

**MODELING AND MEASUREMENT OF TORQUED
PRECESSION IN RADIO PULSARS**

A thesis submitted in fulfilment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

of

RHODES UNIVERSITY

by

ADRIAN JOHN TIPLADY

October 2004

Abstract

The long term isolated pulsar monitoring program, which commenced in 1984 using the 26 m radio telescope at the Hartebeeshoek Radio Astronomy Observatory (HartRAO), has produced high resolution timing residual data over long timespans. This has enabled the analysis of observed spin down behaviour for 27 braking pulsars, most of which have dataspans longer than 14 years. The phenomenology of observed timing residuals of certain pulsars can be explained by pseudo periodic effects such as precession. Analytic and numerical models are developed to study the kinematic and dynamic behaviour of isolated but torqued precessing pulsars. The predicted timing residual behaviour of the models is characterised, and confronted with timing data from selected pulsars. Cyclic variations in the observed timing residuals of PSR B1642-03, PSR B1323-58 and PSR B1557-50 are fitted with a torqued precession model.

The phenomenology of the observed timing behaviour of these pulsars can be explained by the precession models, but precise model fitting was not possible. This is not surprising given that the complexity of the pulsar systems is not completely described by the model.

The extension of the pulsar monitoring program at HartRAO is used as motivation for the design and development of a new low cost, multi-purpose digital pulsar receiver. The instrument is implemented using a hybrid filterbank architecture, consisting of an analogue frontend and digital backend, to perform incoherent dedispersion. The design of a polyphase filtering system, which will consolidate multiple processing units into a single filtering solution, is discussed for future implementation.

Acknowledgements

The success of the pulsar timing program at HartRAO is due to the dedicated efforts of Dr. Claire Flanagan and Dr. George Nicholson. I would like to thank them for their support and encouragement, and for use of the HartRAO pulsar timing data. I thank the staff at HartRAO who have provided invaluable technical assistance.

For his guidance, encouragement and mutual respect, I thank my supervisor and mentor, Prof. Justin Jonas. The profound and selfless support of Prof. Fabio Frescura has, if anything, opened my mind to a way of scientific thought that no book could ever teach, for which I am grateful.

To the staff and students of the Rhodes University department of Physics and Electronics, your support and friendship at all times will always be appreciated. To others who helped me to begin, to continue, and to complete this work, I thank you.

The financial assistance from the Henderson Post-Graduate Scholarship towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to Rhodes University or the donor. I also would like to acknowledge financial support from HartRAO.

I dedicate this to my family who supported every endeavour of mine in the belief that someday, it would lead to a greater good.

- *'What're quantum mechanics?'*
- *'I don't know. People who repair quantum, I suppose.'*

(Terry Pratchett, Eric)

Contents

| | |
|---|------|
| Abstract | i |
| Acknowledgements | ii |
| List of Figures | vii |
| List of Tables | xiii |
| Chapter 1. Introduction | 1 |
| 1.1. The Nature of Torqued Pulsars | 3 |
| 1.2. Pulsar Observation | 4 |
| 1.3. This Thesis | 4 |
| Chapter 2. Braking Pulsars | 6 |
| 2.1. Introduction | 6 |
| 2.2. Braking Model | 6 |
| 2.3. Distinction between n and n_{obs} | 8 |
| 2.4. Time Dependence in the Braking Law | 9 |
| 2.5. The Generalised Braking Index | 12 |
| 2.6. Timing noise in the Measurement of n_{obs} | 14 |
| Chapter 3. Free Precession | 19 |
| 3.1. Introduction | 19 |
| 3.2. The Rigid Body Model | 19 |
| 3.3. Equations of Motion | 20 |
| 3.4. Axisymmetric Pulsar | 20 |
| 3.5. Asymmetric Pulsar | 21 |
| 3.6. Free Precession in Pulsars? | 24 |
| 3.7. Summary | 25 |
| Chapter 4. Torqued Precession | 26 |
| 4.1. Introduction | 26 |
| 4.2. Dipole Torque | 26 |
| 4.3. Spherical Pulsar | 28 |
| 4.4. Axisymmetric Pulsar | 30 |
| 4.5. Asymmetric Pulsar | 36 |
| 4.6. Torqued Precession in Pulsars? | 49 |

| | |
|--|-----|
| 4.7. Summary | 50 |
| Chapter 5. Timing Residuals | 54 |
| 5.1. Introduction | 54 |
| 5.2. Calculation of Phase Residuals: Untorqued Pulsars | 55 |
| 5.3. Calculation of Phase Residuals: Torqued Pulsars | 56 |
| 5.4. Calculation of Observed Phase Residuals | 68 |
| 5.5. Candidate Precessing Pulsars in the HartRAO Catalogue | 68 |
| 5.6. Precession Fits | 69 |
| 5.7. Conclusion | 87 |
| Chapter 6. Pulsar Timing | 89 |
| 6.1. Introduction | 89 |
| 6.2. Dispersion, Scattering and De-Dispersion | 89 |
| 6.3. Critique of the Current Pulsar Processor | 91 |
| 6.4. Review of Existing Pulsar Processors | 94 |
| 6.5. Macroscopic Architecture Justification | 100 |
| 6.6. Parameters for a New Pulsar Processor | 102 |
| Chapter 7. Pulsar Processor Design and Implementation | 107 |
| 7.1. Introduction | 107 |
| 7.2. High Level Architecture | 107 |
| 7.3. Analogue Frontend Architecture | 108 |
| 7.4. Quadrature Mixing | 109 |
| 7.5. Analogue to Digital Conversion | 110 |
| 7.6. Analogue Frontend Assembly | 113 |
| 7.7. Digital Backend Architecture | 115 |
| 7.8. FPGA Design Architecture | 116 |
| 7.9. Input Interface | 117 |
| 7.10. Digital Complex Mixing | 118 |
| 7.11. FIR Filters | 119 |
| 7.12. Sideband Separation | 126 |
| 7.13. Detection and Post-Detection Integration | 126 |
| 7.14. Data Stacking | 127 |
| 7.15. Starter Module | 129 |
| 7.16. Digital Backend Assembly | 129 |
| 7.17. Pulsar Processor Assembly | 132 |
| 7.18. Operation | 132 |
| Chapter 8. Diagnostic Tests and Field Trials | 134 |
| 8.1. Statistical Analysis and Diagnostic Tests | 134 |
| 8.2. Bandpass Separation | 145 |

| | |
|---|-----|
| 8.3. Field Trials | 145 |
| 8.4. Conclusion | 147 |
| Chapter 9. Conclusion and Recommendations | 157 |
| 9.1. Braking Pulsars | 157 |
| 9.2. Remarks on Modeling Precession | 158 |
| 9.3. Future Pulsar Processor Implementation | 159 |
| Appendix A. Pulsar Mechanics | 164 |
| A.1. Determination of Euler Equations | 164 |
| A.2. Co-ordinate Transformations | 164 |
| Appendix B. Precession Model Fits | 166 |
| B.1. PSR B1642-03 | 166 |
| B.2. PSR B1323-58 | 167 |
| B.3. PSR B0736-40 | 173 |
| B.4. PSR B1557-50 | 175 |
| Appendix C. Signal Processing | 177 |
| C.1. Quadrature Mixing and Sideband Separation | 177 |
| C.2. FIR Filter Coefficients | 178 |
| Appendix D. Low Frequency Concept Design | 180 |
| D.1. Design Schematics | 181 |
| D.2. Flowcharts | 184 |
| Appendix E. Prototype Schematic Designs | 187 |
| E.1. Analogue Frontend Module | 188 |
| E.2. Digital Backend Module | 189 |
| Appendix F. Prototype Printed Circuit Board Silkscreens (See accompanying CD) | 190 |
| Appendix G. Microcontroller Algorithms | 191 |
| Appendix H. FPGA Design Schematics: FIR Filter System | 198 |
| H.1. FPGA Schematics | 198 |
| Appendix I. FPGA Design Schematics: Polyphase System | 202 |
| I.1. 8 Point Fast Fourier Transform | 202 |
| I.2. FFT Weights | 208 |
| Appendix J. Eurocard Standard Schematic Designs | 210 |
| J.1. Analogue Frontend Module | 211 |
| J.2. Digital Backend Module | 212 |
| J.3. Motherboard | 213 |
| J.4. Starter Module | 214 |

| | |
|---|-----|
| Appendix K. Eurocard Standard Printed Circuit Board Silkscreens (See accompanying CD) | 215 |
| Appendix L. Code and Setup Files (See accompanying CD) | 216 |
| Appendix. Bibliography | 217 |

List of Figures

| | | |
|--------|--|----|
| 1.0.1 | Pulsar discoveries in major surveys | 2 |
| 3.4.1 | Precessing space and body cone | 22 |
| 3.5.1 | Rotational stability of asymmetric body | 23 |
| 4.4.1 | ω_{\blacksquare} for axisymmetric pulsar, $\chi = \frac{\pi}{2}$ | 33 |
| 4.4.2 | β and α for axisymmetric pulsar, $\chi = \frac{\pi}{2}$ | 34 |
| 4.4.3 | ω_{\blacksquare} for axisymmetric pulsar, $0 < \chi < \frac{\pi}{2}$ | 35 |
| 4.4.4 | β and α for axisymmetric pulsar, $0 < \chi < \frac{\pi}{2}$ | 36 |
| 4.4.5 | ω_{\sim} for axisymmetric pulsar, $\omega(0) \approx \omega_3(0)$, $0 < \chi < \frac{\pi}{2}$ | 37 |
| 4.4.6 | ω_{\sim} for axisymmetric pulsar, $\omega(0) \approx \omega_1(0)$, $0 < \chi < \frac{\pi}{2}$ | 38 |
| 4.4.7 | ω_{\sim} for axisymmetric pulsar, variable vacuum radius r_0 , $\chi = \frac{\pi}{4}$ | 39 |
| 4.5.1 | ω_{\sim} for asymmetric pulsar, variable vacuum radius r_0 , $\chi = 0$ | 42 |
| 4.5.2 | ω_{\blacksquare} for asymmetric pulsar, $\omega(0) \approx \omega_2(0)$, $\chi = 0$ | 43 |
| 4.5.3 | β and α for asymmetric pulsar, $\omega(0) \approx \omega_3(0)$, $\chi = \frac{\pi}{2}$ | 44 |
| 4.5.4 | ω_{\blacksquare} for asymmetric pulsar, $\omega(0) \approx \omega_3(0)$, $0 < \chi < \frac{\pi}{2}$ | 46 |
| 4.5.5 | ω_{\sim} for asymmetric pulsar, $\omega(0) \approx \omega_3(0)$, $0 < \chi < \frac{\pi}{2}$ | 47 |
| 4.5.6 | ω_{\sim} for asymmetric pulsar, $\omega(0) \approx \omega_1(0)$, $0 < \chi < \frac{\pi}{2}$ | 48 |
| 4.5.7 | ω_{\blacksquare} for asymmetric pulsar, $\omega(0) \approx \omega_2(0)$, $0 < \chi < \frac{\pi}{2}$ | 49 |
| 4.5.8 | ω_{\sim} for asymmetric pulsar, variable vacuum radius r_0 , $\omega(0) \approx \omega_2(0)$, $\chi = 0$ | 50 |
| 4.5.9 | ω_{\sim} for asymmetric pulsar, variable vacuum radius r_0 , $\omega(0) \approx \omega_2(0)$, $\chi = \frac{\pi}{4}$ | 51 |
| 4.5.10 | ω_{\sim} for asymmetric pulsar, variable vacuum radius r_0 , $\omega(0) \approx \omega_2(0)$, $\chi = \frac{\pi}{2}$ | 52 |
| 5.2.1 | $\Delta\Phi_{\text{analytical}} - \Delta\Phi_{\text{numerical}}$ for an untorqued asymmetric pulsar | 55 |
| 5.2.2 | Timing residuals of an untorqued axisymmetric pulsar, $\theta(0) = 0.1405$ | 57 |
| 5.2.3 | Timing residuals of an untorqued axisymmetric pulsar, $\theta(0) = 1.29$ | 58 |
| 5.3.1 | Phase residuals and ΔP of a torqued asymmetric pulsar, $\theta(0) = 0.1405$ | 59 |
| 5.3.2 | $\frac{d(\Delta P)}{dt}$ and $\theta(t)$ of a torqued asymmetric pulsar, $\theta(0) = 0.1405$ | 60 |
| 5.3.3 | Phase residuals and ΔP of a torqued asymmetric pulsar, $\theta(0) = 1.4716$ | 61 |
| 5.3.4 | $\frac{d(\Delta P)}{dt}$ and $\theta(t)$ of a torqued asymmetric pulsar, $\theta(0) = 1.4716$ | 62 |

| | | |
|--------|---|---------------|
| 5.3.5 | Phase residuals and ΔP of a torqued asymmetric pulsar, $\theta(0) = 0.5611$ | 63 |
| 5.3.6 | $\frac{d(\Delta P)}{dt}$ and $\theta(t)$ of a torqued asymmetric pulsar, $\theta(0) = 0.5611$ | 64 |
| 5.3.7 | Phase residuals and ΔP of a torqued asymmetric pulsar, $\theta(0) = 1.2152$ | 65 |
| 5.3.8 | $\frac{d(\Delta P)}{dt}$ and $\theta(t)$ of a torqued asymmetric pulsar, $\theta(0) = 1.2152$ | 66 |
| 5.3.9 | Phase residuals of a torqued asymmetric pulsar, $\theta(0) = 0.4405$ | 67 |
| 5.4.1 | Phase residuals of a torqued asymmetric pulsar, $\theta(0) = 0.1405$ | 69 |
| 5.5.1 | Phase residuals of PSRs B0736-40, B0740-28, B0835-41 and B0959-54 | 70 |
| 5.5.2 | Phase residuals of PSRs B1323-58, B1323-62, B1358-63 and B1557-50 | 71 |
| 5.5.3 | Phase residuals of PSRs B1642-03 and B1749-28 | 72 |
| 5.6.1 | Manual fit of PSR B1642-03 timing residuals | 73 |
| 5.6.2 | Manual fit of PSR B1642-03 timing residual spectrums | 74 |
| 5.6.3 | Chi squared surfaces for fitting of PSR B1642-03 timing residuals | 76 |
| 5.6.4 | Best fit of PSR B1642-03 timing residuals | 77 |
| 5.6.5 | Best fit of PSR B1642-03 timing residual spectrums | 78 |
| 5.6.6 | Manual fit of PSR B1323-58 timing residuals | 80 |
| 5.6.7 | Manual fit of PSR B1323-58 timing residual spectrums | 81 |
| 5.6.8 | Manual fit of PSR B0736-40 timing residuals | 83 |
| 5.6.9 | Manual fit of PSR B0736-40 timing residual spectrums | 84 |
| 5.6.10 | Manual fit of PSR B1557-50 timing residuals | 85 |
| 5.6.11 | Manual fit of PSR B1557-50 timing residual spectrums | 86 |
| 6.3.1 | P-DM pulsar distribution: pulse smearing of HartRAO pulsar processor | 91 |
| 6.3.2 | Minimum detectable flux density of HartRAO pulsar processor | 92 |
| 6.3.3 | Detectable pulsar subpopulation of HartRAO pulsar processor | 93 |
| 6.4.1 | Medicina Pulsar Observing System. | 94 |
| 6.4.2 | Coherent Online Baseband Receiver for Astronomy (COBRA) | 96 |
| 6.4.3 | Princeton Mark IV Pulsar Processor | 97 |
| 6.4.4 | Coherent Arecibo-Berkeley Pulsar Processor (ABPP) | 98 |
| 6.4.5 | Coherent Arecibo-Berkeley Pulsar Processor (ABPP): Digital Filter Board | 98 |
| 6.4.6 | Coherent Arecibo-Berkeley Pulsar Processor (ABPP): Dedisperser Board | 99 |
| 6.6.1 | P-DM distribution: pulse smearing of new pulsar processor | 102 |
| 6.6.2 | Detectability of catalogued pulsar population | 103 |
| 6.6.3 | Detectability of simulated pulsar population | 104 |
| 6.6.4 | Minimum detectable flux density for new pulsar processor | 105 |
| 6.6.5 | Detectable pulsar subpopulation of new pulsar processor | 106 |

| | | |
|--------|--|-----|
| 7.2.1 | Macroscopic design structure of new pulsar processor | 108 |
| 7.3.1 | Analogue frontend signal flow | 109 |
| 7.5.1 | Power distribution of analogue, quantised and error signals | 111 |
| 7.5.2 | SNR of quantised signal | 112 |
| 7.5.3 | Linearity of digital system | 113 |
| 7.6.1 | Schematic design for analogue frontend module | 114 |
| 7.7.1 | Digital backend signal flow | 115 |
| 7.8.1 | FPGA signal flow | 116 |
| 7.8.2 | Top heirarchial structure for FPGA design | 117 |
| 7.10.1 | Complex mixer z plane rotation | 118 |
| 7.11.1 | FIR filter coefficients | 122 |
| 7.11.2 | FIR filter passbands | 123 |
| 7.11.3 | Floating point and quantised filter passbands | 124 |
| 7.11.4 | Unsigned digit-serial multiplier | 124 |
| 7.11.5 | LUT implemented adder tree | 126 |
| 7.14.1 | Simplified flowchart for upper sideband data capture | 128 |
| 7.15.1 | Simplified flowchart for starter module | 130 |
| 7.16.1 | Schematic design for digital backend module | 131 |
| 8.1.1 | Probability mass distributions of input signal | 135 |
| 8.1.2 | Probability mass distributions of mixer input signal | 136 |
| 8.1.3 | Cumulative distributions of mixer input signal | 137 |
| 8.1.4 | Probability mass distributions of mixer output signal | 137 |
| 8.1.5 | Probability mass distributions of real filter output signal | 138 |
| 8.1.6 | Smoothed 8×4096 point DTFT of real filter output | 139 |
| 8.1.7 | Probability mass distributions of imaginary filter output signal | 139 |
| 8.1.8 | Smoothed 8×4096 point DTFT of imaginary filter output | 140 |
| 8.1.9 | Probability mass distributions of adder output signal | 141 |
| 8.1.10 | Smoothed 8×4096 point DTFT of adder output | 141 |
| 8.1.11 | Probability mass distributions of subtracter output signal | 142 |
| 8.1.12 | Smoothed 8×4096 point DTFT of subtracter output | 142 |
| 8.1.13 | Probability mass distributions of squarer output signal | 144 |
| 8.1.14 | Probability mass distributions of accumulated squared signal | 144 |
| 8.2.1 | Complex mixer tests: 0 degrees | 146 |
| 8.2.2 | Composite pulsar processor passband | 147 |

| | | |
|--------|---|-----|
| 8.3.1 | PSR B0736-40 pulse profile integration | 148 |
| 8.3.2 | PSR B0740-28 pulse profile integration | 149 |
| 8.3.3 | PSR B0833-45 pulse profile integration | 149 |
| 8.3.4 | PSR B1054-62 pulse profile integration | 150 |
| 8.3.5 | PSR B1240-64 pulse profile integration | 150 |
| 8.3.6 | PSR B1323-58 pulse profile integration | 151 |
| 8.3.7 | PSR B1557-50 pulse profile integration | 151 |
| 8.3.8 | PSR B1642-03 pulse profile integration | 152 |
| 8.3.9 | PSR B1641-45 pulse profile integration | 152 |
| 8.3.10 | PSR B1929+10 pulse profile integration | 153 |
| 8.3.11 | PSR B1939+16 pulse profile integration | 153 |
| 8.3.12 | Noise performance of pulsar processor | 154 |
| 8.3.13 | Individual channel bandwidth tests | 155 |
| 9.3.1 | FPGA schematic design of top heirarchy for polyphase filtering system | 160 |
| 9.3.2 | FPGA schematic design of polyphase filter | 161 |
| 9.3.3 | FPGA schematic design of 4 point FFT | 162 |
| B.1.1 | Calculated period derivative of PSR B1642-03 | 166 |
| B.1.2 | Calculated $\theta(t)$ of PSR B1642-03 | 166 |
| B.1.3 | Calculated $\varpi(t)$ of PSR B1642-03 | 167 |
| B.2.1 | Calculated period derivative of PSR B1323-58 | 167 |
| B.2.2 | Calculated $\theta(t)$ of PSR B1323-58 | 168 |
| B.2.3 | Calculated $\varpi(t)$ of PSR B1323-58 | 168 |
| B.2.4 | Chi squared surface for fitting of PSR B1323-58, constant [B-Pfp] | 169 |
| B.2.5 | Chi squared surface for fitting of PSR B1323-58, constant [$x - \theta(0)$] | 169 |
| B.2.6 | Calculated phase residual for best fit of PSR B1323-58 | 170 |
| B.2.7 | Calculated ΔP for best fit of PSR B1323-58 timing residuals | 170 |
| B.2.8 | Calculated $\frac{d(\Delta P)}{dt}$ for best fit of PSR B1323-58 timing residuals | 171 |
| B.2.9 | Spectrum of phase residual calculated for best fit of PSR B1323-58 | 171 |
| B.2.10 | Spectrum of ΔP calculated for best fit of PSR B1323-58 timing residuals | 172 |
| B.2.11 | Spectrum of $\frac{d(\Delta P)}{dt}$ calculated for best fit of PSR B1323-58 timing residuals | 172 |
| B.3.1 | Calculated period derivative of PSR B0736-40 | 173 |
| B.3.2 | Calculated $\theta(t)$ of PSR B0736-40 | 173 |
| B.3.3 | Calculated $\varpi(t)$ of PSR B0736-40 | 174 |

| | | |
|-------|--|-----|
| B.3.4 | Chi squared surface for fitting of PSR B0736-40, constant [B-Pfp] | 174 |
| B.3.5 | Chi squared surface for fitting of PSR B0736-40, constant [$\chi - \theta(0)$] | 175 |
| B.4.1 | Calculated period derivative of PSR B1557-50 | 175 |
| B.4.2 | Calculated $\theta(t)$ of PSR B1557-50 | 176 |
| B.4.3 | Calculated $\omega(t)$ of PSR B1557-50 | 176 |
| D.1.1 | LFC Schematic design (Sheet 1 of 3) | 181 |
| D.1.2 | LFC Schematic design (Sheet 2 of 3) | 182 |
| D.1.3 | LFC Schematic design (Sheet 3 of 3) | 183 |
| D.2.1 | LFC Flowchart: Real filter | 184 |
| D.2.2 | LFC Flowchart: Imaginary filter | 185 |
| D.2.3 | LFC Flowchart: Sideband separation | 186 |
| E.1.1 | Prototype analogue frontend schematic design | 188 |
| E.2.1 | Prototype digital backend schematic design | 189 |
| G.1.1 | Flow chart for USB Data capture: Sheet 1 of 3 | 192 |
| G.1.2 | Flow chart for USB Data capture: Sheet 2 of 3 | 193 |
| G.1.3 | Flow chart for USB Data capture: Sheet 3 of 3 | 194 |
| G.1.4 | Flow chart for LSB Data capture: Sheet 1 of 3 | 195 |
| G.1.5 | Flow chart for LSB Data capture: Sheet 2 of 3 | 196 |
| G.1.6 | Flow chart for LSB Data capture: Sheet 3 of 3 | 197 |
| H.1.1 | FPGA schematic of input interface | 198 |
| H.1.2 | FPGA schematic of digital complex mixer | 199 |
| H.1.3 | FPGA schematic of accumulator | 200 |
| H.1.4 | FPGA schematic of output interface | 201 |
| I.1.1 | FPGA schematic of 8 point FFT | 202 |
| I.1.2 | FPGA schematic of 2 point FFT | 203 |
| I.1.3 | FPGA schematic of 4 point FFT weights: Set 1 of 4 | 204 |
| I.1.4 | FPGA schematic of 4 point FFT weights: Set 2 of 4 | 205 |
| I.1.5 | FPGA schematic of 4 point FFT weights: Set 3 of 4 | 205 |
| I.1.6 | FPGA schematic of 4 point FFT weights: Set 4 of 4 | 206 |
| I.1.7 | FPGA schematic of 2 point FFT weights: Set 1 of 2 | 207 |
| I.1.8 | FPGA schematic of 2 point FFT weights: Set 2 of 2 | 208 |
| J.1.1 | Schematic for Analogue frontend: Eurocard standard | 211 |

LIST OF FIGURES

xii

| | | |
|-------|--|-----|
| J.2.1 | Schematic for Digital backend: Eurocard standard | 212 |
| J.3.1 | Schematic for Motherboard: Eurocard standard | 213 |
| J.4.1 | Schematic for Starter: Eurocard standard | 214 |

List of Tables

| | | |
|--------|---|-----|
| 2.6.1 | Calculation of ξ for 27 HartRAO pulsars | 17 |
| 5.6.1 | PSR B1642-03 fitting parameters | 75 |
| 5.6.2 | PSR B1642-03 derived pulsar parameters | 75 |
| 5.6.3 | PSR B1642-03 best fit fitting parameters | 75 |
| 5.6.4 | PSR B1642-03 best fit derived pulsar parameters | 79 |
| 5.6.5 | PSR B1323-58 fitting parameters | 79 |
| 5.6.6 | PSR B1323-58 derived pulsar parameters | 79 |
| 5.6.7 | PSR B0736-40 best fit fitting parameters | 82 |
| 5.6.8 | PSR B0736-40 best fit derived pulsar parameters | 82 |
| 5.6.9 | PSR B1557-50 fitting parameters | 82 |
| 5.6.10 | PSR B1557-50 derived pulsar parameters | 87 |
| 7.5.1 | $\frac{\text{RMS}}{\text{FSD}}$ and maximum SNR values for different values of b | 112 |
| 9.1.1 | Fitted pulsar parameters | 158 |
| 9.3.2 | 4 point and 2 FFT phasor weights | 162 |
| C.2.1 | Quantised filter coefficients for real and imaginary filters | 179 |
| I.2.1 | Weights for 4 point and 2 point FFT | 209 |

CHAPTER 1

Introduction

The discovery of a rapidly pulsating radio source in 1967 [79], and subsequent confirmation of the rotating neutron star model [69] proposed by Baade and Zwicky in 1934, was the first step in a field that has provided such diverse applications as solid state physics, particle physics, magneto-electrodynamics and general relativity. Since that initial discovery, the number of catalogued pulsars has increased to approximately 1500 at present [81], and the field has produced three Nobel laureates in Anthony Hewish (1974), Russell Hulse and Joseph Taylor (1993).

During the past decade, pulsar science has become one of the major research fields in modern astronomy and is a science driver for the construction of future, large scale telescopes such as the Square Kilometer Array [66]. Figure 1.0.1 illustrates the growth in pulsar related publications, complemented by the rapid growth in the number of catalogued pulsars.

These highly compact, ‘point masses’ have become ideal laboratories for experiments in a variety of fields of research, with applications found not only in the radio pulsar system itself, but the surrounding interstellar medium [64] and large scale galactic magnetic field [74].

The observation of a sudden discontinuity in the rotation period of the Vela pulsar in 1969 [148, 149], a phenomenon known as a glitch (Vela has since undergone nine such glitches [58, 154]), and subsequent recovery was interpreted as a large scale exchange of angular momentum within the neutron star [20, 21]. Glitches have since been used as probes into the neutron star interior [106]. The possibility of the existence of exotic matter within the neutron star core, such as pion condensates, raises fundamental questions in the field of particle physics.

The large variety of observed emission mechanism signatures has resulted in the classification of a number of classes of neutron stars, namely:

- Radio Pulsars (PSR): The most populous class of pulsar, their pulsed emission is generally observed in the radio frequency band. However, a number of pulsars have shown to emit in the optical, infrared, X-ray and γ -ray energy spectrum. Powered by their own rotational energy, two distinctive evolutionary paths divide these pulsars into two subclasses: normal and millisecond pulsars. Normal pulsars tend to have rotation periods of between 0.03 and 8.5 seconds, and spin down parameter $\dot{P} \geq 10^{-17} \text{ s.s}^{-1}$. However, normal pulsars in globular clusters appear to be spinning up due to their acceleration in the gravitational potential of the cluster [176, 23]. Millisecond pulsars are identified as a separate and much older population with distinct evolutionary histories. Having gone through a period of ‘spin up’, their spin periods are less than a nominal value of 20 ms [41]. Often termed as ‘recycled’ pulsars, mass and angular momentum transfer from a binary companion is seen to be the cause of spin up. They are characterised

by very low surface magnetic fields, $\mathbf{B} \approx 10^8 - 10^9$ G, and subsequent small spin down rates of $\dot{P} \leq 10^{-19} \text{ s.s}^{-1}$.

- γ -Ray Bursters (GRB): First detected in satellite measurements between 1969 and 1972 [87], they are characterised by short duration γ -ray bursts of ≈ 10 s. Recent observations by BATSE (Burst and Transition Experiment) and the Beppo-SAX satellite provides good evidence that the bursts originate at extra-galactic distances, and are associated with star forming regions, or supernovae. To date, over 3000 GRBs have been observed, distributed evenly across the sky.
- Soft γ -Ray Repeaters (SGR): This small class, emitting multiple γ -ray bursts, is distinct from the GRBs with only five observed, all radio quiet [177]. Strong evidence of being neutron stars is supported by their association with supernova remnants [91, 122]. The source of SGRs has been explained by both magnetar models [60, 88], where the neutron stars are born with very short periods and high magnetic fields $> 10^{14}$ G, and accretion based models [141].

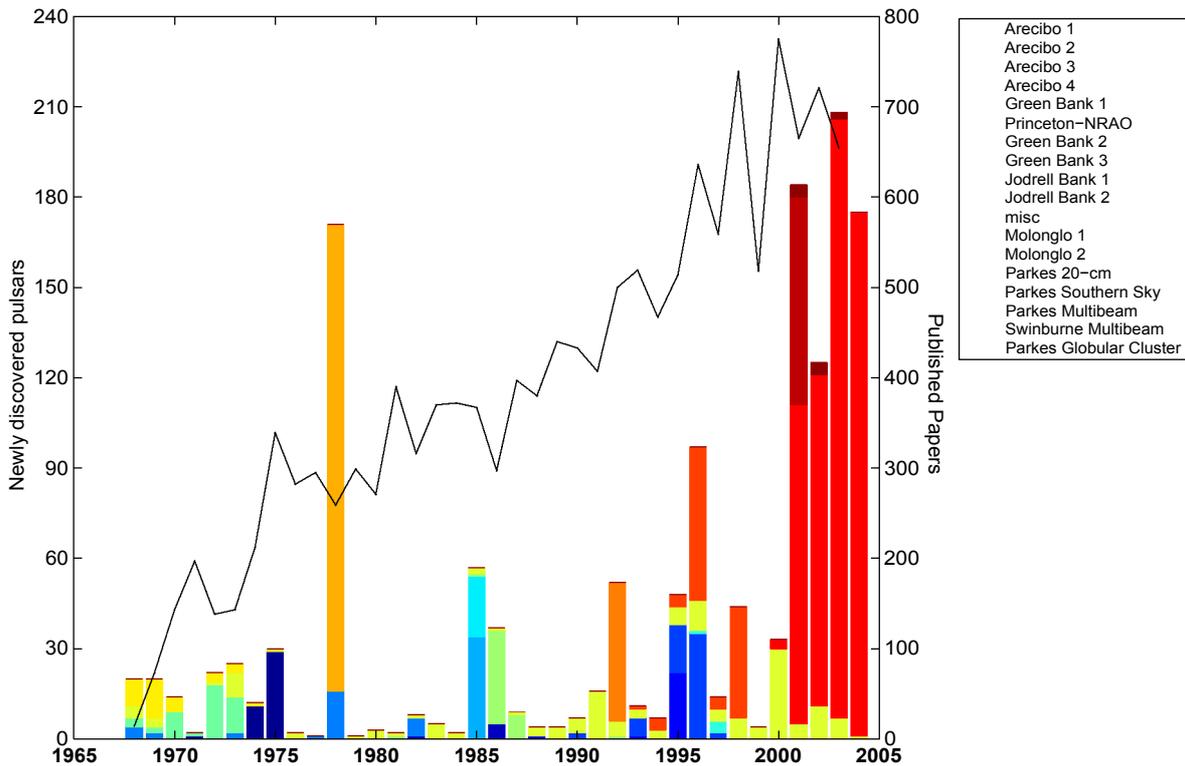


FIGURE 1.0.1. Newly discovered pulsars in each major survey according to the ATNF Pulsar catalogue [81]. The black line corresponds to the number of papers published with the keyword ‘pulsar’, according to the NASA Astrophysics Data System at http://adsabs.harvard.edu/abstract_service.html

- X-Ray Pulsars and Binaries: Most of the pulsating X-ray sources have been associated with binary systems, where the release of energy from accreting matter onto the neutron star surface is emitted as X-rays. They are classed in terms of companion mass size. The High Mass X-ray Binaries (HMXBs) have a companion mass size greater than $10M_{\odot}$. With periods of $0.069 < P < 835$ s, the timing stability of HMXBs is far less than that of radio pulsars. They are found to show a spin up, and characterised by variable X-ray emission. The Low Mass X-ray Binaries (LMXBs), with companion masses less than $2M_{\odot}$, are difficult to observe with little, if any, light observable from the companion star. The main source of light is due to the liberation of gravitational energy. This heats the surrounding gas, which subsequently gives off X-rays.
- Anomalous X-Ray Pulsars (AXP): Much like the SGRs, the AXPs can emit multiple, regular X-ray bursts with a spectrum that appears softer than that of the SGR emission.

With accurate spin down models, analysis of timing residual behaviour allows for the study of the internal and external structure of radio pulsar systems [92, 47, 46]. This includes magnetospheric fluctuations [36], vortex creep and internal temperature models [10] and interactions between the stellar crust and superfluid interior [29]. Observational studies have resulted in accurate measurements of the neutron star mass, radius, temperature, age and internal composition [96], whilst the rotational stability of millisecond pulsars in binary orbits has been shown to be consistent with Einstein's theory of general relativity [34, 50, 163, 167]. The newly discovered double pulsar system PSR J0737-3039 provides the opportunity for fundamental relativistic tests of gravitation [109].

1.1. The Nature of Torqued Pulsars

The spin down phenomenon of pulsars has led to strong interest in the modeling of pulsar magnetic fields and the surrounding magnetosphere. With no viable alternative, the rotating magnetic dipole model has dominated our understanding of isolated pulsar dynamics. However, this model is at variance with a spate of observational evidence. The theoretical braking law that describes the spin down of a rotating magnetic dipole is dependent on the rotation frequency to the third power (commonly referred to as the braking index). Observed spin down behaviour is in regular disagreement with theory [4, 102], with possible reasons including the different scaling of torque due to outflowing plasma [27] and magnetic field distortion [143]. Polarimetric results of millisecond pulsars have shown behaviour that departs significantly from that expected of a rotating magnetic dipole [139]. Possible time dependencies have been considered in the braking law [115] as a result of a time dependent braking torque [4]. In many cases, however, dominant timing noise has shown to be the most likely cause of inconsistent values for measured braking indices [19, 37].

Although timing noise often has the form of stochastic fluctuations, cyclic variations about the mean spin down trend have been observed in numerous pulsars. This is thought to be the result of precession [54, 114, 156, 164]. The verification of a precession mechanism would pose difficult questions of current understanding of neutron star interiors and beam models.

While most candidate precessing pulsars are seemingly fitted with untorqued models, this is contradictory to the observation of a pulsar that is spinning down. The effect of the magnetic field on precession has been thoroughly described by a number of authors [102, 103, 124, 125, 126]. A computational model, which produces theoretically observable residuals, would be an important tool to aid the understanding of the dynamical behaviour of torqued pulsars.

1.2. Pulsar Observation

The efficacy of the detection and monitoring of pulsars is limited by receiver instrumentation. The success of the Parkes Multibeam survey, which has more than doubled the number of catalogued pulsars [80, 90, 120, 131] (see figure 1.0.1), lies in the multi-beam feed array and modern electronic data processing. See D'Amico (1998) [48] for a current summary of major pulsar surveys.

The low cost, digital electronics of the modern era have paved the way for sophisticated, multi-purpose reprogrammable solutions such as COBRA [39] and the Mark IV pulsar processor [146]. They have allowed for the introduction of re-programmable solutions in the guise of software receivers, field programmable gate arrays (FPGAs) and microcontrollers. Since their introduction in the mid 1980's, FPGAs have become viable alternatives to custom hardware, suited to perform a variety of computational tasks such as digital signal processing [72, 73, 144] and radar processing [72]. Their use has changed communication technology and will, in the not too distant future, become standard hardware components in modern observatories.

Since the commencement of the pulsar timing program at Hartebeeshoek Radio Astronomy Observatory (HartRAO) in 1984, few modifications to the pulsar timing receiver instrumentation have been made, although significant improvement of the telescope has occurred with the installation of cryogenically cooled receivers. Hence, with the inability to detect newly discovered, low flux objects, the pulsar timing program has produced long dataspans of timing residual behaviour of 27 pulsars. To date, the primary science of the timing program has been the monitoring of glitches in the Vela pulsar and PSR B1641-45 [61, 63].

1.3. This Thesis

This thesis aims to examine the dynamical behaviour of isolated torqued pulsars in order to investigate evidence of precession in the HartRAO pulsar sample. A scientific case will be developed to argue for new digital instrumentation for the extension of the monitoring of precessing pulsar candidates at HartRAO.

The 20 years span of timing residual data at HartRAO affords the opportunity to study the long term spin down behaviour of 27 pulsars. However, measurement of braking indices for these pulsars have shown to be inconsistent with known braking law theories [37]. Chapter 2 discusses the theoretical and observational anomalies of the braking law for torqued pulsars.

The apparent cyclic variations in the HartRAO timing residual data for a subset of pulsars suggests precessional behaviour. Although already identified as possible candidates for free precession, the notion of a spinning down pulsar implies that a torqued model is appropriate.

Chapters 3 and 4 present the results of numerical modeling for untorqued and torqued precessional behaviour, with a comparison of the theoretical predictions and observed timing residuals at HartRAO following in chapter 5.

The ability of HartRAO to continue its timing program, and to expand its detectable pulsar subpopulation, is limited by the current instrumentation. Chapter 6 makes a scientific argument for the construction of a new pulsar processor at HartRAO, with the design and development of the processor discussed in chapter 7. Subsequent testing and results for the new instrument are discussed in chapter 8. Chapter 9 consists of a general conclusion and discussion of future work, followed by appendices A-L. Appendices F, K and L are found on the accompanying CD.

CHAPTER 2

Braking Pulsars

2.1. Introduction

Consideration of a torqued precession model when analysing cyclic behaviour in measured timing residuals requires a knowledge of the braking mechanism of pulsars. Whilst a rotating magnetic dipole is the most common model used for a braking pulsar, and has been adopted for the purposes of this thesis (see chapter 4), it is important to determine whether pulsar observations provide any evidence as to the nature of pulsar braking mechanisms. This chapter expands on what we know about the measurement of braking pulsars, and the interpretation of noisy data. Deviations from theoretically predicted results are described in terms of plausible physical mechanisms as well as interpretative errors.

2.2. Braking Model

The rotation rate of isolated pulsars has been observed to vary with time. The dominant trend is one of slow spin down. Superimposed on this spin down are smaller effects, which may be classified into three categories: sudden and dramatic spin up events known as ‘glitches’, stochastic fluctuations about the mean periodicity, known as ‘timing noise’, and highly correlated oscillations about the mean spin down trend, the cause of which may be interpreted as systematic, non-random behaviour of the pulsar, such as precession or the presence of a binary companion.

The dominant spin down trend is assumed to evolve according to a theoretically predicted spin down law,

$$(2.2.1) \quad \dot{w} = -Kw^n$$

where w is the rotational angular velocity, the exponent n is called the braking index and K a positive coefficient dependent, in general, on the various parameters that determine the magnitude of the braking torque. The coefficient K has no established name and will be called the *torque coefficient* for the purposes of this thesis.

The standard model of a pulsar is a spinning magnetic dipole with moment m rigidly fixed in the crustal surface of the star [126, 83]. The pulsar rotates with angular velocity w about an axis inclined at an angle a to the magnetic axis, radiating energy at the rotation frequency by magnetic dipole radiation. The radiation reaction torque on a spinning dipole is ([174], pp 482-484)

$$\tau = \frac{2m^2 \sin^2 a}{3c^3} w^3 \dot{w}$$

where c is the speed of light. The equation of motion $\dot{\tau} = \frac{dJ}{dt}$, J being the angular momentum, results in a change of angular velocity

$$(2.2.2) \quad \dot{W} = -\frac{2m^2 \sin^2 \alpha}{3Ic^3} W^3$$

which is of the form shown in equation 2.2.1, with $n = 3$ and $K = \frac{2m^2 \sin^2 \alpha}{3Ic^3}$. Similarly, quadrupole and gravitational radiation requires $n = 5$, while multi-pole radiation will have even higher indices [11].

In reality, the assumption of a single braking mechanism is incorrect since several braking mechanisms will operate simultaneously. The spin down law should therefore be written as a linear superposition of braking law terms

$$\dot{W} = -\sum_i K_i W^{n_i}$$

where, in general, K_i and n_i ($n_i \neq 1$) are different. To a good approximation, one of the terms in this expansion will dominate the others by several orders of magnitude, reducing the expansion to equation 2.2.1.

To test the braking law by direct measurement, equation 2.2.1 could be rewritten in the form

$$\log(-\dot{W}) = \log K + n \log W$$

where a plot of the measured values $\log(-\dot{W})$ vs $\log W$ results in the determination of K and n using linear regression. This method, however, fails because measured values for \dot{W} show fluctuations dominated by timing noise. The large scatter on data points makes the determination of K and n unreliable.

To determine n , experimenters differentiate equation 2.2.1

$$\dot{W} = -KnW^{n-1} \quad \ddot{W} = nW^{-1} \dot{W}^2$$

to obtain

$$(2.2.3) \quad n_{obs} = \frac{\dot{W}\ddot{W}}{W^{-2}}$$

where n_{obs} is the braking index calculated from observed parameters. The common use of equation 2.2.3 by astronomers has effectively redefined the concept of the braking index, which is not necessarily equivalent to that of the theoretical braking index n used in equation 2.2.1. For K and n constant, equation 2.2.3 can be solved for W , resulting in the equivalence of n_{obs} and n . However, there is reason to believe that either one, or both, of K and n are time dependent [4, 27]. For either $K \neq 0$ or $\dot{n} \neq 0$, equation 2.2.3 does not follow from 2.2.1 and the two parameters n and n_{obs} should be regarded as distinct. The definition of n_{obs} in equation 2.2.3 is therefore independent of any form of spin down law that a pulsar may obey, and does not presuppose a spin down law of the form given in equation 2.2.1, unless by the ideal approximation

$\dot{\kappa} = \dot{n} = 0$. The concept of a spin down law should not admit a braking index as defined by equation 2.2.1, and so we have no alternative but to admit equation 2.2.3 as a definition.

To maintain this distinction, n will be called the *theoretical braking index*, while n_{obs} the *observed braking index*. Given the definition of n_{obs} in equation 2.2.3, and the fact that ω is a function of time, we have no *a priori* reason to assume that n_{obs} will be constant.

2.3. Distinction between n and n_{obs}

Equations 2.2.1 and 2.2.3 are equivalent for κ and n constant. It is simple to confirm this relation by calculating n_{obs} from equation 2.2.1 (see section 2.2). It is necessary to show the converse: given n_{obs} constant, then equation 2.2.3 must imply a spin down law of the general type in equation 2.2.1.

Assuming $n_{obs} = n_0$ constant, the differential equation 2.2.3 is solved by substitution of $u = \ln \omega$

$$n_0 = \frac{\dot{\omega}\omega}{\omega^2} = 1 + \frac{\ddot{u}}{\dot{u}^2} = 1 - \frac{d}{dt} \frac{c}{\dot{u}}$$

which integrates to

$$\dot{u} = \frac{1}{-(n_0 - 1)t + c}$$

where c is a constant of integration.

This function has two solution branches selectable by choice of initial condition, defined at time $t = 0$. Since we know the pulsar to be spinning down, $\dot{u} < 0$ making $c < 0 \Rightarrow c = -k^2$, k a constant. Then

$$\dot{u} = \frac{1}{(n_0 - 1)t + k^2}$$

and, provided $n_0 \neq 1$, this integrates to

$$u = \frac{1}{n_0 - 1} \ln[(n_0 - 1)t + k^2] + b$$

where b is a constant of integration. The solution is valid for all t such that $(n_0 - 1)t + k^2 > 0$. Expressed in terms of ω ,

$$\omega = e^{\frac{b}{(n_0 - 1)t + k^2}}$$

which is an explicit solution for $\omega = \omega(t)$. To obtain the corresponding spin down law, ω is differentiated to give

$$\begin{aligned} \dot{\omega} &= -e^{\frac{b}{(n_0 - 1)t + k^2}} \left[\frac{1}{(n_0 - 1)t + k^2} \right]^{-1} \\ &= -e^{-(n_0 - 1)t + k^2} \omega^{n_0} \end{aligned}$$

which is of the form $\dot{\omega} = -\kappa\omega^n$. It has been shown that for n_{obs} constant and not equal to 1, the spin down law is necessarily of the form in equation 2.2.1 and $n = n_{obs}$.

If $n_{obs} = n_{obs}(t)$, the above result does not follow.

$$\begin{aligned}
 \ddot{\omega} &= n_{obs}(t) \dot{\omega} \\
 \dot{\omega} &= \frac{1}{1 - \int n_{obs}(t) dt + c} \\
 \Rightarrow \dot{u} &= \frac{1}{1 - \int n_{obs}(t) dt + c} \\
 \Rightarrow u &= \int \frac{1}{1 - \int n_{obs}(t) dt + c} dt + b \\
 \Rightarrow \omega &= e^{\int \frac{1}{1 - \int n_{obs}(t) dt + c} dt + b} \\
 \Rightarrow \dot{\omega} &= \frac{1}{1 - \int n_{obs}(t) dt + c} e^{\int \frac{1}{1 - \int n_{obs}(t) dt + c} dt + b} \\
 &= \frac{1}{1 - \int n_{obs}(t) dt + c} \omega
 \end{aligned}$$

This is as far as the calculation can be taken due to the indefinite integral $\int n_{obs}(t) dt$. This does not satisfy the form of the spin down law in equation 2.2.1 unless $n_{obs}(t)$ satisfies the equation

$$\xi(t) = e^{\int \xi(t) dt + b}$$

for $\mathcal{V}(t)$ where $\xi(t) = \frac{1}{1 - \int n_{obs}(t) dt + c}$. In general, therefore, if n_{obs} is time dependent, the spin down law will not necessarily have the form in equation 2.2.1. This result can be expressed in terms of a theorem:

THEOREM 2.3.1. *If the spin down law has the form*

$$\dot{\omega} = -\kappa \omega^n$$

the theoretical and observed braking indices are equivalent if and only if κ and n are constant.

COROLLARY 2.3.2. *If the observed braking index*

$$n_{obs} = \frac{\ddot{\omega}(t)\omega(t)}{\dot{\omega}(t)^2}$$

is constant and not equal to 1, then the spin down law has the form $\dot{\omega} = -\kappa \omega^n$ with κ constant and $n = n_{obs}$.

2.4. Time Dependence in the Braking Law

Timing irregularities, such as glitches (including glitch recovery) and timing noise, have made long timing spans necessary for the reliable measurement of $\dot{\omega}$. To date, only four pulsars have had their observed braking indices measured to within theoretically expected magnitudes [82]. Only the Crab pulsar has long enough data spans available to ensure a reliable measurement of n_{obs} . Separating this data into four 5 year time spans to obtain four values for n_{obs} , Lyne et al (1993) [115] found an appreciable time dependency in n_{obs} . The data, however, spanned merely 2% of the total lifetime of the pulsar. Statistically, this would not constitute a well sampled population. The determination of n_{obs} is dominated by timing noise and the observed time dependency could be a sampling effect. Baykal et al (1999) [19] were able to compile over 30 years of timing data for PSR 0823+26, PSR 1706-16, PSR 1749-28 and PSR 2021+51. However,

the data was highly contaminated with frequency derivative noise and led to unphysical braking indices of -10726, 1739, 301 and -887 respectively.

Observed braking indices are in disagreement with the theoretical value of $n = 3$ expected for magnetic dipole radiation [102, 4]. This could be due to a different scaling of the torque with frequency other than that of a magnetic dipole radiating *in vacuo*, possibly caused by out-flowing plasma, removing a significant amount of angular momentum [27], or strong magnetospheric currents, which produces a non-dipolar magnetic field. Magnetic field distortion can also take place via mass accretion, requiring $M_a \geq 10^{-5} M_p$ where M_a is the accreted mass [143]. However, time dependencies in the torque function κ are highly likely to be the cause of deviation of the observed braking index from 3. Inspection of the form of κ , shown in equation 2.2.2, indicates that any number of physical effects may be the cause of time dependencies: magnetic axis wandering [119], non-dipolar magnetic field structure [26], variable magnetic moment as a result of magnetic field decay or a change in the angle between m and ω [27], time dependencies in the moment of inertia of the star, plasma removal from the magnetosphere and propeller torques from supernova fallback [127]. This list is not exhaustive, but is indicative of the type of physical mechanisms that affect pulsar spin down. It is suggested by Link & Epstein (1997) [102] that the persistent, permanent increase seen in the spin down rate $\dot{\omega}$ of the Crab pulsar during post-glitch recovery is due to sudden, glitch induced reorientations of the star's magnetic axis. The discrete effect on α would cause variation in the magnetic field and torque [104]. The variation in the moment of inertia of the star can be attributed to shearing material across the equator, caused by starquakes [107]. The shearing action results in the accumulation of crustal material (approximately 100 μm high) near shearing faults. This breaks axial symmetry of the principal moments of inertia, shifting I by a small angle, and causes the star's angular momentum vector and rotational vector to become misaligned, resulting in precession. Another possible source of a changing I is the accumulation of superfluid vortices in capacitor regions, coupling a significant fraction of the superfluid to the crust [8]. It is thus unreasonable to expect κ and n to be constant.

First consider a time dependent braking index with constant κ . Assume that if the braking index has a time dependence, then it is slow and

$$n(t) = n_0 + \varepsilon f(t)$$

with $\varepsilon \ll 1$. For a constant κ

$$(2.4.1) \quad n_{obs} = \frac{\ddot{\omega}\omega}{\dot{\omega}^2} = n(t) - \frac{\varepsilon}{\kappa} \omega^{1-n} \dot{f}(t) \ln \omega$$

The observed braking index n_{obs} is different from the theoretical braking index $n(t)$ by a quantity of first order smallness. Since the measured braking indices vary substantially from the theoretically expected value of 3, time dependence in n is therefore not the dominant factor responsible.

Consider now a constant braking index n and $\kappa = \kappa(t)$. Equation 2.2.1 gives

$$(2.4.2) \quad n_{obs} = \frac{\ddot{\omega}\omega}{\dot{\omega}^2} = n + \frac{\dot{\kappa}(t)\omega}{\kappa(t)\dot{\omega}}$$

which can lead to a significant difference in the theoretical and observed braking indices. These differences would arise from either a discrete event, such as a glitch, or a continuous variation in the torque coefficient.

Allen and Horvath (1997) [4] estimate the total amount of continuous and discrete variation in the torque coefficient that would be necessary to account for the values of observed braking indices, using a theoretical value of $n = 3$. They estimate a continuous variation of κ to be

$$\frac{\dot{\kappa}}{\kappa} = (n_{obs} - 3) \frac{\dot{\omega}}{\omega} \approx 1.9 \times 10^{-4} \text{ yr}^{-1}$$

while a discrete variation is

$$\frac{\Delta \kappa}{\kappa} \approx 3 \times 10^{-5} \text{ yr}^{-1}$$

The discrete variation accounts for only 15% of the total variation. In the rotating dipole model, κ is a function of the magnetic field strength B , magnetic alignment angle α and the moment of inertia I (see equation 2.2.2). Recalculating equation 2.2.3 for $\kappa = \kappa(t)$, the generalised dependency of n_{obs} on these parameters is

$$(2.4.3) \quad n_{obs} = n + \frac{\dot{\kappa} \omega}{\kappa \dot{\omega}} = n + \frac{\dot{\omega}}{\omega} \frac{I}{I} - 2 \frac{\dot{\alpha}}{\alpha} - 2 \frac{\dot{B}}{B}$$

Equation 2.4.3 can therefore be used to set limits on the variation of parameters in κ to fit measured data.

We consider the motion of the magnetic axis with a varying magnetic inclination angle α . Discrete variation would correspond to an average growth rate of 1.5×10^{-5} rad/yr. The magnetic inclination angle for the Crab pulsar, α_{Crab} , would have moved just 0.82 degrees in its lifetime. Romani & Yadigaroglu (1995) [151] show that $\alpha_{Crab} \approx 70^\circ$, giving a required increase of about 10^{-4} rad/yr. An increase of this magnitude would result from continuous variation. A mechanism for this variation in α is that proposed by Ruderman's plate tectonic theory [153]. The drifting of magnetic poles from the rotation axis to the equator, which is a consequence of the breakdown of the solid structure (leading to starquake induced glitches), causes a net increase in the spin down rate. The movement of crustal plates occurs continuously during the spin down lifetime of the pulsar [152], resulting in continuous variation in κ . Further mechanisms for time dependencies in α could include post glitch reorientation of the magnetic axis and the tendency for the star to become an aligned rotator through the minimisation of its rotational energy state.

It has been argued by a number of authors that neutron star magnetic fields evolve over a timescale of a few million years [42, 71, 134]. The measurement of $n_{obs} < 3$ could imply growth in the external magnetic field, possibly due to ohmic diffusion of a trapped field beneath the crust, where $B_{ext} \approx 10^8 - 10^9$ G and $B_{int} > 10^{12}$ G. In the case of the Crab pulsar, this would be possible if the pulsar underwent fast neutrino cooling, resulting in the rapid freezing and corresponding decrease in the electrical resistivity of the crust. After approximately 20 years, the growth rate of the magnetic field slows considerably, and the magnetic field enters the saturation regime [133]. This scenario is further complicated by the action of accreting material onto the surface of the pulsar, which can act to initially screen the surface magnetic

field and is dependent upon the rate of accretion (comparable to that of ohmic diffusion at $\dot{M}_c \approx 10^{-10} M_\odot \text{ yr}^{-1}$) [44].

Ohmic dissipation leads to a decrease in magnetic field strength on a typical timescale of

$$\tau_\mu = \frac{4\sigma R^2}{\pi c}$$

where σ is taken as the electrical conductivity of the crust, R the radius of the star and c the speed of light [93]. Predicted values range from $\tau_\mu \approx 4 \times 10^6 \text{ yr}$ [140] to $\tau_\mu > 10^8 \text{ yr}$ [35], depending on the assumed electrical conductivity of the neutron star interior. The dependence of the magnetic field on the thermal evolution of neutron stars is well studied, raising the possibility of thermal and rotational instabilities in neutron stars due to crust-core interactions [24, 33, 95, 160, 159]. Evolution of the magnetic field due to the Hall effect also gives reasonable timescales [44].

Since $n_{obs} \approx 2.5$, $\dot{n}_{obs} < 0$, for the Crab pulsar [115], the magnetic field could be interpreted to be increasing, particularly at times of glitches. The characteristic age of the pulsar is therefore increasing slowly with time. The evolution of the pulsar on the $\dot{P} - P$ diagram is toward the region generally occupied by magnetars [110]. If it is supposed that $n = 1$ for a Crab-like pulsar, the evolution time would be approximately 5000 years. On this timescale, the supernova remnant from which the pulsar originated may still be visible.

2.5. The Generalised Braking Index

The interpretation of n_{obs} is such that it is defined by a spin down law of any form. In the special case where n_{obs} is constant, it coincides with the theoretical braking index defined in the spin down law of the form given by equation 2.2.1. If it is time dependent, this interpretation falls away. However, n_{obs} can be naturally interpreted in terms of a dimensionless time series in ω . This also allows for an interpretation of the second braking index and other higher order braking parameters.

