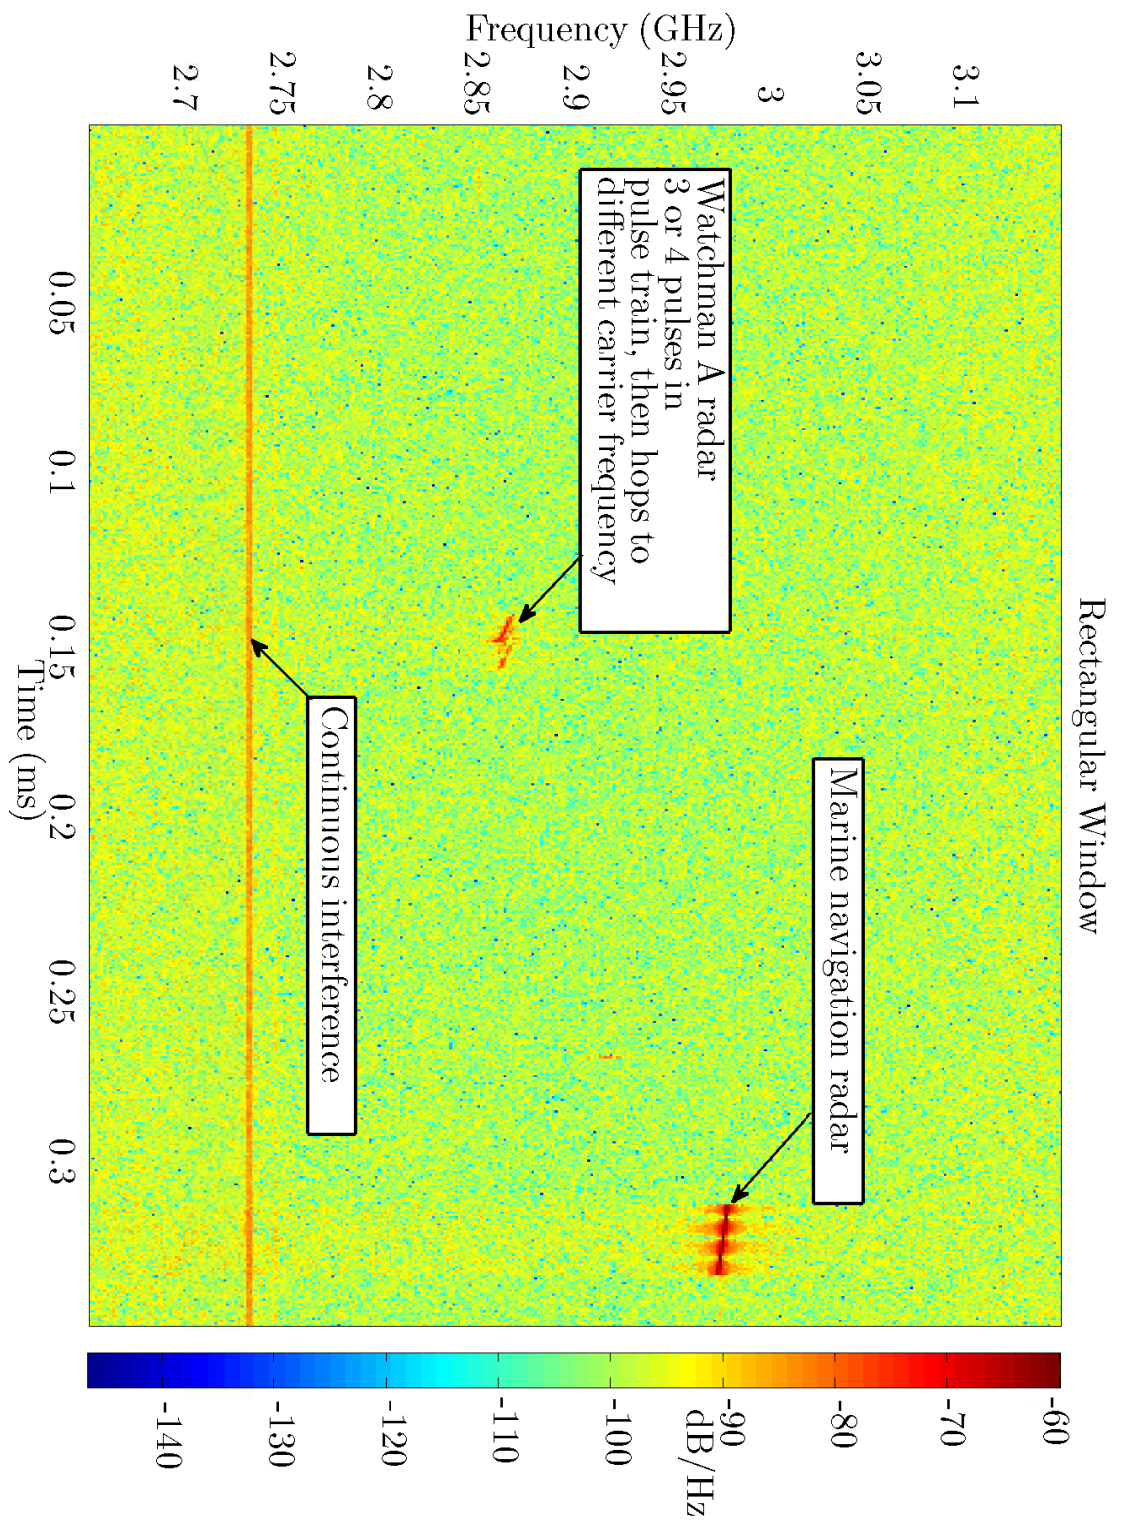


Figure A.5.7 Sample file 11 rectangular-windowed spectrogram

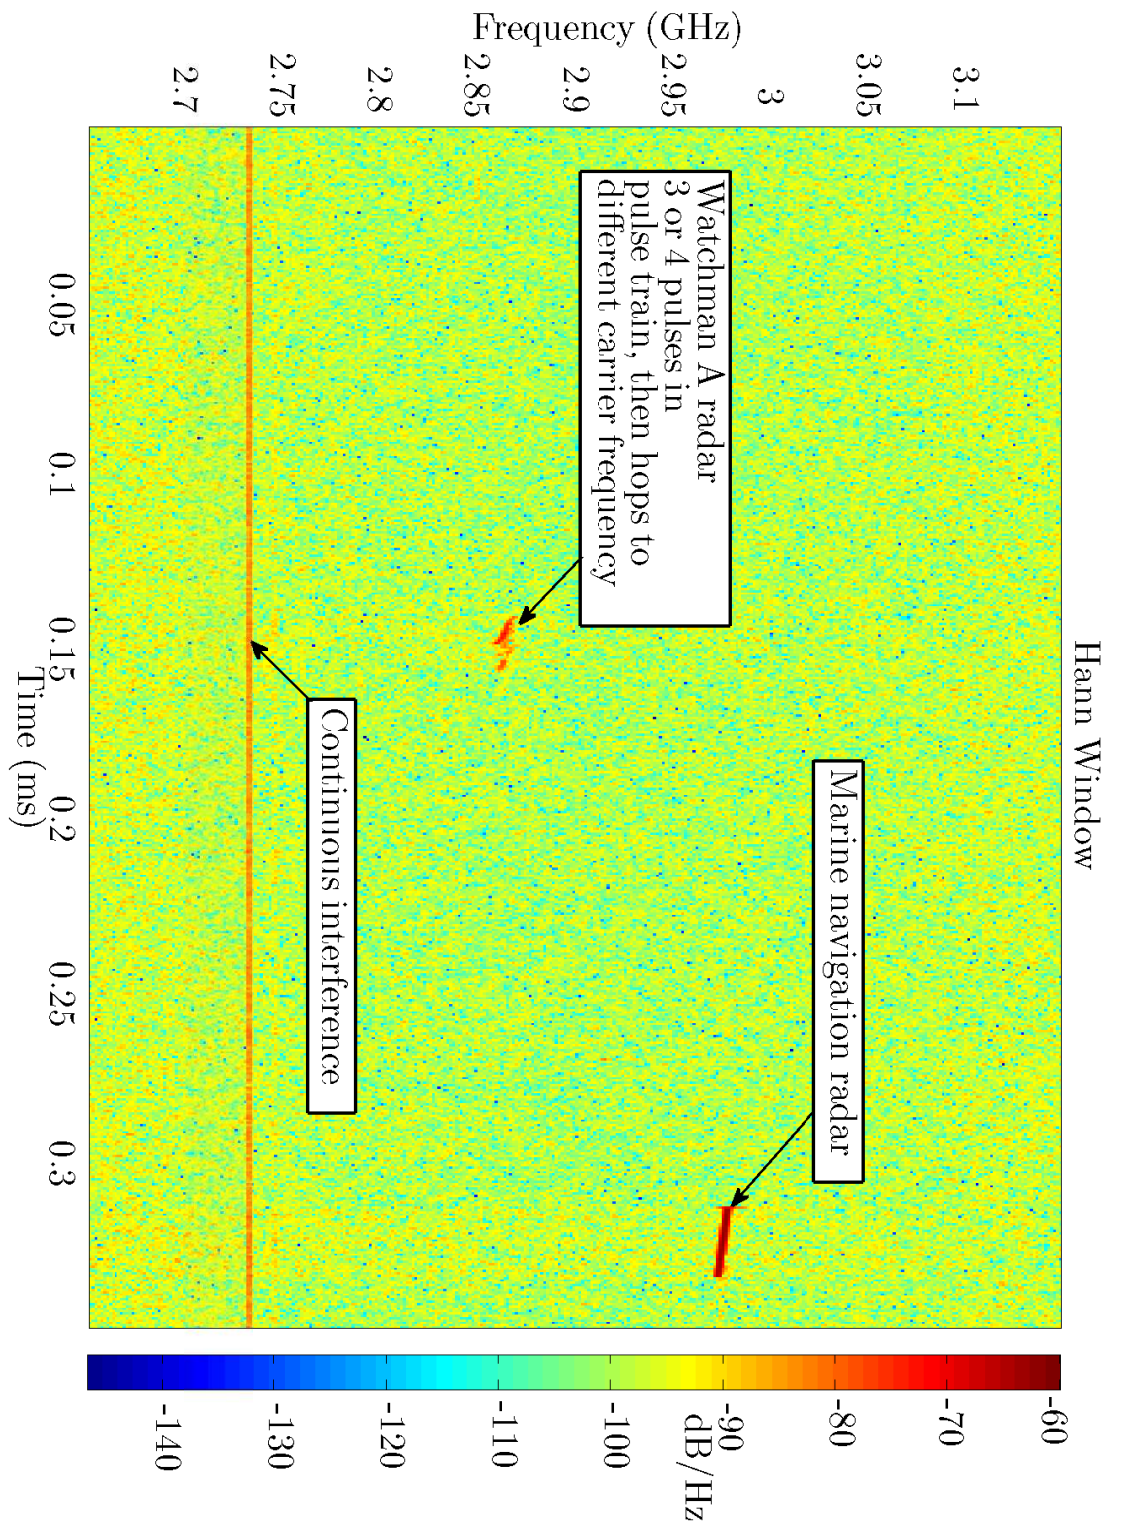


Figure A.5.8 sample file 11 Hann-windowed spectrogram

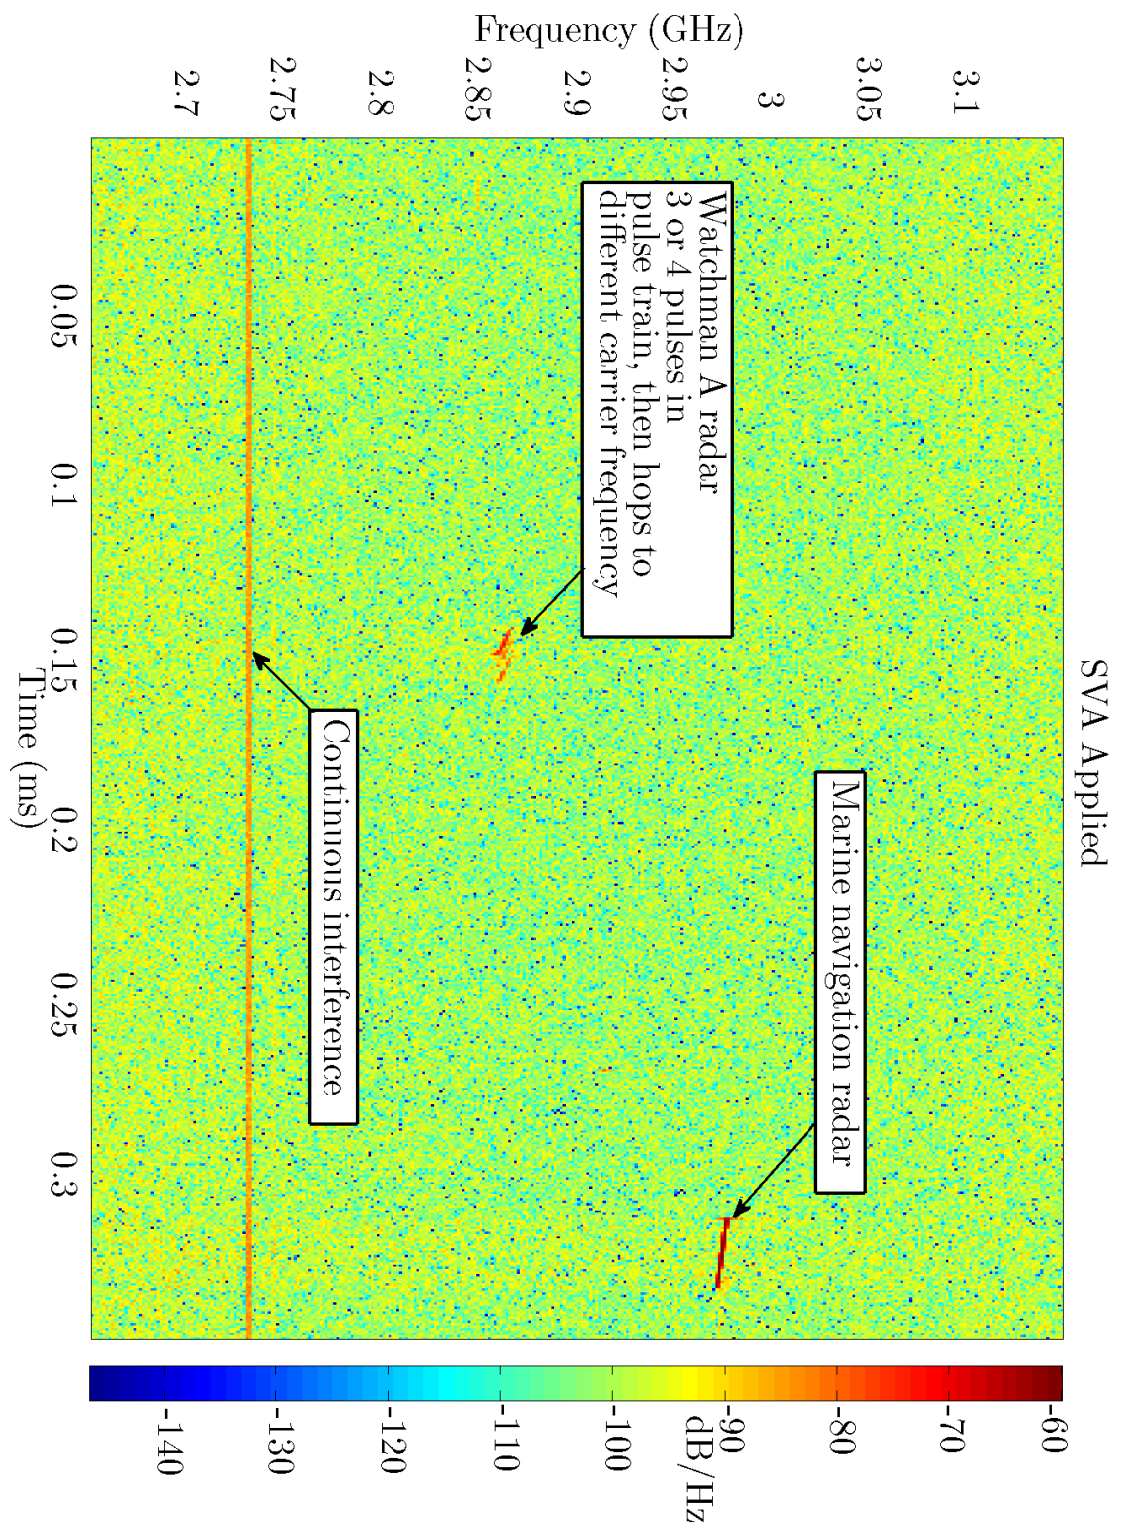


Figure A.5.9 Sample file 11 SVA spectrogram

A.6 Summary of Results

The results of the helicopter trials (for clarity not all are shown) can be summarised by three main points, which confirmed the claims about SVA made in chapter 2:

1. The SVA algorithm maintains the frequency-domain representation resolution of that of the rectangular window function. The use of the Hann window, which is a bell-shaped window function, broadens the mainlobe of any spectral peaks.
2. The SVA algorithm provides good sidelobe attenuation. The attenuation is similar to that of a Hann window function. However a rectangular window function results in noticeable spectral leakage due to relatively high sidelobes.
3. The processing gain of the Hann window function is poorer than that of the rectangular window function and SVA. This is directly related to the NEB of the window functions. For the Hann-windowed spectrograms, the spectral peaks are noticeable lower in relation to the noise floor, however the SVA algorithm does not significantly raise the noise floor in the same manner.

To illustrate further and reinforce these three advantages of the SVA algorithm, some captured sample data from the flight were extracted and are shown below.

Fig A.6.1 contains periodograms formed using $N_{\text{FFT}} = 2048$ and $N = 1024$. These periodograms display a snapshot of the estimated frequency content of the data samples at that time interval. The periodograms are analogous to taking slices through full spectrograms. In the plots the full digitised IF frequency range is not drawn so that the plots can concentrate on a range of interest which contains a signal. Fig A.6.1 was produced from the captured data samples in the file `sample_200809_0007.txt`. The specific 1024 data samples were taken after $\tau \approx 799 \mu\text{s}$, where $\tau = 0 \mu\text{s}$ was defined as the start of a sample capture file. The periodograms show the mainlobe of a spectral peak from the SVA technique is similar to that of the rectangular-windowed FFT mainlobe. However the mainlobe of the spectral peak when a Hann window function was used before the FFT is approximately double that of the rectangular-windowed periodogram and SVA-periodogram. This illustrates point one in the above summary list.

The periodograms also show that the SVA technique and the Hann window function have good sidelobe attenuation in comparison to the rectangular window function, which thus illustrates point two from the above summary list.

Fig A.6.2 was extracted from a particular area of three spectrograms which were calculated with the use of the different windowing techniques: (1) SVA; (2) rectangular window function and (3) Hann window function. It was produced from the sample capture file `sample_200809_0006.txt`. A chirped pulse with a bandwidth spanning a few megahertz is visible. Panel (a) illustrates high sidelobe spectral leakage as a result

of the use of a rectangular window function (point one made in the above summary list). Panel (b) depicts the coarser frequency resolution caused by the use of a Hann window function and reinforces point two from the above summary list. Panel (c) shows the SVA spectrogram has good frequency resolution, good sidelobe attenuation and that the noise floor is not raised significantly in comparison to panel (a), which supports point three made in the above summary list.

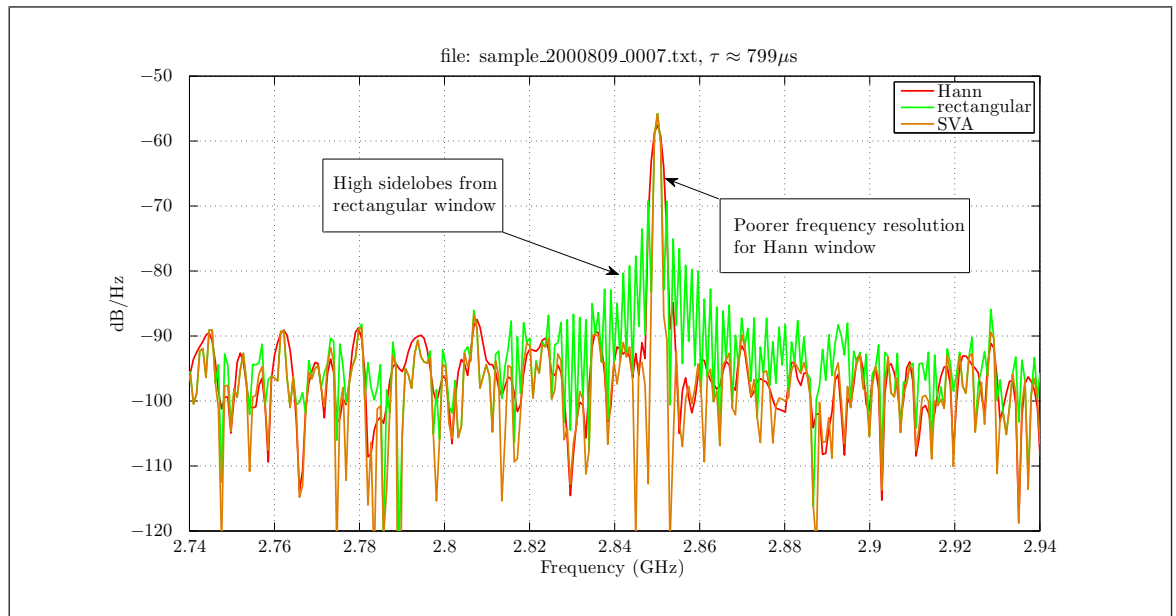


Figure A.6.1 Twice-zero-padded periodogram spectral estimates using the various windowing techniques

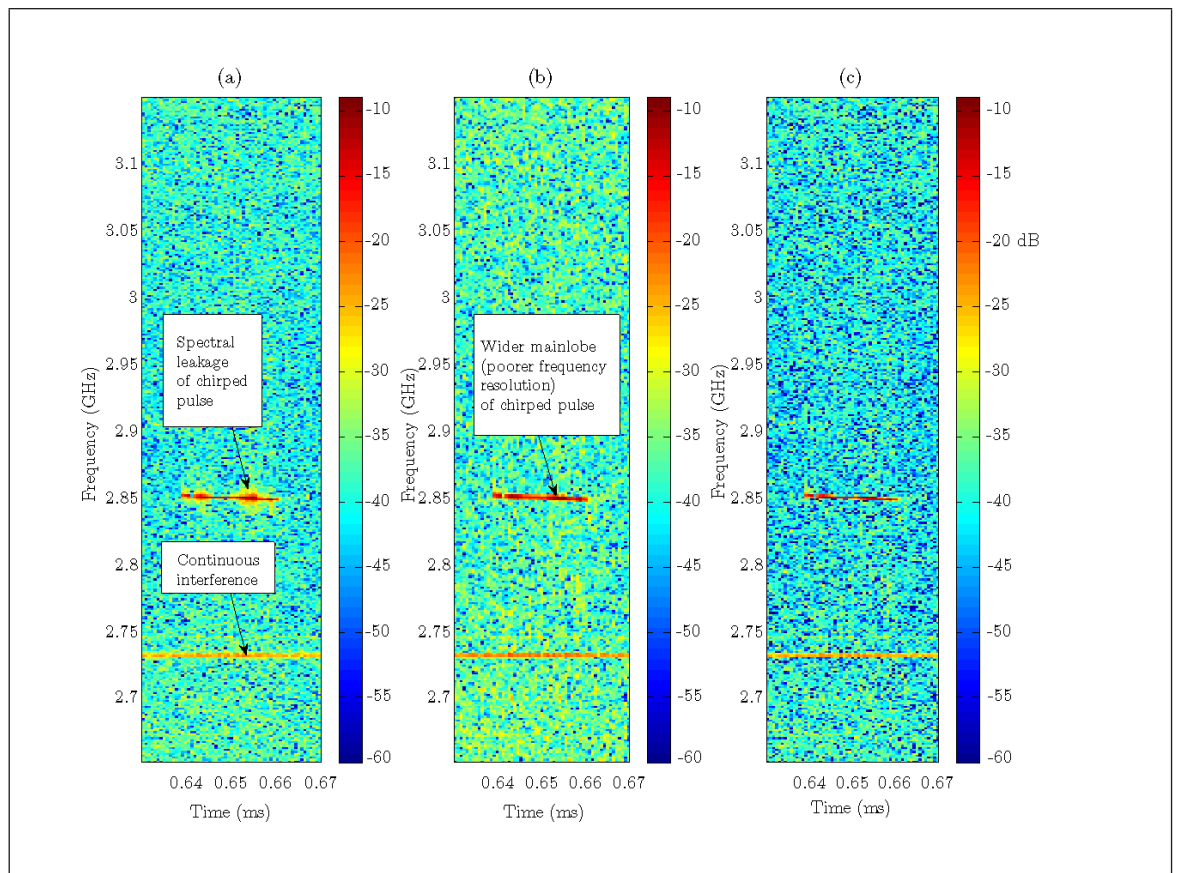


Figure A.6.2 Spectrogram extracts of a chirp with the use of: (a) rectangular windowing; (b) Hann windowing; (c) use of SVA

APPENDIX B

Complex-Parametered SVA

B.1 Unconstrained Complex SVA

In a similar formulation to the SVA technique, the cost function (C) to be minimised is formed from the complex-filtered 3-tap equation and its complex conjugate

$$\begin{aligned}X_a(k) &= X(k) - \beta X(k+1) - \beta^* X(k-1) \\X_a^*(k) &= X^*(k) - \beta^* X^*(k+1) - \beta X^*(k-1)\end{aligned}$$

Given explicitly, the cost function to be minimised is

$$\begin{aligned}C &= X_a(k)X_a^*(k) \\&= (X(k) - \beta X(k+1) - \beta^* X(k-1)) \\&\quad \times (X^*(k) - \beta^* X^*(k+1) - \beta X^*(k-1))\end{aligned}\tag{B.1.1}$$

Upon expanding the brackets in Eq (B.1.1), the following terms are obtained

$$\begin{aligned}C &= |X(k)|^2 - \beta^* X^*(k+1)X(k) - \beta X^*(k-1)X(k) - \beta X(k+1)X^*(k) \\&\quad - \beta^* X(k-1)X^*(k) + |\beta X(k+1)|^2 + |\beta X(k-1)|^2 \\&\quad + \beta^2 X(k+1)X^*(k-1) + (\beta^*)^2 X^*(k+1)X(k-1)\end{aligned}\tag{B.1.2}$$

Since β , $X(k)$, $X(k+1)$ and $X(k-1)$ are all complex numbers, they can be represented in terms of real and imaginary parts, specifically

$$\begin{aligned}\beta &= \beta_R + i\beta_I \\X(k) &= X_R(k) + iX_I(k) \\X(k+1) &= X_R(k+1) + iX_I(k+1) \\X(k-1) &= X_R(k-1) + iX_I(k-1)\end{aligned}\tag{B.1.3}$$

Therefore the β terms in Eq (B.1.2) can be substituted using Eq (B.1.3) to give

$$\begin{aligned} C = & |X(k)|^2 - (\beta_R - i\beta_I)X^*(k+1)X(k) - (\beta_R + i\beta_I)X^*(k-1)X(k) \\ & - (\beta_R + i\beta_I)X(k+1)X^*(k) - (\beta_R - i\beta_I)X(k-1)X^*(k) \\ & + (\beta_R + i\beta_I)(\beta_R - i\beta_I)|X(k+1)|^2 + (\beta_R + i\beta_I)(\beta_R - i\beta_I)|X(k-1)|^2 \\ & + (\beta_R + i\beta_I)^2X(k+1)X^*(k-1) + (\beta_R - i\beta_I)^2X^*(k+1)X(k-1) \end{aligned}$$

To find the minimum value of the cost function (C) with respect to β , the first step is to search for stationary points by taking the first partial derivatives. The partial derivative of C with respect to the real part of β is

$$\begin{aligned} \frac{\partial C}{\partial \beta_R} = & -X^*(k+1)X(k) - X^*(k-1)X(k) - X(k+1)X^*(k) \\ & - X(k-1)X^*(k) + (\beta_R + i\beta_I)|X(k+1)|^2 + (\beta_R + i\beta_I)|X(k-1)|^2 \\ & + (\beta_R - i\beta_I)|X(k+1)|^2 + (\beta_R - i\beta_I)|X(k-1)|^2 \\ & + 2(\beta_R + i\beta_I)X(k+1)X^*(k-1) + 2(\beta_R - i\beta_I)X^*(k+1)X(k-1) \\ = & -2\Re\{X^*(k+1)X(k) + X^*(k-1)X(k) - (\beta_R + i\beta_I)|X(k+1)|^2 \\ & - (\beta_R + i\beta_I)|X(k-1)|^2 \\ & - 2(\beta_R + i\beta_I)X(k+1)X^*(k-1)\} \end{aligned} \quad (\text{B.1.4})$$

Taking the real parts of Eq (B.1.4) gives

$$\begin{aligned} \frac{\partial C}{\partial \beta_R} = & -2(X_R(k+1)X_R(k) + X_I(k+1)X_I(k) + X_R(k-1)X_R(k) \\ & + X_I(k-1)X_I(k)) + 2\beta_R(|X(k+1)|^2 + |X(k-1)|^2) \\ & + 4(\beta_R X_R(k+1)X_R(k-1) - \beta_I X_I(k+1)X_R(k-1) \\ & + \beta_I X_R(k+1)X_I(k-1) + \beta_R X_I(k+1)X_I(k-1)) \end{aligned} \quad (\text{B.1.5})$$

The β_R and β_I terms in Eq (B.1.5) can be collected together to give

$$\begin{aligned} \frac{\partial C}{\partial \beta_R} = & 2\beta_R \left[(X_R(k+1) + X_R(k-1))^2 + (X_I(k+1) + X_I(k-1))^2 \right] \\ & - 4\beta_I (X_I(k+1)X_R(k-1) - X_R(k+1)X_I(k-1)) \\ & - 2 \left(X_R(k)(X_R(k+1) + X_R(k-1)) \right. \\ & \left. + X_I(k)(X_I(k+1) + X_I(k-1)) \right) \end{aligned} \quad (\text{B.1.6})$$

Since β_R and β_I are the variables with which C is being varied, Eq (B.1.6) can be

rewritten as

$$\begin{aligned} \frac{\partial C}{\partial \beta_R} &= 2A\beta_R - 4B\beta_I - 2E, \quad \text{where} \\ A &= (X_R(k+1) + X_R(k-1))^2 + (X_I(k+1) + X_I(k-1))^2, \\ B &= X_I(k+1)X_R(k-1) - X_R(k+1)X_I(k-1), \\ E &= X_R(k)(X_R(k+1) + X_R(k-1)) + X_I(k)(X_I(k+1) + X_I(k-1)). \end{aligned} \quad (\text{B.1.7})$$

Equation (B.1.7) is mentioned in section 2.8.1, the derivation of which is the subject of this appendix. It is one equation which involves two unknown parameters: β_R and β_I . In order to find another equation involving these unknowns the first partial derivative of C with respect to β_I must be taken. This produces

$$\begin{aligned} \frac{\partial C}{\partial \beta_I} &= i(X^*(k+1)X(k) - X^*(k-1)X(k) - X(k+1)X^*(k) + X(k-1)X^*(k)) \\ &\quad - (\beta_R + i\beta_I)|X(k+1)|^2 + (\beta_R - i\beta_I)|X(k+1)|^2 + (\beta_R - i\beta_I)|X(k-1)|^2 \\ &\quad - (\beta_R + i\beta_I)|X(k-1)|^2 + 2(\beta_R + i\beta_I)X(k+1)X^*(k-1) \\ &\quad - 2(\beta_R - i\beta_I)X^*(k+1)X(k-1)) \\ &= -2\Im\{X^*(k+1)X(k) + X(k-1)X^*(k) + (\beta_R - i\beta_I)|X(k+1)|^2 \\ &\quad + (\beta_R - i\beta_I)|X(k-1)|^2 \\ &\quad + 2(\beta_R + i\beta_I)X(k+1)X^*(k-1)\} \end{aligned} \quad (\text{B.1.8})$$

Taking the imaginary parts of Eq (B.1.8) gives

$$\begin{aligned} \frac{\partial C}{\partial \beta_I} &= -2(X_R(k+1)X_I(k) - X_I(k+1)X_R(k) - X_R(k-1)X_I(k) \\ &\quad + X_I(k-1)X_R(k)) + 2\beta_I(|X(k+1)|^2 + |X(k-1)|^2) \\ &\quad + 4(\beta_R X_R(k+1)X_I(k-1) - \beta_I X_I(k+1)X_I(k-1) \\ &\quad - \beta_I X_R(k+1)X_R(k-1) - \beta_R X_I(k+1)X_R(k-1)) \end{aligned} \quad (\text{B.1.9})$$

In a similar step to before, the β_R and β_I terms in Eq (B.1.9) can be collected together to give

$$\begin{aligned} \frac{\partial C}{\partial \beta_I} &= 2\beta_I \left[(X_R(k+1) - X_R(k-1))^2 + (X_I(k+1) - X_I(k-1))^2 \right] \\ &\quad - 4\beta_R (X_I(k+1)X_R(k-1) - X_R(k+1)X_I(k-1)) \\ &\quad - 2 \left(X_R(k)(X_I(k-1) - X_I(k+1)) \right. \\ &\quad \left. + X_I(k)(X_R(k+1) - X_R(k-1)) \right) \end{aligned} \quad (\text{B.1.10})$$

Also in a similar stage, Eq (B.1.10) can be rewritten as

$$\begin{aligned}\frac{\partial C}{\partial \beta_I} &= 2D\beta_I - 4B\beta_R - 2F, \quad \text{where} \\ D &= (X_R(k+1) - X_R(k-1))^2 + (X_I(k+1) - X_I(k-1))^2, \\ F &= X_R(k)(X_I(k-1) - X_I(k+1)) + X_I(k)(X_R(k+1) - X_R(k-1)).\end{aligned}\tag{B.1.11}$$

Equation (B.1.11) is also mentioned in section 2.8.1, again the derivation is the subject of this appendix.

It is possible to combine the two partial derivatives with two unknowns to find a stationary point solution. This is possible because there are no cross-terms between β_R and β_I in Eqs (B.1.7) or (B.1.11). The stationary point is given by

$$(\beta_R, \beta_I) = \left(\frac{2BF + DE}{AD - 4B^2}, \frac{2BC + AF}{AD - 4B^2} \right)\tag{B.1.12}$$

To uncover the nature of this stationary point it is necessary to inspect the second partial derivatives, which are

$$\begin{aligned}\frac{\partial^2 C}{\partial \beta_R^2} &= 2A \\ \frac{\partial^2 C}{\partial \beta_R \partial \beta_I} &= -4B \\ \frac{\partial^2 C}{\partial \beta_I^2} &= 2D \\ \frac{\partial^2 C}{\partial \beta_I \partial \beta_R} &= -4B\end{aligned}$$

This generates the Hessian matrix

$$H = \begin{bmatrix} 2A & -4B \\ -4B & 2D \end{bmatrix}$$

This is tested for definiteness, which provides information on the nature of the stationary point

$$\begin{aligned}\beta^t H \beta &= \begin{bmatrix} \beta_R & \beta_I \end{bmatrix} \begin{bmatrix} 2A & -4B \\ -4B & 2D \end{bmatrix} \begin{bmatrix} \beta_R \\ \beta_I \end{bmatrix} \\ &= 2A\beta_R^2 - 8B\beta_R\beta_I + 2D\beta_I^2 \\ &= 2A \left(\beta_R - \frac{4B}{2A}\beta_I \right)^2 + \left(2D - \frac{16B^2}{2A} \right)\end{aligned}\tag{B.1.13}$$

The only part of Eq (B.1.13) which can be negative is the second bracket. This is

because squares are always non-negative and therefore the first bracket is always non-negative and the β_I^2 is always non-negative. In addition to this the definition of A in Eq (B.1.7) is composed of a sum of squares (also the definition of D in Eq (B.1.11) is composed of a sum of squares). Therefore the relationship between $4AD$ and $16B^2$ determines the nature of the stationary point. To be a minimum point $AD > 4B^2$ must be satisfied. If the second bracket is expressed in terms of its original constituent constants this gives

$$\begin{aligned} AD - 4B^2 = & \left[(X_R(k+1) + X_R(k-1))^2 + (X_I(k+1) + X_I(k-1))^2 \right] \\ & \times \left[(X_R(k+1) - X_R(k-1))^2 + (X_I(k+1) - X_I(k-1))^2 \right] \\ & - 4(X_I(k+1)X_R(k-1) - X_R(k+1)X_I(k-1))^2 \end{aligned} \quad (\text{B.1.14})$$

The smallest possible value for the function in Eq (B.1.14) is zero and occurs if $X_I(k+1) = X_R(k-1)$ or if $X_R(k+1) = X_I(k-1)$. When that happens the surface given by Eq (B.1.13) will degenerate from a bowl to a valley. The Hessian matrix in Eq (B.1) is positive semidefinite and therefore the stationary point will always be a minimum or, in limiting cases, a valley.

B.2 Lagrangian Multiplier complex SVA

It is stated in section 2.8.1 that the minimum found in Eq (B.1.12) does not necessarily correspond to a time-domain raised-cosine window function that is always non-negative because the minimisation was not constrained to a particular set of answers. A simple extension of the minimisation technique is to use a Lagrangian multiplier with the constraint $\alpha = 0.5$ or equivalently $\beta_R^2 + \beta_I^2 = 0.25$. The Lagrangian is then given by

$$L(\beta_R, \beta_I, \lambda) = C(\beta_R, \beta_I) - \lambda (0.25 - \beta_R^2 - \beta_I^2) \quad (\text{B.2.1})$$

The partial derivatives of the Lagrangian from Eq (B.2.1) are given by

$$\frac{\partial L}{\partial \beta_R} = 2A\beta_R - 4B\beta_I - 2E + 2\lambda\beta_R \quad (\text{B.2.2a})$$

$$\frac{\partial L}{\partial \beta_I} = 2D\beta_I - 4B\beta_R - 2F + 2\lambda\beta_I \quad (\text{B.2.2b})$$

$$\frac{\partial L}{\partial \lambda} = \beta_R^2 + \beta_I^2 - 0.25 \quad (\text{B.2.2c})$$

At the constrained minimum all of the partial derivatives will be equal to zero. If Eq (B.2.2a) and Eq (B.2.2b) are set equal to zero and combined to eliminate λ the result is

$$(D - A)\beta_R\beta_I - 2B(\beta_R^2 - \beta_I^2) + E\beta_I - F\beta_R = 0 \quad (\text{B.2.3})$$

Since $(D - A)$ is just another constant, it can be defined as

$$G = 1/4(D - A) = -(X_R(k+1)X_R(k-1) + X_I(k+1)X_I(k-1)) \quad (\text{B.2.4})$$

If Eq (B.2.4) is substituted into Eq (B.2.3) and the complex co-ordinate system is changed to polar co-ordinates using $\beta_R = \alpha \cos(\phi)$ and $\beta_I = \alpha \sin(\phi)$, the simplified result is

$$2\alpha G \sin(2\phi) - 2\alpha B \cos(2\phi) + E \sin(\phi) - F \cos(\phi) = 0 \quad (\text{B.2.5})$$

Finally if the last partial derivative of the Lagrangian in Eq (B.2.2c) is used, which is essentially the same as substituting $\alpha = 0.5$ into Eq (B.2.5), the following is obtained

$$G \sin(2\phi) - B \cos(2\phi) + E \sin(\phi) - F \cos(\phi) = 0 \quad (\text{B.2.6})$$

Equation (B.2.6) has to be solved numerically as there is no analytical solution.

B.3 Constrained complex SVA

From the equations of the simplistic Lagrangian multiplier constraint it is natural to consider relaxing the constraint to the desired constraint of $\alpha \leq 0.5$ and using the KKT conditions to find a constrained solution. The necessary KKT conditions are

$$\left. \begin{aligned} \frac{\partial C(\beta_R, \beta_I)}{\partial \beta_R} + \lambda \frac{\partial g(\beta_R, \beta_I)}{\partial \beta_R} &= 0 \\ \frac{\partial C(\beta_R, \beta_I)}{\partial \beta_I} + \lambda \frac{\partial g(\beta_R, \beta_I)}{\partial \beta_I} &= 0 \\ \lambda(0.25 - g(\beta_R, \beta_I)) &= 0 \\ g(\beta_R, \beta_I) &\leq 0.25 \\ \lambda &\leq 0 \end{aligned} \right\} \quad \text{KKT conditions} \quad (\text{B.3.1})$$

where $g(\beta_R, \beta_I) = \beta_R^2 + \beta_I^2 = 0.25$ describes the upper limit of the constraint.

As with the Lagrangian multiplier technique there is no analytical solution for the series of equations in the KKT conditions, however it is possible to show that the stationary point that can be found numerically is a minimum. This is from consideration of the bordered-Hessian matrix given by

$$H(C, g) = \begin{bmatrix} 0 & \frac{\partial g}{\partial \beta_R} & \frac{\partial g}{\partial \beta_I} \\ \frac{\partial g}{\partial \beta_R} & 2A & -4B \\ \frac{\partial g}{\partial \beta_I} & -4B & 2D \end{bmatrix} = 2 \begin{bmatrix} 0 & \beta_R & \beta_I \\ \beta_R & A & -2B \\ \beta_I & -2B & D \end{bmatrix} \quad (\text{B.3.2})$$

The determinant of the bordered-Hessian matrix in Eq (B.3.2) is given by

$$\det(H) = -8 (\beta_R^2 D + \beta_I^2 A + 4B\beta_I\beta_R) \quad (\text{B.3.3})$$

When Eq (B.3.3) is expressed as squares, the following is obtained

$$\det(H) = -8D \left(\beta_R + \frac{16B}{8D}\beta_I \right)^2 + \left(-8A + \frac{256B^2}{8D} \right) \beta_I^2 \quad (\text{B.3.4})$$

The first term in Eq (B.3.4) is always negative due to the positive nature of squares. The first bracket is a square and D is a sum of squares. Therefore the second term in Eq (B.3.4) determines whether the determinant is positive, negative or zero. From inspection of the second bracket in Eq (B.3.4) it is clear that to be negative the condition $-64AD < 256B^2$ must apply. This implies the condition $AD > 4B^2$, which is the same result as for the unconstrained minimisation case and therefore the same conclusions can be drawn about the nature of the stationary point: the stationary point is always a minimum unless $X_I(k+1) = X_R(k-1)$ or $X_R(k+1) = X_I(k-1)$ in which case it will be a minimum in a valley. Examples of the use of this constrained minimisation are given in section 2.8.3.

APPENDIX C

Model-Based Design

C.1 Introduction

This section aims to explain the context in which the research for this thesis took place. That is, the context in which modern electronic systems are built.

The fast pace of technological and market changes in the electronics industry have stimulated the development of tools and design methodologies to help minimise the risks involved in this market. These tools allow designs to be created at a high-level of abstraction, where reusable intellectual property (IP) blocks are integrated together at a system level.

Tools which place a greater emphasis on the software-modelling-domain, rather than the hardware-domain, have an advantage that they can be more easily changed and applied to re-configurable devices like FPGAs many times over. The design methodologies used tend to be graphical and model-based as opposed to the traditional waterfall design methodology where the design process is strictly sequential as depicted in Fig C.1.1.

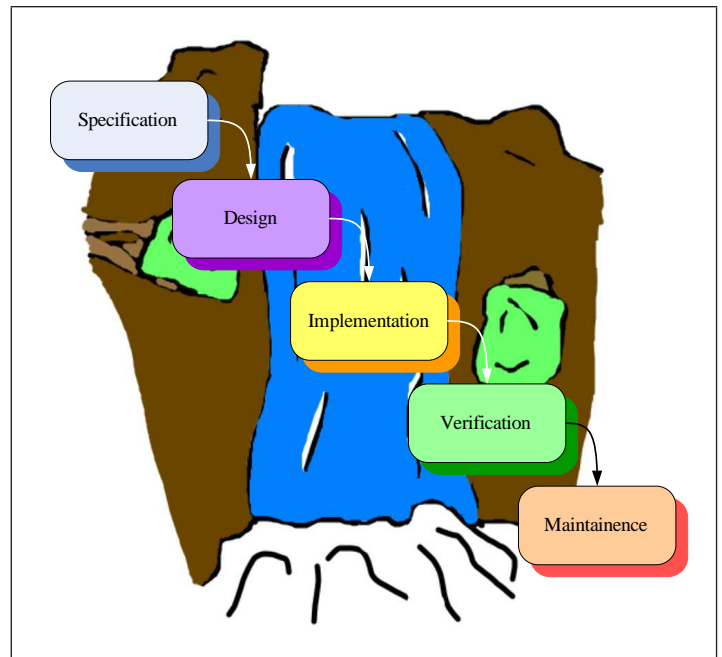


Figure C.1.1 The waterfall design methodology

C.2 MBD Principles

MBD tools invariably involve the creation of a system from building blocks which represent advanced functionality. These blocks are built up to form the system in a hierarchy. Thus it is possible to work at a very high level of abstraction to get an overall feel for the system design or alternatively it is possible to drill-down and work at a lower level of abstraction as desired.

At a high-level of abstraction a graphical model is an important trait when systems are

large and complicated. If, for example, traditional text-based editors or mathematical equation modelling tools were used, systems can become very complicated and hard to decipher. It would be difficult to conceptualise a whole system at once with a text-based editor or mathematical equation, however it is easy to get a grasp of a system with a top-level graphical model diagram. Also the fine control over code provided by a text-based editor is not entirely lost in MBD because the lowest level blocks in the hierarchy can be explored and since these block simply represent code they can be viewed in a text-editor if desired.

The modular nature of MBD aids designers to reduce development time by the facilitation of design re-use of various IP blocks in different designs. This fits in neatly with the rise of re-configurable hardware such as FPGAs to further reduce design times. For commercial businesses the reduction in design time is important as it allows them to meet time-to-market windows more successfully, especially when the trend is for these windows to shrink.

The link with FPGAs in MBD is a strong one. MBD permits the involvement of hardware at a very early stage in the design development. The process is called hardware-in-the-loop. The inclusion of hardware early in the design process allows errors which would only appear once the design had been transferred to hardware to be located promptly. Therefore costs which are involved in designs changes occur towards the start of design development and hence are smaller than they would have otherwise been. It also allows the designer to draw some rough estimates on the characteristics of the design in the hardware such as its footprint or speed early in the design development. The involvement of hardware early in the design development is probably one of the greatest advantages of MBD and the overall effect is a large reduction in the test and verification time. This is particularly true of MBD tools where it is possible to run a hardware co-simulation. Hardware co-simulation involves design execution in simulation in the modelling tool **and** on the FPGA at the same time such that any discrepancies between the two are apparent. More information on MBD and its application in DSP can be found in references such as [138]. The MBD tools and methodology used in the course of the research in this thesis are now be described.

C.3 SELEX Galileo MBD Methodology

The MBD methodology, for the purposes of algorithm research, which is promoted in SELEX Galileo and followed in this thesis involves targeting FPGAs with the use of: (1) MATLAB[®], (2) Simulink[®] and (3) Xilinx[®] System Generator[®] tools. In short, it involves the translation from the highest-level golden reference to FPGA through the addition of detail at each translation.