If $\omega(t)$ is continuously differentiable to order k , $k > 2$, and is smooth, then $\omega(t)$ will possess a power series expansion about any given value t_0 of t of the form

$$(2.5.1) \quad \omega(t) = \omega(t_0) + \dot{\omega}(t_0)(t - t_0) + \frac{1}{2!} \ddot{\omega}(t_0)(t - t_0)^2 + \dots$$

where dot notation represents differentiation by time, with all coefficients evaluated at time $t = t_0$. Both dependent and independent variables are dimensioned and are not suitable for mathematical arguments nor for use as parameters designed to measure physical characteristics independent of the unit of measurement used. Therefore, equation 2.5.1 is written in a dimensionless form. This is done by first removing $\omega(t_0)$ as a common factor

$$\omega(t) = \omega(t_0) \left[1 + \frac{\dot{\omega}(t_0)}{\omega(t_0)} (t - t_0) + \frac{1}{2!} \frac{\ddot{\omega}(t_0)}{\omega(t_0)} (t - t_0)^2 + \dots \right]$$

The series inside the square brackets is clearly dimensionless. In particular, the term $\frac{\dot{\omega}(t_0)}{\omega(t_0)} (t - t_0)$ is a normalised dimensionless measure of time, in which the time t is measured in units of $\frac{\omega(t_0)}{\dot{\omega}(t_0)}$. This ratio is a natural timescale, or unit of time, which is defined by the spin down law and is generally related to a characteristic spin down time. Although commonly referred to as the

‘characteristic age’ of the pulsar, this timescale is not directly related to the actual age of the pulsar.

It is convenient now to introduce a dimensionless time parameter τ defined as

$$\tau = \frac{\dot{\omega}(t_0)}{\omega(t_0)} t$$

and similarly

$$\tau_0 = \frac{\dot{\omega}(t_0)}{\omega(t_0)} t_0$$

Then equation 2.5.1 becomes

$$(2.5.2) \quad \omega(t) = \omega(t_0) \left[1 + (\tau - \tau_0) + \frac{1}{2!} \frac{\ddot{\omega}(t_0)\omega(t_0)}{\dot{\omega}^2(t_0)} (\tau - \tau_0)^2 + \dots \right]$$

Equation 2.5.2 is in dimensionless form and the coefficients of the series in brackets carry information about the spin down law in a simple way. The coefficient of $\frac{(\tau - \tau_0)^2}{2!}$ is the observed braking index evaluated at time t_0 . The second, third and higher order braking indices, to order k , are obtained from a generalised equation in the form

$$n^{(k)} = \frac{1}{\omega} \frac{d^k \omega}{dt^k} \frac{\omega}{\dot{\omega}} = \frac{\omega^{k-1} d^k \omega}{\dot{\omega}^k dt^k}$$

where the k^{th} braking index is defined as the coefficient of the term $\frac{(\tau - \tau_0)^k}{k!}$ in the dimensionless power series expansion. In general, the braking indices $n^{(k)}$ will be functions of time, while those obtained from equation 2.5.2 will be $n^{(k)}$ evaluated at time t_0 .

The implication of equation 2.5.2 is that a single measurement of n_{obs} is simply a measurement of the coefficient of the quadratic term in equation 2.5.2. This does not constitute a test of the spin down law and nothing can be inferred about its form. No assumption was taken with regards to the form of the spin down law, yet it was possible to obtain a generalised formula for the braking indices to any order simply by writing out a time series expansion. This is illustrated further by the form of the solution for n_{obs} in equation 2.4.2, obtained for a constant theoretical braking index n and varying torque function κ . Assuming $\kappa = \kappa(t)$, a power series expansion is obtained

$$\kappa(t) = \kappa(t_0) \left[1 + \frac{\dot{\kappa}(t_0)\omega(t_0)}{\kappa(t_0)\dot{\omega}(t_0)} (\tau - \tau_0) + \frac{1}{2!} \frac{\ddot{\kappa}(t_0)\omega(t_0)^2}{\kappa(t_0)\dot{\omega}(t_0)^2} (\tau - \tau_0)^2 + \dots \right]$$

where $\tau = \frac{\dot{\omega}(t)}{\omega(t)} t$ and $\tau_0 = \frac{\dot{\omega}(t_0)}{\omega(t_0)} t_0$. The second term in equation 2.4.2 is therefore the coefficient of the quadratic term in the power series expansion of a generalised function $\kappa = \kappa(t)$. Again, higher order terms can be obtained to determine the full time dependency of $\kappa(t)$.

To establish the form of the spin down law, one needs to measure a number of coefficients using existing data spans, or at least until there is a sufficient number indicating those coefficients of the higher order terms are negligible. To date, this has not been the case, with attempts at measuring the second braking index resulting in values greater than the first braking index [27] and showing a power of n dependence. This is a clear manifestation of the problem of noisy data. As higher order coefficients are determined from the data, more noise is fitted. The strong presence of fitted noise decreases the possibility of fitting long term spin down trends.

The assumption that the only fault with the spin down law is that $n \notin 3$ but a number $1 < n < 3$ is therefore premature. We have no indication at present as to the form of the braking law, and will not until a long enough data span is available to establish the form of time dependence of n_{obs} .

2.6. Timing noise in the Measurement of n_{obs}

For pulsars monitored at HartRAO, the measurement of n_{obs} , obtained by polynomial fits of phase residuals, resulted in values inconsistent with any known theory of braking indices [37]. Only two, PSR B0450-18 and PSR B1727-47, of the 27 regularly monitored pulsars had $0 < n_{obs} < 10$, with the rest ranging between large positive and negative values. The inconsistent results, together with physically implausible measurements of negative values for $\dot{\omega}$, shows evidence for the fitting of timing noise through higher order polynomial fits of the form in equation 2.5.2 (Section 2.6.2 discusses a mechanism that could result in large positive and negative measurements of n_{obs}). Chukwude (2002) [37] was able to show strong influence of timing noise absorption levels on measured values of $\dot{\omega}$ in HartRAO data by correlating the measured value for $\dot{\omega}$ with the rms phase residuals obtained from second and third order polynomial fits. The determination of the coefficients in equation 2.5.2, be it by polynomial fitting or measurement of ω , $\dot{\omega}$ and $\ddot{\omega}$, should therefore not be attempted to be reconciled with any existing braking index theories.

An assumption that the measured values for $\dot{\omega}$ are not heavily contaminated by timing noise would require the presence of a physical process that is independent of the long term braking trend of the pulsar. A stationary stochastic process with non-zero mean could result in the measurement of a net $\dot{\omega}$, the value of which would appear anomalous using existing braking index theories [43]. The stochastic process would originate from a noisy torque, either internal or external, which is not modeled in the braking law of equation 2.2.1. Reported measurements of $\dot{\omega}$ using the HartRAO data show that 13 of the 27 regularly monitored pulsars have negative measured values, possibly the result of noisy torques.

2.6.1. Noisy Rotational Motion. The theoretical basis for the effects of a noisy torque on the rotational behaviour of a pulsar is taken from the development of brownian motion for rotating rigid bodies [38, 123].

Following the pulsar clock model of Cordes (1980) [40], a pulsar can be treated as a rotating sphere subject to a noisy torque. The origin of this torque is not important in this discussion and can be represented as an angular collision torque $\boldsymbol{\tau}(\mathbf{t})$. The rate at which stochastic events occur in $\boldsymbol{\tau}(\mathbf{t})$ can be characterised by a correlation time $\Delta\mathbf{t}$, which will be of the same order as the response time to that event.

The nature of $\boldsymbol{\tau}(\mathbf{t})$ is such that the rotation of the body has an effect on the torque. This implies the presence of a slowly varying component $\bar{\boldsymbol{\tau}}(\mathbf{t})$, which acts to bring the rotating body into equilibrium together with a rapidly fluctuating component with an expectation value of 0. $\bar{\boldsymbol{\tau}}(\mathbf{t})$ is expanded in terms of a power series where the first non-vanishing term is linear in $\boldsymbol{\omega}(\mathbf{t})$,

to give the general form

$$\tilde{\tau}(\omega(t)) = \beta\omega(t)$$

where β is a constant.

Now consider the process

$$\omega(t) = \frac{d\theta(t)}{dt}$$

that satisfies the equation

$$\frac{d\omega(t)}{dt} + \beta\omega(t) = n(t)$$

where $n(t)$ is a white noise process with zero mean $E\{n(t)\} = 0$ and constant frequency spectrum $S_n(\omega) = \mu$, and β a frictional constant. $\omega(t)$ describes the angular velocity of the body. If we consider a massive body with moment of inertia I that is acted upon by a frictional force proportional to the angular velocity of the body, with no other external torques explicitly expressed, the equation can be rewritten

$$(2.6.1) \quad \dot{\omega}(t) + \xi\omega(t) = \tau(t)$$

where $\dot{\omega}(t) = \frac{d\omega(t)}{dt}$, $\xi = \beta I$ and $\tau(t) = \frac{n(t)}{I}$ is an angular collision force. Choosing a cartesian coordinate system, the derivation is performed in one dimension. The results are the same for the other two basis vectors.

The microscopic nature of $\tau(t)$ can be complicated, but by assuming the differential dt to be large compared to the response of each collision, but small enough to cause a small variation in the process $\omega(t)$, then both $\omega(t)$ and $\tau(t)$ can be treated as moving averages, written in terms of instantaneous values $\omega_j(t)$ and $\tau_i(t)$ as

$$\omega(t) = \frac{1}{\Delta t} \int_t^{t+\Delta t} \omega_j(\alpha) d\alpha$$

$$\tau(t) = \frac{1}{\Delta t} \int_t^{t+\Delta t} \tau_i(\alpha) d\alpha$$

The rotating body would therefore have a delayed response to the stochastic process, and behave under the action of an averaged process. $\omega(t)$ and $\tau(t)$ still satisfy equation 2.6.1, shown by considering the instantaneous values $\omega_j(t)$ and $\tau_i(t)$ in the equation

$$\dot{\omega}_j(t) + \xi\omega_j(t) = \tau_i(t)$$

and integrating.

ξ would not only be dependent on I , but also on the viscosity of the surrounding mediums of the pulsar crust, both internal and external. It is interpreted as the result of a preferred sign stochastic event and leads to a net measurement of $\dot{\omega}$ that is not explained by braking law theory. If the dominant stochastic process were to originate in the interaction of the magnetosphere with the pulsar, such as a differential rotation within the magnetosphere [145], the pulsar should be damped and $\xi < 0$. Were it to originate in the crust-core interaction, through the effects of viscosity and friction due to the scattering of electrons by neutron and proton vortices, then in

order to bring the system into equilibrium with the faster rotating core, $\dot{\nu} > 0$. This would lead to a time dependency in $\dot{\nu}$, since we know the long term dynamics of a pulsar is to slow down.

2.6.2. Internal Frictional Instabilities . The effect of a frictional crust-core interface is well studied. It is known to lead to thermal and rotational instabilities that are largely dependent on the differential crust-core angular velocity, temperature sensitivity and equations of state [24, 33, 95, 99, 159, 160, 171]. The friction dissipates rotational energy as heat. Once the interior of a neutron star reaches an isothermal state, friction becomes an important heat source, which can effectively alter the thermal evolution of the pulsar and the interior rotational structure. With a high sensitivity to temperature, the frictional crust core interface is unstable [160]. This instability in the coupling develops a limit cycle, which effectively has two states: a high temperature coupled state and a low temperature decoupled state [159].

In the coupled state the internal frictional torque is large, leading to a decrease in the differential crust core angular velocity, while angular momentum is dissipated as heat. The frictional torque reaches a maximum, causing $\dot{\nu}$ to increase rapidly. Once the differential crust core angular velocity is sufficiently small, the friction decreases quickly, with the cooling rate of the neutron star dominating the heating rate. As the star cools, it enters a low temperature decoupled state where the slow down trend of the pulsar obeys the standard magnetic dipole braking.

In the decoupled state, the differential angular velocity is large and the cycle is repeated with a dominant heating rate. These rotational instabilities lead to the measurement of anomalous values for $\dot{\nu}$, both negative and positive [159], and even positive measurements of $\dot{\nu}$.

Values for the frictional torque coefficient $\dot{\nu}$ can be calculated from measured parameters as follows. By assuming a spin down law of the form 2.2.1, with $n = 3$, predicted values for $\dot{\nu}$ can be calculated. Observed values for $\dot{\nu}$ are assumed to follow a spin down law of the form 2.6.1. The value $\dot{\nu}$ in equation 2.6.1 is calculated by determining the residual between predicted and measured values of $\dot{\nu}$. Using observed measurements of $\dot{\nu}$ for the 27 HartRAO pulsars analysed by Chukwude (2002) [37], table 2.6.1 lists the results of the determination of $\dot{\nu}$ for these pulsars. Detailed documents describing the data processing and timing program at HartRAO are described by Chukwude (2002) [37] and Flanagan (1995) [62].

$\dot{\nu}_{obs}$ corresponds to the observed values of $\dot{\nu}$ obtained by second and third order polynomial fits of observed phase residuals, while $\dot{\nu}_{pre}$ is the predicted value of $\dot{\nu}$ obtained by assuming $n = 3$. Errors for $\dot{\nu}$ are largely dependent on measurement errors for $\dot{\nu}_{obs}$.

Using a reasonable approximation of 2×10^4 yrs for the interior of a neutron star to reach an isothermal state [171], none of the pulsars in table 2.6.1 would be excluded from a sample group of possible frictional instabilities. Using the thermal evolutionary models of Shibasaki and Mochizuki (1995) [160], the onset of thermal rotational instability is reached on the order of 10^7 years. However, it is possible that stars of this age are too insensitive to temperature, with thermal rotational instabilities occurring in stars younger than 10^5 years [95]. The calculated characteristic age of the pulsars is true only for $n = 3$. In periods of rotational instability, marked by very large negative and positive measured braking indices, the calculated age based

on observed parameters would be a meaningless quantity. Whilst Shibasaki and Hirano (1995) [159] predict positive measured values of $\dot{\omega}$ accompanying the anomalous values of $\ddot{\omega}$ and n_{obs} , this does not seem to be the case for the HartRAO pulsars. The pulsars would therefore be entering high temperature decoupled states for $\ddot{\omega} \gg 0$ or leaving high temperature decoupled states for $\ddot{\omega} \ll 0$. Given that the duty cycle of the rotational instability is a few percent, the probability that all pulsars monitored at HartRAO are moving into or leaving a state of frictional instability is highly improbable, and we would expect perhaps only two or three of the monitored pulsars to have anomalous values for $\ddot{\omega}$.

Thermal rotational instabilities can hesitantly be excluded as a possible cause for the measured values of $\ddot{\omega}$ due to the sample statistics. There is, however, still room for a large frictional internal torque that occurs over a longer period of time. The observed parameters might suggest this.

TABLE 2.6.1. Calculation of ξ for 27 HartRAO pulsars

| Pulsar [J2000] | P [s] | \dot{P} [10^{-15}ss^{-1}] | Err | $\ddot{\omega}_{obs}$ [$10^{-25} \text{Hz.s}^{-2}$] | Err | $\ddot{\omega}_{pre}$ [$10^{-25} \text{Hz.s}^{-2}$] | t_{age} [10^6yrs] | ξ [10^{-13}s^{-1}] |
|-------------------|-----------------|--|-----|--|-----|--|------------------------------------|---------------------------------------|
| 0452-1759 | 0.5489390994991 | 5.75246 | 7 | 0.013 | 4 | 0.0060 | 1.51 | -0.36677 |
| 0738-4042 | 0.3749197185831 | 1.62787 | 3 | -8.389 | 6 | 0.0015 | 3.65 | 724.511 |
| 0742-2822 | 0.166762352983 | 16.8138 | 4 | -39.334 | 5 | 1.8298 | 0.157 | 68.0824 |
| 0837-4135 | 0.7516228351393 | 3.542096 | 16 | 0.366 | 9 | 0.0009 | 3.36 | -58.2328 |
| 1001-5507 | 1.43659803499 | 51.9079 | 8 | -13.85 | 3 | 0.0271 | 0.438 | 551.748 |
| 1056-559 | 0.4224478120999 | 3.571362 | 12 | -0.303 | 8 | 0.0051 | 1.87 | -15.3946 |
| 1136+1551 | 1.1879138426560 | 3.733837 | 5 | 0.0622 | 4 | 0.0002 | 5.04 | -23.4131 |
| 1224-6407 | 0.2164772303537 | 4.954019 | 4 | -0.096 | 8 | 0.0726 | 0.692 | 1.59465 |
| 1243-6423 | 0.3884813984385 | 4.50044 | 5 | -0.774 | 3 | 0.0104 | 1.37 | 26.3029 |
| 1326-5859 | 0.4779911691929 | 3.228641 | 6 | 2.472 | 4 | 0.0029 | 2.35 | -174.729 |
| 1327-6222 | 0.529917304148 | 18.88576 | 9 | -1.554 | 3 | 0.0719 | 0.445 | 24.1756 |
| 1359-6038 | 0.1275037184166 | 6.338324 | 12 | 5.048 | 14 | 0.5814 | 0.319 | -11.4563 |
| 1401-6357 | 0.84279410356 | 16.6632 | 8 | 8.553 | 6 | 0.0140 | 0.801 | -363.995 |
| 1430-6623 | 0.785441562099 | 2.7757 | 3 | -0.197 | 3 | 0.0005 | 4.48 | 43.8906 |
| 1453-6413 | 0.1794852776830 | 2.745765 | 4 | 1.561 | 10 | 0.0391 | 1.04 | -17.8557 |
| 1456-6843 | 0.2633768354079 | 0.098824 | 4 | -0.029 | 2 | 0.0001 | 42.2 | 20.3672 |
| 1559-4438 | 0.2570563022579 | 1.019427 | 4 | 0.217 | 6 | 0.0018 | 4.0 | -13.9467 |
| 1600-5044 | 0.1926008276350 | 5.06234 | 2 | -0.75 | 7 | 0.1076 | 0.603 | 6.28426 |
| 1644-4559 | 0.455064996897 | 20.10576 | 7 | 4.1687 | 5 | 0.1288 | 0.359 | -41.6111 |
| 1645-0317 | 0.387690159752 | 1.7777 | 2 | -0.299 | 4 | 0.0016 | 3.46 | 25.4179 |
| 1709-1640 | 0.65305499092 | 6.2908 | 10 | 6.394 | 9 | 0.0043 | 1.64 | -433.188 |
| 1731-4744 | 0.829791467779 | 163.59359 | 8 | 2.403 | 3 | 1.4056 | 0.0804 | -4.19955 |
| 1752-2806 | 0.562559701417 | 8.1190 | 2 | -0.223 | 2 | 0.0111 | 1.1 | 9.12536 |
| 1932+1059 | 0.2265179059111 | 1.156007 | 9 | -1.584 | 7 | 0.0035 | 3.1 | 70.4603 |
| 1935+1616 | 0.3587397964580 | 6.002093 | 3 | 0.120 | 2 | 0.0234 | 0.947 | -2.07106 |
| 2048-1616 | 1.961574563855 | 10.95891 | 2 | 0.0054 | 9 | 0.0005 | 2.84 | -1.72839 |

Errors are 2σ formal standard errors and denote the least significant digits.

2.6.3. External Frictional Torques. The fluctuation of pair production processes in the magnetospheric gap, with a typical timescale of $10 \mu\text{s}$, causes a variation in the dipole torque, which can be shown to cause red noise in pulse frequency derivatives [36, 19]. These processes create a rapid variation in the current braking torque $\Delta T = (\mathbf{S} \mathbf{J} \times \mathbf{B})$ where $\mathbf{S} \mathbf{J}$ is the perturbed magnetospheric current. This rapid variation has an intrinsic frictional nature that could lead to anomalous values for $\dot{\omega}$. The magnitude of these variations is speculative and more knowledge is needed on the structure of the rotating magnetosphere.

Evidence in the data available at HartRAO seems to support the contamination of $\dot{\omega}$ by timing noise [37]. This does not rule out the possibility of frictional torques influencing the rotation of pulsars. To determine the affect of friction on the spin down of pulsars, a long data span is required with little or no timing noise. Pulsars that exhibit smooth cyclic variations in their timing residuals, possibly due to precession, are ideal candidates for an investigation of this nature.

The fitting of a torqued precession model could be used to determine accurate models for the dynamics of spinning down pulsars. The behaviour of torqued pulsars is largely influenced by mechanical and magnetic stresses. The absence of relatively large stochastic fluctuations in the timing residuals of candidate precessing pulsars provides the possibility to analyse spin down behaviour using uncontaminated data. In the following chapters, the kinematics and dynamics of pulsar precession are discussed. Chapter 3 gives an overview of what we know about free precession in order to maintain consistency. Chapter 4 discusses torqued precession in depth, detailing both previous work and results from new numerical simulations. The precession models are used as a basis for the determination of observed timing residuals, which are correlated with timing residual data from HartRAO in chapter 5.

CHAPTER 3

Free Precession

3.1. Introduction

Certain pulsars display long term cyclic behaviour in their timing residuals, with periods varying between thirty days to over one thousand days. In most cases the timing residuals are not coherently periodic. Precession is a likely cause of the observed periodicities.

The pulsar interior is, in reality, not frictionless. Elastic energy strain is dissipated in the crust due to non-hydrostatic deformation [70], whilst imperfect coupling between the crust and core causes energy losses [105, 157]. Both sources of friction tend to dampen any initial precession over a time scale shorter than one precession period. The occurrence of free precession would require either a complete coupling, or decoupling, of the crust and superfluid, or a mechanism allowing for the continual re-excitement of the precession [85]. Thus, challenges for theories on neutron star interiors, the dynamics of crust-core interactions and beam models are posed by the observation of apparent long term precession in pulsars such as the Crab pulsar [114], Vela pulsar [54], PSR 1828-11 [164] and PSR 1642-03 [156]. Timing residual observations of PSR 1828-11, showing strong harmonics at periods of 167 days, 250 days, 500 days and 1000 days with an amplitude of 1ns, show no obvious signs of damping. Disagreements of timing residual observations with precession models, and the existence of long term precession, should be taken as an indication that current models of the neutron star interior are incomplete.

The problem of torque free pulsar precession is well studied [25, 45, 68, 84, 85, 157], and so only an overview of the field is made here to maintain consistency. Chapter 4 gives a more in depth analysis of torqued precession, detailing recent and new work.

3.2. The Rigid Body Model

Evidence of a superfluid interior demands modification of the adopted rotating rigid body model, which requires that each point in the body is fixed relative to all neighbouring points. Since observational measurements are linked to the rigidly rotating crust of the neutron star, the rigid body model is modified to treat only this component. This requires one of two cases: the superfluid interior is completely decoupled from, or fully coupled to, the crustal exterior. Partial coupling would require that pinned superfluid vortices do not interact with any uncoupled component. This is improbable given the presence of proton-neutron vortex interactions and the coupling of superfluid and non-superfluid components [147]. Both the fully coupled and decoupled approximations are problematic.

The absence of coupling between the superfluid interior and crust can explain the observation of long term precession, which would be damped within a single precession period for imperfectly

pinned crusts [155, 157]. It would, however, also dismiss a number of theories used to explain glitches and post glitch recovery [5, 6, 9, 7, 8, 86, 105]. For a pulsar displaying precession, be it due to an external force or internal dynamics, the resultant internal force could be sufficient to unpin the superfluid vortices and prevent further coupling, thus satisfying the uncoupled approximation. Link and Cutler (2002) [101] estimate the magnus force required for a glitch on Vela to be approximately $10^{15} \text{ dyn.cm}^{-1}$. This is a lower limit for a non precessing pulsar. By comparison, the minimum magnus force of a precessing pulsar $> 10^{17} \text{ dyn.cm}^{-1}$, indicating that the vortices in a precessing pulsar should unpin and remain unpinned. Observational proof would require that precessing pulsars do not glitch or, if they do, do not follow standard glitch recovery models.

The fully coupled body is limited by the lack of a glitch recovery model, implying the absence of differential rotation between the crust and core, and is not considered to be a plausible physical reality.

The rigid body model is a reasonably accurate approximation for the purpose of timing measurements and is supported by the observed rotational behaviour of pulsars [29, 53].

The following sections outline the standard mathematical approach to the problem of precession. Using Euler's equations for rigid body motion, solutions are obtained for the kinematical behaviour of precessing axisymmetric and asymmetric pulsars.

3.3. Equations of Motion

Choose a set of axes e_i ($i = 1, 2, 3$), which are fixed in a rigid body such that the moment of inertia tensor $I_{ij} = 0$ for $i \neq j$. The torque free Euler equations for rigid body motion are (see appendix A.1)

$$(3.3.1) \quad I_1 \dot{\omega}_1 + (I_2 - I_3) \omega_2 \omega_3 = 0$$

$$(3.3.2) \quad I_2 \dot{\omega}_2 + (I_3 - I_1) \omega_1 \omega_3 = 0$$

$$(3.3.3) \quad I_3 \dot{\omega}_3 + (I_1 - I_2) \omega_1 \omega_2 = 0$$

where the ω_i are the angular velocity components of the pulsar in the body frame e_i and I_i are the moments of inertia.

3.4. Axisymmetric Pulsar

e_3 is chosen to be the axis of symmetry. For an oblate body $I_3 > I_2 = I_1$. Defining $e = \frac{I_3 - I_1}{I_1}$, Euler's equations reduce to

$$(3.4.1) \quad \dot{\omega}_1 + e \omega_2 \omega_3 = 0$$

$$(3.4.2) \quad \dot{\omega}_2 - e \omega_1 \omega_3 = 0$$

$$(3.4.3) \quad \dot{\omega}_3 = 0$$

where $e > 0$. For a prolate body, $e < 0$.

3.4.1. **General Solution.** There is no *a priori* reason why the angular velocity ω^\sim should be aligned with a principal moment of inertia at the birth of a pulsar. Were this to be the case, the motion of the pulsar would be trivial. For the solution to equations 3.4.1 - 3.4.3, ω^\sim is generalised to lie at an angle β to \mathbf{e}_3 . Solving

$$(3.4.4) \quad \dot{\omega}_1 = -\varepsilon\omega_3\dot{\omega}_2$$

$$(3.4.5) \quad \Rightarrow \omega_1 = |\omega| \sin \beta \cos(\varepsilon\omega_3 t + \xi)$$

$$(3.4.6) \quad \omega_2 = |\omega| \sin \beta \sin(\varepsilon\omega_3 t + \xi)$$

$$(3.4.7) \quad \omega_3 = |\omega| \cos \beta = \omega_3(0)$$

The projection of ω^\sim onto the axis of symmetry \mathbf{e}_3 is constant. The magnitude of ω^\sim is also constant, and so the body spins with constant speed at fixed angle β to \mathbf{e}_3 . However, the direction of ω^\sim is not constant since $\omega_\perp = \omega_1\mathbf{e}_1 + \omega_2\mathbf{e}_2$ rotates in the $\mathbf{e}_1 - \mathbf{e}_2$ plane. As seen from the body frame, ω^\sim precesses with retrograde motion about \mathbf{e}_3 on a cone at half angle β to \mathbf{e}_3 . This is called the body frame precession.

In terms of the Euler angles, $\omega^\sim = \varphi\tilde{\mathbf{E}}_3 + \theta\tilde{\mathbf{e}}_1 + \psi\mathbf{e}_3$ (See appendix A.2). In an inertial space frame with basis vectors $\tilde{\mathbf{E}}_i$, ω^\sim has the form

$$\omega^\sim = -\varepsilon\omega_3 \sin \varphi \sin \theta \tilde{\mathbf{E}}_1 + \varepsilon\omega_3 \cos \varphi \sin \theta \tilde{\mathbf{E}}_2 + (\varphi - \varepsilon\omega_3 \cos \theta) \tilde{\mathbf{E}}_3$$

where the angles φ , θ and ψ are the standard Euler angles, commonly referred to as the precession angle, the nutation angle and the spin angle respectively.

To an observer in the space frame, the angular momentum vector \mathcal{J}^\sim appears fixed in space. The vector \mathbf{e}_3 precesses about \mathcal{J}^\sim on a cone of half angle θ at an angular frequency φ . This is the inertial frame precession. ω^\sim is in the plane of \mathcal{J}^\sim and \mathbf{e}_3 and so precesses about \mathcal{J}^\sim on a cone of half angle $\beta - \theta$ at an angular frequency φ . Figure 3.4.1 is an illustration of the body and space cones. The intersection of the two cones is aligned with ω . During the precessing motion of the pulsar, the body cone will roll along the outside of the space cone without slipping. Thus, two precessional motions occur. Firstly, the body frame precession, which is the retrograde motion of ω^\sim about \mathbf{e}_3 . Secondly, the inertial frame precession, which is the motion of \mathbf{e}_3 about $\tilde{\mathcal{J}}$.

3.5. Asymmetric Pulsar

Choose the body axes \mathbf{e}_i such that $I_3 > I_2 > I_1$. Defining $\varepsilon_1 = \frac{I_3 - I_2}{I_1}$, $\varepsilon_2 = \frac{I_3 - I_1}{I_2}$ and $\varepsilon_3 = \frac{I_2 - I_1}{I_3}$, equations 3.3.1 - 3.3.3 become

$$\dot{\omega}_1 + \varepsilon_1\omega_2\omega_3 = 0$$

$$\dot{\omega}_2 - \varepsilon_2\omega_1\omega_3 = 0$$

$$\dot{\omega}_3 + \varepsilon_3\omega_1\omega_2 = 0$$

The body axes \mathbf{e}_1 and \mathbf{e}_3 are known as the extreme axes since they are aligned with the minimum and maximum moments of inertia I_1 and I_3 respectively. \mathbf{e}_2 is known as the intermediate axis.

3.5.1. General Solution. For $\boldsymbol{\omega}(0) \approx \omega_3(0)\mathbf{e}_3$ or $\boldsymbol{\omega}(0) \approx \omega_1(0)\mathbf{e}_1$, the solution to Euler's equations approximate that of the axisymmetric case, with a small amplitude oscillation in the solutions of ω_3 or ω_1 respectively. There is relatively little transfer of angular momentum between the angular velocity components, and so the body is defined to be in a stable rotational state.

Assume $\boldsymbol{\omega}(0) \approx \omega_2\mathbf{e}_2$ with a slight perturbation of ω_1 and ω_3 . For $\boldsymbol{\omega}(0) \approx \omega_2(0)$, $\dot{\omega}_2(0) = \varepsilon_2\omega_1(0)\omega_3(0) \approx 0$. Attempting a trial solution

$$\begin{aligned} \omega_1 &= A_1 e^{\lambda t} \\ \omega_3 &= A_3 e^{\lambda t} \\ \Rightarrow \lambda A_1 + \varepsilon_1 \omega_2 A_3 &= 0 \\ \lambda A_3 + \varepsilon_3 \omega_2 A_1 &= 0 \end{aligned}$$

and writing the ratio

$$\begin{aligned} \frac{A_3}{A_1} &= \frac{\lambda}{\varepsilon_1 \omega_2} = \frac{\varepsilon_3 \omega_2}{\lambda} \\ \Rightarrow \lambda &= +\omega_2 \sqrt{\varepsilon_1 \varepsilon_3} \\ \Rightarrow A_3 &= +A_1 \frac{\varepsilon_3}{\varepsilon_1} \end{aligned}$$

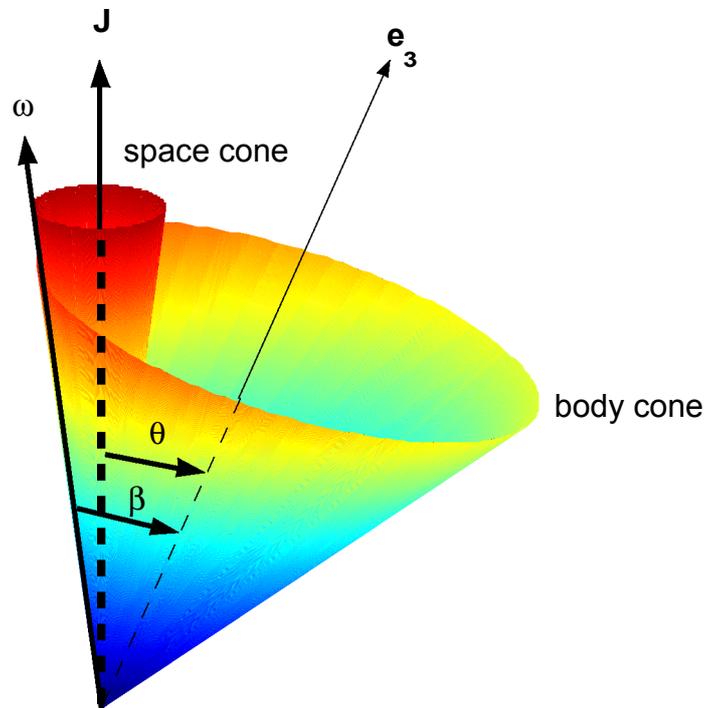


FIGURE 3.4.1. Space and body cone. The precessional motion of the pulsar is described by the larger body cone rolling along the outside of the smaller space cone without slipping. The line of intersection is aligned with $\boldsymbol{\omega}$.

The general solution is the superposition of these two solutions

$$\begin{aligned} \Rightarrow W_1 &= C_1 e^{\omega_2 \sqrt{E_1 E_3} t} + C_2 e^{-\omega_2 \sqrt{E_1 E_3} t} \\ W_3 &= \frac{\epsilon_3}{\epsilon_1} C_1 e^{\omega_2 \sqrt{E_1 E_3} t} + C_2 e^{-\omega_2 \sqrt{E_1 E_3} t} \end{aligned}$$

where $C_1 = \frac{1}{2} (\sqrt{\frac{E_1}{E_3}} W_3(0) - W_1(0))$ and $C_2 = \frac{1}{2} (\sqrt{\frac{E_1}{E_3}} W_3(0) + W_1(0))$.

Since $\epsilon_1, \epsilon_3 > 0$, W_1 and W_3 are a superposition of exponentially increasing and decreasing functions. These solutions only hold for $\omega_1 \omega_3$ small. By the conservation of angular momentum, W_2 decreases with the increase of these angular velocity components. A large transfer of angular momentum occurs, making the rotational state unstable. Once W_1 and W_3 reach a maximum, the cycle is reversed and W_2 increases with the decay of W_1 and W_3 . The unstable rotation represents a tumbling action of the pulsar.

A more complete solution by Landau and Lifshitz (1960) [94] gives the solutions to the angular velocities of an untorqued asymmetric rigid body as

$$\begin{aligned} W_1 &= \frac{\sqrt{2EI_3 - J^2}}{I_1(I_3 - I_1)} \operatorname{cn}(7) \\ W_2 &= \frac{\sqrt{2EI_3 - J^2}}{I_2(I_3 - I_2)} \operatorname{sn}(7) \\ W_3 &= \frac{\sqrt{J^2 - 2EI_1}}{I_3(I_3 - I_1)} \operatorname{dn}(7) \end{aligned}$$

where E is the total mechanical energy of the system, cn , sn and dn are Jacobian elliptic functions and $7 = t \sqrt{\frac{(I_1 - I_2)(I_2 - 2EI_1)}{I_1 I_2}}$.

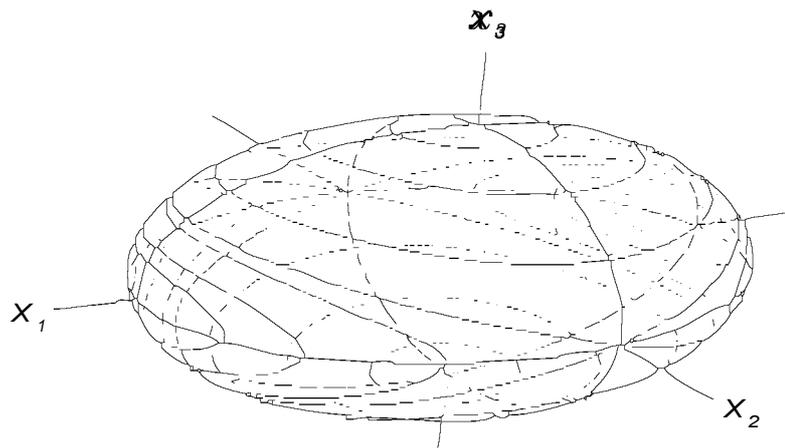


FIGURE 3.5.1. Rotational stability of asymmetric body. The axes \mathbf{x}_i are aligned with the body axes \mathbf{e}_i . The path traced out by the angular momentum vector \mathbf{J} is plotted. The extreme axes \mathbf{x}_1 and \mathbf{x}_3 are attractor points for \mathbf{J} .

Figure 3.5.1 describes the motion of the angular momentum vector \mathbf{J} relative to the moments of inertia (adapted from Landau and Lifshitz (1960) [94]). The axes \mathbf{n}_i are aligned with the body axes \mathbf{a}_i . When near \mathbf{a}_1 or \mathbf{a}_3 , which act as attractor points in the configuration phase space, \mathbf{J} precesses about the respective axes. The intermediate axis \mathbf{a}_2 is not an attractor and so, given an appropriate initial condition, \mathbf{J} passes through \mathbf{a}_2 , travelling a large distance between the corresponding poles. This behaviour is descriptive of the nature of stability of the rigid body about its respective moments of inertia where small deviations about \mathbf{a}_1 and \mathbf{a}_3 result in a stable rotational state, but deviations about \mathbf{a}_2 result in large variation in the motion of \mathbf{J} .

Since the paths form closed loops, the motion is periodic in the body frame. However, viewing from the inertial space frame, solutions for the Euler angles show no definite periodicity.

3.6. Free Precession in Pulsars?

The observation of free precession requires a re-analysis of the action of the neutron superfluid in the crust-core interface. Jones et al (2001) [83] argues that, in order to be consistent with precessing candidates, a crust only precession requires 10^{-10} of the total stellar mass to be pinned superfluid. If one considers a single precessing crust-core component, then this factor increases to 10^{-8} . Both of these quantities are reasonable even though they are orders of magnitude less than those predicted by glitch theories [157].

Free precession models have been used to fit the observed timing residuals of PSR B1642-03, where Shabanova et al (2001) [156] were able to recover smooth, cyclic variations. The 1000 day periodicity in brightness of the pulse profile leading component, reported by Blaskiewicz (1991) [28], is evidence for a free precession model, which is fitted with $\theta(0) \approx 0.5$ rad.

PSR 1828-11 has been modeled as a freely precessing, asymmetric pulsar [164] with periods of 250, 500 and 1000 days. The absence of damping due to vortices in the superfluid interior is explained by a strong decoupling of the crust and core [101]. The pulse profile and amplitude changes in the period derivative support a variation of approximately 0.3 of the angle x between \mathbf{m} and \mathbf{a}_3 (with $x = 60^\circ$).

In his thesis, Chukwude (2002) [37] presents evidence for free precession in PSR B1323-58 and PSR B1358-63, with PSR B1323-58 described by two harmonically related sine functions of periods ≈ 2400 and 1200 days. Timing residuals for PSR B1358-63 are fitted with three sinusoids with periods of ≈ 3600 , 1900 and 1000 days, which could be harmonically related given the quoted errors.

Using photometric results, a simple free precession model has been fitted to both the Crab pulsar and Her X-1 neutron star with periods of 60 seconds and 35 days respectively [31, 32]. A powerful test of the model is the observed proportionality between the increases in precession period and rotation period of the Crab pulsar [32]. It is arguable, however, that given an intrinsic slow down in the pulsar rotation rate, whether the precession can strictly be considered to be torque free.

3.7. Summary

The free precession model of axisymmetric bodies is the simplest of precessional behaviour, with the precession period dependent on the mass distribution of the pulsar. The two rotational states of the asymmetric model may provide an explanation of observed anomalous timing residual behaviour, with an asymmetric free precession model having been fitted to PSR 1828-11 [164]. The unstable rotational state of an asymmetric body results in pseudo-glitch phenomenon. This would probably be accompanied with a periodic disappearance of the pulsar signal, as it is highly unlikely that the beam would be visible for the entire duration of the precession cycle. Whilst there are no known cases of this occurring, suggesting small precession angles in all cases, this may be due to selection effects where pulsar surveys have optimised sensitivities to non-precessing objects.

The analysis of precessing pulsars in terms of an analytic framework has laid the basis for an investigation into the dynamics of precession and its dependencies on relevant physical parameters. The method used to determine analytic solutions for free precession is adopted in the treatment of torqued precession in chapter 4.

CHAPTER 4

Torqued Precession

4.1. Introduction

Pulsars are not torque free, but instead have a number of mechanisms that could account for a loss of angular momentum. They include magnetic dipole (and higher order multi-pole) radiation, magnetospheric drag and gravitational radiation. Long term periodicities in timing residuals determined from HartRAO data raise the possibility of torqued precession being responsible for the observed periodic behaviour.

The analysis of a magnetic dipole model, with orthogonal components in the near field radiation zone, shows that precession may be a requirement in order for a pulsar to reach a local energy eigenstate. If the distortions of the magnetic field are larger than that of the crust, then a steady state solution for the rotation of the star is highly unlikely [175]. For an untorqued body, the global minimum energy eigenstate occurs when the angular velocity and angular momentum are aligned with the largest principal moment of inertia. However, the action of orthogonal components in the magnetic field of a torqued body tends to cause alignment of the angular velocity with the magnetic axis. Therefore, in the case of an axisymmetric or asymmetric body, the coupling between the magnetic axis and angular momentum results in the local energy eigenstate displaying precession. Whilst the pulsar may reach this local state in a reasonably short time, it will tend to move towards a global energy eigenstate of an aligned rotator as the pulsar ages. The effect of radiative precession has been implicated in the bumpy spin down of magnetars [125], although evidence for similar features in the spin down of ordinary radio pulsars is inconclusive.

The effect of a magnetic stresses, which in general do not act in the same direction as crustal stresses, is such that the rotational behaviour of a torqued axisymmetric body resembles that characterised by an asymmetric body, resulting in incoherent periodic timing residuals.

The following sections outline the mathematical formulism for a spinning magnetic dipole, and then solve the angular velocities with a generalised $\mathbf{v}(0)$ for a spherical, axisymmetric and asymmetric pulsar. Where analytical solutions are not possible, numerical simulations are used to examine the rotational behaviour.

4.2. Dipole Torque

The electromagnetic torque resulting from a spinning magnetic dipole is commonly used in the analysis of pulsar spin dynamics.

Consider a pulsar with magnetic axis \mathbf{M} fixed in the rigid crust. Choose a set of body axes \mathbf{e}_i ($i = 1, 2, 3$) such that \mathbf{M} lies in the $\mathbf{e}_1 - \mathbf{e}_3$ plane at an angle X to \mathbf{e}_3 . This construction

involves a slight loss of generality. For both spherical and axisymmetric bodies, where \mathbf{e}_3 is made to lie along the axis of symmetry, an infinite number of axes can be chosen, rotated about \mathbf{e}_3 , such that the basis vectors \mathbf{e}_2 lie along the principal moments of inertia. \mathbf{e}_2 can always be rotated such that $m\sim$ lies in the $\mathbf{e}_1 - \mathbf{e}_3$ plane without loss of generality. In the asymmetric case, however, there is a unique set of axes that lie along the principal moments of inertia and $m\sim$ is not constrained to lie in the $\mathbf{e}_1 - \mathbf{e}_3$ plane. For a number of cases the Euler equations are intractable and generalising $m\sim$ further increases complexity. However, restricting $m\sim$ to lie in the $\mathbf{e}_1 - \mathbf{e}_3$ plane will not cause the solutions to deviate too far from the generalised case.

A spinning magnetic dipole has a near field and far field radiation zone. The Euler equations are written separately with near field torque components:

$$(4.2.1) \quad \dot{\omega}_1 + \frac{(\varepsilon - \varepsilon^0)}{1 + \varepsilon^0} \omega_2 \omega_3 = \mathbf{KG}(x_0) |\omega| \omega_2 (\omega_1 \sin \chi + \omega_3 \cos \chi) \cos \chi$$

$$(4.2.2) \quad \dot{\omega}_2 - \frac{\varepsilon}{1 + \varepsilon^0} \omega_1 \omega_3 = \mathbf{KG}(x_0) \frac{1}{1 + \varepsilon^0} |\omega| (-\omega_1 \cos \chi + \omega_3 \sin \chi)$$

$$(4.2.3) \quad \times (\omega_1 \sin \chi + \omega_3 \cos \chi)$$

$$(4.2.4) \quad \dot{\omega}_3 + \frac{\varepsilon^0}{1 + \varepsilon} \omega_1 \omega_2 = -\mathbf{KG}(x_0) \frac{1}{1 + \varepsilon} |\omega| \omega_2 (\omega_1 \sin \chi + \omega_3 \cos \chi) \sin \chi$$

and with far field torque components:

$$(4.2.5) \quad \dot{\omega}_1 + \frac{(\varepsilon - \varepsilon^0)}{1 + \varepsilon^0} \omega_2 \omega_3 = \mathbf{KF}(x_0) |\omega|^2 (-\omega_1 \cos \chi + \omega_3 \sin \chi) \cos \chi$$

$$(4.2.6) \quad \dot{\omega}_2 - \frac{\varepsilon}{1 + \varepsilon^0} \omega_1 \omega_3 = -\mathbf{KF}(x_0) \frac{1}{1 + \varepsilon^0} |\omega|^2 \omega_2$$

$$(4.2.7) \quad \dot{\omega}_3 + \frac{\varepsilon^0}{1 + \varepsilon} \omega_1 \omega_2 = -\mathbf{KF}(x_0) \frac{1}{1 + \varepsilon} |\omega|^2 (-\omega_1 \cos \chi + \omega_3 \sin \chi) \sin \chi$$

where $\mathbf{K} = \frac{27r_0^2 B_0^2}{\mu_0 c^3}$, $F(x_0) = \frac{x_0^4}{5(x_0^6 - 3x_0^4 + 36)} + \frac{1}{3(x_0 + 1)}$, $G(x_0) = \frac{3(x_0^2 + 6)}{5x_0(x_0^6 - 3x_0^4 + 36)} + \frac{-2x_0^2}{15x_0(x_0 + 1)}$, $\varepsilon = \frac{r_0}{r}$, $\varepsilon^0 = \frac{r_0}{r_0}$ [126]. \mathbf{K} is used as a scaling factor for the timescales of numerical simulations found later in this chapter. The functions $F(x_0)$ and $G(x_0)$ reflect the structure of the far field and near field radiation zones respectively [55], with $x_0 = \frac{r_0 \omega}{c}$. The far field torque components are responsible for the familiar braking torque, whilst the near field torque components are responsible for a precessional motion that is independent of the mass distribution of the pulsar. r_0 is the ‘vacuum radius’, defined by Melatos (1997) [124] to be the innermost point within the magnetosphere where the plasma can become three-dimensional and field-aligned flow breaks down. This radius is generally much larger than $\frac{r_0}{c}$ the radius of the pulsar, but is somewhat less than the light-cylinder radius, defined as $r_L = \frac{c}{\omega}$.

In representing the pulsar as a rotating dipole, it is assumed that the star is internally frictionless. The pinning and unpinning action of superfluid vortices is therefore neglected. We assume that the star rotates *in vacuo*.

4.2.1. Effect of Orthogonal and Spin Down Torque. The flow of energy for a rotating magnetic dipole is calculated by the Poynting vector $\mathbf{E}\sim \times \mathbf{B}\sim^*$. The magnetic and electric fields

\tilde{B} and \tilde{E} are written in spherical polar coordinates as

$$\begin{aligned} \tilde{B} &= -\frac{k^3 m_0}{4\pi\epsilon_0 c^2} \frac{2i}{(kr)^2} - \frac{2}{(kr)^3} \cos\theta \hat{r} + \frac{1}{kr} + \frac{i}{(kr)^2} - \frac{1}{(kr)^3} \sin\theta \hat{\theta} e^{i(kr-\omega t)} \\ \tilde{E} &= \frac{\mu_0 k^2 \omega m_0}{4\pi} \left[\frac{1}{kr} + \frac{i}{(kr)^2} \sin\theta \right] e^{i(kr-\omega t)} \hat{\psi} \end{aligned}$$

[174], resulting in the r dependence of the time averaged Poynting vector, with constants omitted, to be

$$\tilde{E} \times \tilde{B}^* \propto \left[\frac{1}{(kr)^2} + \frac{i}{(kr)^5} \sin^2\theta \hat{r} - \frac{i}{(kr)^3} + \frac{i}{(kr)^5} \sin 2\theta \hat{\theta} \right]$$

As only the real parts contribute to the time averaged flow of energy, the terms proportional to $\frac{1}{(kr)^2}$ represent the spin down torque. These terms describe the far field radiation zone and are parallel with the radius vector \hat{r} . The imaginary components contribute to the instantaneous value of the Poynting vector and are proportional to the near field radiation zone. As there is no mechanism for energy absorption in the near field zone, the net average flow is zero (the zone is non-radiative) but the instantaneous form of the terms give rise to near field components that oscillate as $e^{-i\omega t}$.

4.3. Spherical Pulsar

The body frame basis vectors \mathbf{e}_i are chosen such that the magnetic axis m is aligned with \mathbf{e}_3 , making $\chi = 0$ for all cases. Equations 4.2.1 - 4.2.4 and 4.2.5 - 4.2.7, with $\mathbf{l}_1 = \mathbf{l}_2 = \mathbf{l}_3$, simplify to

$$\begin{aligned} \dot{\omega}_1 &= KG(x_0)|\omega|\omega_2\omega_3 - KF(x_0)|\omega|^2 \omega_1 \\ \dot{\omega}_2 &= -KG(x_0)|\omega|\omega_1\omega_3 + KF(x_0)|\omega|^2 \omega_2 \\ \dot{\omega}_3 &= 0 \end{aligned}$$

Consider a generalised $\omega(0)$ such that $\omega_1(0) \notin \omega_2(0) \notin \omega_3(0) \notin 0$. The solutions for the angular velocities are affected by a radiating component (far field torque) and a non-radiating component (near field torque). As these torques are orthogonal and will act over different timescales, ω_i is solved for each torque regime separately.

4.3.1. General Solution: Near Field. $|\omega|$ is shown to be constant

$$\frac{d}{dt}(\omega_1^2 + \omega_2^2 + \omega_3^2) = \frac{d|\omega|^2}{dt} = 0$$

This is true for all near field solutions. For $\dot{\omega}_3 = 0 \Rightarrow \omega_3 = \omega_3(0)$, and ω_1 and ω_2 display a quadrature relation

$$\begin{aligned} \dot{\omega}_1 &= KG(x_0)|\omega|\omega_3\dot{\omega}_2 \\ &= -(KG(x_0)|\omega|\omega_3)^2 \omega_1 \\ \Rightarrow \omega_1 &= A \cos(KG(x_0)|\omega|\omega_3 t + \xi) \\ \omega_2 &= A \sin(KG(x_0)|\omega|\omega_3 t + \xi) \end{aligned}$$

where $A = |\omega| \sin \beta$ is a constant of the motion. Under the action of far field torques, $A = A(t)$ and decays with time.

4.3.2. General Solution: Far Field. With $\dot{\omega}_3 = 0 \Rightarrow \omega_3 = \omega_3(0)$, $|\omega|$ is determined

$$\begin{aligned} \frac{d|\omega|^2}{dt} &= -2KF(x_0)|\omega|^2(\omega_1^2 + \omega_2^2) \\ \text{let } u &= \omega_1^2 + \omega_2^2 \\ \Rightarrow \frac{du}{dt} &= -2KF(x_0)(u + \omega_3^2)u \\ \Rightarrow u &= \frac{\omega_3^2}{-1 + C_1 \omega_3^2 e^{2KF(x_p)\omega_3^2 t}} \\ \Rightarrow |\omega| &= \frac{\omega_3}{\sqrt{1 + C_1 \omega_3^2 e^{2KF(x_p)\omega_3^2 t}}} \end{aligned}$$

Substituting u into $\dot{\omega}_1$ and $\dot{\omega}_2$ and solving for ω_1 and ω_2

$$(4.3.1) \quad \omega_1 = \frac{C_2}{\sqrt{C_1 \omega_3^2 e^{2KF(x_p)\omega_3^2 t} - 1}}$$

$$(4.3.2) \quad \omega_2 = \frac{C_3}{\sqrt{C_1 \omega_3^2 e^{2KF(x_p)\omega_3^2 t} - 1}}$$

where

$$\begin{aligned} C_1 &= \frac{|\omega(0)|^2}{\omega_3^2(0)u(0)} \\ C_{2,3} &= \omega_{1,2}(0) \sqrt{C_1 \omega_3^2(0) - 1} \end{aligned}$$

4.3.3. General Solution: Dynamics. The orthogonal components of the near field radiation zone cause ω to precess about m in a manner that is indistinct from that precessional motion caused by the axisymmetric mass distribution of a rigid body. The form of solutions under the action of both torques are the same as the exact solutions determined by Davis and Goldstein (1970) [52].

There are radiative mechanisms not predicted by the rotating dipole model, which may lead to a solution for an aligned rotator where $\omega_3 \neq \omega_3(0)$. Consider the case of an aligned rotator with $\omega = \omega_3 \mathbf{e}_3$. No braking torque exists according to standard rotating dipole models. However, the action of the magnetic field lines, which twist above the polar caps, form a spiraling helix. The twisting helix structure, in the case of a spherical or axisymmetric rotating body, can lead to a continuous outflow of energy, angular momentum and magnetic flux [100, 108]. This radiation should have consequences on the form of braking law used. The twisting structure may cause inflation of the field lines, decreasing the magnitude of the initial magnetic field [1]. The inflation, and subsequent deflation, of magnetic field lines causes oscillatory behaviour in the braking torque, whilst higher order multi-poles would become more prominent in the dipolar magnetic field structure.

The precession period is dependent on $|\omega|$ under the action of far field torques, and increases as

$$\Delta P_{\text{prec}} = \frac{2\pi C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} E(x_0) \Delta t}{-1 + C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} G(x_0)}$$

The precession angle β , the angle between the vectors \mathbf{e}_3 and $\boldsymbol{\omega}$, evolves as

$$\beta = \arccos \frac{\omega_3}{|\omega|} = \arccos \left(\frac{\omega_3}{C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} - 1} \right) = \alpha$$

where α is the angle between m and $\boldsymbol{\omega}$.

As $\chi = 0$ for all cases, $\alpha = \beta \rightarrow 0$ for $t \rightarrow \infty$. Thus, given any set of initial conditions, the spherical pulsar tends toward an aligned rotator. Whilst it may appear to mimic the behaviour of the untorqued axisymmetric system, the characteristic time constant is different.

4.4. Axisymmetric Pulsar

Choose the body frame basis vectors \mathbf{e}_i such that \mathbf{e}_3 is aligned with the axis of symmetry. With $I_3 > I_1 = I_2 \Rightarrow \varepsilon = 0$, equations 4.2.1 - 4.2.4 and 4.2.5 - 4.2.7 become

$$(4.4.1) \quad \dot{\omega}_1 + \varepsilon \omega_2 \omega_3 = KG(x_0) |\omega| \omega_2 \cos \chi (\omega_1 \sin \chi + \omega_3 \cos \chi)$$

$$(4.4.2) \quad \dot{\omega}_2 - \varepsilon \omega_1 \omega_3 = -KG(x_0) |\omega| (\omega_1 \cos \chi - \omega_3 \sin \chi) (\omega_1 \sin \chi + \omega_3 \cos \chi)$$

$$(4.4.3) \quad \dot{\omega}_3 = -KG(x_0) \frac{1}{1 + \varepsilon} |\omega| \omega_2 \sin \chi (\omega_1 \sin \chi + \omega_3 \cos \chi)$$

and

$$(4.4.4) \quad \dot{\omega}_1 + \varepsilon \omega_2 \omega_3 = -KF(x_0) |\omega|^2 \cos \chi (\omega_1 \cos \chi - \omega_3 \sin \chi)$$

$$(4.4.5) \quad \dot{\omega}_2 - \varepsilon \omega_1 \omega_3 = -KF(x_0) |\omega|^2 \omega_2$$

$$(4.4.6) \quad \dot{\omega}_3 = KF(x_0) \frac{1}{1 + \varepsilon} |\omega|^2 \sin \chi (\omega_1 \cos \chi - \omega_3 \sin \chi)$$

respectively. General solutions to Euler's equations are solved for $\chi = 0$, $\chi = \frac{\pi}{2}$ and $0 < \chi < \frac{\pi}{2}$ below.

4.4.1. General Solution: Near Field ($\chi = 0$). Solutions to Euler's equations are

$$\omega_1 = A \cos([KG(x_0) |\omega| - \varepsilon] \omega_3 t + \xi)$$

$$\omega_2 = A \sin([KG(x_0) |\omega| - \varepsilon] \omega_3 t + \xi)$$

$$\omega_3 = \omega_3(0)$$

The precession periodicity is dependent on both the magnetic and crustal stresses.

4.4.2. General Solution: Far Field ($\chi = 0$). Solutions to Euler's equations are

$$\dot{\omega}_3 = 0 \Rightarrow \omega_3 = \omega_3(0)$$

while $|\omega|$ is given by

$$|\omega| = \frac{\omega_3}{\sqrt{C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} - 1}}$$

where

$$C_1 = \frac{\omega_2^2(0) + |\omega(0)|^2}{\omega_3^2(0)|\omega(0)|^2}$$

ω_1 and ω_2 have no analytical solution. The rotational behaviour of the aligned axisymmetric pulsar is similar to the spherical pulsar since the magnetic and crustal stresses act in the same direction. We can therefore expect similar behaviour in the solutions for the decay of ω_1 and ω_2 as described by equations 4.3.1 and 4.3.2, although they will oscillate in quadrature due to the action of crustal and magnetic stresses.

4.4.3. General Solution: Dynamics ($\chi = 0$). The total mechanical energy of the pulsar decays as $|\omega|$. With ω_3 constant, the precession period is dependent on $|\omega|$ and increases according to

$$P_{\text{prec}} = \frac{2\pi \Theta C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} - 1}{[KG(x_0)\omega_3 - \epsilon \Theta C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} - 1]\omega_3}$$

The behaviour of the precession angle β ($= \alpha$ for $\chi = 0$) indicates that the pulsar becomes an aligned rotator

$$\beta = \alpha = \arccos(\Theta C_1 \omega_3^2 e^{2KF(x_0)\omega_3^2 t} - 1)$$

with $\omega(t) \rightarrow \omega_3(0)\mathbf{e}_3$ as $t \rightarrow \infty$. The effect of the axisymmetric mass distribution and near field radiation zone causes ω to precess about \mathbf{e}_3 (since $m^* = m\mathbf{e}_3$). Although the precessional behaviour is similar to that of the spherical pulsar, the precession period is a function of both $G(x_0)$ and ϵ .

Since $P_{\text{prec}} \propto \frac{1}{|\omega|}$, in the limit $\chi \rightarrow 0$ the torqued axisymmetric system would satisfy a test for precession of the Crab pulsar where an observed 60 second precession periodicity is said to be increasing at the same rate as the rotation period [32, 31]. This would require, however, a disproportionately high magnetic field strength for the near field radiation zone.

4.4.4. General Solution: Near Field ($\chi = \frac{\pi}{2}$). The solution for $|\omega|$ is

$$\frac{d|\omega|^2}{dt} = 2KG(x_0)|\omega|\omega_1\omega_2\omega_3 \frac{\epsilon}{1+\epsilon}$$

As $\epsilon \rightarrow 0$, $|\omega|$ remains a constant of the motion. For $\epsilon \neq 0$, $|\omega|$ oscillates, but not necessarily with a well defined period.

ω_1 , ω_2 and ω_3 have solutions that exhibit quadrature phase relationships.

$$\begin{aligned} \dot{\omega}_1 &= -\epsilon(\omega_2\dot{\omega}_3 + \dot{\omega}_2\omega_3) \\ &= \epsilon KG(x_0)|\omega| \frac{1}{1+\epsilon} \omega_2^2 - \omega_3^2 - \epsilon^2 \omega_3^2 \omega_1 \end{aligned}$$

The term inside the square bracket varies slowly and, to order ϵ , can be approximated as constant

$$\Rightarrow \omega_1 = A \cos \left[\epsilon KG(x_0)|\omega| \frac{1}{1+\epsilon} \frac{\omega_2^2}{\omega_3^2} - \epsilon^2 \omega_3^2 t + \xi \right]$$

In a similar manner,

$$W_2 = B \sin \left(\frac{\varepsilon K G(x_0) I W_1 W_3^2 + \frac{1}{1+\varepsilon} W_1^2 + \frac{1}{1+\varepsilon} (K G(x_0) I W_1)^2 + \varepsilon W_3^2}{\varepsilon K G(x_0) I W_1 W_3^2 + \frac{1}{1+\varepsilon} W_1^2 + \frac{1}{1+\varepsilon} (K G(x_0) I W_1)^2 + \varepsilon W_3^2} t + S \right)$$

$$W_3 = C \sin \left(\frac{\varepsilon K G(x_0) I W_1 (W_1^2 - W_2^2) + \frac{1}{1+\varepsilon} (K G(x_0) I W_1)^2}{\varepsilon K G(x_0) I W_1 (W_1^2 - W_2^2) + \frac{1}{1+\varepsilon} (K G(x_0) I W_1)^2} t + S \right)$$

4.4.5. General Solution: Far Field ($x = \frac{\pi}{2}$). The magnitude of the angular velocity

$$I W_1 = \int e^{-2K F(x_0)} \frac{1}{\left(\frac{1}{1+\varepsilon} \omega_3^2 + \omega_2^2 \right)} dt$$

decays as in previous cases. However, there is no analytical solution because of the indefinite integral. Substituting IWI into equations 4.4.4 - 4.4.6 gives

$$W_1 = -\varepsilon W_2 W_3$$

$$W_2 = -K F(x_0) e^{-2K F(x_0)} \int \left(\omega_2^2 + \frac{1}{1+\varepsilon} \omega_3^2 \right) dt W_2 + \varepsilon W_1 W_3$$

$$W_3 = -K F(x_0) e^{-2K F(x_0)} \int \left(\omega_2^2 + \frac{1}{1+\varepsilon} \omega_3^2 \right) dt W_3$$

4.4.6. General Solution: Dynamics ($x = \frac{\pi}{2}$ and $0 < x < \frac{\pi}{2}$). Euler's equations are coupled in three ways. Crustal stresses couple W_1 and W_2 to order ε and leads to Eulerian precession of the kind seen in a torque free system. Similarly, magnetic stresses in the near field radiation zone couple the angular velocity components, leading to radiative precession. IWI, which in general is not a constant of the motion, couples all the angular velocity components. The two precession mechanisms, under circumstances in which the precessional periods are of the same order, couple and lead to a complicated behaviour.

Figures 4.4.1 and 4.4.2 are numerical simulations of a torqued axisymmetric pulsar with $\varepsilon = 10^{-9}$, $r_0 = 10^5$ [m], $B_0 = 3 \times 10^{10}$ [T], $R = 10^4$ [m], $W_3(0) = 10 W_{1,2}(0)$ [rad.s⁻¹] and $x = \frac{\pi}{2}$ [rad], ensuring sufficient coupling of the Eulerian and radiative precession. To maintain this degree of coupling, any change to the magnetic field is constrained by the proportionality relationship $\frac{\Delta \varepsilon^2}{\varepsilon^2} \propto \frac{\Delta B_0}{B_0}$. Using the listed values, $K \sim 10^{-12}$. If 3×10^{10} [T] is to be considered the upper limit for B_0 , then the lower limit for a coupled precession period is of the order of 1000 days. Several authors have raised the possibility of very strong internal and surface magnetic fields $> 10^{10}$ T [133, 170, 179], and the discovery of the pulsars PSR J1718-37184 and PSR J1847-0130, with surface magnetic fields of almost 10^{10} T, suggests that this magnitude may be realistic [118]. Weaker magnetic fields, with the same oblateness ε , would not have sufficient effect on the precession period, but would still affect the rotational behaviour of the pulsar through the magnetic stresses.

The decay of W_3 has superimposed on it an oscillatory behaviour subject to magnetic stresses. The precession period, which in general will increase, is dependent on W_3 and given by

$$P_{\text{prec}} = \frac{4}{\varepsilon} \int \frac{1}{I W_1^2 - W_1^2 + \frac{\varepsilon K}{\omega G(x_0)} W_3^2} dt \left(\frac{1}{\varepsilon K} \frac{\omega G(x_0) W_1^2 - W_3^2 + I W_1^2}{W_1^2 + W_2^2 + \frac{\varepsilon K}{\omega G(x_0)} W_3^2} \right)$$

where $E(k)$ is a complete elliptic integral of the first kind [126]. This behaviour is observed in figure 4.4.1, which shows the precession period increasing with a slight superimposed oscillation. The behaviour of ω is characteristic of an asymmetric body, with oscillations in all three angular velocity components.

The nature of the precession angle β , shown in figure 4.4.2, is such that ω will precess and nutate about a vector ρ in the $e_1 - e_3$ plane. Defining the angles $\rho - e_3 = \gamma$ and $\rho - \hat{m} = \zeta$, where $\hat{m} = m\hat{m}$, the nature of precession (Eulerian, radiative or coupled) is indicated by $\gamma \ll \zeta$, $\gamma \gg \zeta$ or $\gamma \sim \zeta$ respectively. Since $\gamma = \gamma(\epsilon, K, G(x_0), \omega)$ and $\zeta = \zeta(\epsilon, K, G(x_0), \omega)$, then for $K = K(t)$, $\zeta \neq 0$ for $\frac{dK}{dt} > 0$ and $\gamma \neq 0$ for $\frac{dK}{dt} < 0$.

For $\chi = \frac{\pi}{2}$, the nature of precession is dominated by the ratio $\frac{\epsilon}{KG(x_0)|\omega|}$. For $\frac{\epsilon}{KG(x_0)|\omega|} \gg 1$, Eulerian precession dominates with $\gamma \ll \zeta$ while for $\frac{\epsilon}{KG(x_0)|\omega|} \ll 1$, $\gamma \gg \zeta$ and radiative precession dominates. This does not hold for $0 < \chi < \frac{\pi}{2}$, nor does it hold at the birth of the pulsar when the nature of precession is sensitive to $\omega(0)$, with \hat{m} and e_3 acting as attractors in a configuration phase space.

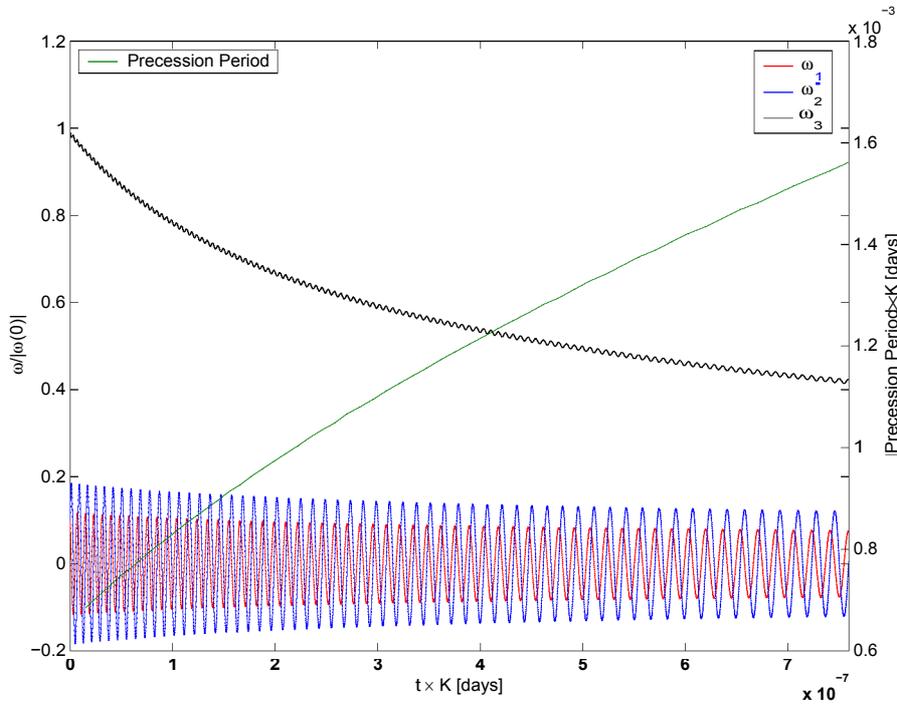


FIGURE 4.4.1. Numerical solution to Euler's equations for axisymmetric pulsar, $\chi = \frac{\pi}{2}$.

In the more general case, $0 < \chi < \frac{\pi}{2}$, analytical solutions cannot be obtained for Euler's equations. Figure 4.4.3 shows numerical solutions of ω_i for $\chi = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$. By increasing χ from 0 to $\frac{\pi}{2}$, the ratio of the Eulerian precession period to that of radiative precession changes, as does the coupling angle (that angle at which the two precessional periods are equal for a generalised $\omega(0)$).

Numerical solutions show that, for $\chi = \frac{\pi}{12}$ and $\frac{\pi}{6}$, the two precessing actions are coupled, with the ratio of the associated periods approaching unity. For $\chi = \frac{\pi}{3}$ and $\frac{5\pi}{12}$, the pulsar behaves

as an asymmetric body dominated by the mass distribution about e_3 . The double maxima in the superimposed oscillation of cv_3 are a result of the offset magnetic axis in, which cv couples to every half precession period. $X = \frac{7r}{4}$ results in a precessional mode change, which is highlighted by a dramatic increase in the precession period. This is followed by a slow return to the precession period function characteristic of that shown for $X = \frac{7r}{3}$ and $\frac{57r}{12}$. The mode change is precipitated by two factors: the increase in $G(x_0)$ (due to a decrease in $|cv|$) and the pulsar evolving towards an aligned rotator. Hence, the pulsar shifts from an Eulerian dominated precession to a coupled, or radiative, dominated precession. This change occurs over a single precession period and could last between a few minutes to a number of years.

Figure 4.4.4 shows the behaviour of α and Q for the same set of Euler solutions. Depending on whether the Eulerian precession and radiative precession are coupled sufficiently ($X = \frac{7r}{12}$ and $\frac{7r}{6}$) or not ($X = \frac{7r}{3}$ and $\frac{57r}{12}$), Q will oscillate between decaying minima and maxima or will exhibit two alternating minima as the amplitude of oscillation increases respectively. The continued oscillation of α and Q indicates that the system will tend towards a minimal state of precession. The progression of the system to become an aligned rotator, $\alpha = 0$, has a complicated dependency on the initial value $cv(0)$, the angle X , the ratio $\frac{G(x_0)}{F(x_0)}$ and K . These dependencies are examined by further simulation.

The rotational behaviour of the pulsar is best described by studying the behaviour of \dot{v} , shown in figure 4.4.5 for $0 < X < \frac{7r}{2}$. The magnetic axis ni is indicated by the red line, which is superimposed on a set of body axes e_i . For all values of X , cv precesses about a vector pi at an angle $-y$ to e_3 , or angle C to in . For $X = \frac{7r}{12}$ or $\frac{7r}{6}$, radiative precession dominates. The behaviour

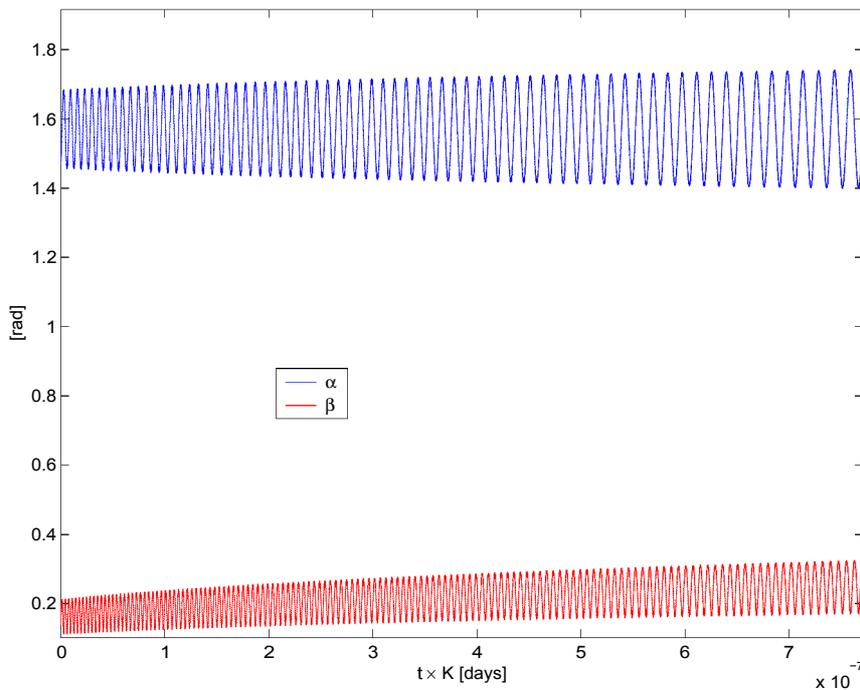


FIGURE 4.4.2. Numerical solution of β and α for axisymmetric pulsar, $X = \frac{7r}{2}$.

of \tilde{w} is characteristic of that of an asymmetric body in an unstable ($X = \frac{\pi}{12}$ and $X = \frac{\pi}{6}$) or stable ($X = \frac{\pi}{3}$ and $\frac{5\pi}{12}$) rotational state (see figure 4.5.5). For $X = \frac{\pi}{3}$ and $\frac{5\pi}{12}$, the system is dominated by Eulerian, or coupled, precession, with $r_y \ll 1$. A dramatic reorientation of \tilde{p} occurs for $X = \frac{\pi}{4}$, due to a shift from Eulerian dominated precession to a coupled or radiative dominated precession. This results in anomalous behaviour in observed timing residuals, and a decrease in the precession period is expected afterwards.

The dependence of the nature of $\tilde{w}(0)$ is due to the numerous energy eigenstates. For $X = \frac{5\pi}{12}$, the system can still behave as one dominated by radiative precession, that is $r_y \gg 1$, if $\alpha(0)$ is sufficiently small. Figure 4.4.6 shows the behaviour of the pulsar for initial conditions described by $w_1(0) = 10w_{2,3}(0)$. Radiative precession dominates for all values of X .

The resultant behaviour due to a change in r_0 , and thus the ratio $\frac{G(x_0)}{F(x_0)}$, is shown in figure 4.4.7, with a tendency for the system to be dominated by Eulerian precession with an increase in r_0 . Since the long term behaviour of the pulsar is such that $|w_j|$, and so x_0 , decreases, the ratio $\frac{G(x_0)}{F(x_0)}$ increases with time. This is to be expected since the near field is non-radiative.

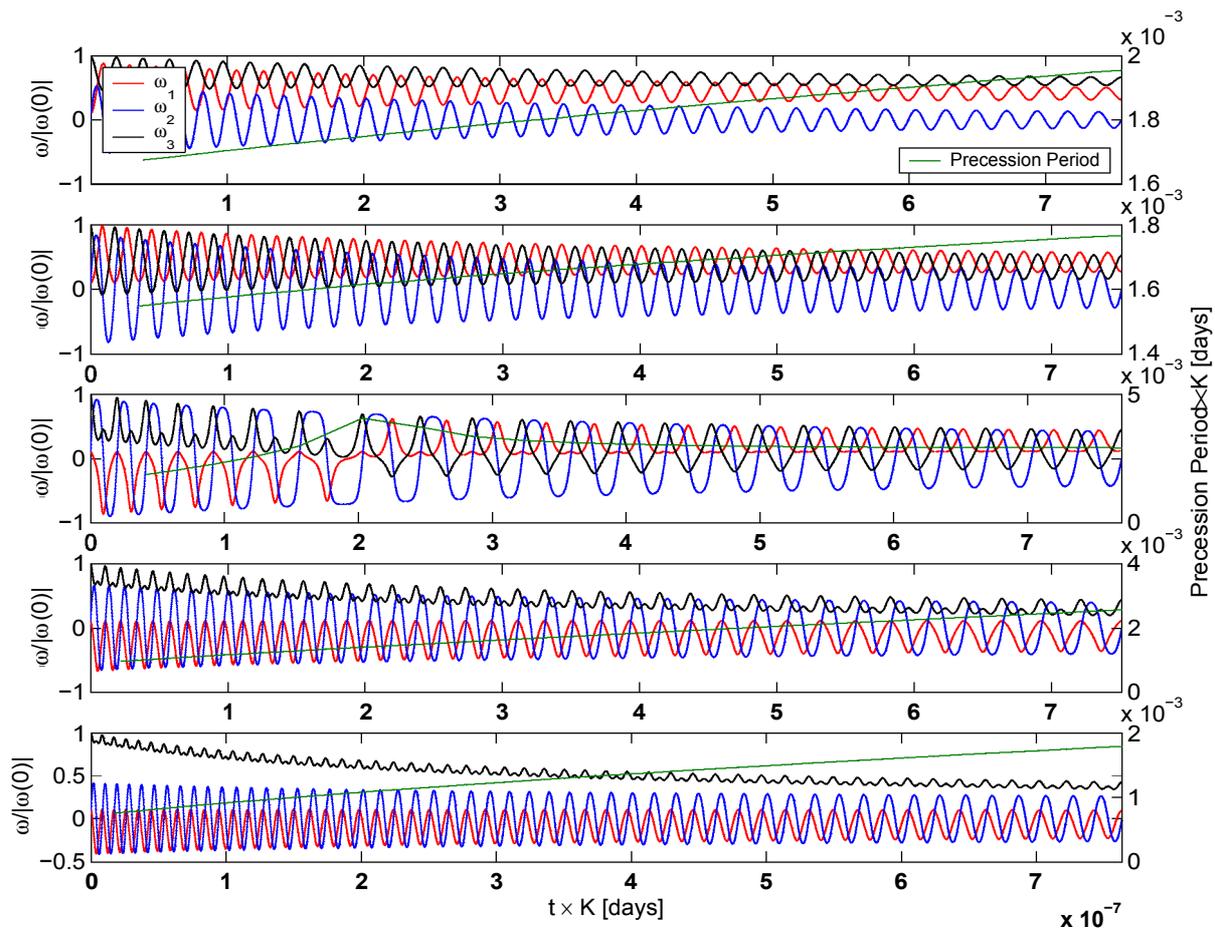


FIGURE 4.4.3. Numerical solution to Euler's equations for axisymmetric pulsar. From top to bottom $X = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$.

4.5. Asymmetric Pulsar

The body frame basis vectors e_i are chosen such that $I_1 < I_2 < I_3$. Euler's equations are given by equations 4.2.1 - 4.2.4 and 4.2.5 - 4.2.7. It is convenient to define the quantities $e_1 = e - e^1$, $e_2 = \frac{\xi}{1+\xi}$, and $e_3 = \frac{\xi^0}{1+\xi}$. Following the solution for an untorqued asymmetric body, equations 4.2.1 - 4.2.4 and 4.2.5 - 4.2.7 are solved using the same initial conditions for $\varphi(0)$ as in section 3.5.

The solutions to the untorqued asymmetric pulsar indicate that stable precession can take place about the maximum and minimum moments of inertia. The angle γ is therefore redefined as the angle between p and that axis φ would preferentially precess about in an untorqued system given the initial condition $\varphi(0)$.

4.5.1. General Solutions ($x = 0$).

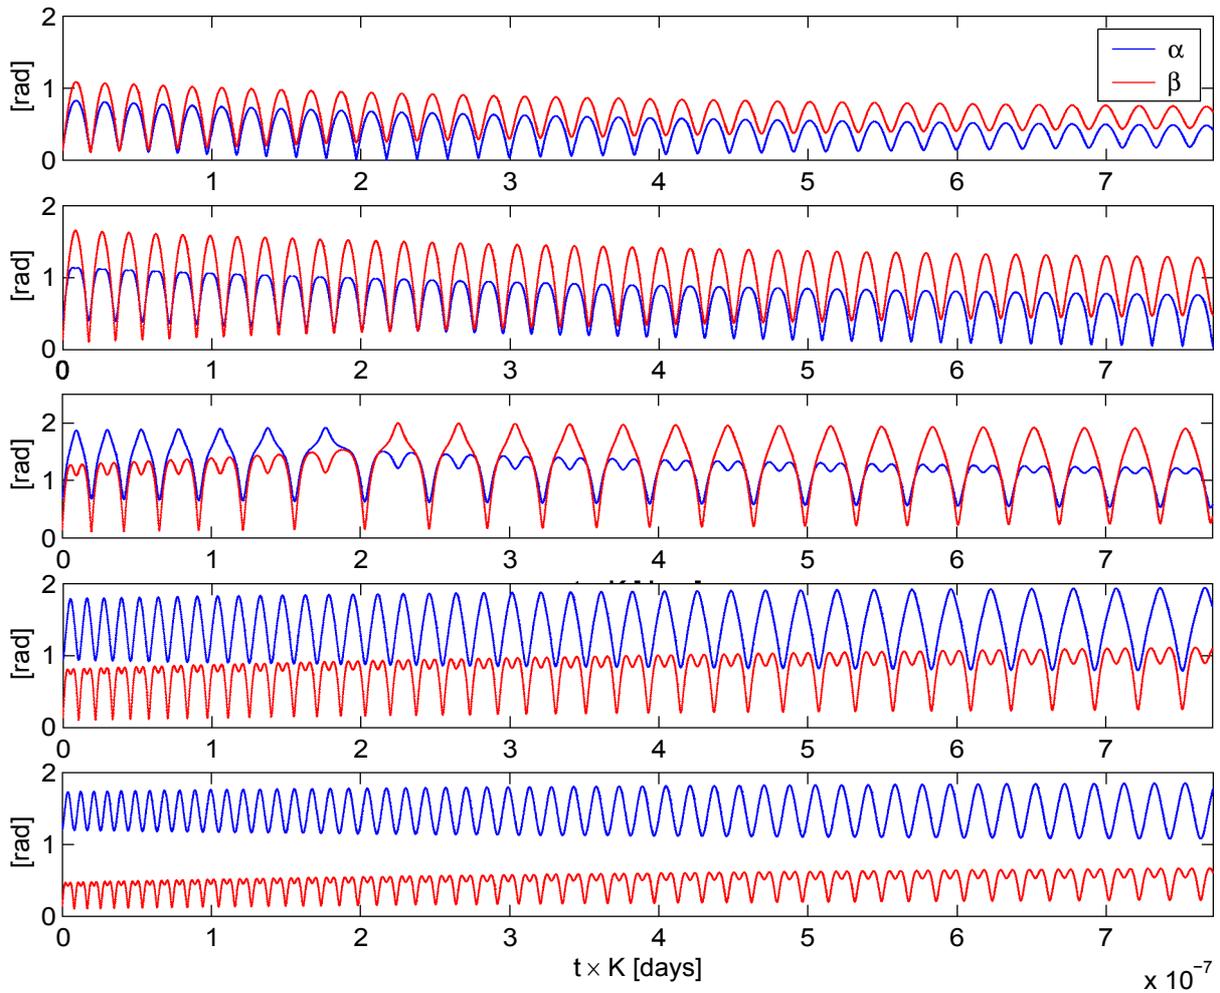


FIGURE 4.4.4. Numerical solution of α and β for axisymmetric pulsar. From top to bottom, $\times = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$.

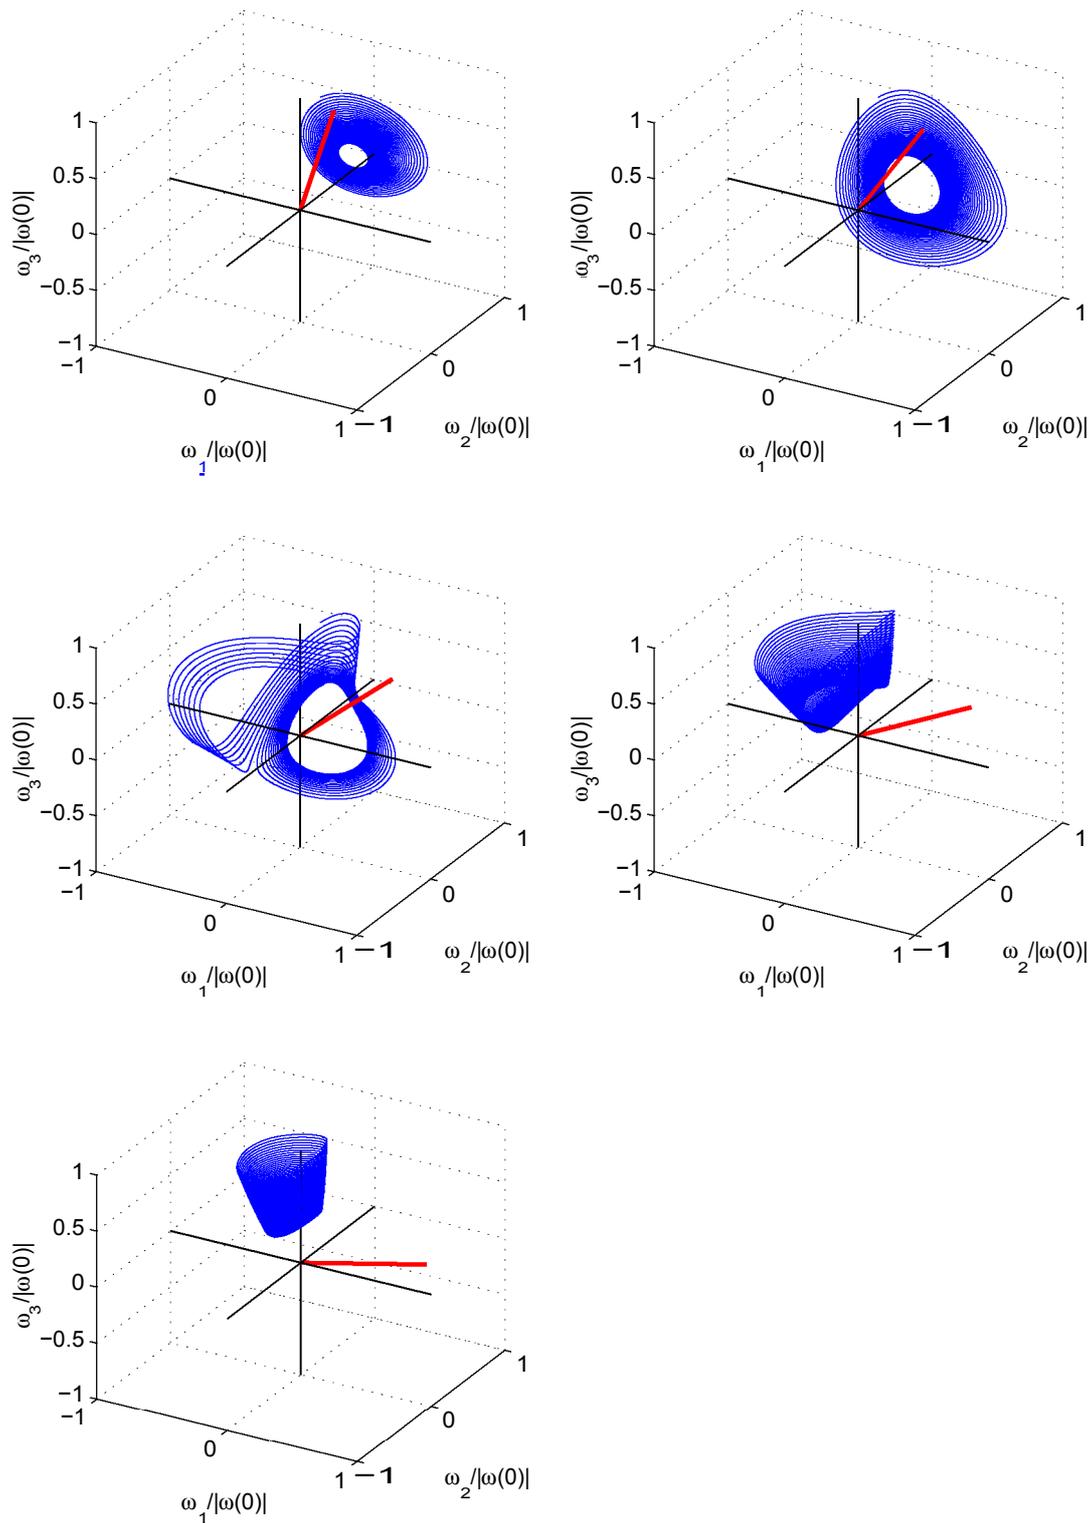


FIGURE 4.4.5. Numerical solution of ω for axisymmetric pulsar, $\omega_3(0) = 10\omega_{12}(0)$. From top to bottom, left to right, $X = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$ (ω indicated by red line).

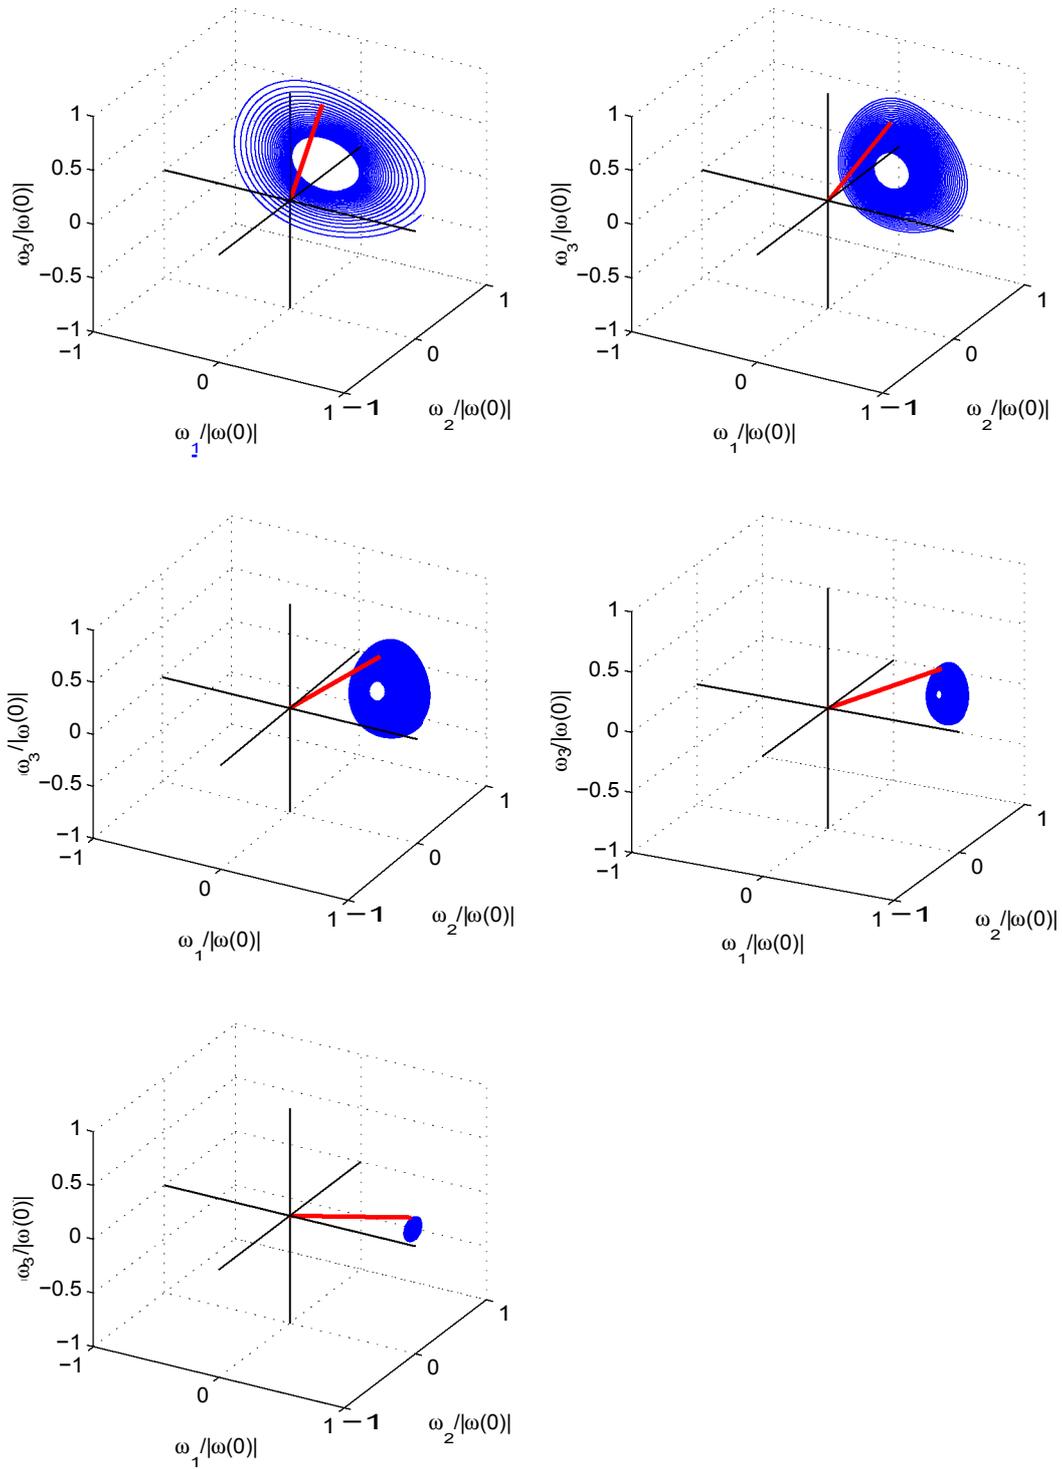


FIGURE 4.4.6. Numerical solution of ω for axisymmetric pulsar, $\omega_1(0) = 10\omega_{2,3}(0)$. From top to bottom, left to right, $\chi = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$ (ω indicated by red line).

4.5.1.1. *Near Field* ($\mathbf{w}(0) \approx w_3(0)\mathbf{e}_3$). The approximation $w_3 \approx w_3(0)$ reduces Euler's equations to

$$\begin{aligned} w_1 + E_1 w_2 w_3 &= \mathbf{KG}(x_0)|w|w_2 w_3 \\ w_2 - E_2 w_1 w_3 &= \frac{1}{1+E} \mathbf{KG}(x_0)|w|w_1 w_3 \\ w_3 &= 0 \end{aligned}$$

Solving

$$\begin{aligned} w_1 &= \frac{\mathbf{KG}(x_0)|w|w_3 - E_1 w_3 w_2}{\mathbf{KG}(x_0)|w|w_3^2 E_2 + \frac{1}{1+E} - \frac{1}{1+E} \mathbf{KG}(x_0)|w| - E_1 E_2} w_1 \\ &\approx \frac{\mathbf{KG}(x_0)|w|w_3 - E_1 w_3 w_2}{\mathbf{KG}(x_0)|w|w_3^2 E_2 + \frac{1}{1+E} - \frac{1}{1+E} \mathbf{KG}(x_0)|w| - E_1 E_2} w_1 \end{aligned}$$

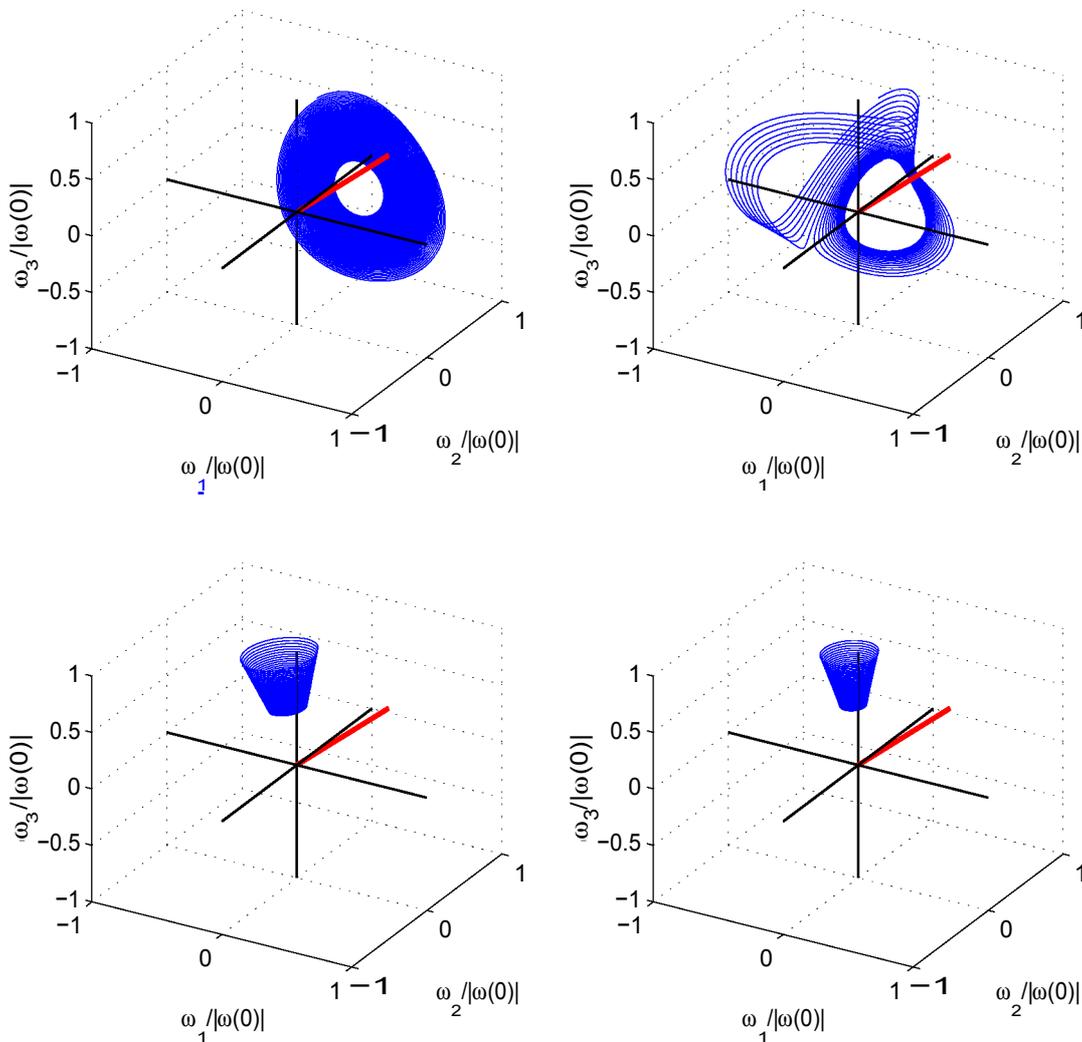


FIGURE 4.4.7. Numerical solution of \mathbf{w}^* for axisymmetric pulsar. From top to bottom, left to right, $r_0 = 5$ km, 10 km, 50 km, 100 km. $\chi = \frac{\pi}{4}$ (m indicated by red line).

to order ε . It may be shown that the terms inside the square brackets are almost constant,

$$\begin{aligned} \Rightarrow \omega_1 &= A \sin\left(\underbrace{KG(x_0)|\omega|}_{\varepsilon_2 + \frac{\varepsilon}{1+\varepsilon}} \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| - \varepsilon_1 \varepsilon_2 \omega_3 t + \xi_1 \right) \\ \dot{\omega}_2 &= \varepsilon_2 \omega_1 \omega_3 \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| \omega_1 \omega_3 \\ &= \varepsilon_2 \omega_3 - \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| \omega_3 \times \\ &\quad \left(\underbrace{KG(x_0)|\omega|}_{\varepsilon_2 + \frac{\varepsilon}{1+\varepsilon}} \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| - \varepsilon_1 \varepsilon_2 \omega_3 t + \xi_1 \right) \\ \Rightarrow \omega_2 &= \left| \frac{1}{KG(x_0)|\omega| \left(\varepsilon_2 + \frac{\varepsilon}{1+\varepsilon} \right) \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| - \varepsilon_1 \varepsilon_2} \right| \times \\ &\quad \left(\underbrace{KG(x_0)|\omega|}_{\varepsilon_2 + \frac{\varepsilon}{1+\varepsilon}} \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| - \varepsilon_1 \varepsilon_2 \omega_3 t + \xi_1 \right) \\ &\quad A \cos\left(\underbrace{KG(x_0)|\omega|}_{\varepsilon_2 + \frac{\varepsilon}{1+\varepsilon}} \frac{1}{1+\varepsilon} \cdot KG(x_0)|\omega| - \varepsilon_1 \varepsilon_2 \omega_3 t + \xi_1 \right) \end{aligned}$$

Relaxing the approximation $\dot{\omega}_3 = 0$, ω_3 is calculated to oscillate slightly.

4.5.1.2. *Far Field* ($\boldsymbol{\omega}(0) \approx \omega_3(0)\mathbf{e}_3$). Solving for $|\omega|$

$$\frac{d|\omega|^2}{dt} = -2 \underbrace{KF(x_0)|\omega|^2}_{\omega_1^2 + \frac{1}{1+\varepsilon} \omega_2^2} + (\varepsilon_1 - \varepsilon_2 + \varepsilon_3) \omega_1 \omega_2 \omega_3$$

The term $(\varepsilon_1 - \varepsilon_2 + \varepsilon_3) \omega_1 \omega_2 \omega_3$ causes the pulsar to precess. This motion, like the action of the near field torques, is non-radiative and acts orthogonal to the radial component of the far field torques. The period of oscillation decays, as well as the amplitude, as a result of terms dependent on $F(x_0)$ only. To obtain this decay envelope, Euler's equations are separated into short term dynamics due to Eulerian precession, and long term spin down.

$$\begin{aligned} \Rightarrow \frac{d|\omega|^2}{dt} &\approx -2KF(x_0)|\omega|^2 \left(\omega_1^2 + \frac{1}{1+\varepsilon} \omega_2^2 \right) \\ \text{let } u &= |\omega|^2 \\ \Rightarrow \frac{du}{dt} &\approx -2KF(x_0)u(\omega_3^2 - u) \frac{\mu}{1 + \frac{1}{1+\varepsilon}} \\ \Rightarrow u &= \frac{\omega_3^2}{1 + C_1 \omega_3^2 e^{\frac{2KF(x_0)(\omega_3^2 - u)}{1+\varepsilon} t}} \\ \Rightarrow |\omega| &= \sqrt{\frac{\omega_3}{1 + C_1 \omega_3^2 e^{\frac{2KF(x_0)(\omega_3^2 - u)}{1+\varepsilon} t}}} \\ \text{with } C_1 &= \frac{\omega_3^2(0) - |\omega(0)|^2}{|\omega(0)|^2 \omega_3^2(0)} \end{aligned}$$

The solution predicts that $|\omega| \rightarrow \omega_3$ as $t \rightarrow \infty$, with ω_1 and ω_2 decaying to zero. Since $m \sim$ is aligned with $\omega_3 \mathbf{e}_3$, the approximation $\dot{\omega}_3 \approx 0$ is valid. Even with this approximation, Euler's equations still remain intractable. It is possible, however, to obtain qualitative understanding

as to the decay functions of the angular velocities. Using the same reasoning of approximation, we can obtain the decay envelope of the angular velocity components

$$\begin{aligned} \Rightarrow \omega_1 &= C_1 e^{-KG(x_0)|\omega|t} \\ \omega_2 &= C_2 e^{-\frac{KE(x_0)|\omega|^2 t}{1+E^2}} \\ \text{with } C_{1,2} &= \omega_{1,2}(0) \end{aligned}$$

The solution to the angular velocity components are stretched exponentials and describe the decaying envelopes of the near field solutions for ω_1 and ω_2 .

4.5.1.3. *Dynamics* ($\omega(0) \parallel z \sim \omega_3(0) \mathbf{e}_3$). The approximate solution due to the near field torques of the asymmetric pulsar is accurate only for certain values of r_0 . The rotational behaviour of the pulsar is dependent on a net moment of inertia \tilde{I}_{net} , which takes into account the coupling of crustal and magnetic stresses. Wasserman (2003) [175] defines \tilde{I}_{net} as

$$\begin{aligned} I_{ij} &= \int d^3r \rho(\mathbf{r}) (r^2 \delta_{ij} - r_i r_j) \\ \rho(\mathbf{r}) &= \rho_0(r) + \rho_\omega(r, \mathbf{r}; \tilde{\mathbf{L}}) + \rho_B(r, \mathbf{r}; \tilde{\mathbf{B}}) \end{aligned}$$

where $\rho_0(r)$ is the spherical profile of an untorqued star, and $\rho_\omega(r, \mathbf{r}; \tilde{\mathbf{L}})$ and $\rho_B(r, \mathbf{r}; \tilde{\mathbf{B}})$ are distortions due to rotation and magnetic fields for a two component star.

An increase in $\mathbf{G}(x_0)$ due to a reduction in r_0 causes a decrease in the magnitude of \tilde{I}_{net} , which is illustrated by the behaviour of ω in figure 4.5.1 where, from top to bottom, left to right, $r_0 = 5$ km, 10 km, 50 km and 100 km. For $r_0 = 5$ km, the magnetic stress due to the near field torque is such that $I_{2net} > I_{1net} > I_{3net}$, while for $r_0 = 50$ km and 100 km, we have $I_{3net} > I_{2net} > I_{1net}$. Both cases mimic that of the stable rotator, with the path of ω having a flat disk like nature due to the absence of radiation along \mathbf{e}_3 . For $r_0 = 10$ km, $I_{2net} > I_{3net} > I_{1net}$. This is a transitory period for the pulsar and as the magnetic field decreases, it will evolve into one of two possible stable rotations, or local energy eigenstates.

Solutions for $\omega(0) \parallel z \sim \omega_1(0) \mathbf{e}_1$ show the pulsar to behave similarly, with the shift between stable and unstable rotational states dominated by r_0 .

The same dependency exists on \mathbf{B}_0 , where an increase in the magnetic field results in an increase in the magnetic stresses, and so decrease in the magnitude of \tilde{I}_{net} .

4.5.1.4. *Near Field* ($\omega(0) \parallel z \sim \omega_2(0) \mathbf{e}_2$). Given an initial condition, which should lead to rotational instability, solutions are obtained as for an untorqued body

$$\begin{aligned} \omega_1 &= \frac{C_1 e^{-\omega_2 t} + C_2 e^{-\omega_1 t}}{\varepsilon_3} \\ \omega_2 &= \frac{C_1 e^{-\omega_2 t} + C_2 e^{-\omega_1 t}}{\varepsilon_1 - KG(x_0)|\omega|} \end{aligned}$$

4.5.1.5. *Far Field* ($\omega(0) \parallel z \sim \omega_2(0) \mathbf{e}_2$). There are no analytical solutions for the far field case. However, the far field torque has the effect of stabilising the rotation about an extreme moment of inertia. Both α and β will tend to a small, non-zero angle and will nutate.

4.5.1.6. *Dynamics* ($\boldsymbol{\omega}(0) \approx \omega_2(0)\mathbf{e}_2$). With the dependency of the rotational behaviour on $\tilde{\mathbf{I}}_{net}$, the initial value of $\boldsymbol{\omega}(0)$ no longer plays such a dominant role in the stability of the pulsar, shown by the stable rotation in figure 4.5.2 where the effect of magnetic stresses is such that $I_{2net} > I_{3net} > I_{1net}$.

4.5.2. General Solutions ($\chi = \frac{\pi}{2}$).

4.5.2.1. *Near Field* ($\boldsymbol{\omega}(0) \approx \omega_1(0)\mathbf{e}_1$). The qualitative rotational behaviour is much the same as for $\chi = 0$. Using the same approximations as for the case $\chi = 0$, the near field solutions are

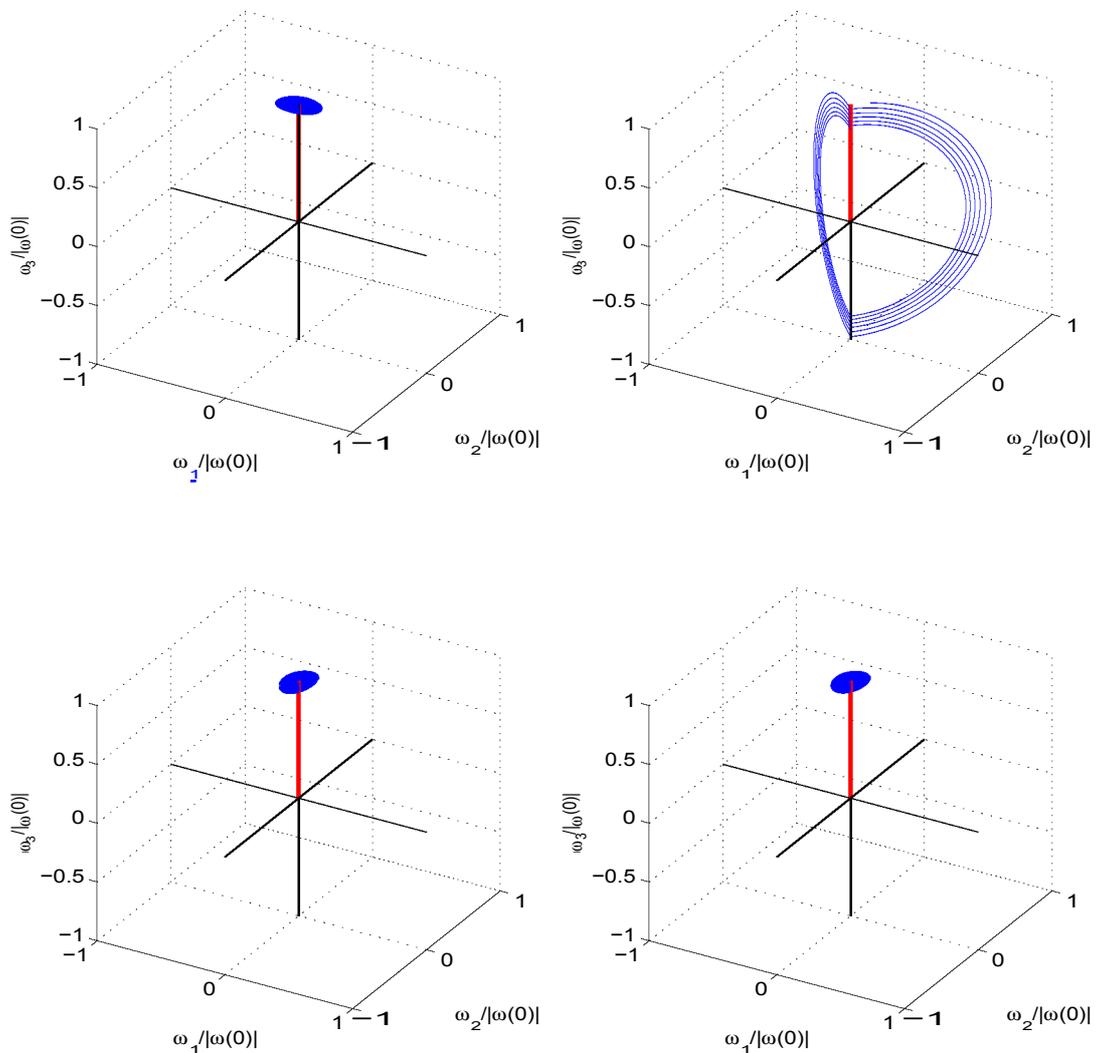


FIGURE 4.5.1. Numerical solution of $\boldsymbol{\omega}$ for asymmetric pulsar. From top to bottom, left to right, $r_0 = 5 \text{ km}, 10 \text{ km}, 50 \text{ km}, 100 \text{ km}$, $\chi = 0$ (red line)

given by

$$\begin{aligned} \omega_1 &= \omega_1(0) \\ \omega_2 &\approx A \sin \left(\frac{\partial \sigma}{\partial \omega} \frac{\partial}{\partial \omega} \frac{\epsilon_2 \pm \epsilon_3}{(1+\epsilon)(1+\epsilon^*)} \frac{1}{KG(x_0)|\omega|} + \epsilon_2 \epsilon_3 \omega_1 t + \xi \right) \\ \omega_3 &\approx \left\| \frac{1}{1+\epsilon} \frac{KG(x_0)|\omega| + \epsilon_3}{KG(x_0)|\omega| \frac{\epsilon_2 + \epsilon_3}{(1+\epsilon)(1+\epsilon^*)} + \frac{1}{(1+\epsilon)(1+\epsilon^*)} KG(x_0)|\omega| + \epsilon_2 \epsilon_3} \right\| \times \\ &A \cos \left(\frac{\partial \sigma}{\partial \omega} \frac{\partial}{\partial \omega} \frac{\epsilon_2 \pm \epsilon_3}{(1+\epsilon)(1+\epsilon^*)} \frac{1}{KG(x_0)|\omega|} + \epsilon_2 \epsilon_3 \omega_1 t + \xi \right) \end{aligned}$$

4.5.2.2. *Far Field* ($\omega(0) \approx \omega_1(0)\mathbf{e}_1$). Since m lies along ω_1 , the approximation $\dot{\omega}_1 \approx 0$ is valid. The solutions are calculated to be

$$\frac{d|\omega|}{dt} = (\epsilon_2 - \epsilon_3)\omega_1\omega_2\omega_3 - KF(x_0)|\omega|^2 \frac{1}{1+\epsilon}\omega_2^2 + \frac{1}{1+\epsilon\omega^2}$$

As in previous solutions, the first term is dependent on the free precession of the pulsar, and acts orthogonal to the radial action of the far field torques. With a large difference in time scales,

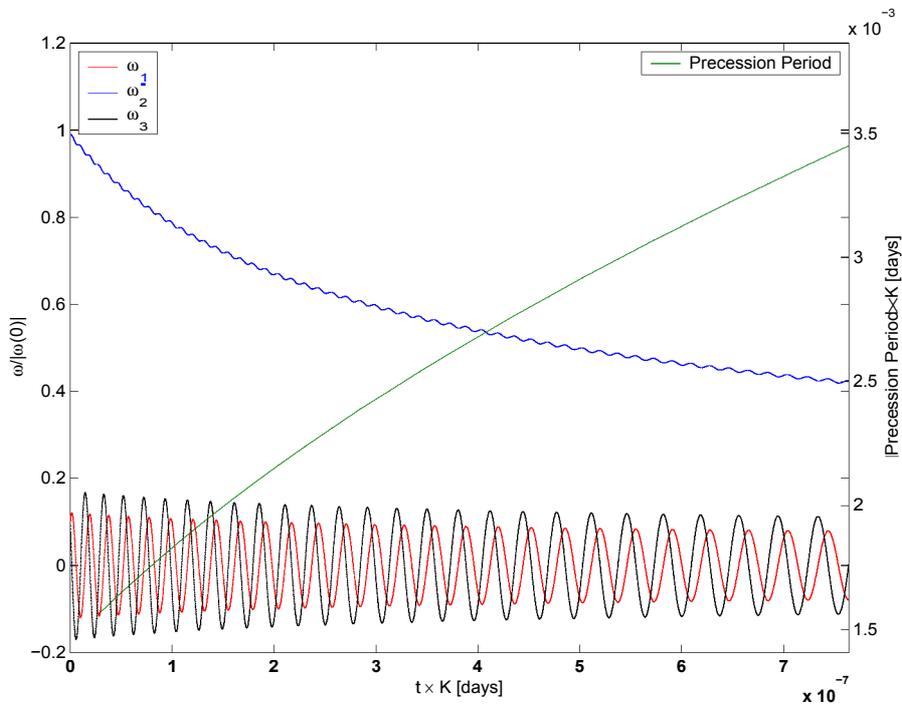


FIGURE 4.5.2. Numerical solution to Euler's equations for asymmetric pulsar, $\omega_2(0) = 10\omega_3(0)$, $\chi = 0$

the crustal stresses are ignored. Then

$$\begin{aligned} \text{let } u &= |\omega|^2 \\ \Rightarrow \frac{du}{dt} &\approx -KF(x_0)u(u - \omega_1^2) \frac{1}{1+\varepsilon} + \frac{1}{1+\varepsilon} \omega_1^2 \\ \Rightarrow u &= \frac{\omega_1^2}{1 + C_1 \omega_1^2 e^{-\frac{KF(x_0)\omega_1^2(2+\varepsilon)}{(1+\varepsilon)(1+\varepsilon)}t}} \\ \Rightarrow |\omega| &= \frac{\omega_1}{\sqrt{1 + C_1 \omega_1^2 e^{-\frac{KF(x_0)\omega_1^2(2+\varepsilon)}{(1+\varepsilon)(1+\varepsilon)}t}}} \\ \text{with } C_1 &= \frac{\omega_1(0)^2 - |\omega(0)|^2}{\omega_1(0)^2 |\omega(0)|^2} \end{aligned}$$

$|\omega| \rightarrow \omega_1$ as $t \rightarrow \infty$, with ω aligning with m . Recalling that $\dot{\omega}_1 \approx 0$ and approximating $\frac{1}{1+\varepsilon} \approx \frac{1}{1+\varepsilon}$.

$$\begin{aligned} \omega_2 &\approx \frac{2C_2}{C_1\omega_1^2 + e^{KF(x_0)\omega_1^2 t}} \times \frac{1}{1+\varepsilon} \cos(\sqrt{\varepsilon_2\varepsilon_3}\omega_1 t + \xi) \\ \omega_3 &\approx -\frac{\varepsilon_3}{\varepsilon_2} \frac{2C_2}{C_1\omega_1^2 + e^{KF(x_0)\omega_1^2 t}} \times \frac{1}{1+\varepsilon} \sin(\sqrt{\varepsilon_2\varepsilon_3}\omega_1 t + \xi) \end{aligned}$$

4.5.2.3. *Dynamics* ($\omega(0) \approx \omega_1(0)\mathbf{e}_1$). ω_1 oscillates, resulting in an oscillation in the precession period, but does not decay since $m \approx m\mathbf{e}_1$. A superimposed oscillation occurs on the decay envelope of ω_2 and ω_3 .