Fig C.3.1 describes the methodology. The feedback arrows on the left of the diagram show that the models are verified against previous stages to prevent error propagation

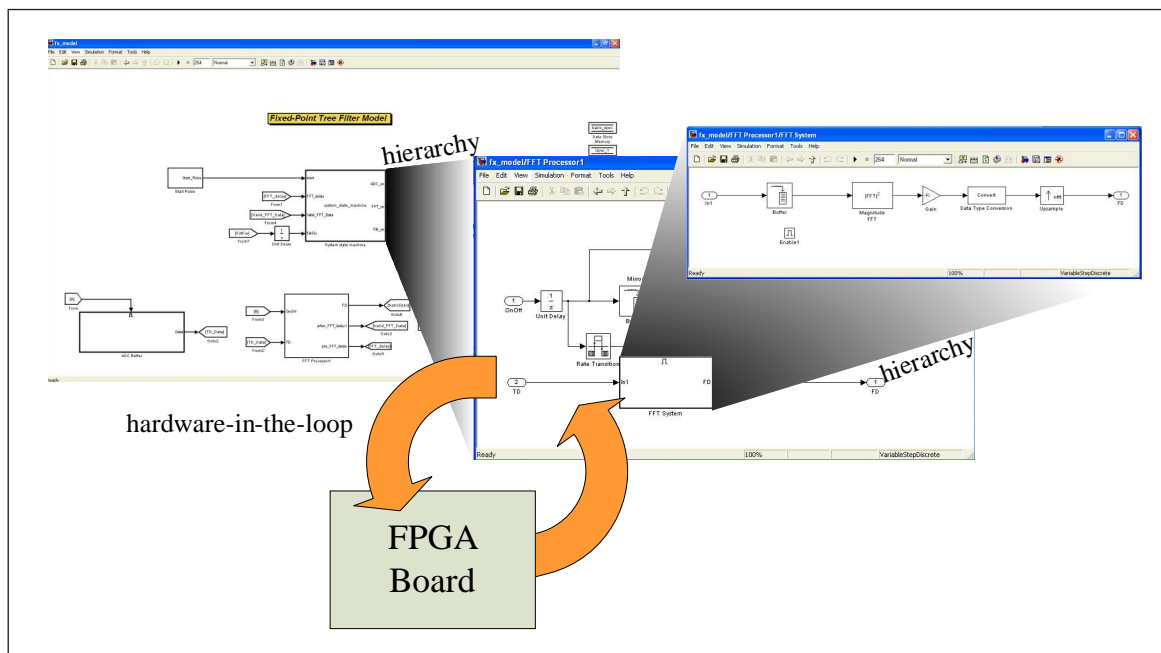


Figure C.2.1 The MBD methodology

down the design flow. The numbers on the right-hand side show a ranking of relative level of abstraction, 1 denotes the highest level and 3 the lowest. The circular arrows at the bottom-right of the diagram indicate the iterative refinement of the FPGA design after inclusion of hardware in the loop. The golden reference is located at the

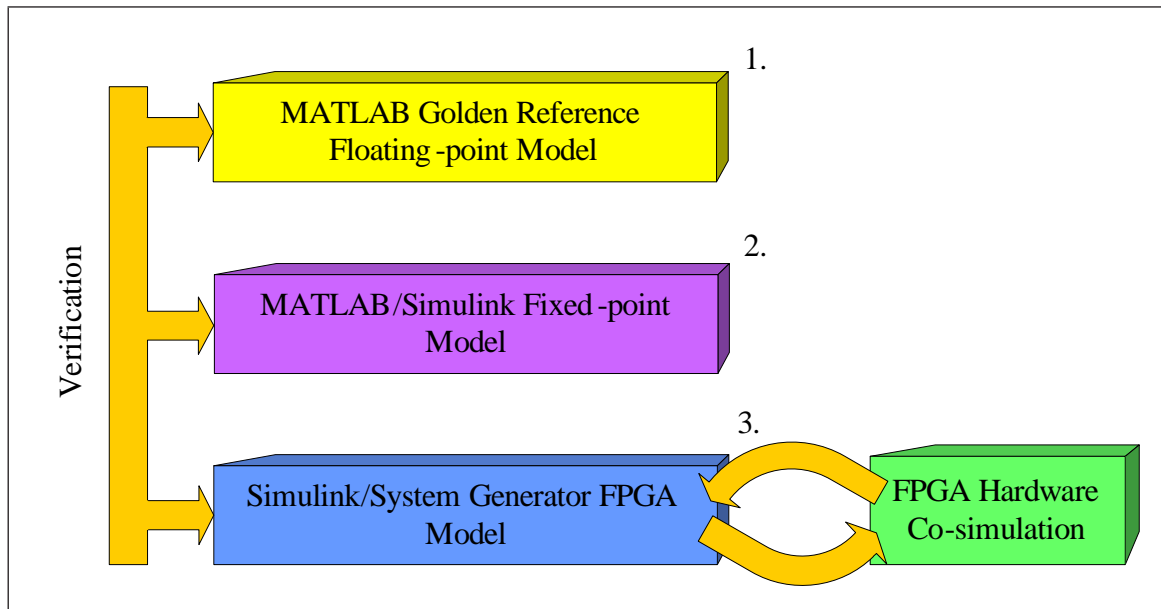


Figure C.3.1 The SELEX Galileo MBD methodology

highest level of abstraction. As a model it captures all system requirements and is the model against which all lower level models are verified. It is built with the use of MATLAB[®] functions and scripts in floating-point arithmetic which provides a high degree of accuracy for a low amount of effort.

For the next model down in abstraction, the golden reference is translated into a

MATLAB[®]/Simulink[®] model which utilises fixed-point arithmetic. Typically, with a real-life system implementation in mind, there will be restrictions on the number of bits of precision which the fixed-point model can use. The construction of the fixed-point model within these limitations makes it possible to get an understanding of the error level achievable between the floating-point golden reference and the fixed-point model and whether or not the algorithm will remain useful in a real-life system in fixed-point precision on FPGA. It is important at this stage to design the fixed-point model as well as possible in order to ease the translation to the next stage of abstraction. One way of doing this is to have a good knowledge of the target hardware. For example what mathematical operations are possible and how the fixed-point arithmetic is handled by the hardware. A good knowledge of the target hardware allows the fixed-point model to be designed with the hardware in mind.

Provided the fixed-point model was well-designed, the translation to the next level of abstraction should be straightforward. The model here consists of a set of building blocks which represent basic operations and functions that can be carried out on the target hardware. The preferred tools for this task are Simulink[®]/Xilinx[®] System Generator[®]. The System Generator[®] blocks are the ones which visually represent the FPGA operations. Everything coded in these blocks can be literally translated to FPGA. Alternative tools include Synplify DSP[®] from Synplicity[®] and DSP Builder[®] from Altera[®] if the target hardware is not Xilinx[®]. The results from this FPGA model may differ from the previous fixed-point model due to the differences in which certain higher-level IP blocks are implemented. For example division may be coded with the use of one particular algorithm in fixed-point and coded with the use of a different algorithm in the FPGA model which would lead to slight discrepancies between the two models. However as long as both models can be verified against the golden reference this should not matter.

The use of System Generator[®] allows for continuous test and verification of the algorithm actually running on the Xilinx[®] target hardware through a hardware co-simulation mode.

C.4 System Generator[®]

System Generator[®] essentially abstracts hardware description languages (HDLs) to a high-level and takes care of the design flow through the invocation of other Xilinx[®] ISE design suite tools so that the inner workings of the production of an FPGA bit file can be treated as a black box process. A list of some of its benefits are described here:

- ability to synthesise HDL code in both Verilog and VHDL languages;
- fits in well with the Xilinx[®] ISE design suite;
- includes libraries of IP which can perform specific functions such as division or

coordinate rotation digital computer (CORDIC);

- can import other core IP as HDL black boxes into System Generator®;
- can handle all the FPGA design stages such as synthesis, map and place and route to create an FPGA downloadable bit file through invocation of various tools from the Xilinx® ISE design suite;
- the FPGA model built with System Generator® blocks creates a model which is cycle-true and bit-true to the actual FPGA implementation;
- can specify exactly the fixed-point precision and overflow handling at each stage of the algorithmic data flow;
- algorithmic verification can be performed using hardware-in-the-loop.

The overall design flow followed by System Generator from start to finish is shown in Fig C.4.1. Perhaps the most important point from the above list is the hardware-in-the-loop ability. Such is the support from System Generator® for hardware co-simulation that the results from the system implementation on the FPGA can be displayed alongside the results from any other model of the system from the MBD methodology shown in Fig C.3.1, hence the overall design flow from golden reference to FPGA is accelerated considerably and the FPGA implementation is easily verified.

Lack of hardware co-simulation is the primary reason why the HDL Coder® tool from Mathworks® was not chosen as the MBD tool for the work in this thesis. Naturally for good hardware co-simulation with the use of a Xilinx® FPGA, it is easier to work with the Xilinx® tool because then it is not necessary to work with device drivers manually. Hence the choice of System Generator® for the work in this thesis. An example of how to set up a hardware co-simulation with System Generator® and a Xilinx® FPGA is described in [139].

Other methodologies exist to map to FPGAs such as hand-coding directly in VHDL or Verilog or the direct use of CORE generator. However the disadvantages of these methods are more severe than the disadvantages of System Generator®. For example, hand-coding in VHDL or Verilog can be very time-consuming and proper usage of CORE generator requires a good understanding of the signals and parameters in order to match them to your specifications.

For completeness, the disadvantage of System Generator® should be mentioned: the tool may not always use the minimum possible FPGA resources (however this is a small price to pay for a high level of automation in direct FPGA map design).

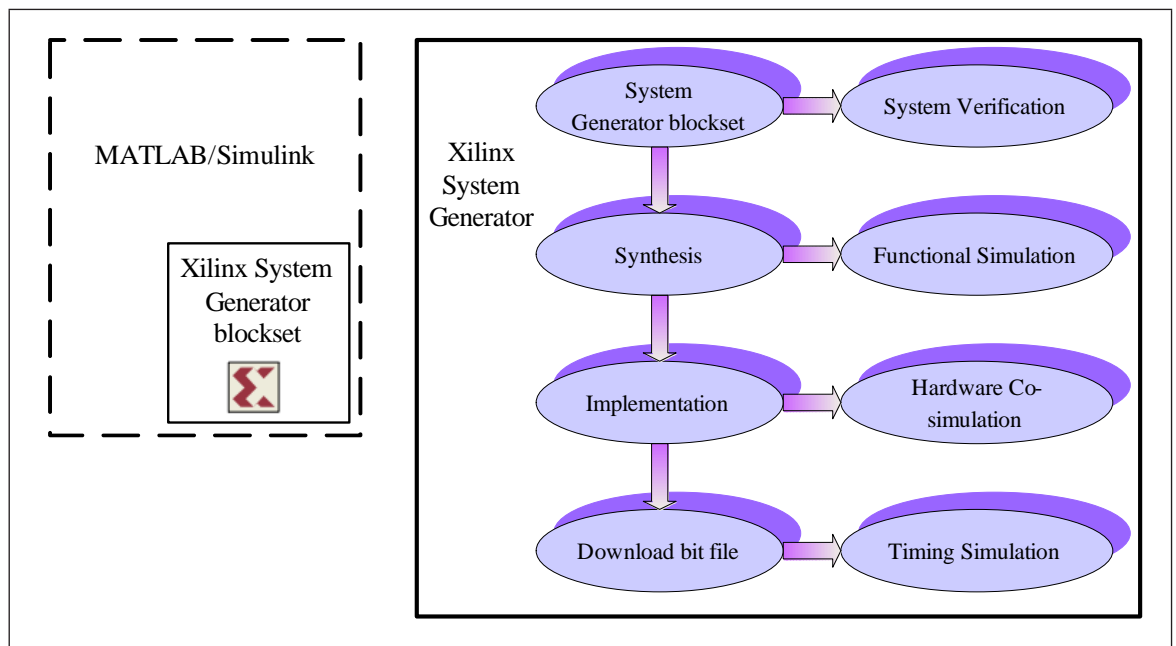


Figure C.4.1 System Generator® design flow from blockset to bitstream

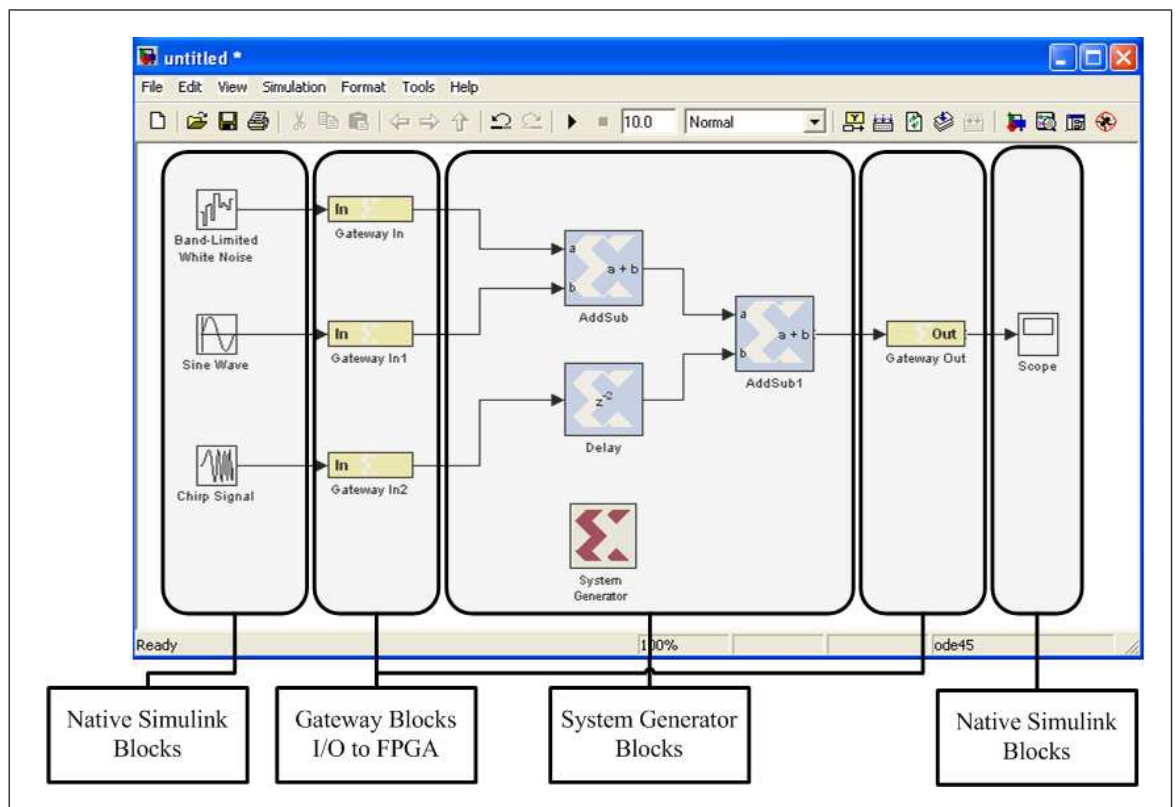


Figure C.4.2 Example System Generator® system

The relationship between Simulink[®], System Generator[®] and FPGA hardware becomes clear upon examination of Fig C.4.2. This shows a made-up example system for illustrative purposes. Simulink[®] code (and Embedded MATLAB[®] code) exists in the Simulink[®] model window as various sources, process blocks and sinks. The System Generator[®] blocks which represent the IP and which eventually are downloaded onto an FPGA must exist within special “Gateway” blocks which act as interfaces between native Simulink[®] blocks and System Generator[®] blocks. The System Generator[®] blocks are prominent because they are light blue in colour and are labelled with the Xilinx[®] logo. The Gateway blocks also mark a data type conversion between Simulink[®] data types and the custom fixed-point data types on the FPGA. Essentially the Gateway blocks signify the boundaries of the FPGA. One other important point which concerns Fig C.4.2 is the actual System Generator[®] block represented by a red Xilinx logo. The purpose of this block is to set certain parameters for the target hardware and how far System Generator[®] should progress along the FPGA design flow from Fig C.4.1.

C.5 Xilinx ISE Suite

As mentioned in the previous section, System Generator[®] can take a design all the way from FPGA model to bit file download. It does so by the invocation of various Xilinx[®] ISE tools in the background. For example Xilinx[®] iMPACT is invoked in the course of the FPGA download process. If desired, the incremental design process shown in the bottom-left pane of Fig C.5.1 can be stepped through with Xilinx[®] ISE Project Navigator tool. The two main stages in the design flow of this tool are synthesis and implementation. The reports are also selectable for view in the tool (see the middle pane and right-hand pane of Fig C.5.1).

C.5.1 Synthesis Stage

The Project Navigator tool creates a synthesis report, a synthesis errors report, a warnings report and a summary upon completion of the synthesis stage. After completion of the synthesis stage, it is possible to see the register transfer level (RTL) diagrams of any part of the design. The reason for a well planned System Generator[®] model becomes clear from inspection of the RTL diagrams because the Xilinx[®] ISE tools tend to replicate the partitions and functional blocks in RTL. For example, with the aid of the model from appendix F, there is a striking similarity between the System Generator[®] model from Fig F.2.6 and the RTL which describes it in Fig C.5.2.

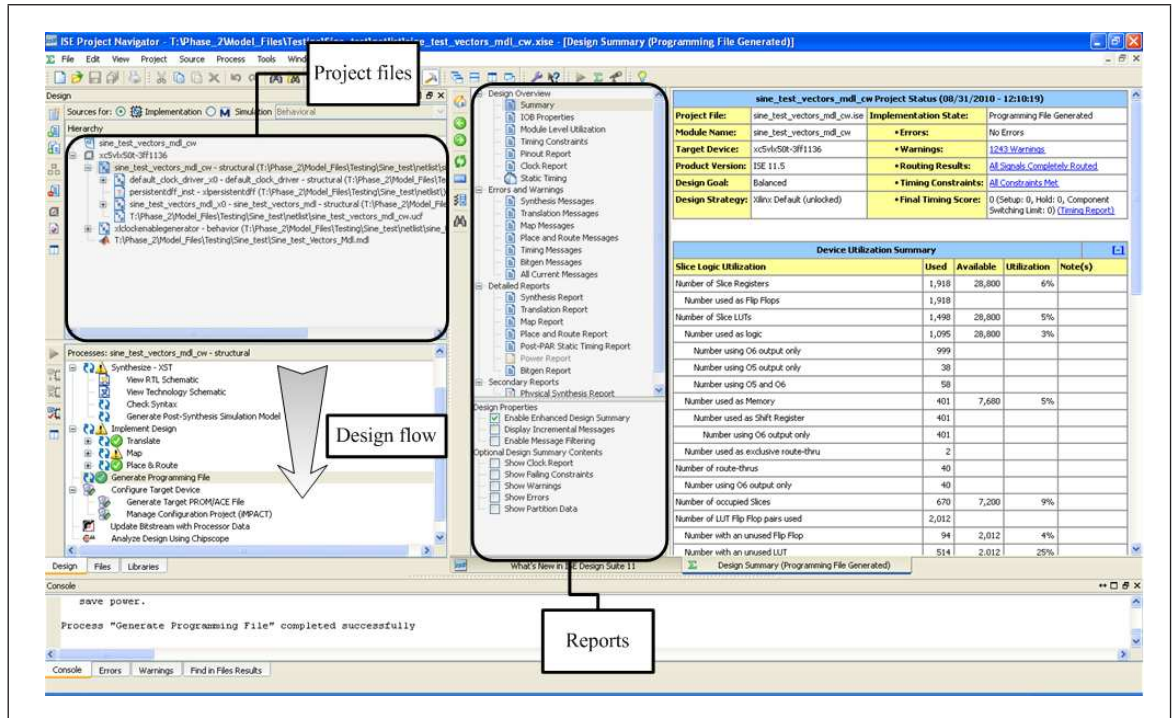


Figure C.5.1 Xilinx® ISE Project Navigator in use

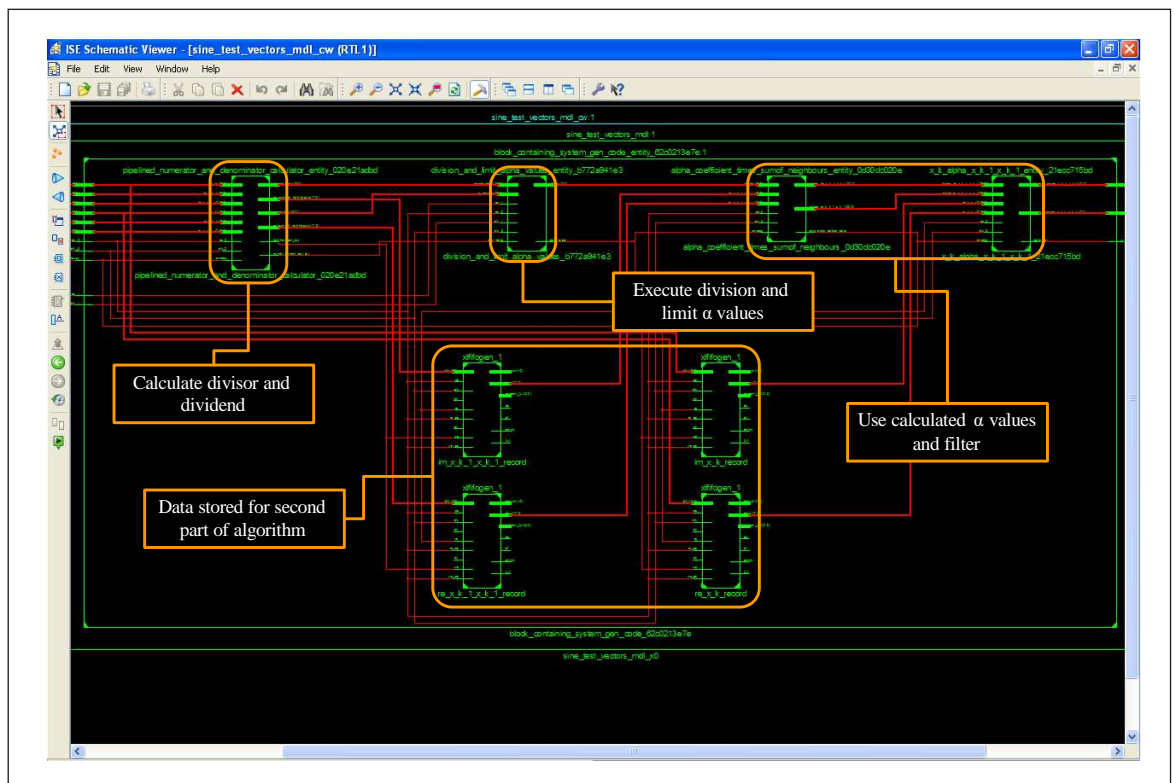


Figure C.5.2 RTL diagram of phase 2 SVA model. Labels have been added for clarity

C.5.2 Implementation Stage

This stage maps synthesised code to target hardware and performs a place-and-route on the design. Upon completion of this stage the following reports and summaries are created:

- translation report and messages
- map report and messages
- place and route report and messages
- post-PAR Static Timing report and timing messages
- IOB properties
- module level utilization
- timing constraints
- pinout report
- clock report.

Any part of the placed and routed design can be viewed through the use of Xilinx® FPGA editor. For example, Fig C.5.3 shows part of the phase 2 model from appendix F after the place and route stage in the FPGA editor. This example is highly zoomed-in, and it is difficult to attribute any specific structure to the SVA algorithm at this low level. However it is possible to see an individual DSP48E embedded multiplier as a red box surrounded by Xilinx® slices to the left and right, which are shown as transparent boxes. The route between slices is visible as light-blue lines and the route between DSP48E embedded multipliers is visible as green lines.

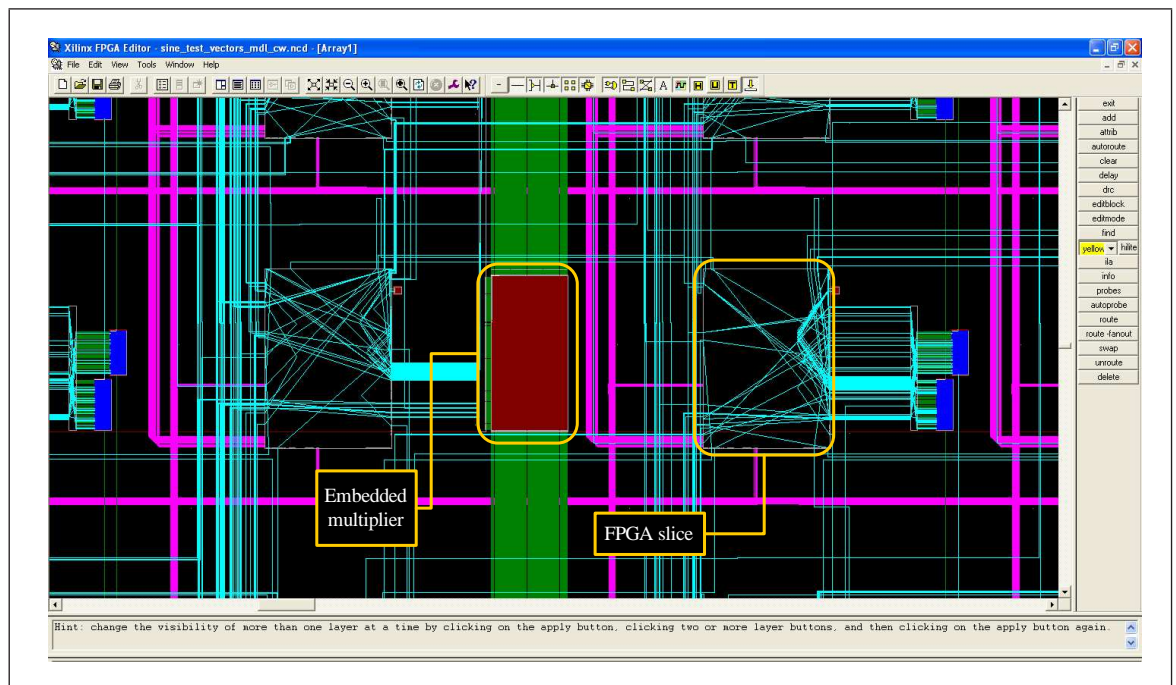


Figure C.5.3 Close-up view of FPGA after the place and route of phase 2 model

APPENDIX D

Fixed-Point SVA

The results obtained using a floating-point software model of SVA in a simulated DRx in chapter 3 indicated that SVA is a promising technique which enables DRx algorithm designers to avoid the usual limitations imposed when using fixed-weight window functions. These results are summarised in Table 3.6.3. In all the scenarios tested, SVA performed better than or almost as well as the best conventional windowing technique. In scenarios dominated by sidelobe spectral leakage, the detectors which utilised SVA came out on top, whilst in scenarios dominated by the noise floor, SVA approached the rectangular window performance. This is important because the conventional approach to DFT processing in DRxs designed for EW is not to use a rectangular window but rather to use a moderate-dynamic range window function like a Hann window function, despite the associated increase in NEB and wider 3.0-dB and 6.0-dB mainlobe. Since SVA reduces sidelobes, mitigates the windowing loss (i.e. the increase in NEB) and maintains the rectangular-window frequency resolution, it is definitely advantageous to use it over conventional moderate-dynamic range windows.

It was appropriate to continue the floating-point work by taking the next step in the MBD methodology. This next step involved an SVA implementation working within stricter, more realistic, fixed-point constraints. This determined whether the advantages offered by SVA as a high-level floating-point algorithm were still present in implementable fixed-point arithmetic.

D.1 Implementation Decisions

At the fixed-point stage in any MBD methodology it is useful to make decisions concerning the implementation with some foresight because if the algorithm is shaped towards eventual FPGA implementation at this stage, it minimises the difficulty at subsequent stages in the MBD methodology. There were two particularly important decisions made and these were: (a) what language to use to describe the algorithm in fixed-point and (b) what hardware the algorithm should eventually target. Factors which affected the decisions were:

- the software tools available;
- the prescribed SELEX Galileo MBD methodology;

- the FPGA resources available.

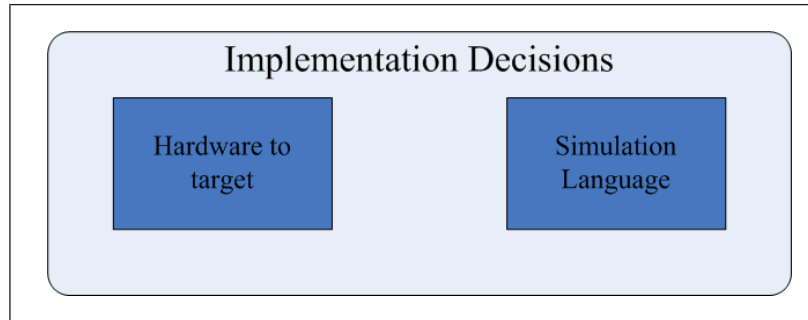


Figure D.1.1 Implementation decisions

D.1.1 Target Hardware

Although writing the SVA algorithm to be completely hardware-independent would have its merits, the available FPGA hardware at the time within the group at SELEX Galileo was Xilinx[®]-manufactured which meant that the fixed-point code would be no worse-off written with a Xilinx[®] application in mind. Indeed, by keeping the Xilinx[®] target hardware in mind the code was able to take advantage of Xilinx[®]-specific features and optimisations later in the MBD methodology. The target hardware used was an experimental board from the Xilinx[®] Virtex-5[®] family. One of the attractive features of the Xilinx[®] Virtex-5[®] FPGA fabric is it incorporates so-called DSP48E slices. This feature allows designers to create fast DSP designs. Therefore the fixed-point code which was written at this stage in the MBD methodology emulated the basic arithmetic operations possible in the slices, which are:

- 48 bit input addition/subtraction (or equivalently two 24 bit additions/subtractions);
- multiply-accumulate (MACC) using 25×18 multiplication and addition/subtraction of result with 48 bit value;
- 18 bit barrel shifts;
- 48 bit counter;
- 25×18 bit two's complement multiplication.

More complicated arithmetic operations had to be built from these above basic operations. For example, the simulated detectors which calculated a variable component of threshold (detector C2 in section D.3 and detector C3 in section D.4) required a non-standard operation in the form of a square-root. The usual solution to this is to use the well-known square root iterative algorithm which consists of successive multiplication and subtractions, however as is described later, the simulations in this appendix used an alternative method which resulted in calculation of the square-root of a magnitude-squared in a more indirect fashion. This alternative method has an accuracy advantage

over the usual direct square-root iterative algorithm because the output results are not limited to half the number of useful bits as the input.

D.1.2 Simulation Language

There are many different options for which language to choose for writing any potential fixed-point algorithm with pros and cons for each. The main options which were considered are listed below:

The C language Coding in C would produce highly portable code because the arithmetic in American National Standards Institute (ANSI) C is very simplistic and makes no assumptions about the handling of fixed-point numbers. For example “floor” rounding, wrapping on overflow and fixed wordlengths are all used by default. However ANSI C does not natively provide for non-standard data types which were required in the fixed-point models such as 18 bit, 36 bit and 48 bit unusual wordlengths.

MATLAB[®] and Simulink[®] A fixed-point model coded using these tools would involve extensive use of the Fixed-Point Toolbox and Embedded Toolbox. Coding in MATLAB[®] and Simulink[®] using fixed-point data types would be the natural progression from the previous the floating-point simulations. It is also no accident that the MBD methodology advocated at SELEX Galileo advises a fixed-point MATLAB[®] implementation following on from a floating-point one. The Fixed-Point Toolbox is highly flexible and would allow the fixed-point arithmetic to emulate the arithmetic and wordlengths of a DSP48E slice. However the disadvantage of fixed-point MATLAB[®] is that it is still a fairly high level language and may not translate as easily as other languages to different hardware. In other words, it is less portable.

VHDL or other HDL The level of abstraction of this language is very low which would allow a great degree of control over the hardware. However coding using VHDL immediately would be a jump too far in the MBD methodology, which is designed to make incremental steps in abstraction.

On balance, the decision was made to code the fixed-point SVA algorithm in fixed-point MATLAB[®].

D.2 Coding Methodology

The task to code, simulate and verify the fixed-point SVA algorithm was approached by conversion of the modules in detectors A1, B1, C1 and E1 to handle fixed-point arithmetic. With this approach, each module could be tested in its own testbench to ensure functionality before the separate modules were integrated into one large simulation with its own accompanying testbench. Additionally, this approach allowed

considerations to be made to optimise each module according to its functionality. For example the fractional bit handling in each module was more easily formulated and monitored given the task of the module. When the modules were translated to fixed-point, there was also a literal translation of the new code generated using floating-point data types to further verify the fixed-point code functionality. As an example, Fig D.2.1 shows the fixed-point output of the SVA module could be checked against a floating-point translation of the same code for the purpose of verifying functionality. Panel (a) shows that after the magnitude-squared is taken of the SVA output the two different datatypes are almost on top of each other in the plot and more explicitly panel (b) shows the actual error of the fixed-point SVA module for this particular data input. The amplitude of a least significant bit (LSB) for the real or imaginary part of a frequency bin of the fixed-point output is overlaid for comparison¹.

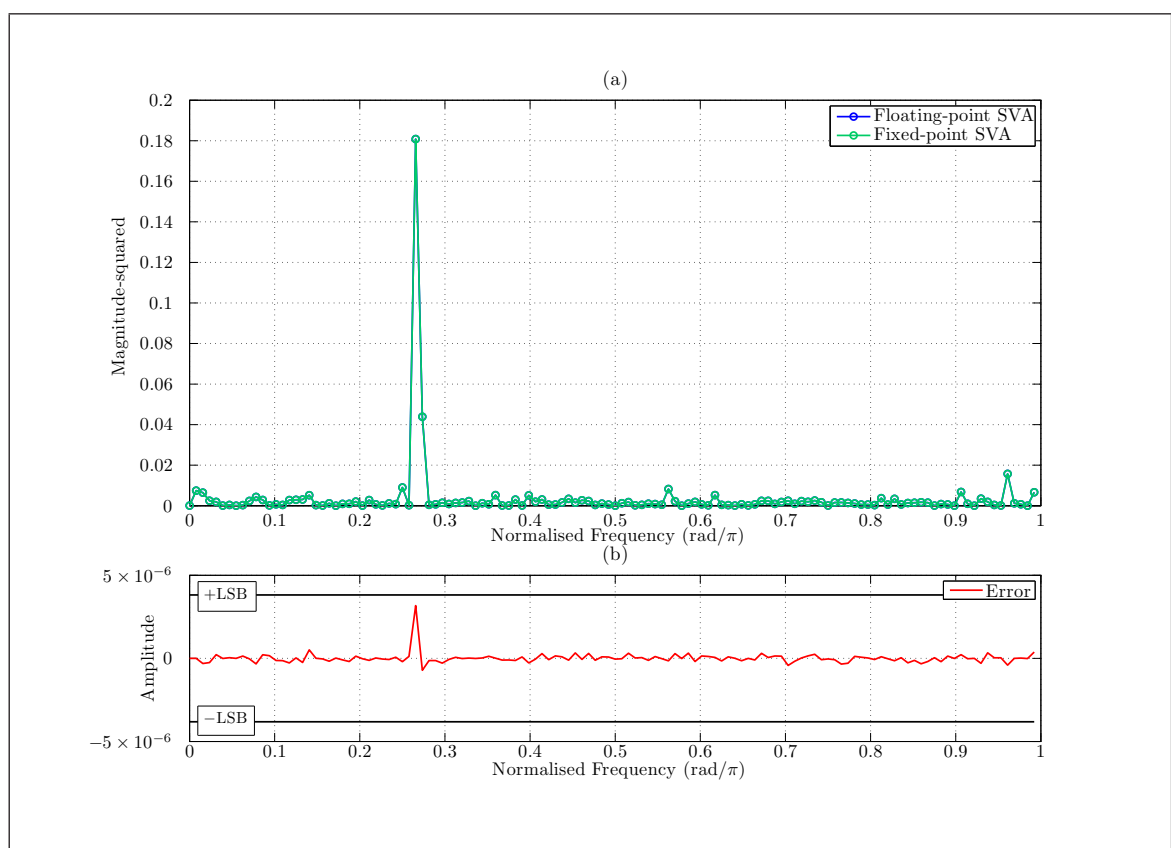


Figure D.2.1 The outputs from the SVA block: (a) the output of the SVA block run in fixed-point and in the floating-point translation; (b) the difference between the two outputs

The specifics of the fixed-point arithmetic used in the fixed-point models were:

Rounding Mode = Floor This was done as it is the most elementary and natural method of rounding a number for a fixed-point processor and therefore it is the most computationally efficient method of rounding. Essentially this involves solving the problem of a number with more fractional bits being cast to a number with fewer fractional bits by allowing the LSBs to “drop off” the LSB end.

¹For reference, a sinusoidal signal and AWGN at $\text{SNR}_{\text{in}} = 0 \text{ dB}$ were used in producing Fig D.2.1

Overflow Mode = Wrap When a maximum overflow error occurs, this is the simplest technique to deal with it by allowing significant bits to “drop off” the most significant bit (MSB) end. Therefore, as a quick and efficient technique, the wrap technique was used throughout.

Product Mode = Keep LSB This is another method which was chosen for its computational simplicity. If a product is formed from two inputs which cannot be represented in the multiplier register, then bits will “drop off” the MSB end. In section D.5.1 it is explained that this characteristic was put to good use by employing a common fixed-point arithmetic number representation.

Multiplier Register Size = 43 bit This unusual wordlength is the output wordlength of the multiplier in a DSP48E slice and therefore the fixed-point models reflected this.

Sum Mode = Keep LSB This method of dealing with a sum is computationally simple. If a sum is formed from two inputs which cannot be represented in the accumulator register, then under this method bits “drop off” the MSB end. Therefore this method, like the others, was chosen to maintain a simple fixed-point arithmetic scheme.

Accumulator Register Size = 48 bit This is the wordlength of the accumulator in a DSP48E slice and therefore was used in the fixed-point models.

Cast Before Sum = No This is the general method used in FPGA boards as a sum is usually cast to the accumulator register wordlength size after the summation operation has taken place and therefore the fixed-point models mimicked this characteristic.

D.3 Coding Phase 1

The actual coding of the fixed-point models took place in two phases so that the changes were more incremental and verification could take place more often. The first phase of conversion of the floating-point simulated DRx modules to fixed-point was to code the post-FFT processing in fixed-point. The modules before this partition, such as the DCW module and the FFT module, were left in the floating-point domain. As was the case during the tests of SVA in the purely floating-point models (or golden reference models), the models built during the first phase were such that the detector performance could be evaluated with and without the use of SVA. However the post-FFT processing was calculated in fixed-point. Essentially in the first phase of coding detectors A1, B1, C1 and E1 became detectors A2, B2, C2 and E2. Modules such as the DCW module and FFT module remained as floating-point code in anticipation of the FFT being a complicated module to code in fixed-point and to provide some interim results. Detectors A2, B2, C2 and E2 are shown in Figs D.3.1–D.3.4 respectively. As

mentioned in the preceding section, when the modules were converted to fixed-point a literal floating-point translation of the code was included for verification purposes. Therefore in Figs D.3.1–D.3.4 modules contained within the domain labelled as “Fixed-point” could also be run in floating-point as an extra check on their functionality.

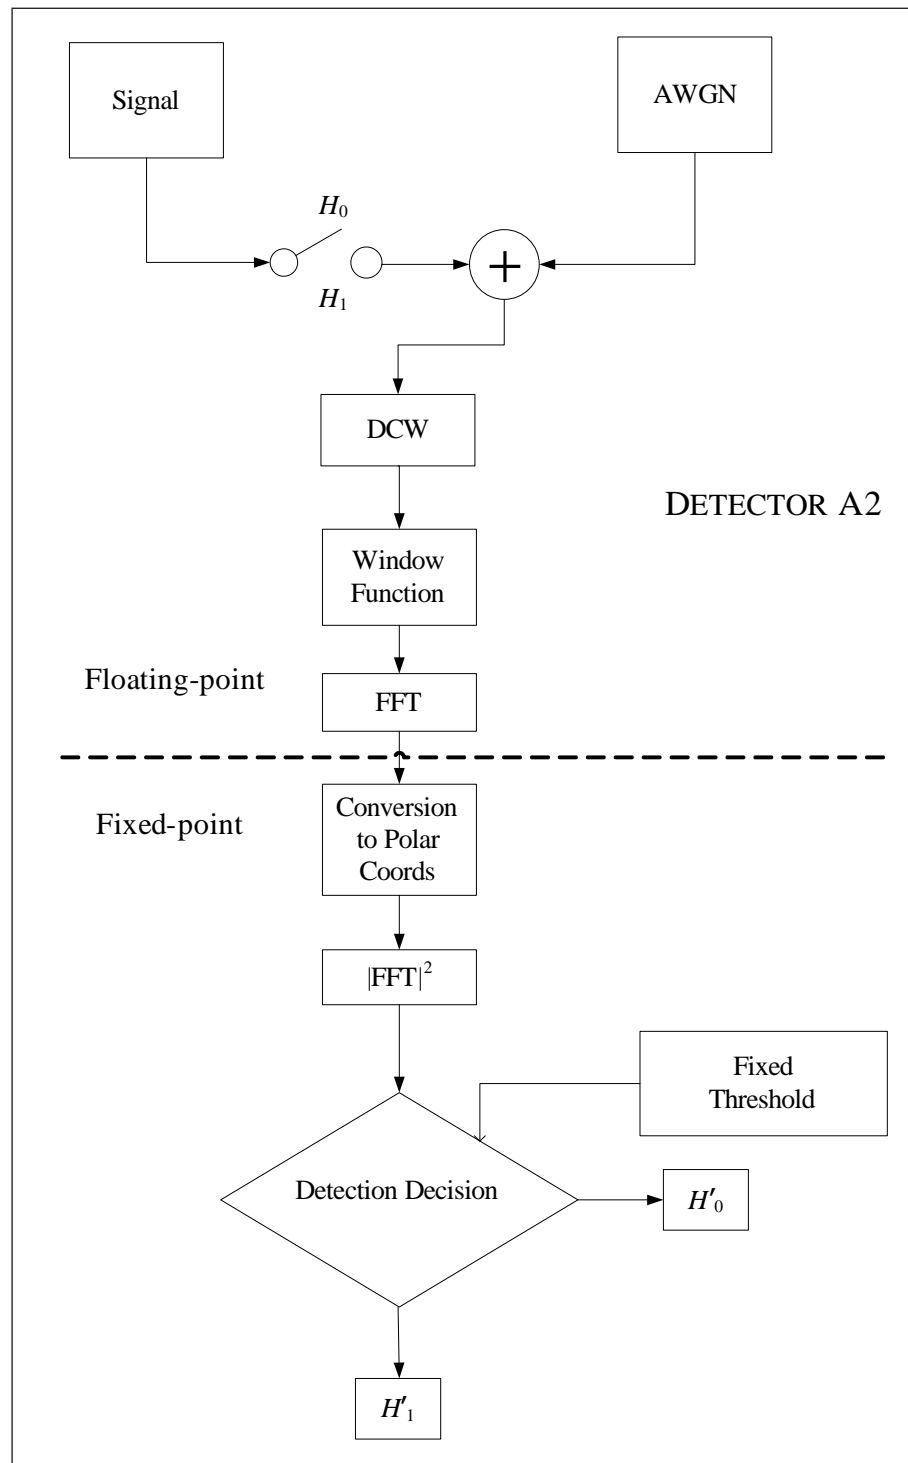


Figure D.3.1 Architecture of the conventional windowing detector A2. The partition between floating-point and fixed-point is shown

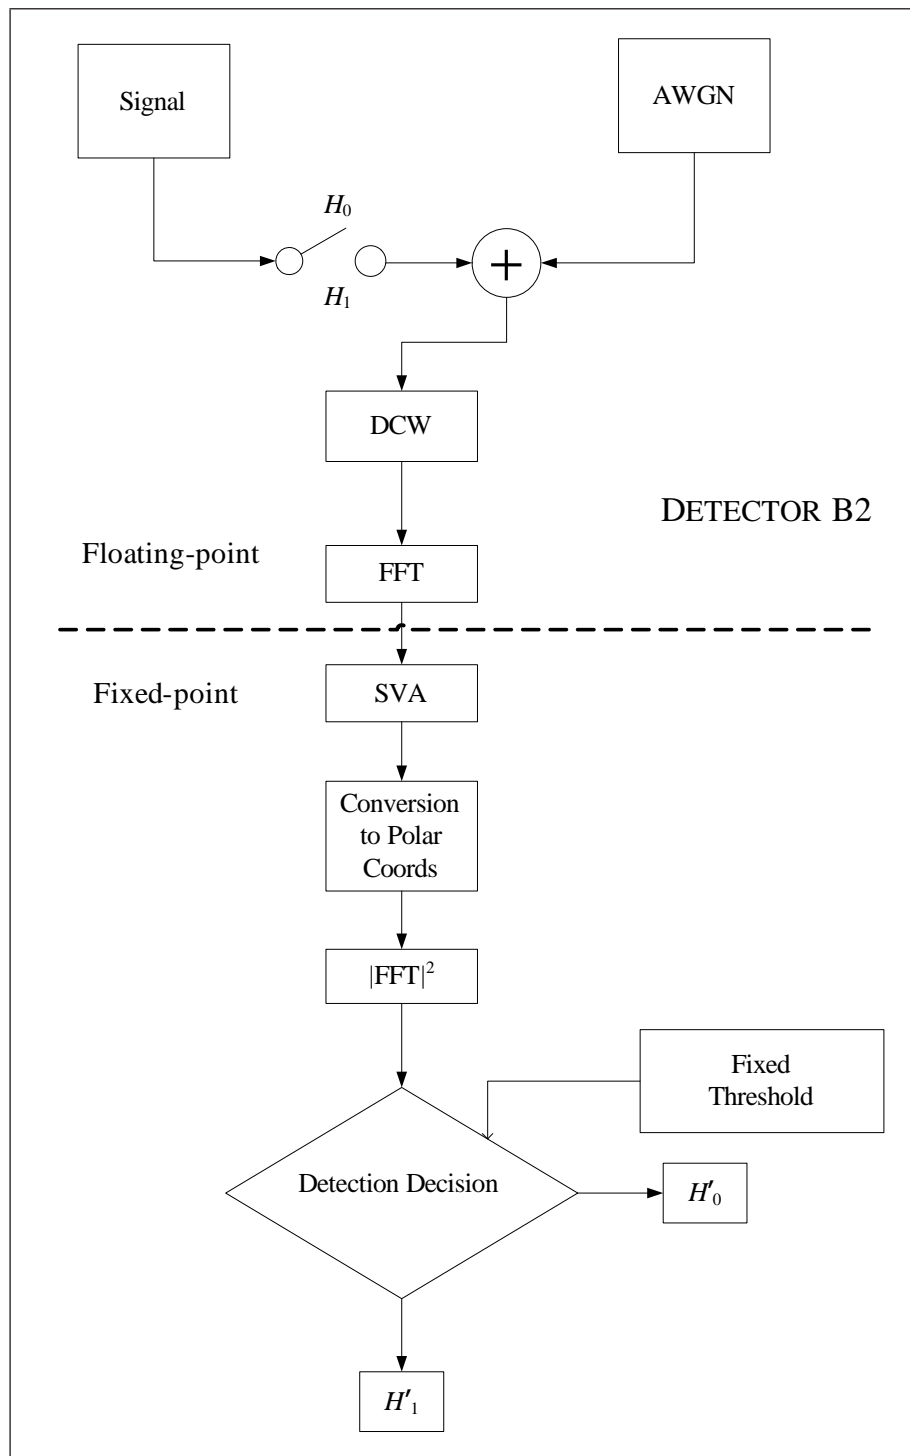


Figure D.3.2 Architecture of the detector B2, which utilises SVA. The partition between floating-point and fixed-point is shown

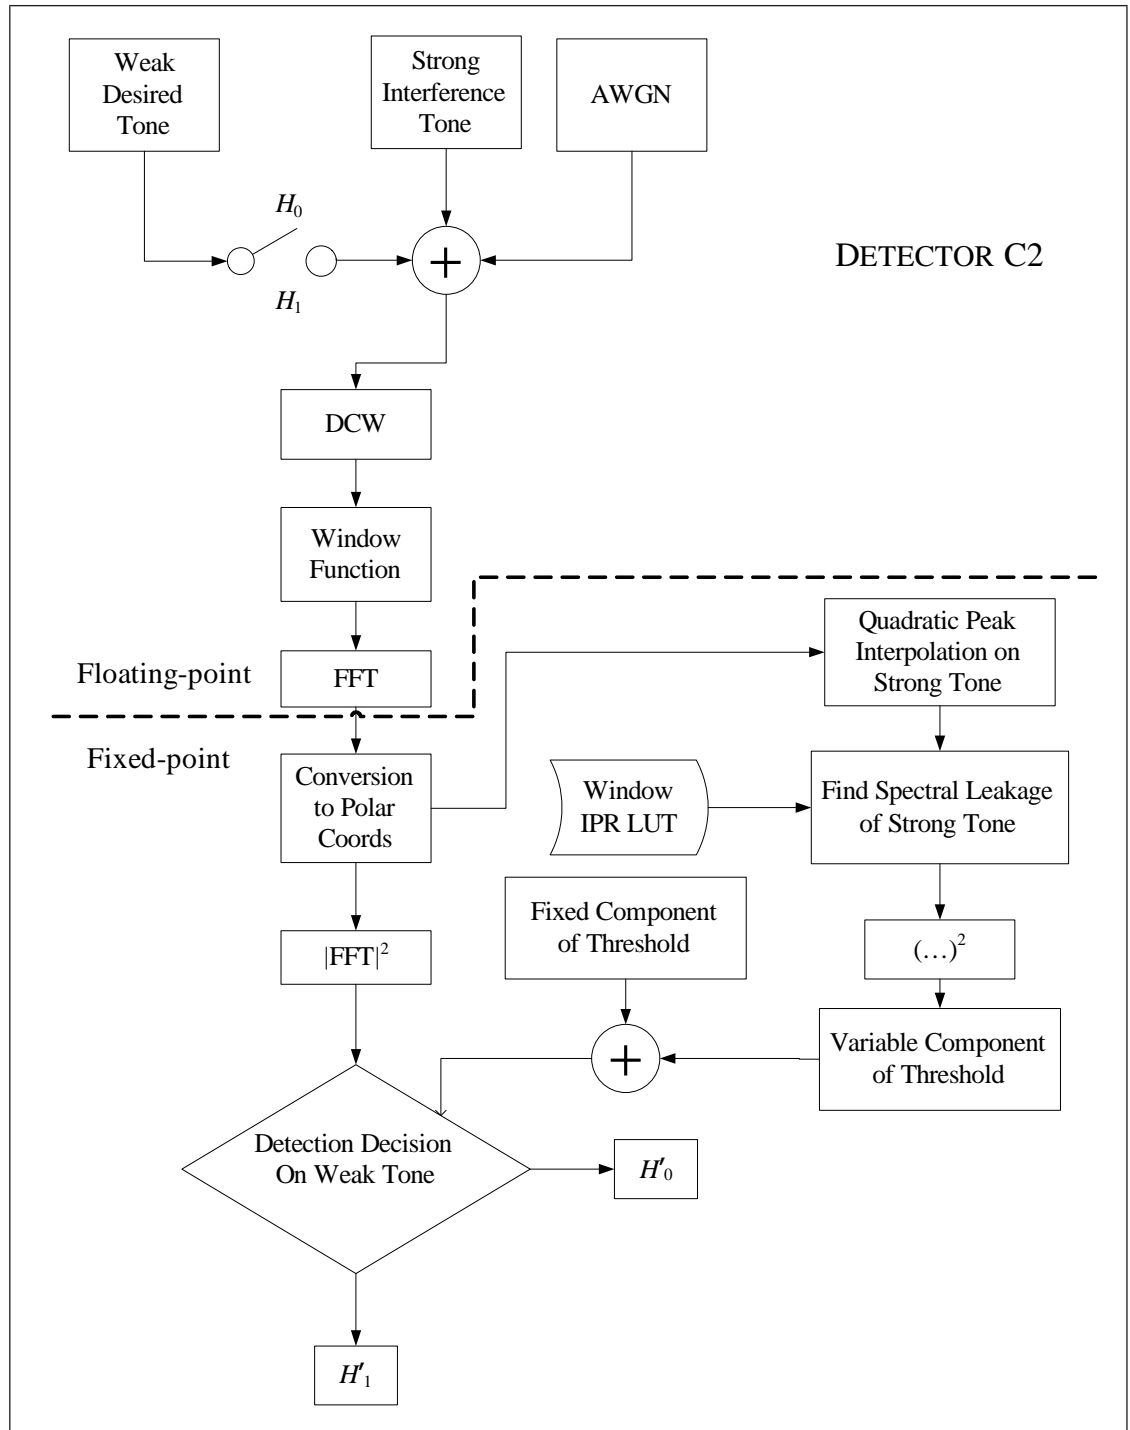


Figure D.3.3 Architecture of the conventional windowing detector C2, which includes a variable component of threshold. The partition between floating-point and fixed-point is shown

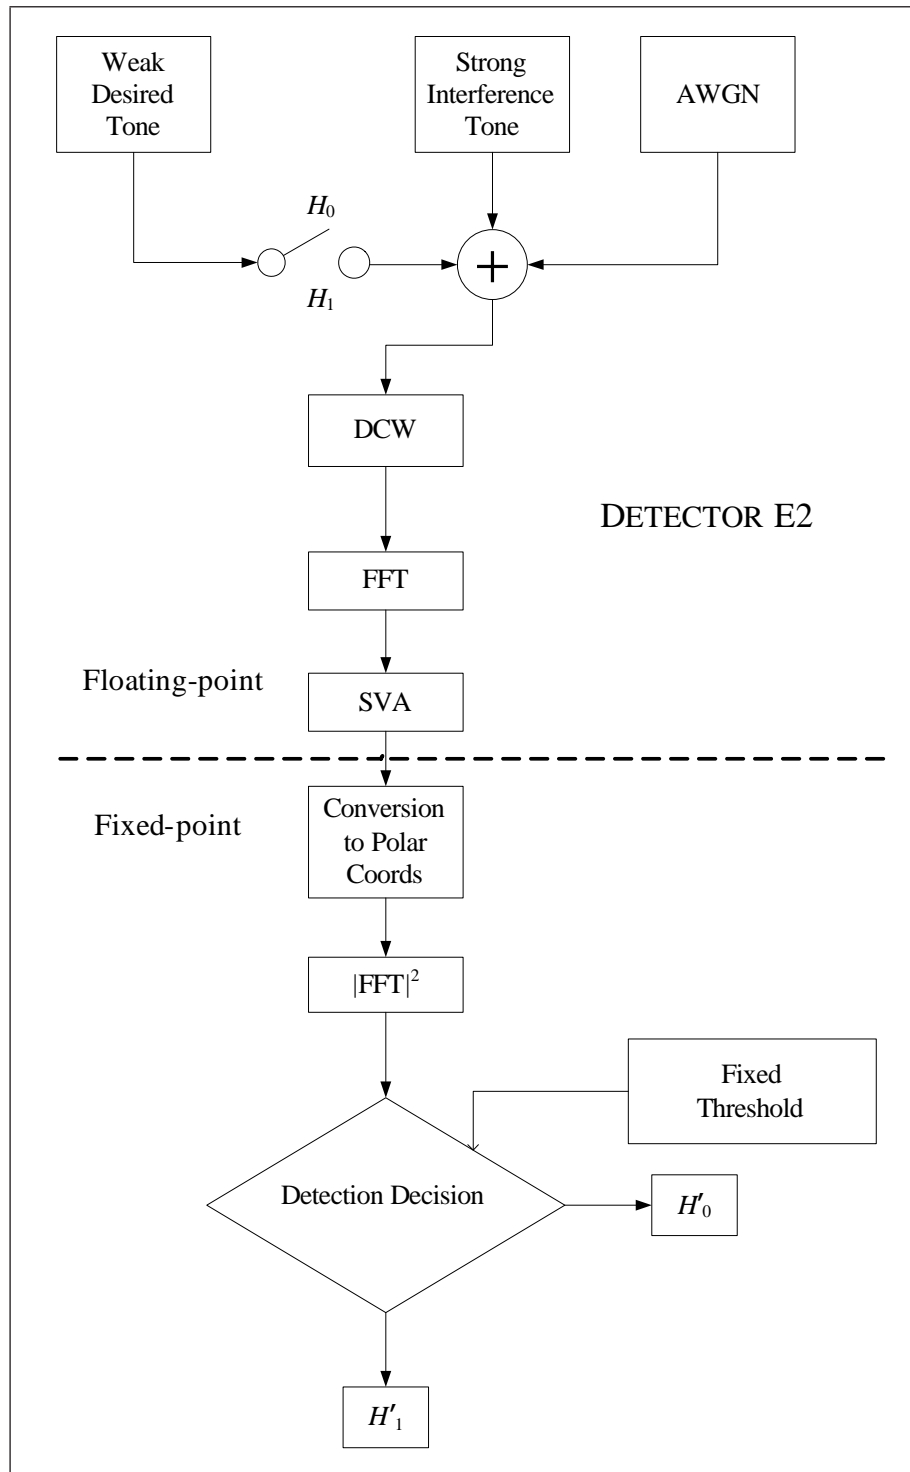


Figure D.3.4 Architecture of the SVA detector E2, which does not include a variable component of threshold. The partition between floating-point and fixed-point is shown

D.4 Coding Phase 2

Naturally, the second phase of coding was to complete the conversion of the detectors from floating-point to fixed-point via translation of the other modules. In other words, the modules within the floating-point partition in Figs D.3.1–D.3.4, which represent processes in a real-life detector, were converted to fixed-point. From the diagrams these modules were the DCW and FFT modules. The DCW module represents an idealised gain control and ADC in a real-life detector ADC.