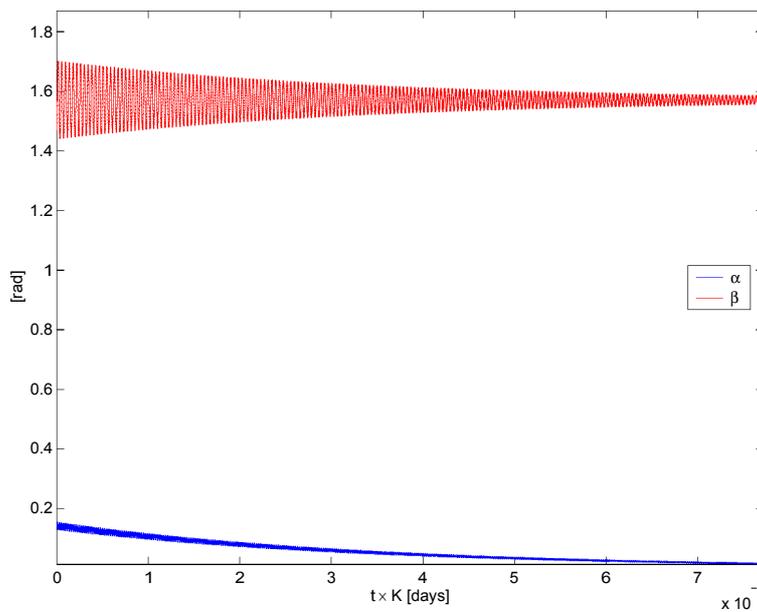


FIGURE 4.5.3. Numerical solution of β and α for asymmetric pulsar, $\chi = \frac{\pi}{2}$, $\omega_3(0) = 10\omega_{1/2}(0)$.

The behaviour of α and β , described in figure 4.5.3, indicates that the pulsar tends to become an aligned rotator with $\beta \rightarrow \frac{\pi}{2}$ and $\alpha \rightarrow 0$. The amplitude of nutation will decay, as will the nutation period.

4.5.2.4. *Near Field* ($\omega(0) \approx \omega_2(0)\mathbf{e}_2$). Following the method used to solve for the untorqued asymmetric pulsar in section 3.5, the solutions for the torqued case are

$$\begin{aligned}\omega_1 &\approx -C_1 e^{\omega_2 \int \varepsilon_1(\varepsilon_3 + \frac{1}{1+\varepsilon}) KG(x_0)\omega_2 t} + C_2 e^{-\omega_2 \int \varepsilon_1(\varepsilon_3 + \frac{1}{1+\varepsilon}) KG(x_0)\omega_2 t} \\ \omega_3 &\approx \frac{\varepsilon_3 \pm \frac{1}{1+\varepsilon} KG(x_0)|\omega|}{\varepsilon_1} C_1 e^{\omega_2 \int \varepsilon_1(\varepsilon_3 + \frac{1}{1+\varepsilon}) KG(x_0)\omega_2 t} + C_2 e^{-\omega_2 \int \varepsilon_1(\varepsilon_3 + \frac{1}{1+\varepsilon}) KG(x_0)\omega_2 t}\end{aligned}$$

with $\omega_2 \approx \text{constant}$. The pulsar behaves as in the untorqued case for suitable values of r_0 .

4.5.2.5. *Far Field* ($\omega(0) \approx \omega_2(0)\mathbf{e}_2$). No analytical solutions can be obtained for the far field torque.

4.5.2.6. *Dynamics* ($\omega(0) \approx \omega_2(0)\mathbf{e}_2$). ω_3 and ω_2 oscillate in quadrature, whilst ω_1 oscillates about a constant value. $|\omega|$ decays exponentially with a superimposed oscillation. The effect of α and β depends upon the timescale of the near field torques. If it is shorter than that of Eulerian precession ($\gamma \gg \zeta$), $\alpha \rightarrow 0$ and $\beta \rightarrow \frac{\pi}{2}$, becoming an aligned rotator, both with decaying nutational amplitudes and periods. If Eulerian precession dominates ($\gamma \ll \zeta$), then $\alpha \rightarrow 0$ or $\frac{\pi}{2}$, depending on initial conditions.

4.5.3. General Solutions ($0 < \chi < \frac{\pi}{2}$).

4.5.3.1. *Dynamics* ($\omega(0) \approx \omega_3(0)\mathbf{e}_3$). Equations 4.2.1 - 4.2.4 and 4.2.7 - 4.2.5 have no analytical solutions for an arbitrary m . Numerical simulations of Euler's equations are shown in figure 4.5.4. The solutions have similar characteristic behaviour to that of the axisymmetric pulsar in figure 4.4.3. This can be explained by the coupling of the magnetic and crustal stresses, which cause the axisymmetric pulsar to behave as an asymmetric body. The vector paths of ω are shown in figure 4.5.5 for $\chi = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}$ and $\frac{5\pi}{12}$. The effect of the asymmetry of the pulsar, in comparison to the axisymmetric system, is to 'squash' the disk-like nature of the vector path along the axis of compression, in this case \mathbf{e}_1 . Thus, on average, the angle between $\hat{\rho}$ and $\hat{\omega}$ decreases for $\gamma < \frac{\pi}{4}$ and increases for $\gamma > \frac{\pi}{4}$. Comparing the vector path of ω for $\chi = \frac{\pi}{4}$ in Figures 4.4.5 and 4.5.5, the initial cone of precession (the cone traced out by ω) has a smaller half angle in the asymmetric case. This half angle is, however, larger in the asymmetric case after the precession mode change as γ increases beyond $\frac{\pi}{4}$.

The mode change for $\chi = \frac{\pi}{4}$ is not observed for the initial condition $\omega(0) \approx \omega_1(0)\mathbf{e}_1$, shown in figure 4.5.6. The mode change will occur when $\omega(0)$ is such that the magnetic axis is not constrained within a pseudoperiodic loop or circumscription of the path of ω . Thus, the mode change will not occur in the case $\omega(0) \approx \omega_1(0)\mathbf{e}_1$ as long as $\chi < \beta$ for all time.

4.5.3.2. *Dynamics* ($\omega(0) \approx \omega_2(0)\mathbf{e}_2$). The effect of the magnetic axis on \tilde{I}_{net} , and subsequent rotational stability, is seen in figure 4.5.7. For $\chi = \frac{\pi}{12}$, the magnetic stresses have sufficiently adjusted \tilde{I}_{net} such that $I_{1net} < I_{3net} < I_{2net}$, resulting in a stable rotational state. With an increase in χ , I_{3net} increases while I_{1net} decreases such that $I_{1net} < I_{2net} < I_{3net}$ and the system returns to an unstable rotational state for $\chi = \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}$ and $\frac{5\pi}{12}$.

The simulations indicate mode jumping in c_i , resulting in the pseudo-glitch phenomenon discussed by Melatos (2000) [126].

Under the action of far field torques, free precession will be damped, and the pulsar becomes a stable oscillator with $-y \gg C$. For $-y \ll C$, $0 \rightarrow x$ as $t \rightarrow \infty$. For $-y \gg C$, 0 increases such that \mathfrak{c} tends towards an extreme moment of inertia. α oscillates about a non-zero mean, with decaying nutational amplitude and period.

The dependency on r_0 is examined by simulation of the systems for different values of r_0 at $x = 0$, $\frac{\pi}{4}$ and $\frac{\pi}{2}$, shown in Figures 4.5.8, 4.5.9 and 4.5.10 respectively. For $x = 0$, the pulsar has a stable rotation for $r_0 = 5$ km and 10 km. With an increase in r_0 , the ratio $\frac{G(x_0)}{F(x_0)}$ decreases and so the magnitude of \mathbf{I}_{net} increases, resulting in unstable rotational behaviour for $r_0 = 50$ km and 100 km. For $x = \frac{\pi}{4}$ and $\frac{\pi}{2}$, the decrease in $\frac{G(x_0)}{F(x_0)}$ results in a more unstable rotation. However, this instability is pronounced for $x = \frac{\pi}{4}$ since the magnetic axis will decrease \mathbf{I}_{net} substantially for $x = \frac{\pi}{2}$. The decrease in \mathbf{I}_{net} results in the compression along \mathfrak{c}_1 of the vector path of \mathfrak{c} and effectively flattens the disk-like structure.

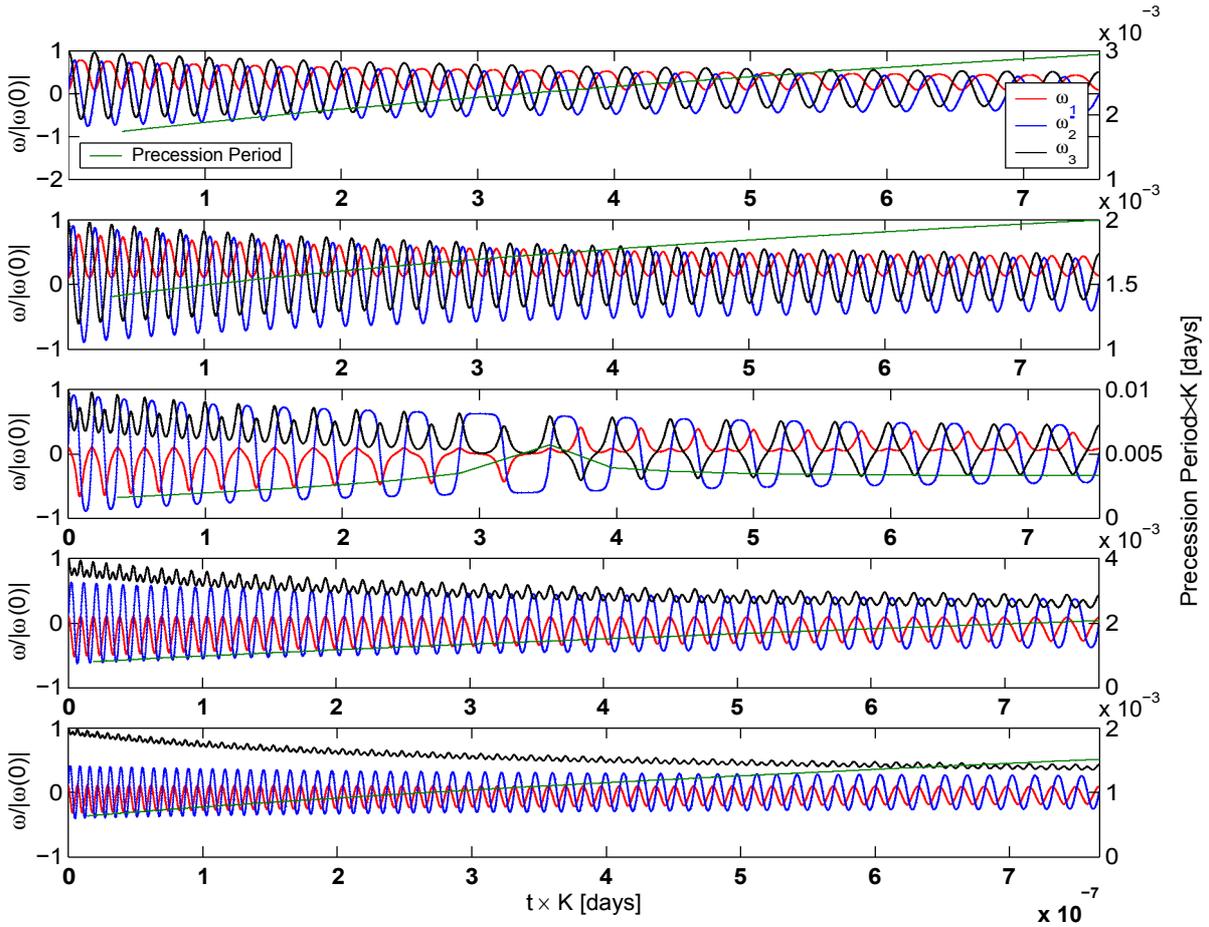


FIGURE 4.5.4. Numerical solution to Euler's equations for asymmetric pulsar, $W_3(0) = 10W_{12}(0)$. From top to bottom, $x = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$.

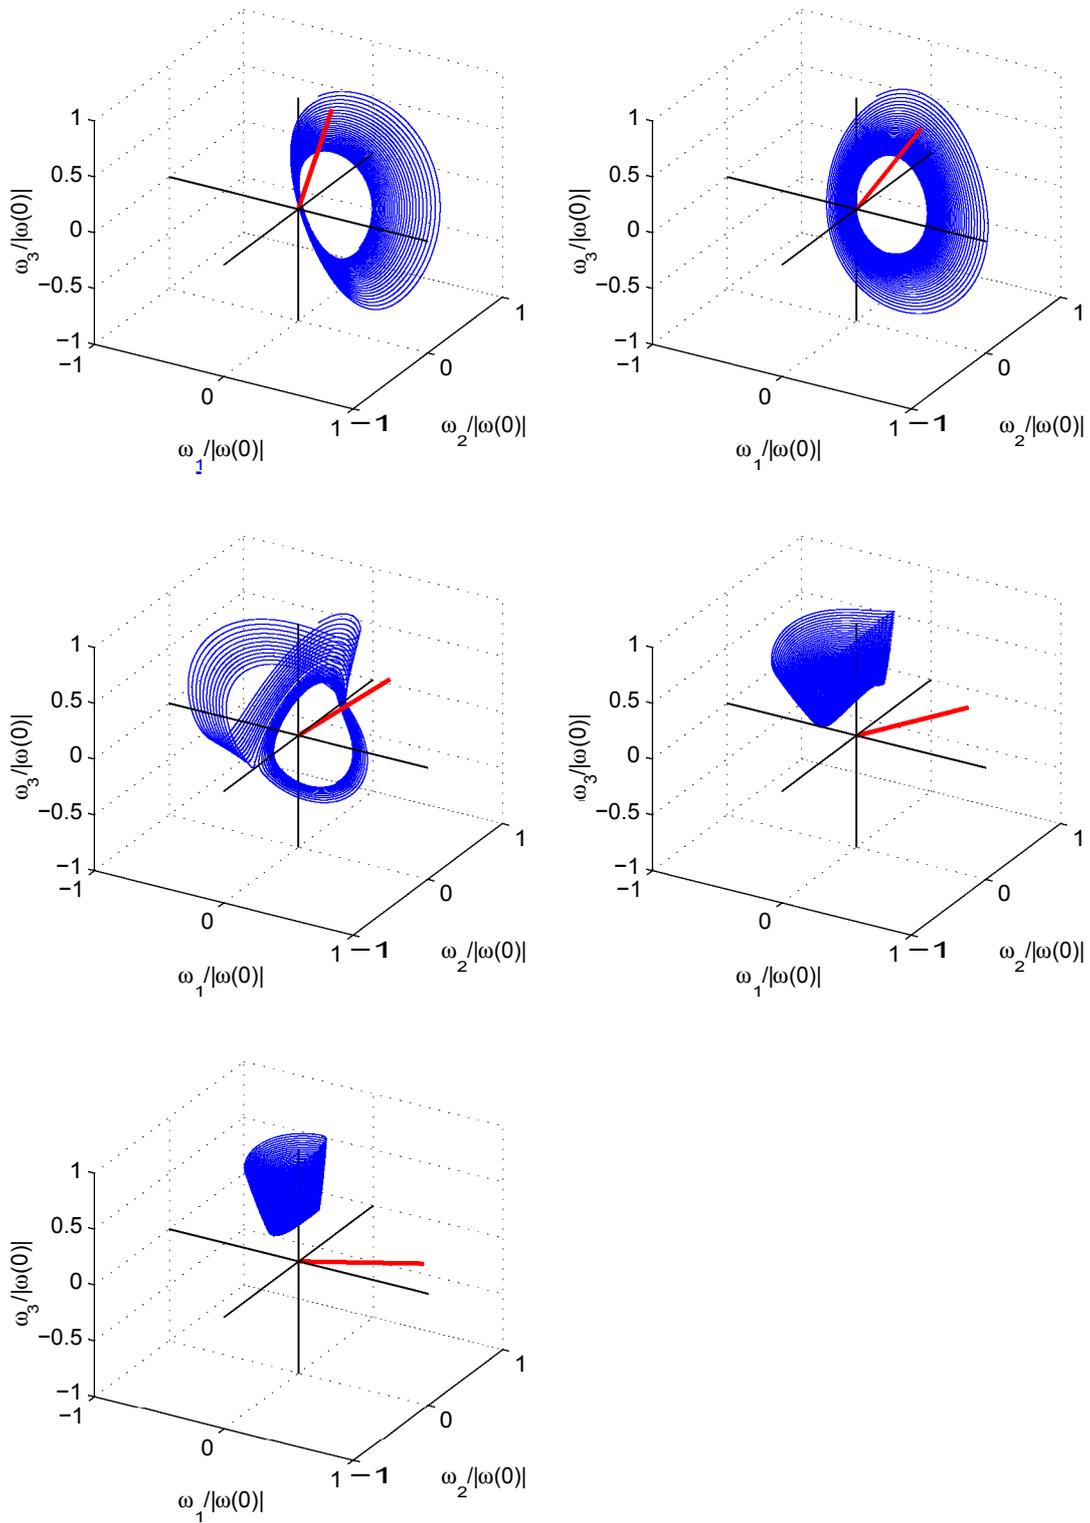


FIGURE 4.5.5. Numerical solution of ω for asymmetric pulsar, $\omega_3(0) = 10\omega_{12}(0)$. From top to bottom, left to right, $\chi = 12, 6, 4, \frac{\pi}{3}, \frac{5\pi}{12}$ (ω indicated by red line).

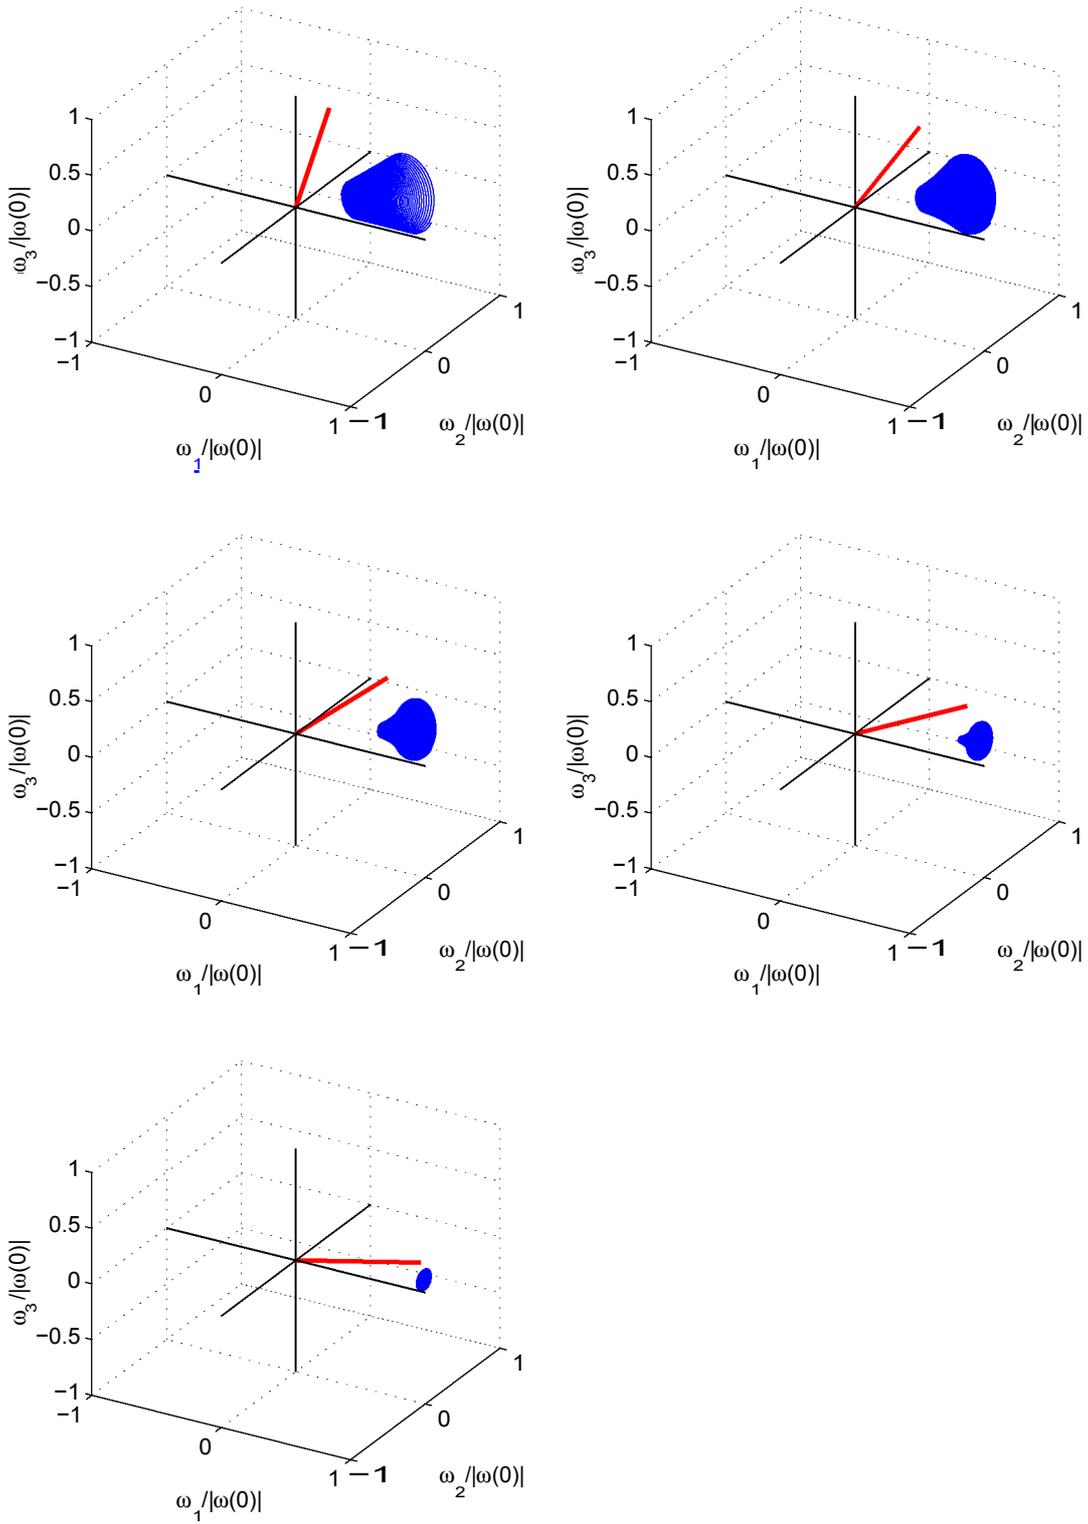


FIGURE 4.5.6. Numerical solution of ω for asymmetric pulsar, $\omega_1(0) = 10\omega_{2,3}(0)$. From top to bottom, left to right, $\chi = \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}, \frac{5\pi}{12}$ (\mathbf{m} indicated by red line).

4.6. Torqued Precession in Pulsars?

Uncontrovertable evidence of torqued precession in isolated pulsars is best described as ‘thin’. However, a number of candidates for free precession, such as PSR 1828-11, can be mimicked convincingly by a torqued system model. The coupling of Eulerian and radiative precession in many circumstances leads to degenerate solutions in the timing residual solution space. Link & Epstein (2001) [103] were able to model PSR 1828-11 with a magnetic dipole consisting of a near field and far field radiation zone. By including time dependencies in the magnetic dipole torque through the addition of a rotating dipole vector, torqued precession can attribute for the 1000 day period, as well as the entire observed Fourier spectrum and pulse shape variations [150].

The two anomalous X-Ray pulsars 1E 1048.1-5937 and 1E 2259+586, which are known to spin down irregularly, can be modeled as isolated axisymmetric neutron stars with magnetic dipoles of surface field strength $\geq 10^{10}$ T [125]. The high magnetic field modeled here should not be used to exclude most observed radio pulsars since asymmetric modeling of the magnetic field can lead to a decrease in the expected field strength.

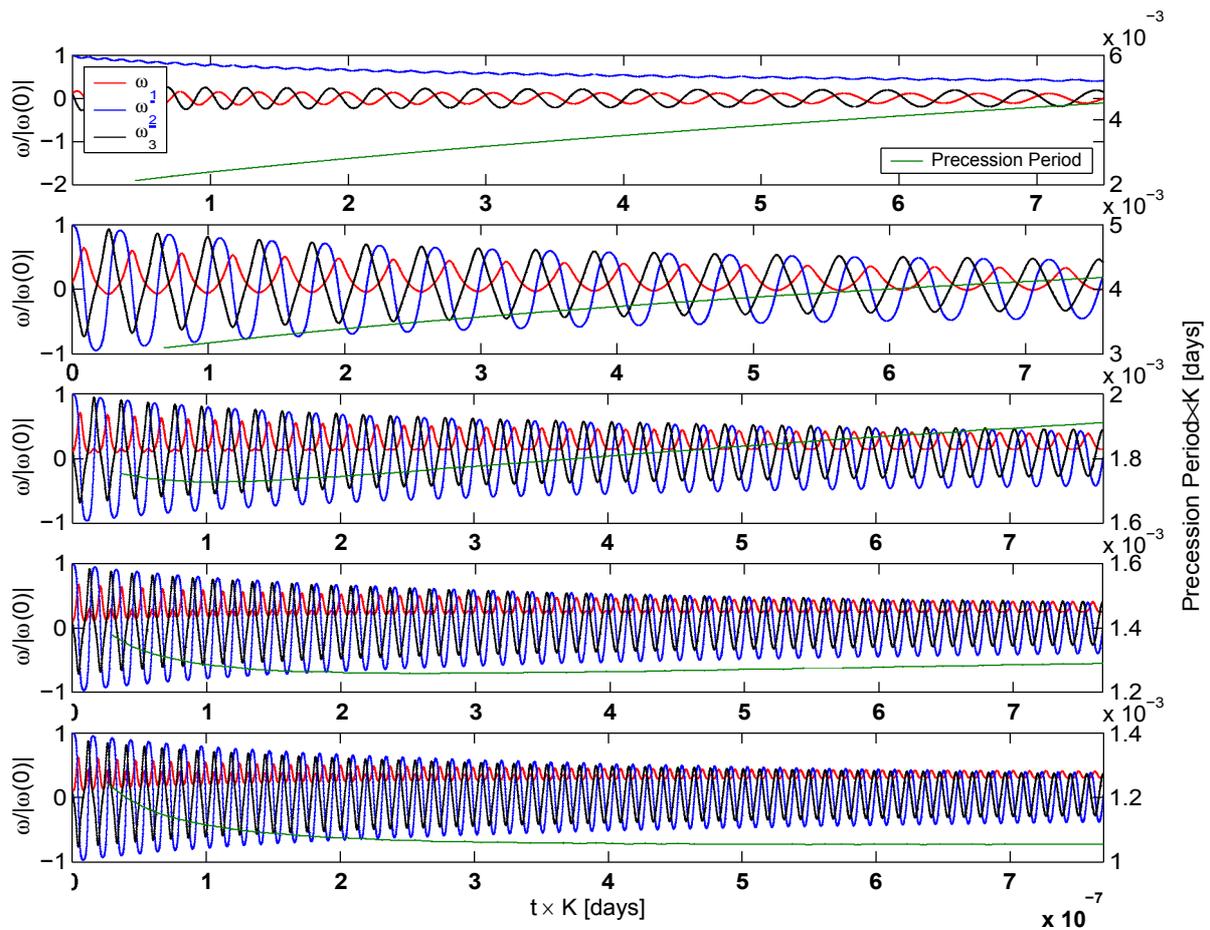


FIGURE 4.5.7. Numerical solution to Euler's equations for asymmetric pulsar, $W_2(0) = 10W_{13}(0)$. From top to bottom, $X = \begin{matrix} 7 & 57 \\ 12 & 6 & 4 & 3 & 12 \end{matrix}$.

The Crab pulsar, with a reported 60 second precession period [32], has been fitted with a simple free precession model. A torqued precessional model, with a very high surface magnetic field $\geq 10^{14}$ T, could reproduce the observed precession period.

With the long dataspan available at HartRAO, a number of pulsars are expected to be ideal candidates to test against precession models. Although simple fitting has been attempted in the past [37], only free precession has been considered. Due to the degenerate timing residual solution space, it should be possible to fit a torqued precessional model equally well. This would provide a more satisfactory dynamical solution to the behaviour of pulsars.

4.7. Summary

The analytic and numerical analysis of torqued precession describes the rotational behaviour of braking pulsars, and the dependency of precession on critical physical parameters. The use of

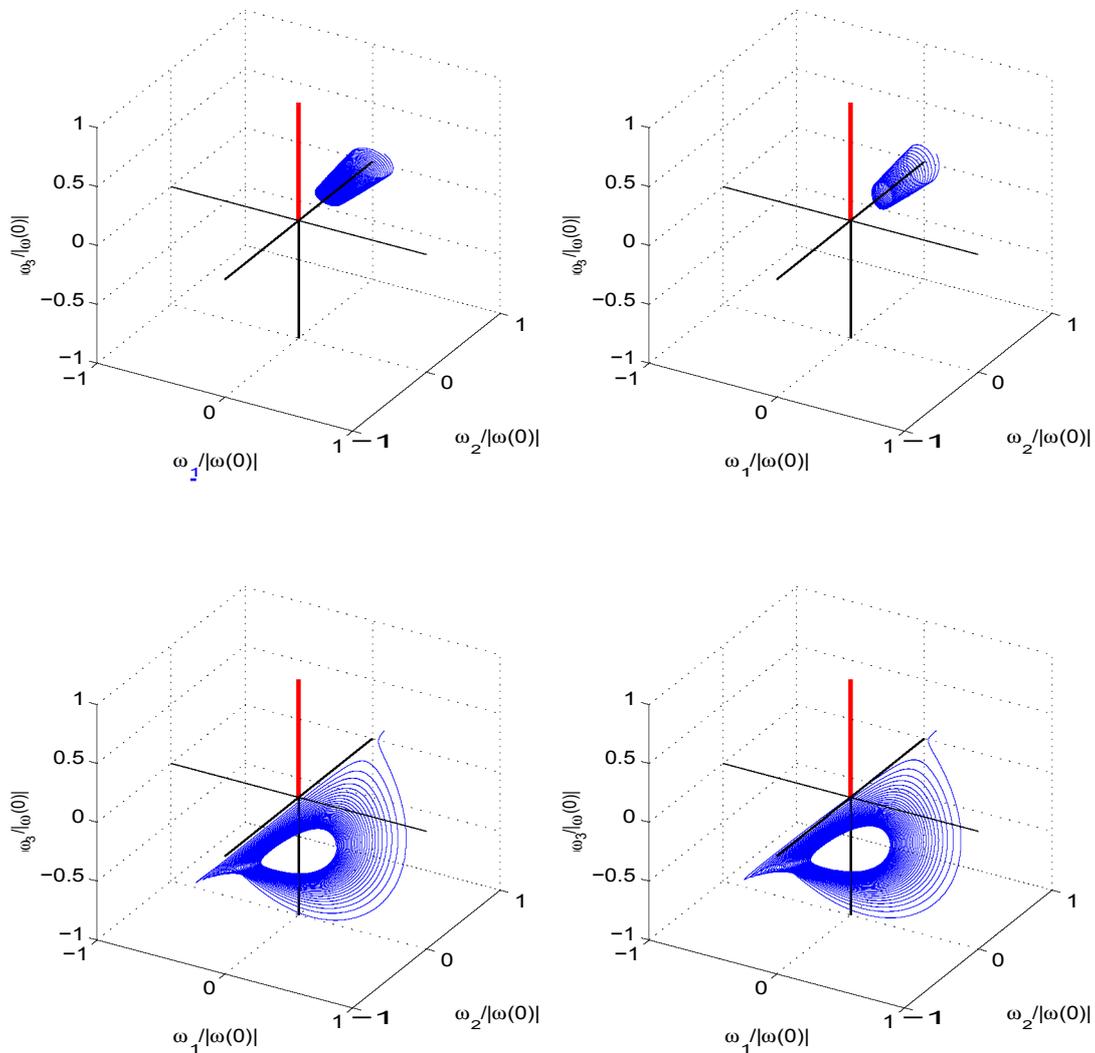


FIGURE 4.5.8. Numerical solution of W for asymmetric pulsar. From top to bottom, left to right, $r_0 = 5$ km, 10 km, 50 km and 100 km. $X = 0$, $W_2(0) = 10W_{1,3}(0)$.

a magnetic dipole model introduces a second point of attraction in the configuration phase space, which increases the dependence of precessional behaviour on initial conditions. The coupling of magnetic and crustal stresses leads to rotational behaviour characteristic of asymmetric bodies, and is not unique for a given set of parameters.

Using an upper limit of 10^9 T for \mathbf{B}_0 results in coupled precessional periods of the order of 1000 days, with $\epsilon \approx 10^{-5}$. Shorter periods would require arguments for stronger surface magnetic fields and near field radiation zones, which do not alter the magnitude of the far field torque.

The torqued spherical system causes a precessional motion indistinct from that of the untorqued, axisymmetric pulsar. Clearly defined analytical solutions show that the pulsar becomes an aligned rotator for all initial conditions.

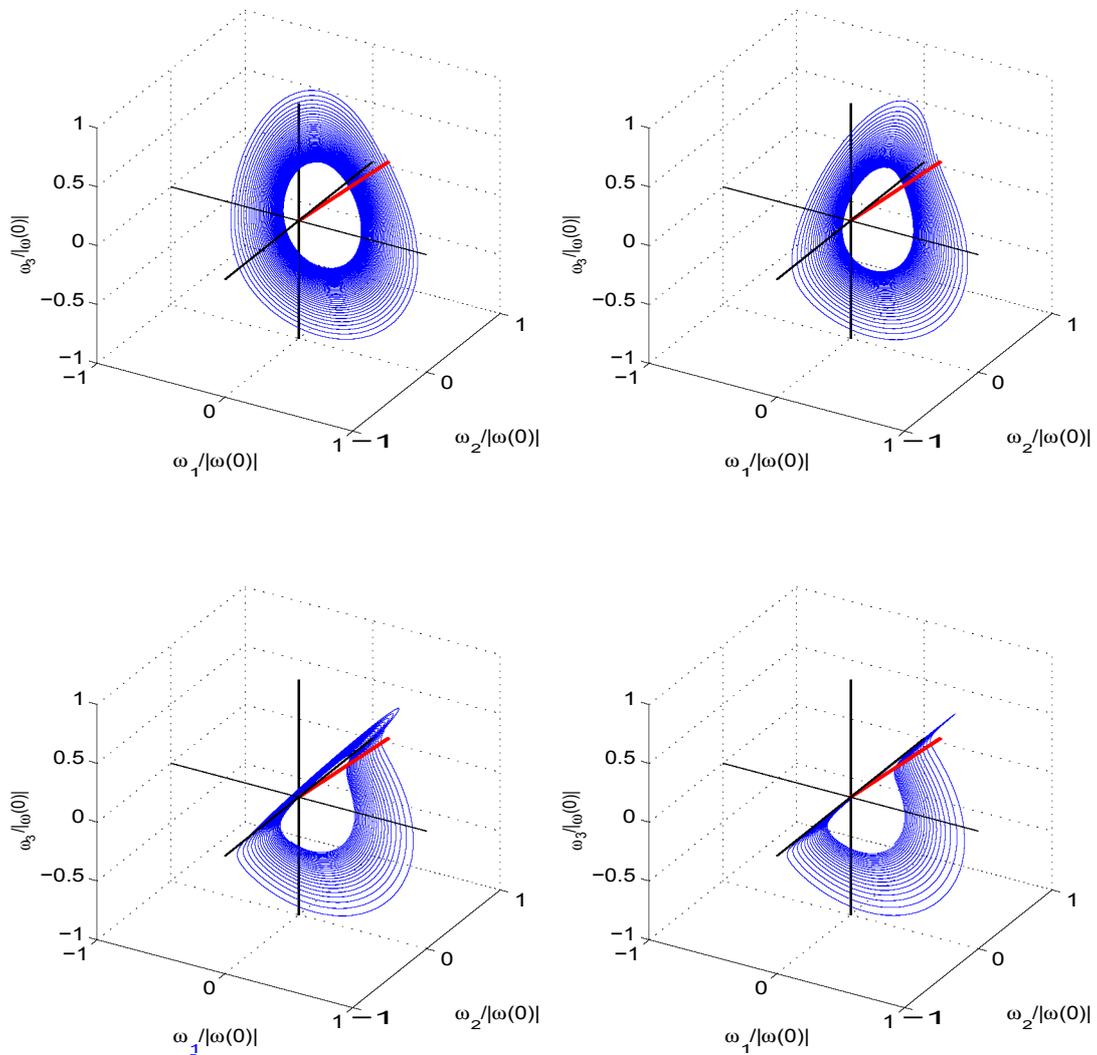


FIGURE 4.5.9. Numerical solution of W for asymmetric pulsar. From top to bottom, left to right, $r_0 = 5$ km, 10 km, 50 km and 100 km. $X = 4$, $W_2(0) = 10W_{1,3}(0)$.

The rotational behaviour for the torqued axisymmetric pulsar is similar to that of the untorqued axisymmetric system. For $x = 0$, the pulsar has a precession period dependent on both crustal and magnetic stresses. In general, the energy eigenstate is one of precession. The nature of precession, Eulerian or radiative, is dependent on K , r_0 , x , 1c and $cv(0)$ in a complicated manner. The action of the far field torque causes the rotational behaviour of the pulsar to tend towards stability, and to become an aligned rotator. The resultant time dependency in the precessional behaviour can cause a change in the nature of precession if the pulsar is not initially dominated by radiative precession. This modal change causes anomalies in the precession period, and can result in the pulsar becoming unobservable.

The rotational stability of the asymmetric body, and by argument all possible cases of torqued precession, is shown to be dependent on the net inertial effect of magnetic and crustal stresses. The magnetic field decreases the magnitude of f_{net} , dependent on x , K and r_0 . In an

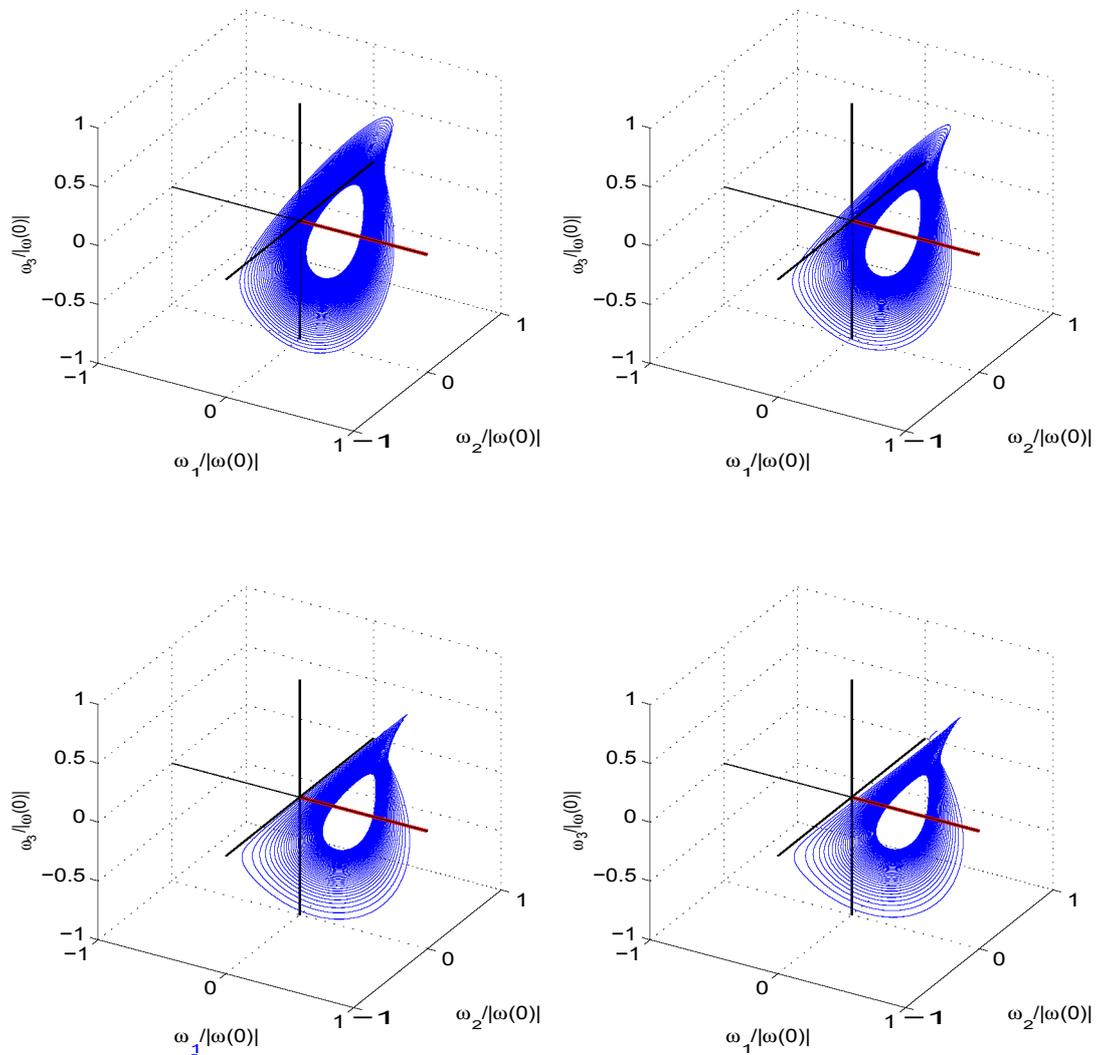


FIGURE 4.5.10. Numerical solution of W for asymmetric pulsar. From top to bottom, left to right, $r_0 = 5$ km, 10 km, 50 km and 100 km. $X = \frac{\pi}{2}$, $W_2(0) = 10W_{1,3}(0)$.

unstable rotational state, the far field torques act to stabilise the rotation about net maximum or minimum moment of inertia. With the decay of the magnetic field, the pulsar may be subject to two phases of rotational instability in its lifetime. This instability results in two possible phenomenon: the intermittent visibility of the pulsar beam, and observation of both magnetic poles, which do not necessarily have the same geometry.

Much like the axisymmetric case, the precessional mode change is characteristic of the behaviour of the torqued asymmetric pulsar. This result is used to best describe the effect of the asymmetry on ω where the vector path of ω is compressed along the axis of crustal and magnetic compression.

Given the degenerate solution space for the fitting of precession models, the availability of evidence for torqued precession is minimal. With the long, uninterrupted data spans required, the sample of HartRAO pulsars may best provide a candidate to which a torqued precession model could be fitted.

CHAPTER 5

Timing Residuals

5.1. Introduction

Pseudo-cyclic modulation in the timing residuals of a subset of the pulsars monitored at HartRAO is speculated to be a result of torqued precessional behaviour. The period of modulation would require an oblateness of the order of 10^{-9} . With the possibility of crustal deformation due to spin down and a magnetic field, precession is expected in a large number of pulsars.

Calculation of timing residuals requires the use of the Euler angles φ , θ and ψ , which are defined by expressing the angular velocity $\vec{\omega}$ as the time derivative of the three independent rotations

$$\vec{\omega} = \omega_1 \mathbf{e}_1 + \omega_2 \mathbf{e}_2 + \omega_3 \mathbf{e}_3 = \dot{\varphi} \tilde{\mathbf{E}}_3 + \dot{\theta} \mathbf{n}_1 + \dot{\psi} \mathbf{e}_3$$

where $\tilde{\mathbf{E}}_3$, \mathbf{n}_1 and \mathbf{e}_3 are the relevant basis vectors in an inertial frame, nodal frame and body frame respectively. In the body's frame of reference (see appendix A.2),

$$\omega_1 = \dot{\varphi} \sin \psi \sin \theta + \dot{\theta} \cos \psi$$

$$\omega_2 = \dot{\varphi} \cos \psi \sin \theta - \dot{\theta} \sin \psi$$

$$\omega_3 = \dot{\varphi} \cos \theta + \dot{\psi}$$

Solving for $\dot{\varphi}$, $\dot{\theta}$ and $\dot{\psi}$ gives

$$(5.1.1) \quad \dot{\varphi} = \frac{\omega_1 \sin \psi + \omega_2 \cos \psi}{\sin \theta}$$

$$(5.1.2) \quad \dot{\theta} = \omega_1 \cos \psi - \omega_2 \sin \psi$$

$$(5.1.3) \quad \dot{\psi} = \omega_3 - \frac{\omega_1 \sin \psi + \omega_2 \cos \psi}{\tan \theta}$$

where numerical integration produces the Euler angles. To ensure numerical stability, the integration step size should be a small fraction of the pulsar rotation period. The number of rotations per day is of the order of 10^6 . With observed precession periods of the order of 1000 days, the number of calculated points would exceed 10^9 . The computer processing power to perform this calculation is not available. However, since the long term spin down of the pulsar is not required for the calculation of timing residuals, which produces an angle $-\pi < \Delta 4b < \pi$, the Euler angles are expressed in terms of trigonometric functions. These could take on a different form for the untorqued and torqued cases.

The measured phase modulation of the pulsar signal for $\theta > \chi$ is described by

$$(5.1.4) \quad \Delta 4b = 4b - \left(\varphi - \frac{\pi}{2} \right) = \arctan \left(\frac{1}{\cos \theta} \frac{\cos \psi \tan \chi}{\tan \theta - \sin \psi \tan \chi} \right)$$

where χ describes the azimuthal angle of the magnetic axis m [83]. For $\theta < \chi$, the timing residuals are

$$(5.1.5) \quad \Delta t_1 = t_1 - (\varphi + \psi) = \arctan \frac{(\cos \theta - 1) \sin \psi \sin \chi - \sin \theta \cos \chi}{\psi \sin \chi + (\cos \theta \cos \psi \sin \chi - \sin \theta \cos \chi) \tan \psi}$$

The terms $(\varphi - \frac{\pi}{2})$ and $(\varphi + \psi)$ are angular displacements at a time t resulting from two different periods for the pulsar. The pulsar period is dependent on whether the angle of magnetic inclination χ is greater than or less than the wobble angle θ .

5.2. Calculation of Phase Residuals: Untorqued Pulsars

For torque free motion, we assume that the angular momentum vector \tilde{J} is aligned with \tilde{E}_3 . The Euler angles for this case are [175]

$$(5.2.1) \quad \cos \psi = \frac{J_2}{J \sin \theta}$$

$$(5.2.2) \quad \sin \psi = \frac{J_1}{J \cos \theta}$$

$$(5.2.3) \quad \theta = \arccos \frac{J_3}{J}$$

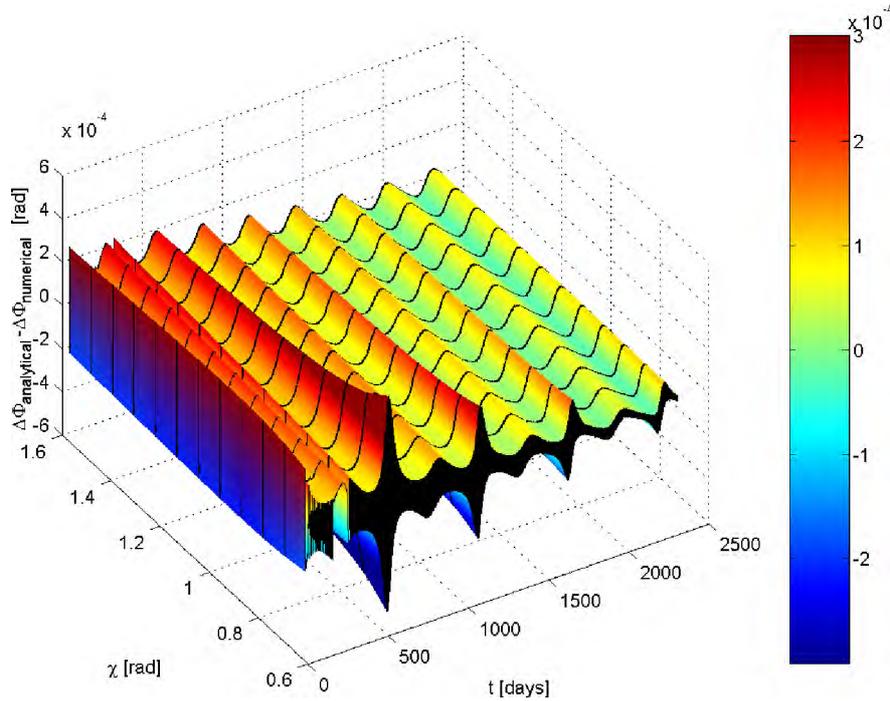


FIGURE 5.2.1. Calculation of $\Delta\Phi_{\text{analytical}} - \Delta\Phi_{\text{numerical}}$ for an asymmetric pulsar.

Phase residuals $\Delta t_1^{\text{numerical}}$ and $\Delta t_1^{\text{analytical}}$ were calculated for an untorqued asymmetric pulsar, with $0 < \chi < \frac{\pi}{2}$, using two different methods. $\Delta t_1^{\text{numerical}}$ were calculated using Euler angles obtained by numerical integration of equations 5.1.1-5.1.3. $\Delta t_1^{\text{analytical}}$ were calculated using the trigonometric expressions of the Euler angles in equations 5.2.1-5.2.3. Figure 5.2.1

shows $\Delta 4b_{\text{analytical}} - \Delta 4b_{\text{numerical}}$. The residual error is due to the limited degree of accuracy of the numerical simulation (10^{-4}) and a constant phase offset between the two surfaces.

5.2.1. Axisymmetric Pulsars. Following the format of published results on pulsar precession [164, 175], figure 5.2.2 are the phase residuals $\Delta 4b$, $\Delta P = \frac{d(\Delta t)}{dt}$ (where $\Delta t = \frac{\Delta \Phi}{-g^{2\pi}}$ \times Pulsar Period) and $\frac{d(\Delta P)}{dt}$ predicted for an axisymmetric, untorqued pulsar, $e = 10^{-4}$ and $\theta(0) = 0.1405$ rad. Semi-analytic solutions have been discussed by Jones & Andersson (2001) [83]. For the case $\theta \ll x$, $\Delta 4b = -\frac{e}{\tan x} \cos \theta$, while for $\theta \gg x$, $\Delta 4b = \frac{x}{\sin \theta} \cos \theta$, shown in figure 5.2.3.

The precession period scales linearly with the oblateness parameter e .

5.2.2. Asymmetric Pulsars. As no analytical solutions exist for the asymmetric case, we rely on numerical simulations to provide a generic template from which timing residuals can be fitted. Given the inherent similarity in timing residuals of the untorqued asymmetric case and torqued case, only two categories are considered: that of the untorqued axisymmetric pulsar, and the torqued pulsar (which, due to the coupling of crustal and magnetic stresses, has a rotational behaviour characteristic of an asymmetric body).

5.3. Calculation of Phase Residuals: Torqued Pulsars

Calculation of timing residuals for the untorqued case required the assumption that \tilde{J} is aligned with \tilde{E}_3 . In the torqued case, this approximation holds provided the spin down timescale is far greater than that of the precession and that, given the radial nature of the far field radiation torque, \tilde{J} will decrease in magnitude only. Over a single precession period, these approximations are valid.

As discussed in the previous section, calculation of timing residuals is performed for the asymmetric pulsar only. The timing residual model takes into account variations in the braking torque due to precession, and subsequent variations in P .

5.3.1. Asymmetric Pulsars. Calculations have revealed four generic templates from which timing residuals can be fitted and explained by a torqued precessional model. Figures 5.3.1, 5.3.2, 5.3.3, 5.3.4, 5.3.5, 5.3.6, 5.3.7 and 5.3.8 are timing residuals for $\theta(t) < x$, $\theta(t) > x$ (stable rotation), $\theta(t) < x$ and $\theta(t) > x$ (unstable rotation). The terms stable and unstable rotation are used loosely, and simply imply a value of $\text{modulo}[\theta(0), \frac{\pi}{2}] \approx 0$ for stable rotation, whilst for unstable rotation this does not hold true. Given the degeneracy of the timing residual solution space, no automatic algorithms exist to obtain parameters from the data. In general, models and simulations predict that the amplitude of phase residual scales with $\theta(0)$, the shape of the phase residual is determined by x , whilst the precession period is determined by B (including the coupling distance between the near field and far field structures) and the oblateness parameters e and e' . To maintain constant coupling between the Eulerian and radiative precession, the period scales as $\frac{\Delta B}{B} \propto \Delta \frac{e'^2}{e^2}$.

The timing residuals for torqued (and untorqued) pulsars have been calculated for an oblate body. The characteristic precessional behaviour is reversed in time for a prolate body.

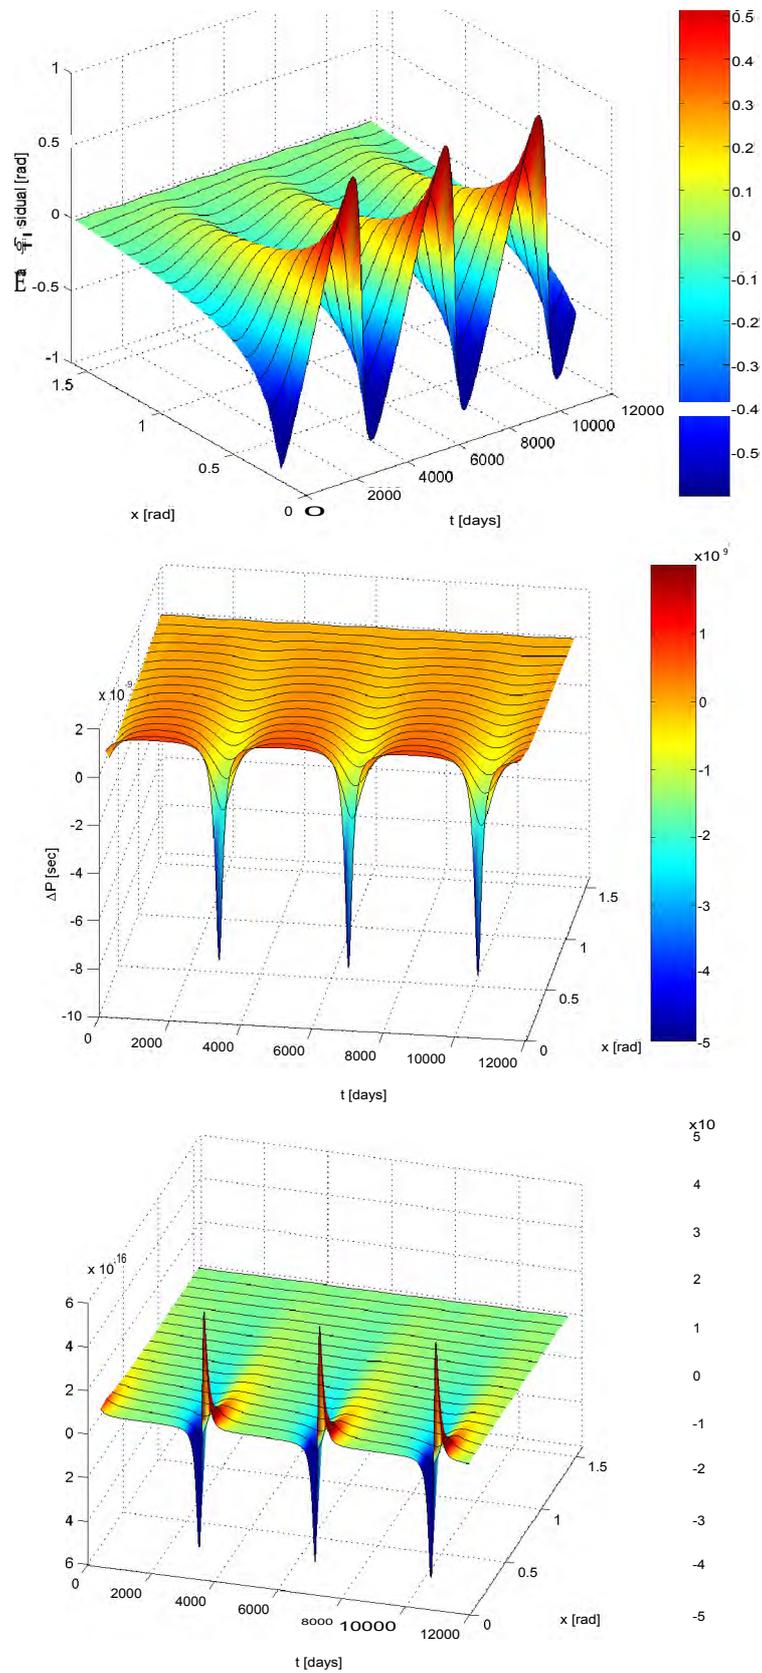


FIGURE 5.2.2. Timing residuals for an untorqued, axisymmetric pulsar, $\alpha(0) = 0.1405$ rad. From top to bottom are Phase, AP and $d(\Delta P)$.

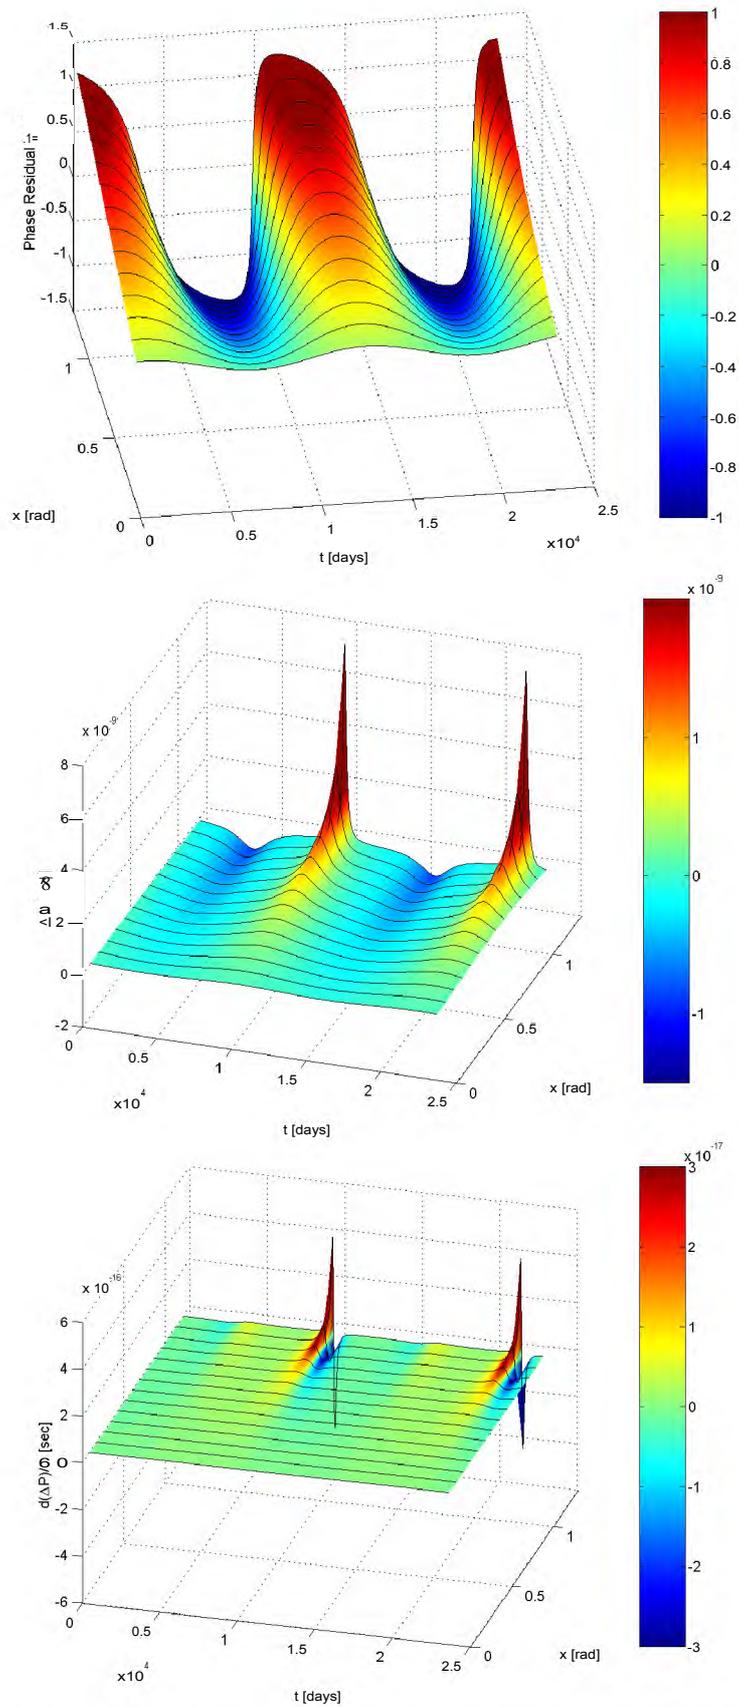


FIGURE 5.2.3. Timing residuals for an untorqued, axisymmetric pulsar, $\theta(0) = 1.29$ rad. From top to bottom are Phase, ΔP and $d(\Delta P)/dt$

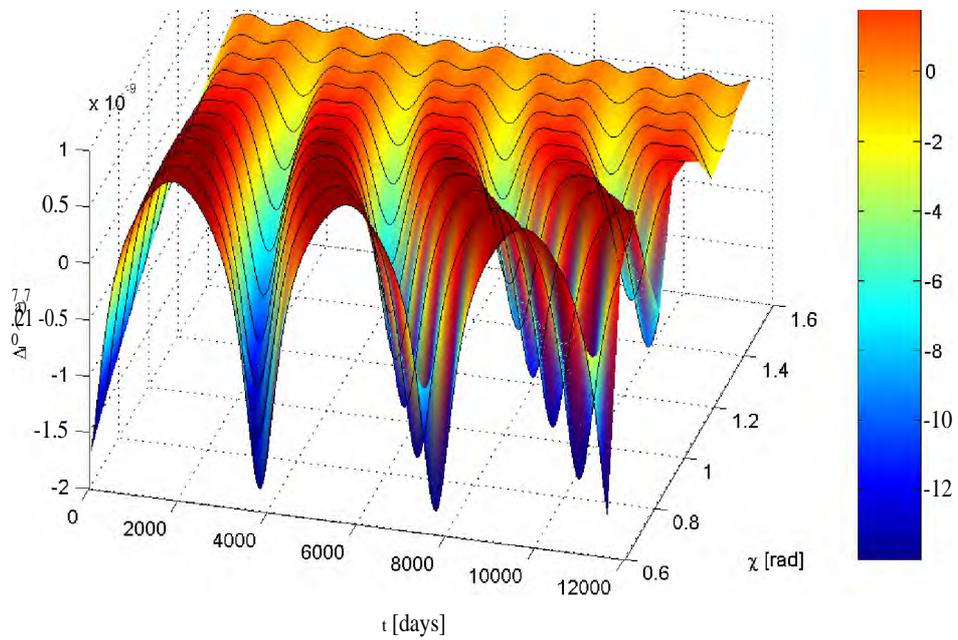
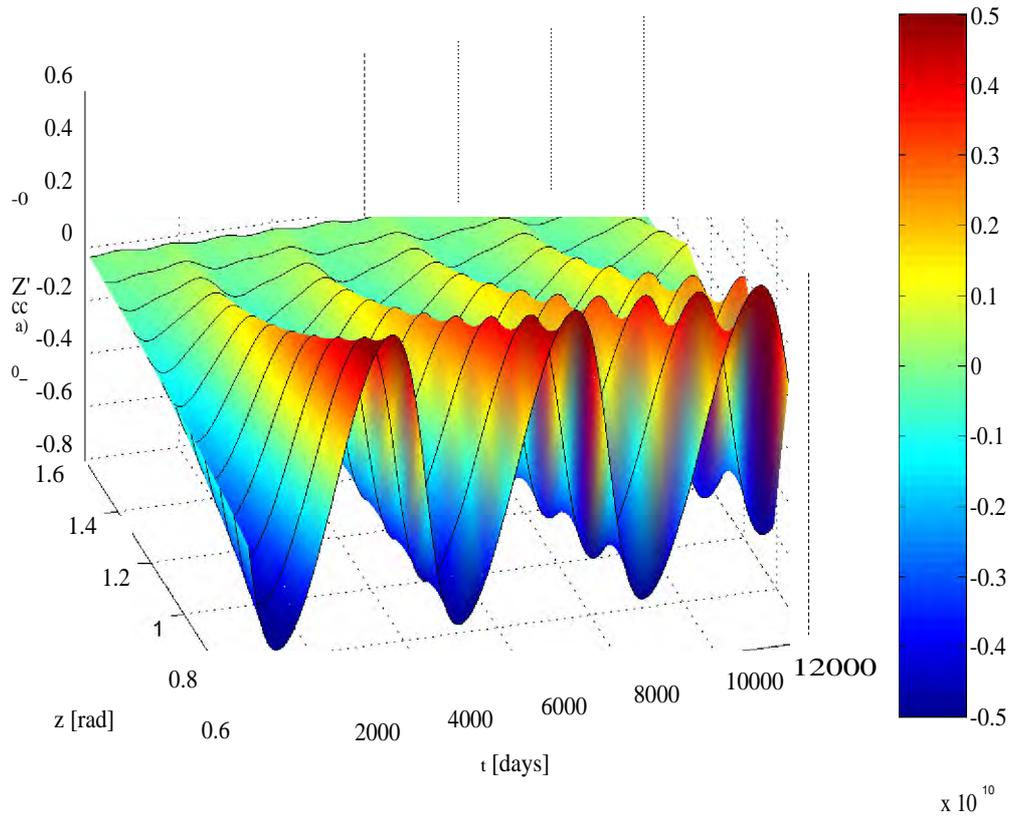


FIGURE 5.3.1. Timing residuals for a torqued, asymmetric pulsar, $O(0) = 0.1405$ rad. Top and bottom figures are Phase and AP respectively.

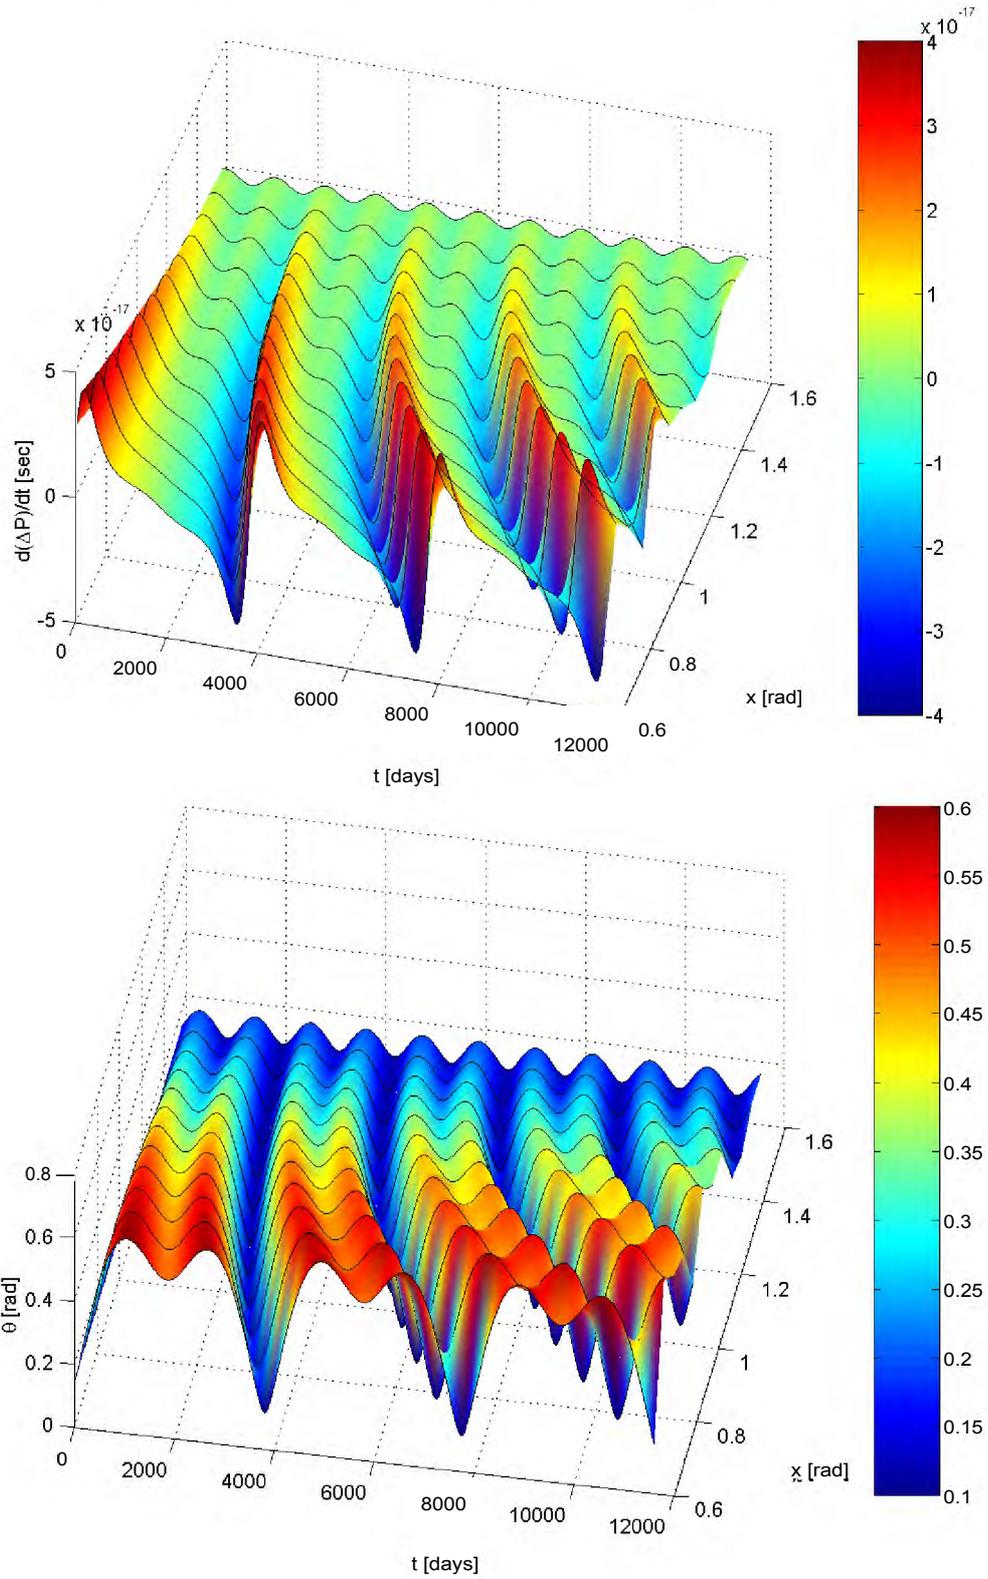


FIGURE 5.3.2. Timing residuals for a torqued, asymmetric pulsar, $\theta(0) = 0.1405$ rad. Top and bottom figures are $\frac{d(\Delta P)}{dt}$ and $e(t)$ respectively.

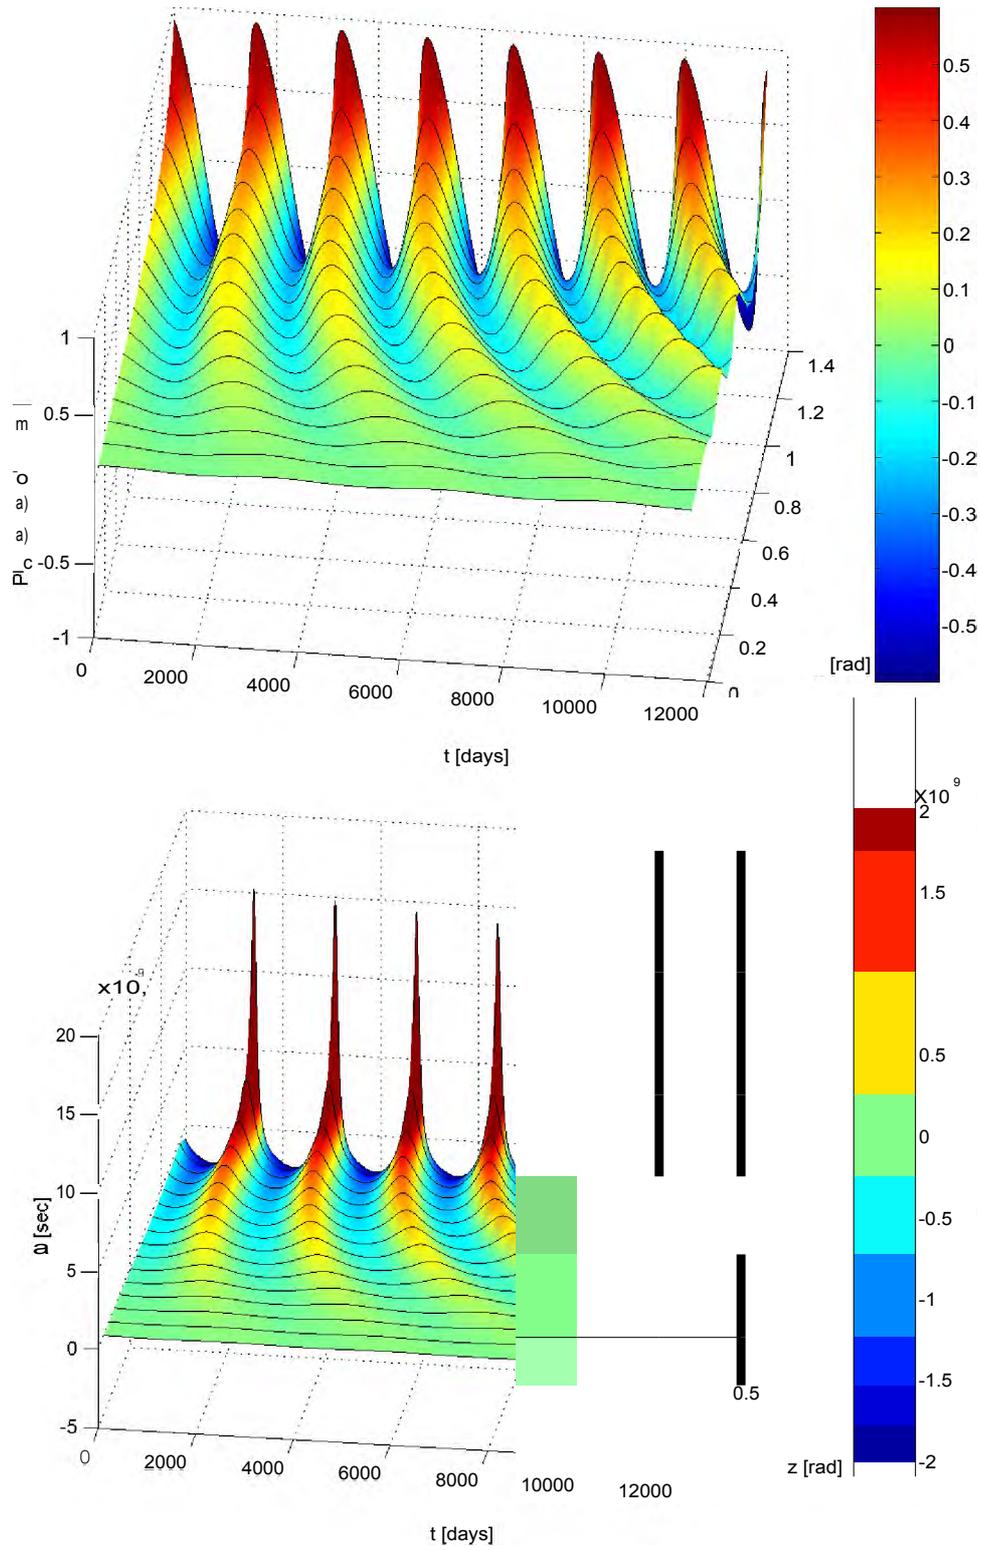


FIGURE 5.3.3. Timing residuals for a torqued, asymmetric pulsar, $\phi(0) = 1.4716$ rad. Top and bottom figures are Phase and AP respectively.

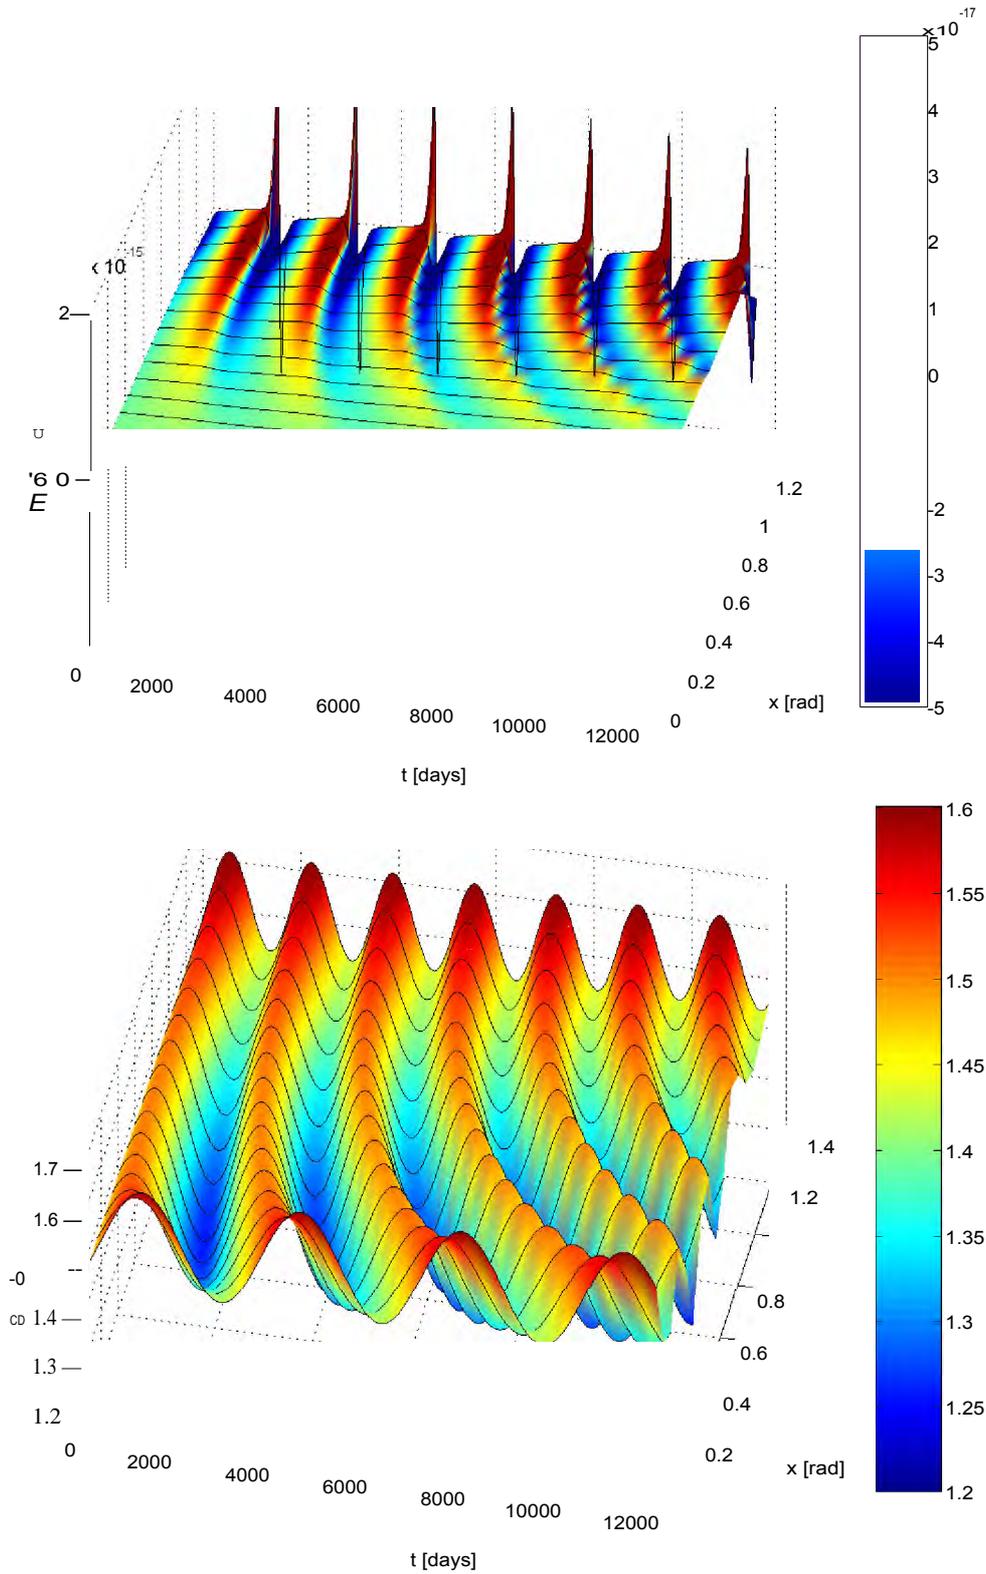


FIGURE 5.3.4. Timing residuals for a torqued, asymmetric pulsar, $\Omega M = 1.4716$ rad. Top and bottom figures are τ and $e(t)$ respectively.

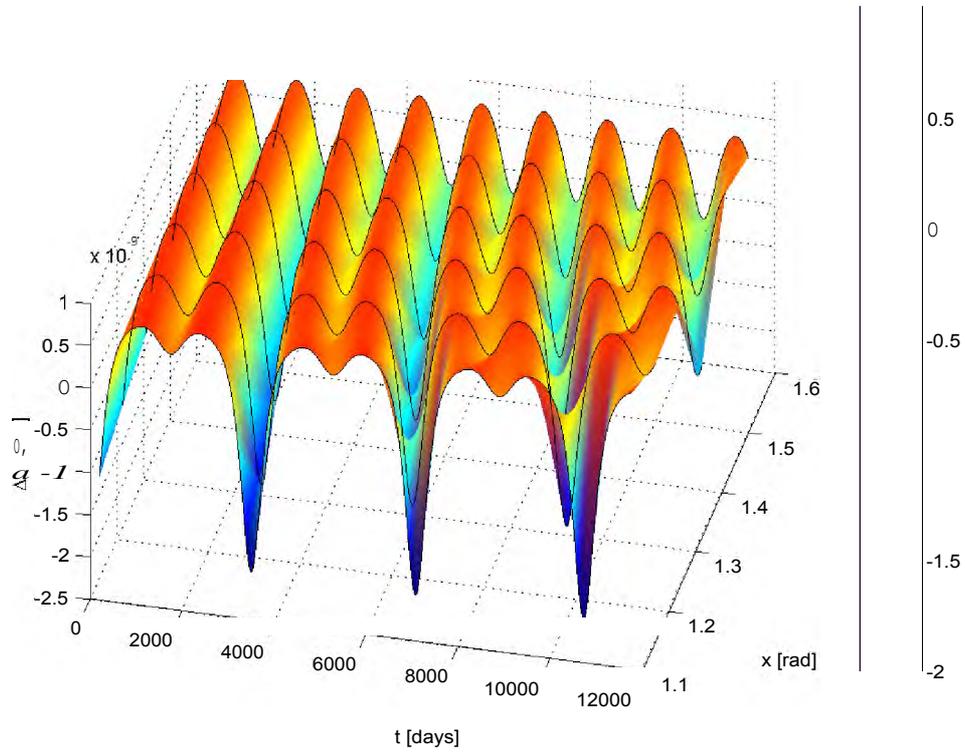
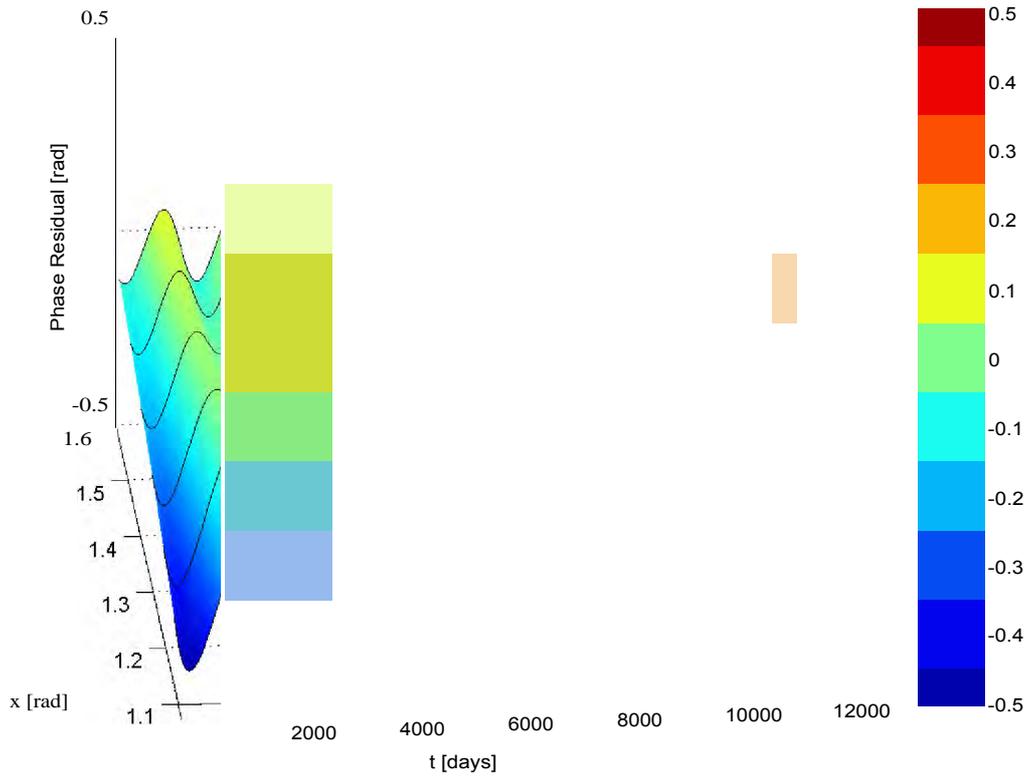


FIGURE 5.3.5. Timing residuals for a torqued, asymmetric pulsar, $OM = 0.5611$ rad. Top and bottom figures are Phase and AP respectively.

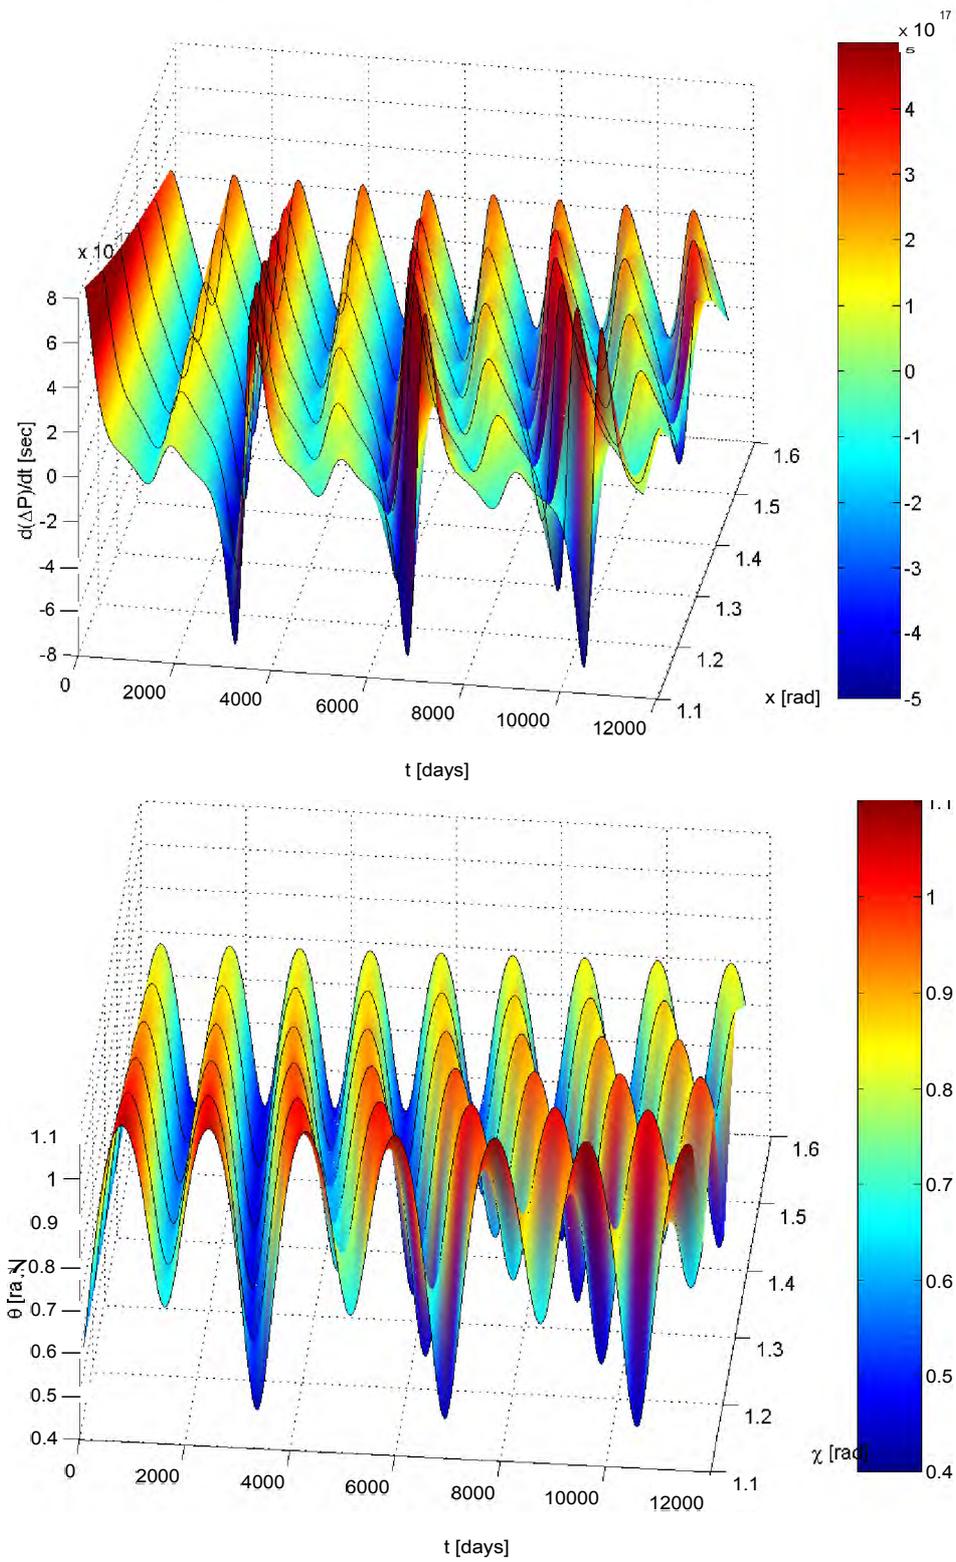


FIGURE 5.3.6. Timing residuals for a torqued, asymmetric pulsar, $\theta(0) = 0.5611$ rad. Top and bottom figures are $\frac{d\theta}{dt}$ and $\theta(t)$ respectively.

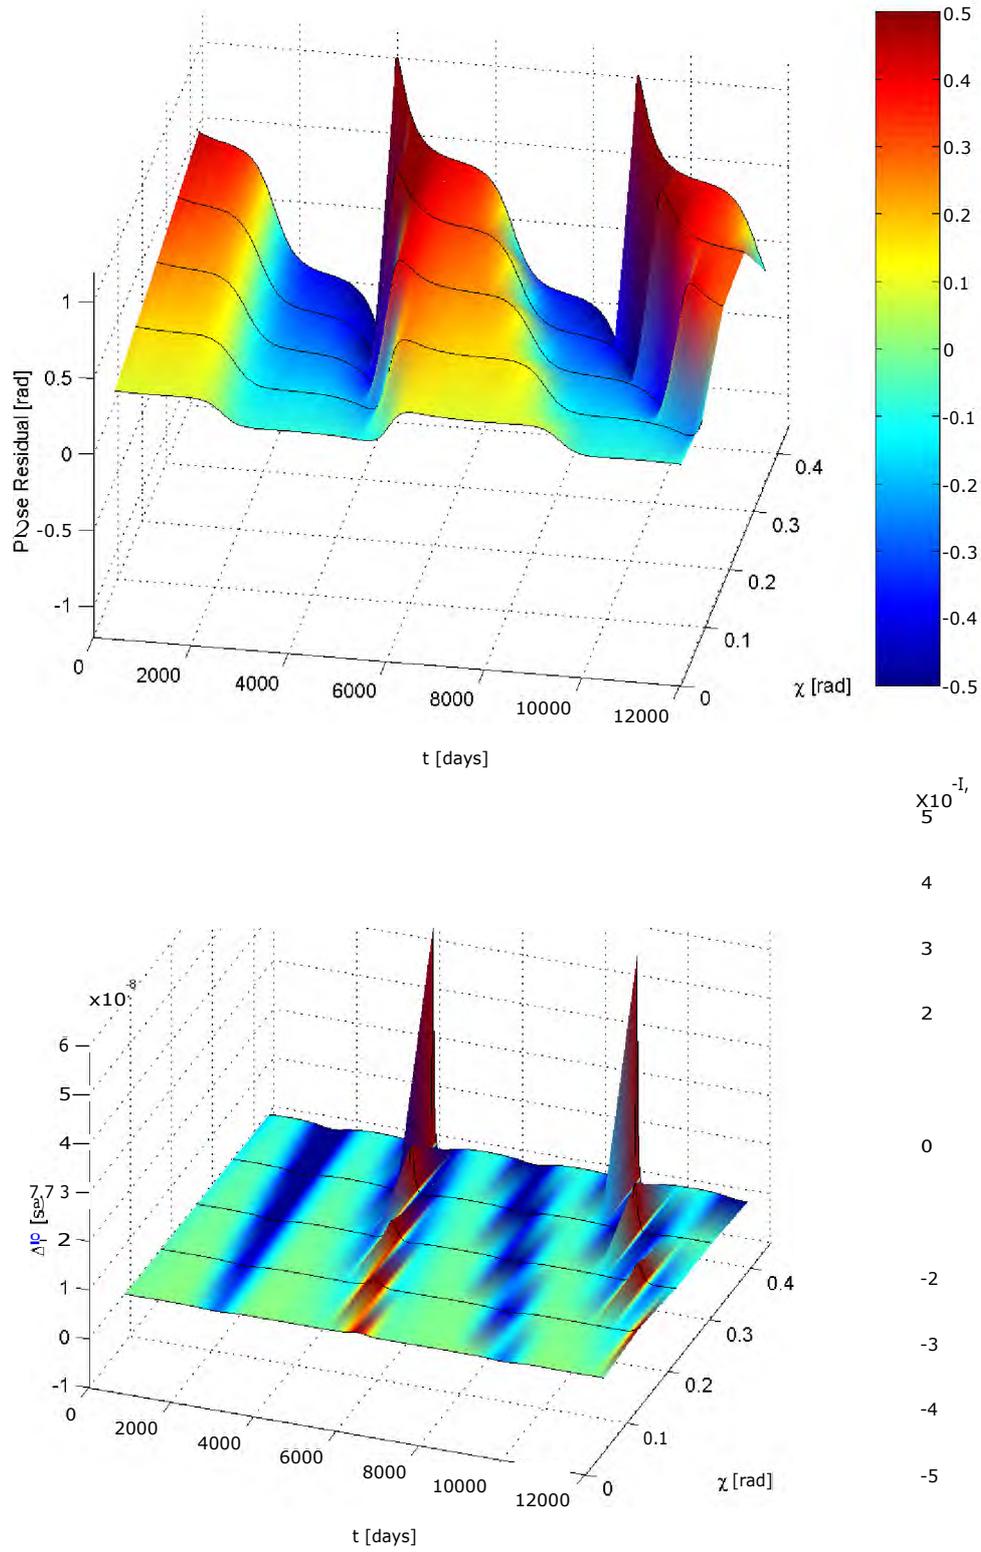


FIGURE 5.3.7. Timing residuals for a torqued, asymmetric pulsar, $\theta(0) = 1.2152$ rad. Top and bottom figures are Phase and AP respectively.

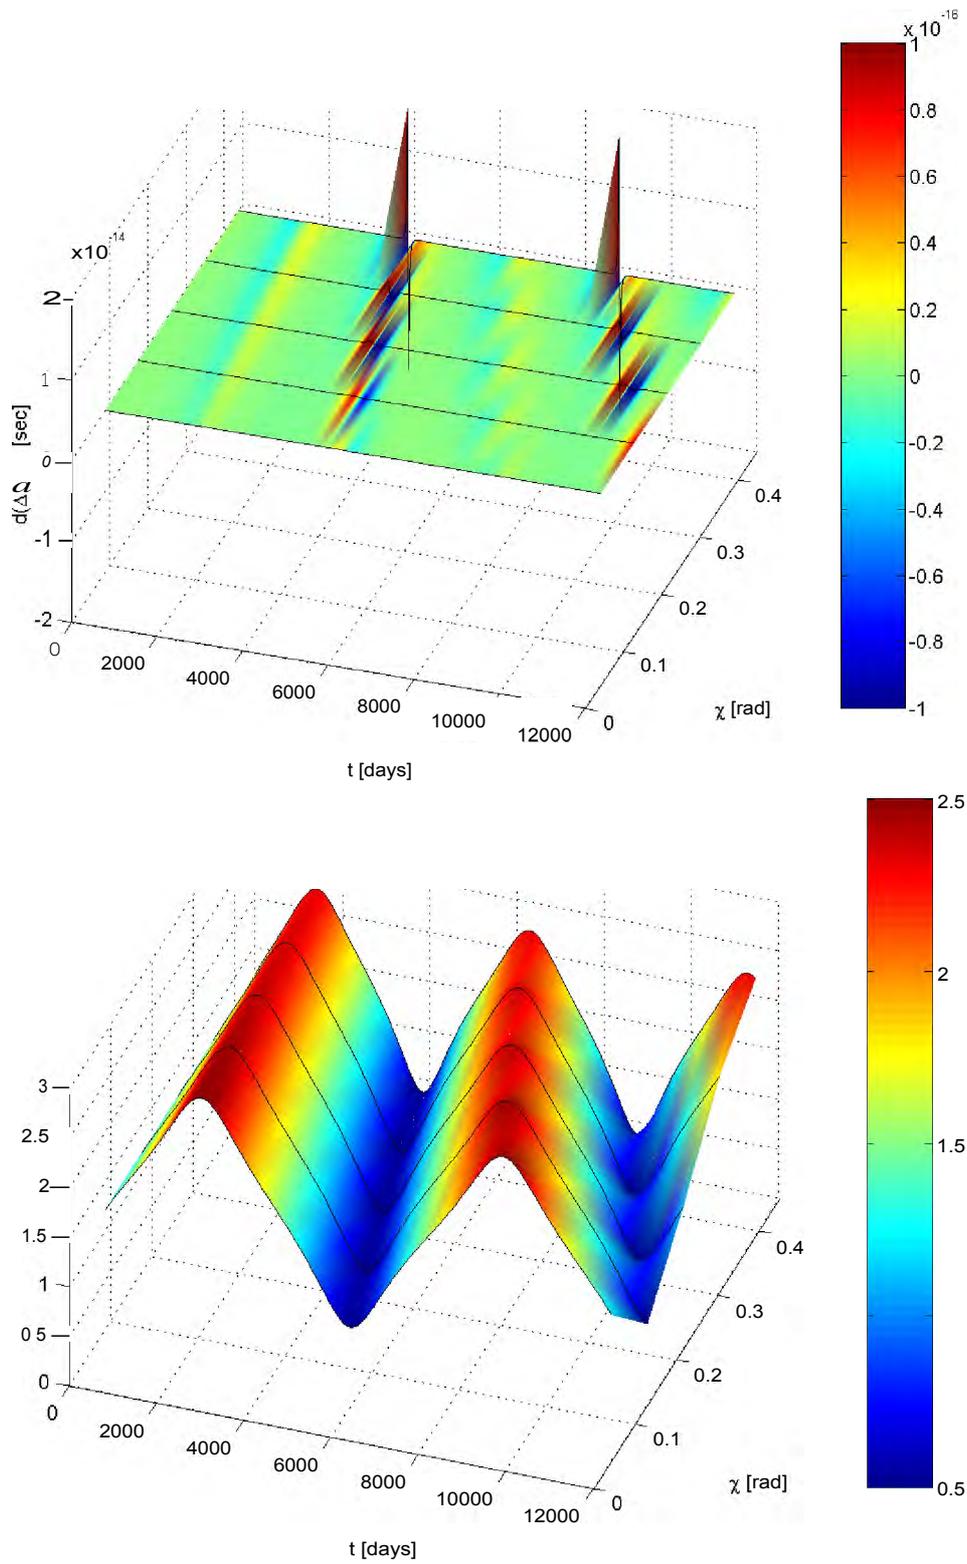


FIGURE 5.3.8. Timing residuals for a torqued, asymmetric pulsar, $OM = 1.2152$ rad. Top and bottom figures are $\underline{\underline{d(AP)}}$ and $e(t)$ respectively.

$\theta(t)$ is not constant and, in general, oscillates between two values θ_1 and θ_2 . For the case $\theta_1 < \chi < \theta_2$, calculation of timing residuals leads to a discontinuous piecewise function, shown in figure 5.3.9. As χ and $\theta(t)$ intersect, there is a sudden change in periodicity of the pulsar, from φ to $\psi + \varphi$ and vice versa. This leads to discontinuous jumps in the phase residual.

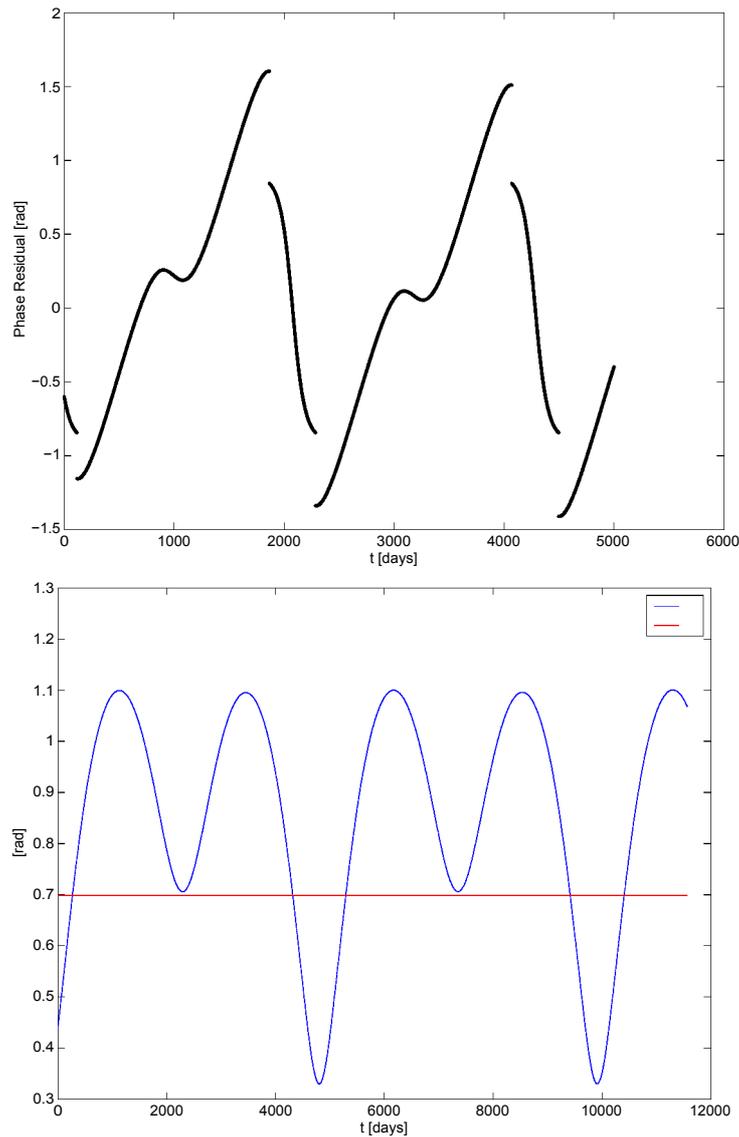


FIGURE 5.3.9. Phase residual for a torqued, asymmetric pulsar, $\theta(0) = 0.4405$ rad. Top figure is the phase residual. Bottom figure indicates $\theta(t)$ and χ .

5.4. Calculation of Observed Phase Residuals

Phase residuals are calculated at HartRAO by subtracting an ephemeris model from measured times of arrival

$$(5.4.1) \quad \Delta 4b = 4b - \left[4b_0 + v_0(t - t_0) + \frac{1}{2}\dot{v}(t - t_0)^2 + \frac{1}{6}\ddot{v}(t - t_0)^3 \right] = 4b - 4b_{spindown}$$

where $4b$ is the measured phase, $4b_{spindown}$ is a third order polynomial describing the spin down of the pulsar. This polynomial (ephemeris) is updated after each observation.

Given that the standard spin down model does not allow for an oscillating \dot{P} , the calculation of a predicted time of arrival uses a mean value $\langle \dot{P} \rangle$ about which \dot{P} oscillates. This trend would appear in the residual data after subtraction of $4b_{spindown}$. Were this trend to be subtracted, careful consideration would be needed so as not to remove ‘real’ precessional effects. In general, the changes in the period derivative are of the order of 10^{-17} for a normal pulsar. The effect of this varying torque on phase residuals has been studied numerically by Melatos (1999) [125]. Using a dipole model, Jones & Andersson (2001) [83] determined the effect of a varying torque on the phase residuals of an axisymmetric pulsar, $\theta < \chi$, to be

$$(5.4.2) \quad \Delta 4b \sim 2\kappa\omega^3 \sin \chi \cos \chi \sin \psi \frac{\theta}{\psi}$$

where κ is the torque function of the spin down law in equation 2.2.1, and ω the spin angular frequency. Using equation 5.4.2 as a simple approximation, the effect of a varying torque on timing residuals for a torqued asymmetric pulsar with $\theta(0) = 0.1405$ is shown in figure 5.4.1. On average, the magnitude of phase residual is 10% of that of the calculated residual in figure 5.3.1. It should be taken into account that the observed residuals will include this effect of \dot{P} modulation.

5.5. Candidate Precessing Pulsars in the HartRAO Catalogue

Figures 5.5.1, 5.5.2 and 5.5.3 are the observed phase residuals of 10 candidate precessing pulsars in the HartRAO catalogue.

The residuals diverge for more recent epochs because of the method used to derive the ephemeris model. PSRs B0736-40, B0740-28 and B0959-54 in figure 5.5.1 and PSR B1358-63 in figure 5.5.2 are the most effected by this divergence. For the purposes of this thesis, residual data affected by this ‘edge effect’ is excluded from the fitting.

Offset points from the main locus of points are due to errors in the predicted times of arrival and timing noise.

Care was taken to ensure that the sample of candidates was selected on the basis of intrinsic periodicity, rather than a systematic effect resulting from polynomial fitting.

5.6. Precession Fits

Calculated residuals were fitted to data by the adjustment of 6 parameters, namely

$$\begin{aligned}
 B &\rightarrow \text{magnetic field strength} \\
 Pfp &\rightarrow \text{Precession parameter such that } \varepsilon = \frac{\text{Pulsar Period}}{Pfp} \\
 \theta(0) &\rightarrow \arccos \frac{I_3}{J} \\
 \chi &\rightarrow \text{magnetic axis inclination} \\
 z &\rightarrow I(1) = I(2) \times (1 - z \times \varepsilon) \\
 pro &\rightarrow \{-1; 1\} = \{\text{prolate; oblate}\}
 \end{aligned}$$

In most cases, $z = 1$ and $pro = 1$ (oblate pulsar). $I(3)$ is taken to be $1 \times 10^{38} \text{ kg.m}^2$.