The remaining modules which did not represent processes in a real-life detector were left in the floating-point domain because they represent continuous, real-world processes. These modules were basically the signals and background noise modules. These processes in front of the ADC in the real world would be: (a) receive by an antenna and (b) mix, amplify and filter by an RF chain (the superhet is a component this).

With fully converted fixed-point detectors a direct comparison was made to the original floating-point results obtained in chapter 3 and hence whether the use of feasible fixed-point arithmetic allowed SVA to keep its advantages over conventional windowing. For clarity the detectors which were modelled entirely in fixed-point are depicted in Figs D.4.1–D.4.4.

The next section gives a detailed account of the actual considerations which were made during the two phases of coding.

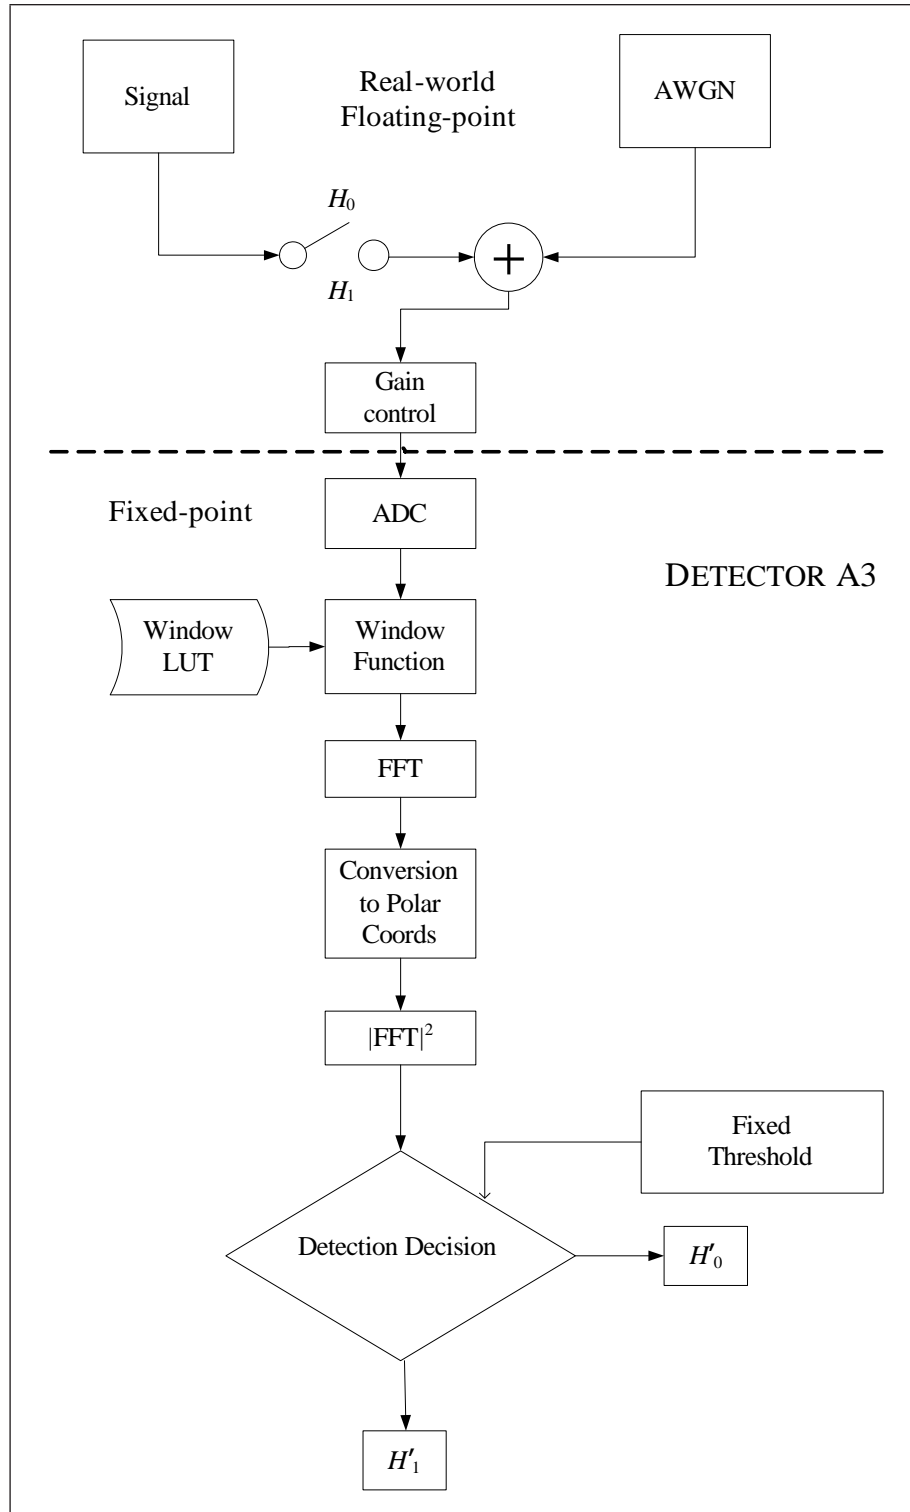


Figure D.4.1 Architecture of the fixed-point conventional windowing detector A3

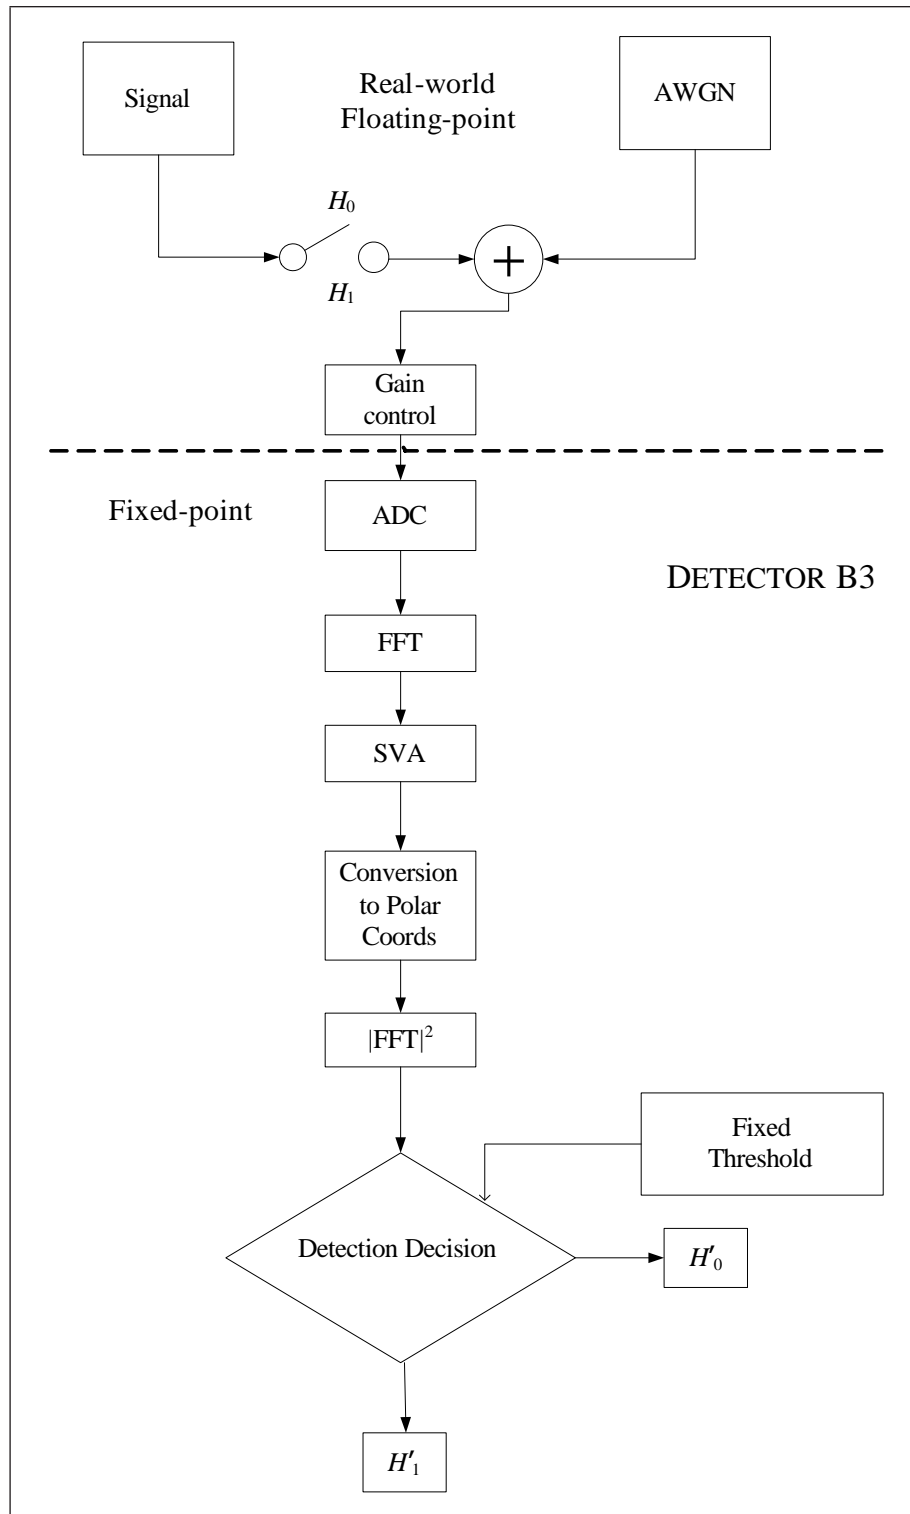


Figure D.4.2 Architecture of the fixed-point detector B3, which utilises SVA

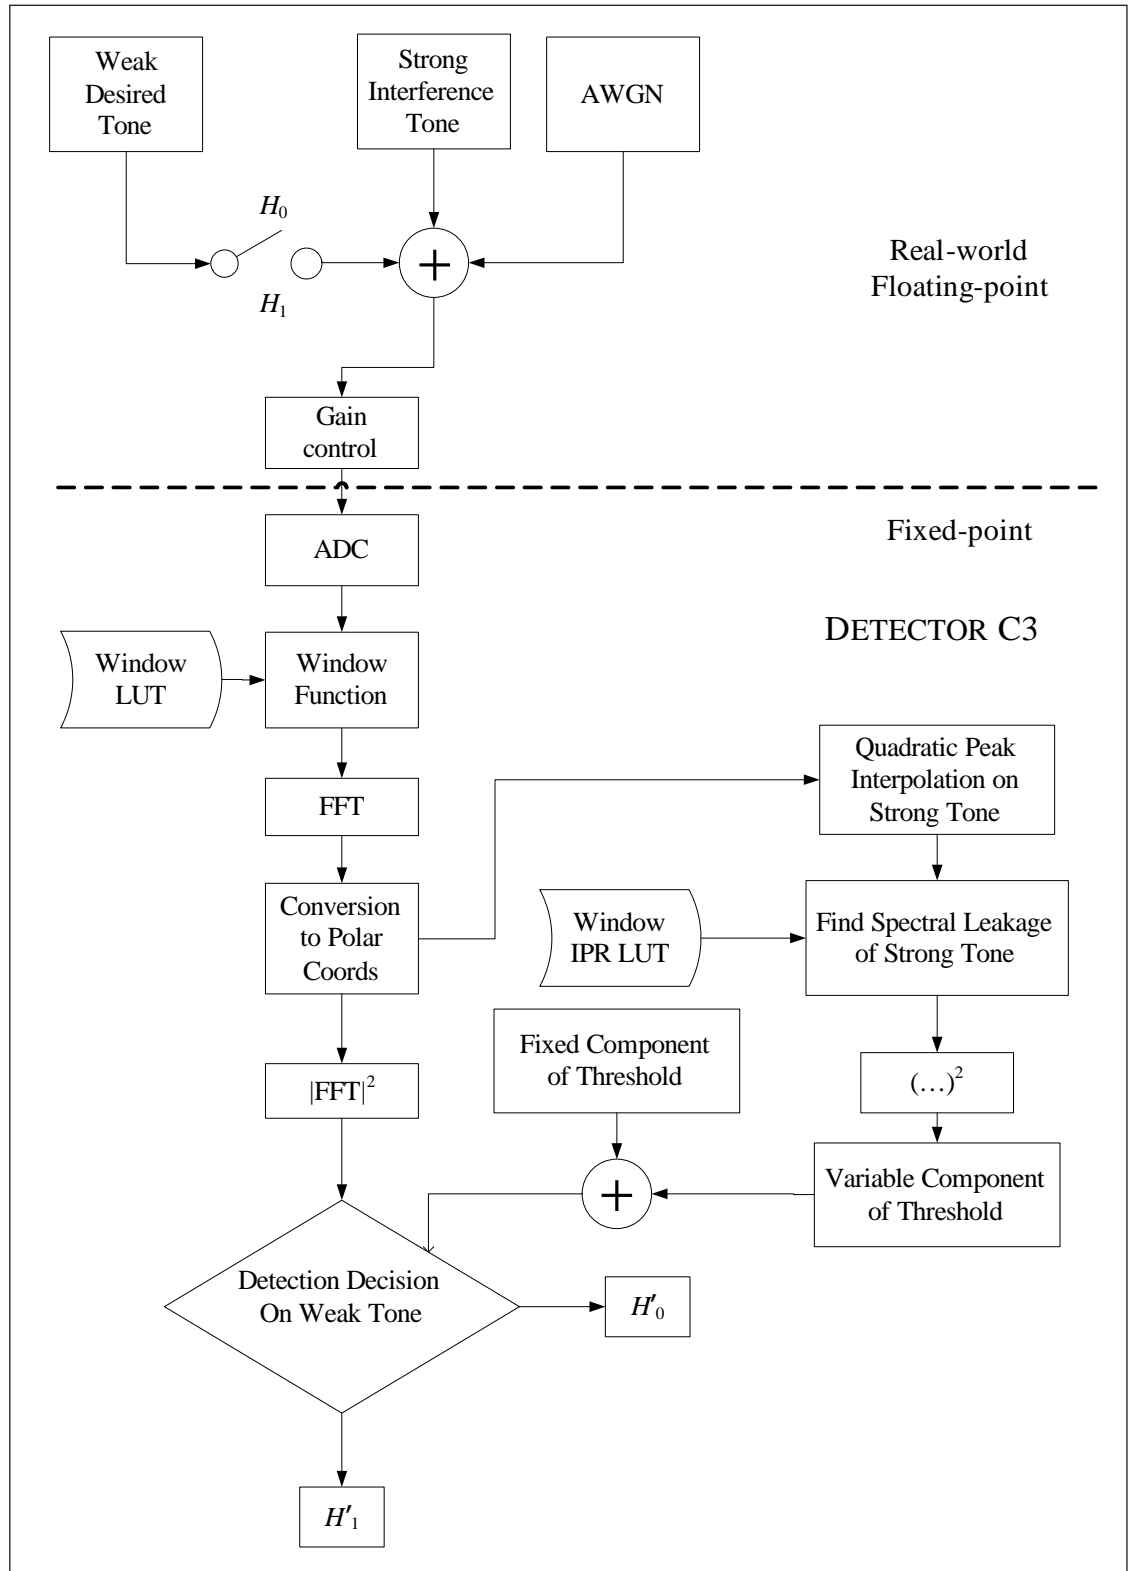


Figure D.4.3 Architecture of the fixed-point conventional windowing detector C3

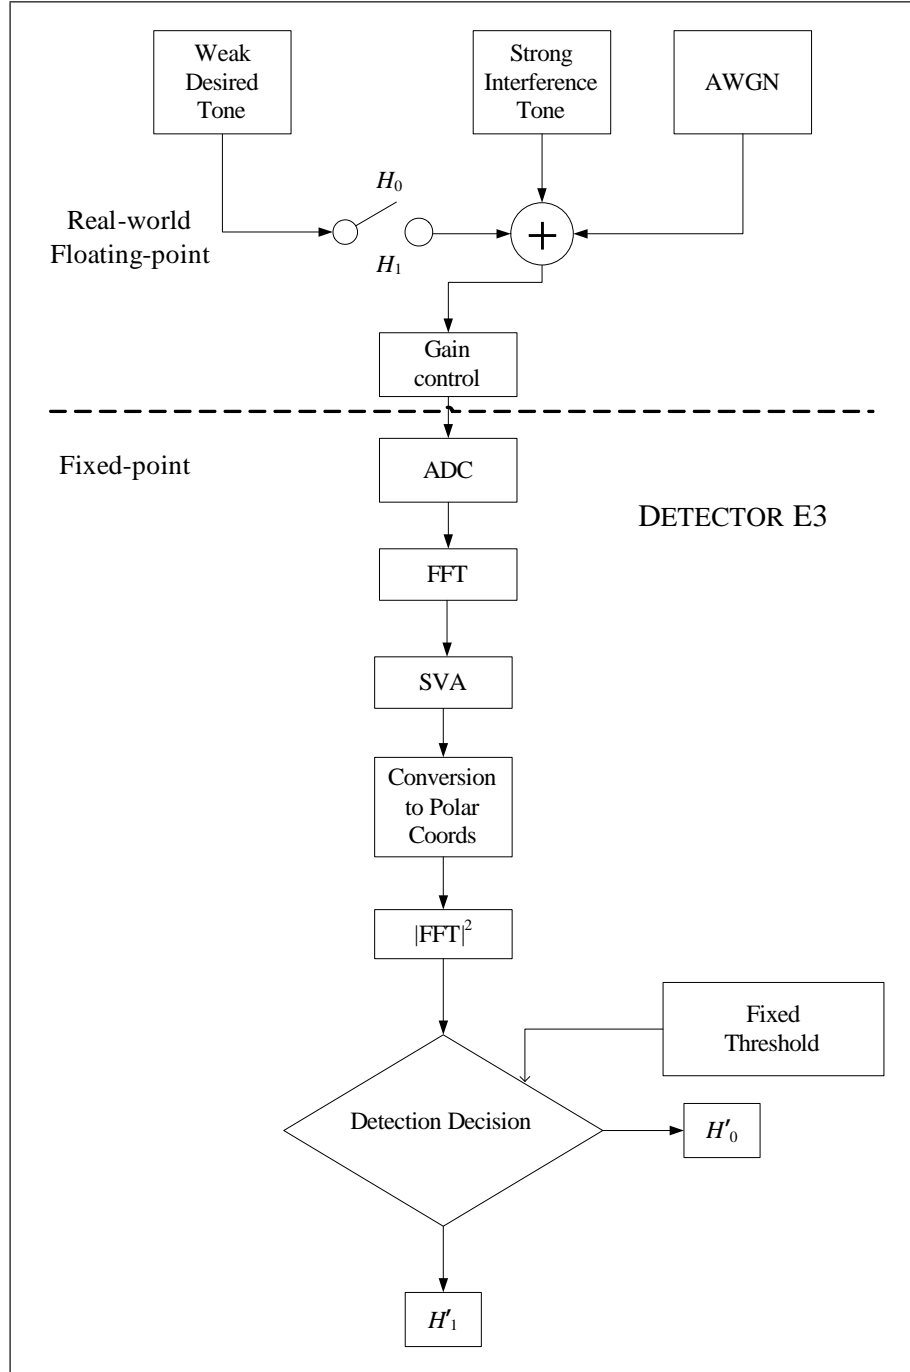


Figure D.4.4 Architecture of the fixed-point SVA detector E3, which does not have a variable component of threshold

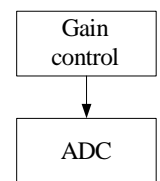
D.5 Coding Explanation

In this section the important information and considerations that were taken into account in production of the fixed-point DRx models are described. This information is given to provide an understanding of the depth the models reached towards a Virtex-5[®] implementation. Therefore, what follows is a description of the major fixed-point modules from Figs D.4.1–D.4.4.

D.5.1 Gain Control & ADC

Note: this fixed-point module description is applicable for the detectors A3, B3, C3 and E3

Firstly the input for the simulated detectors was created in the continuous, floating-point domain in a testbench. The input consisted of either a pulsed sinusoidal radar signal and AWGN or one desired pulsed sinusoidal signal, one interference pulsed sinusoidal signal and AWGN depending on whether detectors A2/A3 and B2/B3 or C2/C3 and E2/E3 were being tested. In either case, the desired pulsed sinusoidal radar signal was present in the H_1 case but not in the H_0 case.



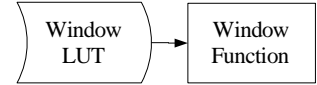
The gain control module served to adjust the received signal power to the dynamic range covered by the ADC quantisation levels and prevent clipping. In a real-life DRx the usual practice is to specify a “headroom” relative to the largest operational signal anticipated. This signal is usually a sinusoid of a specific power. Therefore for waveforms with more complicated structures such as AWGN (with exponential magnitude-squared statistics), the headroom accommodates the tail of the distribution such that overflows are extremely unlikely. Suitable headroom values are not considered in this thesis because the input waveforms can be controlled in the simulation environment.

As a conceptual construct the ADC was defined to produce fixed-point words which consisted solely of fractional bits and no integer bits. This was done to make the fixed-point calculations somewhat easier further down the data flow because for example when multiplications were required the words which consisted only of fractional bits could not produce products that overflowed. The wordlength chosen was 18 bit because this is the main data port width in a DSP48E slice.

D.5.2 Window Function

Note: this fixed-point module description is only applicable for the detectors A3 and C3

The application of a weighted window function in the time-domain involves a simple multiplication between each sample in the DCW and the indexed window function coefficient. In



floating-point MATLAB® the coefficients can be calculated on-the-fly using functions such as `chebwin(...)` or `hann(...)`. This method of calculating on-the-fly would create unnecessary work in a real-life DRx, so rather than proceed down that route the fixed-point code that was created mimics a method of storing the window coefficients in a LUT and multiplying them with the corresponding DCW sample in a pipelined fashion. Assuming one clock cycle for looking-up a coefficient, three clock cycles for a fully-pipelined multiplication and using an embedded multiplier in a DSP48E slice, the total latency using one DSP48E slice for this task would be $4N$ sample clock cycles for a detector with a non-zeropadded FFT.

D.5.3 FFT Filter Bank

Note: this fixed-point module description is applicable for the detectors A3, B3, C3 and E3

An FFT can be implemented in fixed-point in many different guises, each with their own benefits and drawbacks. However the FFT algorithm used

FFT

was the decimation-in-time unit-stride algorithm 1.6.2 from page 45 of [140] with the inclusion of a scaling factor of $1/2$ at each butterfly stage. This was done to counteract the FFT gain, so there was no danger of the intermediate products within the algorithm causing a fixed-point overflow error. The eventual FFT filter bank output had a scaling of $1/N$ in both real and imaginary parts. As a result of the scaling and with the number of data samples at $N = 128$ and the number of fractional bits at $\text{frac} = 18$, the error bound for the FFT bins was

$$\begin{aligned} \text{Error} &= \frac{1}{2} \log_2(N) 2^{-\text{frac}} \\ &= 1.3352 \times 10^{-5} \end{aligned}$$

One of the advantages of this fixed-point FFT algorithm was that the twiddle factors (complex roots of unity, W_N), were pre-computed and stored in a LUT, which allowed different twiddle factors to be called upon depending on the FFT length. This was important as the simulations examined the effect of zeropadding the DCW and therefore different twiddle factors were required to be stored and recalled as necessary. The floating-point translation of the fixed-point code used the MATLAB® built-in FFT algorithm.

D.5.4 SVA Filter

Note: this fixed-point module description is applicable for the detectors B2/B3 and E2/E3

In chapter 3, SVA is occasionally referred to as an adaptive FIR filtering technique. Under this interpretation the FIR filters would be composed of three coefficients (i.e. taps) and two FIR filtering operations would be required per frequency bin: one for the real part and one for the imaginary part. If the FIR filter coefficients were constant, it would be a simple case of designing a detector to use pipelined FIR filters. However the complication with the SVA technique is that the coefficients for the FIR filters at each FFT frequency bin have to be calculated from the FFT output before any FIR filter-like operations can take place. In other words, the value for two of the taps has to be calculated from Eq (2.4.1) before any FIR filter-like operations can occur. The final fixed-point code generated for the SVA technique takes this two stage process into account.

SVA

The calculation of the α coefficients from Eq (2.4.1) is a trivial affair in floating-point, but in fixed-point the situation is more difficult because the dividend and divisor are complex and division must be done with the use of a fixed-point division algorithm. During complex division, the real and imaginary parts of the dividend are multiplied by the complex conjugate of the divisor, then real and imaginary parts are collected together and divided by the magnitude-squared of the divisor. This process is described in Eq (D.5.1) below

$$\frac{a + ib}{c + ie} = \frac{(ac + be) + i(bc - ae)}{(cc + ee)} = \frac{(ac + be)}{c^2 + e^2} + i \frac{(bc - ae)}{c^2 + e^2} \quad (\text{D.5.1})$$

where a and c are the real parts of two complex numbers and b and e are the imaginary parts. However to calculate the α coefficients **only the real part** of the right-hand side of Eq (D.5.1) is actually retained. In other words, only the following in Eq (D.5.2) is required to calculate the α coefficients

$$\begin{aligned} \alpha &= \frac{(ac + bd)}{c^2 + e^2} \\ &\equiv \frac{\Re\{X(k)\} \Re\{X(k+1) + X(k-1)\} + \Im\{X(k)\} \Im\{X(k+1) + X(k-1)\}}{[\Re\{X(k+1)\} + \Re\{X(k-1)\}]^2 + [\Im\{X(k+1)\} + \Im\{X(k-1)\}]^2} \end{aligned} \quad (\text{D.5.2})$$

Therefore for each α coefficient, the dividend calculation would involve two multiplications and one summation, which would consume seven sample clock cycles in one DSP48E slice in a fully-pipelined MACC operation. This assumes that the summation of $X(k+1) + X(k-1)$ has been calculated previously and the results can be reused. Similarly, to calculate the divisor, again assuming $X(k+1) + X(k-1)$ had been calculated elsewhere, two multiplications and one summation are needed. This would also involve seven sample clock cycles if implemented in one DSP48E slice as

a fully-pipelined MACC operation. In addition to this, if the calculation of dividend and divisor is designed to take place within a pipelined architecture using one DSP48E slice for the dividend and one for the divisor, then the total latency to calculate the dividends and divisors for each frequency bin α parameter would be increased by a factor of N_{FFT} .

Before the actual division operation takes place, it is possible to do a series of tests on the dividend to judge whether further calculations are necessary. In effect it is possible to optimise the SVA block. Details on these optimisations can be found in appendix E. Although not used here these optimisations unfortunately do not lend themselves well to a buffered, synchronous pipelined implementation of SVA because they rely on the latency of the constituent components having a fixed delay. The SVA optimisations explained in appendix E result in components having a variable delay and therefore would work only in an architecture involving “acknowledge” and “request” signals between the components. In a hardware implementation the potential optimisations could not speed-up the design but they could achieve some power savings by avoiding unnecessary computations.

D.5.5 Cartesian to Polar Co-ordinates Conversion

The vector output from the FFT block consisted of complex-valued numbers, composed of real and imaginary parts in a Cartesian format. However this is unhelpful because generally ESM DRxs base detection decisions on the received power or envelope of an input waveform. This point was covered in section 1.6. In the floating-point detectors in chapter 3 the received power and envelope were easily calculated using the MATLAB[®] commands `abs(...)`² and `abs(...)` respectively on the FFT output. Using these commands the magnitude-squared and magnitude of the complex numbers was calculated. In detector C2/C3, the magnitude of the FFT output was required for an accurate interpolation of the spectral peak frequencies, whereas, strictly-speaking for detectors A2/A3, B2/B3 and E2/E3 only the magnitude-squared was required. However for code-reuse and comparison reasons all the detectors calculated the magnitude of the FFT output first and then went on to calculate the magnitude-squared for detection purposes.

Conversion to Polar Coords

To undertake the fixed-point equivalent of calculating the magnitude and magnitude-squared of a complex number, a brute force technique would be to multiply each real and imaginary part by itself, sum them, then take the square root. However this is an inefficient method to calculate a magnitude in fixed-point. This is because squaring fixed-point values doubles the number of required bits. Adding the numbers would then increase the number of bits required again by one bit. The square-root operation would allow some of the bits to be dropped, but a square-root operation in itself is non-trivial in fixed-point arithmetic. Therefore the brute-force technique is inefficient because some of the calculated bits are destined to be ignored and it involves a non-trivial

square-root operation. With this information it seemed sensible to use a different route to obtain the magnitude (and subsequently, magnitude-squared) of a complex number and this was achieved by through utilisation of the CORDIC algorithm.

The beauty of CORDIC algorithms is all functions and transformations are based on rotation of a vector in the first or fourth quadrant in an Argand diagram. This is done by known angles stored in a LUT through the use of only simple shift and summation operations in an iterative fashion (and usually one multiplication at the start or end of the CORDIC algorithm to counteract any gain of the algorithm).

The simplest explanation of the CORDIC algorithm is to consider an attempt to approximate the cosine or sine function of an input angle. In this case, a suitable starting point to understand CORDIC is to envisage an iterative vector rotation process for a general vector (a_0, b_0) . A useful mental picture for rotation of a general vector is an unforced pendulum gradually swinging towards its resting position, which corresponds to the vector at the desired angle. The n^{th} iteration of rotation with the use of a rotation matrix is described mathematically as

$$\begin{aligned} \begin{pmatrix} a_{n+1} \\ b_{n+1} \end{pmatrix} &= \begin{pmatrix} \cos \theta_n & -\sin \theta_n \\ \sin \theta_n & \cos \theta_n \end{pmatrix} \begin{pmatrix} a_n \\ b_n \end{pmatrix}, \quad n = 0, 1, 2, \dots, M-1 \\ &= \frac{1}{\sqrt{1 + \tan^2 \theta_n}} \begin{pmatrix} 1 & -\tan \theta_n \\ \tan \theta_n & 1 \end{pmatrix} \begin{pmatrix} a_n \\ b_n \end{pmatrix} \end{aligned} \quad (\text{D.5.3})$$

where M is the total number of iterations and θ is the angle the vector makes against the real axis. A value for M is selected prior to the calculations. If the rotation angles $(\tan \theta_n)$ are limited to powers of two $(\pm 2^{-n})$, then Eq (D.5.3) can be simplified into two equations. These two equations require only shifts (multiplications by a power of two) and summations to implement in fixed-point arithmetic. This is highly favourable because these operations are highly computationally efficient. These equations in vector format are

$$\begin{pmatrix} a_{n+1} \\ b_{n+1} \end{pmatrix} = K_n \begin{pmatrix} a_n - \delta_n b_n 2^{-n} \\ b_n + \delta_n a_n 2^{-n} \end{pmatrix}, \quad \text{where } K_n = \frac{1}{\sqrt{1 + 2^{-2n}}} \quad (\text{D.5.4})$$

and $\delta_n = \pm 1$

K_n is the CORDIC algorithm gain after the n^{th} rotation and therefore after M iterations the total gain K is

$$K = \prod_{n=0}^{M-1} K_n$$

The result after M iterations needs to be adjusted to account for this gain through multiplication by $1/K$. Fortunately $1/K$ can be easily stored in a LUT as a pre-computed constant. The value of δ_n in Eq (D.5.4) is equivalent to the choice of direction

of rotation for the pendulum and is chosen depending on the function or transformation that is desired to be approximated by the CORDIC algorithm. In the above example of the approximation of the cosine or sine of a given angle, δ_n is chosen to attempt to match the rotated vector to a vector at the desired angle (the resting position of the pendulum) and this introduces the **third** CORDIC equation: the angle accumulator

$$z_{n+1} = z_n - \delta_n \theta_n \quad (\text{D.5.5})$$

For example, to calculate the cosine and sine of an angle, δ_n is declared to be positive for positive z_n and negative for negative z_n .

All CORDIC algorithms use three fundamental equations as the basis for their computation, this can be seen in [141]. The various CORDIC algorithms differ only in their choice of the value of the parameter δ_n at each iterative step and another parameter m . These change according to which function or transformation CORDIC is used to approximate. If $m = 1$ the CORDIC algorithm will calculate the trigonometric group of functions; if $m = -1$ then CORDIC can be used to calculate the hyperbolic functions and if $m = 0$ CORDIC can be used for “linear” transformations. The linear transformations allow multiplication and division to be calculated using only shifts and summations. To summarise the general form of the three CORDIC equations is

$$a_{n+1} = a_n - m \delta_n b_n 2^{-n} \quad (\text{D.5.6a})$$

$$b_{n+1} = b_n + \delta_n a_n 2^{-n} \quad (\text{D.5.6b})$$

$$z_{n+1} = z_n - \delta_n \theta_n \quad (\text{D.5.6c})$$

The CORDIC algorithms store pre-computed values for the scaling factor ($1/K$) and the iterative rotation angles used in the angle accumulator equation (these are $\theta_n = \arctan(2^{-n})$ in the example of the cosine and sine of an input angle).

The class of CORDIC algorithm for calculation of the magnitude of a complex SVA-filtered frequency bin is the trigonometric class. The calculation was achieved by iterative rotation of the vector (X_I, X_Q) in an attempt to force it to sit on the x -axis in an Argand diagram. The initial conditions for Eqs (D.5.6a)–(D.5.6c) were: $a_0 = X_I$, $b_0 = X_Q$, $z_0 = 0$, $m = 1$ and $\delta_n = 1$ if b_n is negative, otherwise $\delta_n = 0$. The SVA-filtered frequency bin in real and imaginary part form was converted to magnitude and phase form. After M iterations the approximation to the magnitude of X was located in the memory register holding the current value for Eq (D.5.6a) and was given by

$$a_M \approx K \sqrt{X_I^2 + X_Q^2} \quad (\text{D.5.7})$$

The result in Eq (D.5.7) was then scaled by multiplication with the stored constant value for $1/K$.

The accuracy of the unscaled magnitude result improves by two bits for every iteration of the CORDIC algorithm, this meant that for a complex number input consisting of 18 bit real and imaginary parts, the CORDIC process was halted after 9 iterations and the scaling multiplication was applied. Finally to calculate the magnitude-squared of the complex number, the magnitude result was multiplied by itself.

D.5.6 Quadratic Peak Interpolation

Note: this fixed-point module description is only applicable for detectors C2/C3

The use of detectors which employ conventional fixed weight window functions necessitates mitigation against sidelobe spectral leakage and mainlobe width caused by the particular window function

Quadratic Peak
Interpolation on
Strong Tone

used. It has already been shown in section 3.4 that for detectors using SVA the spectral leakage from one signal into another is negligible and can be safely ignored. This is not true for the conventional window function detectors. It is necessary to estimate the sidelobe spectral leakage from any strong signal peaks into the surrounding frequency bins to prevent spurious detections due to leakage.

This meant it was important to get a good estimate of the true magnitude and frequency of interference signal peaks. There are many methodologies for the estimation of the true peak from available FFT data, a selection of these are described in [142]. The most accurate method to obtain an improved estimate of a spectral peak is to zeropad the DCW before the FFT operation. However as zeropadding is relatively computationally complex and there are diminishing returns from zeropadding to more than double the original DCW length. It was desirable to keep computational complexity to a minimum, hence a simple interpolation technique of plotting a quadratic through the data points was used in all of the fixed window function detector simulations, both zeropadded and non-zeropadded. As was described in section 3.4, a QIFFT is a better quick estimation technique than zeropadding alone. Therefore the quadratic peak interpolation block in the floating-point detector C1 was translated to fixed-point in detectors C2 and C3.

In fixed-point, Eq (3.4.4) could be implemented in hardware using two FIR filters (one for the dividend and one for the divisor) and one division operator. FIR filters lend themselves well to implementation on FPGA. However the alternative approach which was used in the fixed-point coding of detectors C2 and C3 was to calculate the dividend and divisor through a combination of arithmetic shifts and accumulation. Specifically these operations were:

- The dividend was calculated by shifting the fixed-point numbers p and r one bit to the left and then r was subtracted from p .
- The divisor was calculated by shifting the fixed-point number q one bit to the right; the result was subtracted from p and stored in an accumulator and r was summed into the accumulator.

Once the calculation of Δf had been accomplished, the calculation of Q in Eq (3.4.5) was evaluated in fixed-point. This too involved arithmetic shifts and MACCs.

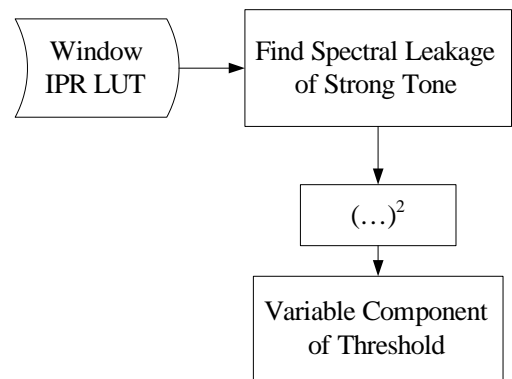
As an additional note, in an actual DRx there would need to be a search for the largest magnitude peaks in the spectral estimate as the sources of most of the spectral leakage. Since the magnitude of the FFT output is not sorted in any way, searching for the maximum value would be an $O(N_{\text{FFT}})$ computation. This would require approximately 2 clock cycles for each comparison, so that the search in total would take $2N_{\text{FFT}}$ clock cycles.

The usefulness of quadratic peak interpolation is demonstrated in chapter 3 in the floating-point simulations where it combined well with zeropadded FFT output data. This was expected to continue to be the case in the fixed-point domain.

D.5.7 Spectral Leakage & Variable Threshold Component

Note: this fixed-point module description is only applicable for detectors C2/C3

With a good estimate of a spectral peak frequency and magnitude and knowledge of which window function was applied prior to the FFT, it was possible to estimate the level of sidelobe spectral leakage and mainlobe width from the peak into other frequency bins by masking the envelope of the magnitude of the frequency-domain response of the window function around the peak, exactly as described in Fig 3.4.3.



The most accurate method of calculation would be to use the theoretical equations of the magnitude of the frequency-domain responses such as Eqs (3.4.1)–(3.4.3). However, such calculations would be extremely complicated in fixed-point as they invariably involve trigonometric functions and much use would be required of fixed-point techniques such as those in [143]. The simpler method that was used was to store 2^{18} pre-calculated values of the magnitude of the frequency-domain response for each window function tested. The stored values were scaled to fit the estimated magnitude of the spectral peak. The sidelobe spectral leakage and mainlobe width from the estimated spectral peak were, in turn, estimated at other frequency bins from the frequency difference between the scaled response peak and the bin of interest. To form the variable component of threshold the scaled, recalled magnitude value was squared by multiplication with itself.

Finding the spectral leakage in fixed-point was undertaken using the following procedure which was chosen because it would be easily implemented in hardware using the operations in the brackets:

- find the frequency difference between the estimated spectral peak and bin of interest (subtraction);
- use the frequency difference as an address for the LUT;
- scale the stored frequency-domain response magnitude value (multiplication);
- square the magnitude value (multiplication).

The fixed-point SVA detectors (E2 and E3 in the interference-present scenarios) did not use or require a variable component of threshold due to the conclusion in section 3.4 that the effect of sidelobe spectral leakage is negligible with the use of SVA. Therefore although the complexity of SVA is more complicated than conventional windowing, detectors such as E2 and E3 can recover some efficiency over detectors such as C2 and C3 by foregoing the variable-component-of-threshold calculations.

D.5.8 Fixed Threshold Component

This value of the fixed component of threshold was found from previous P_{fa} against P_d simulations. The detectors were simulated over many trials under the H_0 condition (1×10^5). This entailed AWGN as the input for detectors A2/A3 and B2/B3 and AWGN plus interference for detectors C2/C3 and E2/E3. The fixed component of threshold was set to a value such that the probability of false alarm was empirically $P_{fa} = 0.01$. This value for P_{fa} was selected so that it matched the value in the floating-point simulations and therefore the end results could be compared.

Fixed Component
of Threshold

D.5.9 Detection Decision

The comparator part of the detector checked the magnitude-squared of the frequency bin of interest against the total threshold, which was composed of the fixed component only in detectors A2/A3 and B2/B3 and the fixed and variable components in detectors C2/C3 and E2/E3.

If the value in the frequency bin was larger than the threshold then a signal was declared to be present (case H'_1) otherwise no signal of interest was declared to be present (case H'_0). No special coding was required in the fixed-point coding because in an actual FPGA implementation, the comparator would simply require some logic circuitry and a subtraction operation.

Detection Decision

D.6 Results

In short the results demonstrated the success of the fixed-point SVA implementation. This is because when the models were examined using Monte Carlo simulations,

the ROC curve results that were returned showed no difference, save for a negligible amount of error, to those obtained with the use of floating-point models from chapter 3. Although fixed-point results were obtained for all the floating-point scenarios from chapter 3, here results are shown only for an “easy” scenario (full DCW) and a more challenging scenario (interference in full DCW, $\text{SNR}_{\text{out},B} = 15 \text{ dB}$, frequency separation $2f_s/N$) to demonstrate the negligible difference after translation to the fixed-point system. Additionally, results obtained with the use of detectors A2, B2, C2 and E2 have been omitted due to their close similarity to the results from detectors A3, B3, C3 and E3 and because they are of lower interest owing to the hybrid digital floating-point/fixed-point nature of the detector models.

The criterion by which the results were judged in order to decide whether the implementable fixed-point detectors A3, B3, C3 and E3 were successful or not was that the ROC curve results should match those obtained from the corresponding floating-point detectors A1, B1, C1 and E1. The notion was if they matched the floating-point results it would show that the translation of SVA to FPGA-implementable fixed-point code does not have a detrimental effect on the performance of SVA as a windowing technique and the advantages of SVA could be enjoyed in an realistic DRx design.

D.6.1 Full DCW

Detectors A3 and B3 were relevant for this scenario. In order to compare the ROC curve results from the fixed-point detectors A3 and B3 with the results from their floating-point counterparts A1 and B1, the full DCW scenario had to be set-up identically to that in section 3.2. Briefly this entailed:

- complex exponential signal and AWGN;
- random frequency, f_T , of the signal over the frequency bin interval;
- random initial phase, ϕ_0 , of the signal;
- probability of false alarm, $P_{fa} = 0.01$;
- non-overlapped DCWs;
- detection on each DCW;
- window functions used in detector A3: Chebychev, Hann and rectangular;
- non-zeropadded and zeropadded FFTs used;
- 1×10^5 trials in a simulation.

As stated previously and shown in Fig D.2.1 for the SVA module, literal floating-point translations of the fixed-point code were used to further verify the functionality of the modules in the fixed-point detectors. Once the functionality was satisfactorily verified, the main Monte Carlo simulations were undertaken. The ROC curve results

for the fixed-point detectors are shown in Fig D.6.1. When compared back to the floating-point results in Fig 3.2.6 of section 3.2, it is clear that the results are virtually identical, except for some negligible error caused by quantisation. As the focus of the translation to fixed-point arithmetic was whether SVA would be able to perform as well in implementable fixed-point as it could in floating-point, in Fig D.6.2 the fixed-point SVA result are overlaid on the floating-point results from Fig 3.2.6.

The same conclusions could be drawn over the relative detector performances in Fig D.6.1 as were drawn in section 3.2. For example the rectangular-windowed fixed-point detector was the best and the fixed-point SVA detector recovered most of the windowing loss over the Chebychev-windowed and Hann-windowed detectors. However this was not the aim of the fixed-point simulations. The important conclusion to draw is that with the use of implementable fixed-point arithmetic, which could be easily mapped onto a Xilinx[®] Virtex-5[®] FPGA, the detectors have the same relative performance as the floating-point detectors save for a negligible amount of quantisation error.

If, for example, the Hann-windowed detector performed differently and had been better in fixed-point than in floating-point, it would have proved that there would be a problem with the implementation of SVA in a real-life DRx in an ESM system. The MBD methodology chain of SVA would be terminated at this point. In the actual event, the full DCW results show there are no adverse effects of the implementation of SVA in a fixed-point system which can be easily mapped onto a gate-array.

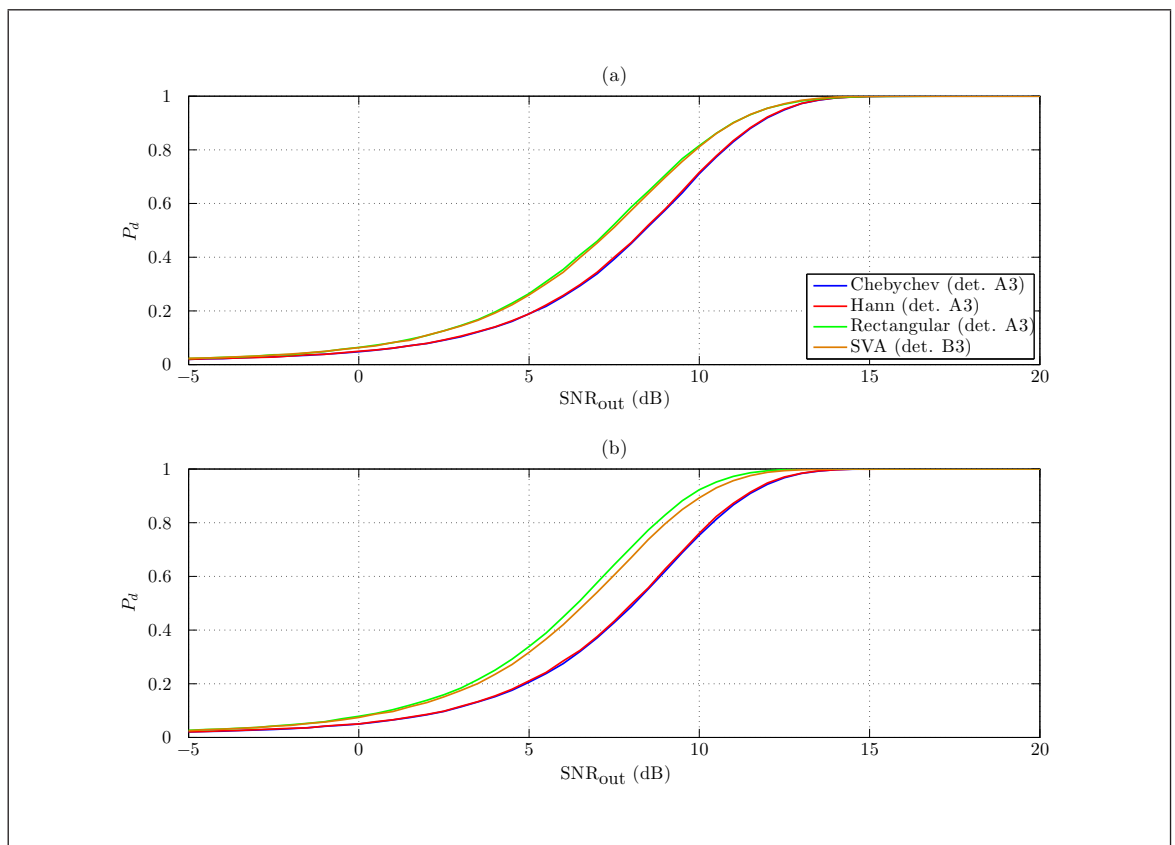


Figure D.6.1 ROC curve results from fixed-point detectors A3 and B3 for: (a) the non-zero padded case; (b) the zero padded case

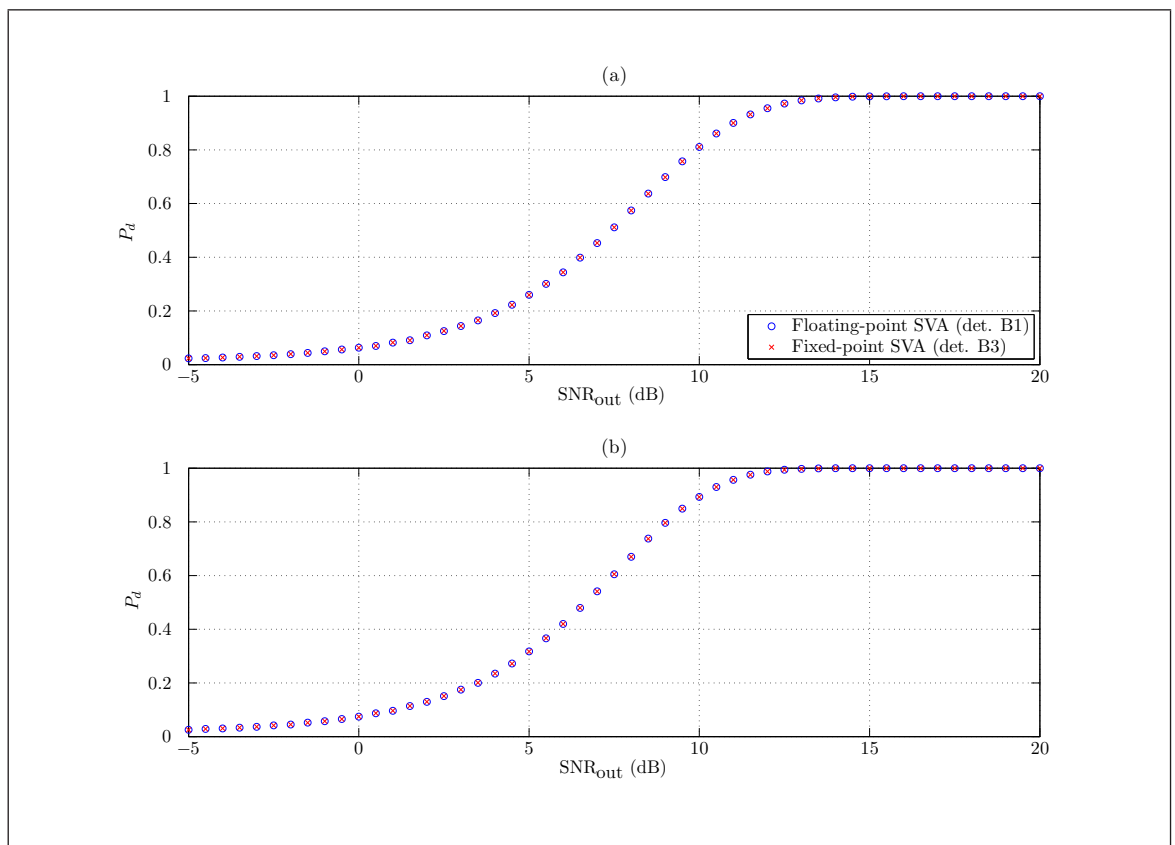


Figure D.6.2 Full DCW scenario detector B3 ROC data points (\times) overlaid on floating-point detector B1 points (\circ) for: (a) the non-zero padded case; (b) the zero padded case

D.6.2 Interference in Full DCW

In this scenario detectors C3 and E3 were relevant. Again the scenario was built using the same characteristics as the corresponding floating-point scenario in section 3.4. As a summary these were:

- a desired complex exponential, S_A , an interfering complex exponential, S_B , and AWGN;
- signals with random frequencies, $f_{T,A}$ and $f_{T,B}$, each randomly chosen over a frequency bin interval at the start of each test run;
- random initial phases of the signals, $\phi_{0,A}$ and $\phi_{0,B}$;
- probability of false alarm, $P_{fa} = 0.01$;
- non-overlapped DCWs;
- detection on each DCW;
- window functions used in detector C3: Chebychev, Hann and rectangular;
- non-zeropadded and zeropadded FFTs used;
- 1×10^5 trials in a simulation;
- average frequency separation between $f_{T,A}$ and $f_{T,B}$ was $2f_s/N$ and $5f_s/N$.

Once the functionality of the individual modules in the fixed-point detectors had been verified through a comparison with the outputs obtained from the literal floating-point translation of the code, the main Monte Carlo simulations were undertaken. Fig D.6.3 shows the ROC curve results for the $\text{SNR}_{\text{out},B} = 15\text{ dB}$ with an average frequency separation of $2f_s/N$.

In addition the specific SVA detector ROC curve results are isolated and shown in Fig D.6.4.

The situation is much the same as in the full DCW scenario in that the results are identical to those from the corresponding floating-point simulations, save for a negligible error caused by quantisation.

The crucial conclusion to draw from the results is that with the use of SVA in a feasible fixed-point detector, it maintains its sidelobe spectral leakage minimisation property in a difficult scenario. This furthers the argument that SVA is a viable algorithm to be implemented on FPGA.

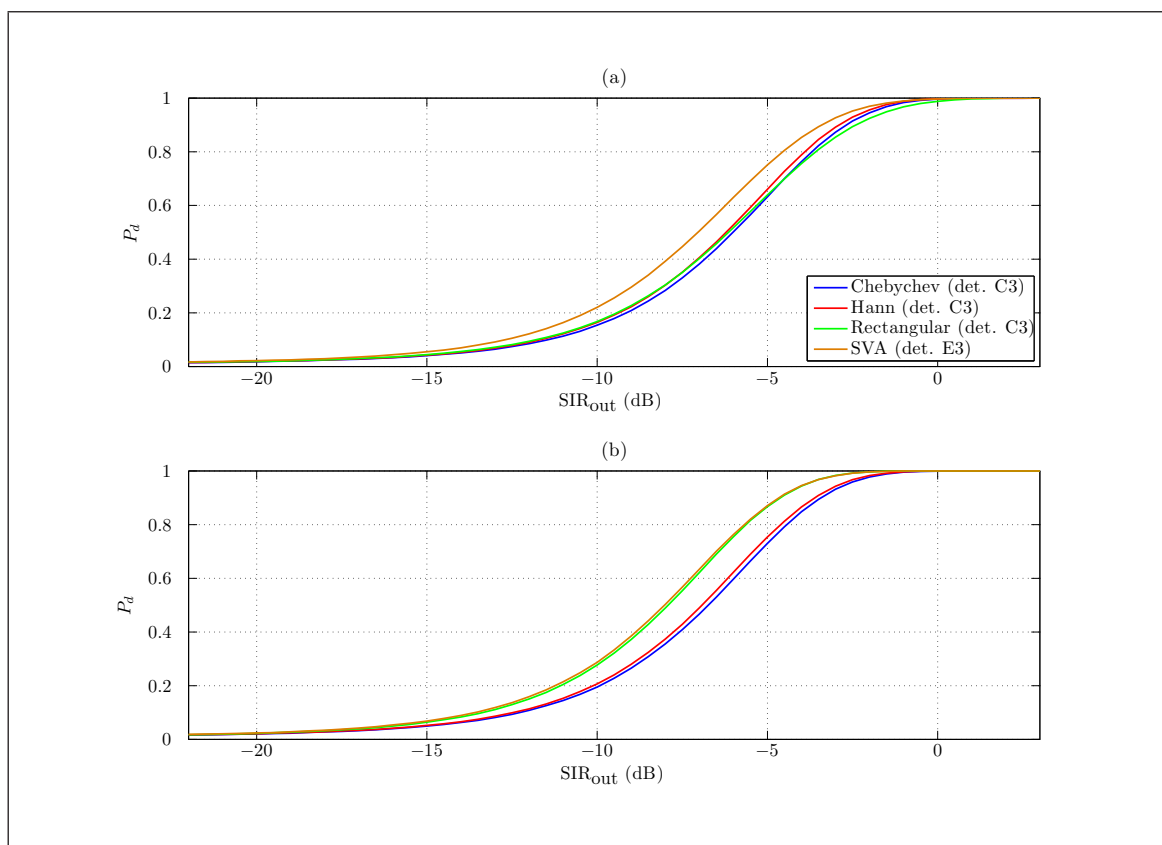


Figure D.6.3 ROC curve results from the fixed-point detectors C3 and E3 for $\text{SNR}_{\text{out}} = 15$ dB and average separation of $2f_s/N$ for: (a) the non-zeropadded case; (b) the zeropadded case

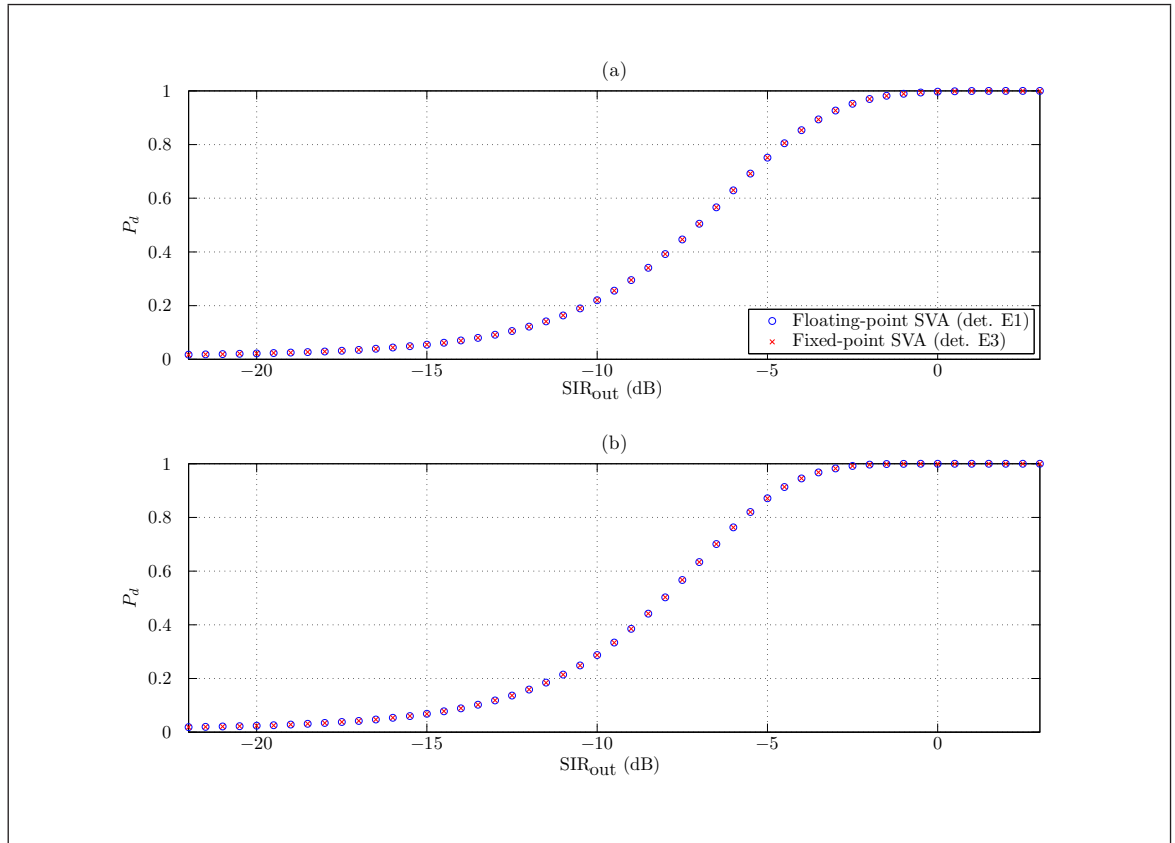


Figure D.6.4 Interference full DCW scenario detector E3 ROC data points (\times) overlaid on floating-point detector E1 points (\circ) at $SNR_{out,B} = 15$ dB and average separation of $2f_s/N$ for: (a) the non-zero padded case; (b) the zero padded case

D.7 Conclusions

The study of the use of SVA in a fixed-point arithmetic system, which was designed with implementation in mind, showed that the advantages of SVA in a simulated DRx are still present.

The fixed-point system used in this section was tailored to suit the Xilinx[®] Virtex-5[®] DSP capabilities and as such the results provide support for the notion that a real-life DRx using SVA is attainable. The results paved the way for the next stage in the MBD flow, that is, to implement on FPGA.

APPENDIX E

SVA Optimisations

A valid argument about a disadvantage of the fixed-point implementation of SVA is:

The division operator in the SVA algorithm is relatively, computationally intensive compared to conventional FIR filters which have the ability to employ MACC methods on FPGA boards. The division operator to calculate the α parameter in the SVA algorithm would be required for **every** FFT frequency bin if the desire was to apply SVA to the entire FFT.

However, there are certain optimizations which could be implemented and which appear when the calculations of fixed-point SVA are analysed. These could be used to reduce the overall computational burden of SVA. The fixed-point coding in detectors B2/B3 and E2/E3 did not utilize these optimizations, since the primary goal was to translate to fixed-point and then later consider where optimizations could be made. Another reason against the inclusion of the optimisations in detectors B2/B3 and E2/E3 was that the optimisations would lead to a variable delay for the α parameter calculation part of the SVA functional block. This would be acceptable if the eventual implementation on FPGA consisted of functional blocks which have “acknowledge” and “request” signals between them. However the optimisations detailed in this appendix are less suitable for a buffered, synchronous pipeline architecture because this would rely on functional blocks with a fixed latency.

The optimisations could be utilised in a synchronous pipeline if the target optimisation criterion was to reduce power consumption. In this situation delays would be introduced within the functional blocks of the SVA block and they would be forced to await request signals. In essence, the latency of the SVA functional block would be fixed. Therefore the optimisations in this appendix would be unlikely to improve the SVA algorithm latency, however they are included nonetheless because they have power consumption relevance. Power consumption is an important factor for an ESM DRx because heat dissipation in a mobile platform can be problematic.

E.1 Optimisations

The basic equation for an SVA-filtered FFT is given in Eq (2.4.2) with the value of the α coefficient at each FFT frequency bin calculated from Eq (2.4.1). Hence, there are two main stages in the execution of SVA:

1. The calculation of the α coefficients for each FFT frequency bin.
2. The application of the α coefficients in an operation which resembles a 3-tap FIR filter at each bin.

Therefore in order to execute Eq (2.4.2), the α coefficients from Eq (2.4.1) are prerequisite. Complex-valued $X(k)$ can be decomposed into real and imaginary parts as $X(k) = X_I(k) + iX_Q(k)$. From this representation Eq (2.4.1) can be broken down into its divisor (D) and dividend (E) inside the \Re operator from which the real and imaginary parts can be collected together as thus

$$\begin{aligned}
 D &= X(k+1) + X(k-1) \\
 &= [X_I(k+1) + iX_Q(k+1)] + [X_I(k-1) + iX_Q(k-1)] \\
 &= [X_I(k+1) + X_I(k-1)] + i[X_Q(k+1) + X_Q(k-1)] \\
 E &= X(k) \\
 &= X_I(k) + iX_Q(k)
 \end{aligned}$$

To perform a complex division, the E and D need to be multiplied by the complex conjugate of D , which is given by

$$\begin{aligned}
 DD^* &= ([X_I(k+1) + X_I(k-1)] + i[X_Q(k+1) + X_Q(k-1)]) \\
 &\quad \times ([X_I(k+1) + X_I(k-1)] - i[X_Q(k+1) + X_Q(k-1)]) \\
 &= [X_I(k+1) + X_I(k-1)]^2 + [X_Q(k+1) + X_Q(k-1)]^2 \quad (\text{E.1.1})
 \end{aligned}$$

and the dividend multiplied by the complex conjugate of the divisor is

$$\begin{aligned}
 ED^* &= (X_I(k) + iX_Q(k)) ([X_I(k+1) + X_I(k-1)] - i[X_Q(k+1) + X_Q(k-1)]) \\
 &= X_I(k) [X_I(k+1) + X_I(k-1)] - iX_I(k) [X_Q(k+1) + X_Q(k-1)] \\
 &\quad + iX_Q(k) [X_I(k+1) + X_I(k-1)] + X_Q(k) [X_Q(k+1) + X_Q(k-1)] \quad (\text{E.1.2})
 \end{aligned}$$

The revised divisor part of the equation for the α coefficient, DD^* , in Eq (E.1.1) is definitely positive and real because it is the magnitude-squared of a complex number. However the dividend in Eq (E.1.2) is complex. When the \Re operator is reintroduced, it can be deduced that only the real part ED^* is required, explicitly

$$\Re\{ED^*\} = X_I(k) [X_I(k+1) + X_I(k-1)] + X_Q(k) [X_Q(k+1) + X_Q(k-1)] \quad (\text{E.1.3})$$

Overall, considering the real and imaginary parts of $X(k)$, the mathematical equation to find α at each FFT frequency bin is

$$\alpha = \frac{X_I(k) [X_I(k+1) + X_I(k-1)] + X_Q(k) [X_Q(k+1) + X_Q(k-1)]}{[X_I(k+1) + X_I(k-1)]^2 + [X_Q(k+1) + X_Q(k-1)]^2} \quad (\text{E.1.4})$$

Since DD^* is positive and real, the numerator of Eq (E.1.4) determines the sign of α . From Eq (2.4.2), if $\alpha \leq 0$, then the whole FIR filter part of SVA can be ignored. Therefore an optimization would be to calculate the real part of the numerator from Eq (E.1.3) and test if it is negative or zero. This procedure is described in pseudo-code in Table E.1.1, where it is assumed that the intermediate result in the DSP48E slice register consists of a signed 48 bit word.

MACC Operation	Intermediate Result in Register
Multiply	$X_I(k)X_I(k+1)$
Accumulate	$X_I(k)X_I(k+1)$
Multiply	$X_I(k)X_I(k-1)$
Accumulate	$X_I(k)X_I(k+1) + X_I(k)X_I(k-1)$
Multiply	$X_Q(k)X_Q(k+1)$
Accumulate	$X_I(k)X_I(k+1) + X_I(k)X_I(k-1)$ $+ X_Q(k)X_Q(k+1)$
Multiply	$X_Q(k)X_Q(k-1)$
Accumulate	$X_I(k)X_I(k+1) + X_I(k)X_I(k-1)$ $+ X_Q(k)X_Q(k+1) + X_Q(k)X_Q(k-1)$
Logic Operation	Pseudo-code
Sign Test OR all zeros	if bitget(acc_reg,48) or (acc_reg==0) then (...) end if

Table E.1.1 Operations to find the numerator for α and determine if it less than or equal to zero

In the calculation of ED^* , one quantity is calculated which also features in Eq (2.4.2) and should be stored for later use. This is the “sum of the neighbour frequency bins”. Explicitly: $X(k+1) + X(k-1) = [X_I(k+1) + X_I(k-1)] + i[X_Q(k+1) + X_Q(k-1)]$.

If the sign test in Table E.1.1 is true then the FIR filter stage of SVA would be skipped for that bin. If the sign test is false then it would be necessary to carry out the full calculation of α by execution of the E/D division, which from the previously calculated result in Table E.1.1, would be:

$$\alpha = \frac{\text{acc_reg}}{[X_I(k+1) + X_I(k-1)]^2 + [X_Q(k+1) + X_Q(k-1)]^2} \quad (\text{E.1.5})$$

The next step in the optimisations would be to test whether the quotient from E/D is greater than or equal to $1/2$. Such a test could be performed by appending the first fractional bit of the quotient to the integer bits of the quotient and applying a

short-circuit OR operation to these bits. If this test is returned as true, then rather than multiply the “sum of the neighbour frequency bins” by $1/2$ as in Eq (2.4.2), a barrel shift one place to the right should be performed on the “sum of the neighbour frequency bins” (this is depicted in Table E.1.2). A barrel shift is computationally much simpler than a multiplication and therefore would allow a saving to be made on either computational time or power consumption depending on whether SVA is implemented as a procedure or as a pipeline.

If the test whether the quotient is greater than or equal to $1/2$ is returned as false then the fractional bits (except the MSB) from the quotient would be used to multiply with the real and imaginary parts of the “sum of the neighbour frequency bins”. The result from this multiplication would be subtracted from the real and imaginary parts of the frequency bin under test.

Real Part			
Operation	$\alpha \geq 0.5$ Intermediate Result	Operation	$0 < \alpha < 0.5$ Intermediate Result
Shift ¹	$X_I(k+1)$ $+X_I(k-1) \gg 1$	Multiply	$\alpha (X_I(k+1)$ $+X_I(k-1))$
Subtract	$X_I(k)$ $-(X_I(k+1)$ $+X_I(k-1) \gg 1)$	Subtract	$X_I(k)$ $-\alpha (X_I(k+1)$ $+X_I(k-1))$
Imaginary Part			
Operation	$\alpha \geq 0.5$ Intermediate Result	Operation	$0 < \alpha < 0.5$ Intermediate Result
Shift	$X_Q(k+1)$ $+X_Q(k-1) \gg 1$	Multiply	$\alpha (X_Q(k+1)$ $+X_Q(k-1))$
Subtract	$X_Q(k)$ $-(X_Q(k+1)$ $+X_Q(k-1) \gg 1)$	Subtract	$X_Q(k)$ $-\alpha (X_Q(k+1)$ $+X_Q(k-1))$

Table E.1.2 Calculations to filter FFT output bins with SVA 3-tap FIR filter when $\alpha > 0$

A summary of the optimisations is given in the following pseudo-code, where the E/D quotient is assumed to have a 48 bit wordlength with $\text{frac} = 18$ fractional bits:

```

if (bitget(48)) or (acc_reg== 0) then
    // alpha_numerator<= 0
     $X_a(k) = X(k)$ ; // No FIR stage applied
else
     $S_r = X_I(k+1) + X_I(k-1)$ ; // Sum of real parts of neighbours
     $S_i = X_Q(k+1) + X_Q(k-1)$ ; // Sum of imaginary parts of neighbours
     $\alpha = \text{acc\_reg} / (S_r^2 + S_i^2)$ ;
    if or (bitget( $\alpha$ ,48:18)) then
        // checking if  $\alpha \leq 0.5$ 
        acc_re=  $S_r \gg 1$ ;
        acc_im=  $S_i \gg 1$ ;
         $X_{Ia}(k) = X_I(k) - \text{acc\_re}$ ;
         $X_{Qa}(k) = X_Q(k) - \text{acc\_im}$ ;
    else
        //  $0 < \alpha < 0.5$ 

```

¹In Table E.1.2 \gg is used to symbolise a barrel shift of a binary number one place to the right

```

    accre =  $S_r \times \alpha$ ;
    accim =  $S_i \times \alpha$ ;
     $X_{Ia}(k) = X_I(k) - \text{acc}_{re}$ ;
     $X_{Qa}(k) = X_Q(k) - \text{acc}_{im}$ ;
  end if
end if

```

An alternative description of the optimisations is given by the flowchart in Fig E.1.1:

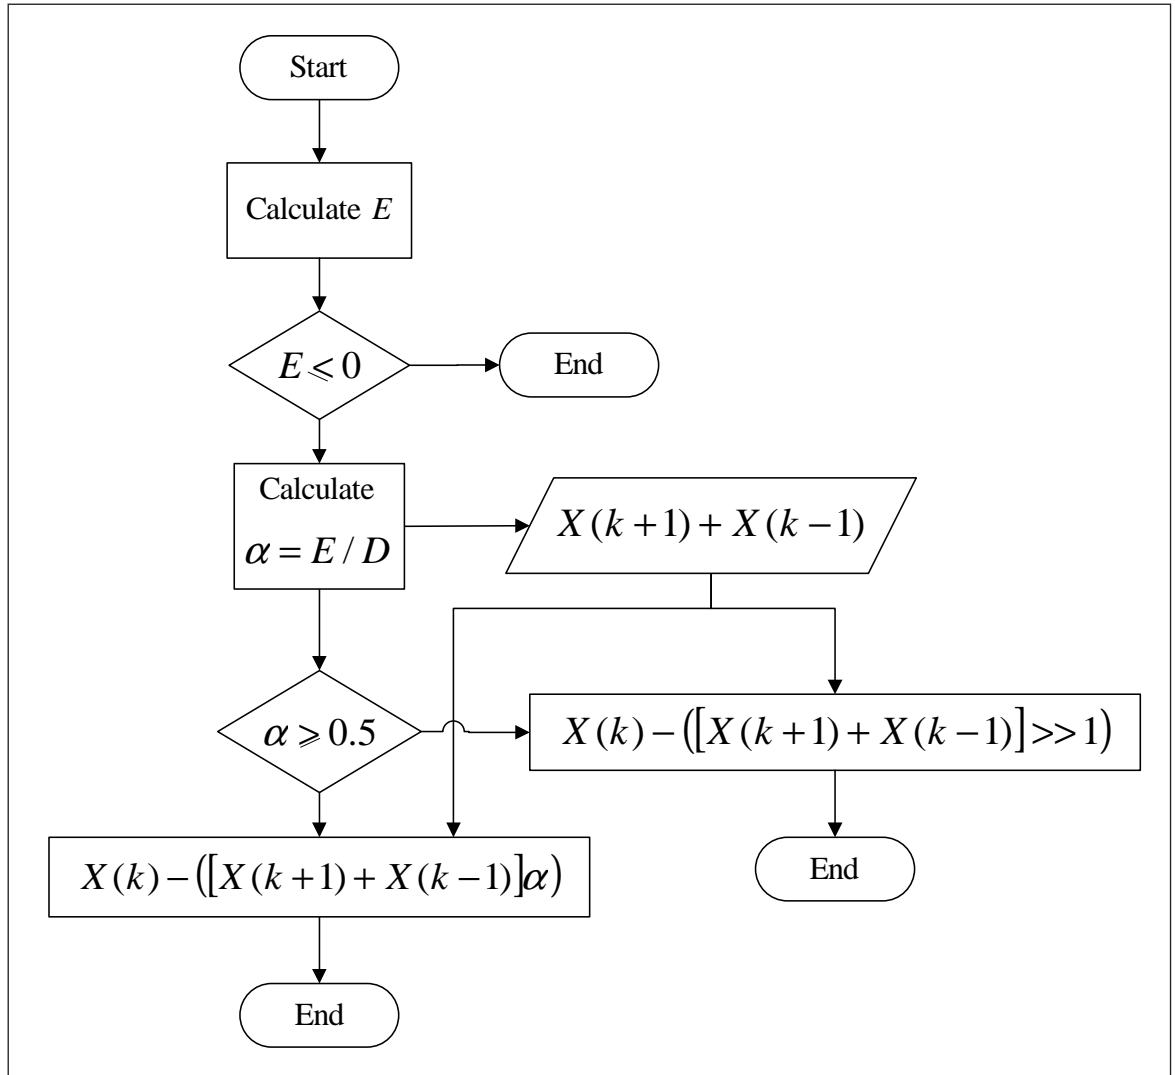


Figure E.1.1 Flowchart of the α parameter optimisations

APPENDIX F

SVA-on-FPGA

The evaluation of the SVA algorithm in a variety of scenarios in floating-point arithmetic revealed that SVA has advantages over conventional windowing techniques.

Further investigation of SVA in the fixed-point domain confirmed that these advantages were still present in a **realisable** fixed-point arithmetic system. The results obtained indicated that the fixed-point arithmetic, at the fixed-point precision studied, converged to the floating-point results. Additionally, it showed that the relative detector performance of the different window techniques did not alter after the conversion to fixed-point arithmetic. The fixed-point SVA simulations were designed with Xilinx® Virtex-5® FPGA implementation in mind.

The logical progression from the fixed-point simulations was the next step in the MBD methodology, that is, to translate the fixed-point code to an FPGA implementation. One of the important concerns in the SELEX Galileo MBD methodology is to ensure that fixed-point simulations are well designed to allow an easier transition down to the next lower level of abstraction. This instruction was followed thoroughly to aid the translation of SVA to FPGA implementation.

The main aim of coding SVA-on-FPGA was not to uncover its mathematical properties because the fixed-point simulations had already proved that an SVA detector can retain its floating-point performance in FPGA-like arithmetic. Rather, the primary aim was to discover physical information about SVA-on-FPGA such as the fabric resources required, the speed at which it can run, its latency and throughput.

F.1 Implementation Decisions

F.1.1 Target Hardware

There were various reasons why a Xilinx® Virtex-5® board was chosen as the FPGA board:

- the fixed-point simulations had already been designed with Virtex-5® implementation in mind;
- the hardware resources available at SELEX Galileo;

- one of the MBD tools used, System Generator[®], enables generation of the HDL code shaped for Xilinx[®] boards.

Incidentally it is possible to generate generic HDL for other FPGA manufacturer boards using System Generator[®]. However since System Generator[®] is a Xilinx[®] tool, it therefore produces optimised HDL code for Xilinx[®] boards. The specific Xilinx[®] board used was an evaluation card with model designation Virtex-5[®] LXT ML505 (device XC5VLX50T). For SVA-on-FPGA, the important features of this general purpose board were:

Block RAM According to the Virtex-5[®] datasheet[144], a ML505 board has a total of 2160 kB available in block random access memory (RAM).

DSP48E Slices These resources are tailored for high-speed DSP. They feature a 25×18 two's complement multiplier and a 48 bit adder. An ML505 board contains 48 of these DSP resources.

F.1.2 Simulation Language

A recommendation from the SELEX Galileo MBD methodology is to build models for FPGA implementation with the use of Simulink[®] and System Generator[®] which includes a Xilinx[®] blockset library in Simulink. This set of tools permits hardware-in-the-loop test and verification via a Joint Test Action Group (JTAG) header. Therefore the tools chosen for creating the eventual FPGA bit file download were Simulink[®] and System Generator[®].

F.2 Coding Methodology

In a similar fashion to the procedure followed for the fixed-point SVA model, the process of SVA-on-FPGA creation was broken down into distinct phases.

F.2.1 Phase 1

The first phase of SVA-on-FPGA was to translate the fixed-point SVA code into a graphical representation in Simulink with the use of the Xilinx[®] blockset. The decision was made to leave support functions like the FFT out of the code and represent them with Simulink blocks. This was because the objective was to uncover the characteristics of SVA-on-FPGA and this was achieved with the SVA algorithm isolated on the FPGA.

Fig F.2.2 is a screenshot of the model produced during phase 1. At this level the diagram shows the model was composed mostly of native Simulink blocks. The obvious exceptions are the System Generator[®] tool icons. Fig F.2.3 was obtained by drilling down into “Block containing Xilinx blocks” in Fig F.2.2. This diagram shows the actual Xilinx[®] blocks used to implement SVA-on-FPGA. The boundaries of the FPGA are

defined by the “Gateway in” and “Gateway out” blocks.

As mentioned in section D.4 SVA naturally lends itself to a two stage process for implementation:

1. calculation and limitation of the α parameters (which involves a division operation);
2. use of the α parameters in a type of FIR filter.

The two stages of the algorithm are apparent in Fig F.2.3. The labeled boxes: 1. calculate the dividend and divisor for the first stage and 2. execute the division and limit α values for the second stage.

The phase 1 design was simple. Factors and parameters such as N_{FFT} were kept to a mere fixed length of 32. The phase 1 model also opted to use a buffered, synchronous pipeline architecture. That is, registers were placed between stages in the pipeline and they were clocked synchronously. This system worked well as it was built around an FFT processor where the output was essentially downsampled by a factor of N . A diagram of the pseudo-algorithm showing the pipeline-flow is shown in Fig F.2.1. The pipelined implementation of SVA ran at the same speed as the input to the FFT processor and operated on one FFT frequency bin at a time before the arrival of the next output FFT frame.

The pipeline architecture produced savings on FPGA hardware resource usage because not all FFT frequency bins had to be operated on in parallel. Once the pipeline was filled with data, all the hardware resources allocated to the pipeline were fully employed. Naturally pipeline usage resulted in a better throughput of the SVA system than if no pipeline had been implemented. However, the drawback of the pipeline was the increased latency over an implementation which would operate on all frequency bins in parallel.

The pipeline architecture was chosen as a trade-off between FPGA resource usage and latency. An example of the problems which arise if the pipeline is removed is: if $N_{\text{FFT}} = 32$ and all the frequency bins are operated on in parallel, 640 DSP48E slices would be required¹. This is far in excess of the 48 DSP48E slices available on the ML505 Virtex-5[®] Evaluation Platform. The Xilinx[®] blocks which were used to generate HDL code for the pipeline and later re-order the data within it are highlighted in Fig F.2.3.

The pipeline was not a simple flow from beginning to end. It required extra effort to accommodate the division function. The IP chosen for the division function derived from the Xilinx[®] blockset library as this was the most likely to be an efficient implementation. The difficulty with this division IP block was that it could not be placed as a piece of logic between registers in a simple pipeline. This was because the IP is implemented internally as a loop to minimise FPGA resource requirements. It is however an efficient implementation of a loop because it can receive and output data at bursty intervals through a carousel-like operation. In order to place this piece of IP within the pipeline architecture it was necessary to achieve a constant number of clock cycles between data going in to the divider IP and calculated data coming out. The condition for this, according to [145], was “the number of cycles between new inputs must be mutually prime with the number of pipe stages in the iterative engine”. In the phase 1 model the data received from the pipeline in front of the divider was held with the use of a first-in-first-out (FIFO) until the divider was ready for it. The divider was clocked at three times the FFT input data rate which was mutually prime with the number of pipe stages in the implemented divider. The clock rate of the divider ensured data was output (after being downsampled) at the original pipelined data rate.

An alternative to the increase in the clock rate of the divider would be to use three times as many divider blocks or to slow the entire system down to the “natural” throughput of the divider block. Therefore either three times as many resources would be required to perform the division (along with extra handling of the data to ensure the inputs would be evenly distributed across all the dividers) or a slower data throughput would need to be accepted (but with greater control over the data flow via a state machine). More discussion on this trade-off is given in section F.5 where it is suggested that it would be necessary to choose from one of these alternatives to eliminate the multiple clock rates from the system. In brief this is because it is undesirable for the detection component of a DRx in an ESM system to run with the use of multiple clock rates.

A further point to note from Fig F.2.3 is a box which highlights data stored for the second part of the algorithm. The reason for the inclusion of this in the model was to recall data from the start of the data flow to be used again in the second stage of the algorithm to prevent repeat calculation. FIFOs were used as the elements of memory in this design, however it is noted in section F.5 this is a potential weakness. Briefly

¹This number was obtained with the use of Xilinx[®] Resource Estimator on the model multiplied by the number of frequency bins

this is because if a FIFO within a pipeline architecture becomes corrupt, the system would thereafter fail. An improved design would be to use two banks of RAM in place of each FIFO. Instead of data removed from the bottom of the pile in a FIFO, data would be written to one bank whilst read from the other bank. At the end of each bank read/write operation the functionality of the banks would switch so that the one previously read from would be written to and *vice-versa*. The trade-off in this case would be an increase in system stability for an increase in FPGA resource usage and increase in latency (as the first RAM bank would require to be filled before it could be read). A two-RAM-bank implementation would probably be included in a real-life DRx but since the aim in this SVA-on-FPGA investigation was mainly to characterise the SVA algorithm on FPGA, the RAM-banks task has been left as a future improvement.

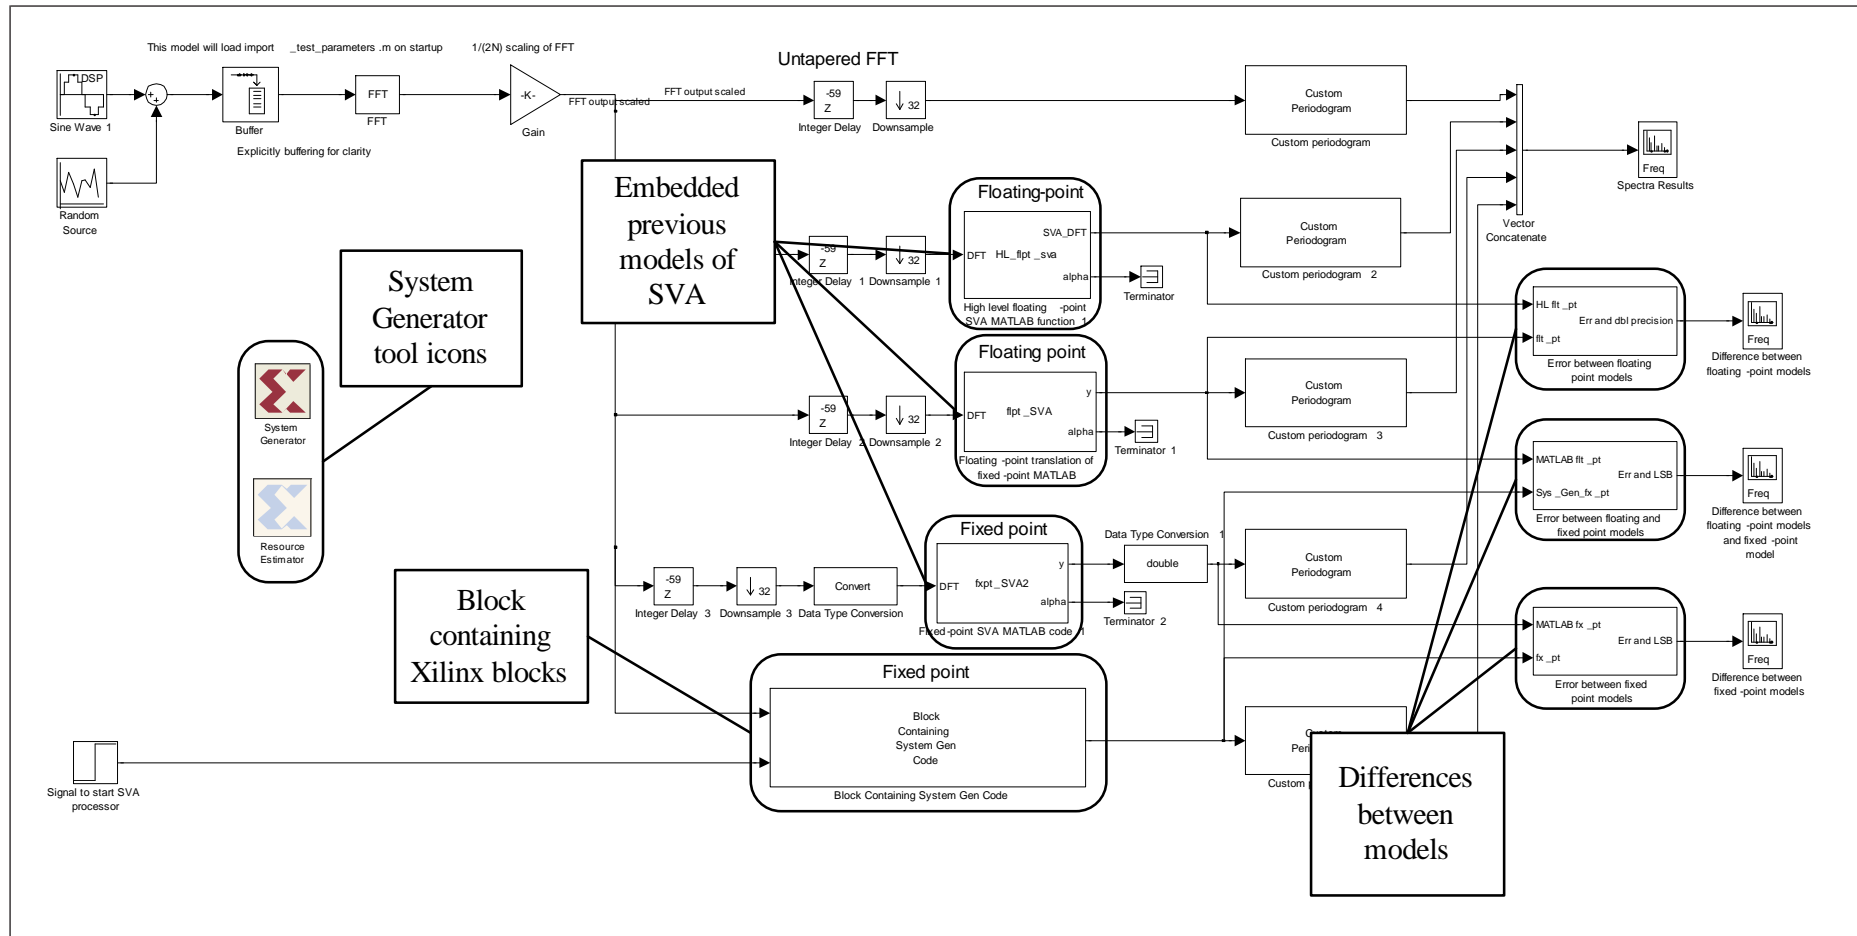


Figure F.2.2 Phase 1 screenshot of the highest level in the Simulink® hierarchy

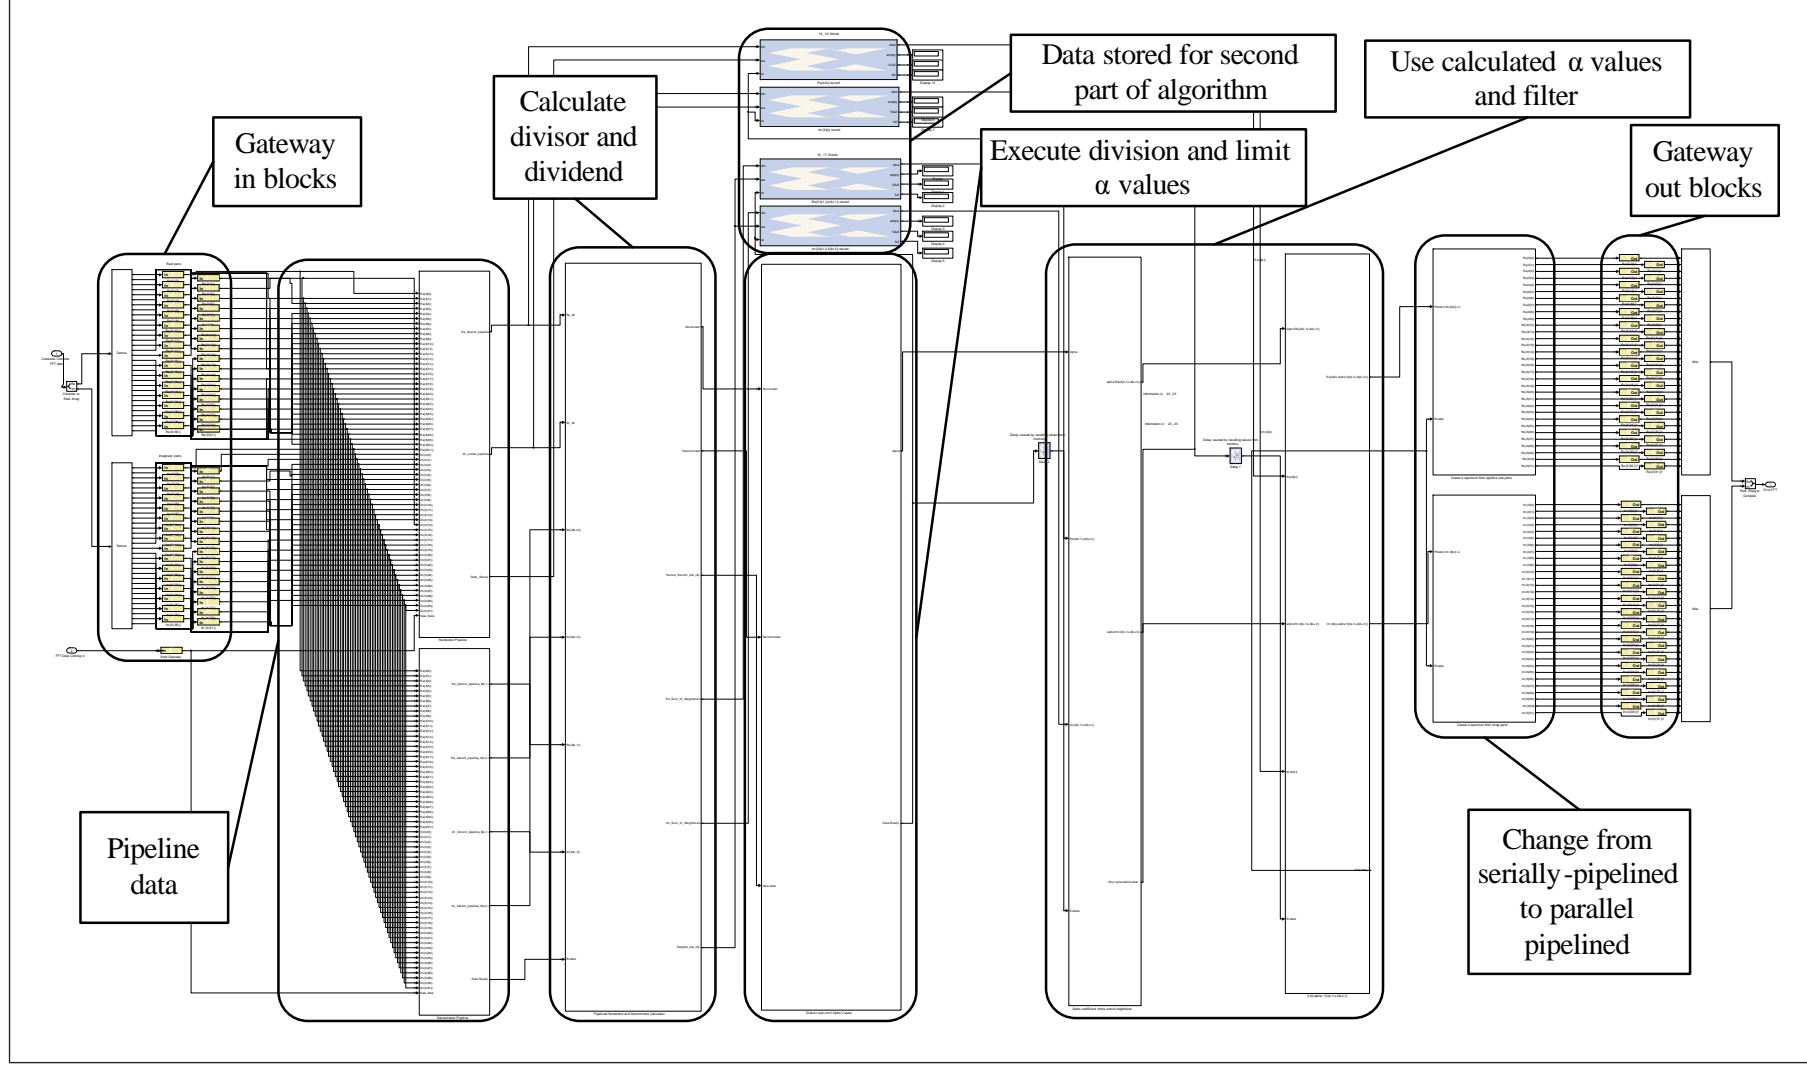


Figure F.2.3 Phase 1 screenshot of the top level of the Xilinx® blocks

F.2.2 Phase 2

In the second phase of production of SVA-on-FPGA, the Xilinx® ISE design tools were used extensively to give an indication of the FPGA resources required by the algorithm. This information was used to refine the phase 1 model to fit in the target hardware.