Due to the multi-variate solution space and limited computer processing power, data fitting was not performed using an automatic least-squares method. Where possible, however, a modified least-squares method was used with manual manipulation of certain parameters. Alternately, fitting was performed entirely manually. Using this procedure, full optimisation of variables was not possible. A primary goal was set to replicate the phenomenology of observed

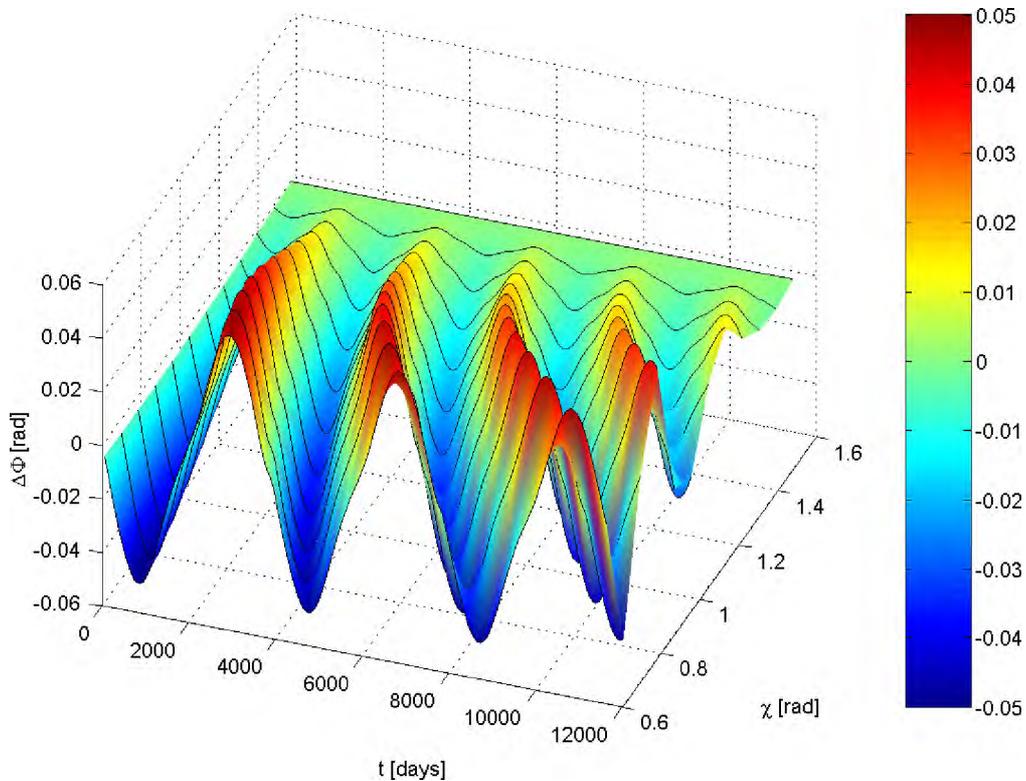


FIGURE 5.4.1. Phase residual of torqued asymmetric pulsar, $\theta(0) = 0.1405$ rad, due to varying braking torque only.

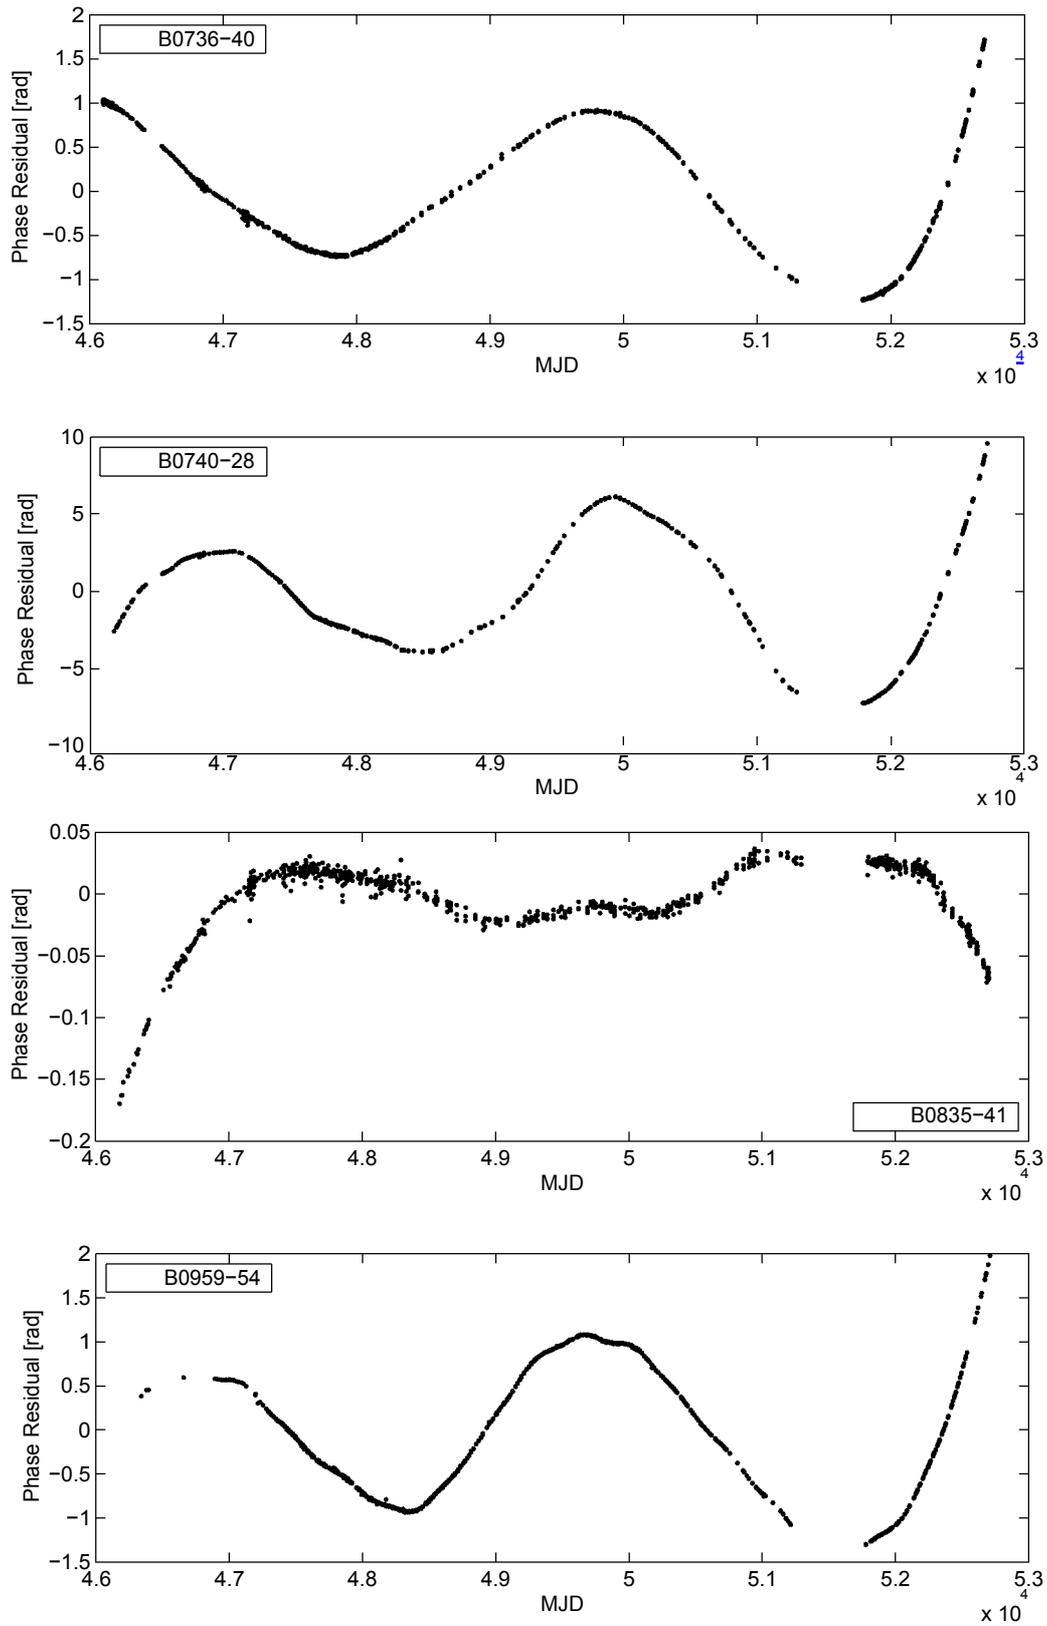


FIGURE 5.5.1. Phase residuals of, from top to bottom, PSRs B0736-40, B0740-28, B0835-41 and B0959-54.

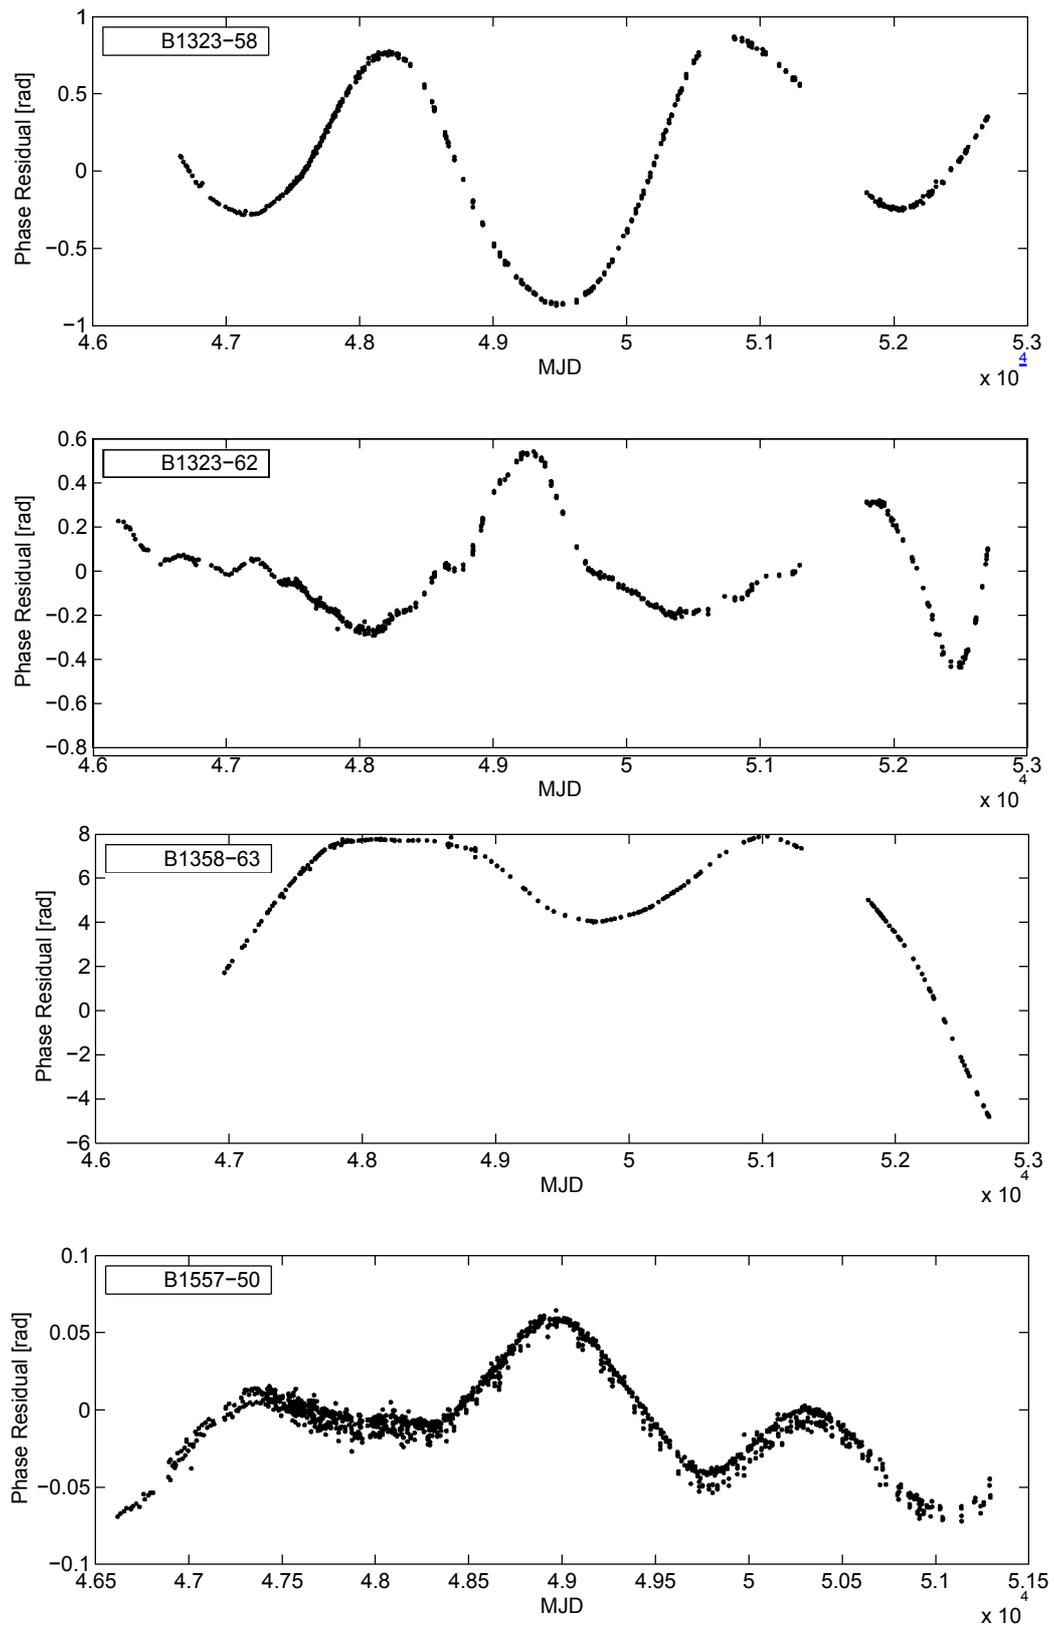


FIGURE 5.5.2. Phase residuals of, from top to bottom, PSRs B1323-58, B1323-62, B1358-63 and B1557-50.

precession behaviour (precession profile) rather than to provide an exact fit to timing residual data.

As some precession profiles show evidence of a linearly decreasing trend, a 1st order polynomial fit is included in the precession model to force stationarity. The measured residuals were fitted using a piecewise, least-squares cubic spline to determine values for ΔP and $\frac{d(\Delta P)}{dt}$.

5.6.1. B1642-03. The timing residuals exhibit convincing evidence of precession in this pulsar, and it has previously been identified by other authors as a precession candidate [156]. PSR B1642-03 has been fitted by manual adjustment of parameters, listed in table 5.6.1. Derived pulsar parameters are listed in table 5.6.2. The results of the fitted model are shown in Figures 5.6.1, 5.6.2 and appendix B.1.

The composite ‘double-hump’ shape in the measured phase residual profile, shown in figure 5.6.1, implies a certain degree of asymmetry with $\chi > \theta(0)$ (see figure 5.3.1). The larger amplitude of the leading ‘hump’, in comparison to the lagging ‘hump’, is typical of a prolate body. The calculated $\langle P \rangle$ is an order of magnitude greater than the published value of 1.7777×10^{-15} (listed in chapter 2). This could be due either to an incorrect coupling distance between the near field and far field radiation zones (χ_0 is kept constant for fitting), or non-dipolar magnetic field structures.

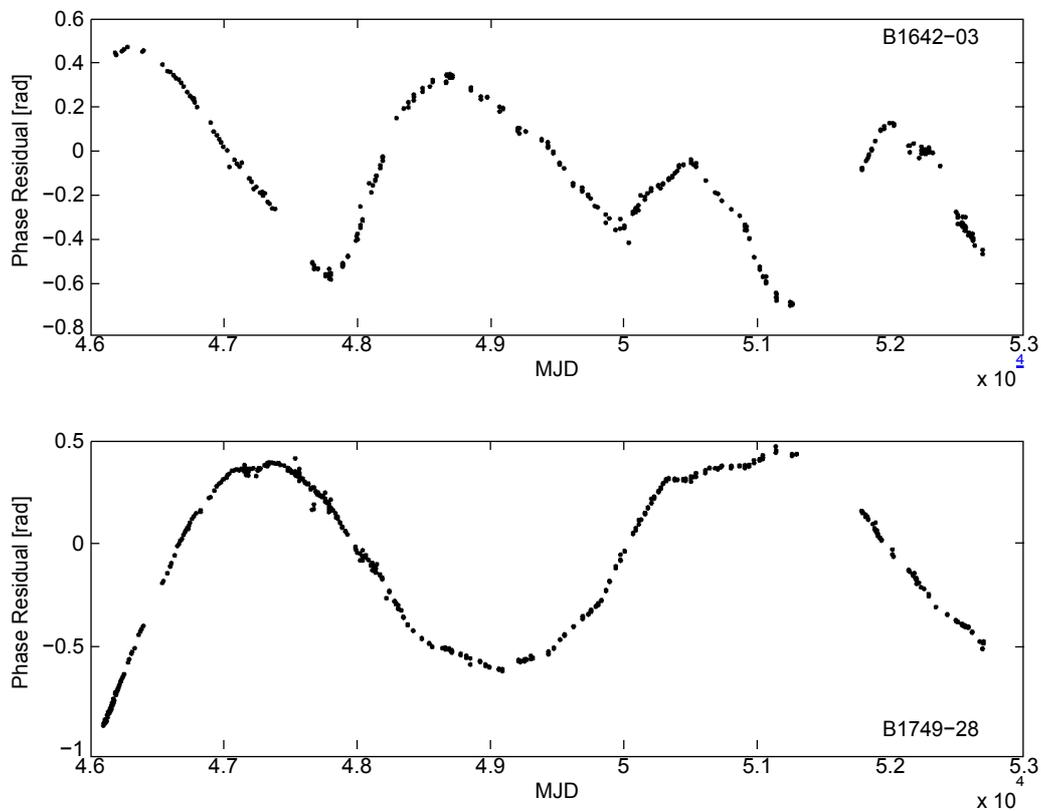


FIGURE 5.5.3. Phase residuals of, from top to bottom, PSRs B1642-03 and B1749-28

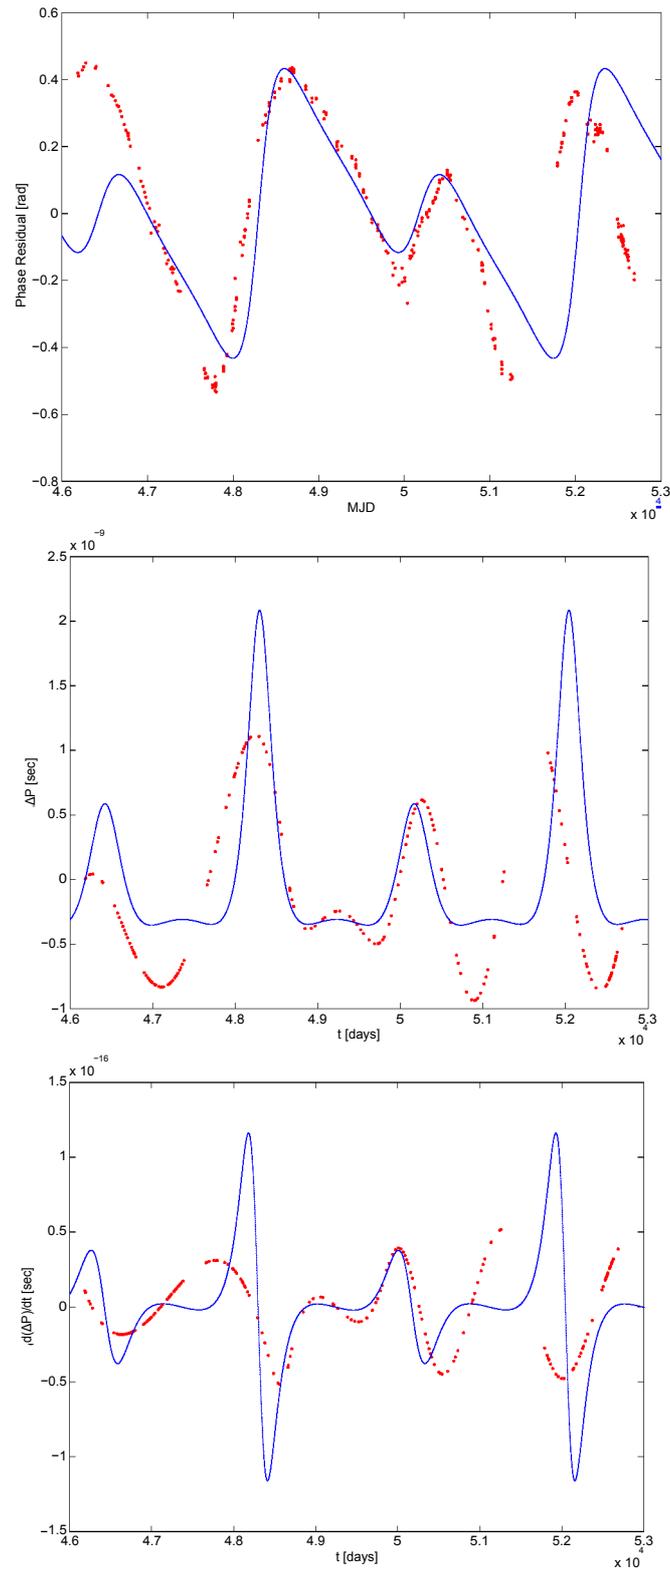


FIGURE 5.6.1. Manual fit of observed timing residuals of PSR B1642-03. From top to bottom are Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

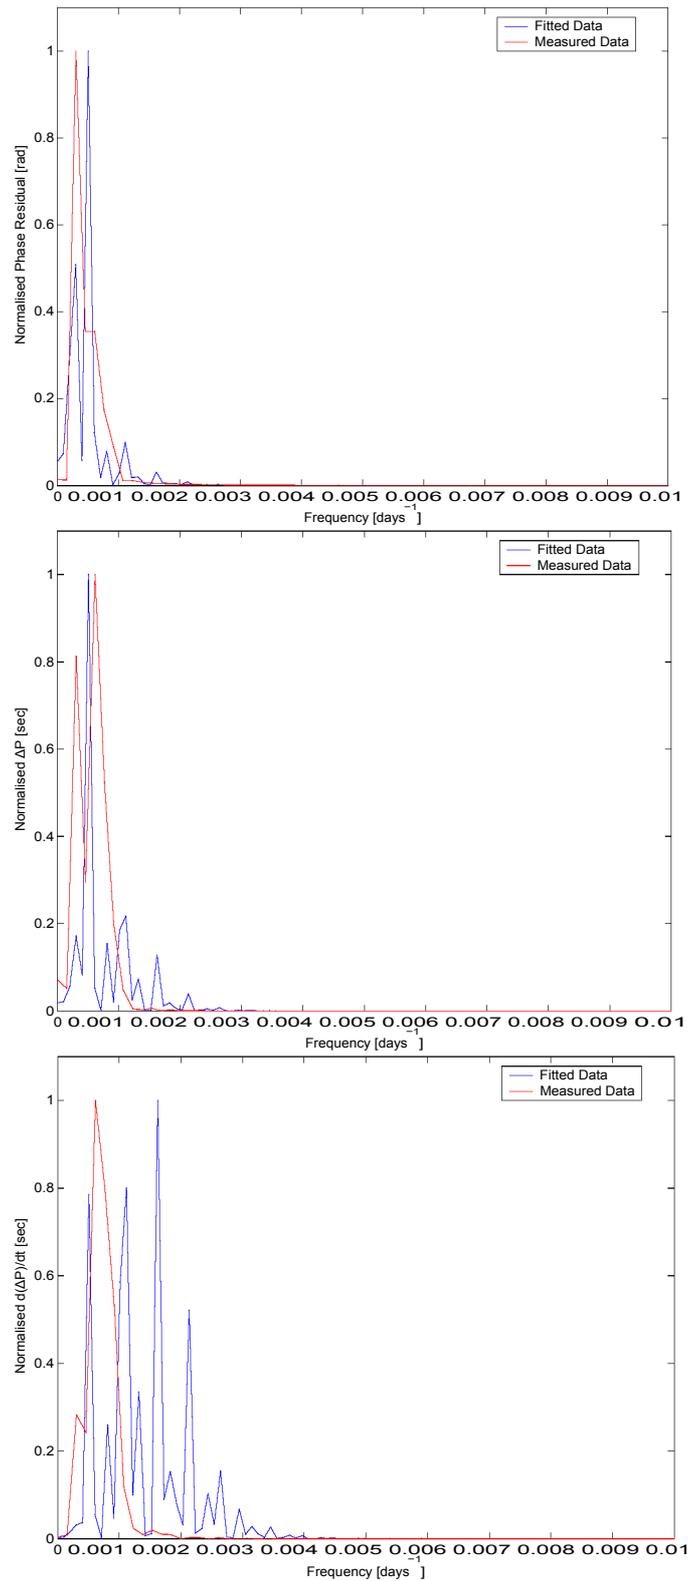


FIGURE 5.6.2. Manual fit of observed timing residuals of PSR B1642-03. From top to bottom are spectrums of Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

To determine whether the torqued precession model could provide a good fit of timing residual data, a Chi-squared surface was calculated. Due to the computing power required, only two parameters were fitted at a time, namely $[B - Pfp]$ and $[\chi - \theta(0)]$. While not entirely mutually exclusive, the dominant effect of loci in the $[B - Pfp]$ plane is to change the precession period, whilst loci in the $[\chi - \theta(0)]$ plane change the precession profile shape. This produced two surfaces, shown in figure 5.6.3, for constant $[\chi - \theta(0)]$ and constant $[B - Pfp]$. The large deviations from the moderately flat surfaces are a result of calculated phase residuals for $\theta(t_1) < \chi < \theta(t_2)$. Simulation of both sets of parameters shows no evidence of an obvious global minimum in the Chi-squared surfaces. Given the restricted subspace of parameters used in the simulations, it may be that the fit could be improved by a more rigorous search in the entire parameter space. Limited computing resources do not allow this. Furthermore, edge effects, an inaccurate ephemeris and the effect of \dot{P} induced timing residual modulations (see section 5.4) could all place a limit on the achievable RMS error.

The best fit parameters obtained from the simulations were used to attempt a further fit of the timing residual data, shown in Figures 5.6.4 and 5.6.5. The fitted parameters and pulsar parameters are shown in tables 5.6.3 and 5.6.4.

TABLE 5.6.1. PSR B1642-03 fitting parameters

| Parameter | Value |
|-------------|-----------------------|
| B | 1.1×10^9 [T] |
| Pfp | 2100 [days] |
| $\theta(0)$ | 0.937 [rad] |
| χ | 80.1 [deg] |
| z | 1 |
| pro | -1 (prolate) |

TABLE 5.6.2. PSR B1642-03 derived pulsar parameters

| Parameter | Value |
|---------------------|--|
| ϵ | $-2.136740298456790 \times 10^{-9}$ |
| $[\mu]$ | $[1.00000000427348, 1.00000000213674, 1] \times 10^{38}$ [kg.m] ² |
| $[\omega(0)]$ | $[9.23420107691926, 9.23420107691926, 9.59774907941657]$ [rad.s] ⁻¹ |
| $\langle P \rangle$ | 6.2544×10^{-4} [s.s] ⁻¹ |

TABLE 5.6.3. PSR B1642-03 best fit fitting parameters

| Parameter | Value |
|-------------|-----------------------|
| B | 2.1×10^9 [T] |
| Pfp | 1900 [days] |
| χ | 83.1 [deg] |
| $\theta(0)$ | 0.937 [rad] |
| z | 1 |
| pro | -1 (prolate) |

The phase residual spectrum, shown in figure 5.6.5, seems to provide a better fit to the measured data than the spectrum shown in figure 5.6.2. However, the amplitude of phase residual, shown in figure 5.6.4, is significantly smaller than the measured data. The model does not allow for a simple increase in the amplitude without considerable alteration to the shape of

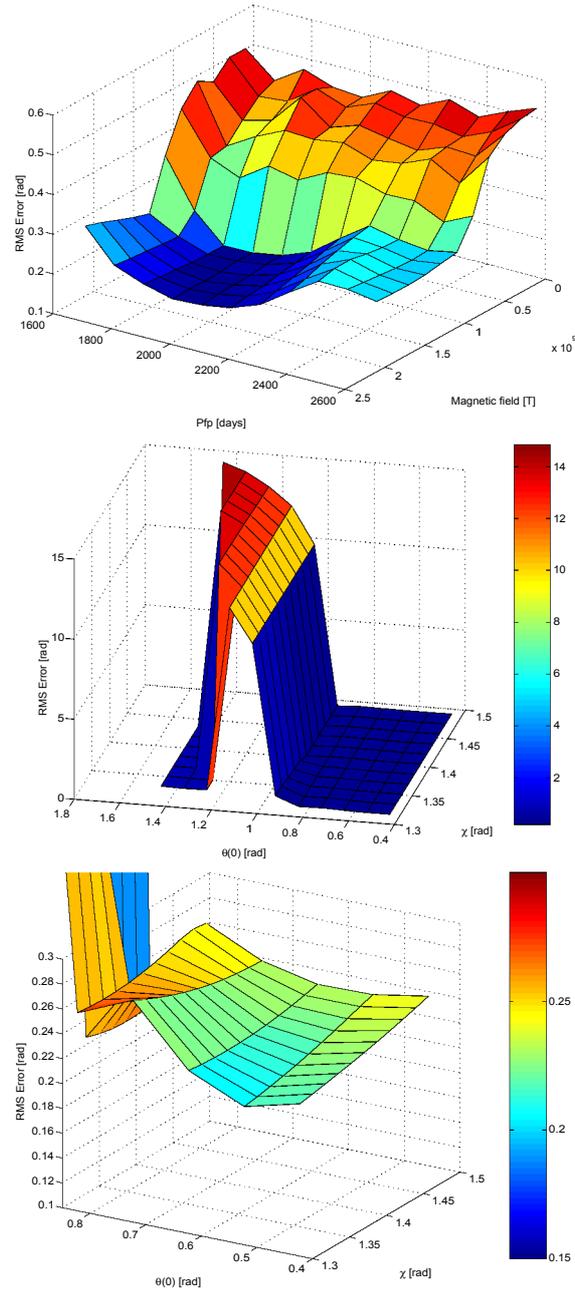


FIGURE 5.6.3. Chi-squared surfaces for fitting of observed timing residuals of PSR B1642-03. Top figure is for constant $\chi = 80.1$ [deg], $\theta(0) = 0.937$ [rad], middle figure is for constant $\mathbf{B} = 1.1 \times 10^9$ [T], $\mathbf{Pfp} = 2100$ [days] with a closeup in the bottom figure.

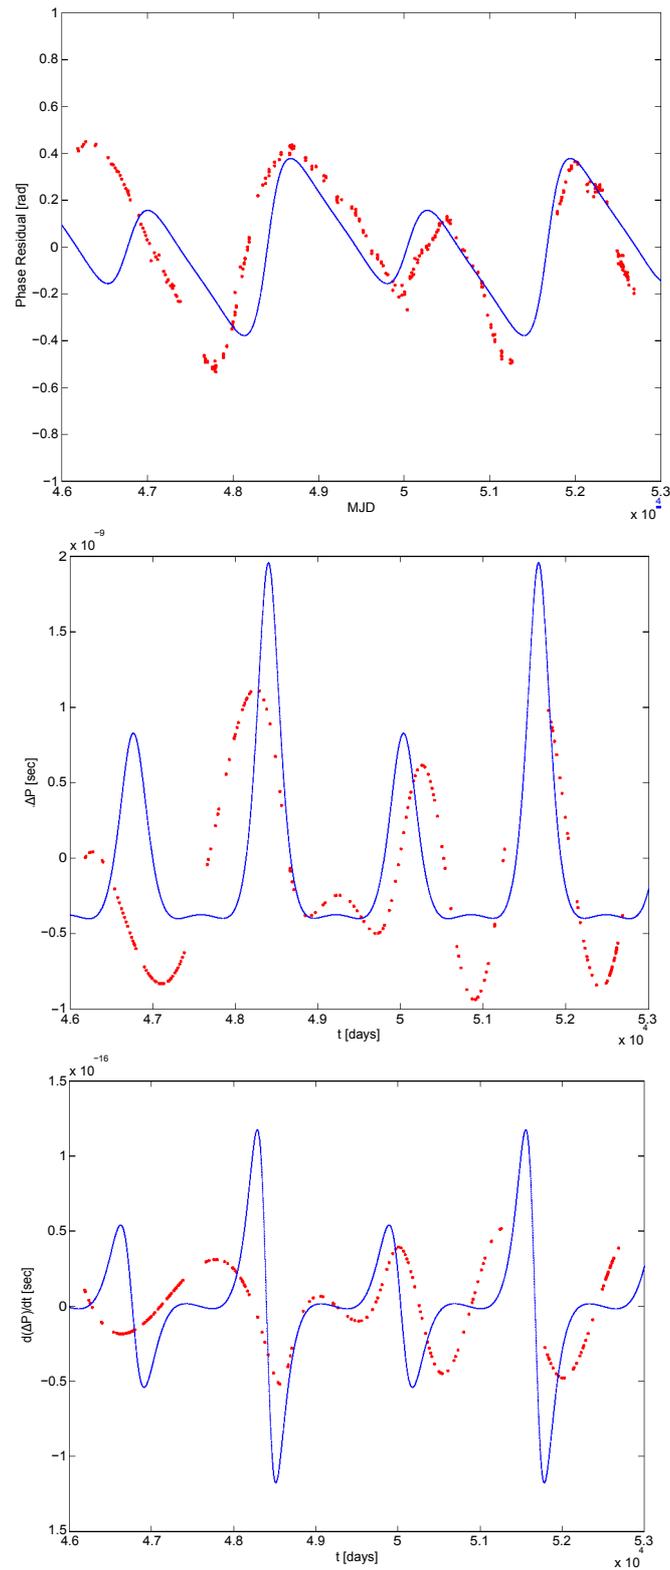


FIGURE 5.6.4. Best fit of observed timing residuals of PSR B1642-03. From top to bottom are Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

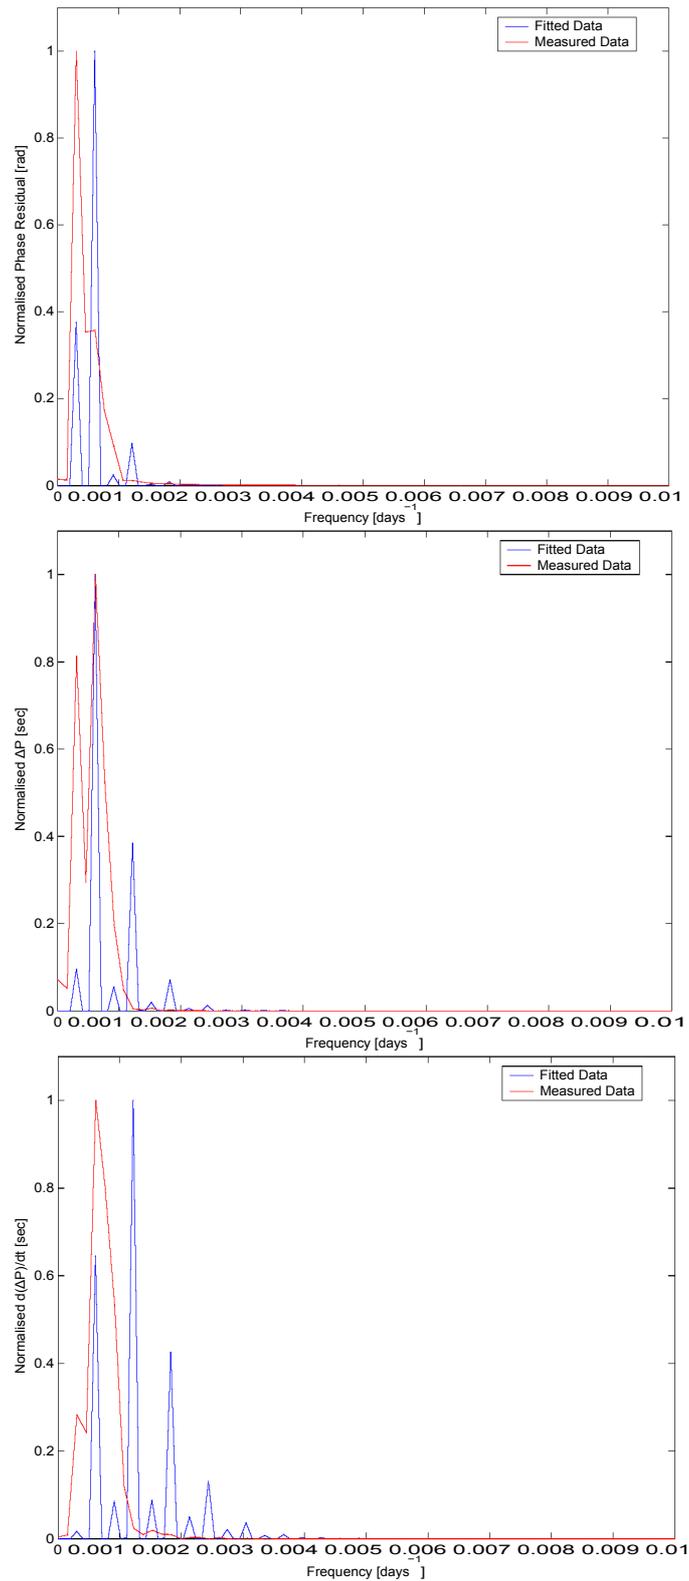


FIGURE 5.6.5. Best fit of observed timing residuals of PSR B1642-03. From top to bottom are spectrums of Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

the profile. The larger amplitude seen in the measured data could be a result of P' modulation (see figure showing calculated values for P' in appendix B.1).

5.6.2. B1323-58. Fits of observed timing residuals of PSR B1323-58 using a torqued precession model are shown in Figures 5.6.6 and 5.6.7, and appendix B.2. Fitted and derived pulsar parameters are shown in tables 5.6.5 and 5.6.6.

The observed phase residuals in figure 5.6.6 show a ‘double-humped’ profile, which implies asymmetry in the pulsar, with $\chi > 0(0)$ (See figure 5.3.1). The increase in amplitude of the observed residuals for the epoch after MJD 50000 is unexplained by the precession model. A longer data span will indicate whether this is the result of an inaccurate ephemeris, or a real effect. Following the same procedure as used in section 5.6.1, the search for minimums in the Chi-squared surfaces for the $[B - Pfp]$ and $[\chi - 0(0)]$ parameters resulted in best fits of the observed timing residuals, shown in appendix B.2. Compared to tables 5.6.5 and 5.6.6, the best fit procedure resulted in the adjustment of the parameters $B = 0.6 \times 10^8$ [T], $\chi = 89.99$ [deg] and $\langle P \rangle = 1.835 \times 10^{-16}$ [s.s⁻¹].

5.6.3. B0736-40. Figures 5.6.8 and 5.6.9, and appendix B.3, are the results are a best fit procedure of observed timing residuals of PSR B0736-40. The Chi-squared surfaces, adjusting the parameters $[B - Pfp]$ and $[\chi - 0(0)]$, are shown in appendix B.3.

TABLE 5.6.4. PSR B1642-03 best fit derived pulsar parameters

| Parameter | Value |
|---------------------|---|
| ϵ | $-2.361660329873295 \times 10^{-9}$ |
| $[\tilde{M}]$ | $[1.00000000472332, 1.00000000236166, 1] \times 10^{38}$ [kg.m ²] |
| $[\omega(0)]$ | $[9.23420107582665, 9.23420107582665, 9.59774908151903]$ [rad.s ⁻¹] |
| $\langle P \rangle$ | 2.29×10^{-13} [s.s ⁻¹] |

TABLE 5.6.5. PSR B1323-58 fitting parameters

| Parameter | Value |
|-------------------------------|-----------------------|
| B | 2.6×10^8 [T] |
| Pfp | 49000 [days] |
| χ | 89.89 [deg] |
| $\theta(0)$ | 0.022 [rad] |
| z | 100 |
| pro | 1 (oblate) |

TABLE 5.6.6. PSR B1323-58 derived pulsar parameters

| Parameter | Value |
|---------------------|--|
| ϵ | $1.129041877345285 \times 10^{-10}$ |
| $[\tilde{M}]$ | $[0.999999988596677, 0.99999999887096, 1] \times 10^{38}$ [kg.m ²] |
| $[\omega(0)]$ | $[0.20447143599235, 0.20447143599235, 13.14180102584652]$ [rad.s ⁻¹] |
| $\langle P \rangle$ | 3.446×10^{-15} [s.s ⁻¹] |

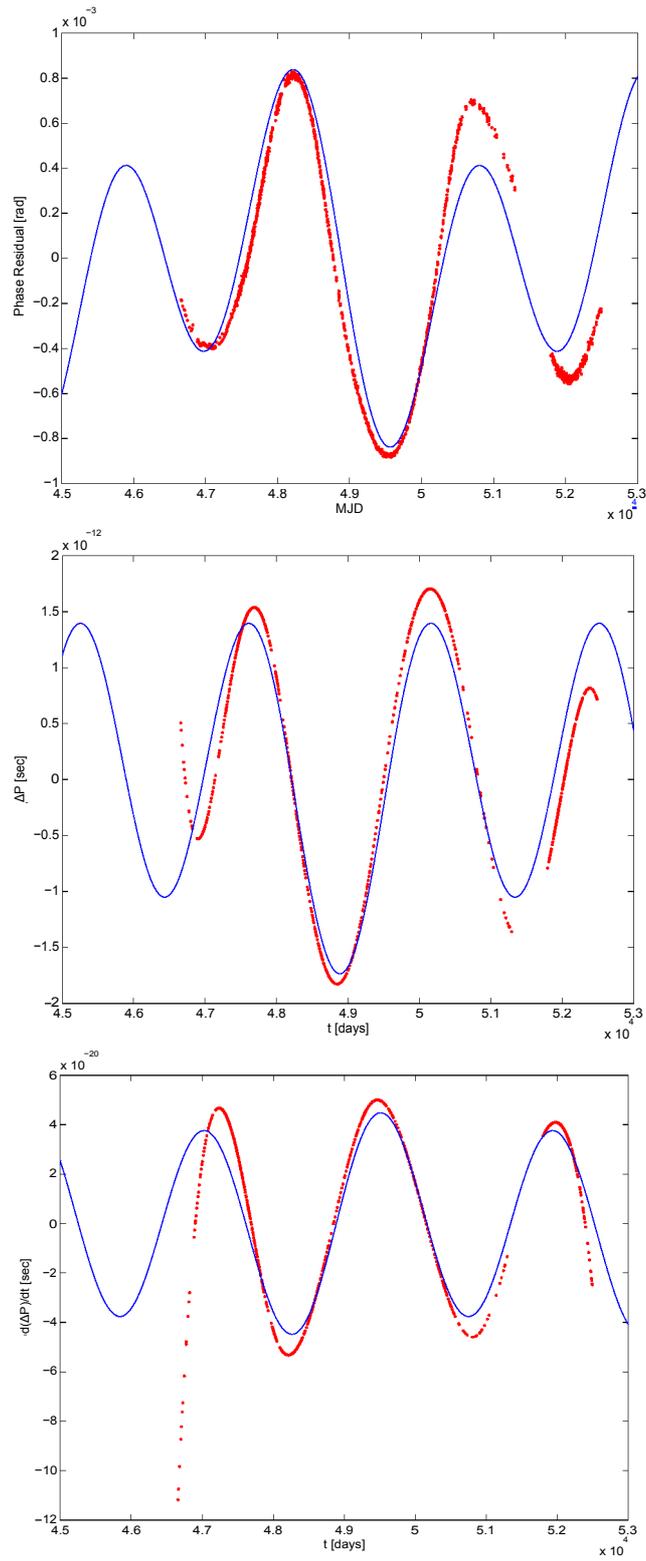


FIGURE 5.6.6. Manual fit of observed timing residuals of PSR B1323-58. From top to bottom are Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

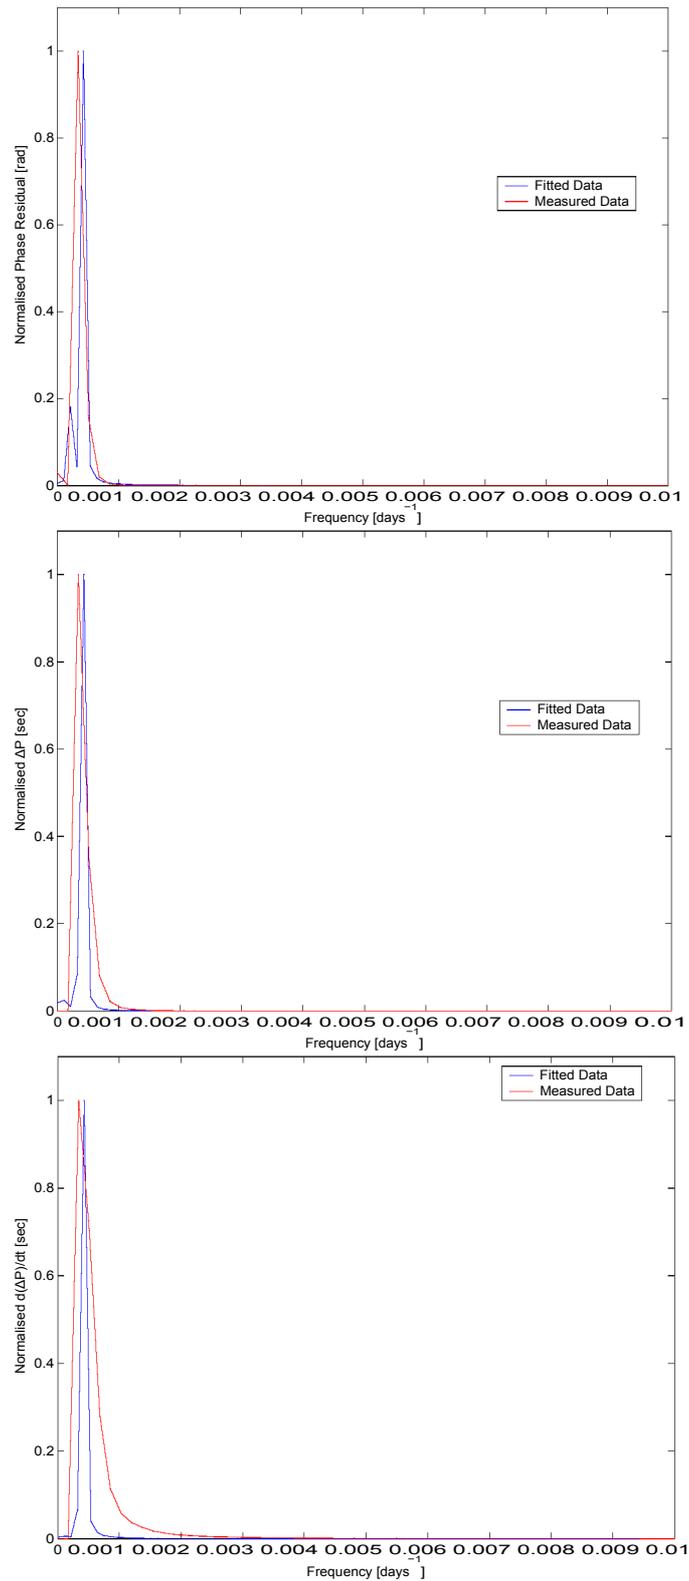


FIGURE 5.6.7. Manual fit of observed timing residuals of PSR B1323-58. From top to bottom are spectrums of Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$

The amplitude of the observed phase residual, as well as the symmetric properties of the precession profile, in figure 5.6.8 cannot be reproduced by either the torqued or untorqued precession model. The observed phase residual behaviour is highly likely to be an artefact of polynomial fitting.

The Fourier spectrum of the fitted timing residuals show strong harmonic structure in phase, ΔP and $\frac{d(\Delta P)}{dt}$, whereas the observed data has wide, sloping spectrums. The fundamental frequency is accurately reproduced, but higher harmonics are distorted in the measured data due to timing noise.

TABLE 5.6.7. PSR B0736-40 best fit fitting parameters

| Parameter | Value |
|-------------|--------------------------|
| B | 1.5×10^{-9} [T] |
| P fp | 3800 [days] |
| χ | 12.00 [deg] |
| $\theta(0)$ | 0.096 [rad] |
| z | 1 |
| pro | 1 (oblate) |

TABLE 5.6.8. PSR B0736-40 best fit derived pulsar parameters

| Parameter | Value |
|---------------------|--|
| ϵ | $1.084837148677951 \times 10^{-9}$ |
| [I] | $[0.999999997830326, 0.999999998915163, 1] \times 10^{38}$ [kg.m ²] |
| $[\omega(0)]$ | $[1.13587499860042, 1.13587499860042, 16.68158357164555]$ [rad.s ⁻¹] |
| $\langle P \rangle$ | 1.258×10^{-14} [s.s] |

5.6.4. B1557-50. Observed timing residuals of PSR B1557-50 were fitted by the manual adjustment of parameters, shown in Figures 5.6.10 and 5.6.11, and appendix B.4. Fitting and derived pulsar parameters are shown in tables 5.6.9 and 5.6.10. The ‘triple-humped’ profile in the observed timing residual, shown in figure 5.6.10, cannot be explained by the current version of the torqued precession model. The use of a shorter data span, which includes data from earlier epochs up to MJD 50200, was used for fitting procedures.

TABLE 5.6.9. PSR B1557-50 fitting parameters

| Parameter | Value |
|-------------|---------------------------|
| B | 2.25×10^{-8} [T] |
| P fp | 4460 [days] |
| χ | 89.84 [deg] |
| $\theta(0)$ | 0.010 [rad] |
| z | 1 |
| pro | 1 (oblate) |

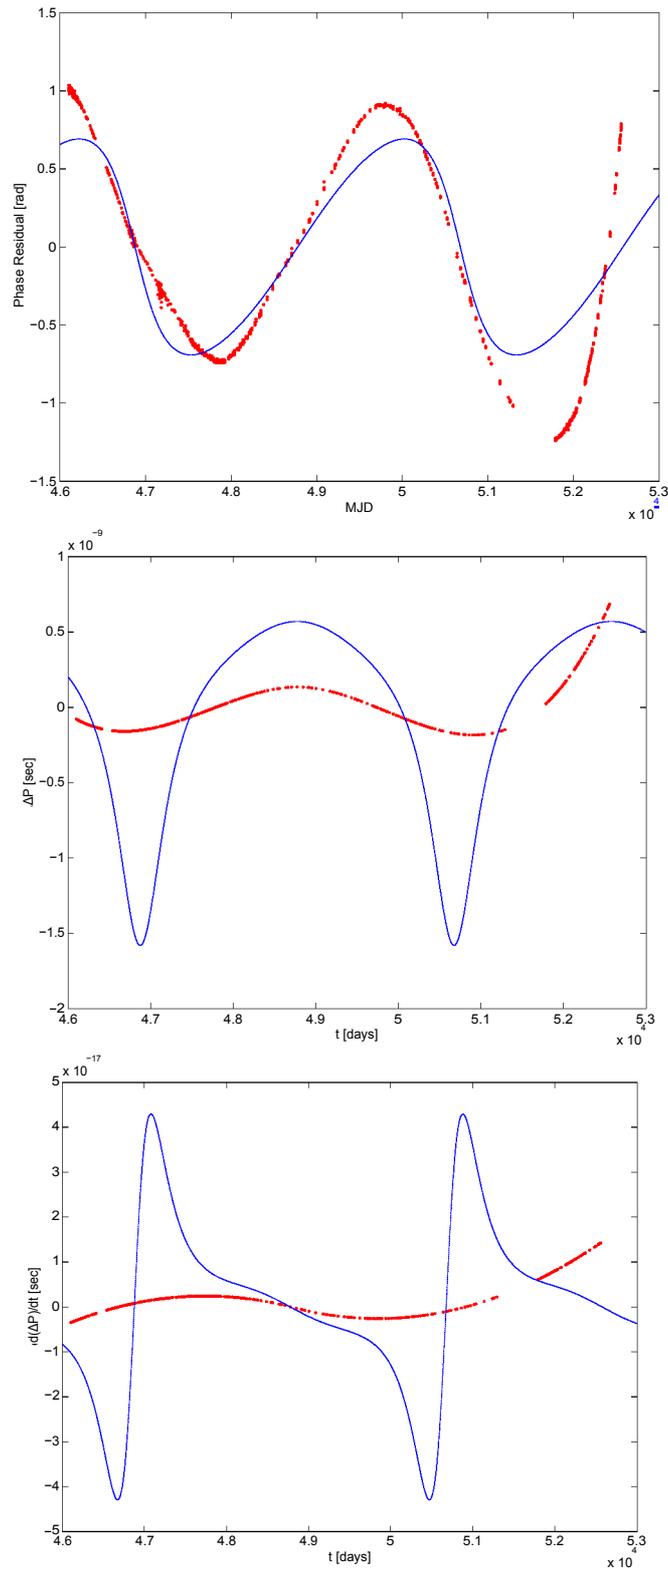


FIGURE 5.6.8. Manual fit of observed timing residuals of PSR B0736-40. From top to bottom are Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

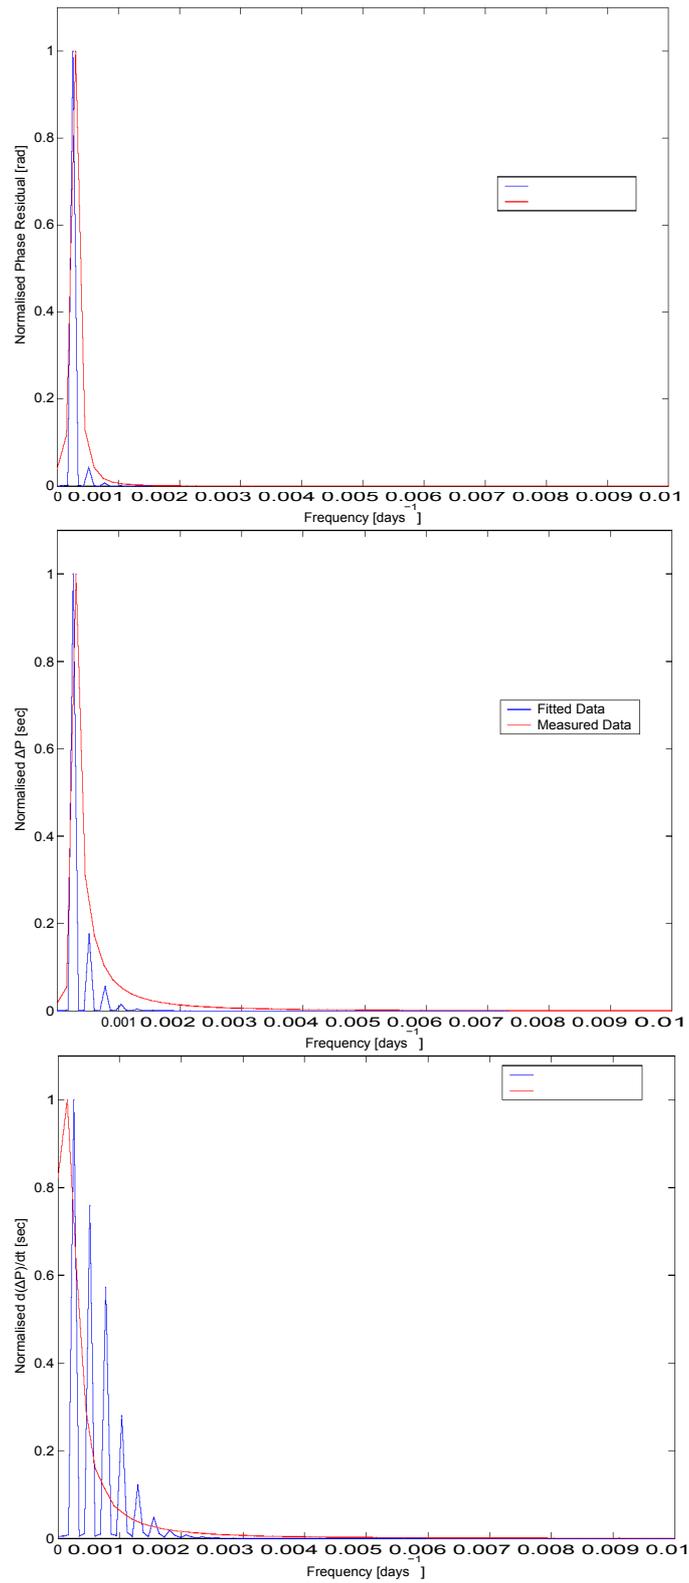


FIGURE 5.6.9. Manual fit of observed timing residuals of PSR B0736-40. From top to bottom are spectrums of Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

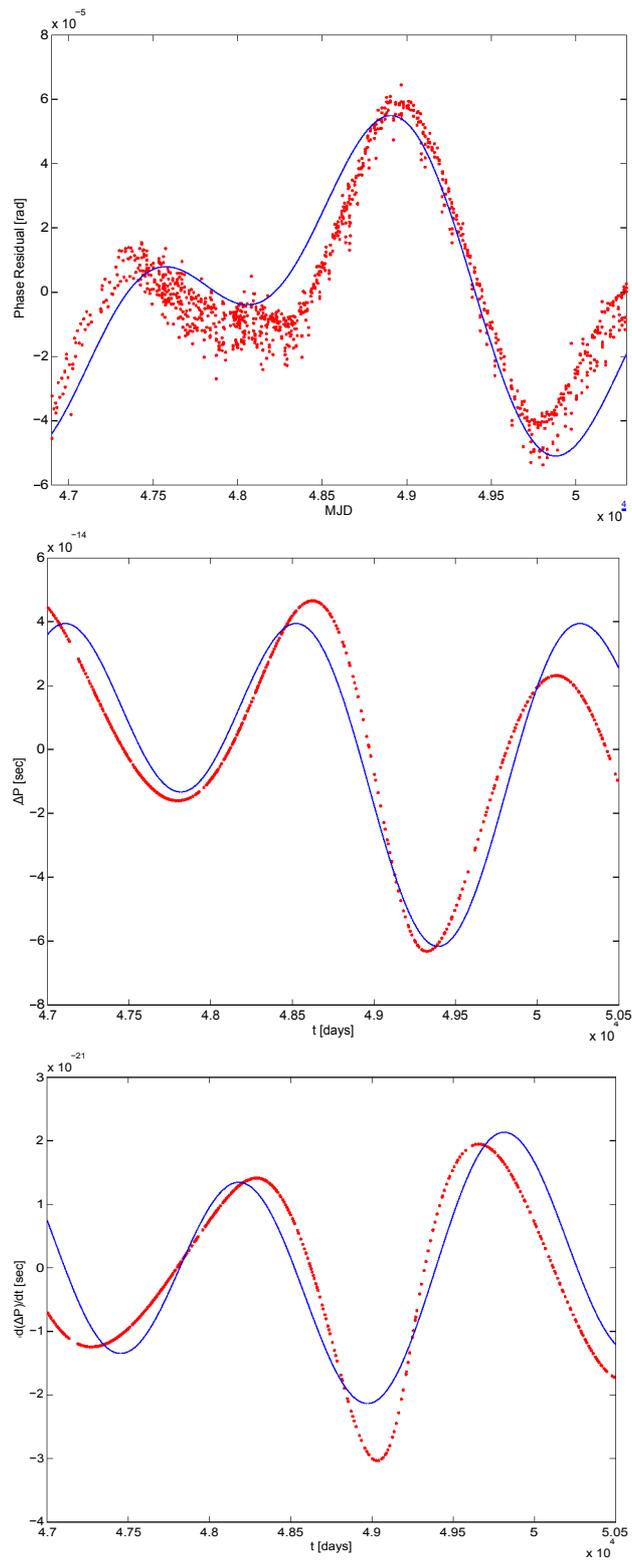


FIGURE 5.6.10. Manual fit of observed timing residuals of PSR B1557-50. From top to bottom are Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

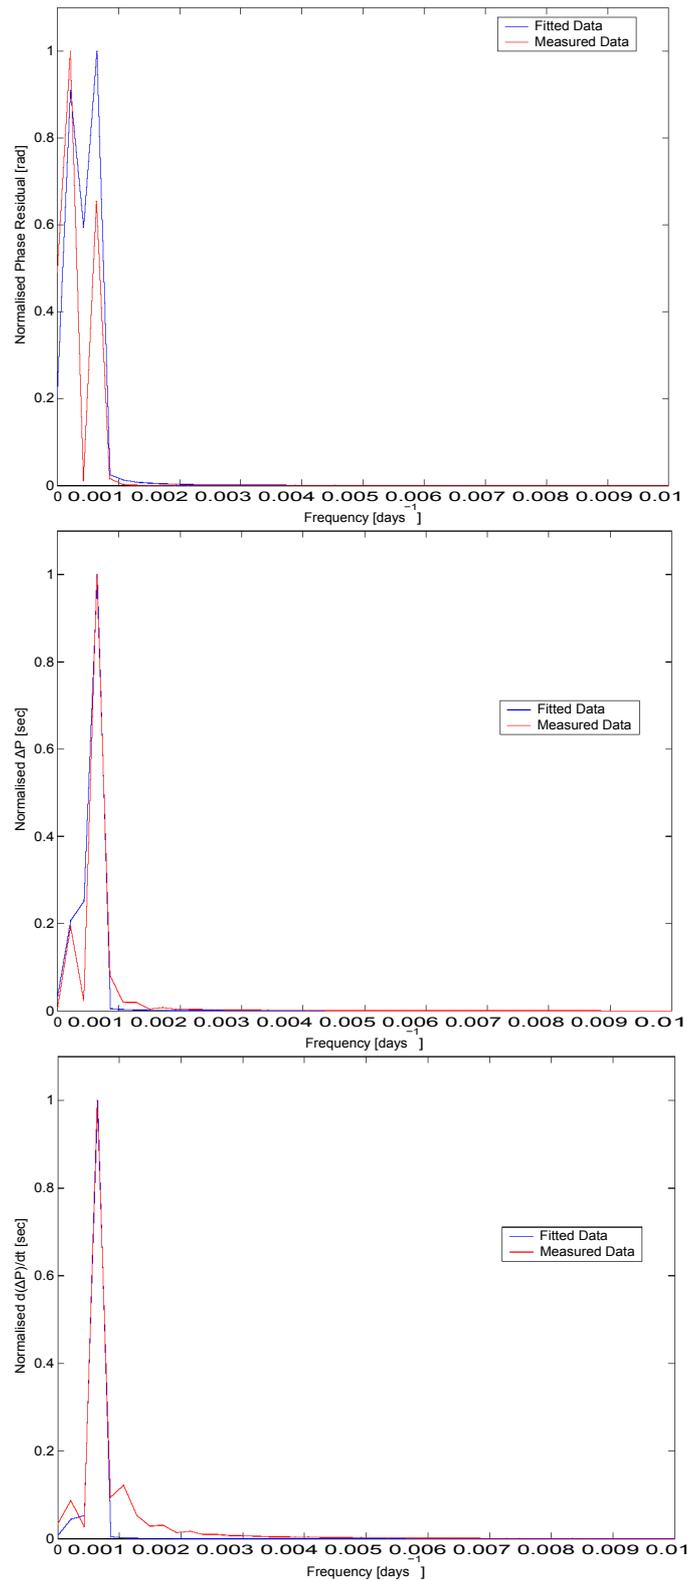


FIGURE 5.6.11. Manual fit of observed timing residuals of PSR B1557-50. From top to bottom are spectrums of Phase Residual, ΔP and $\frac{d(\Delta P)}{dt}$.