An extract of the synthesis report obtained from the Xilinx® ISE design suite for the phase 1 model is listed below in Table F.2.1. Table F.2.1 clearly shows that the phase 1

Device utilization summary:

Selected Device : 5vlx50tff1136-1

Slice Logic Utilization:

Number of Slice Registers:	4123 out of 28800	14%
Number of Slice LUTs:	2768 out of 28800	9%
Number used as Logic:	1157 out of 28800	4%
Number used as Memory:	1611 out of 7680	20%
Number used as SRL:	1611	

Slice Logic Distribution:

Number of LUT Flip Flop pairs used:	4782	
Number with an unused Flip Flop:	659 out of 4782	13%
Number with an unused LUT:	2014 out of 4782	42%
Number of fully used LUT-FF pairs:	2109 out of 4782	44%
Number of unique control sets:	8	

IO Utilization:

Number of IOs:	2307	
Number of bonded IOBs:	2306 out of 480	480% (*)

Specific Feature Utilization:

Number of BUFG/BUFGCTRLs:	1 out of 32	3%
Number of DSP48Es:	10 out of 48	20%

WARNING:Xst:1336 - (*) More than 100% of Device resources are used

Table F.2.1 Extract from phase 1 model synthesis report

design used too many “bonded IOBs” or in other words too many IO pins. The effect of this was to compel the mapping stage to fail in the design flow. Upon closer inspection the reason for the design device over-utilisation in the phase 1 model became clear. The design attempted to transfer the full FFT output onto the FPGA board from an off-board FFT processor. In actual fact this transfer would be an artificial situation because in a real-life DRx the SVA module would be accommodated in the same FPGA as the FFT processor and the route from one to the other would not require IO pins. Therefore the solution in the phase 2 model was to move the pipeline creation out of the SVA module.

The refined phase 2 model is shown in Figs F.2.5 and F.2.6. Fig F.2.5 is a top-level system description and Fig F.2.6 shows the SVA algorithm. The diagrams show the pipeline creation and reorganisation were moved out of the FPGA in the phase 2 model. The phase 2 model was also used in the subsequent test phase, where it allowed a comparison to be made between the different implementations of the SVA algorithm to ensure the functionality remained the same. These implementations were:

- high-level floating-point SVA
- fixed-point SVA
- low-level floating-point translation of fixed-point SVA
- System Generator[®] SVA.

Within the phase 2 model there was also the provision to allow a direct comparison to be made against Hann-windowed and rectangular-windowed FFTs which operated on the same data.

The phase 2 model, had a much more favourable synthesis report. This design successfully mapped onto the target device and successfully completed the place and route process. The device utilisation report, shown in Table F.2.2, for the phase 2 model was extracted from the place and route report rather than the synthesis report. Af-

Device Utilization Summary:

Number of BUFGs	1 out of 32	3%
Number of DSP48Es	24 out of 48	50%
Number of External IOBs	147 out of 480	30%
Number of LOCed IOBs	0 out of 147	0%
Number of RAMB18X2s	1 out of 60	1%
Number of RAMB18X2SDPs	4 out of 60	6%
Number of RAMB36SDP_EXPs	2 out of 60	3%
Number of Slice Registers	1918 out of 28800	6%
Number used as Flip Flops	1918	
Number used as Latches	0	
Number used as LatchThrus	0	
Number of Slice LUTS	1498 out of 28800	5%
Number of Slice LUT-Flip Flop pairs	1896 out of 28800	6%

Table F.2.2 Extract from phase 2 model place and route report

ter the phase 2 model successfully transitioned through the place and route stage, hardware co-simulation became possible. This meant it became a hardware-in-the-loop process during which the model and its hardware co-simulation were tested with specially-designed test vectors. This process is depicted in Fig F.2.4, where the hardware co-simulation is a parallel branch to the phase 2 model.

The next phase discusses in more detail tests of the phase 2 model and its hardware co-simulation.

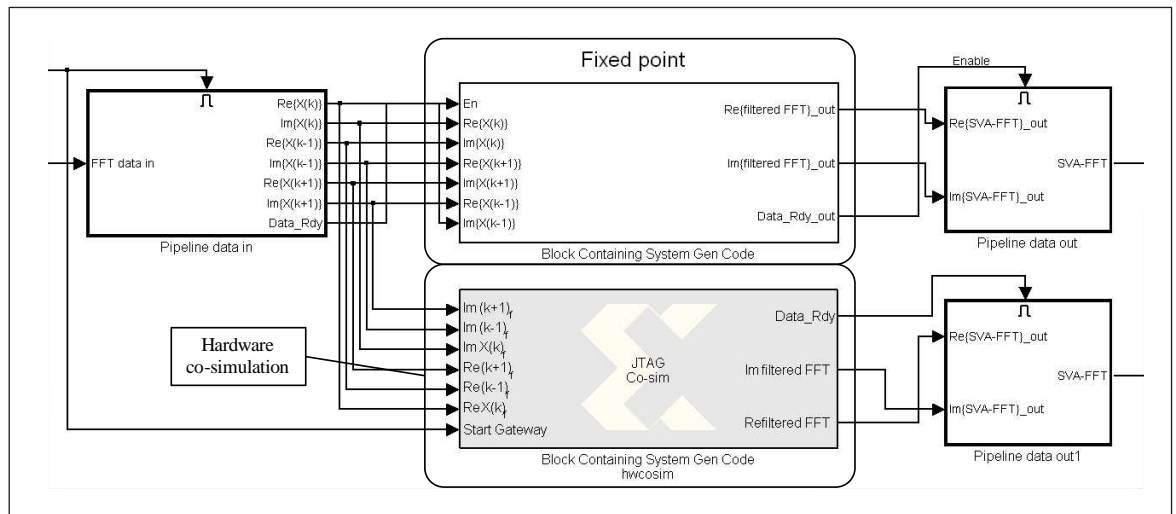


Figure F.2.4 SVA hardware co-simulation

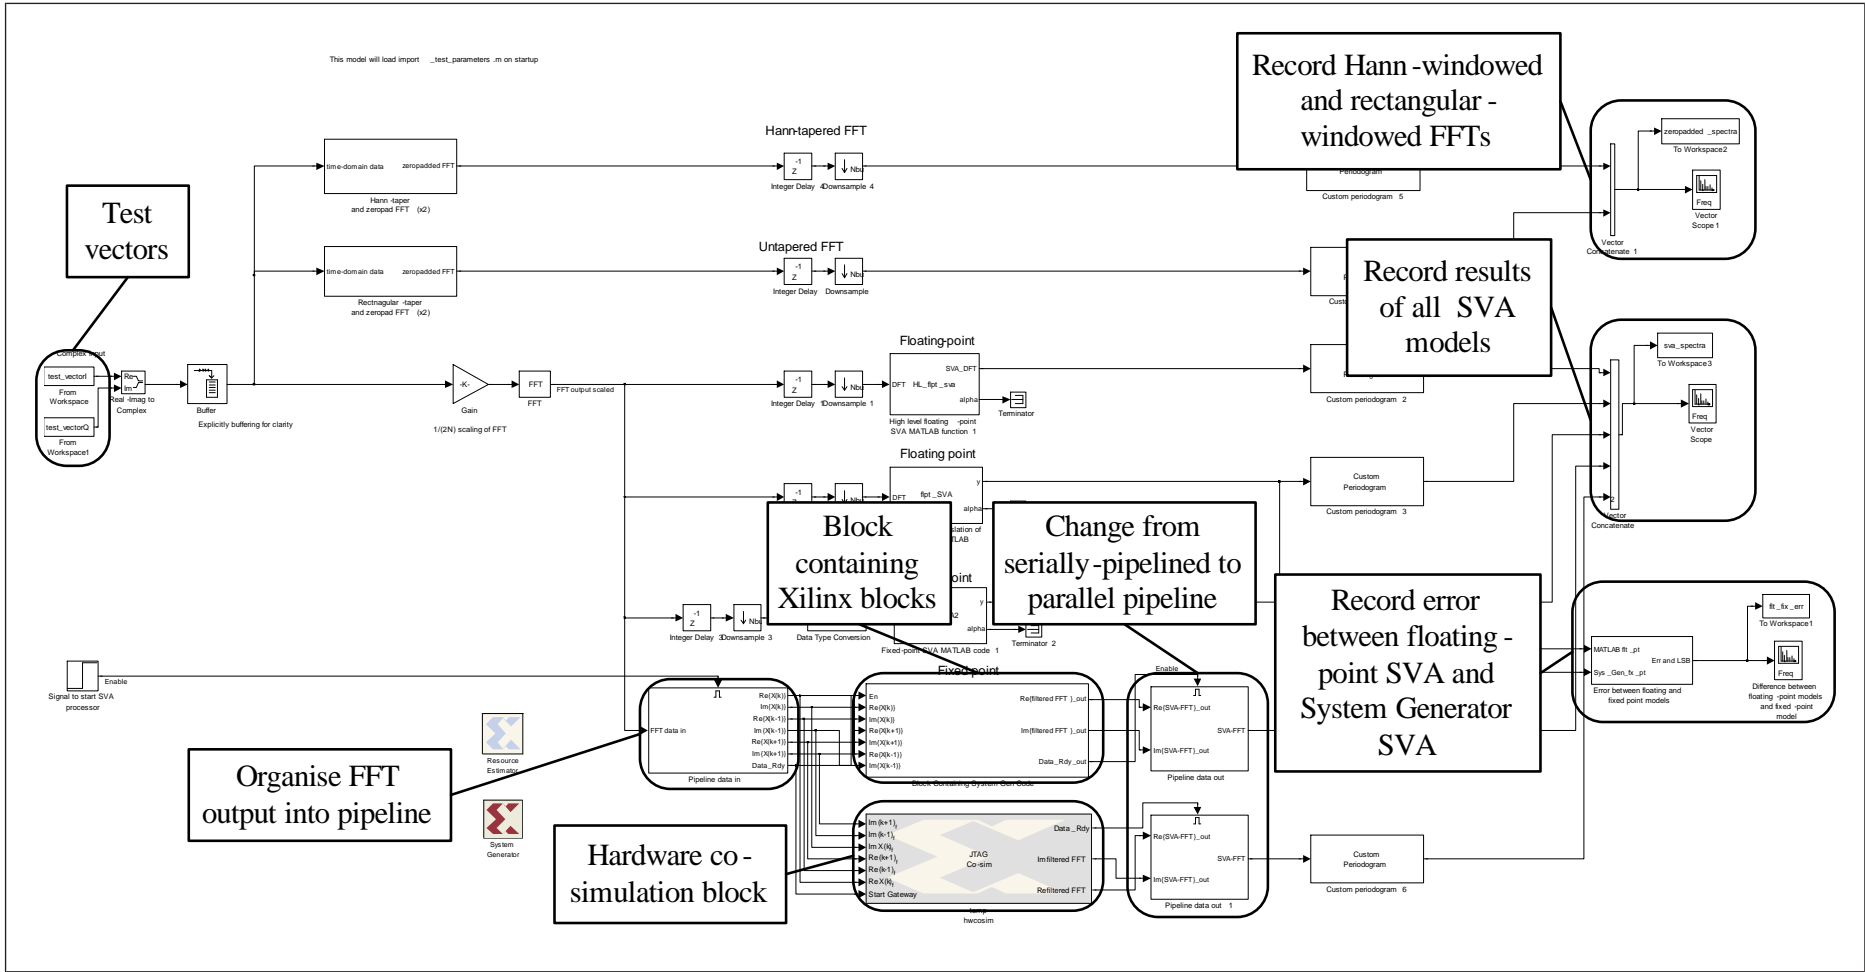


Figure F.2.5 Top-level of the phase 2 system

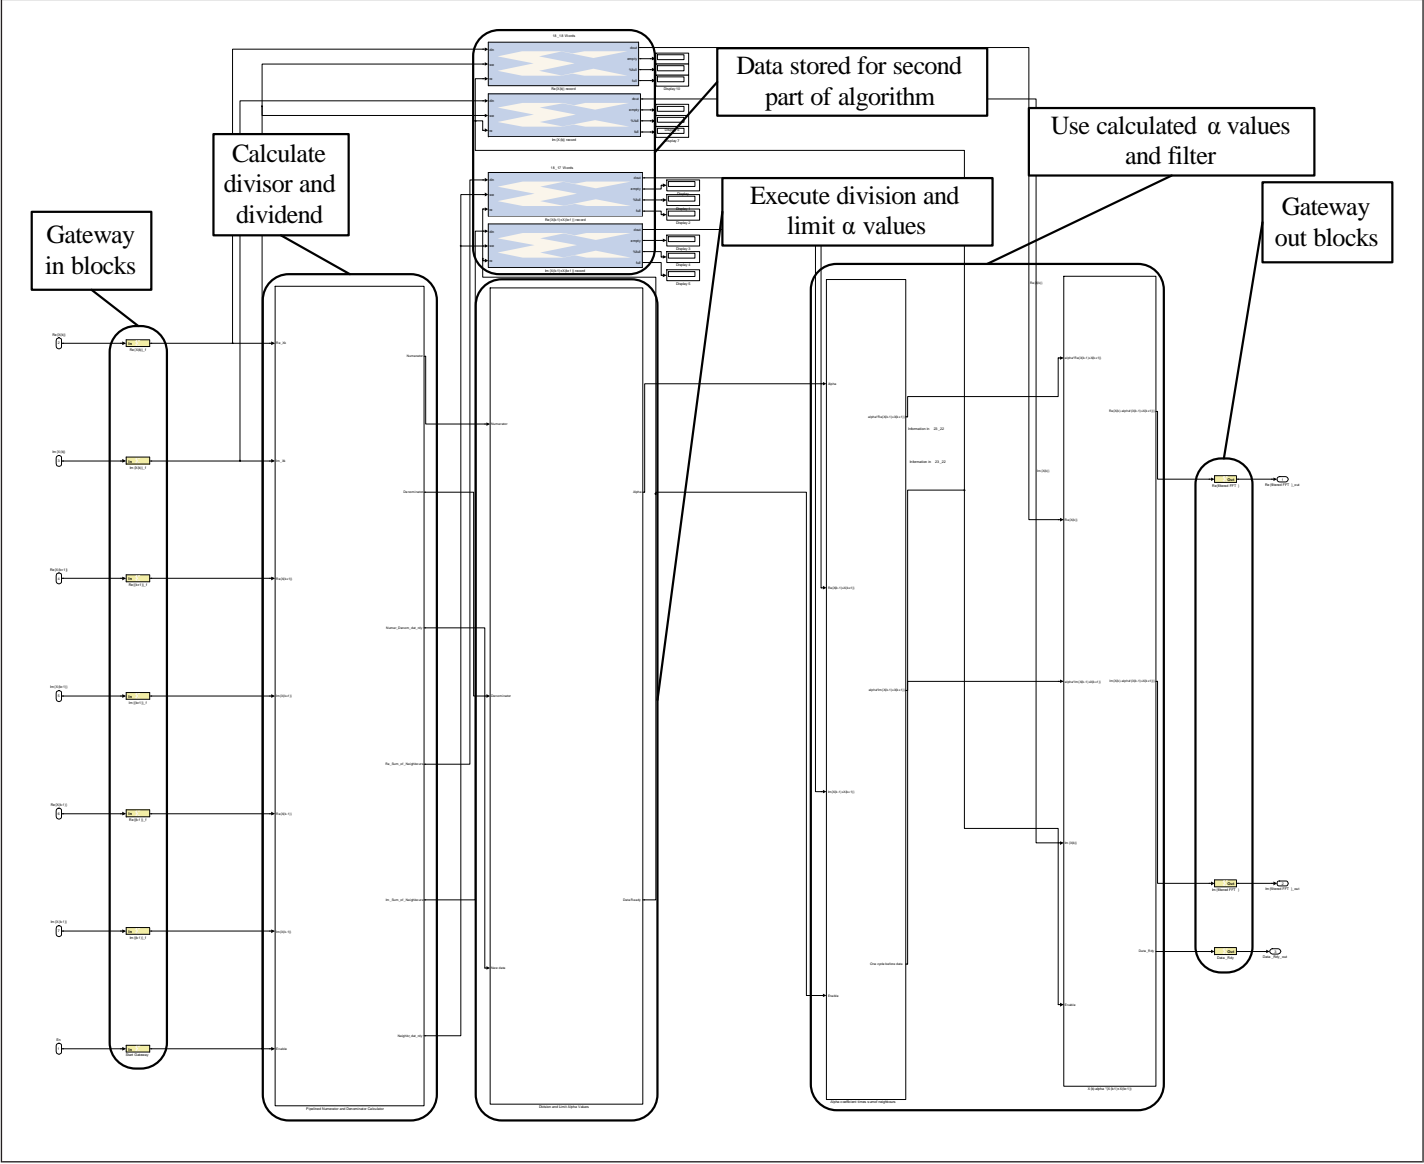


Figure F.2.6 Phase 2 Xilinx® blocks which perform SVA-on-FPGA

F.2.3 Phase 3

Test and verification are the most important and time consuming stages of implementation of any algorithm in FPGA. The use of a MBD methodology facilitated a reduction in the verification time because in the transition from one level of abstraction to the next the models were verified against the previous layer. In addition the graphical synthesisable models were generally bit-true and cycle-true representations of the actual algorithm in FPGA and if there were any discrepancies between the model and the actual implementation in FPGA, the use of hardware co-simulation brought attention to them. Nevertheless the test and verification stages are still extremely important in MBD and this section describes the methods used to test and verify the phase 2 model and its hardware co-simulation.

Tests occurred in two different types of test:

1. Particular test vectors were used to explore the differences between the magnitude-squared of the output from the SVA implementations and the rectangular-windowed and the Hann-windowed FFTs.
2. The key Monte Carlo simulations from section 3.2 and 3.4 were repeated on FPGA.

F.2.4 First Type of Test

SVA-on-FPGA, which hereafter in this section is taken to collectively refer to the phase 2 model and its hardware co-simulation counterpart, was tested against previous SVA models and the extreme values the α parameter can assume ($\alpha = 0$ and $\alpha = 0.5$). Essentially, the extreme α values created a comparison between: (1) SVA-on-FPGA; (2) rectangular-windowed results and (3) Hann-windowed results. According to the multi-apodization interpretation and after appropriate normalisation, the SVA-on-FPGA spectral values should always be **less than or equal to** the envelope of the spectral values produced by the rectangular-windowed FFT or Hann-windowed FFT for the same input data.

The test vectors used for the comparison were sinusoidal waveforms. They were chosen because they represent the extreme case to test the dynamic range of SVA-on-FPGA and additionally they represent the signal type which the SVA algorithm is best suited to detect. The sinusoidal test vectors acted as corner cases and therefore it was not necessary to repeatedly test SVA-on-FPGA with an exhaustive set of signal types.

In total SVA-on-FPGA was tested against the:

- high-level floating-point SVA from chapter 3;
- fixed-point SVA from section D;

- low-level floating-point SVA translation of the fixed-point code from section D;
- floating-point rectangular-windowed FFT;
- floating-point Hann-windowed FFT.

All of the first types of test were conducted with the use of $N = 128$. The first type of test is also referred to as the difference tests in this chapter.

F.2.5 Sinusoidal Input (Complex Exponential)

The Fourier transform is the matched filter for sine/cosine waves. Therefore waveforms with a similar structure will produce the strongest response at the output of a Fourier transform processor. In any envisaged detector which uses SVA, the SVA operates on the output of an FFT, which is a type of discrete Fourier transform. Therefore in the first type of test when sinusoids were processed by the FFT, they produced the strongest response at the output of the FFT and hence tested the SVA algorithm in an extreme case.

One of the purposes of testing SVA-on-FPGA with sinusoids was to check that there were no flaws in the design after it had been converted from the fixed-point implementation. In particular, the fixed-point design and SVA-on-FPGA were created to ensure there were no overflow errors whilst underflows were minimised. In effect a full scale (FS) sinusoid was used with no provision for headroom.

Specifically, the sinusoidal tests involved frequency increments of a complex exponential through the interval f_s/N^2 . The results for all bins would be identical. In the simplest case of no offset and no window function (i.e. a rectangular window) the projection of the frequency response of the complex exponential occurred at an FFT bin centre and had a magnitude equal to N times the time-domain amplitude of the complex exponential. However when there was a frequency offset between the FFT bin and the complex exponential the cusping loss caused the apparent magnitude to reduce from this maximum (see section 1.4).

As SVA-on-FPGA was a translation of the fixed-point SVA, the previous fixed-point decisions from section D.4 were valid. For example the number of fractional bits (18 bit) was chosen to be the same as the wordlength.

The time-domain amplitude of the complex exponential was pegged to the ADC FS value of $1 - 2^{-17}$ and the gain post-FFT was set to $\frac{1}{2 \sum_{n=0}^{N-1} |w(n)|^2}$, where $w(n)$ are the window coefficients. In real-life the FS value would represent the maximum voltage level of the ADC.

The post-FFT gain was set to ensure that the real and imaginary parts of the output of

²The frequency was changed such that the offset from a bin centre was incremented by $0.1 \times f_s/N$ through $[-0.5, 0.5] \times f_s/N$

the FFT were not be able to exceed the value $|0.5 - 2^{-18}|$, thus 18 bit words with 18 bit fractional bits. The post-FFT gain is similar to the situation where each individual butterfly within the FFT is divided by two in order to ensure the algorithmic gain is 1 (i.e. 0 dB). The phase of the complex exponential test vector was selected at random at the start of each test run.

In the course of these difference tests spectral estimates were recorded for the:

- Hann-windowed FFT;
- rectangular-windowed FFT;
- high-level floating-point SVA;
- fixed-point SVA-filtered FFT;
- low-level floating-point translation of the fixed-point SVA;
- SVA-on-FPGA.

The Hann-windowed and rectangular-windowed FFTs were twice-zero padded to achieve a better picture of their overall spectral response. This was done simply to give an accurate visual comparison to the SVA spectral estimates. In the results which follow in Fig F.2.7 the frequency-domain response envelope of the rectangular and Hann window functions are also drawn over the top as the upper bound for the rectangular-windowed or Hann-windowed FFTs.

The predicted outcome was that the SVA results should be **less than or equal to** the rectangular-windowed spectral response envelope and the Hann-windowed spectral response envelope centred on the frequency of the complex exponential. The envelopes are important theoretical constructs to consider because although some spectral values that result from the use of a rectangular-windowed or Hann-windowed FFT can **appear** below the envelopes, this occurs because the FFT samples only a finite number of points from the DTFT. In a real-life DRx, the frequency of a target sinusoid signal would initially be unknown and it would not be possible to state with any certainty that there had been no cusping loss and no measured spectral leakage, therefore the safe option must be assumed which is the spectral response envelope of the rectangular-windowed or Hann-windowed FFT.

One set of plots is shown in Fig F.2.7 for a particular frequency offset value. The results show that the target signal frequency does not line up perfectly with the nearest FFT bin. Panel (a) compares the various SVA models from the MBD methodology against the spectral estimation results from the application of a rectangular-windowed FFT to the same data. The first positive conclusion from the test results was that the phase 2 SVA model and its hardware co-simulation were always in perfect agreement. The implication was that the phase 2 model of SVA translated in a bit-true and cycle-true manner on FPGA with the use of the Xilinx[®] tools. Therefore it was correct to group

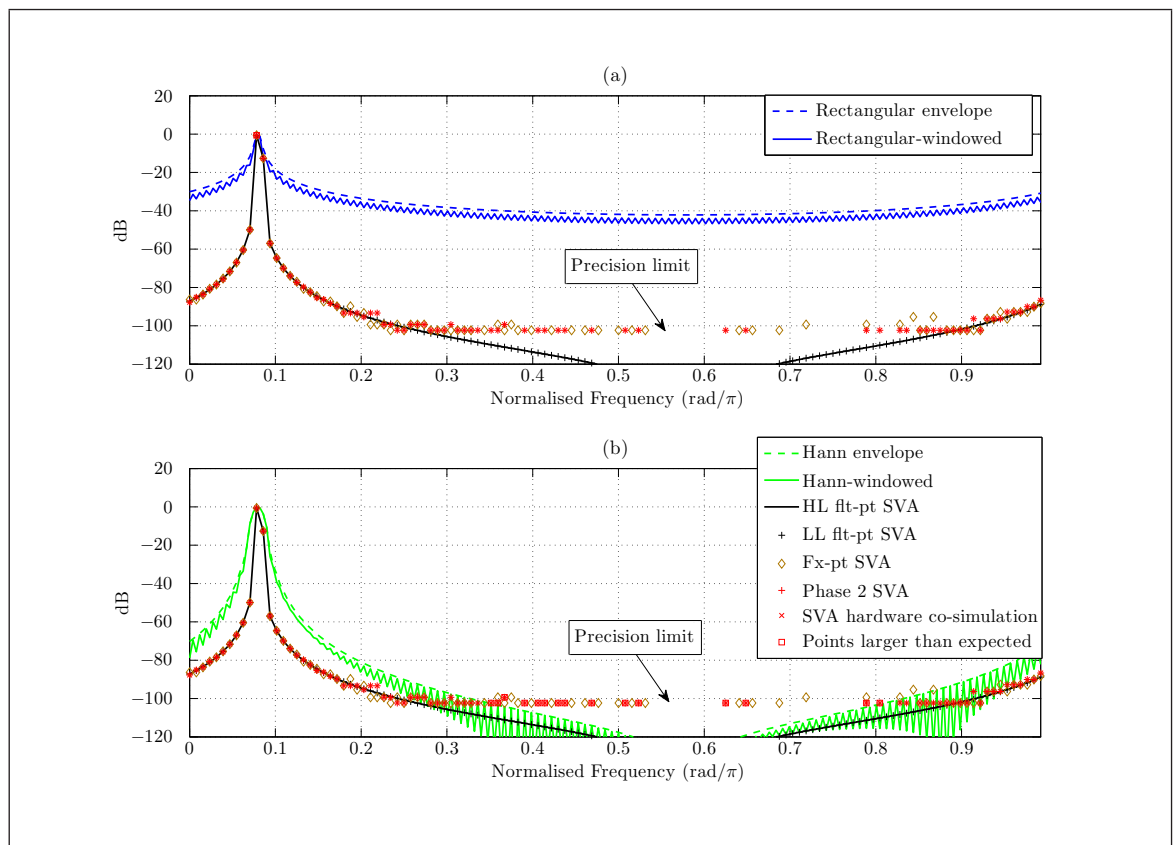


Figure F.2.7 Spectral estimates of sinusoidal wave input to SVA models compared against: (a) the rectangular-windowed spectral estimate; (b) Hann-windowed FFT spectral estimate

the phase 2 model and its hardware co-simulation as SVA-on-FPGA since there was no distinction between the results.

The twice-zero-padded rectangular-windowed spectral estimate and the overlaid frequency-domain response envelope of the rectangular window centred on the actual target frequency are drawn in panel (a). Any data point obtained from SVA-on-FPGA which did not meet the **less than or equal to** criterion is marked with a red box. The only instances where this occurred in all tests against the rectangular-windowed FFT was near to the spectral peak of the sinusoidal signal. However the difference between the envelope and SVA-on-FPGA was extremely small and less than one LSB and can safely be ignored.

Panel (b) of Fig F.2.7 shows the same selected SVA model results in comparison to the Hann-windowed spectral estimate and its envelope. One of the advantages of SVA over the Hann window function is clear, that is, the mainlobe of the spectral peak is narrower. However there are many red boxes at the ~ -100 dB level where the floating-point Hann-windowed spectral estimate appears better than SVA-on-FPGA. This lower bound effect of SVA-on-FPGA is also present in panel (a) however the test did not draw a red box around these points because they are still much lower than the sidelobes of the rectangular window envelope. This effect is explained by the

approximate dynamic range for SVA-on-FPGA:

$$\text{DR} = 20 \log_{10} \left(\frac{2^{-1}}{2^{-18}} \right) \approx 102 \text{ dB} \quad (\text{F.2.1})$$

The spectral peak of the sinusoid has a maximum possible frequency-domain magnitude of ~ 0 dB, minus any cusping loss. Therefore the value from equation (F.2.1) is an approximate theoretical dynamic range limit expected for SVA-on-FPGA. It is approximate because full precision was not used in the full datapath of SVA-on-FPGA and thus some accuracy was lost in places and this appears in Fig F.2.7 as data points slightly above the ~ 102 dB line.

The fixed-point SVA data points do not exactly match the SVA-on-FPGA data points because the division operation was implemented differently in each case. These small discrepancies were only noticeable at the very low dynamic range.

F.2.6 Sinusoid (Complex Exponential) Plus AWGN Input

In this difference test the input was changed to consist of a sinusoid with a frequency which occurred at a bin centre and AWGN. The addition of bandlimited AWGN to the input waveform caused a rise in the power at each frequency bin of the spectral estimates over the previous difference test. If this was unaccounted for, it could force the value of the real or imaginary part of the FFT frequency bins after the post-FFT gain to be greater than the $|0.5|$ limit and thereby cause an overflow. Therefore in this difference test a cautious approach was taken so that the time-domain amplitude of the sinusoid plus AWGN input was restricted such that it could not exceed $1 - 2^{-17}$. This was more cautious than necessary as AWGN power is theoretically spread evenly across all frequencies in the band.

In a real-life system the approach is usually taken to leave some headroom between the ADC FS and the largest sinusoidal signal which is expected. The headroom accommodate a large proportion of the tail of the AWGN distribution such that the likelihood of overflow becomes very small.

The parameters of the sinusoid were kept constant as the AWGN power level was varied. The noise power level was varied such that the SNR_{in} of the waveform before the FFT varied from 30 dB to -15 dB, with -1 dB increments between AWGN power levels.

The increase in the AWGN power level caused a reduction in the SNR_{in} but also a reduction in the dynamic range of values which SVA-on-FPGA needed to process. This was because whilst the maximum values were recognised with the post-FFT gain, the minimum values were determined by the approximate noise floor power level. As this floor was raised, the minimum value required to be represented by the fixed-point

precision became closer to the maximum value required to be represented and therefore the dynamic range required was reduced.

Similar to the previous difference test, the spectral estimates from the rectangular-windowed FFT and Hann-windowed FFT were twice-zero padded to get a better picture of their frequency response. With AWGN present, there were no theoretical spectral response envelopes for the Hann-windowed and rectangular-windowed FFTs, hence they are omitted from Fig F.2.8. In Fig F.2.8 the large sidelobes of the rectangular-

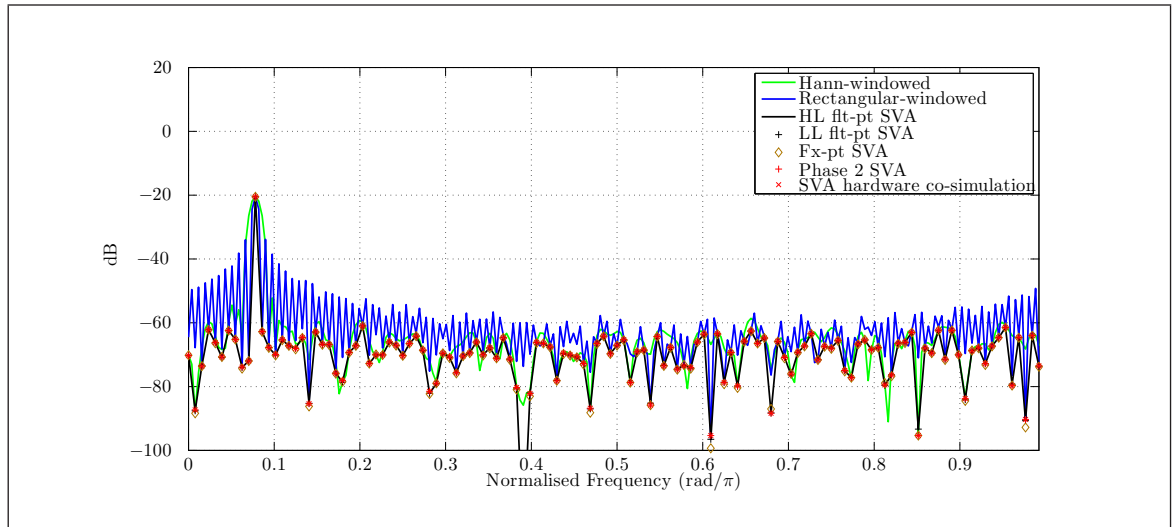


Figure F.2.8 Spectral estimates of a sinusoid plus AWGN input to SVA models compared against rectangular-windowed spectral estimate and Hann-windowed FFT spectral estimate

windowed spectral estimate and the wide mainlobe of the Hann-windowed spectral estimate are both evident. The SVA models overcome these two problems.

Another important point from the results was the phase 2 SVA model and its hardware co-simulation were again identical, which provided further evidence that the FPGA implementation was bit-true and cycle-accurate. Corresponding data points for all SVA models were very similar. This was because all the SVA models, whether floating-point, fixed-point or in FPGA, operated within a representable dynamic range. Most significantly SVA-on-FPGA was found to be nowhere greater than the spectral estimate values obtained from the rectangular-windowed or Hann-windowed FFT. In the results the lower limit of the fixed-point precision dynamic range ~ -100 dB was no longer visible as the noise floor component generally raised the spectral estimate points above this value.

From the results across the tested SNR_{in} range it was concluded that SVA-on-FPGA performed as well as the equivalent floating-point implementations when presented with a sinusoidal signal and AWGN input.

F.2.7 Two Sinusoids (Complex Exponentials) Plus AWGN Input

The last difference test examined the situation where multiple sinusoidal signals were present. This was an important situation to test because in the floating-point results in section 3.4 the SVA detector clearly outperformed the conventional windowing detectors. The only alteration to the previous “difference tests” for this set of test vectors was that the test vector input was composed of two complex exponential inputs and an AWGN component.

The parameters of one of the sine waves was kept constant as a control whilst the parameters of the other sine wave was changed to test two schemes of interest where the SVA algorithm would be expected to outperform the rectangular-windowed spectral estimate and/or the Hann-windowed spectral estimate. These two schemes were:

1. Two sine waves of similar amplitude (and also power) which were close in frequency. The idea was to see the advantages of the use of SVA in a sidelobe spectral leakage-dominated regime, much like that which produced the results from Fig 3.5.6.
2. Two sine waves of disparate amplitude (and also power) which were spaced apart in frequency. The idea behind this was to confirm the appearance of SVA spectral estimates in a noise-dominated scenario.

In these test vectors the noise floor was set so that the SNR_{in} (with respect to the control complex exponential) was high at 30 dB. This was to ensure scenario 1 from the above list was spectral-leakage-dominated.

The actual frequencies which were used for the two sine waves in units of frequency bins were: (1) $10.3 \times f_s/N$ for the control sinusoidal wave; (2) $11.8 \times f_s/N$ for the second sinusoidal wave in the first scheme and (3) $14.3 \times f_s/N$ for the second sinusoidal wave in the second scheme. The reason for the seemingly arbitrary coefficients was to ensure the frequencies were not harmonics of each other. The power ratio of the sinusoidal waves was SIR_{in} 0 dB and -34 dB, where the second sinusoid was considered as the signal.

The actual amplitudes were not of great importance however for the same reasons given in the previous difference tests, the maximum value that the superposition of the two sinusoidal waves plus AWGN was permitted to take was $1 - 2^{-17}$. Again, in a real-life DRx a headroom would be specified to mitigate against large power inputs.

As with the previous difference tests, the rectangular-windowed and Hann-windowed spectral estimates were twice-zero padded to show clearly the outline of their frequency response.

An example of the results from these test vectors is shown in Fig F.2.9. Panel (a)

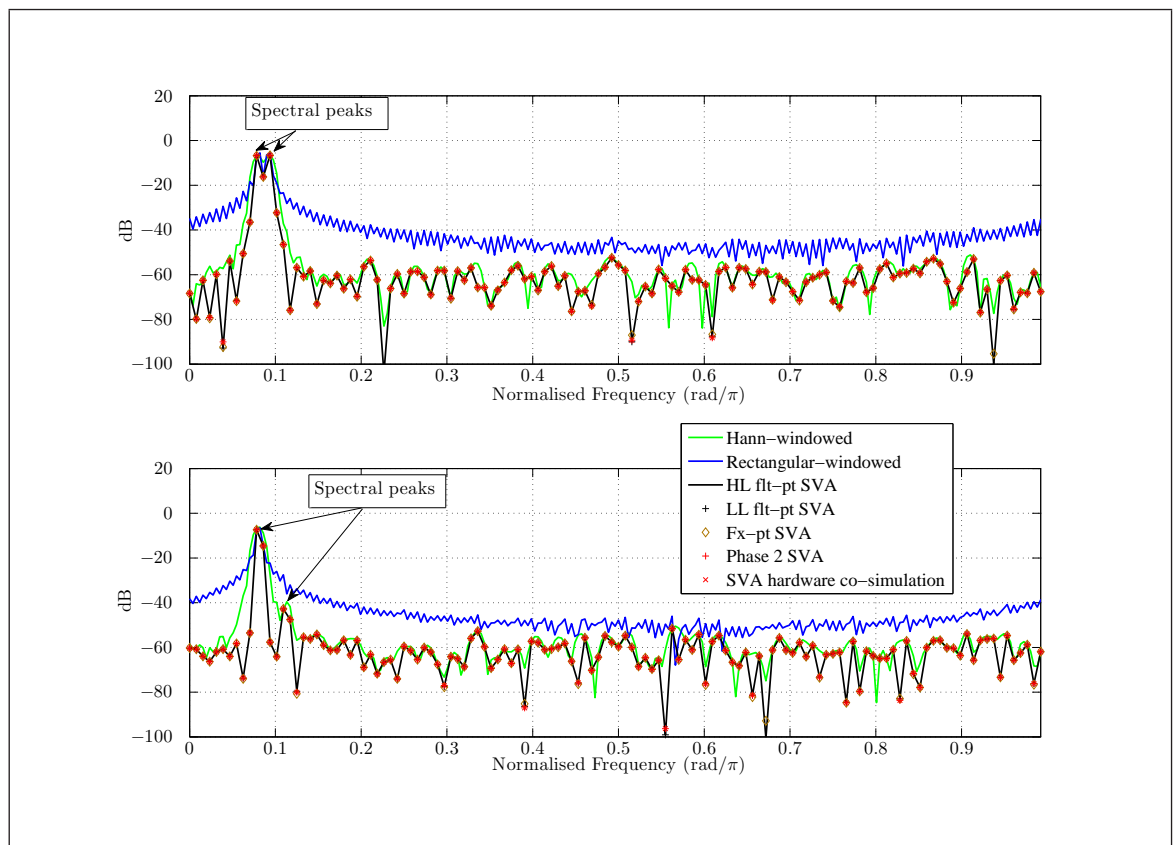


Figure F.2.9 Spectral estimates of two sinusoids plus AWGN input to SVA models compared against rectangular-windowed spectral estimate and Hann-windowed FFT spectral estimate for: (a) first scheme; (b) second scheme

shows the results of the first scheme from the above list and panel (b) shows the results of the second scheme. In panel (a) the two spectral peaks are distinguishable from the spectral estimates obtained with the use of the SVA model or the rectangular window. In panel (b) two spectral peaks are discernable from the spectral estimates obtained with the use of the SVA models or the Hann window. This further demonstrates the advantages of the use of SVA in general, but more importantly it shows that SVA-on-FPGA continued to be an improvement over conventional windows in an FPGA.

F.3 Second Type of Test

The first type of verification test was concerned with the functionality of SVA-on-FPGA through the examination of differences between the SVA models and windowed FFTs. The second type of verification test was concerned with the use of the SVA module on the actual FPGA in Monte Carlo simulations.

The Monte Carlo simulations repeated key scenarios from the floating-point SVA chapter 3. This was achieved solely with the use of the SVA module in FPGA whilst support functions remained in Simulink[®]. The results from the first type of verification test had shown that the phase 2 model and its hardware co-simulation produced identical

results and therefore affirmed that the hardware co-simulation was bit-true and cycle-accurate. Therefore there was no need to involve the phase 2 model in the Monte Carlo simulations.

The simulations were conducted through the execution of the hardware co-simulation over a large number of iterations. It was provided with data from Simulink® and the results were collected with the use of Simulink®.

The two key Monte Carlo simulations evaluated using SVA-on-FPGA were:

1. full DCW
2. high SNR_{out} -regime interference full DCW.

F.3.1 Monte-Carlo Simulation 1

The results from Fig 3.2.6 implied some conclusions about the SVA algorithm in general. Fig 3.2.6 indicates that SVA recovers the window loss suffered by bell-shaped windows. The first Monte Carlo simulation attempted to prove SVA-on-FPGA would demonstrate this beneficial feature by repeating the scenario under which the results in Fig 3.2.6 were obtained.

The ROC curve produced from SVA-on-FPGA was compared against ROC curves produced from: (1) a rectangular-windowed detector (2) a Hann-windowed detector and (3) a Chebychev-windowed detector.

Fig F.3.1 shows the results for this scenario. The graph shows the rectangular-windowed detector gave the best detector performance in this particular scenario and that the SVA-on-FPGA mostly recovered the window loss. Therefore SVA-on-FPGA successfully retained its window-loss-recovery characteristic seen in floating-point and fixed-point arithmetic. The next Monte Carlo simulation demonstrated SVA-on-FPGA was able to accomplish this whilst sidelobes were minimised.

F.3.2 Monte-Carlo Simulation 2

The second Monte Carlo simulation was designed to demonstrate the sidelobe spectral leakage minimisation properties of SVA were still present in SVA-on-FPGA. The scenario under which the ROC curve results from Fig 3.4.9 were obtained was particularly good as a demonstration of this property of SVA. Therefore this simulation was repeated with SVA-on-FPGA.

The simulation design and parameters mirrored those which produced the results in Fig 3.4.9. The results are shown in Fig F.3.2. The success of SVA-on-FPGA was apparent as the results were virtually the same as those obtained from floating-point SVA aside from negligible differences caused by quantisation.

The results for this scenario supported the idea that SVA-on-FPGA minimised sidelobe spectral leakage. Coupled with the recovery of the window loss, it was concluded that this implementation of SVA on FPGA afforded the same advantages seen in floating-point arithmetic simulations. Hence SVA-on-FPGA was verified.

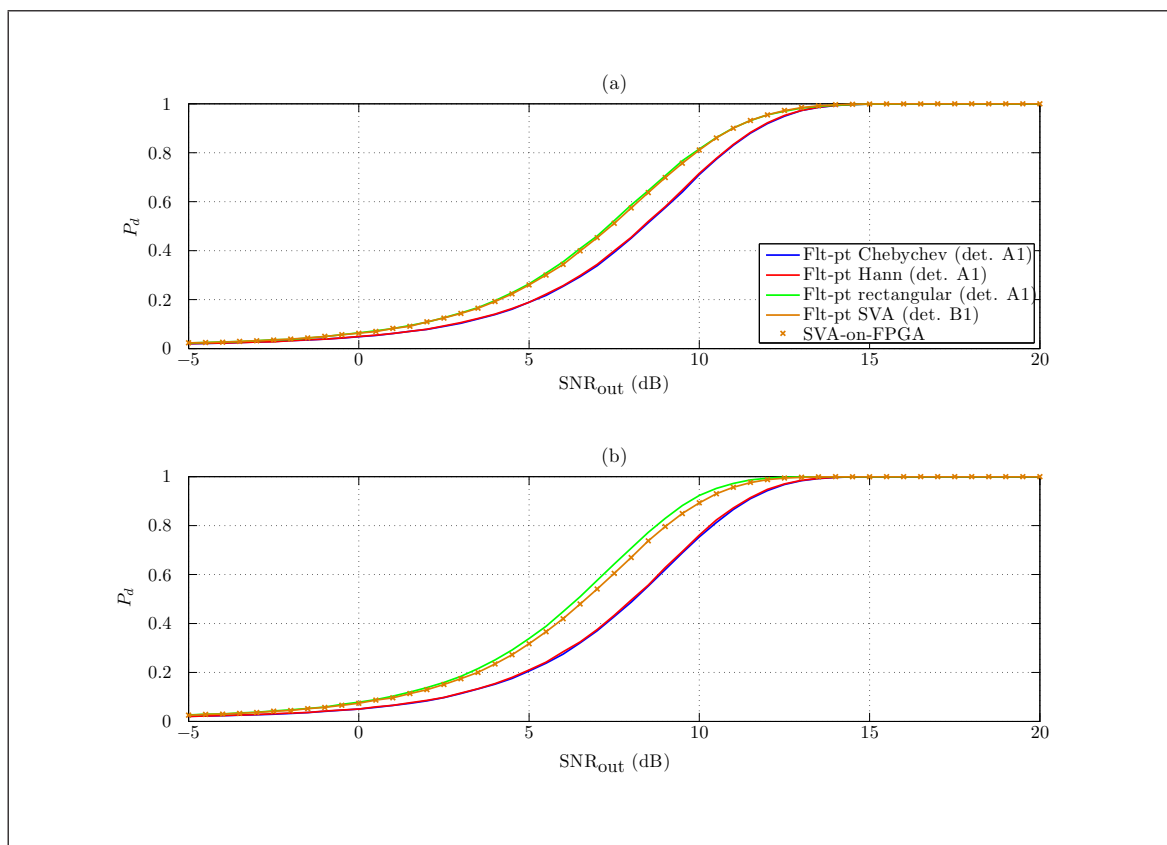


Figure F.3.1 SVA-on-FPGA full DCW scenario for: (a) the non-zeropadded case; (b) the zeropadded case

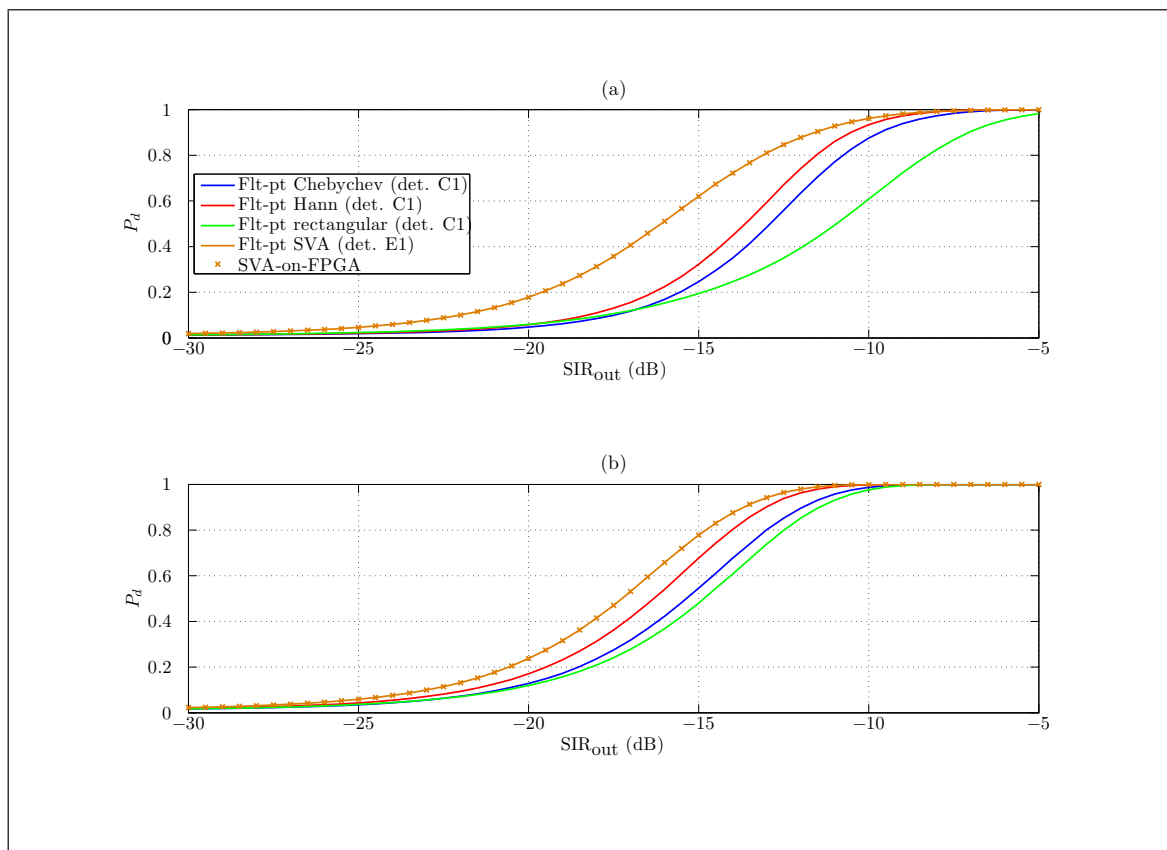


Figure F.3.2 SVA-on-FPGA interference full DCW scenario for: (a) the non-zeropadded case; (b) the zeropadded case

F.4 Reports

The completion of the MBD methodology down to FPGA implementation for the SVA algorithm permitted estimates of the characteristics of the SVA algorithm on a Xilinx® FPGA. This allowed assessments to be made over the viability of the algorithm for incorporation into real-life DRxs. The creation of the phase 2 model provided approximate numbers for these characteristics. The numbers were approximate because the design would be optimised before actual inclusion in a real-life DRx. The information of interest was contained in the reports which are generated after the implementation stage. The two reports of main interest were the “Place and Route” report and the “Post-PAR Static Timing” report. The next two sections describe the information of interest which was extracted from these reports.

F.4.1 Place and Route Report

The place and route report provided answers about the footprint of the SVA algorithm. Previously in section F.2 some important information was extracted from the place and route report and shown in Table F.2.2. This was extracted after the use of the phase 2 design.

The numbers in Table F.2.2 refer to the resource requirements of the **SVA algorithm only** because this was the only process running on the device. The two important factors from Table F.2.2 are: (1) the number of DSP48E embedded multipliers being used (24) and (2) the number of external IOBs being used (147). The number of embedded multipliers being used was relatively high (50% for the ML505 evaluation board). In fact seventeen of these multipliers were used just to perform the division. This could be significant but is dependent on the actual number of embedded multipliers on the target hardware and how many embedded multipliers would be required for the other processes and algorithms required to run alongside the SVA algorithm in a real-life DRx as part of an ESM system.

The significant 30% of external IOBs utilised was expected because 18 bit words (both real and imaginary parts of the data) were sent to and received from the FPGA board. This number is somewhat less of a concern because in a real-life DRx the SVA algorithm would be unlikely to be used in isolation on an FPGA board and therefore there would not be a requirement to get 18 bit words (with both real and imaginary parts) on and off the FPGA via the IO pins.

F.4.2 Post-PAR Static Timing Report

At the bottom of the post place and route static timing report there was a key table of information which contained the maximum clock frequency at which the phase 2 model design could run. This information is displayed in Table F.4.1. This maximum

Timing summary:

Timing errors:	0 Score: 0 (Setup/Max: 0, Hold: 0)
Constraints cover	11802 paths, 0 nets, and 7898 connections
Design statistics:	
Minimum period:	4.924 ns ³ (Maximum frequency: 203.087 MHz)
Maximum path delay from/to any node:	4.924 ns

Table F.4.1 SVA-on-FPGA post place and route static timing report extract

possible clock rate of ~ 200 MHz appeared to be ideal, especially as it was obtained with the use of a design for which there had been no attempt at critical path analysis. Therefore there remains scope for improvement on this rate.

One undesirable feature of the phase 2 model design was that to ensure the divider block could keep pace with the rest of the synchronous pipelined architecture, it was clocked at a rate five times faster than the rate at which new time-domain samples appeared. Therefore although the SVA algorithm clock rate was ~ 200 MHz, actual SVA-processed spectral estimates would be produced at a rate of $203.087/5N_{\text{FFT}}$. This highlighted the nature of the use of a multi-clock system and suggested that in order to produce a single clock tree it might be necessary to move away from the pipelined architecture and towards an explicit state-machine controlled system. Such a system would require handshakes between data-processing elements if the resource utilisation were to be kept roughly constant (i.e. maintain the use of only one divider block).

F.5 Suggested Design Improvements

There are many points upon which the phase 2 model could be improved in the future if optimisation of the model is to take place. The following is a list of known issues which it would be desirable to address or at least consider further:

Use of FIFOs The use of FIFOs in a pipelined architecture is a risky strategy because if corruption occurs in the queue the whole system would produce incorrect results from that moment onwards (the so-called garbage-in-garbage-out (GIGO) phenomenon). An improved memory layout would be to use two banks of memory for each FIFO buffer. Data would be written to one bank and read from the other bank in one period and then once the write-to bank was full (and simultaneously the read-from bank was fully read) the roles of the banks would switch. This process would ensure that if a bank of memory became corrupted, then only one spectral estimate at the end of the algorithm would be erroneous. The penalty to pay for a more complicated memory structure would be greater

³The minimum period statistic assumes all single cycle delays. Analysis completed Mon Jan 18 13:16:06 2010

latency as the initial bank of memory is written to before any read operations occur and the greater memory resource requirements.

Dropping of LSBs after word growth In the phase 2 design LSBs were dropped after the summation of the neighbour bins operation so that single 18 bit wordlength multipliers could easily be used. However the use of full precision would lead to answers with better accuracy. The downside to such an approach would be greater resource usage in the FPGA, more complicated data handling and greater latency in the system as carry-out bits would need to be dealt with correctly. Additionally, from section F.2.5, the current dynamic range with which the phase 2 model could cope was found to be ~ 100 dB which in most circumstances would be sufficient and negate the need for full precision.

Use of the faster clock rate for divider block The divider block could not be pipelined to take in fresh data each clock cycle. This is because the division algorithm is inherently a loop-based operation. The Xilinx[®] divider block is, however, optimised to accept data like a carousel and can achieve a constant rate of accepting new data if it is clocked fast enough compared to the rate of new data. This was the approach used in the SVA-on-FPGA model where the divider block was clocked at a rate five times faster than the rest of the data flow. However, a multi-rate system with a complicated clock tree similar to this description is less desirable than a simpler, single clock tree. To achieve a single clock design either:

- The number of divider blocks would need to be increased fivefold to deal with the pipelined workload. This would not be possible in the Xilinx ML505 FPGA evaluation board because such a design would require at least $7 + 17 \times 5 = 92$ embedded multipliers (the evaluation board contains 48 embedded multipliers). Even in a larger FPGA board this solution may not be preferable because it leaves few embedded multiplier resources available for other processes and algorithms which would be required in a real-life DRx.
- The pipelined architecture would be removed and replaced by combinatorial elements whilst the data flow between them would be controlled by a finite state machine (FSM). Such a system would be more robust to corruption than the synchronous pipeline architecture. The cost of such a system would be a slower data throughput as data would effectively be reduced to the rate of the divider block (actually the data rate would be less than the divider block rate because of the extra handshake required between data-processing elements). A pseudo-algorithm of how this type of architecture might be implemented is shown in Fig F.5.1.

Testing over ethernet The tests of SVA-on-FPGA were achieved with the use of a

JTAG connection on the ML505 board. The switch to testing over ethernet would allow larger blocks of data to be processed at once at a much faster rate. The utilisation of larger data blocks would also allow the system to run for longer periods of time. Therefore it would not be reset as often and any errors that might be wiped during reinitialisation would be easier to discover. Obviously testing over ethernet is a slightly more involved procedure than testing over JTAG, however the Edinburgh FPGA group of SELEX Galileo already use this method of test and have re-usable IP which could be incorporated into future SVA-on-FPGA testbenches.

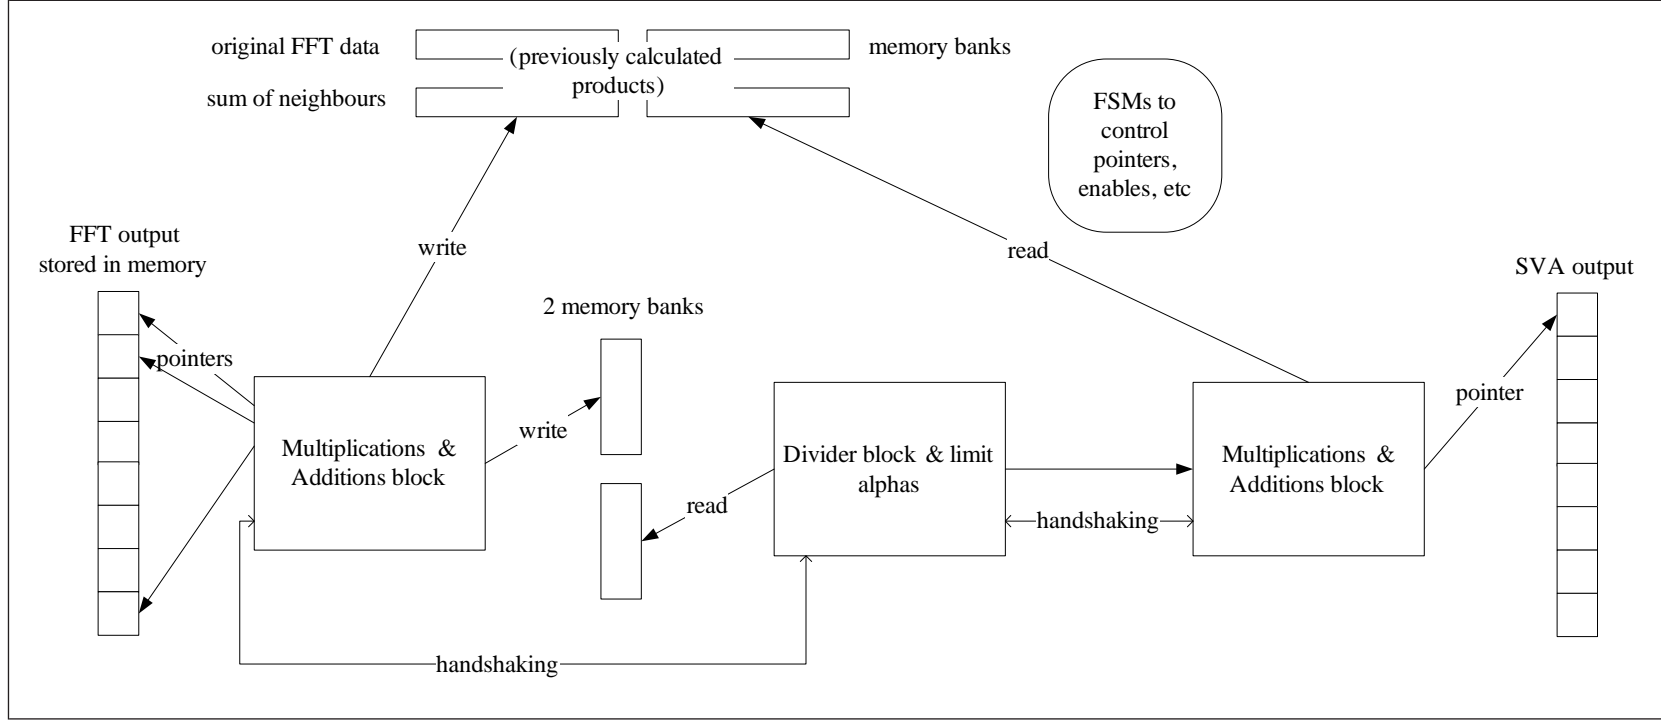


Figure F.5.1 Pseudo-algorithm of the suggested FSM-controlled SVA-on-FPGA

APPENDIX G

Navajo Trials

G.1 General Information

The results from the floating-point MTM and periodogram smoothing simulations indicated that the detection of broader-band LPI radar signals could be improved beyond that attainable with the use of conventional algorithms. The periodogram smoothing techniques and MTMs were tested and shown to result generally in an increase in the number of detections in comparison to the conventional techniques.

In a similar case to the SVA algorithm it was desirable to examine whether the increased number of detections for at least some of the algorithms of interest persisted when tested on real-life signal data.

The “Navajo trials” on 27th May 2010 presented an opportunity to operate the algorithms of interest on real-life sample data captured with a superhet and DRx combination and to demonstrate some of the possibilities open to SELEX Galileo through the use of ELB antennas.

The general idea behind this experiment was to capture sample data which contained broader-band television (TV) signals transmitted from known locations and determine which of the nonparametric spectral estimation techniques resulted in a greater number of detections.

G.2 Constraints

Some of the constraints from the helicopter trials in appendix A no longer applied or were removed for the Navajo trials. Specifically, the hardware was able to alternate between EW dwells in-flight and the GPS unit, to time-stamp capture events, was operational.

G.3 Aims

As demonstrated in chapter 5, the MTMs and periodogram smoothing techniques can be more useful in the search for signals which are broader-band than traditional narrowband pulsed sinusoids. This is because for larger FFT lengths conventional detection methods may over-resolve the spectral estimate. Therefore in a real-life DRx

implementation, the MTMs and periodogram smoothing techniques are predicted to be useful in the detection of LPI radar signals which occupy a relatively large bandwidth. However in the absence of a known emitter of LPI radar signals other known broader-band signals were used to test some of the algorithms of interest. The signals used for this purpose were PAL-I and DVB-T encoded TV signals which in the UK are broadcast in the ultra high frequency (UHF) band [470, 862] MHz. These TV signals are significant because they occupy a relatively large bandwidth of 8 MHz and are continuous and therefore arguably resemble a CW LPI radar signal. The aim was to show an increase in the number of detections with the use of some of the algorithms of interest over the basic conventional ESM DRx procedure of: $\text{DCW} \rightarrow \text{FFT} \rightarrow |\text{FFT}|^2 \rightarrow \text{detect on ordinates}$.

G.4 Experiment

Important parameters about the main transmitter and repeater stations in the vicinity of the flight plan were known from Ofcom[146].

The UHF band allocated for TV signals in the UK ([470, 862] MHz) is subdivided into 8 MHz “channels”. Under the analogue PAL-I system (panel (a) of Fig G.4.1) each of these channels carries only one TV station, however under the digital DVB-T system (panel (b) of Fig G.4.1) each of these channels carries multiple TV stations and “televised-radio” stations multiplexed together. The digital signals are broadcast at lower effective radiated power (ERP) levels partly because the digital encoding carries error correction to compensate for weaker reception.

The EW dwells 67 and 68¹ were the most appropriate frequency ranges which the superhet and DRx system monitored. As in the Helicopter trials, the EW dwells were achieved through a combination of frequency mixing and analogue filtering to obtain a constant bandwidth centred on a particular IF. The IF was digitised by means of the ADC in the DRx.

The digital samples were sent to a laptop connected to the DRx for analysis. The laptop performed the frequency channelisation by means of an STFT from which a spectrogram was produced. Before the actual flight, predictions were made for the maximum range at which the TV signals could be detected by the laptop connected to the DRx. Therefore data was captured where digital TV signals were on the boundary of detection (minimum operational sensitivity (MOS)). The key behind these predictions was the one-way receiver equation

$$\text{Pow}_r = \frac{\text{Pow}_t G_T G_{\text{ESM}} \lambda^2}{(4\pi R_{\text{ESM}})^2} \quad (\text{G.4.1})$$

where:

¹These numbers are indices for frequency ranges which are not shown here

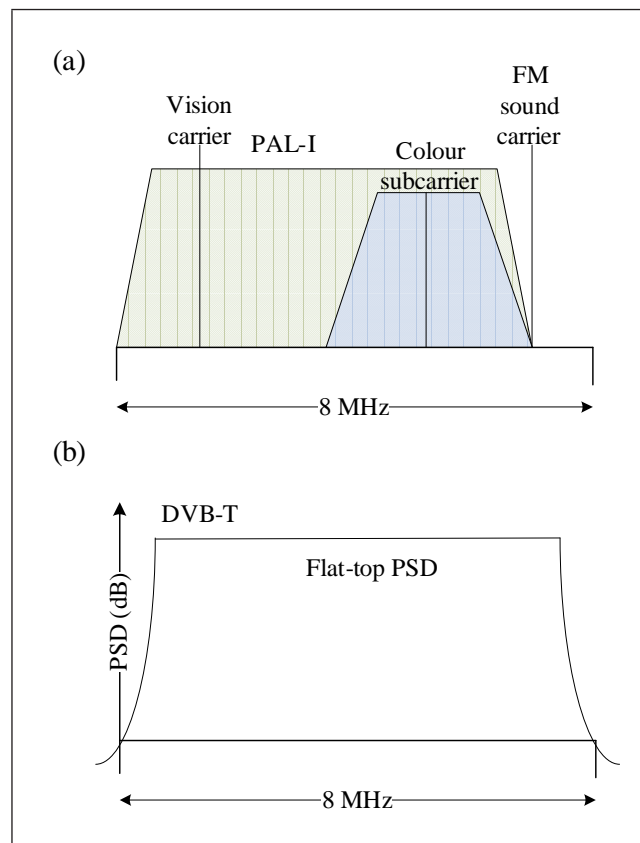


Figure G.4.1 Stylised diagrams of UK broadcast TV signal encoding schemes: (a) PAL-I; (b) DVB-T

- Pow_r is the power received by a receiver;
- Pow_t is the transmitted power;
- G_T is the gain of the transmit antenna;
- G_{ESM} is the gain of the ESM antenna and receiver system;
- λ is the received signal wavelength;
- R_{ESM} is the distance between the transmitter and ESM system.

Equation (G.4.1) is essentially a one-way version of the radar range equation. It was manipulated to derive an approximation for an expression of the SNR in the laptop connected to the DRx. In decibels this was²

$$\text{SNR} = \text{ERP} + 30 + G_T + 20 \log_{10} \lambda - 20 \log_{10}(4\pi R_{\text{ESM}}) - (N_p - G_{\text{ESM}}) \quad (\text{G.4.2})$$

where antenna gain G_T was in the units dBi. An explanation of the terms in Eq (G.4.2) is:

SNR

This was the ratio of signal power to noise power in decibels. This term replaced Pow_r in Eq (G.4.1).

²The SNR in Eq (G.4.2) is not SNR_{out} which would compensate for the FFT processing gain

ERP + 30

Converted the ERP of the transmitter from units of dBW to dBm. This term replaced Pow_t in Eq (G.4.1).

 G_T

Large TV transmitters radiate almost uniformly in all directions towards the horizon. However for elevation angles greater than this the gain is reduced. As the Navajo aircraft flew above the height of the transmitters an approximate reduction of 5 dB was used for this term.

 $-N_p$

The total thermal noise floor power which was passed to the ADC. It is included in Eq (G.4.2) to balance the SNR on the left-hand side.

 G_{ESM}

The gain of the ESM antenna and receiver system. The term receiver gain covers both the gain of the RF chain and the gain of the DRx. The antenna gain component of this was fairly uniform across a broad mainlobe.

The noise floor power in the laptop connected to the DRx was estimated from

$$N_p = k_B T_0 B_N N_f L_{sc} \quad (\text{G.4.3})$$

where:

- k_B is Boltzmann's constant;
- T_0 was the absolute standard temperature 290 Kelvin;
- B_N was the bandwidth of the FFT frequency bin filter;
- N_f was the cascaded noise factor of the total receiver chain;
- L_{sc} was a worst-case loss factor to account for the scalloping loss of the window function used.

This was then converted to logarithms to become a noise floor power (N_p) in units of dBm. The bandwidth of the frequency bin filter was approximated from: $B_N = f_s/N$. Similarly the sensitivity of the laptop connected to the DRx was estimated from

$$\text{Sensitivity} = k_B T_0 B_N N_f L_{sc} s_{pf} \text{SNR} \quad (\text{G.4.4})$$

where SNR was set at a ratio value that would typically be required for an operator to spot a detection in a spectrogram and s_{pf} was a spreading factor associated with broader-band signals whose power is approximately evenly spread across their frequency.

In the course of the experiment the operator was the user of the laptop which was

used to process the captured sample data from the DRx. The SNR level at which the operator was able to distinguish a signal from the noise floor in a spectrogram is typically [8, 12] dB. A range is specified here because the ability to distinguish weak signals from the noise floor is somewhat subjective.

The effect of the spreading factor was to increase the power required in a spectrogram before the broader-band signals became noticeable. Its effect is illustrated in Fig G.4.2. If the areas under the signal peaks in (a) and (b) are considered to be the same, the peak power density is reduced in (b) by the spreading effect of broader-band signals. The size of the effect in the experiment was $s_{pf} = 8N/f_s$.

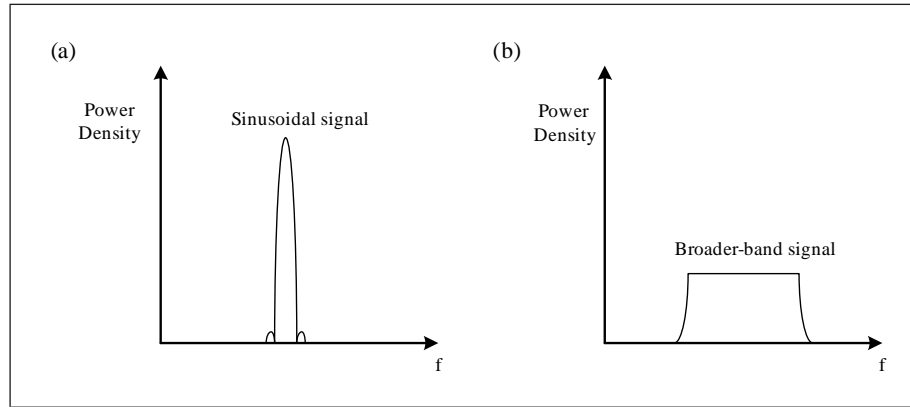


Figure G.4.2 Stylised effect of spreading factor

The overall receiver noise factor N_f and receiver gain G_{ESMr} (not G_{ESM} as this would include the antenna gain) were calculated by a cascade of partitions in the RF/IF chain and DRx. The partitions chosen were those shown in Fig G.4.3³. Each element was characterised by a gain and a noise figure. The overall receiver gain in Fig G.4.3 was calculated from the summation of the individual elemental gains in decibels

$$G_{\text{ESMr}} = \sum_{n=1}^6 G_{\text{ESMr},n} \quad (\text{G.4.5})$$

This formed part of the G_{ESM} term from Eq (G.4.2). The overall receiver noise figure was calculated from Friis formula[147] for cascaded noise factors

$$N_f = N_{f,1} + \frac{N_{f,2} - 1}{G_{\text{ESMr},1}} + \frac{N_{f,3} - 1}{G_{\text{ESMr},1} G_{\text{ESMr},2}} + \dots + \frac{N_{f,6} - 1}{\prod_{n=1}^6 G_{\text{ESMr},n}} \quad (\text{G.4.6})$$

where $N_{f,n}$ has a lowercase f to signify noise factor and the gains are power ratios not in decibels. To convert to a noise figure, logarithms were taken $N_F = 10 \log_{10} N_f$. The noise figures for the passive elements were simply the negative of the gains in decibels.

³The actual characteristics of gain and noise figure of the receiver system are not shown here

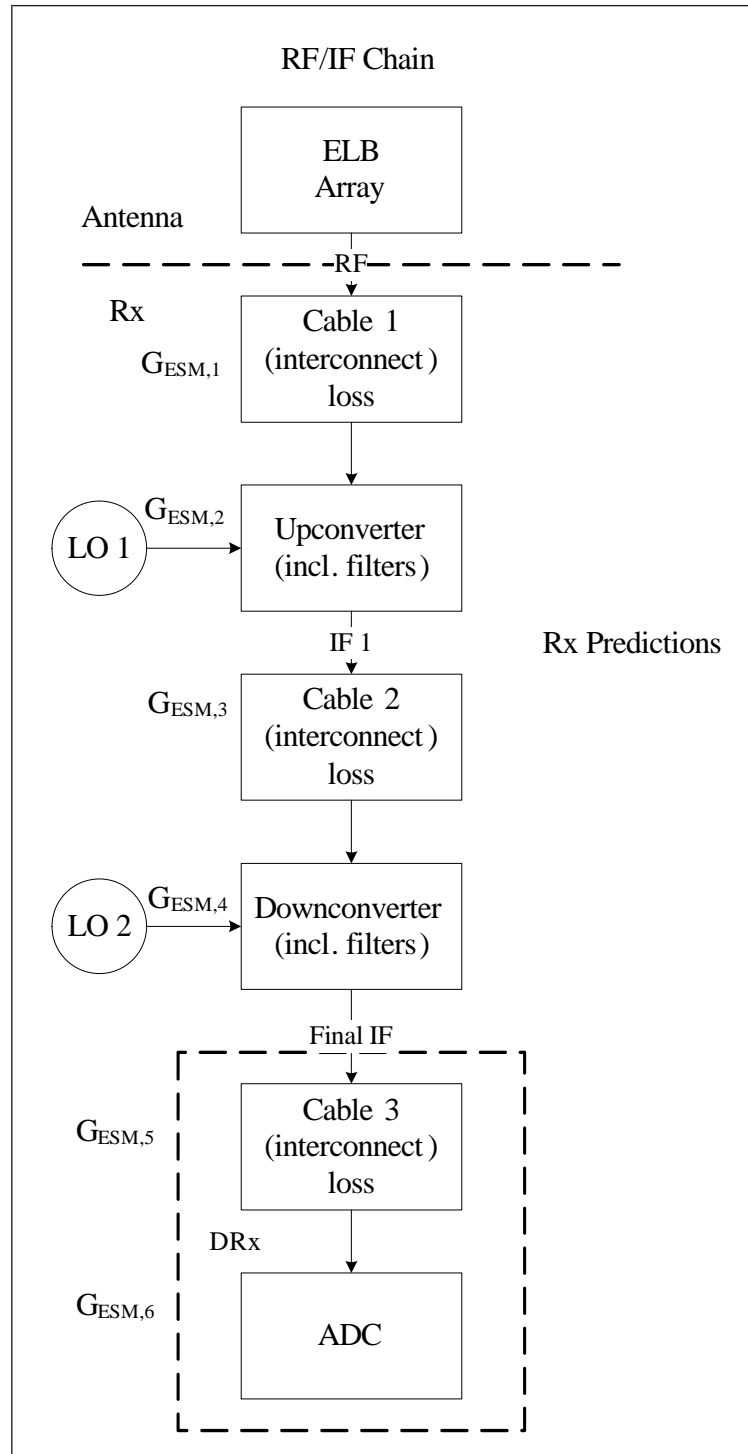


Figure G.4.3 RF/IF chain to the DRx

The partitions in Fig G.4.3 are arbitrary constructs which were chosen at logical physical boundaries. The elements could be further decomposed into sub-elements with corresponding gains and noise figures if desired for traceability purposes.

In the event, the predictions for the range at which the laptop operator would be able to distinguish the digital TV signals from the noise background in a spectrogram created from the sample capture files proved accurate provided the ELB antennas were directed towards the transmitter of interest (i.e. when the estimate used for G_{ESM} in Eq (G.4.2) was valid)⁴.

The analysis accumulated the number of detections when different MTMs, periodogram smoothing techniques and conventional techniques were used to examine if the algorithms of interest provided a detection improvement from field data.

G.5 Results

The sample capture files were rapidly checked in-flight with the use of MATLAB[®] on a laptop. Spectrograms such as that shown in Fig G.5.1 were formed to verify the file contained digital TV signals. Fig G.5.1 was created from the first 2^{19} data points from the “X” channel (one of the antennas) of file `s_270510_0036.txt`. As per the examples from appendix A, the z-axis of the spectrogram was in units of dB/Hz but a conversion to units of dBm/Hz can be made through Eq (A.5.1). The analogue TV signals are clearly visible as CW signals with two closely-spaced carriers (as in Fig G.4.1). For informative purposes, the analogue channels from the Sandy Heath transmitter have been labeled in the figure. Digital TV signals appear as power smeared out over 8 MHz bandwidth. Also for illustrative purposes, one digital signal multiplex from the Oxford transmitter has been marked on Fig G.5.1.

⁴the actual ranges are not shown here

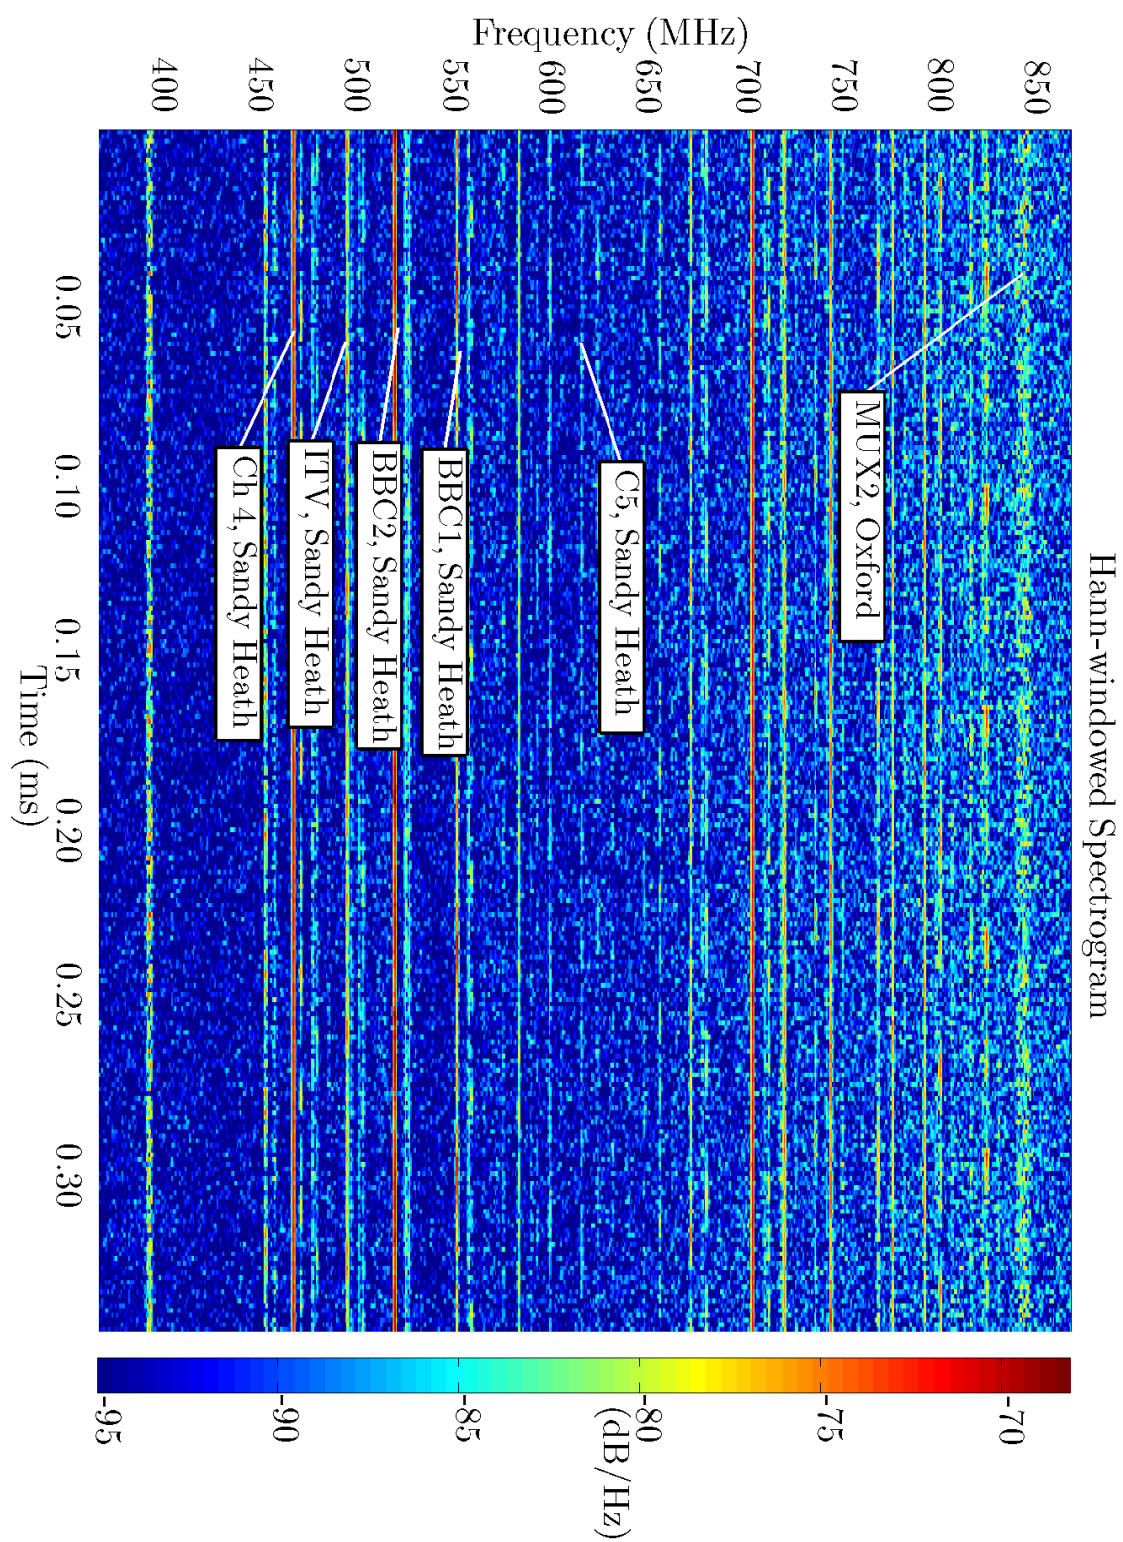


Figure G.5.1 Example Hann-windowed spectrogram from Navajo trials

In the offline process, sample capture files which were known to contain digital TV signals were isolated and spectrograms were created from these files. Areas in these spectrograms which contained noise-only data were used to estimate the noise floor power density. This noise floor power density was crucial for the establishment of a threshold of $P_{fa} = 1 \times 10^{-6}$. From the sample capture file a weak digital TV signal was isolated from the rest of the data via a digital filter applied to the file. The filter used was a bandpass filter whose bandwidth was larger than the largest hypothesis bandwidth to be tested. This concept is illustrated in Fig G.5.2.

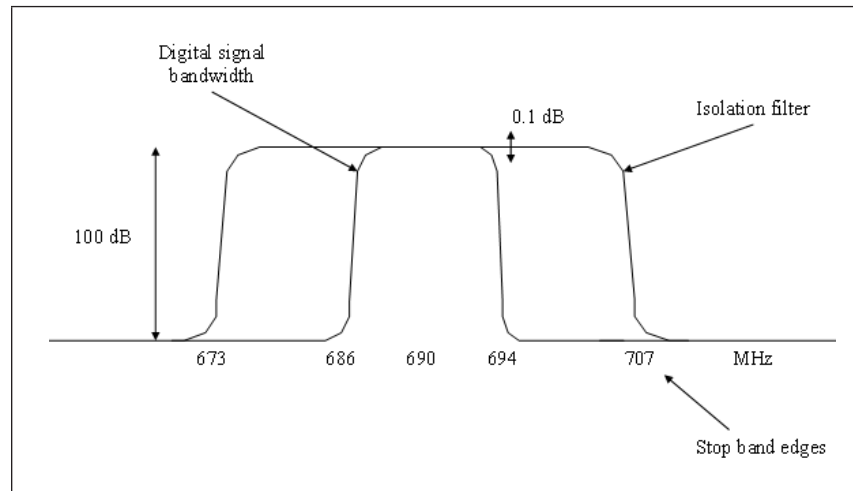


Figure G.5.2 Isolation of a digital TV signal on which to perform detection experiment

The MTMs and periodogram smoothing techniques which could be accurately modelled with the use of the scaled chi-squared distribution for an AWGN input (Eq (4.2.2)) were the subject of this experiment. Those chosen for the experiment were the:

- unity-weighted DPSS MTM (detector G1);
- unity-weighted sine taper MTM (detector H1);
- SMP (detector J1 variant);
- TSMP (detector L1 variant).

The number of detections which occurred when these techniques were used to process an isolated digital TV signal from a sample capture file of interest was counted. The plot formed was very similar to that of Fig 4.3.3 of hypothesis bandwidth against frequency, although in this case the plot was formed from real (non-complex) sample data. The basis for the techniques was an FFT of length $N_{\text{FFT}} = 2048$. The number of detections obtained with the use of the above techniques at various hypothesis bandwidths could be compared against the number obtained with the conventional modified periodogram technique. This was because the limiting case of a hypothesis bandwidth test for the SMP or TSMP (detectors J1 or L1 variants) is one frequency bin wide which is equivalent to the results from a modified periodogram.

Example results for the number of detections are shown in Figs G.5.3 to G.5.6. The results were produced from 69 DCWs.

In all figures labels have been set to show the number of detections if only the modified periodogram (detector A1 variant) was considered. The extreme hypothesis of one frequency bin width for the DPSS MTM (detector G1) and sine taper MTM (detector H1) is also a modified periodogram because it is formed from the application of the first DPSS taper and first sine taper to a DCW before the FFT and magnitude-squared process.

The frequency span in Figs G.5.4–G.5.6 shortens with increasing bandwidth hypothesis. This was a consequence of their computation method and the source of the data. The sine taper MTM (detector H1) used was the more computationally-efficient frequency-domain formulation and the smoothing methods involved the summation and average of ordinates. The sample data was essentially obtained from a bandpass filter where the data at the edges was less reliable and therefore a linear rather than circular convolution was used for the sine taper MTM calculation which thus shortened the frequency axis at greater bandwidth hypotheses. The frequency axis also shortened in the smoothing method results because the ordinates could not be summed past the band edges.

If the modified periodogram (detector A1 variant) was used in the experiments, only the number of detections marked at the left-hand edge of the plots would be achieved. However through the summation of additional ordinates or through the average of more eigenspectra a greater number of detections was obtained. The number of detections peaked near the actual centre frequency and bandwidth of the digital TV signal under test. In comparison to the MTM techniques which produced Figs G.5.3 and G.5.4, SMP technique (detector J1 variant) generated slightly fewer detections but a peak is present at the approximate frequency and bandwidth. This trend, which was repeated for other sample capture files, suggested that the processing gain attainable from the MTM over the smoothing techniques is slight and eclipsed by the additional computational effort required.

The TSMP (detector L1 variant) could not test every bandwidth hypothesis which the other techniques could because it summed ordinates in powers of two. However the experiment showed that this does not greatly degrade the performance of the technique because broader-band signals which the DRx is likely to process are spread over only a few frequency bins even at high N_{FFT} values. The CW digital TV signals demonstrate this and it would be unlikely to find a spread-spectrum LPI radar signal with a broader bandwidth than this. The peak in Fig G.5.6 does not exactly match the approximate bandwidth and this is a consequence of the coarser bandwidth hypothesis test of the tree smoothing technique.

The small degradation in performance of the SMP and TSMP (detectors J1 and L1 variants) over the MTMs matched the pattern of results for the full DCW in section 5.2.

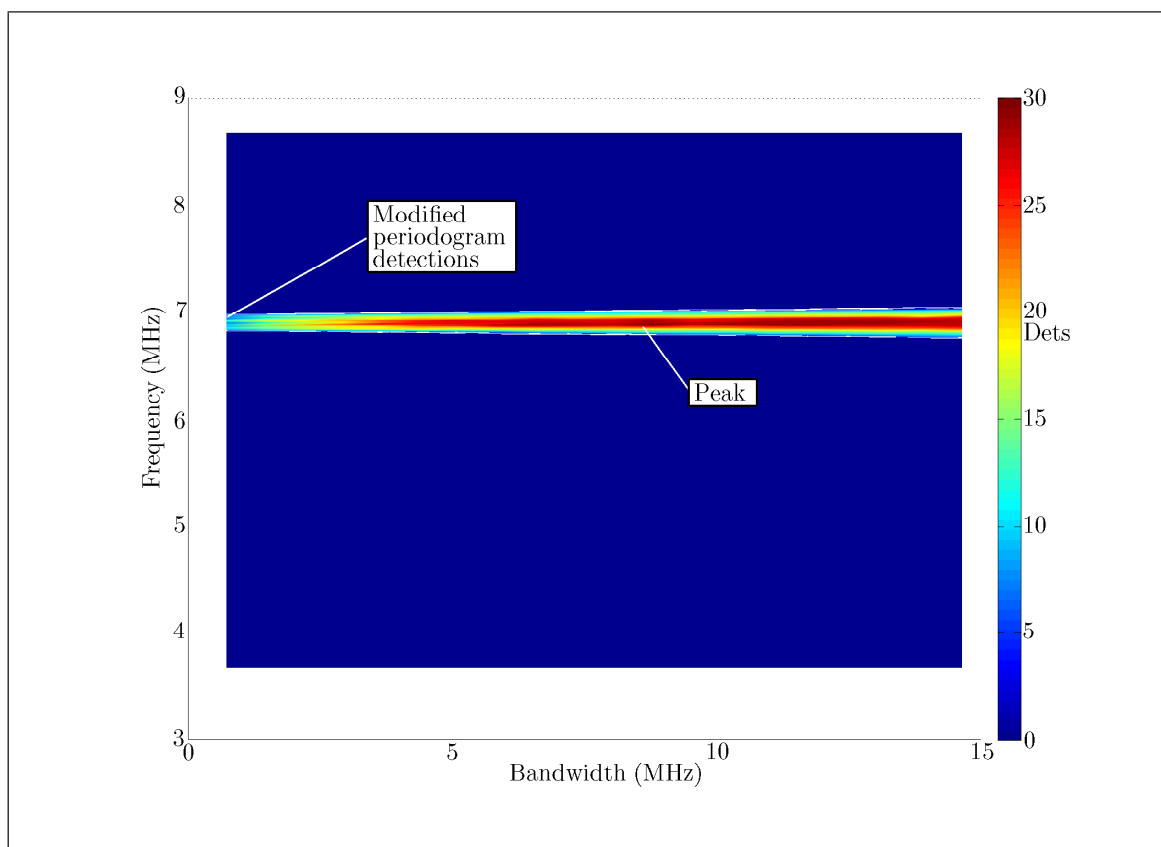


Figure G.5.3 Accumulated detections with the DPSS MTM (detector G1)

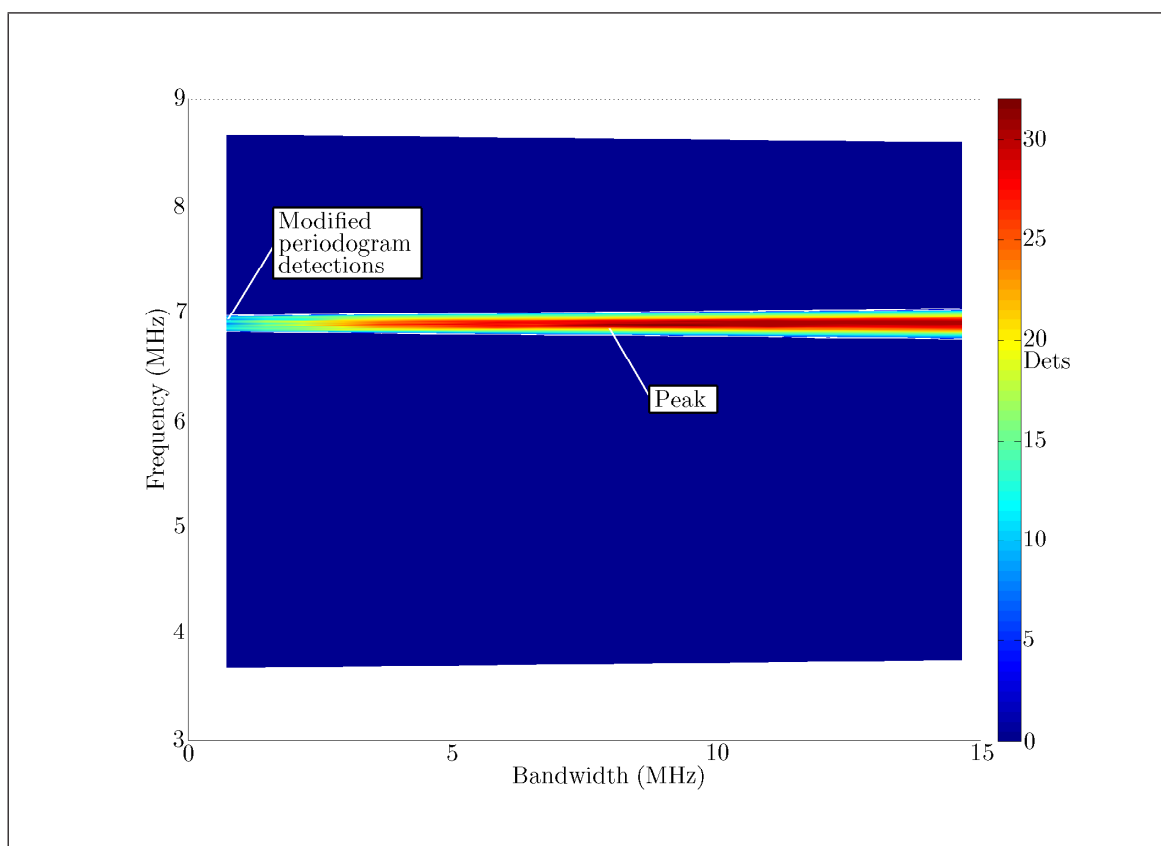


Figure G.5.4 Accumulated detections with the sine taper MTM (detector H1)

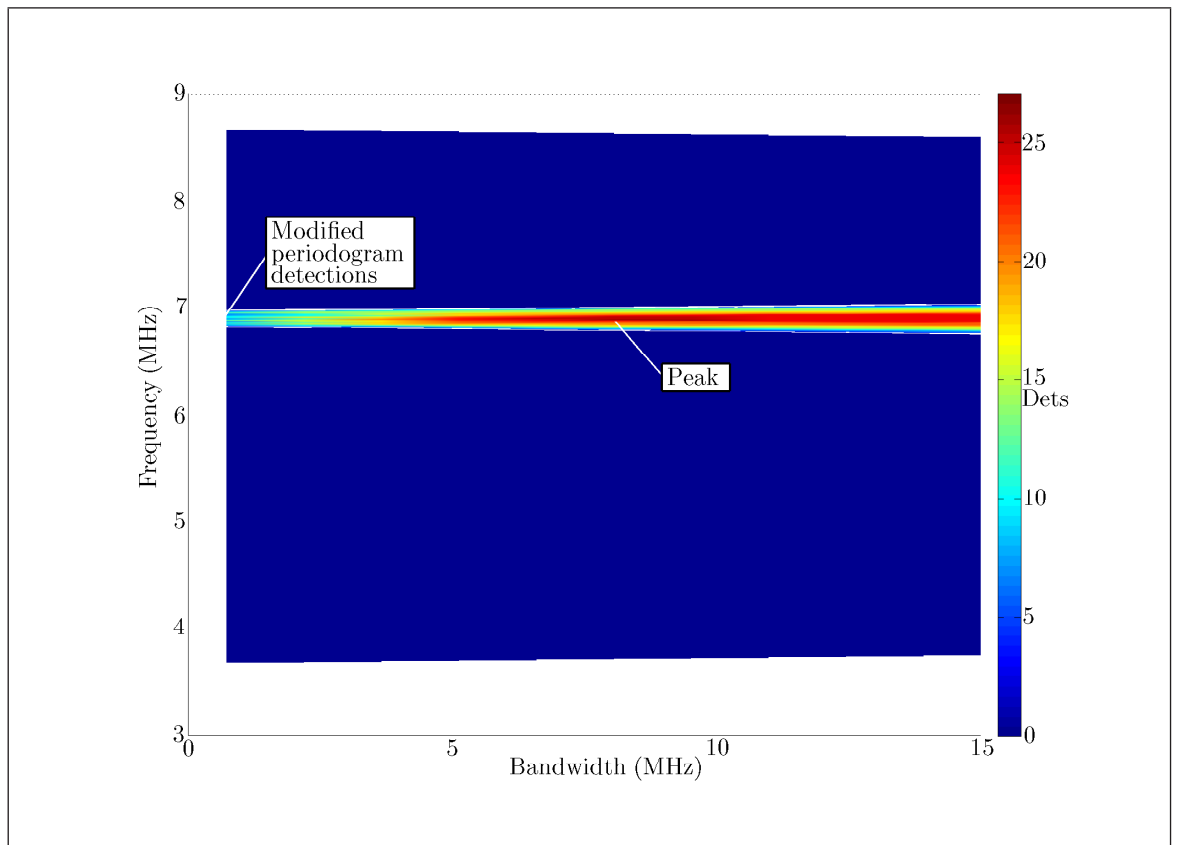


Figure G.5.5 Accumulated detections with the SMP (detector J1 variant)

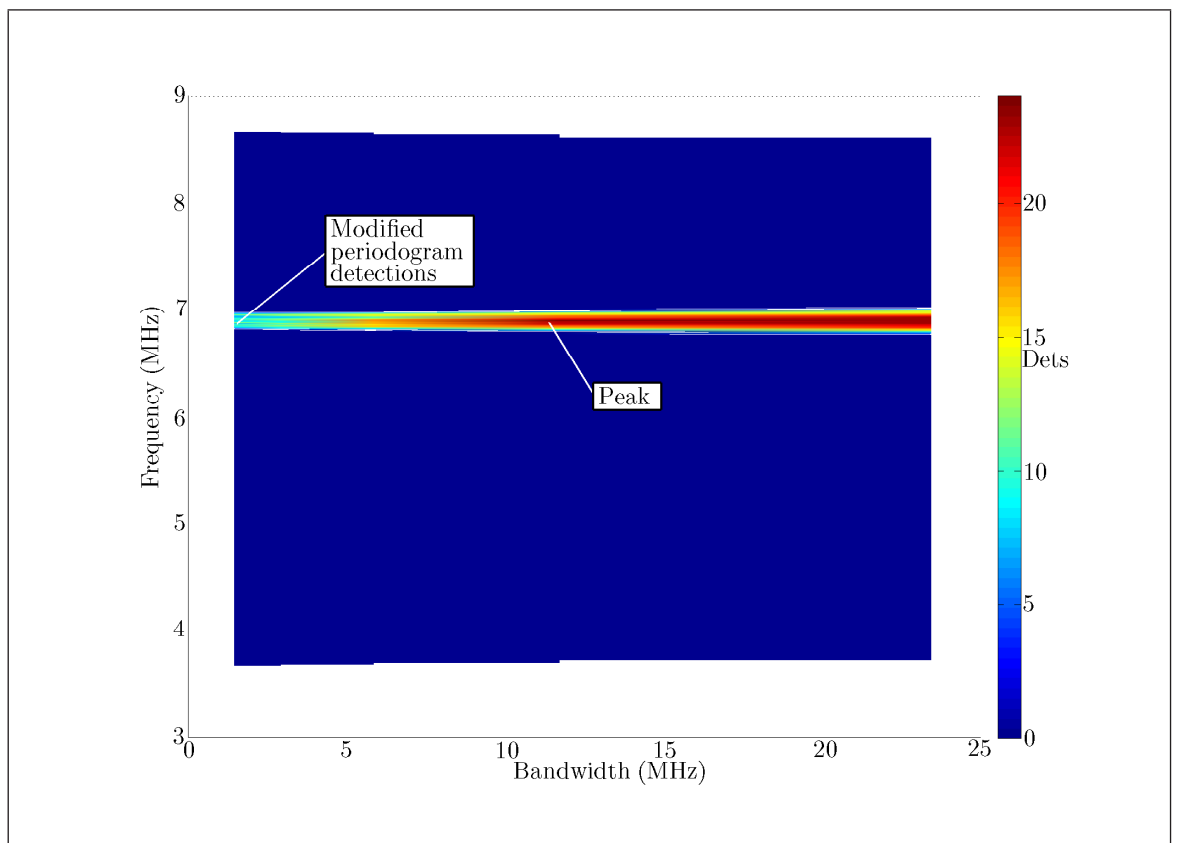


Figure G.5.6 Accumulated detections with the TSMP (detector L1 variant)

G.6 Summary of Results

This experiment was important as it was performed on real-life data from actual hardware. The experiment was used to compare different MTMs and smoothing techniques against each other and against the conventional modified periodogram technique in terms of the number of detections from a range at which the signals of interest were on the verge of detection. The results showed that the differences between the MTMs and smoothing techniques tested were small, but they were all considerably better than the modified periodogram technique.