5.6.5. B0740-28. Observed timing residuals of PSR B0740-28 cannot be fitted with a torqued precession model. Both the amplitude, and shape of profile of the observed phase residuals (in particular, the concave slope between MJD 49000 and MJD 50000 followed by a convex slope between MJD 50000 and MJD 51000), are unexplained by the model.

5.6.6. B0845-41. Due to large edge effects, the observed timing residuals of PSR B0835-41 cannot be fitted at this time.

5.6.7. B0959-54. The shape and amplitude of the observed phase residual time series of PSR B0959-54 cannot be explained by the precession model.

5.6.8. B1323-62. The observed phase residuals of PSR B1323-62 is highly contaminated with timing noise. There is evidence of an incorrect ephemeris, and so the data cannot be fitted with a precession model.

5.6.9. B1358-63. While observed timing residuals during the epochs MJD 47000-50000 and MJD 50000-52000 of PSR B1358-63 can each be explained by the precession model (in both amplitude and profile shape), the data span cannot be fitted in its entirety. The apparent edge effect could either be real timing residual behaviour, or a result of an inaccurate ephemeris due to a shortened dataspan.

5.6.10. B1749-28. The profile shape of the observed phase residuals of PSR B1749-28 cannot be explained by the precession model.

5.7. Conclusion

Of the candidate pulsars, precession models give plausible fits to the timing residuals of PSRs B1642-03, B1323-58, B0736-40 and B1557-50. The observed timing residuals of PSR B1642-03 were fitted by manual adjustment of the fitting parameters. A least-squares optimisation of the parameters $[B - Pfp]$ and $[\chi - \theta(0)]$ resulted in ‘flat’ Chi-squared surfaces. It is therefore arguable whether using an automatic procedure will result in a better fit of the observed timing residuals. A comparison of the spectral fits do show better agreement with the measured data. The harmonic structure in the fitted precession model is buried in the measured data where wide, sloping spectral lines are dominant. This is due to noisy and insufficient dataspans. The manual adjustment of fitting parameters produced a better replication of the phenomenology of precession in the observed timing residuals.

TABLE 5.6.10. PSR B1557-50 derived pulsar parameters

| Parameter | Value |
|---------------------|---|
| ϵ | $4.998153017433773 \times 10^{-10}$ |
| $[I \dot{\Omega}]$ | $[0.99999999000369, 0.99999999500185, 1] \times 10^{38} \text{ [kg.m}^2\text{ s}^{-1}]$ |
| $[\omega(0)]$ | $[0.23067444456833, 0.23067444456833, 32.62120535867313] \text{ [rad.s}^{-1}]$ |
| $\langle P \rangle$ | $6.404 \times 10^{-15} \text{ [s.s}^{-1}]$ |

Although the manual adjustment of parameters for the fitting of observed phase residuals of PSR B1323-58 displayed a ‘double-humped’ precession profile, the best fit model produced a single frequency modulation. Spectral fits of the best fit model compare more favourably to the measured data, but it cannot explain the increasing and decreasing amplitude of observed phase residuals. This change in amplitudes could be a result of variations in $\dot{\mathbf{P}}$.

The fitting of observed timing residuals of PSR B0736-40 was performed using a best fit procedure. The precession profile in the observed phase residuals cannot be explained by a torqued, or untorqued precession model and is most likely the result of post-observation data processing.

Manual fitting of observed timing residuals of PSR B1557-50 reproduced accurate frequencies in the Fourier spectrums for phase, $\Delta\mathbf{P}$ and $\frac{d(\Delta\mathbf{P})}{dt}$.

PSRs B1642-03, B1323-58 and B1557-50 appear to be good candidates for the fitting of a torqued precession model. The continued monitoring of these pulsars will produce more accurate ephemeris and increased dataspan, allowing for the observation of long term precession behaviour. The use of an HPC (High Performance Computer) cluster allows for multi-variate optimisation of the model, which is presently limited by the degree of optimisation possible.

CHAPTER 6

Pulsar Timing

6.1. Introduction

Due to intermittent technical problems with the current pulsar instrument, and future research requirements highlighted in previous chapters, the continuation of the pulsar monitoring program at HartRAO is becoming increasingly dependent on the construction of a new digital pulsar processor. This chapter argues for the design of a new instrument, based on the architecture of existing processor designs and operation parameters chosen to optimise the detection of pulsars contained in the Australia Telescope National Facility (ATNF) Pulsar Catalogue [81].

6.2. Dispersion, Scattering and De-Dispersion

6.2.1. Interstellar Medium and Receiver Effects. Much of the underlying mathematical detail of this section can be found in Hankins & Rickett (1975) [76] and Backer et al (1990) [15].

The tenuous electron plasma between an observer and a distant radio pulsar source has two dominant effects on the received pulsar signal. Firstly, inhomogeneities in the medium cause interstellar scattering. The resulting multiple signal paths spread the pulsar signal out in time, causing temporal modulations in the power spectrum of the signal as well as a shift in the time of arrival. Secondly, the free electrons subject the signal to non-linear dispersion.

Examining the effect of dispersion in the Fourier domain, the pulsar signal is represented as narrow band signal upon which the interstellar medium acts as a dispersive filter with transfer function

$$H(\nu) = \exp^{-i\kappa(\nu)z}$$

where z is distance and $\kappa(\nu)$ is obtained from the dispersion relation of a tenuous plasma.

In the time domain, the spreading of the signal is modeled by convolution of the pulsar signal with the impulse response function $h(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} H(\nu) e^{i2\pi\nu t} d\nu = F^{-1}\{H(\nu)\}$. This results in a differential group delay between signals at two frequencies ν_1 and ν_2

$$(6.2.1) \quad \Delta t = \frac{DM}{2.41 \times 10^{-4}} \left(\nu_1^{-2} - \nu_2^{-2} \right) \text{ [s]}$$

where the frequencies $\nu_{1,2}$ are in MHz and $DM = \int n_e dl$ is the dispersion measure [76].

Time smearing of the detected pulse signal across a fixed receiver bandwidth B (in MHz) is

$$(6.2.2) \quad \Delta t_{DM} = 8.3 \times 10^{-3} \times DM \nu^{-3} B \text{ [s]}$$

where DM is the dispersion measure and ν the observing centre frequency in MHz [111].

The receiver sensitivity is a function of pulse smearing, which is dependent on the parameter space in which incoherent dedispersion based timers can operate [49, 56, 166]. The effect of smearing is minimised by the use of dedispersion techniques [77, 165].

6.2.2. Receiver Sensitivity. The minimum detectable pulsar flux density depends on the effects of the interstellar medium and the parameters of the receiver instrumentation. For $P > W$, theoretical analysis predicts a sensitivity limit for the receiver

$$(6.2.3) \quad S_{\min} = \frac{T_{\text{sys}}}{G} \frac{\alpha Q}{N_p B T} \frac{W}{P - W}$$

where T_{sys} is the receiver system temperature, G the telescope gain, α the desired signal to noise ratio, Q a constant dependent on radiometer type, N_p the number of polarisations, B the total receiver bandwidth, T the integration period, W the effective pulse width and P the pulse period [56, 166]. Equation 6.2.3 becomes a crude approximation for $P \approx W$. This is especially true in millisecond pulsar searches where, even though pulse broadening can cause $W > P$, there is still sensitivity to the pulsar.

The effective pulse width is

$$(6.2.4) \quad W^2 = W_0^2 + \Delta t_{\text{samp}}^2 + \Delta t_{\text{DM}}^2 + \Delta t_{\text{scat}}^2$$

where $W_0 = P \times \frac{\Delta \theta_{\text{pB}}}{2}$ is the intrinsic pulse width, P the pulsar period and $\Delta \theta_{\text{pB}} = 13^\circ \times P^{-1/3}$ the apparent pulsar beamwidth [112]. The expression for apparent beamwidth is not valid for millisecond pulsars, but should still allow for the calculation of sensitivity to the majority of the known catalogued pulsar population. Δt_{samp} is the sampling interval, Δt_{DM} is the smearing across a single channel due to dispersion and $\Delta t_{\text{scat}} = \left(\frac{\text{DM}}{1000} \right)^{3.5} \left(\frac{400}{\text{MHz}} \right)^{-1/4}$ is an approximation of the pulsar broadening due to interstellar scattering [161].

6.2.3. Dedispersion. The removal of dispersion effects to achieve high signal to noise ratio is performed using one of two techniques: coherent or incoherent dedispersion.

The coherent technique removes the dispersion effect of the interstellar medium by passing the pre-detection pulsar signal through a filter with a transfer function $H_f(\nu)$ of the form $\exp(iK(\nu)^z)$. $H_f(\nu)$ is the inverse of the transfer function that characterises the interstellar medium, and allows for microsecond resolution at low frequencies and large DM [75, 165]. The sampled bandwidth is limited by the substantial processing power required for the high data rates and filter complexity.

Incoherent dedispersion requires the division of the total receiver bandwidth into subchannels, narrow enough such that $\Delta t_{\text{DM}} \ll W_0$, where each subchannel adds an appropriate delay into the signal path to correct for dispersive delays (see equation 6.2.1) [15]. Implemented as a filterbank receiver, the channels are stacked using offline processing or in real time to form a composite profile.

6.3. Critique of the Current Pulsar Processor

The pulsar timing program at HartRAO commenced in 1984 using a 10 MHz fixed receiver bandwidth centered on one of two frequencies: 1668 MHz and 2273 MHz. The relatively small $\frac{A_e}{T_{\text{sys}}}$ of the HartRAO 26m dish, where A_e is the effective aperture, as well as the narrow fixed bandwidth of the existing pulsar processor means that the minimum detectable pulsar flux density of 3.2mJy is high by modern standards. At the completion of the current pulsar processor design in early 1984, the known pulsar population was approximately 370 pulsars [113]. Most pulsar searches in that era were designed to focus at low dispersion measure, relatively long period, high flux objects [48]. The HartRAO processor therefore met monitoring observation requirements for a large fraction of the known pulsar population at the time of project commencement. Most of the currently catalogued pulsar population in the ATNF Pulsar Catalogue [81], however, are undetectable by the HartRAO pulsar processor due to selection effects (high dispersion measure, low period, low flux density).

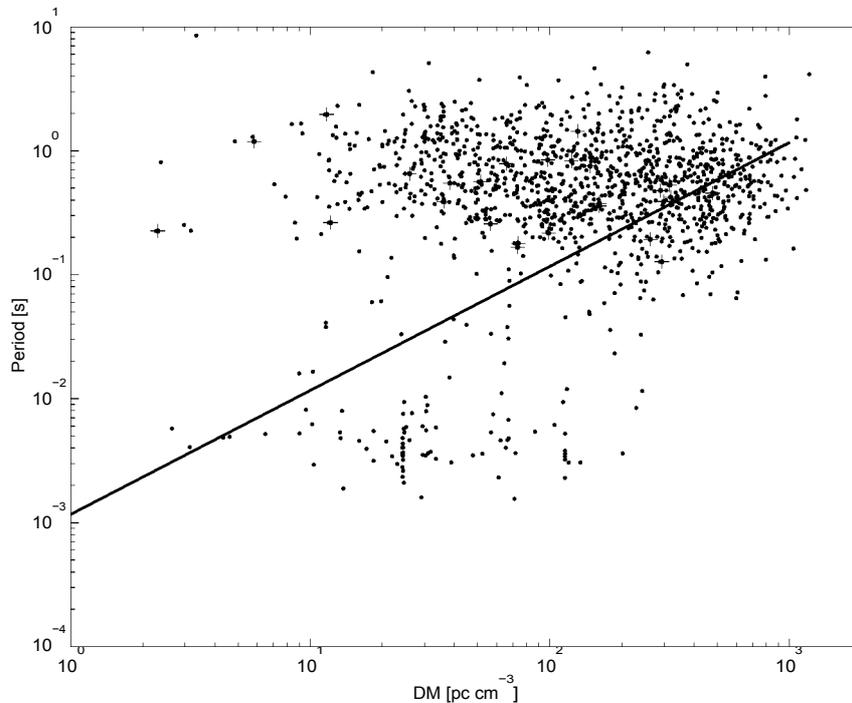


FIGURE 6.3.1. Pulsars from the ATNF Pulsar Catalogue [81] represented on the P-DM plane. The solid line represents the lower limit of pulse detectability of the HartRAO pulsar processor due to pulse smearing. The intrinsic pulse width of pulsars above this line are smeared by less than 50%. Crosses designate pulsars currently monitored at HartRAO.

Pulsars in the ATNF Pulsar Catalogue [81] are represented in the P-DM plane in Figure 6.3.1. Pulsars monitored at HartRAO are designated by crosses. The solid line represents the lower limit of detectability due to dispersion, calculated using equation 6.2.2. The intrinsic pulse widths of pulsars above this line are smeared by less than 50%. For the 10 MHz receiver channel bandwidth used by the HartRAO pulsar processor, a large proportion of high DM, short period

pulsars are not observable due to inadequate pulse resolution. This can be improved with a reduction in channel bandwidth.

The surface of minimum detectable flux density for the current pulsar processor is calculated using equation 6.2.3 and shown in Figure 6.3.2. For HartRAO, $T_{\text{sys}} = 45 \text{ K}$, $G = 0.1$, $\alpha = 1.5$, $\beta = 1.3$, $N_p = 1$, $B = 10 \text{ MHz}$. The number of samples per period is set to 1000 and $\tau = 600 \text{ s}$. While both the number of samples per period and τ are variable parameters that are dependent on the pulsar source, it is necessary to fix them at a reasonable value to calculate the surface. The same values are used for the analysis of a new pulsar processor to ensure consistency.

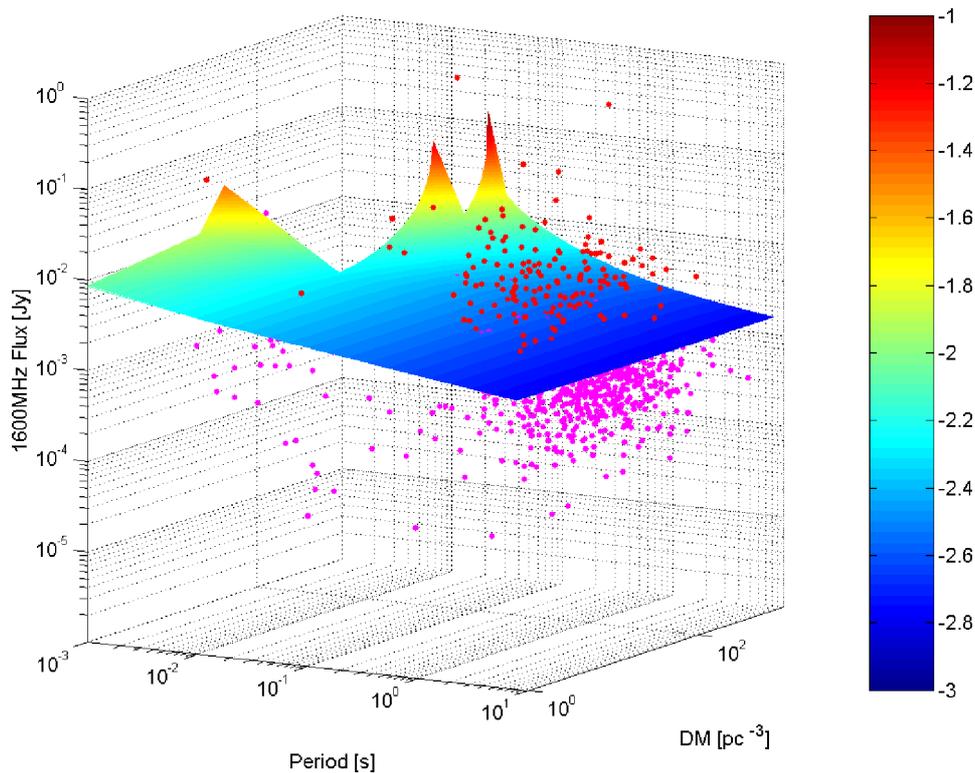


FIGURE 6.3.2. Surface of minimum detectable flux density for current HartRAO pulsar processor in a $P - DM - \text{Flux}_{1600\text{MHz}}$ parameter space. The dots represent pulsars from the ATNF Pulsar Catalogue [81]. Detectable pulsars (red) lie above the surface, while undetectable (violet) lie beneath.

The surface shown in Figure 6.3.2 rises steeply towards the high DM, short period parameter space. Pulse smearing leads to an apparent gap in the surface, referred to as the ‘exclusion zone’ in this thesis, where the effective pulse width W calculated for the parameter space is greater than the pulsar period. Pulsars that lie within this exclusion zone are undetectable. For low dispersion measure and long periods, the current pulsar processor has a minimum detectable flux density of approximately 2×10^{-3} Jansky. Pulsars in the ATNF Pulsar Catalogue [81] are represented as dots, with detectable pulsars (red) located above the surface, and undetectable pulsars (violet) beneath. The published 400 MHz and 1400 MHz fluxes were used to calculate spectral indices in order to extrapolate 1600 MHz flux values. In cases where only one of

the fluxes were available, an average spectral index of -1.8 was used. This value is valid for frequencies greater than 1400 MHz for approximately 90% of the known pulsar population [121]. Those that had no published fluxes were excluded. The total number of pulsars represented in the diagram is 1038.

Figure 6.3.3 represents the detectable pulsar subpopulation, which lies above the detectability surface in Figure 6.3.2, in a P-DM plane. This subpopulation comprises approximately 15% of the total pulsar population represented in Figure 6.3.2.

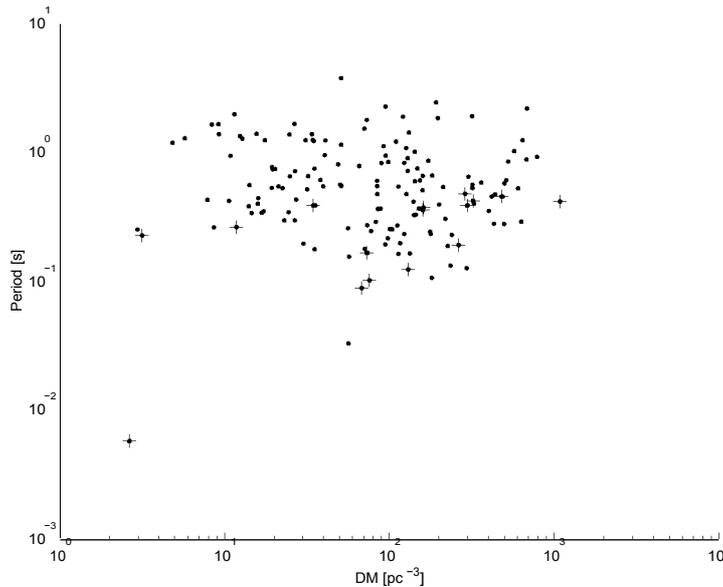


FIGURE 6.3.3. Detectable pulsar subpopulation, which lies above the surface of minimum detectable flux density in Figure 6.3.2, represented in a P-DM plane. Crosses designate pulsars currently monitored at HartRAO.

In order to study the spin down behaviour of pulsars, long continuous data spans are required to avoid ambiguity due to glitches, or microglitches. In recent history, equipment failure and subsequent servicing of legacy architecture and components has meant that regular pulsar observations at HartRAO have been interrupted for long periods of time, leading to large gaps in timing data. Servicing of the existing processor is not expected to be possible much longer, with the ever increasing difficulty of obtaining obsolete parts. To avoid further gaps in timing data, the design and implementation of a new pulsar processor whilst the current version is still in operation would ensure a seamless rollover between old and new instrument. It would also provide a window period for cross-calibration of data obtained between the two instruments operating concurrently.

With new generation radio telescopes, multi-beam systems and greater computer processing power, modern day surveys are discovering pulsars at higher dispersion measure, weaker flux and shorter periods than previously. With such extreme parameters, most newly discovered pulsars are not within the detectable region of the current pulsar processor (see Figure 6.3.2). Although continued observation of the pulsar subpopulation regularly monitored at HartRAO is a main

priority, in order to help strengthen, or disprove, current hypotheses on pulsar precession, timing noise, the braking law and glitch monitoring, the HartRAO pulsar sample needs to be extended. This extension results from an improvement in the surface of minimum detectable flux density, and subsequent increase of the detectable subpopulation of pulsars. Short of constructing a new large telescope, this is only possible with the use of a wider total receiver bandwidth and control over the subsequent increase in pulsar dispersion and degraded pulse resolution. As the current observing bandwidth for the pulsar processor is fixed and has no dedispersion capability, the installation of a new instrument is required.

6.4. Review of Existing Pulsar Processors

Before adopting a macroscopic architecture for the design of a new pulsar processor, it was necessary to review existing timers at other observatories to identify possible pitfalls and recommendations.

6.4.1. Medicina 1.4/1.6 GHz Pulsar Observing System. All technical information and figures for the Medicina Pulsar Observing System are obtained from the relevant website [117].

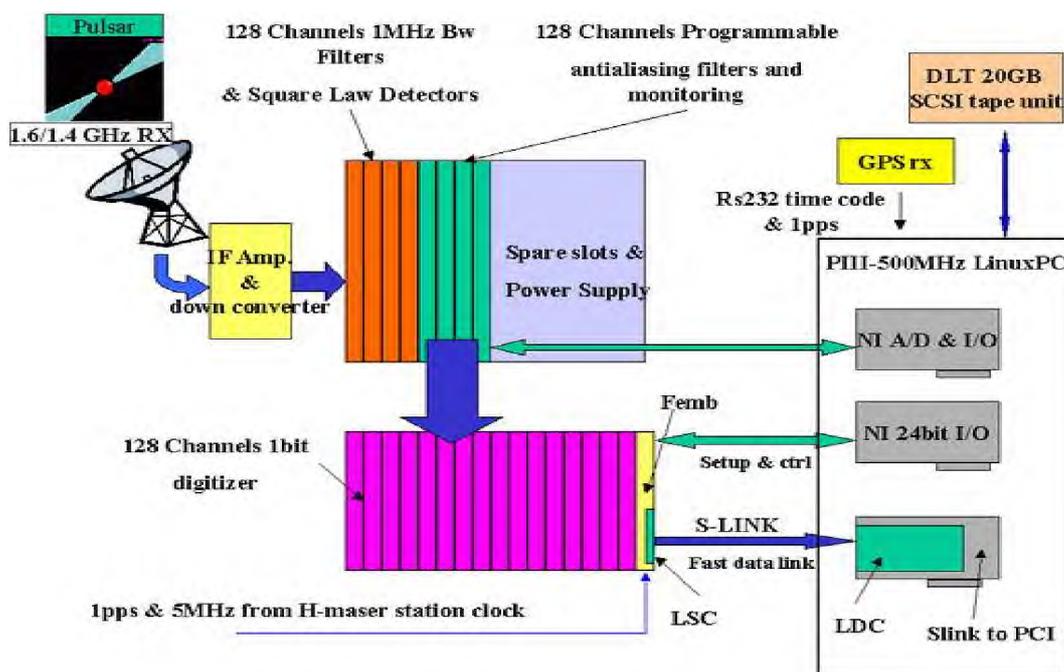


FIGURE 6.4.1. Medicina Pulsar Observing System.

The incoherent dedispersion pulsar observing system at Medicina, illustrated in Figure 6.4.1, is a 128 channel analogue filterbank receiver, covering $2 \times 2 \times 32 \times 1$ MHz of bandwidth (using left and right circular polarisation of two 32 MHz IF bands) that is centered at an observing frequency of either 1.4 GHz or 1.6 GHz. Each 1 MHz channel is square law detected before passing through a programmable anti-aliasing filter. Following this, the signal is quantised by a

1 bit digitiser at a maximum speed of 100 KHz. Data is transferred to a computer at a rate of 12.8 MB/s and stored on 20 GB Direct Linear Tapes. Alignment and stacking of pulse profiles is performed offline.

The use of an analogue filterbank receiver should be avoided at HartRAO. Achieving identical signal attenuation in each filter channel is unnecessarily difficult, but a necessity to maintain a constant gain across the filterbank frequency response. Filter boards would need to be manufactured commercially to obtain the required accuracy across the receiver response, an expense that cannot be budgeted for. The use of narrowband channels has an advantage in that RFI could be restricted to a single channel as opposed to saturating a wideband signal. Each signal path, however, needs to be digitised as soon as possible to minimise any RFI. The high replacement cost of a specialised fast data link between the digitiser and computer does not support the cost effective and efficient design philosophy for the new pulsar processor at HartRAO. Data storage within each narrowband channel would reduce the need for specialised data links. With the high turnover rate of modern computers, the compatibility of proprietary cards would cause difficulty in future upgrades and maintenance.

6.4.2. COBRA. All technical information and figures for COBRA are obtained from the relevant website [39] and online documentation.

COBRA is a coherent online baseband receiver for astronomy, primarily used to perform coherent dedispersion of pulsar data obtained at Jodrell Bank. Illustrated in Figure 6.4.2, the system consists of a hybrid analogue/digital frontend concerned with the signal processing of data, while the backend is a Beowulf cluster of approximately 150 computers, which can perform real time or off line signal processing. The receiver records both left and right circular polarisations, each polarisation having a 100 MHz bandwidth. The complex signal that results from baseband sampling passes through a number of direct memory access (DMA) cards in parallel, splitting each band into smaller sub-bands of 1.25, 2.5, 5 or 10 MHz in order to avoid saturation of the entire bandwidth due to possible sources of RFI. The signal is digitised with 8 bit precision, producing a datastream with an output data rate of over 400 MB/s, which flows into the Beowulf cluster. Each computer has at least 512 MB RAM and 10 GB disk space. The data from each channel is coherently dedispersed in real time before individual profiles are combined. COBRA can also perform incoherent dedispersion over bandwidths of up to 200 MHz wide.

The advantage of COBRA is that it is a fully digital system after baseband sampling, reducing the impact of RFI. Digitisation of the signal path is a priority at HartRAO given the increasingly widespread use of portable communication devices such as cellular phones. Observing frequencies can be adjacent to one another or spread because of the use of programmable local oscillators for individual sampler modules, each of which ranges between 20 MHz and 83.75 MHz in 250 KHz steps. Although the Beowulf cluster is both powerful and upgradable, the capital outlay for such a system far exceeds the budget for a new pulsar processor at HartRAO. While the COBRA design has moved towards that of a software receiver, the multi-purpose system exceeds the design requirements for a pulsar processor at HartRAO. The system relies

on a large commercial component, which could be replaced by cheaper and superior alternatives such as programmable logic.

6.4.3. Princeton Mark IV Pulsar Processor. All technical information and figures for the Mark IV are obtained from the relevant website [146] and Stairs et al (2000) [165].

Operated at Arecibo Observatory, Princeton’s Mark IV system, shown in Figure 6.4.3, is a baseband recorder that covers a 5 or 10 MHz bandwidth centered at an observing frequency of 430 MHz. Left and right polarisations are mixed to baseband, generating a 5 (10) MHz complex signal, which is 4 (2) bit digitised. The resulting 10 MB/s data stream is acquired by a SPARC-20 workstation and stored onto disk or Direct Linear Tapes (DLT) for offline processing. Coherent dedispersion takes place on a custom made 1.25 Gflop parallel processor, which is optimised for computing fast Fourier transforms.

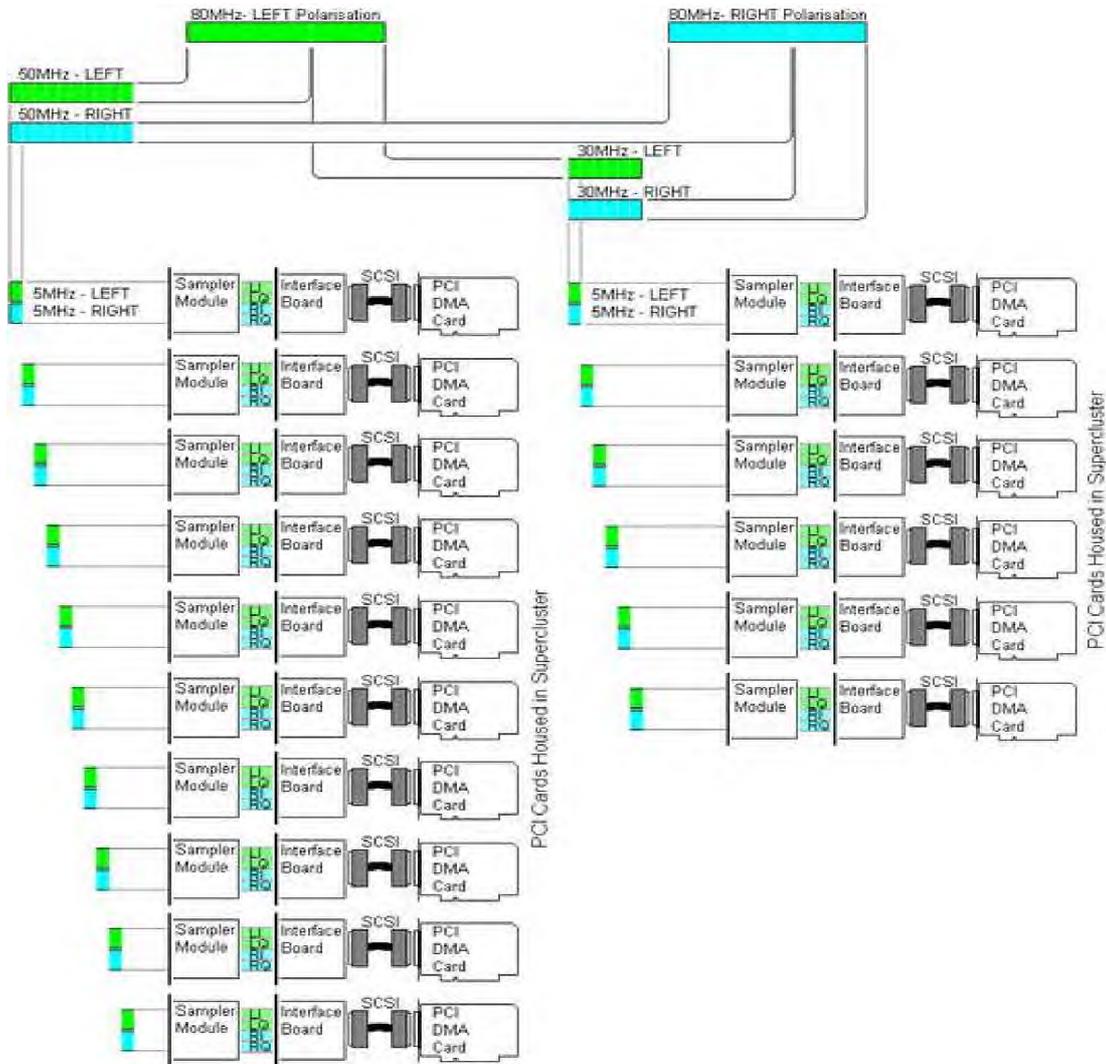


FIGURE 6.4.2. Coherent Online Baseband Receiver for Astronomy (COBRA)

Together with the large collecting area, the observing frequency of 430 MHz at the Arecibo Observatory means that the relatively narrow observing bandwidth of the pulsar processor is sufficient to observe weak pulsars. However, with the high observing frequencies of 1.67 GHz and 2.23 GHz at HartRAO, a much wider observing bandwidth is required to obtain similar signal to noise ratios. Coherent dedispersion would result in the need for greater computer processing power to handle the high data rates. The data generated by the baseband sampler requires sophisticated data links to the processor. Customised circuit boards, parallel processors and tape drives for storage are all expensive components in the context of HartRAO. The need for a large digital component is highlighted by the proven superiority of the Mark IV to the Mark III, an analogue filterbank receiver implementing incoherent dedispersion.

6.4.4. Coherent Arecibo-Berkeley Pulsar Processor (ABPP) . All technical information and figures for the ABPP is obtained from the relevant website [17] and available online documentation.

Illustrated in Figure 6.4.4, the ABPP is a hybrid filterbank receiver, using four analogue baseband mixers followed by 8 way digital divisions in order to spread up to 32 channels across a 112 MHz bandwidth in total power coherent mode. Each analogue mixer produces a quadrature baseband signal, which in turn is distributed to eight digital filter boards (DFBs) where the signal is quantised to 4 bits (See Backer et al (1997) [16] for more detail on the digital filter boards). Each DFB, shown in Figure 6.4.5, consists of a digital mixer (implemented by look up table), Harris digital FIR filter chips, and a 256kB EPROM for data encoding. The signal is passed to a dedispersion board, shown in Figure 6.4.6, which is capable of performing both coherent and incoherent dedispersion. Data acquisition is performed using a dedicated computer.

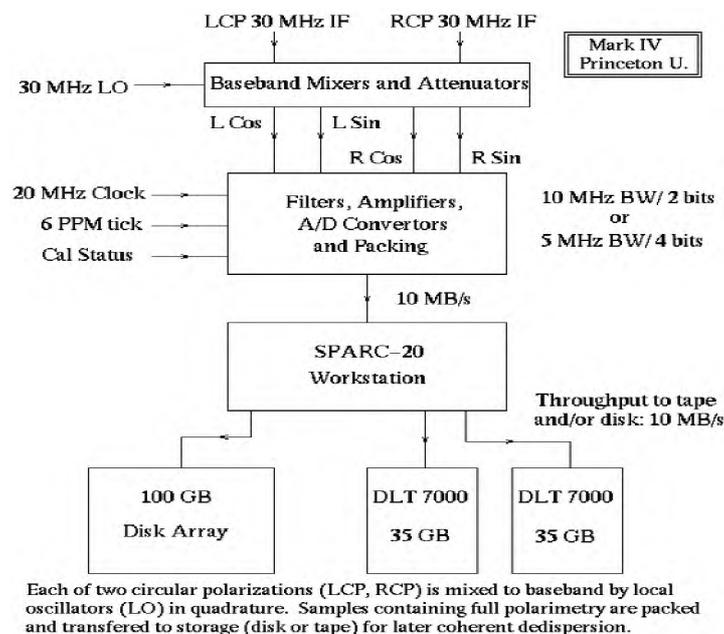


FIGURE 6.4.3. Princeton Mark IV Pulsar Processor

The ABPP is a highly generalised instrument, essentially relying on the replication of two modules: an analogue mixer frontend and a digital backend. During the design stage the boundary interface, or stage of digitisation, between the analogue and digital components can be varied. This movement will have a substantial effect on the cost of implementation, with the cost of analogue components generally far cheaper than that of digital components in both implementation, processing overhead and currency. Given a suitable stage of digitisation, the performance/cost ratio can be maximised. With the use of reprogrammable logic in the digital backend, the system could form the basis of a low cost, multi-purpose radio astronomy instrument at HartRAO. However, the use of specialised FIR filter chips represents a large budgetary component. As each DFB is replicated a number of times, the expense of these chips is too great in the HartRAO context and viable alternatives should be considered.

6.4.5. Navy-Berkeley Pulsar Processor (NBPP). All technical information for the NBPP is obtained from the relevant website [135].

Based on the original design for the ABPP, the NBPP is a 96 channel 175 MHz filterbank receiver with analogue frontend and digital backend. Both the design implementation and

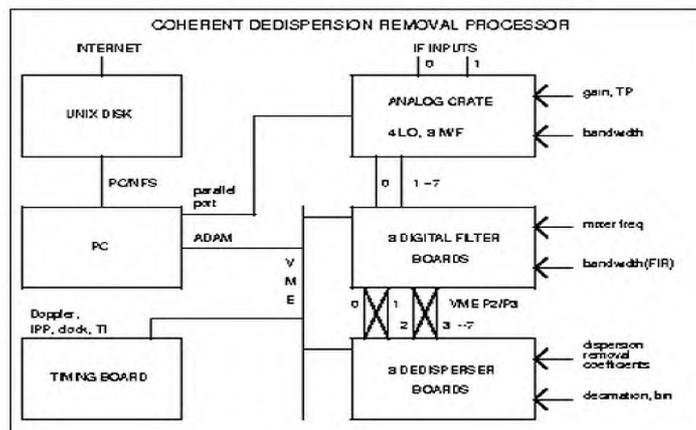


FIGURE 6.4.4. Coherent Arecibo-Berkeley Pulsar Processor (ABPP)

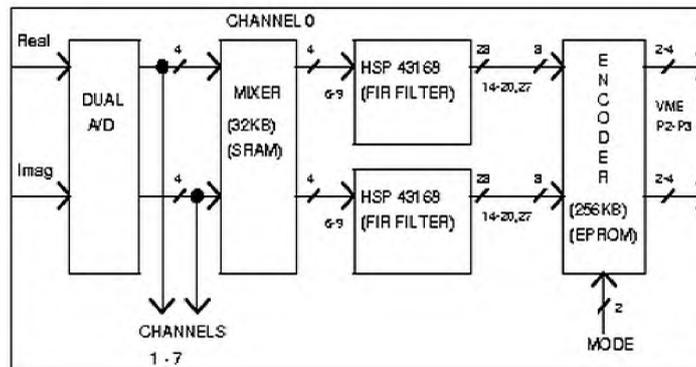


FIGURE 6.4.5. Coherent Arecibo-Berkeley Pulsar Processor (ABPP): Digital Filter Board

components used are the same as those in the ABPP. Improvements over the ABPP include the use of Hilbert transforms in the FIR low pass filters [16] in order to retrieve upper and lower sideband information in the EPROM encoder. Further modification of the FIR filter coefficients have allowed for 192 tap filter implementations.

The use of the Hilbert transform allows for a halving of digital sample rate, with a digital reconstruction of an $N \times 1$ MHz wideband signal possible using an effective half Nyquist data rate.

6.4.6. PuMa. All technical information for the PuMa is obtained from Voute et al (2002) [173].

Installed at Westerbork Synthesis Radio Telescope, the PuMa, or digital Pulsar Machine, was originally designed as a digital filterbank receiver, replacing analogue filters with digital signal processors (DSPs) performing fast Fourier transforms. With a variable number of channels in the filterbank, as many as 32768 channels can be spread over an 80 MHz receiver bandwidth using 192 SHARC DSPs. During the latter stages of development, baseband sampling of two 10 MHz bands was included. Split into two halves, the system can observe in both observational modes, as a filterbank receiver and baseband sampler, simultaneously. The system avoids coarse digitisation by use of a 12 bit analogue to digital converter. This circumvents the problems associated with adverse signal to noise ratios as a result of RFI, highlighted by Voute et al (2002) [173]. Data storage is on hard disks, with both coherent and incoherent dedispersion being performed offline.

PuMa is a versatile instrument. However, the use of 192 DSPs would be a costly design component in the HartRAO budgetary context. The use of a 12 bit digitiser highlights the

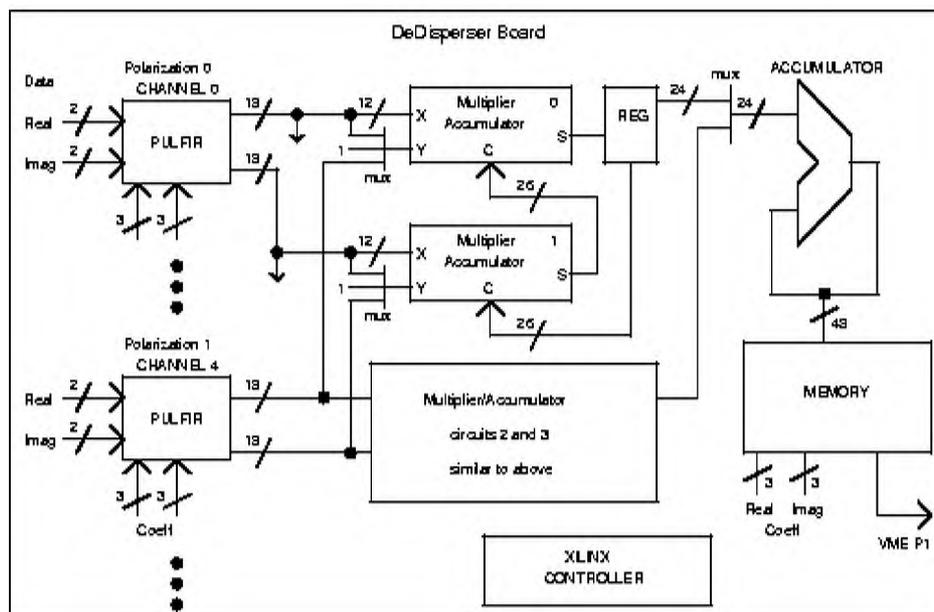


FIGURE 6.4.6. Coherent Arecibo-Berkeley Pulsar Processor (ABPP): Dedisperser Board

problems associated with using coarse inputs signals, and allows good management of saturating input signals. High precision digital signals can be implemented at HartRAO with the low cost, high precision fast commercial analogue to digital converters that are presently available.

6.4.7. Caltech Parkes Swinburne Recorder II (CPSR2). All technical information for the CPSR2 is obtained from the relevant website [18].

Installed at the Parkes radio telescope, the CPSR2 is a baseband recorder, which is used for precision timing studies of pulsars. The recorder accepts four baseband converted 64 MHz RF signals, each of which is 8 bit sampled by a customised high speed digitiser. Due to the four independent IF channel inputs, the digitiser allows for the ability to observe with full polarimetry at two different sky frequencies, limited only by the receiver passband and down conversion mixer chain settings. The resulting packetised data stream (up to 128 MBPS) is fed into two primary processor nodes by means of DMA cards, and then distributed by high speed ethernet into a further cluster of 28 dual processor computers. This cluster has enough processing power to coherently dedisperse data in real time from most low dispersion measure pulsars.

CPSR2 is the most modern functional baseband processor currently in operation. The ability to process four independent IF channels is an important attribute for multi-frequency observations. Although this does not form part of the user requirements for a new pulsar processor at HartRAO, it could form part of a future enhanced specification. CPSR2 highlights the need for a small supercomputer, or computer cluster, to perform coherent dedispersion of pulsars. Budgetary constraints at HartRAO do not allow for this, and may limit us to incoherent dedispersion techniques.

6.5. Macroscopic Architecture Justification

6.5.1. Introduction. Based on the review of existing timers in section 6.4, this section details macroscopic design choices for a new pulsar processor.

6.5.2. Analogue vs. Digital. Modern instrumentation is moving towards highly computerised and digital implementations. Due to the instrumental sensitivity required by the field of radio astronomy, and the necessity to mitigate against radio frequency interference, the need to digitise the signal path as soon as possible is required.

A digital implementation requires the choice between two different digital technologies: reprogrammable logic, in the guise of field programmable gate arrays (FPGAs), or preprogrammable hardware, in the guise of digital signal processors (DSPs).

Although now covering a broad spectrum of engineering applications, the computational requirements of most digital signal processing designs remain the same: a need for real time computational performance coupled with the ability to adapt to constantly changing data sets and conditions. Cost and power have become important parameters to consider as signal processing applications move from a predominantly military orientated market to a commercial one. This has led to a desirability for the ability to implement hardware solutions with the

flexibility of reprogrammability. This inherent ability in FPGAs has allowed them to become viable alternatives to programmable DSPs [72, 73, 144].

Since the mid 1980's, design structures for FPGAs have been developed to perform a variety of computational tasks, and in more recent times FPGAs have become a viable alternative to custom hardware (Application Specific Integrated Circuits, or ASICs), with clock rates approaching 1 GHz. Current FPGAs outperform custom approaches to real time computation as they provide the speed of dedicated circuitry while still containing the flexibility of a programmable system, which allows for optimisation of design structures via incremental adjustments.

The use of pipelined structures, coupled with the parallel nature of many signal processing applications, such as filterbanks, has in certain cases led to FPGA based computing outperforming DSPs by up to an order of magnitude [72]. High gate densities of today's FPGAs has opened a new era of single chip solutions. Advantages include: less off-chip connections due to fewer on-board components and the resolution of signal processing tasks in a single FPGA, reducing interconnect timing problems.

The precision parallelism found in FPGA architectures is suited to high sample rates in distributed computation such as that required in speech and image processing [169] and radar processing [13]. Multiple MAC (multiply and accumulate) units are implemented on a single FPGA, allowing up to 600 billion MAC operations per second, which is two orders of magnitude greater than industry standard DSP chips [137].

The performance of FPGAs seem to consistently be superior in terms of flexibility, processing performance and price/performance ratio [72, 144] and are a natural choice for use in modern day signal processing applications.

6.5.3. Coherent vs. Incoherent Dedispersion. The implementation of a coherent dedispersion system requires the design of a single unit solution. Removal of dispersion effects on target pulsar signals needs unique sets of filter coefficients and reprogrammable filters. Wide total receiver bandwidths, necessary at HartRAO to obtain the desired receiver sensitivity, requires the use of high cost, specialised DSPs. The pulsar processor bandwidth would therefore be restricted by the feasibility of a single unit solution. Subsequent increases in the bandwidth will likely require a complete redesign of the system. The single unit solution is compact, with little replication of parts. This means that servicing of the instrument is relatively simple. The biggest advantage of the coherent dedispersion system is the complete removal of dispersion effects, obtaining microsecond resolution in the integrated pulse profile. However, at the high observing frequencies at HartRAO, the benefits of a coherent dedispersion system would not be fully realised.

Implementation of a filterbank architecture is required for the use of an incoherent dedispersion pulsar observing system. The scalability of the processor means that the wide receiver bandwidths required are obtained by the replication of generalised modules. Further units can be added to the filterbank in order to increase the bandwidth as the need requires. The narrow-band filterbank channels reduces the required processing power for the digital filters, and allows for a reprogrammable solution.

To implement the pulsar processor system in a cost effective way, and to reduce the required processing power, incoherent dedispersion is the most applicable technique for use in the HartRAO observing environment.

6.6. Parameters for a New Pulsar Processor

Two critical parameters to be taken into consideration for the design of a new pulsar processor are cost effectiveness and the improvement in the surface of detectable flux density (see Figure 6.3.2). Cost effectiveness is achieved by appropriate use of the array of COTS (consumer-off-the-shelf) components that are currently available. The low cost and availability of these components is the result of a consumer market, which is driven by the development of cellular phone and satellite technology. Specialised components not found in mainstream industries are to be avoided to keep within the budget requirements of HartRAO.

Improvement of the surface in Figure 6.3.2 implies an increase in the overall system sensitivity and shrinkage of the exclusion zone. Ideally, the surface of detectable flux density should be lowered by an order of magnitude to substantially increase the detectable pulsar subpopulation. Channel bandwidth and total bandwidth are the two most important technical parameters to be decided on at the early stage of the design process as they directly affect the number of observable pulsars (see equation 6.2.3).

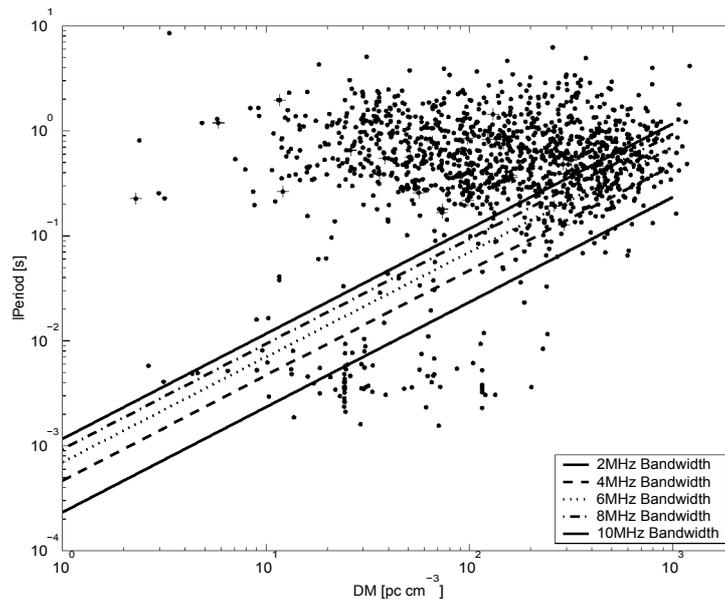


FIGURE 6.6.1. Pulsars from the ATNF Pulsar Catalogue [81] represented on a P-DM plane. The solid lines represent the lower limit of pulse detectability for fixed channel bandwidths due to pulse smearing. The intrinsic pulse width of pulsars above these lines are smeared by less than 50%. Crosses designate pulsars currently monitored at HartRAO.

Reducing the current channel bandwidth of 10 MHz to 2 MHz significantly increases the number of detectable pulsars, as seen in Figure 6.6.1, due to the consequent reduction in pulse smearing.

Figure 6.6.2 illustrates the dependence of the detectable pulsar subpopulation of the ATNF Pulsar Catalogue [81], calculated using equation 6.2.3, on channel bandwidth and total bandwidth. The surface shows the increase in the number of detectable pulsars following an increase and decrease in the total bandwidth and channel bandwidth respectively. The rapid increase of detectable pulsars due to an increase in total bandwidth is a far more dominant effect than the reduction of channel bandwidth. This is due mainly to two reasons. Firstly, the relatively high observing frequencies at HartRAO reduces the effect of dispersion. Therefore, the undetectable pulsar subpopulation that lies within the exclusion zone of Figure 6.6.4 is a small percentage of the total pulsar population represented. Secondly, the current galactic pulsar distribution in the P-DM plane is not necessarily an accurate reflection of the total galactic population. The number of pulsars discovered in the short period, high dispersion measure, low flux parameter space (exclusion zone of the current pulsar processor) should increase with time as the region becomes targeted by more exhaustive pulsar searches. Therefore, the reduction of channel bandwidth should take into account a projected galactic population as opposed to the current ATNF Pulsar Catalogue [81].

Figure 6.6.3 shows the dependence of the detectable pulsar subpopulation of a simulated earth beaming population of 100 000 pulsars¹ on channel bandwidth and total bandwidth. The

¹The author would like to thank Michael Kramer of Jodrell Bank Observatory for use of the simulated pulsar population which was produced for use in Square Kilometer Array simulations.