The results reinforced the floating-point simulation findings from chapter 5. Therefore the notion of the tree-like smoothing technique as the most promising candidate was reemphasised because it was able to obtain an only slightly degraded number of detections for less computational complexity.

APPENDIX H

Fixed-Point Periodogram Smoothing

The floating-point software models of the MTMs and periodogram smoothing techniques established that their detector performance against broader-band LPI radar signals was superior to detectors which relied on the periodogram, the modified periodogram and even the SVA periodogram single ordinate detection. The detector performance of these conventional techniques could be improved somewhat by the utilisation of OR operations between neighbour ordinates but on the whole they resulted in poorer detectors than those which employed the algorithms of interest, i.e. the MTMs and the smoothing techniques.

The DPSS MTM (detector G1) or the unity-weighted sine taper MTM (detector H1) would be the ideal choice for hardware implementation and hence further investigation if hardware resource requirements were of no concern. However the performance difference between the TSSP (detector M1) and the aforementioned detectors G1 and H1 was small in all the scenarios tested and for typical LPI signal parameters. The TSSP is computationally far simpler and furthermore the first hypothesis bandwidth links in neatly to the detection of narrowband sinusoidal signals. Therefore it was appropriate as an engineering decision to further detector M1 along the MBD methodology. The next step of the methodology was to operate detector M1 in a realistic, fixed-point mode and investigate its performance as a fixed-point detector.

H.1 Implementation Decisions

The same factors which had affected the SVA implementation at this point in the MBD methodology also applied here. New code was required only for the tree-like smoothing component of detector M1, however, this did not require complicated operations which would otherwise benefit from Xilinx®-specific applications. The operations involved in tree-like smoothing are simply additions and barrel shifts.

H.2 Coding Methodology

A similar modular approach was used for the conversion of detector M1 as was taken in the fixed-point SVA section. However many of the modules and most importantly

the SVA module, had previously been converted to fixed-point. Therefore the objective was to convert the modules which represented the tree-like smoothing. These modules are shown in Fig H.2.1 and can be observed to involve only additions and barrel shift mathematical operations.

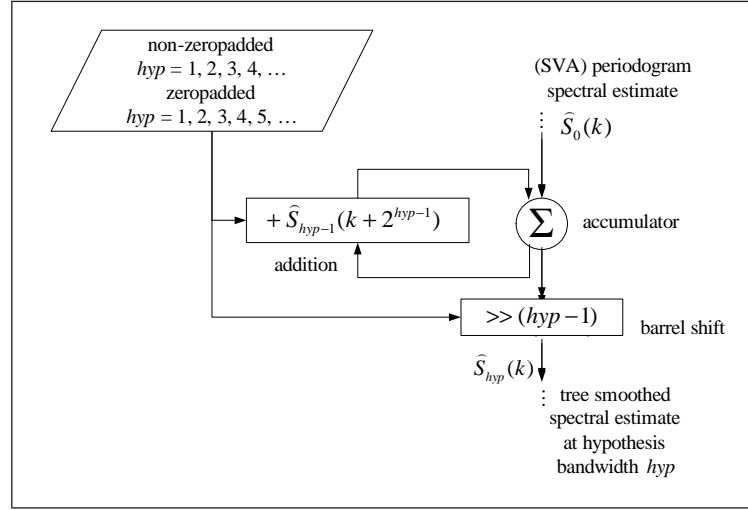


Figure H.2.1 Tree-like smoothing modules to be converted

Since the effort in conversion was concentrated on the tree-like smoothing component of detector M1, the eventual converted code could also be applied to the TSP and TSMP (detectors L1 and R1 variants) simply by an alteration to the front end of the detector from the SVA module to the appropriate window function if this was desired.

The same fixed-point arithmetic properties used in the SVA fixed-point appendix D were reused for reasons of consistency. Overall the code methodology did not involve two distinct phases as it did in appendix D. This was due in part to the ability to reuse the fixed-point code from appendix D. However, literal floating-point translations of any new fixed-point code were still generated for verification purposes.

An example of the verification used is shown in Fig H.2.2. The plots from left to right show the error at each ordinate when the fixed-point implementation results were compared with the floating-point translation. For clarity purposes an FFT length of $N_{\text{FFT}} = 128$ and a periodogram was used in this example and therefore there was a maximum total of $hyp = 7$ layers. The left-hand and right-hand plots display the error results from the two extremes of smoothing: the non-smoothed periodogram and the total power respectively. In this example a time-domain, sampled test vector was applied which contained complex noise with real and imaginary parts generated from a zero-mean, truncated normal distribution with PDF given by

$$f(x, \sigma, a, b) = \frac{\frac{1}{\sigma} \phi_{\text{PDF}}\left(\frac{x}{\sigma}\right)}{\Phi_{\text{CDF}}\left(\frac{b}{\sigma}\right) - \Phi_{\text{CDF}}\left(\frac{a}{\sigma}\right)}$$

where $\phi_{\text{PDF}}(\dots)$ and $\Phi_{\text{CDF}}(\dots)$ are the PDF and cumulative distribution function

(CDF) of the standardised normal distribution respectively. The real and imaginary parts of the complex noise were drawn from a normal distribution which had been truncated such that the amplitude of either the real or imaginary samples could not cause an overflow error in the system. This was because the overflow errors would have masked the true module and system characteristics. In essence the real and imaginary noise from the truncated normal was FS noise. Other types of test vectors were available such as a FS sinusoid or a FS complex exponential.

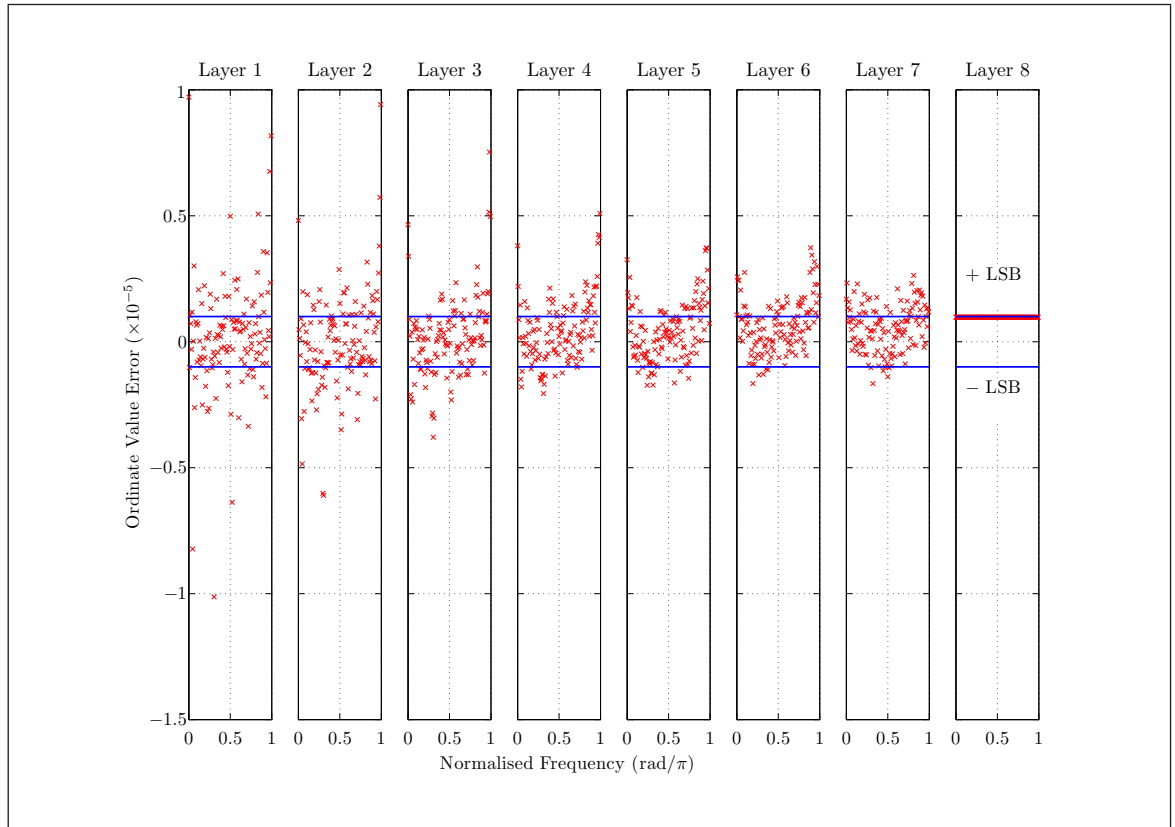


Figure H.2.2 Ordinate value error of the spectral estimates at different layers of the TSP due to fixed-point precision

A schematic of the result after conversion to fixed-point is shown in Fig H.2.3 as detector M2. The SVA module in the diagram is special because it could be switched in or out as desired to produce results for the TSP (detector L2 variant, which does not take into account any spectral leakage from strong interferers) or the TSSP (detector M2). However this was mainly for academic interest because the objective was to produce the TSSP (detector M2). As before in the fixed-point SVA appendix D the modules which remained in the floating-point domain represented continuous, real-world processes. For scenarios where interference was present the trivial change was made in detector M2 to replace the signal and noise generation with that of Fig 5.4.2.

The Simulink[®] diagram for the TSSP detector is shown in Fig H.2.4. It shows the system at its highest abstraction level. In the diagram: (1) the FSM which controlled the process; (2) the registers and memory; (3) the smoothing filter; (4) the periodogram (and SVA) calculation and (5) the quantised inputs.

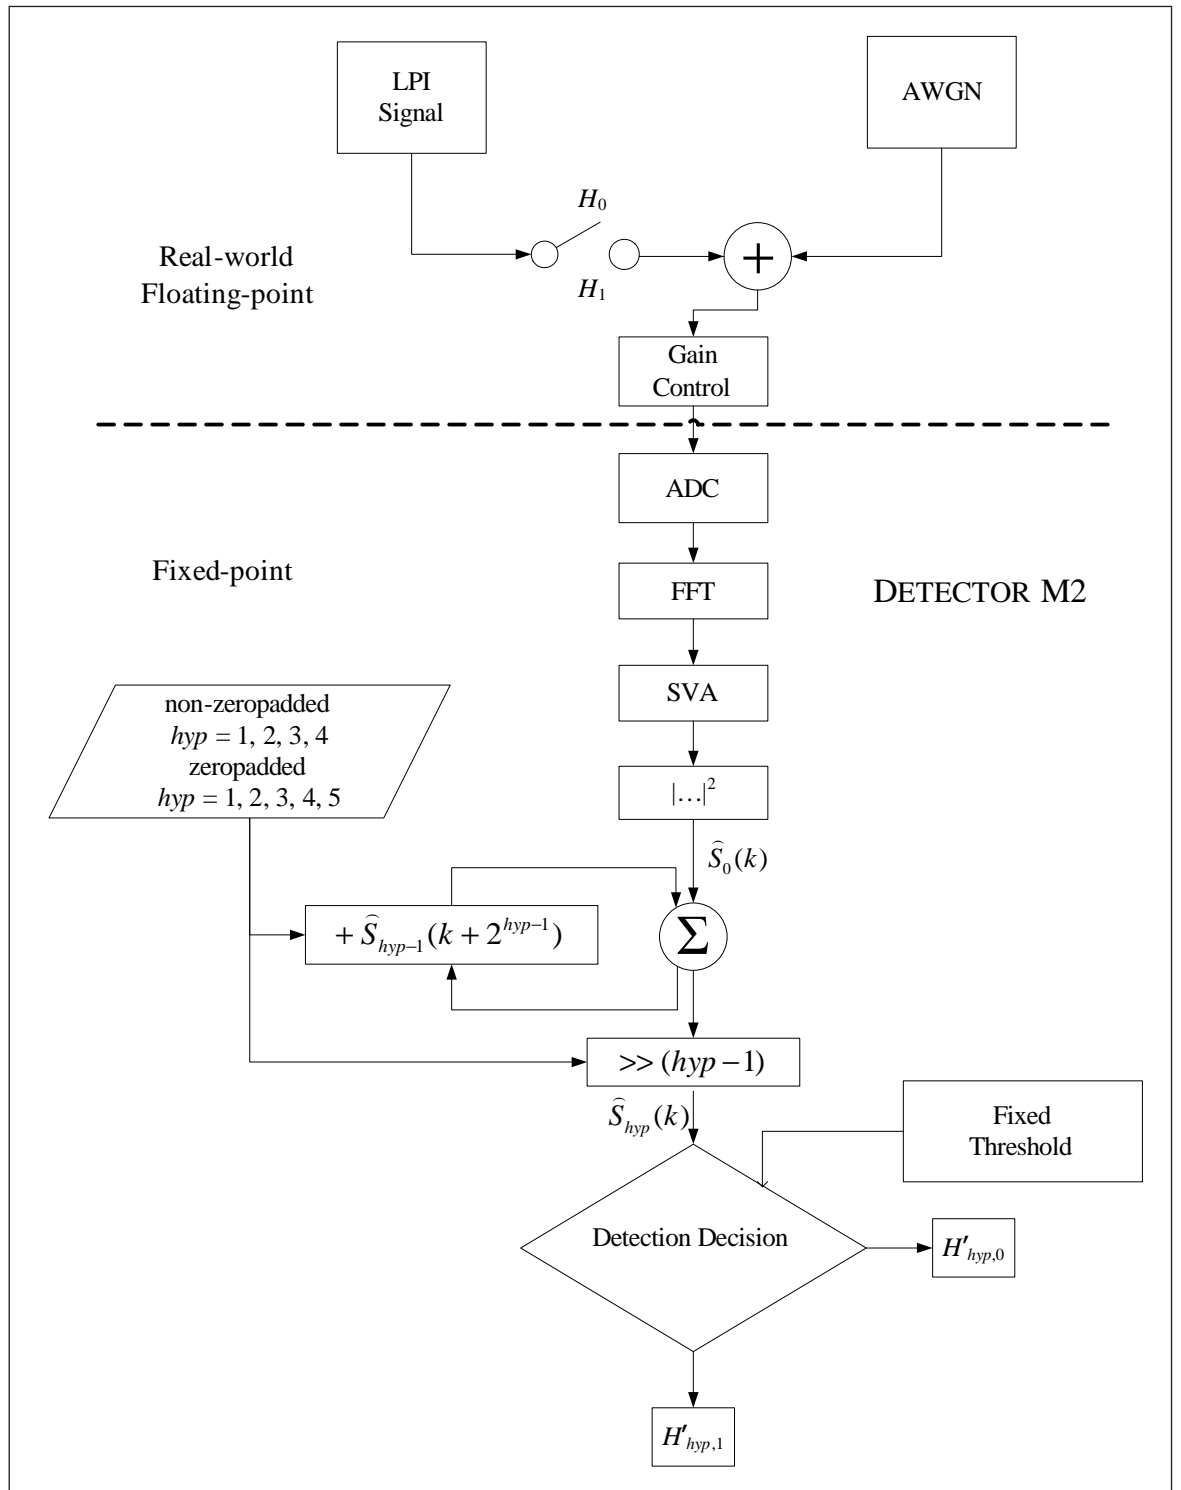


Figure H.2.3 Architecture of the TSSP detector M2 which shows the partition between floating-point and fixed-point

To demonstrate that the same smoothing filter was used multiple times, its instantiation is shown in Fig H.2.5. In the diagram the code to counteract the tendency of the spectral estimate to shift to the right has been highlighted.

The two diagrams suggest that a much greater proportion of the code for this system was written in native Simulink[®] than in the fixed-point SVA code and this indication is true. The reason for the change from a system coded mainly in MATLAB[®] m-code to a system coded mainly in Simulink[®] was purely to facilitate the learning curve of the two facets of the overall MATLAB/Simulink[®] tool.

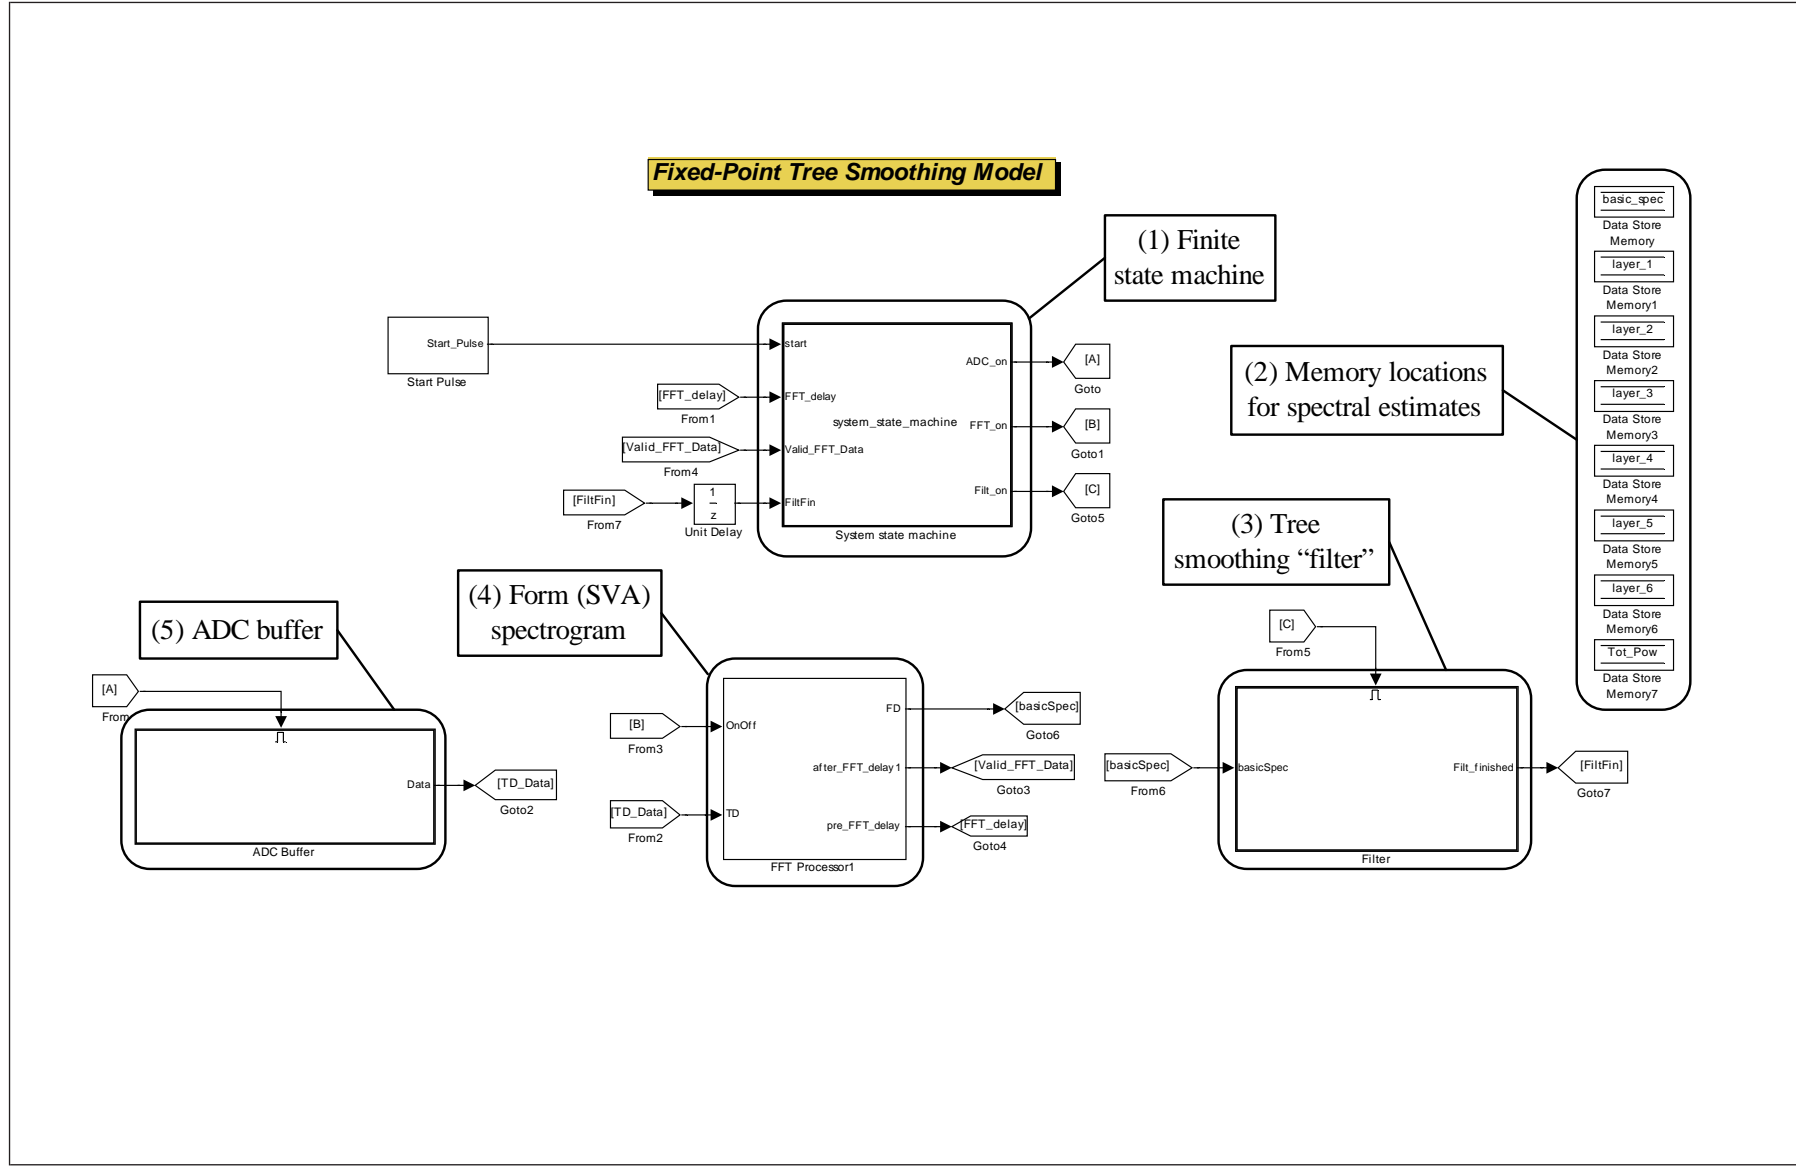


Figure H.2.4 Screenshot of the highest level in the Simulink® hierarchy of the fixed-point TSSP (detector M1)

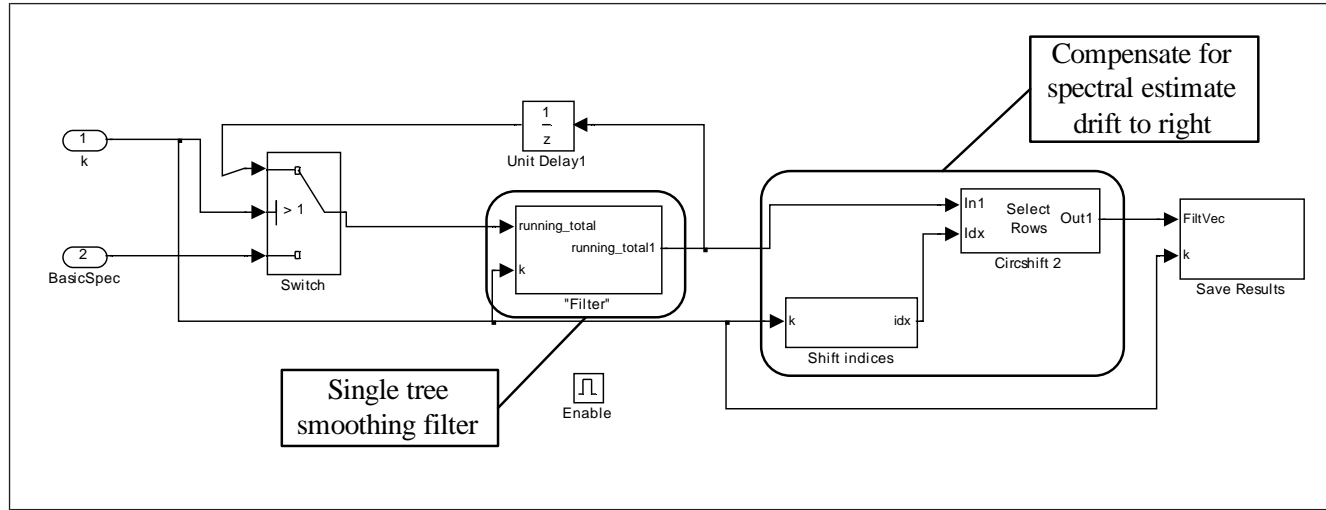


Figure H.2.5 Screenshot of the single instantiation of the tree smoothing filter

H.3 Coding Explanation

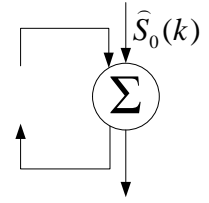
This section describes the considerations for the conversion of the modules which performed the tree-like smoothing from Fig H.2.1. One interpretation of the tree smoothing architecture is as a moving average filter which processes the initial spectral estimate to produce spectral estimates at different layers in accordance with the span of the filter which is constrained to a power of two. However this interpretation does not reveal the tree smoothing computational efficiency or potential for hardware reuse and if the technique were implemented in this way, it would result in a system with a large degree of redundancy.

The fixed-point code was designed such that there was only one filter which accepted two values, summed them and divided by two. This was done for each ordinate in the frequency-domain. Therefore with the correct addresses for the values to sum at each layer, spectral estimates were produced for successive layers through reuse of the spectral estimate from the previous layer. This was done through the use of the same filter and hence would be produced from the same hardware implementation. Such a scheme in real hardware would suffer from a latency which would be l times greater than a scheme which implemented all the moving average filters in parallel but would be highly efficient (where l refers to the total number of layers). Due to the increased latency and the repeated use of the same hardware to perform the smoothing, the decision was made to employ handshakes between blocks to control the flow of data. The options were either to clock the filter at a rate l times faster than the rate of serial input to the FFT block or disable the serial input and FFT block until the filter block was ready to accept new data. The latter option was chosen for this implementation. This is in contrast to the fixed-point SVA implementation where a synchronous pipeline was used and the division block in the SVA algorithm was required to be clocked at a faster rate. The consequence of the choice to implement handshakes between modules required an FSM to control the flow of data.

The best interpretation of how the single filter scheme for one ordinate is achieved is as a combination of a few simple elements: a register, an adder and a barrel shifter. Furthermore the register and adder can be interpreted together as an accumulator.

H.3.1 Register

The nature of the tree smoothed detectors means that the results from the previous spectral estimate were incorporated into the calculation of the next spectral estimate. The initial spectral estimate passed to the accumulator register from the SVA periodogram was based on 18 bit wordlengths. The Virtex-5[®] board includes a 48 bit accumulator register in each DSP48E slice so there was no risk of overflow due to wordlength growth for realistic FFT sizes.



For example, with a relatively long FFT of length $N_{\text{FFT}} = 2048$, the final spectral estimate would be possible at $hyp = 12$ (which would also be a total power estimate) and the full-precision wordlength would be 36 bit.

H.3.2 Adder

The addition block relied on the formation of the summation of the correct two input ordinates from the spectral estimate of the previous layer. In other words the correct counterpart for each ordinate was required to have the overall desired effect of a summation tree which averages the initial spectral estimate in increasing powers of two. This was achieved through the summation of each ordinate with a neighbour 2^{hyp-1} places to the right. The addresses were circular so that ordinates at the right-hand edge were added to those at the left-hand edge.

$$+ \hat{S}_{hyp-1}(k + 2^{hyp-1})$$

To compensate for the drift of the centre frequency of the two ordinates to the right, the addition block was required to specify the correct output address for the result. One peculiarity of the tree smoothed technique as a whole was that the initial spectral estimate was sampled at integer frequency points but all subsequent spectral estimates at greater hyp values were sampled at half-integer frequency points as a result of the formation of summations in powers of two.

H.3.3 Shifter

Each successive layer summed two frequency ordinates. Therefore a common divisor of two was required to form an average for each spectral estimate. This differs from the moving average smoothed techniques formulated in section 4.3 where the divisor could be any odd integer. The common divisor of two was highly beneficial because it could be simply implemented as a right arithmetic barrel shift in fixed-point arithmetic.

$$>> (hyp - 1)$$

H.4 Results

The results obtained from the fixed-point model of the TSSP (detector M2) were a good match to the floating-point model ROC curve results aside from a negligible amount of error due to quantisation effects. All the scenarios from chapter 5 were examined and repeated but only two interesting scenarios for the Barker-13 pulse are shown here: (1) full DCW and (2) full interference DCW.

H.4.1 Full DCW

The parameters used for the fixed-point full DCW scenario were a carbon copy of those which had been used in the floating-point version. Specifically these were:

- Barker-13 pulse with characteristics from Table 5.2.1 and AWGN inputs;
- random initial phase of the signal, ϕ_0 ;
- random frequency, f_T , of the signal over the interval f_s/N_{FFT} ;
- finite duration phase jumps;
- random portion of the LPI signal pulse captured;
- rectangular window function before FFT then SVA applied;
- detector M2 $P_{fa} = 0.01$;
- non-overlapped DCWs;
- detection on each DCW;
- non-zeropadded and zeropadded FFTs used;
- 1×10^5 trials to produce a data point in the graphs.

The results in Fig H.4.1 show that the TSSP (detector M2) can provide a similar level of performance to the floating-point TSSP detector M1 in a feasible fixed-point scheme.

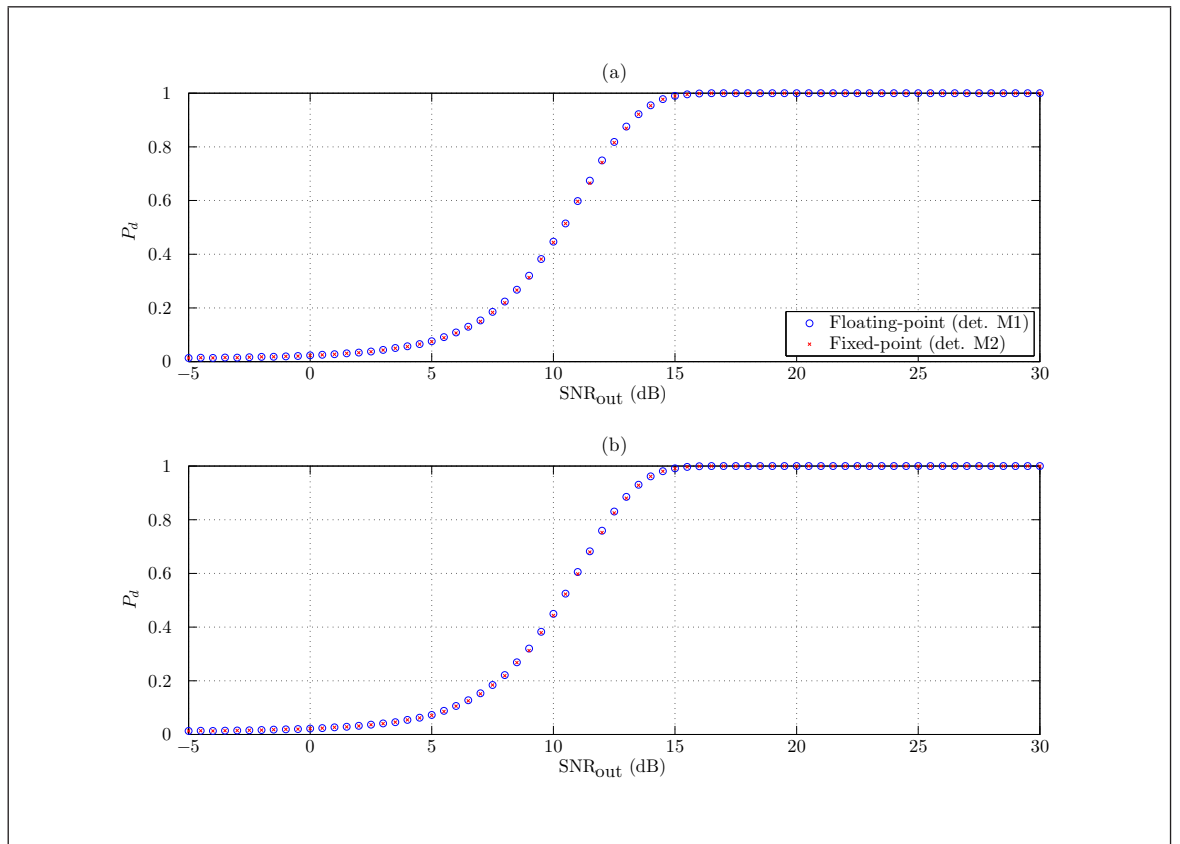


Figure H.4.1 Detector M2 ROC data points (\times) overlaid on floating-point detector M1 points (\circ) for Barker-13 pulse and for : (a) the non-zero padded case; (b) the zero padded case

H.4.2 Interference in Full DCW

The same detector M2 was used in the fixed-point interference full DCW scenario. The scenario was undertaken with identical parameters to the corresponding scenario from the floating-point section 5.4. These included:

- a desired LPI signal, S_A , an interference complex exponential, S_B and AWGN input;
- random frequencies $f_{T,A}$ and $f_{T,B}$ for S_A and S_B respectively, each was randomly chosen over the interval f_s/N_{FFT} at commencement of a test run;
- finite duration phase jumps;
- random portion of the LPI signal captured;
- rectangular window function before FFT and then SVA applied;
- random initial phases for the signals, $\phi_{0,A}$ and $\phi_{0,B}$;
- detector M2 $P_{fa} = 0.01$;
- non-overlapped DCWs;
- detection on each DCW;
- 1×10^5 trials to produce a data point in the graphs;
- average frequency separation between $f_{T,A}$ and $f_{T,B}$ was $5f_s/N$.

The fixed-point ROC curve results are shown overlaid on the corresponding floating-point result in Figs H.4.2. The results indicate that the transition to fixed-point arithmetic did not affect the performance of the detector too greatly except for a small amount of error caused by quantisation effects.

It is not necessary to reiterate the conclusions about the TSSP from the floating-point chapter. The main conclusion from this set of simulations is that detector M2 is realisable in a feasible fixed-point system.

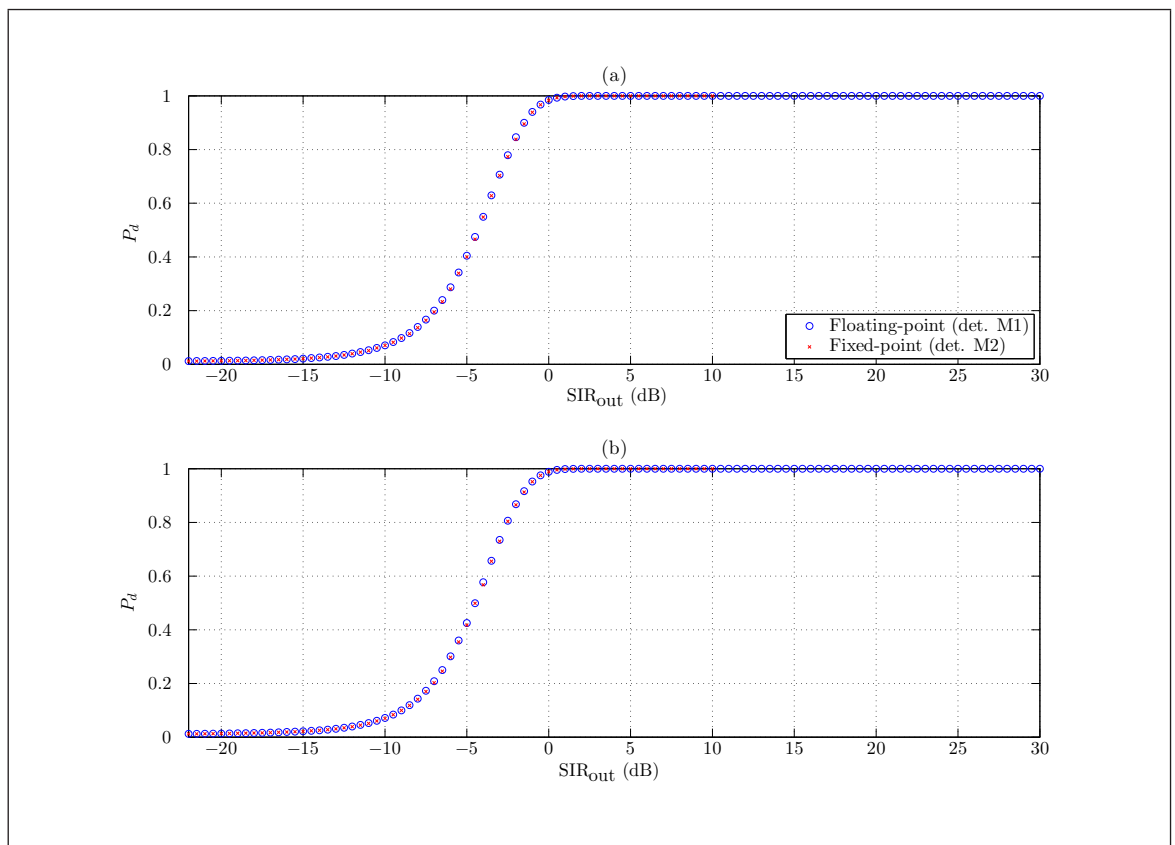


Figure H.4.2 Detector M2 ROC data points (\times) overlaid on floating-point detector M1 points (\circ) at $SNR_{out,B} = 15$ dB for Barker-13 pulse for: (a) the non-zeropadded case; (b) the zeropadded case

H.5 Conclusions

Detector M2 was a fixed-point implementation of the TSSP. The fixed-point arithmetic was fashioned to match the Xilinx[®] Virtex-5[®] capabilities because this facilitated the transition to implementation in FPGA in the next stage of the MBD methodology. The ROC curve results obtained from detector M2 supported the notion that the use of the TSSP in a feasible fixed-point arithmetic system would not be detrimental to an unacceptable extent to the detector performance.

APPENDIX I

TSP-on-FPGA

The floating-point simulations which involved the algorithms of interest in chapter 5 revealed that the TSSP technique is useful in the detection of broader-band LPI signals in the scenarios tested in comparison to conventional algorithms. Although the detector M1, which employed this algorithm, had not enjoyed the best detection performance in all the scenarios tested, this was offset against its other advantageous features and hence it was chosen as the best candidate for implementation on FPGA.

The algorithm was converted to a fixed-point implementation in appendix H. The performance results obtained from a detector which employed the fixed-point algorithm converged to the floating-point results. Furthermore the fixed-point implementation had the option to switch in or out the SVA module, which thus created the TSP (detector L2 variant) or the TSSP (detector M2) respectively.

Due to time constraints and due to the main interest vesting in the tree smoothing component of the detector rather than the SVA component, the TSP (detector L2 variant) was implemented in FPGA. This section discusses the task in detail. Therefore although detector M2 performance was not characterised on FPGA, the physical information about the TSP (detector L2 variant) implementation on FPGA was uncovered. This included information such as the resources required, its speed, its latency and throughput. Thus the main aim, to characterise the tree smoothing component of the detector was satisfied and the main characteristics could be combined with the SVA-on-FPGA characteristics to form a crude estimate of the overall TSSP (detector M2) characteristics.

I.1 Implementation Decisions

The same target hardware and simulation languages were used for TSP-on-FPGA as were used for the SVA-on-FPGA task (Xilinx[®] Virtex-5[®] ML505 device XCVLX50T and simulation languages Simulink[®] and System Generator[®]).

I.2 Coding Methodology

The phased code procedure onto FPGA was followed. The phases reflected the lessons from the SVA-on-FPGA task.

I.2.1 Phase 1

The first phase was the production of a System Generator[®] model. This involved the conversion of the modules from Fig 5.1.7 to fixed-point System Generator[®] blocks. The top-level of the resultant model is shown in Fig I.2.1. For the avoidance of doubt there was no SVA module implemented in this model. The model included implementations of all the modules required for a TSP detector such as an FFT, tree smoothing filter and threshold logic. The analogue, real-world inputs were generated off-board. Although time-constraints did not allow, there is a possibility to add an SVA module in this architecture in the future.

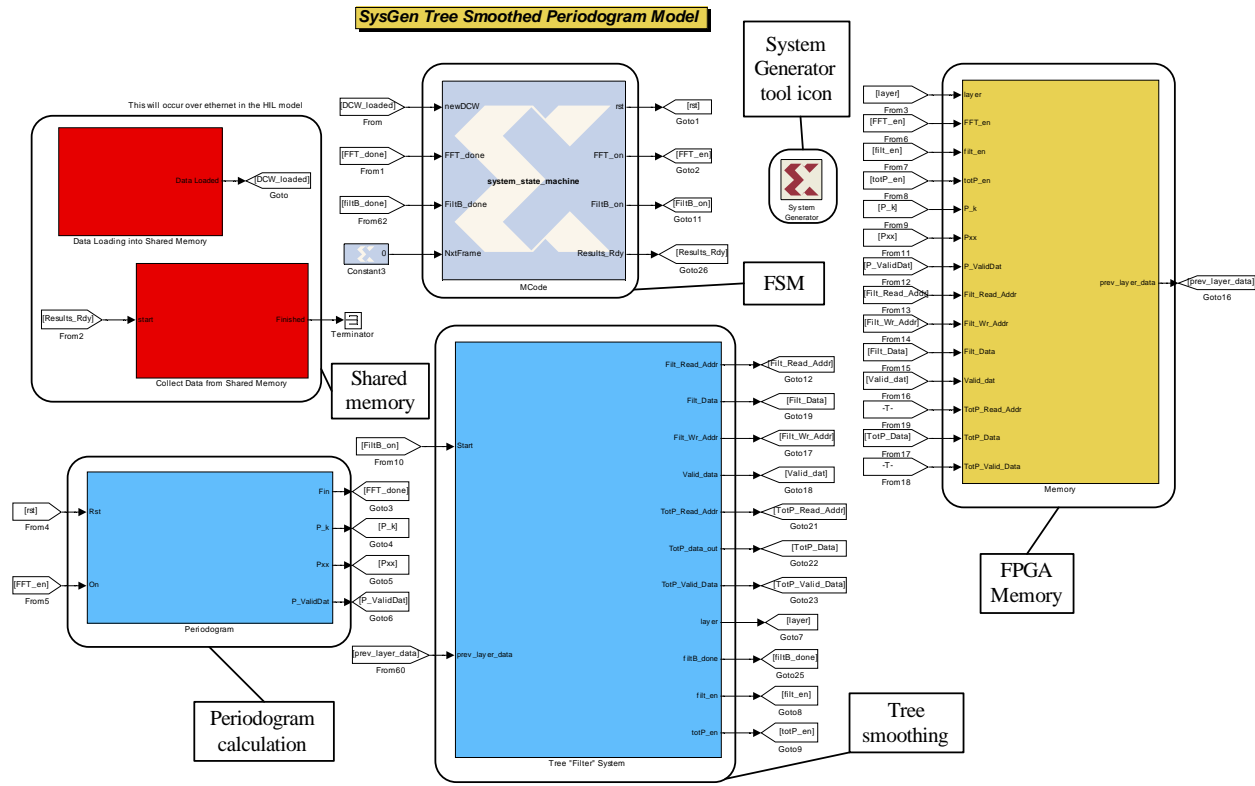


Figure I.2.1.1 Screenshot of the phase 1 TSP-on-FPGA model

Although the resolution of Fig I.2.1 does not permit the fine details of the model to be deduced from the picture, the labels provide some assistance. The picture shows that the multirate clock tree architecture from the SVA-on-FPGA section has been replaced with an FSM-controlled data flow. Whilst the SVA-on-FPGA architecture had the advantage of production of a constant stream of processed output data, this necessitated an undesirable multi-rate clock tree. Therefore, the TSP-on-FPGA detector was coded in a highly-stable FSM-controlled process such that it produced data at regular intervals but would drop samples at the input.

The picture in Fig I.2.1 also shows the memory was not organised into FIFOs. This was done to address the concerns about corrupted FIFOs from section F.5. In the picture, the left-hand memory blocks were shared memory¹ in anticipation of the independent execution of the detector in the laptop or PC co-simulation. The inputs for the FPGA were read from the shared memory at the start of a hardware co-simulation and the FPGA results were written to the shared memory at the end.

In terms of the FFT implementation in the phase 1 model, the output words of the FFT was permitted to use more bits than the input words. In other words the LSBs were not dropped after the effective division by N^2 ². This was in contrast to the SVA-on-FPGA FFT implementation and was implemented in this manner because a larger dynamic range was necessary to detect broader-band, smaller signal peaks.

The phase 1 TSP-on-FPGA model was verified against the previous models from the MBD methodology and found to be equivalent to an acceptable tolerance.

Some of the characteristics of the phase 1 TSP-on-FPGA were extracted from the place and route report. An important excerpt of the report is shown in Table I.2.1. The resource demands given in the table are quite modest, especially when it is noted that the model is the full TSP (detector L2 variant) and not simply the smoothing component. The one exception to this is the number of utilised bonded IOBs. This is quite high at 243. This was a consequence of the incorporation of the consideration for hardware co-simulation in the code. Communication with the laptop or PC over the shared memory required many IO pins. This was not a concern, however, because a real-life system would not require these connections. This is because they are, in effect, only for test purposes.

To achieve an estimation of the resource requirements for a TSSP detector, the values from Table I.2.1 could be added to those from Table F.2.2. Meanwhile, the post place and route static timing report declared an estimated maximum clock frequency of ~ 100 MHz for the phase 1 model. This is shown in Table I.2.2. Clearly, some further critical path analysis would be required if a derivative of this model were to be implemented.

¹Memory was shared between the laptop or PC and the FPGA

²Implemented through division of the individual butterflies by 2

Table I.2.1 Extract from phase 1 TSP-on-FPGA model place and route report

Device Utilization Summary:

Number of BUFGs	1 out of 32	3%
Number of DSP48Es	15 out of 48	31%
Number of External IOBs	243 out of 480	50%
Number of LOCed IOBs	0 out of 243	0%
Number of RAMB18X2s	9 out of 60	15%
Number of RAMB18X2SDPs	4 out of 60	6%
Number of RAMB36SDP_EXPs	8 out of 60	13%
Number of Slice Registers	1515 out of 28800	5%
Number used as Flip Flops	1515	
Number used as Latches	0	
Number used as LatchThrus	0	
Number of Slice LUTS	1059 out of 28800	3%
Number of Slice LUT-Flip Flop pairs	1582 out of 28800	5%

Table I.2.2 Extract from phase 1 TSP-on-FPGA post place and route static timing report

Timing summary:

Timing errors:	0 Score: 0 (Setup/Max: 0, Hold: 0)
Constraints cover	9081 paths, 0 nets, and 6815 connections
Design statistics:	
Minimum period:	10.244 ns ³ (Maximum frequency: 97.618 MHz)

I.2.2 Phase 2

The second phase of production of TSP-on-FPGA was undertaken specifically to address the “testing over ethernet” concern from section F.5. In short, the phase 1 model was modified such that custom SELEX Galileo MATLAB[®] application programming interfaces (APIs) could be used to communicate with the FPGA board from a laptop or PC over an ethernet connection. This ensured that the test and verification occurred continuously and the FPGA was not reset to an initial state between test vectors. It also significantly increased the rate of test because the PC or laptop and FPGA were not forced into lockstep with each other as were during the test phase of SVA-on-FPGA over the JTAG connection.

The screenshot in Fig I.2.2 shows the modifications and wrappers which were applied to the phase 1 model to allow communication with the board over ethernet.

³The minimum period statistic assumes all single cycle delays. Analysis completed Mon May 18 13:16:06 2010

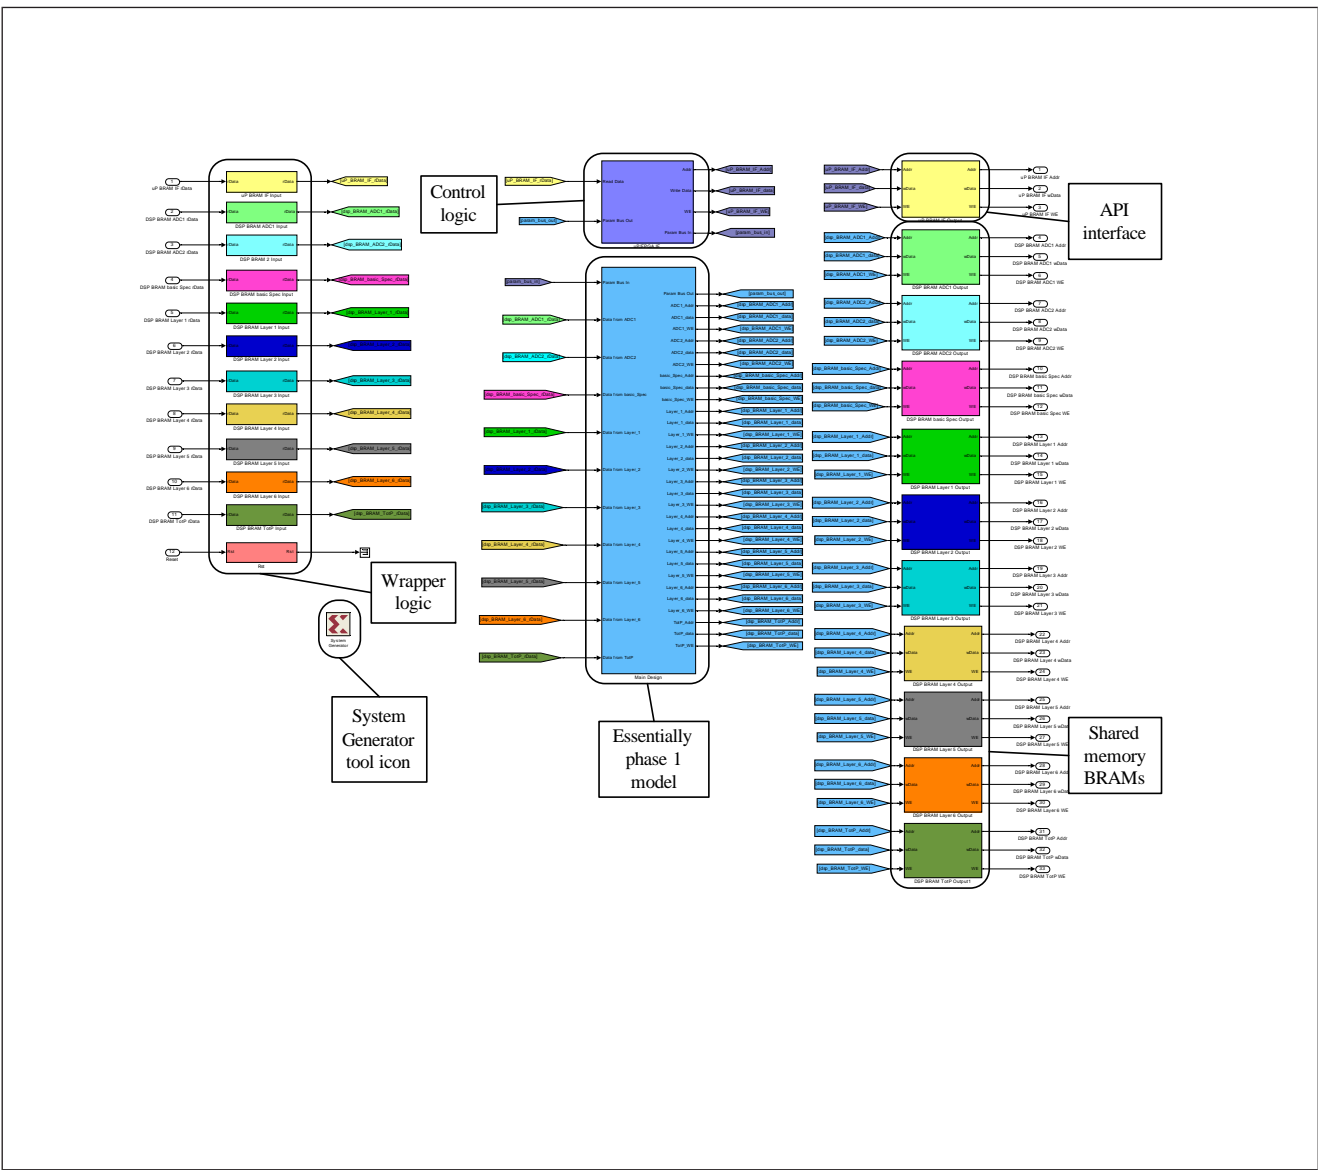


Figure I.2.2 Screenshot of the phase 2 TSP-on-FPGA model

The extra wrappers around the phase 1 model had an impact on the device resource utilisation in the FPGA. This is shown from an extract of the place and route report in Table I.2.3. The resource requirements are still fairly modest, again, except for the bonded IOBs. For the same reason given in the phase 1 model analysis, the number of utilised bonded IOBs was not a concern.

Table I.2.3 Extract from phase 2 TSP-on-FPGA model place and route report

Device Utilization Summary:		
Number of BUFGs	1 out of 32	3%
Number of DSP48Es	15 out of 48	31%
Number of External IOBs	470 out of 480	97%
Number of LOCed IOBs	0 out of 470	0%
Number of RAMB18X2s	9 out of 60	15%
Number of RAMB18X2SDPs	3 out of 60	5%
Number of Slice Registers	1868 out of 28800	6%
Number used as Flip Flops	1868	
Number used as Latches	0	
Number used as LatchThrus	0	
Number of Slice LUTS	1158 out of 28800	4%
Number of Slice LUT-Flip Flop pairs	1937 out of 28800	6%

The post place and route static timing report revealed that the maximum clock frequency of the phase 2 model was estimated to be ~ 150 MHz. The extract is shown in Table I.2.4.

Table I.2.4 Extract from phase 2 TSP-on-FPGA post place and route static timing report

Timing summary:	
Timing errors:	0 Score: 0 (Setup/Max: 0, Hold: 0)
Constraints cover	9391 paths, 0 nets, and 6882 connections
Design statistics:	
Minimum period:	6.859 ns ⁴ (Maximum frequency: 145.794 MHz)

I.3 Monte Carlo Simulation Tests

The main aim of TSP-on-FPGA was not to explore the performance of the TSP algorithm. That task was examined in the floating-point simulations in chapter 5. However some Monte Carlo simulations were undertaken to provide confidence in the implementation through a comparison of the results with the floating-point simulation results.

⁴The minimum period statistic assumes all single cycle delays. Analysis completed Mon May 24 17:55:26 2010

The simulations undertaken were: (1) the full DCW scenario for a Barker-13 signal input and (2) the interference full DCW scenario for a desired Barker-13 signal. The TSP-on-FPGA results are shown overlaid on the corresponding floating-point results in Figs I.3.1 and I.3.2. The floating-point results are simply the same data points reproduced from Figs 5.2.2 and 5.4.14 respectively for the TSP (detector variant L1). The results provided sufficient confidence in the FPGA implementation.

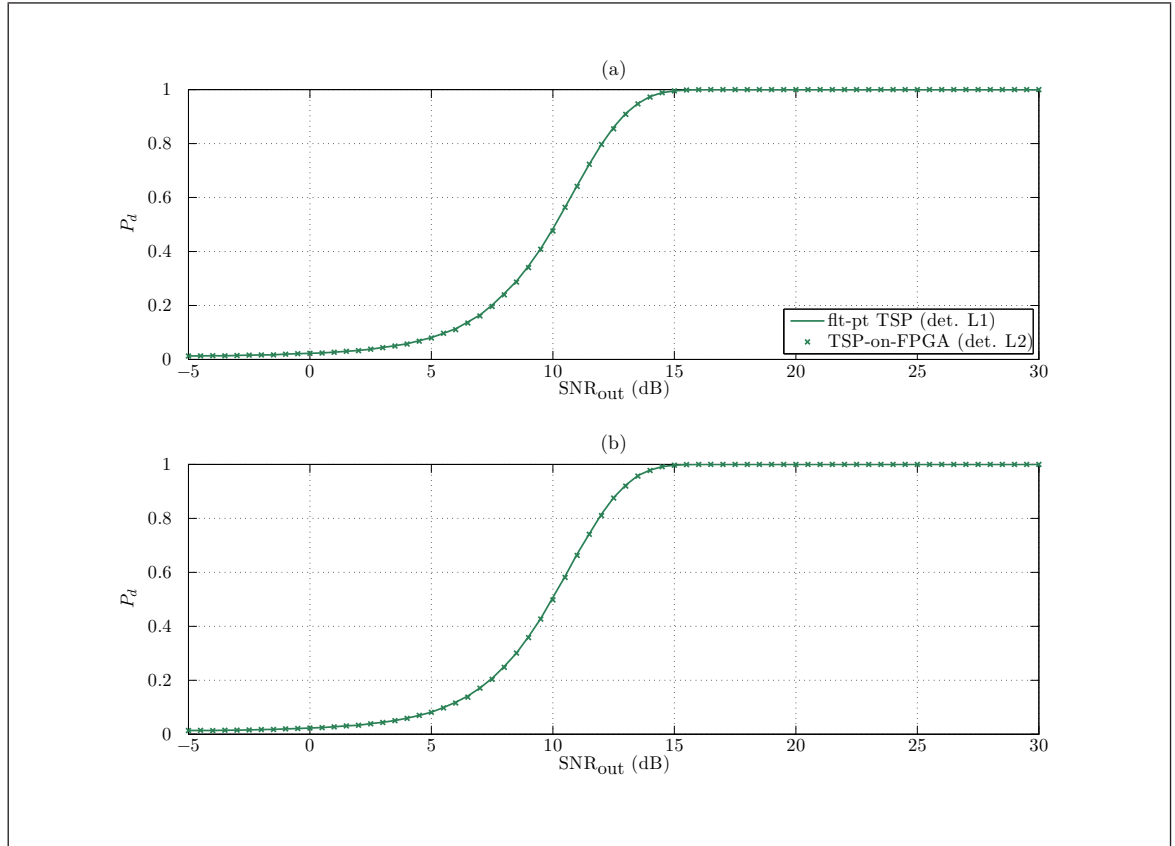


Figure I.3.1 Full DCW scenario ROC curves of floating-point TSP and TSP-on-FPGA for: (a) the non-zeropadded case; (b) the zeropadded case

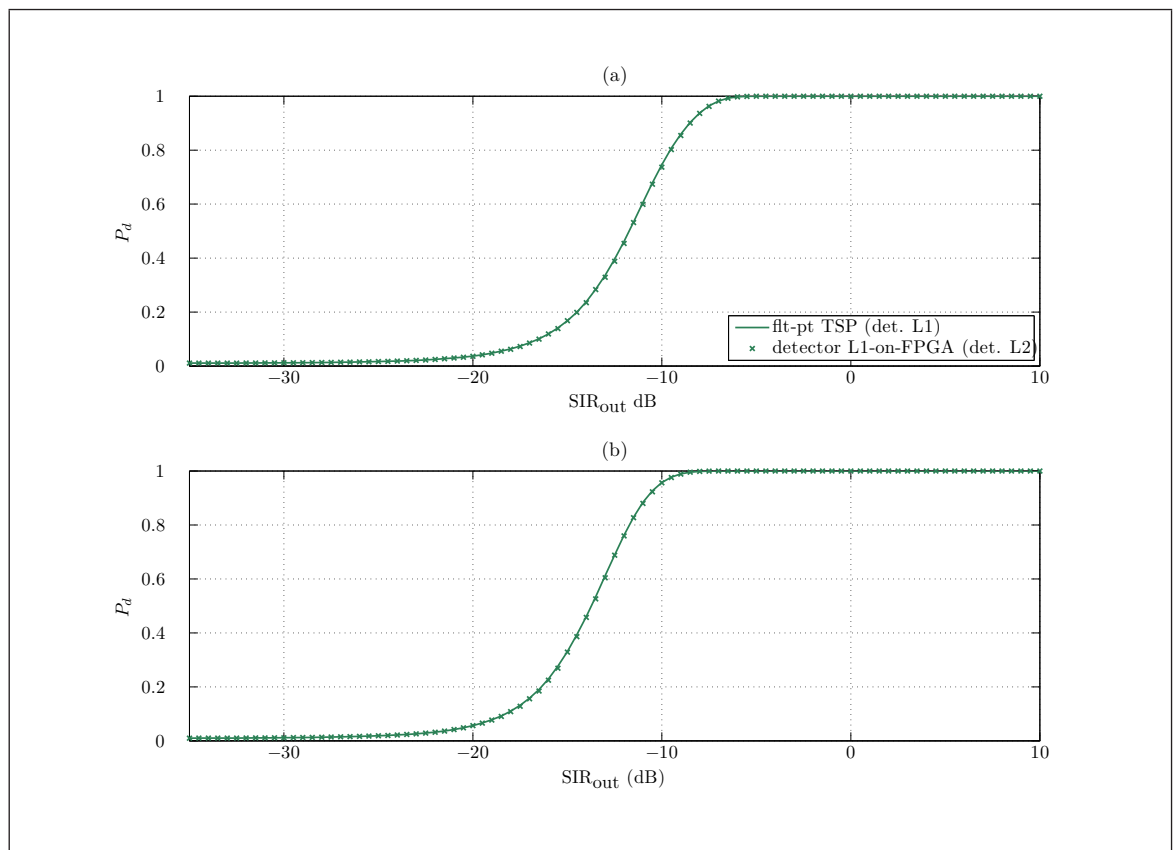


Figure I.3.2 Interference full DCW scenario ROC curves of floating-point TSP and TSP-on-FPGA for: (a) the non-zero padded case; (b) the zero padded case

I.4 Suggested Design Improvements

A desirable goal would be to have a detector with improved detection capability over conventional EW DRxs against narrowband radar signals through the use of SVA and against broader-band LPI signals through the use of a tree smoothing scheme. Therefore it would be ideal to combine the SVA IP and the tree smoothing IP. However the SVA IP was designed for a synchronous pipeline whilst the smoothing IP was designed to drop samples at the input if necessary and be controlled by an FSM. Therefore some extra work would be necessary to translate one of the architecture forms into the other before they could be combined in all-purpose detector M1 on FPGA.

APPENDIX J

Patent Application & Paper

The following pages show the patent application exactly as filed and the scientific paper exactly as submitted.

The patent application covers the use of SVA in a detector. It was filed as patent application number 1104107.6 with reference 001510GB.

The scientific paper covers the use of SVA as a detector. It was submitted for *Defense Applications of Signal Processing* (DASP) 2011.

Glossary

Terminology

chirp class

An internal SELEX Galileo classification system whereby chirps with a slow chirp rate are class A; chirps with a medium chirp rate are class B and chirps with a fast chirp rate are class C.

data capture window

A data capture window is a sampled snapshot of the baseband waveform passed to the digital receiver. The data capture window is typically composed of a number of samples which is a power of two as it is typically processed by a fast Fourier transform.

minimum operational sensitivity

The minimum input signal impinging on an antenna at the system level to produce the minimum detectable signal at a specified output SNR.

Model Based Design

A design methodology to reduce development time and costs and to promote design re-use. This is achieved through construction of models of systems with the use of functional, modular building blocks.

modulation-on-pulse

Pulsed radar waveforms that incorporate a phase or frequency coding scheme for radar coherent integration.

ordinate

The magnitude-squared of a complex sample from the output of one of the DFT channels.

Bibliography

- [1] K. Fuller, “To see and not be seen [radar],” *Radar and Signal Processing, IEE Proceedings F*, vol. 137, no. 1, pp. 1 –10, February 1990.
- [2] G. Schrick and R. Wiley, “Interception of LPI radar signals,” in *Radar Conference, 1990., Record of the IEEE 1990 International*, May 1990, pp. 108 –111.
- [3] J. B. Y. Tsui, *Digital Techniques for Wideband Receivers, Second Edition*. Norwood, MA, USA: Artech House, Inc., 2001.
- [4] G. Lopez-Risueno, J. Grajal, and A. Sanz-Osorio, “Digital channelized receiver based on time-frequency analysis for signal interception,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, no. 3, pp. 879 – 898, July 2005.
- [5] F. Harris, “On the use of windows for harmonic analysis with the discrete fourier transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51 – 83, January 1978.
- [6] A. Nuttall, “Near-optimum detection of random signals of unknown location, structure, extent, and strength,” in *OCEANS 95. MTS/IEEE. Challenges of Our Changing Global Environment. Conference Proceedings.*, vol. 3, October 1995, pp. 1659–1664.
- [7] R. Dillard, “Detectability of Spread-Spectrum Signals,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-15, no. 4, pp. 526 –537, July 1979.
- [8] R. A. Dillard and G. H. Dillard, *Detectability of Spread-Spectrum Signals*. Norwood, MA, USA: Artech House, Inc., 1989.
- [9] H. Urkowitz, “Energy detection of unknown deterministic signals,” *Proceedings of the IEEE*, vol. 55, no. 4, pp. 523 – 531, April 1967.
- [10] G. Dillard, “A moving-window detector for binary integration,” *Information Theory, IEEE Transactions on*, vol. 13, no. 1, pp. 2 – 6, January 1967.
- [11] P. van Genderen, “Recent advances in waveforms for radar, including those with communication capability,” in *Radar Conference, 2009. EuRAD 2009. European*, October 2009, pp. 318 –325.
- [12] J. Townsend, M. Saville, S. Hong, R. Martin, C. Simpson, and O. Mayhew, “Waveform optimization for electronic countermeasure technique generation,” in *Radar Conference, 2008. RADAR '08. IEEE*, May 2008, pp. 1 –6.
- [13] P. E. Pace, *Detecting and Classifying Low Probability of Intercept Radar*. Artech House, Inc., 2009.

- [14] N. Levanon and E. Mozeson, *Radar Signals*, 1st ed. John Wiley & Sons, 2004.
- [15] A. Stove, "Linear FMCW radar techniques," *Radar and Signal Processing, IEE Proceedings F*, vol. 139, no. 5, pp. 343 –350, October 1992.
- [16] A. Stove, A. Hume, and C. Baker, "Low probability of intercept radar strategies," *Radar, Sonar and Navigation, IEE Proceedings -*, vol. 151, no. 5, pp. 249 – 260, October 2004.
- [17] P. Willett, J. Reinert, and R. Lynch, "LPI waveforms for active sonar?" in *Aerospace Conference, 2004. Proceedings. 2004 IEEE*, vol. 4, March 2004, pp. 2236 – 2248 Vol.4.
- [18] J. Fielding, "Polytime coding as a means of pulse compression," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 2, pp. 716 –721, April 1999.
- [19] H. Jiangang, T. Ran, S. Tao, and Q. Lin, "A novel LPI radar signal based on hyperbolic frequency hopping combined with Barker phase code," in *Signal Processing, 2004. Proceedings. ICSP '04. 2004 7th International Conference on*, vol. 3, August 2004, pp. 2070 – 2073 vol.3.
- [20] E. Fowle, "The design of FM pulse compression signals," *Information Theory, IEEE Transactions on*, vol. 10, no. 1, pp. 61 – 67, January 1964.
- [21] J. Johnston and A. Fairhead, "Waveform design and doppler sensitivity analysis for nonlinear FM chirp pulses," *Communications, Radar and Signal Processing, IEE Proceedings F*, vol. 133, no. 2, pp. 163 –175, April 1986.
- [22] J. Costas, "A study of a class of detection waveforms having nearly ideal range, #8212; Doppler ambiguity properties," *Proceedings of the IEEE*, vol. 72, no. 8, pp. 996 – 1009, August 1984.
- [23] N. Levanon and E. Mozeson, "Modified Costas signal," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 40, no. 3, pp. 946 – 953, July 2004.
- [24] R. Barker, "Group synchronization of binary digital systems," *Communication Theory*, pp. 273 –287, 1953.
- [25] S. Golomb, "Two-valued sequences with perfect periodic autocorrelation," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 383 –386, April 1992.
- [26] L. Bomer and M. Antweiler, "Polyphase Barker sequences," *Electronics Letters*, vol. 25, no. 23, pp. 1577 –1579, November 1989.
- [27] M. Friese and H. Zottmann, "Polyphase Barker sequences up to length 31," *Electronics Letters*, vol. 30, no. 23, pp. 1930 –1931, November 1994.

- [28] M. Friese, "Polyphase Barker sequences up to length 36," *Information Theory, IEEE Transactions on*, vol. 42, no. 4, pp. 1248–1250, July 1996.
- [29] A. Brenner, "Polyphase Barker sequences up to length 45 with small alphabets," *Electronics Letters*, vol. 34, no. 16, pp. 1576–1577, August 1998.
- [30] P. Borwein and R. Ferguson, "Polyphase sequences with low autocorrelation," *Information Theory, IEEE Transactions on*, vol. 51, no. 4, pp. 1564–1567, April 2005.
- [31] D. Huffman, "The generation of impulse-equivalent pulse trains," *Information Theory, IRE Transactions on*, vol. 8, no. 5, pp. 10–16, September 1962.
- [32] R. Frank, S. Zadoff, and R. Heimiller, "Phase shift pulse codes with good periodic correlation properties (Corresp.)," *Information Theory, IRE Transactions on*, vol. 8, no. 6, pp. 381–382, October 1962.
- [33] R. L. Frank, "Phase coded communication system," US Patent 3 099 795, July, 1963. [Online]. Available: <http://www.freepatentsonline.com/3099795.html>
- [34] B. Lewis and F. Kretschmer, "A New Class of Polyphase Pulse Compression Codes and Techniques," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-17, no. 3, pp. 364–372, May 1981.
- [35] —, "Correction to: "A New Class of Polyphase Pulse Compression Codes and Techniques"," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-17, no. 5, p. 726, September 1981.
- [36] —, "Linear Frequency Modulation Derived Polyphase Pulse Compression Codes," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-18, no. 5, pp. 637–641, September 1982.
- [37] F. Kretschmer and B. Lewis, "Doppler properties of polyphase coded pulse compression waveforms," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-19, no. 4, pp. 521–531, July 1983.
- [38] P. Rapajic and R. Kennedy, "Merit factor based comparison of new polyphase sequences," *Communications Letters, IEEE*, vol. 2, no. 10, pp. 269–270, October 1998.
- [39] G. Painchaud, J. McKenzie, M. Blanchette, and A. Voy, "An experimental adaptive digital pulse compression subsystem for multi-function radar applications," in *Radar Conference, 1990., Record of the IEEE 1990 International*, May 1990, pp. 153–158.
- [40] N. Zhang and S. Golomb, "Polyphase sequence with low autocorrelations," *Information Theory, IEEE Transactions on*, vol. 39, no. 3, pp. 1085–1089, May 1993.

- [41] Z. Solomon, "Phase coded communication system," US Patent 3 099 796, July, 1963. [Online]. Available: <http://www.freepatentsonline.com/3099796.html>
- [42] D. Chu, "Polyphase codes with good periodic correlation properties (Corresp.)," *Information Theory, IEEE Transactions on*, vol. 18, no. 4, pp. 531 – 532, July 1972.
- [43] T. Felhauer, "Design and analysis of new $p(n,k)$ polyphase pulse compression codes," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 30, no. 3, pp. 865 –874, July 1994.
- [44] R. Ardoino and A. Megna, "LPI Radar detection: SNR performances for a dual channel Cross-Correlation based ESM Receiver," in *Radar Conference, 2009. EuRAD 2009. European*, October 2009, pp. 113 –116.
- [45] A. Stove, "Radar and ESM: The Current State of the LPI Battle," in *1st EMRS DTC Technical Conference*, vol. A14, 2004.
- [46] P. Wu, "On sensitivity analysis of low probability of intercept (LPI) capability," in *Military Communications Conference, 2005. MILCOM 2005. IEEE*, October 2005, pp. 2889 –2895 Vol. 5.
- [47] J. Stephens, "Advances in signal processing technology for electronic warfare," in *Aerospace and Electronics Conference, 1996. NAECON 1996., Proceedings of the IEEE 1996 National*, vol. 1, May 1996, pp. 129 –136 vol.1.
- [48] F. Taboada, A. Lima, J. Gau, P. Jarpa, and P. E. Pace, "Intercept receiver signal processing techniques to detect low probability of intercept radar signals," in *Proceedings of the 5th Nordic Signal Processing Symposium*, Hurtigruta Tromso, Norway, October 2002.
- [49] E. Carpentieri and S. Cuomo, "An adaptive threshold algorithm for detection of pulse radar signals," in *Radar Conference, 2008. RADAR '08. IEEE*, May 2008, pp. 1 –5.
- [50] P. Camuso, G. Foglia, and D. Pistoia, "A comprehensive analysis on detection performances of LPI signals filtering strategies," in *Radar Conference, 2009. EuRAD 2009. European*, October 2009, pp. 125 –128.
- [51] C. Lee, Jim P. Y. (Nepean), "Circuit for LPI signal detection and suppression of conventional pulsed signals," US Patent 6 388 604, May, 2002. [Online]. Available: <http://www.freepatentsonline.com/6388604.html>
- [52] A. K. Nandi and E. E. Azzouz, "Automatic analogue modulation recognition," *Signal Process.*, vol. 46, pp. 211–222, October 1995. [Online]. Available: <http://portal.acm.org/citation.cfm?id=217632.217639>

- [53] J. Lunden and V. Koivunen, "Scaled Conjugate Gradient Method for Radar Pulse Modulation Estimation," in *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, vol. 2, April 2007, pp. II-297–II-300.
- [54] S. Colonnese and G. Scarano, "Transient signal detection using higher order moments," *Signal Processing, IEEE Transactions on*, vol. 47, no. 2, pp. 515–520, February 1999.
- [55] Z. Wang and P. Willett, "A performance study of some transient detectors," *Signal Processing, IEEE Transactions on*, vol. 48, no. 9, pp. 2682–2685, September 2000.
- [56] T. W. Anderson and D. A. Darling, "Asymptotic Theory of Certain "Goodness of Fit" Criteria Based on Stochastic Processes," *The Annals of Mathematical Statistics*, vol. 23, no. 2, pp. 193–212, June 1952.
- [57] —, "A Test of Goodness of Fit," *Journal of the American Statistical Association*, vol. 49, no. 268, pp. 765–769, December 1954.
- [58] M. A. Stevens and R. D. D'Agostino, *Goodness-of-fit-techniques*, 1st ed. Dekker, June 1986.
- [59] R. Prasad, H. Saruwatari, and K. Shikano, "Negentropy based voice-activity detection for noise estimation in very low SNR condition," *IEICE Electronics Express*, vol. 1, no. 16, pp. 495–500, 2004.
- [60] H. L. V. Trees, *Detection, Estimation, and Modulation Theory: Radar-Sonar Signal Processing and Gaussian Signals in Noise*. Melbourne, FL, USA: Krieger Publishing Co., Inc., 1992.
- [61] C. Cornu, C. Ioana, and A. Quinquis, "Characterization of LPI Waveforms using a Polynomial Phase Signal Modelling," in *RADAR 2004 - International Conference on Radar Systems*, October 2004.
- [62] L. Anjaneyulu, N. V. S. N. Sarma, and N. S. Murthy, "Identification of LPI radar signals by higher order spectra and neural network techniques," *Int. J. Inf. Commun. Techol.*, vol. 2, pp. 142–155, June 2009. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1554588.1554601>
- [63] F. Sattar and G. Salomonsson, "On detection using filter banks and higher order statistics," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 36, no. 4, pp. 1179–1189, October 2000.
- [64] Z. Wang and P. Willett, "All-purpose and plug-in power-law detectors for transient signals," *Signal Processing, IEEE Transactions on*, vol. 49, no. 11, pp. 2454–2466, November 2001.

- [65] J. Lehtomaki, M. Juntti, and H. Saarnisaari, "Power-law based intercept receiver with nonorthogonal transforms," in *Military Communications Conference, 2004. MILCOM 2004. IEEE*, vol. 3, October 2004, pp. 1402 – 1408 Vol. 3.
- [66] J. Lehtomaki, "Detection of spread spectrum signals using a power-law based intercept receiver," in *Vehicular Technology Conference, 2004. VTC 2004-Spring. 2004 IEEE 59th*, vol. 3, May 2004, pp. 1480 – 1484 Vol.3.
- [67] R. Kronland-Martinet, "The wavelet transform for analysis, synthesis, and processing of speech and music sounds," *Computer Music Journal*, vol. 12, no. 4, pp. pp. 11–20, 1988. [Online]. Available: <http://www.jstor.org/stable/3680149>
- [68] S. Kadambe and G. Boudreaux-Bartels, "Application of the wavelet transform for pitch detection of speech signals," *Information Theory, IEEE Transactions on*, vol. 38, no. 2, pp. 917 –924, March 1992.
- [69] G. Tzanetakis, G. Essl, and P. Cook, "Audio Analysis using the Discrete Wavelet Transform," in *in Proc. Conf. in Acoustics and Music Theory Applications. WSES*, 2001.
- [70] R. D. Hippenstiel, *Detection Theory: Applications and Digital Signal Processing*. CRC Press, November 2001.
- [71] K. Stoffell, "Implementation of a Quadrature Mirror Filter Bank on an SRC Reconfigurable Computer for Real-time Signal Processing," Master's thesis, Naval Postgraduate School, Monterey, California, September 2006.
- [72] T. Farrell and G. Prescott, "A low probability of intercept signal detection receiver using quadrature mirror filter bank trees," in *Acoustics, Speech, and Signal Processing, 1996. ICASSP-96. Conference Proceedings., 1996 IEEE International Conference on*, vol. 3, May 1996, pp. 1558 –1561 vol. 3.
- [73] D. Copeland and P. Pace, "Detection and analysis of FMCW and P-4 polyphase LPI waveforms using quadrature mirror filter trees," in *Acoustics, Speech, and Signal Processing, 1993. ICASSP-93., 1993 IEEE International Conference on*, vol. 4, April 2002, p. IV.
- [74] F. Millioz and M. Davies, "Detection and segmentation of fmcw radar signals based on the chirplet transform," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, may 2011, pp. 1765 –1768.
- [75] P. Milne and P. Pace, "Wigner distribution detection and analysis of FMCW and P-4 polyphase LPI waveforms," in *Acoustics, Speech, and Signal Processing, 1993. ICASSP-93., 1993 IEEE International Conference on*, vol. 4, April 2002, p. IV.

- [76] J.-Y. Gau, "Analysis of Low Probability of Intercept (LPI) Radar Signals using the Wigner Distribution," Master's thesis, Naval Postgraduate School, Monterey, California, September 2002.
- [77] V. C. Chen and H. Ling, *Time-Frequency Transforms for Radar Imaging and Signal Analysis*. Artech House Print on Demand, January 2002.
- [78] L. Cirillo, A. Zoubir, and M. Amin, "Estimation of FM Parameters Using a Time-Frequency Hough Transform," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 3, May 2006, p. III.
- [79] C. Ioana, A. Jarrot, A. Quinquis, S. Stankovic, and L. Stankovic, "Analysis of time-frequency transient components using phase chirping operator," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 3, May 2006, p. III.
- [80] L. Cirillo, A. Zoubir, and M. Amin, "Parameter Estimation for Locally Linear FM Signals Using a Time-Frequency Hough Transform," *Signal Processing, IEEE Transactions on*, vol. 56, no. 9, pp. 4162–4175, September 2008.
- [81] J. Lunden and V. Koivunen, "Automatic Radar Waveform Recognition," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 1, no. 1, pp. 124–136, June 2007.
- [82] P. Kootsookos, B. Lovell, and B. Boashash, "A unified approach to the STFT, TFDs, and instantaneous frequency," *Signal Processing, IEEE Transactions on*, vol. 40, no. 8, pp. 1971–1982, August 1992.
- [83] S. Kadambe and R. Orr, "Comparative study of the cross-term deleted wigner and cross biorthogonal representations," in *Signals, Systems Computers, 1997. Conference Record of the Thirty-First Asilomar Conference on*, vol. 2, November 1997, pp. 1484–1488 vol.2.
- [84] G. Upperman, T. Upperman, D. Fouts, and P. Pace, "Efficient time-frequency and bi-frequency signal processing on a reconfigurable computer," in *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, October 2008, pp. 176–180.
- [85] T. Upperman, "ELINT Signal Processing using Choi-Williams Distribution on Reconfigurable Computers for Detection and Classification of LPI Emitters," Master's thesis, Naval Postgraduate School, Monterey, California, March 2008.
- [86] A. Sonnenschein and P. Fishman, "Radiometric detection of spread-spectrum signals in noise of uncertain power," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 28, no. 3, pp. 654–660, July 1992.

- [87] W. A. Gardner, "Structural characterization of locally optimum detectors in terms of locally optimum estimators and correlators," *Information Theory, IEEE Transactions on*, vol. 28, no. 6, pp. 924–932, November 1982.
- [88] —, "Non-parametric signal detection and identification of modulation type for signals hidden in noise: An application of the new theory of cyclic spectral analysis," in *1983 - GRETSI - Actes de Colloques*. GRETSI, Groupe d'Etudes du Traitement du Signal et des Images, 1983, pp. 119–124. [Online]. Available: <http://hdl.handle.net/2042/10645>
- [89] W. Gardner, "Measurement of spectral correlation," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 34, no. 5, pp. 1111 – 1123, October 1986.
- [90] W. A. Gardner, "Signal interception: a unifying theoretical framework for feature detection," *Communications, IEEE Transactions on*, vol. 36, no. 8, pp. 897–906, August 1988.
- [91] —, "Exploitation of spectral correlation in cyclostationary signals," in *Spectrum Estimation and Modeling, 1988., Fourth Annual ASSP Workshop on*, August 1988, pp. 1–6.
- [92] W. Gardner and C. Spooner, "Cyclic spectral analysis for signal detection and modulation recognition," in *Military Communications Conference, 1988. MIL-COM 88, Conference record. '21st Century Military Communications - What's Possible?'. 1988 IEEE*, October 1988, pp. 419 –424 vol.2.
- [93] W. A. Gardner, *Introduction To Random Processes*, 2nd ed. Mcgraw-Hill (Tx), 1990.
- [94] —, "Exploitation of spectral redundancy in cyclostationary signals," *Signal Processing Magazine, IEEE*, vol. 8, no. 2, pp. 14–36, April 1991.
- [95] W. Gardner, "Exploitation of spectral redundancy in cyclostationary signals," *Signal Processing Magazine, IEEE*, vol. 8, no. 2, pp. 14 –36, April 1991.
- [96] W. Gardner and C. Spooner, "Signal interception: performance advantages of cyclic-feature detectors," *Communications, IEEE Transactions on*, vol. 40, no. 1, pp. 149 –159, January 1992.
- [97] W. A. Gardner and C. Spooner, "Detection and source location of weak cyclostationary signals: simplifications of the maximum-likelihood receiver," *Communications, IEEE Transactions on*, vol. 41, no. 6, pp. 905–916, June 1993.
- [98] A. Dandawate and G. Giannakis, "Statistical tests for presence of cyclostationarity," *Signal Processing, IEEE Transactions on*, vol. 42, no. 9, pp. 2355 –2369, September 1994.

- [99] S. Enserink and D. Cochran, "On detection of cyclostationary signals," in *Acoustics, Speech, and Signal Processing, 1995. ICASSP-95., 1995 International Conference on*, vol. 3, May 1995, pp. 2004 –2007 vol.3.
- [100] E. April, "The Advantage of Cyclic Spectral Analysis," Defence Research Establishment, Ottawa, Tech. Rep., October 1991.
- [101] A. Lima, Jr., "Analysis of Low Probability of Intercept (LPI) Radar Signals using Cyclostationary Processing," Master's thesis, Naval Postgraduate School, Monterey, California, September 2002.
- [102] O. Yeste-Ojeda and J. Grajal, "Cyclostationarity-based signal separation in interceptors based on a single sensor," in *Radar Conference, 2008. RADAR '08. IEEE*, May 2008, pp. 1 –6.
- [103] G. Upperman, "Implementation of a Cyclostationary Spectral Analysis Algorithm on an SRC Reconfigurable Computer for Real-time Signal Processing," Master's thesis, Naval Postgraduate School, Monterey, California, March 2008.
- [104] P. Ciblat, P. Loubaton, E. Serpedin, and G. Giannakis, "Asymptotic analysis of blind cyclic correlation-based symbol-rate estimators," *Information Theory, IEEE Transactions on*, vol. 48, no. 7, pp. 1922 – 1934, July 2002.
- [105] G. Gelli, L. Izzo, and L. Paura, "Cyclostationarity-based signal detection and source location in non-Gaussian noise," *Communications, IEEE Transactions on*, vol. 44, no. 3, pp. 368 –376, March 1996.
- [106] A. Tkachenko, A. Cabric, and R. Brodersen, "Cyclostationary Feature Detector Experiments Using Reconfigurable BEE2," in *New Frontiers in Dynamic Spectrum Access Networks, 2007. DySPAN 2007. 2nd IEEE International Symposium on*, April 2007, pp. 216 –219.
- [107] P. Marchand, J.-L. Lacoume, and C. Le Martret, "Multiple hypothesis modulation classification based on cyclic cumulants of different orders," in *Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference on*, vol. 4, May 1998, pp. 2157 –2160 vol.4.
- [108] J. Antoni, "Cyclic spectral analysis in practice," *Mechanical Systems and Signal Processing*, vol. 21, no. 2, pp. 597 – 630, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/B6WN1-4M7VB33-1/2/2416eaa2be5b186be5c7823d06afe3f2>
- [109] G. Yeung and W. Gardner, "Search-efficient methods of detection of cyclostationary signals," *Signal Processing, IEEE Transactions on*, vol. 44, no. 5, pp. 1214–1223, May 1996.

- [110] R. S. Roberts, W. A. Brown, and H. H. Loomis, Jr., "Computationally efficient algorithms for cyclic spectral analysis," *IEEE Signal Processing Magazine*, vol. 8, pp. 38–49, April 1991.
- [111] I. Brown, W.A. and J. Loomis, H.H., "Digital implementations of spectral correlation analyzers," *Signal Processing, IEEE Transactions on*, vol. 41, no. 2, pp. 703–720, February 1993.
- [112] R. S. Roberts, W. A. Brown, and J. Loomis, H. H., *A review of digital spectral correlation analysis: Theory and implementation*, ser. Cyclostationarity in Communications and Signal Processing. IEEE Press, 1994.
- [113] J. A. C. Lee and . J. Munson, D. C., "Effectiveness of spatially-variant apodization," in *ICIP '95: Proceedings of the 1995 International Conference on Image Processing (Vol. 1)-Volume 1*. Washington, DC, USA: IEEE Computer Society, 1995, p. 147.
- [114] H. Stankwitz, R. Dallaire, and J. Fienup, "Nonlinear apodization for sidelobe control in SAR imagery," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 31, no. 1, pp. 267–279, January 1995.
- [115] D. Pastina, F. Colone, and P. Lombardo, "Effect of Apodization on SAR Image Understanding," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 45, no. 11, pp. 3533–3551, November 2007.
- [116] G. Thomas, "Time-frequency analysis using sidelobe apodization," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, F. T. Luk, Ed., vol. 4474, November 2001, pp. 25–34.
- [117] G. Thomas and N. Gadhok, "Sidelobe apodization in Fourier imaging," *Signals, Systems and Computers, 2001. Conference Record of the Thirty-Fifth Asilomar Conference on*, vol. 2, pp. 1369–1373 vol.2, 2001.
- [118] R. J. Dallaire and H. C. Stankwitz, "Spatially variant apodization," United States Patent 5,349,359, September 1994.
- [119] R. J. Iannuzzelli, "Adaptive Windowing of FFT Data for Increased Resolution and Sidelobe Rejection," United States Patent US 6,298,363 B1, October 2001.
- [120] G. Thomas, B. Flores, and J. Sok-Son, "SAR sidelobe apodization using the Kaiser window," *Image Processing, 2000. Proceedings. 2000 International Conference on*, vol. 1, pp. 709–712 vol.1, 2000.
- [121] G. Thomas, J. LoVetri, W. Chamma, S. Kashyap, and A. Louie, "Sidelobe apodization for high resolution of scattering centres in ISAR images," in *Antennas and Propagation Society International Symposium, 2001. IEEE*, vol. 2, 2001, pp. 232–235 vol.2.

- [122] X. Xu and R. M. Narayanan, “Enhanced Resolution in SAR/ISAR Imaging Using Iterative Sidelobe Apodization,” *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 537–547, April 2005.
- [123] H. C. Stankwitz and M. R. Kosek, “Super-resolution for SAR/ISAR RCS measurements using spatially variant apodization super-SVA,” *Proceedings 17th Meeting Symp. Antenna Measurement Techniques Association*, pp. 251–256, November 1995.
- [124] M. Abe and J. O. Smith, “Design Criteria for the Quadratically Interpolated FFT Method I: Bias Due to Interpolation,” Stanford University, Department of Music, Technical Report STAN–M–117, 2004.
- [125] S. Haykin, D. Thomson, and J. Reed, “Spectrum sensing for cognitive radio,” *Proceedings of the IEEE*, vol. 97, no. 5, pp. 849–877, May 2009.
- [126] S. Haykin, “The multitaper method for accurate spectrum sensing in cognitive radio environments,” in *Signals, Systems and Computers, 2007. ACSSC 2007. Conference Record of the Forty-First Asilomar Conference on*, November 2007, pp. 436–439.
- [127] B. Farhang-Boroujeny, “Filter Bank Spectrum Sensing for Cognitive Radios,” *Signal Processing, IEEE Transactions on*, vol. 56, no. 5, pp. 1801–1811, May 2008.
- [128] D. Thomson, “Spectrum estimation and harmonic analysis,” *Proceedings of the IEEE*, vol. 70, no. 9, pp. 1055 – 1096, September 1982.
- [129] K. Riedel and A. Sidorenko, “Minimum bias multiple taper spectral estimation,” *Signal Processing, IEEE Transactions on*, vol. 43, no. 1, pp. 188 –195, January 1995.
- [130] E. Sloane, “Comparison of linearly and quadratically modified spectral estimates of Gaussian signals,” *Audio and Electroacoustics, IEEE Transactions on*, vol. 17, no. 2, pp. 133 – 137, June 1969.
- [131] D. B. Percival and A. T. Walden, *Spectral analysis for physical applications: multitaper and conventional univariate techniques*. Cambridge University Press, 1993.
- [132] D. Slepian, “Prolate spheroidal wave functions, Fourier analysis, and uncertainty. V - The discrete case,” *AT&T Technical Journal*, vol. 57, pp. 1371–1430, June 1978.
- [133] D. B. Percival and A. T. Walden, *Wavelet Methods for Time Series Analysis (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press, 2001.

- [134] V. K. Anandan, C. J. Pan, T. Rajalakshmi, and G. Ramachandra Reddy, "Multitaper spectral analysis of atmospheric radar signals," *Annales Geophysicae*, vol. 22, no. 11, pp. 3995–4003, 2004. [Online]. Available: <http://www.ann-geophys.net/22/3995/2004/>
- [135] G. A. Prieto, R. L. Parker, and F. L. Vernon, III, "A Fortran 90 library for multitaper spectrum analysis," *Computers and Geosciences*, vol. 35, pp. 1701–1710, August 2009.
- [136] D. Thomson, L. Lanzerotti, F. Vernon, M. Lessard, and L. Smith, "Solar Modal Structure of the Engineering Environment," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 1085–1132, May 2007.
- [137] A. Walden, "Multitaper estimation of the innovation variance of a stationary time series," *Signal Processing, IEEE Transactions on*, vol. 43, no. 1, pp. 181–187, January 1995.
- [138] G. Nicolescu and P. J. Mosterman, *Model-Based Design for Embedded Systems*, 1st ed., P. J. Mosterman, Ed. CRC Press, November 2009.
- [139] Xilinx, *System Generator for DSP Performing Hardware-in-the-loop With the Spartan-3E Starter Kit*, 1st ed., Xilinx Inc., Xilinx Headquarters: Xilinx, Inc. 2100 Logic Drive San Jose, CA 95124-3400, June 2006.
- [140] C. van Loan, *Computational Frameworks for the Fast Fourier Transform*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1992.
- [141] R. Andraka, "A survey of CORDIC algorithms for FPGA computers," in *FPGA '98: Proceedings of the 1998 ACM/SIGDA sixth international symposium on Field programmable gate arrays*. New York, NY, USA: ACM, 1998, pp. 191–200.
- [142] D. Matovic and C. Tropea, "Spectral peak interpolation with application to LDA signal processing," *Measurement Science and Technology*, vol. 2, no. 11, pp. 1100–1106, 1991. [Online]. Available: <http://stacks.iop.org/0957-0233/2/1100>
- [143] T. Lang and E. Antelo, "CORDIC-Based Computation of arccos and $\sqrt{1-t^2}$," *J. VLSI Signal Process. Syst.*, vol. 25, no. 1, pp. 19–38, 2000.
- [144] Xilinx, *Virtex-5 Family Overview*, 5th ed., Xilinx Inc., Xilinx Headquarters: Xilinx, Inc. 2100 Logic Drive San Jose, CA 95124-3400, February 2009.
- [145] —, *Xilinx LogiCORE Divider Generator v3.0*, 3rd ed., Xilinx Inc., Xilinx Headquarters: Xilinx, Inc. 2100 Logic Drive San Jose, CA 95124-3400, June 2009.
- [146] *Digital Television Transmitters pre-switchover network*, 4th ed., Ofcom, Riverside House, 2a Southwark Bridge Road, London SE1 9HA, February 2010.
- [147] M. I. Skolnik, *Introduction to radar systems*. McGraw-Hill, New York,, 1962.