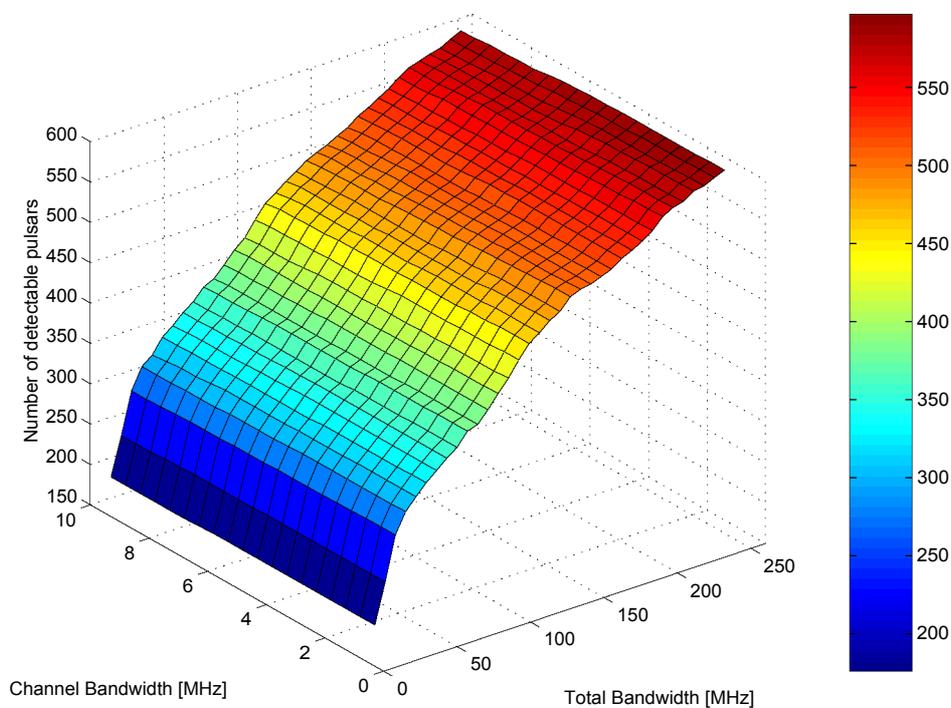


FIGURE 6.6.2. Dependence of detectable pulsar subpopulation of the ATNF Pulse Catalogue [81] on channel bandwidth and total bandwidth.

similarity between Figures 6.6.2 and 6.6.3 strongly suggest that any macroscopic design optimisations to increase the detectable pulsar subpopulation of the ATNF Pulsar Catalogue [81] would result in an increase in the detectable subpopulation of future pulsar catalogues.

Based on previous arguments on technical parameter values, Figure 6.6.4 represents the surface of minimum detectable flux density, calculated using equations 6.2.3 and 6.2.4, with 2 MHz channel bandwidth and 112 MHz total receiver bandwidth. On average, the improvement in sensitivity is of the order of $\sqrt{\frac{112}{10}} = 3.34$. The decrease in minimum flux density results in an almost threefold increase in the size of the detectable pulsar subpopulation, plotted in Figure 6.6.5.

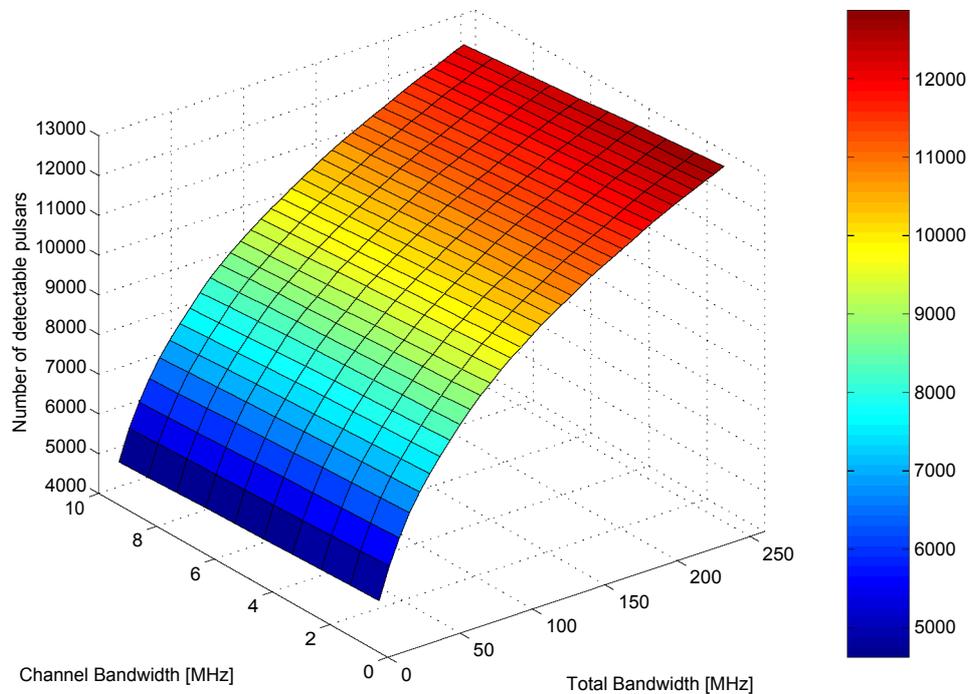


FIGURE 6.6.3. Dependence of detectable pulsar subpopulation of a 100 000 simulated earth beaming pulsar population on channel bandwidth and total bandwidth.

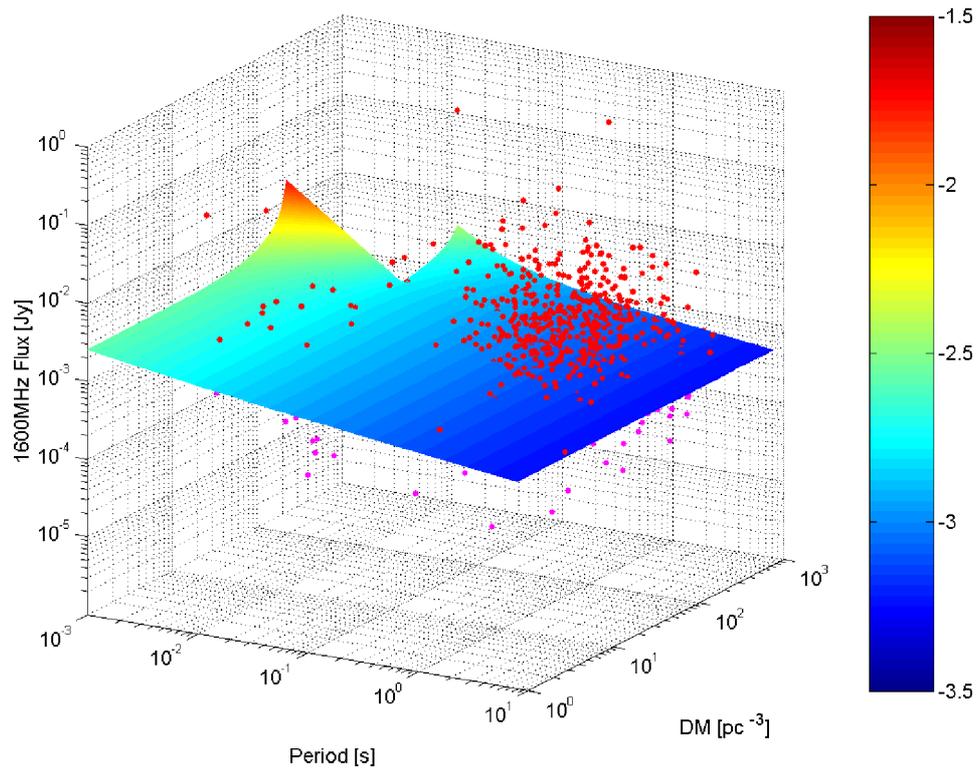


FIGURE 6.6.4. Surface of minimum detectable flux density for proposed new pulsar processor in a $P - DM - \text{Flux}_{1600\text{MHz}}$ parameter space. Channel bandwidth is 2 MHz and total bandwidth is 112 MHz. The dots represent pulsars from the ATNF Pulsar Catalogue [81]. Detectable pulsars (red) lie above the surface, while undetectable (violet) lie beneath.

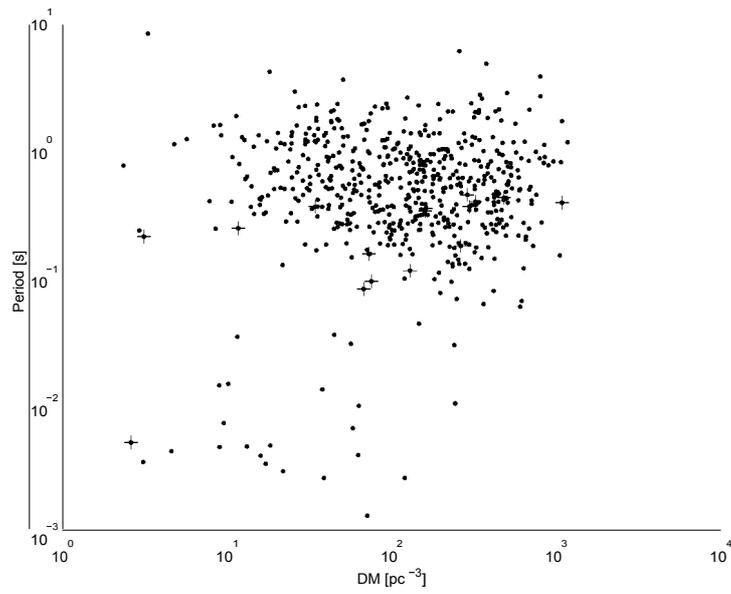


FIGURE 6.6.5. Detectable pulsar subpopulation, which lies above the surface of minimum detectable flux density in Figure 6.6.4, represented in a P-DM plane. An almost threefold improvement of the detectable pulsar subpopulation size is obtained. Crosses designate pulsars currently monitored at HartRAO.

CHAPTER 7

Pulsar Processor Design and Implementation

7.1. Introduction

The design and implementation of a new digital pulsar processor for HartRAO is discussed in this chapter. The instrument is based on the implementation of a digital filterbank receiver, which is used to perform incoherent dedispersion. The system is based largely upon the hybrid approach of Backer et al (1997) [16] and subsequent pulsar processors such as the ABPP (see section 6.4.4).

The pulsar processor was initially implemented to satisfy a low frequency specification ($f_{s,mple} = 1\text{MHz}$), before designing a high frequency ($f_{s,mple} = 32\text{MHz}$) equivalent (see appendix D for schematic designs and flowcharts for the Low Frequency Concept). Construction of the high frequency pulsar processor resulted in two hardware implementations of the system. The first version is designed for installation into a prototype rack, whereas the second version was manufactured to standard eurocard specification (incorporating the DIN 41612 and IEC 603-2 connector standards, IEEE 1101 PC board standards and the DIN 41494 and IEC 297-3 rack standards).

The design and implementation of the prototype pulsar processor is discussed in this chapter, with a short description of the eurocard implementation given in section 7.17.2. Relevant implementation decisions, and the selection of hardware, is a compromise between price, performance and availability in the local market.

7.2. High Level Architecture

7.2.1. Filterbank Architecture for Parallel Digital Signal Processing . The filterbank architecture, usually consisting of one or both of an analysis bank (for downsampling) and a synthesis bank (for upsampling), is primarily used to exploit parallel architectures in signal processing and provide decomposition of wideband signals. The subsequent sub-bands, obtained either by bandpass or low pass filtering, are approximately orthogonal (given sharp rolloffs at the filter passband cutoff frequencies). As such, they can be conveniently processed in parallel.

Following downsampling, the necessary processing of each sub-band signal is decreased by a factor N , where N is the number of sub-bands, which synchronously add up to the total wideband signal. This allows for high bandwidth, complex signal processing without the difficulties associated with high speed digital electronics.

The principle of orthogonality states, given optimum conditions for a transversal filter, each element of the input vector $x[n]$ and estimation errors $e_o[n]$ are orthogonal. Mathematically, this reduces to $E[x[n]e_o[n]] = 0$ where E is the estimation function [78].

This architecture is physically scalable, with a direct correspondence between processed sub-bands and physical processing units. The option of bypassing certain sub-bands means that any narrowband interference can be reduced by disabling the relevant processing units.

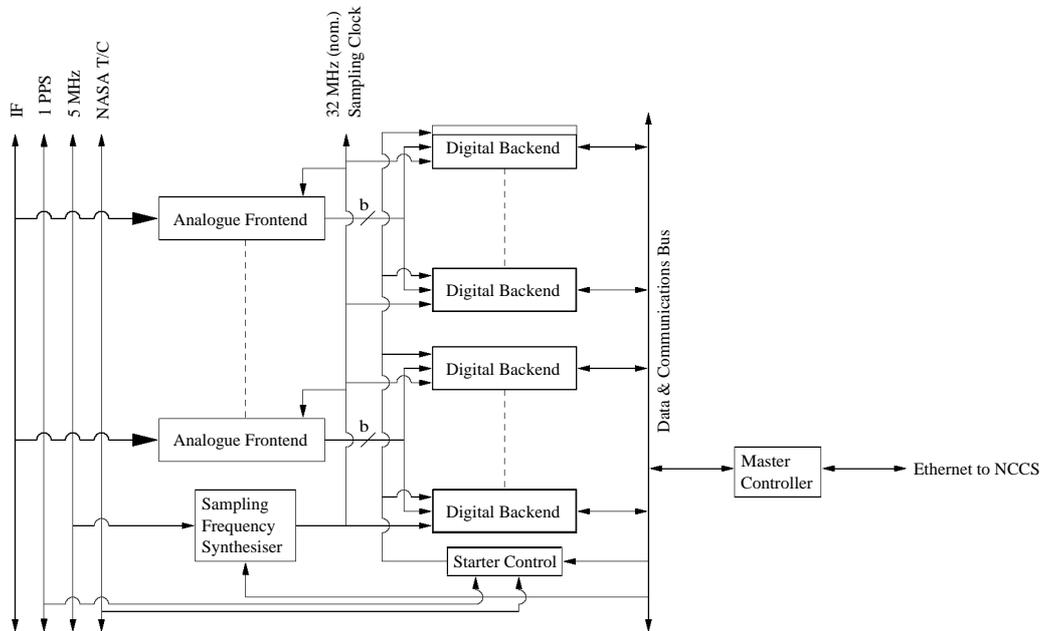


FIGURE 7.2.1. Macroscopic design structure of new pulsar processor

7.2.2. Schematic Design. The processor design, shown in Figure 7.2.1, consists of two hierarchical filterbank layers: an analogue frontend and a digital backend. Each layer is implemented by the replication of a generalised module. The analogue frontend module is a quadrature mixer that selects a 32 MHz IF slice and converts it to a complex baseband signal. This signal is digitised to b bit precision and fanned out to n $[0 : 7]$ digital backend modules, each of which processes 2×2 MHz channels. Individual channels perform on-board data stacking and are connected to a common serial data bus. The processor is configured by a master controller, which is accessible via an ethernet network and communicates with the New Computer Control System (NCCS) at HartRAO [67]. A starter module is used to initiate pulsar observations.

7.3. Analogue Frontend Architecture

7.3.1. Schematic Design. The analogue frontend, shown in Figure 7.3.1, combines a quadrature mixer with a high speed dual analogue to digital converter. The mixer has an input bandwidth that covers the 90 MHz - 230 MHz IF band at HartRAO. The analogue gain control (AGC) has 100 dB dynamic range, adjustable programmatically through the use of the communications bus and a digital to analogue converter. Low pass anti-aliasing filters on the mixer output limit the baseband signals to 16 MHz. Thus, each analogue frontend module supplies a 32 MHz double sideband signal, tunable within the IF band via suitable choice of local oscillator.

The complex output signal, consisting of an in-phase (real) component $I(t)$ and quadrature (imaginary) component $Q(t)$, is digitised by a dual analogue to digital converter, which ensures that the phase jitter between real and imaginary signals is minimised. The signals are Nyquist sampled at a nominal 32 MHz, sufficient for the double sideband baseband signal. The actual sampling frequency is chosen to be commensurate with the pulsar period (see section 7.18).

7.4. Quadrature Mixing

A description of quadrature mixing and complex signals is given on pp 297-301 of Lyons (1997) [116].

Quadrature mixing produces two 90 degree phase separated baseband signals $I(t)$ and $Q(t)$, with Fourier transforms

$$\begin{aligned} F\{I(t)\} &\#> Y_r(\omega) = X_L(\omega) + X_U(\omega) + X_L^*(-\omega) + X_U^*(-\omega) \\ F\{Q(t)\} &\#> Y_i(\omega) = j[X_L(\omega) + X_U(\omega) - X_L^*(-\omega) - X_U^*(-\omega)] \end{aligned}$$

where the subscripts U and L indicate upper and lower sidebands respectively, $X(\omega)$ the Fourier transform of the input signal $x(t)$ and $*$ the complex conjugate operator (see appendix C.1).

7.4.1. Hardware: Quadrature Mixer. The AD6121 from Analog Devices is a CDMA 3V Receiver IF subsystem. It includes a linear IF amplifier with 94.5 dB dynamic range analogue gain control. The quadrature demodulator, including a divide-by-two quadrature generator, mixes 32 MHz slices from an IF signal, in the range 50 – 350 MHz, down to baseband with a ± 1 degree accuracy between complex components and produces a differential output signal with a maximum 700 mV p-p amplitude. The product datasheet can be found at http://www.analog.com/UploadedFiles/Obsolete_Data_Sheets/674735AD6121.pdf. The AD6121 is now obsolete, but can be replaced by the AD8348, which has an IF operating frequency in the range 50 – 1000 MHz. The product datasheet can be downloaded from http://www.analog.com/UploadedFiles/Data_Sheets/53999288331825AD8348_0.pdf.

7.4.2. Hardware: Differential Amplifier. The AD8130 differential amplifier converts the differential output of the quadrature mixer into an unbalanced single-ended signal. It

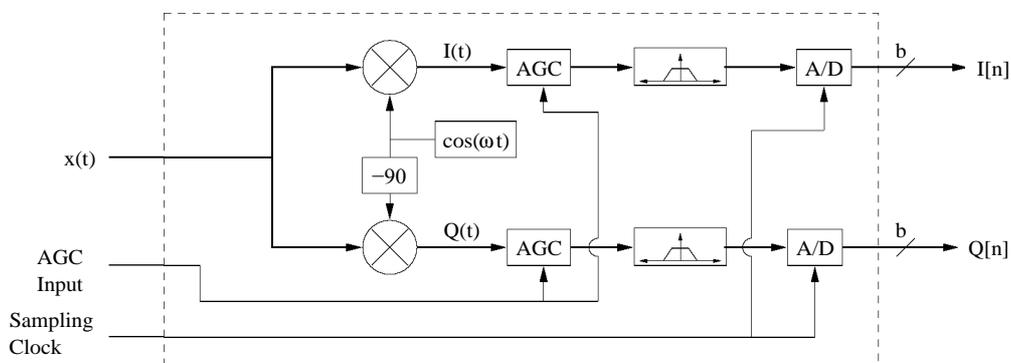


FIGURE 7.3.1. Analogue frontend signal flow diagram

has a differential input operating range of +2.5 V p-p and a 270 MHz bandwidth. Using an active feedback architecture and multiple differential inputs, the gain response of the amplifier is independent of the input signal, which allows for good common-mode rejection. The amplifier is implemented with unity gain. The relevant datasheet can be located at http://www.analog.com/UploadedFiles/Data_Sheets/328430154AD8129_30_0.pdf.

7.5. Analogue to Digital Conversion

7.5.1. Quantisation Errors and Effects. Sampling of the analogue signal rounds the floating point value to the nearest quantised level representable by a fixed binary length data-word. This discrete signal is

$$(7.5.1) \quad Q[x(n)] = x(n) + \varepsilon(n)$$

where, at times $t = nT$, T the sampling period, $x(n)$ is the original continuous signal and $\varepsilon(n)$ is the additive quantisation error signal. $\varepsilon(n)$ has a maximum amplitude of $\pm \frac{1}{2}$ LSB (least significant bit) for an unsaturated signal. Under the following conditions :

- (1) $\{\varepsilon(n)\}$ is a sampled sequence of a random stationary process
- (2) The error sequence is uncorrelated with the analogue samples
- (3) The random variables of the error signal are uncorrelated
- (4) The probability distribution of the error signal is uniform over the range of quantisation error

the quantisation error signal will have a normal distribution ([138], pp. 415).

The standard deviation σ_ε of the quantisation error signal $\varepsilon[n]$ is related to the quantisation resolution Δ by ([138], pp. 416)

$$\sigma_\varepsilon = \sqrt{\frac{\Delta^2}{12}}$$

The signal to noise ratio (SNR) of the quantised signal is

$$(7.5.2) \quad \text{SNR} = 10 \log \frac{\sigma_x^2}{\sigma_\varepsilon^2} = 6.02b + 10.79 + 10 \log(\sigma_x^2) \text{ [dB]}$$

where b is the number of bits in the quantisation process, and σ_x the standard deviation of the input signal normalised to the full scale dynamic range of the digitiser ([138], pp. 416). The SNR therefore increases approximately 6 dB for every bit added to the fixed point binary representation. This assumes, however, that the full dynamic range of the input is optimally used. For a reduced input dynamic range $Ax[n]$, $A < 1$,

$$\text{SNR} = 6.02b + 10.79 + 10 \log(\sigma_x^2) + 20 \log(A) \text{ [dB]}$$

([138], pp. 417). Reducing the input to less than the full dynamic range of the system reduces the SNR. Reduction in the SNR for a fixed bit length b is expected when using an input noise signal.

7.5.2. Effect of Quantisation on Analogue Input Signals. In calculating the power distribution of the analogue signal, quantised signal and the error signal, the effect of root-mean-squared (RMS) input voltage levels on the SNR is determined. Figure 7.5.1 plots the normalised power distribution of these signals as a function of standard deviation σ_x of the input signal. The analogue input signal level is normalised by the ratio $\frac{RMS}{FSD} = 0.2$, where FSD is the full scale dynamic range of a 4 bit analogue to digital converter. The value of $\frac{RMS}{FSD}$ was obtained using Monte-Carlo simulations, which produced a graph of SNR plotted as a function of $\frac{RMS}{FSD}$ for 4 bit quantisation, shown in Figure 7.5.2. While a reduced input of less than 0.2 results in a rapid decline in the SNR due to inefficient use of the full scale dynamic range, the long 'roll off' for a saturating signal is caused by the inability of the quantised signal to carry the full power of the input signal, with the maximum amplitude of the error signal growing beyond $\pm \frac{1}{2} \Delta$.

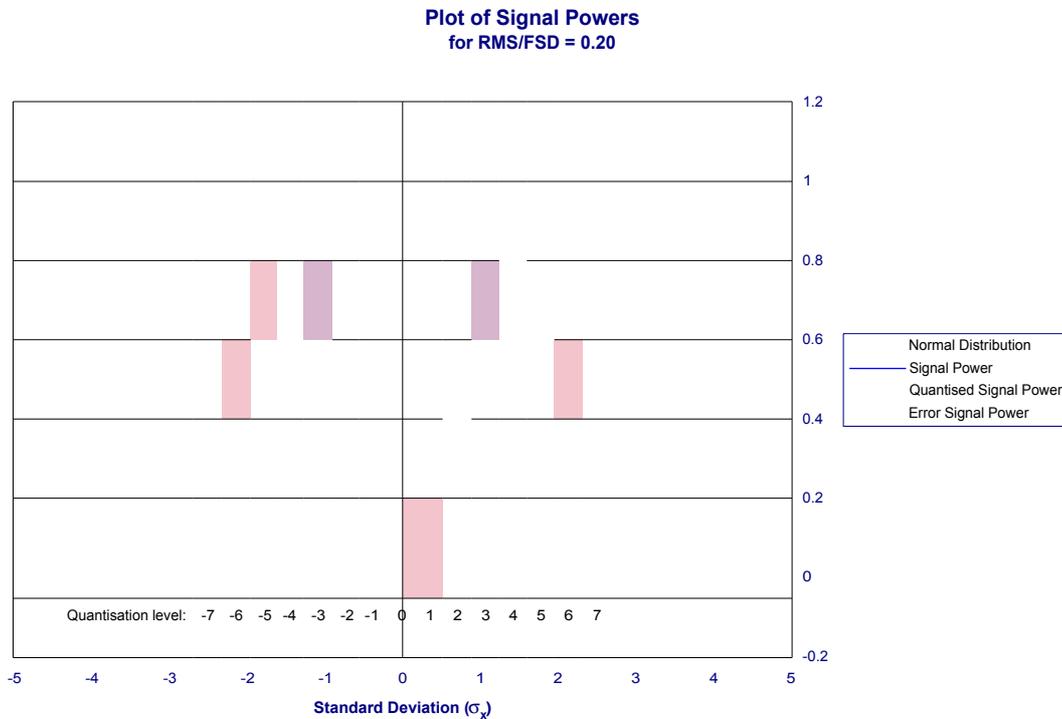


FIGURE 7.5.1. Power of analogue signal, quantised signal and error signal plotted as a function of standard deviation σ_x of the input signal for $\frac{RMS}{FSD} = 0.2$.

Table 7.5.1 is a list of $\frac{RMS}{FSD}$ and maximum SNR values for different bit lengths b . The calculated SNR is approximately 4 dB less than expected using equation 7.5.2, a result of using a Gaussian distributed input noise signal. Fitting a 5th order polynomial gives $\frac{RMS}{FSD} = \text{£} -7.8409b^5 + 13.199b^4 - 8.401b^3 + 24.04b^2 - 51.846b + 169.57 \text{] } \times 10^{-3}$.

Quantisation of the input signal introduces a degree of non-linearity in the power response of the system. The need for the system to have a linear response to small power fluctuations is critical. Figure 7.5.3 is a graph of $\frac{\text{Input signal power}}{\text{Quantised signal power}}$ as a function of $\frac{RMS}{FSD}$. In agreement with

Backer et al 1997 [16], setting $\frac{\text{RMS}}{\text{FSD}}$ to slightly less than the optimal SNR value of 0.2 results in the least non-linearity for small power fluctuations at the input of the system.

The results of this analysis agree with Kouwenhoven and Voute (2001) [89], who have calculated optimal input signals for 1, 1.5, 2, 4 and 8 bit digitisers.

7.5.3. Hardware: Analogue to Digital Converter. The AD9066 is a dual, 6 bit, 60 MSPS analogue to digital converter employing a patented interpolated flash architecture and is optimised for use with quadrature demodulators. The full scale input range of 500 mV is automatically DC biased about a fixed internal reference voltage, producing an offset binary format

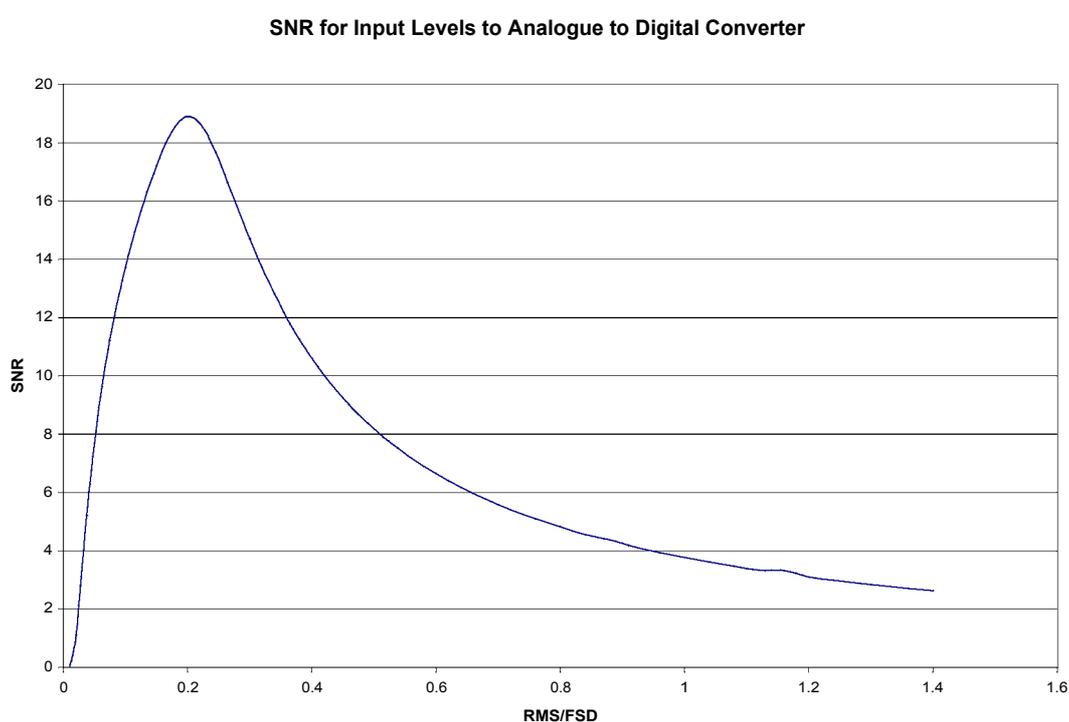


FIGURE 7.5.2. SNR of a 4 bit quantised signal as a function of the normalised analogue input signal level $\frac{\text{RMS}}{\text{FSD}}$.

TABLE 7.5.1. $\frac{\text{RMS}}{\text{FSD}}$ and maximum SNR values for different values of b

| No. of Bits (b) | $\frac{\text{RMS}}{\text{FSD}}$ | SNR [dB] |
|-----------------|---------------------------------|----------|
| 2 | 0.4 | 7.2 |
| 3 | 0.26 | 13.29 |
| 4 | 0.2 | 18.9 |
| 5 | 0.17 | 24.32 |
| 6 | 0.15 | 29.73 |
| 7 | 0.14 | 35.12 |
| 8 | 0.13 | 40.56 |

output. The relevant datasheet can be obtained at http://www.analog.com/UploadedFiles/Data_Sheets/120851866AD9066_a.pdf.

7.6. Analogue Frontend Assembly

7.6.1. Schematic Design. Figure 7.6.1 is the schematic design for the analogue frontend module, illustrating separated functional blocks and labelled components.

All input RF signals are impedance matched and AC coupled. The OFF-BOARD POWER CONNECTIONS employ transient suppression on the unregulated input voltage connections. Separate power supplies are used for the analogue and digital sections of the board to reduce ‘ground bounce’. This phenomenon occurs when a number of pins drive a current in the same direction simultaneously, causing a sudden drop in the ground plane potential.

The complex output signal of the differential amplifiers have off board connections to anti-aliasing low pass filters. The signals return to the ANALOGUE DIGITAL CONVERSION where they are sampled and quantised with 6 bit precision. This signal is truncated to 4 bits (see section 7.6.3).

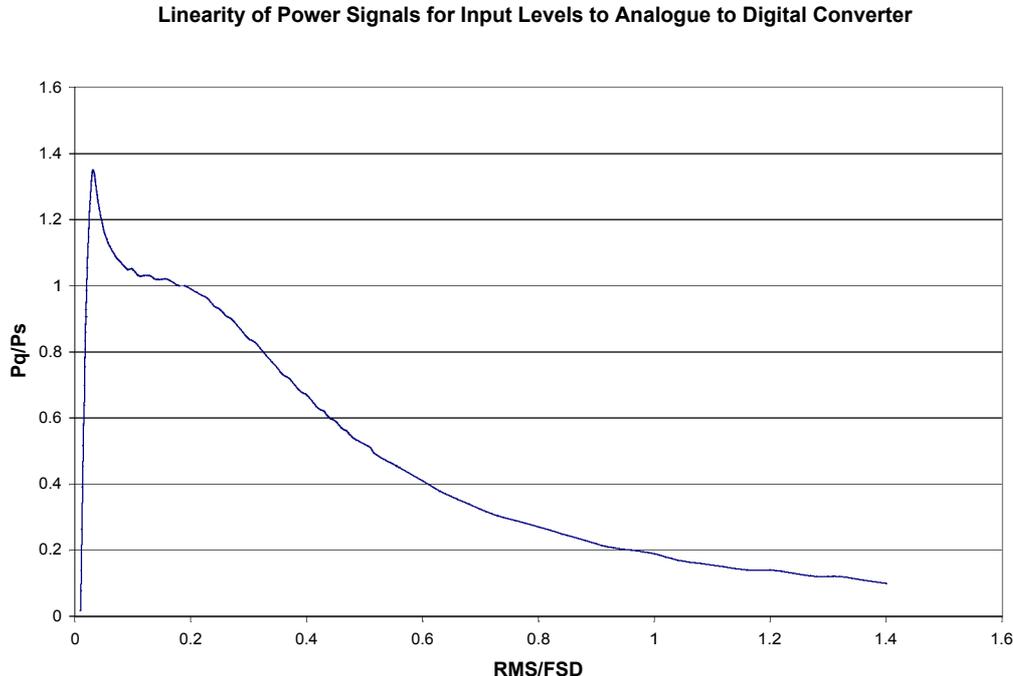


FIGURE 7.5.3. $\frac{\text{Input signal power}}{\text{Quantised signal power}}$ as a function of $\frac{\text{RMS}}{\text{FSD}}$ for 4 bit quantisation. Choosing an input value of marginally less than 0.2 results in the least degree of non-linearity.

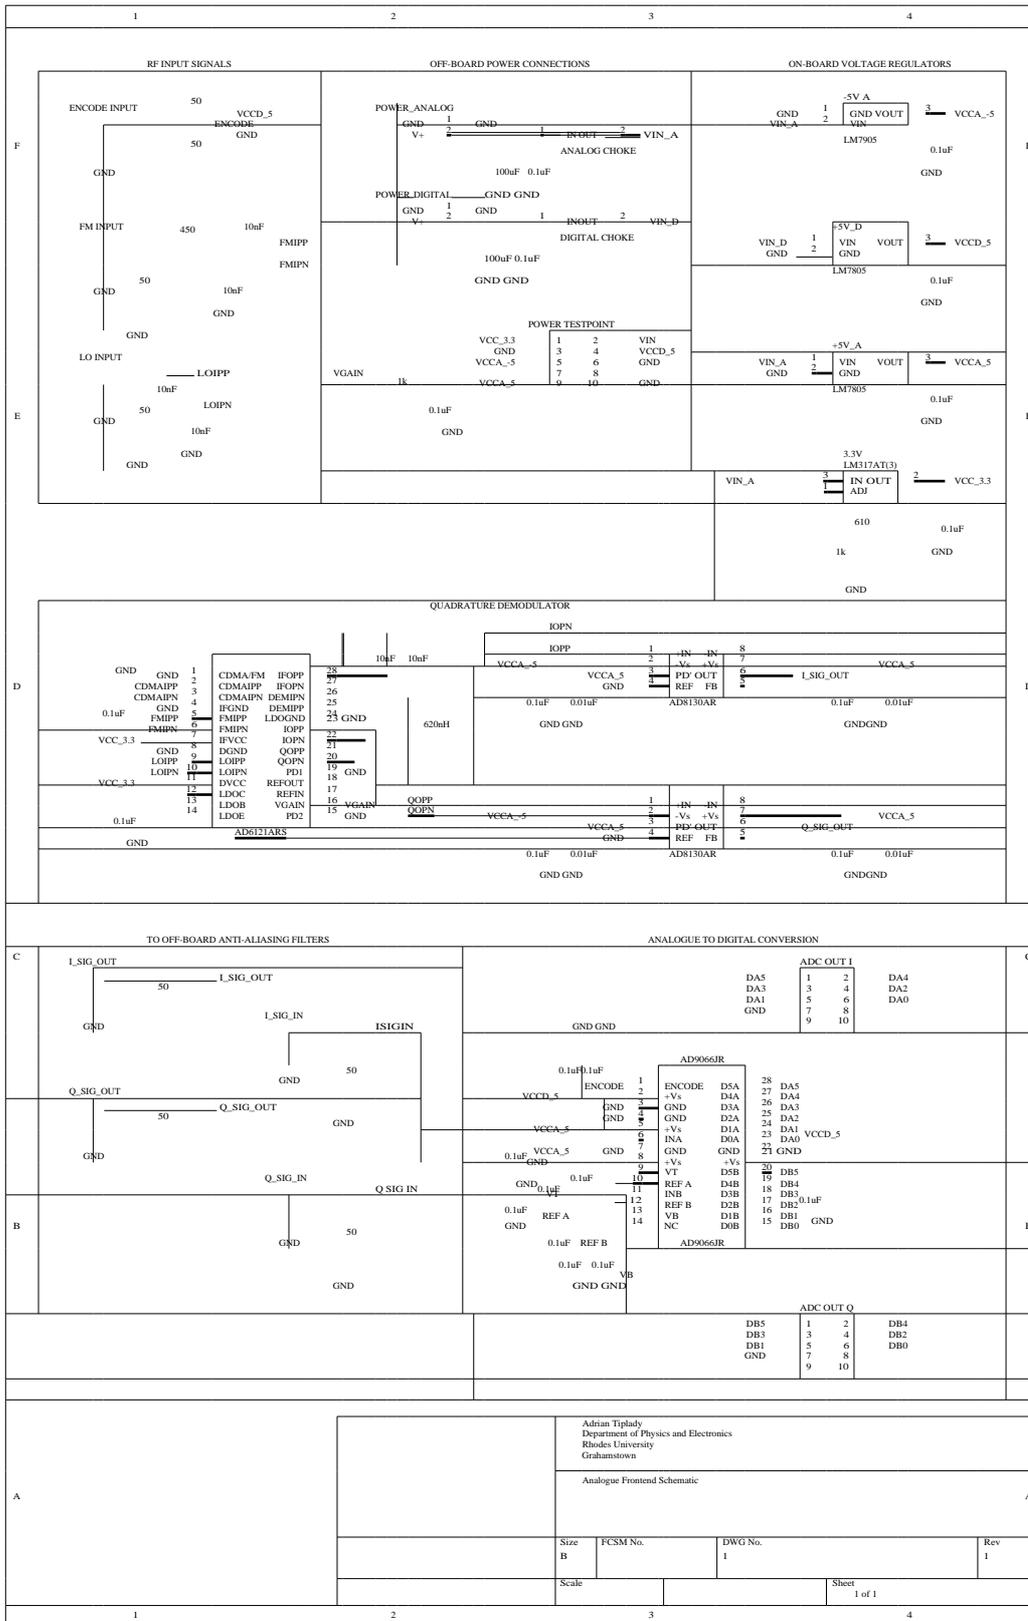


FIGURE 7.6.1. Schematic design for analogue frontend module

7.6.2. Manufacturing. The printed circuit board design guidelines provided by S.A. Signal Electronics [168] were followed in the manufacture of the analogue frontend module. Standard practice was followed to maintain signal integrity. Silkscreens for the analogue frontend printed circuit board are found in appendix F.

7.6.3. Signal Levels and Analysis. With an output range of 700 mV, the quadrature mixer makes optimal use of the dynamic range of the analogue to digital converter. However, by using only bits [1 : 4], the most significant bit can be used to check for saturation of the input signal. Setting the peak amplitude of the input signal to fill the full scale dynamic range of bits [0 : 4] results in an optimum SNR and linearity of the quantised signal.

7.7. Digital Backend Architecture

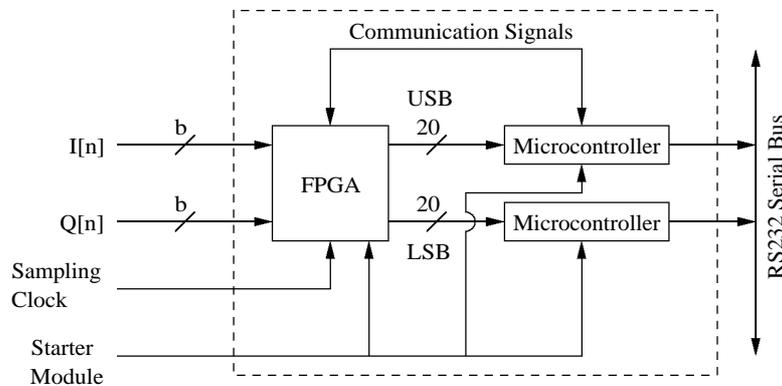


FIGURE 7.7.1. Schematic of the digital backend module. Each module accepts a 6 bit complex signal, which is truncated to 4 bits. Upper and lower sidebands are separated within the FPGA and 20 bit accumulated values, corresponding to 1 bin period, are read and stored by microcontrollers.

7.7.1. Schematic Design. Figure 7.7.1 is a high level schematic for the digital backend module. This unit consists of an FPGA for high speed digital signal processing and two microcontrollers, used for data stacking, temporary data storage and communication with the master controller. The FPGA accepts a 6 bit binary offset complex signal, which is truncated to a two's complement 4 bit signal, using bits [1 : 4]. The most significant bit [5] is used to monitor saturation of the input signal, which is adjusted by setting the analogue gain control of the frontend module. The FPGA produces sideband separated signals that are accumulated for a single bin period.

The accumulated 20 bit value is read by a microcontroller, which performs data stacking during an observation to produce an integrated pulse profile with a maximum of 1000 bins. Communication with a master controller is via an RS232 serial bus, using a unique ID tag for identification. Once the observation is complete, the master controller retrieves accumulated pulse profiles from each microcontroller, after which incoherent dedispersion takes place and a composite profile obtained.

7.7.2. Hardware: FPGA. The Xilinx Spartan-II XC2S200 is well suited to the implementation of FIR filters [3], with selectable I/O standards that reduce the need for off-chip impedance matching, on-board RAM and low skew clock distribution nets. Compatible printed circuit board footprints allow easy replacement and upgrades within the Spartan family. The required nominal operating frequency of approximately 32 MHz is well within the FPGA’s clock specification.

Full documentation of the XC2S200 can be found at <http://direct.xilinx.com/bvdocs/publications/ds001.pdf>.

7.7.3. Hardware: Microcontroller. The Atmel ATmega128 is a low power AVR 8 bit microcontroller that uses RISC architecture, running at a clock rate of 16 MHz. There is 128 kBytes of programmable flash and 4 kBytes of internal SRAM.

Full documentation of the ATmega128 can be found at http://www.atmel.com/dyn/resources/prod_documents/doc2476.pdf.

The Atmel AT90S8535 microcontroller is used in the implementation of the starter module. Full documentation can be found at http://www.atmel.com/dyn/resources/prod_documents/DOC1041.pdf.

7.8. FPGA Design Architecture

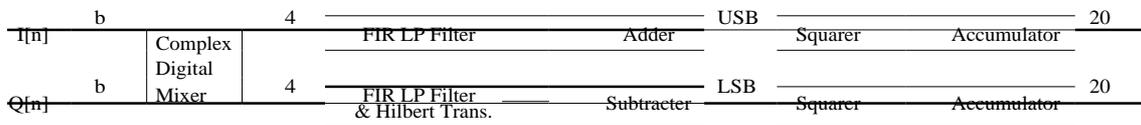


FIGURE 7.8.1. Internal schematic of FPGA

7.8.1. Signal Flow. Signal flow within the FPGA is illustrated in Figure 7.8.1, which consists of a digital complex mixer, FIR low pass filters, adder and subtractor for sideband separation, and squarer and accumulator for detection and post-detection integration. The Xilinx Foundation F3.1i Build 3.1.181 software was used for development and implementation of the FPGA design.

Figure 7.8.2 is the top heirarchial schematic design for the FPGA. The signal flows from left to right,using a parallel structure to ensure a new output per clock cycle. The design consists of seven functional components: clock, input interface, complex mixer, FIR filters, adder and subtractor, squarer and accumulator and the microcontroller interface.

The LOC tag is used on the schematics to indicate input/output pin assignments.

7.8.2. Clock. One of four delay locked loops in the XC2S200 is used. With global low skew clock distribution nets, this component ensures that signals at each clock input are synchronous, reducing timing delays between registers. All external clocked components ie. the analogue to digital converter, are driven by the FPGA.

7.8.3. Microcontroller Interface. The microcontroller interface connects all external communication signals used to initialise the FPGA, namely N[0:13], the Clear FPGA signal and the Enable FPGA signal.

7.8.4. FPGA Implementation Constraints. The implementation of constraints in the floormapping of FPGA designs is critical for such operations as signal processing. Net lengths are kept to a minimum or, where this is not possible, latches are used on longer nets. By using relationally placed macros, optimal structures are obtained for each individual module. With the use of these fixed structures, the FPGA resources used for the implemented design is approximately 52 % (calculated by the number of slices used).

7.9. Input Interface

The input interface is designed to latch the 6 bit complex input signals before removing a DC offset of 2^5 , due to the output format of the analogue to digital converter, in order to turn the signal into a two's complement format. The signal is truncated, using bits [1 : 4], and adjusted

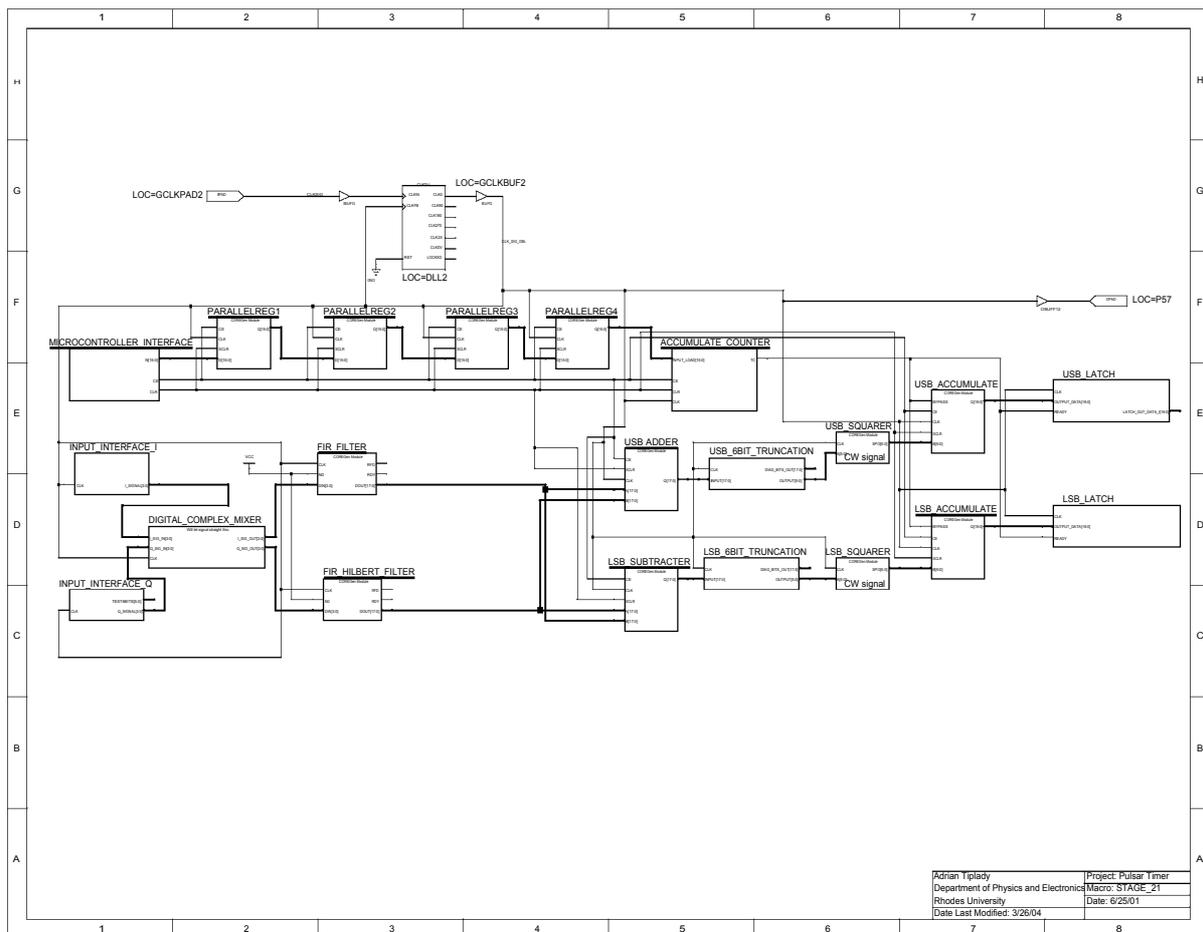


FIGURE 7.8.2. Top heirarchical structure for FPGA design

for truncation errors. These errors result from quantised data that is linearly scaled down by m bits, and are adjusted for by adding $\frac{1}{2}$ LSB of the 4 bit truncated signal to the untruncated 6 bit signal.

7.9.1. FPGA Input Interface Implementation. The schematic design for the input interface is shown in appendix H.1. Subtraction of the DC offset, the two's complement format conversion and truncation of the signal (including error correction) is performed by the use of lookup tables. The initialisation files `input_dc_subtract.coe`, `twos_comp_da_rom.coe` and `input_interface_mid4bits_LUT.coe` can be found in appendix L.

7.10. Digital Complex Mixing

Generalisation of the digital backend module to use a single set of filter coefficients requires the use of a complex mixer to perform the required eight way digital division of the 32 MHz baseband signal. Each narrowband channel is shifted into the FIR filter passband, which allows each digital backend module to employ identical bandpass filters.

Consider the baseband signal $x[n]$ with discrete time Fourier transform (DTFT) $X(\Omega)$. The DTFT $\tilde{X}(\Omega)$ of the shifted signal $\tilde{x}[n]$ is

$$\begin{aligned} \tilde{x}[n] &= \exp^{j\frac{k\pi}{4}n} x[n] \\ &= \cos\left[\frac{k\pi}{4}n\right] + j \sin\left[\frac{k\pi}{4}n\right] x[n] \\ \Rightarrow \tilde{X}(\Omega) &= X(\Omega) * \delta\left(\Omega - \frac{k\pi}{4}\right) \\ &= X\left(\Omega - \frac{k\pi}{4}\right) \end{aligned}$$

$\exp^{j\frac{k\pi}{4}n}$ is a discrete rotating phasor. By choosing an integer k , $-4 < k < 4$, the baseband signal $X(\Omega)$ is shifted about 0 rad in multiples of $\frac{\pi}{4}$. Considering the z plane, the signal is rotated by $\frac{k\pi}{4}$ radians about the unit circle. Figure 7.10.1 shows a $\frac{\pi}{4}$ rotation of $X(\Omega)$ where the $\frac{\pi}{4}$ subband has been shifted into the digital low pass filter passband.

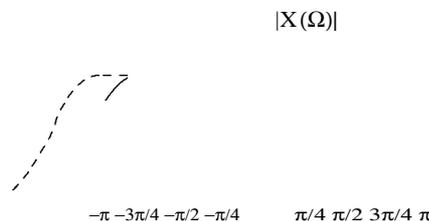


FIGURE 7.10.1. $\frac{\pi}{4}$ rotation of the z plane unit circle by the complex mixer

Implementation of the complex mixer limits k to $-3 < k < 3$ in order to make use of the outer channels $X(\Omega \pm \pi)$ as guard bands between adjacent 32 MHz analogue IF slices.

7.10.1. FPGA Complex Mixer Implementation. The schematic design for the digital complex mixer is found in appendix H.1. While the schematic shows the mixer for $k = 0$, for $k \neq 0$ the three ground signals connected to I SIG LUT INPUT[0:2] and Q SIG LUT INPUT[0:2] are replaced by a 3 bit counter incrementing by k . The set of lookup tables uses the 7 bit input signals (4 bit input signal, 3 bit phasor) to produce the correct 4 bit complex products concatenated onto an 8 bit output bus. The initialisation files `digital_mixer45_lut_i_signal.coe` and `digital_mixer45_lut_q_signal.coe` can be found in appendix L. The relevant real and imaginary components are added together to form digitally mixed complex signals.

7.11. FIR Filters

7.11.1. Considerations in Fixed Point Arithmetic Filter Design . Following Yates (2002) [178], the difference equation for an N tap FIR filter is

$$y[n] = \sum_{i=0}^{N-1} W_i x[n - i]$$

where W_i are the filter coefficients, and $x[n]$ is the input signal. This relation holds for an idealised, floating point system where the mathematical operations of addition, subtraction, multiplication and division are strictly defined over the rational numbers. In a digital system, however, both the input and output data are constrained by a fixed point binary representation. Quantisation of the idealised filter coefficients is performed for use in a digital system.

The quantised coefficient W_i' is

$$W_i' = \text{round}\left(\frac{W_i \times 2^b}{2^{b_w}}\right)$$

where b_w is the length of dataword representing W_i and the function `round()` maps the rational numbers W_i to the closest quantised level represented by a b_w bit dataword. The quantisation error is

$$\varepsilon_i = W_i' - W_i = \frac{\text{round}(W_i \times 2^{b_w})}{2^{b_w}} - W_i$$

where the maximum value for ε_i is $\frac{1}{2} \times \frac{1}{2^{b_w}} = 2^{-b_w-1}$. The maximum two's complement value that can be represented by a dataword of bit length b is $2^{b-1} - 1$. Therefore

$$b_w \dots = \log_2 \left(\frac{2^{b_w} - 1 - 1}{W_i} \right) \quad \mathbf{b} \quad W_i$$

The filter gain for a white noise input signal is

$$(7.11.1) \quad \text{White Noise Gain (WNG)} = \sqrt{\sum_{i=0}^{N-1} W_i^2}$$

If the result of each coefficient-input product fills M bits, the accumulated result for an N tap filter fills a maximum $b_x + \log_2 M$ bits where b_x is the bit length of the input dataword. To avoid overflows in the filter, the scaling factor b_w is determined by

$$b_w = b_{\text{acc}} - b_x - \log_2(\text{WNG})$$

where b_{acc} and b_x are the accumulator and input signal bit lengths respectively [178]. Consideration of fixed point arithmetic enables optimisation of the FPGA implemented filter floorplan, trimming unnecessary logic.

7.11.2. FIR Filter Design. The 4 bit output signals of the complex mixer are input to two 40 tap FIR filters, each of which uses 10 bit fixed point two's complement filter coefficients. Simulation of floating point and quantised filter coefficients indicate a 1% discrepancy in the cutoff frequencies of the two cases (see Figure 7.11.3).

Both sets of filter coefficients are based on a discrete Hamming weighted sinc(x) function

$$(7.11.2) \quad w[n] = 0.54 - 0.46 \cos \frac{2\pi n}{N-1}, \quad 0 < n < N-1$$

where N is the total number of samples across $w[n]$ ([138], pp. 242).

Consider a complex signal $x[n] = x_r[n] + jx_i[n]$, where the components $x_r[n]$ and $x_i[n]$ are processed by a 'real' and 'imaginary' filter respectively. To obtain sideband separation (See section 7.12 and Backer et al (1997) [16]), the imaginary filter transfer function is multiplied by a Hilbert transform, defined as ([138], pp. 358-359)

$$H(\Omega) = \begin{cases} j & -\pi < \Omega < 0 \\ -j & 0 < \Omega < \pi \end{cases}$$

Consider the signal $\hat{x}[n] = g[n] * x[n]$, where $g[n]$ is the impulse response of a low pass filter, and '*' the convolution operator. The signal $\bar{x}[n]$ is defined as

$$\begin{aligned} \bar{x}[n] &= h[n] * \hat{x}[n] \\ &= (h[n] * g[n]) * x[n] \\ \Rightarrow \bar{X}(\Omega) &= H(\Omega) X G(\Omega) X(\Omega) \end{aligned}$$

where $H(\Omega)$ is a Hilbert transform. $h[n] * g[n]$ is anti-symmetric. In other words, the DC gain of the imaginary filter

$$\text{Gain}_{DC} = \sum_i W_i = 0$$

where W_i are the filter coefficients.

Efficient sideband separation is required in order to reduce noise leakage between sidebands. To achieve this, a first approximation is made to match the white noise gain of the real and imaginary filters

$$\sum W_{i(\text{real})}^2 = \sum W_{i(\text{imaginary})}^2$$

This approximation, however, leads to very low noise rejection. With the absence of a DC block in the real filter, the integration of the squared coefficients is taken underneath a larger area. This implies that the peak frequency response of the imaginary filter will be greater than that of the real filter. In order to rectify this, a change in the initial approximation is required.

A second approximation to match frequency responses is done by introducing a DC block into the real filter such that

$$\text{Gain}_{\text{DC (real)}} = 0$$

Simple multiplication by a constant can be used to match white noise gains, although this does not achieve a satisfactory solution. Introducing a single channel DC block into the real filter would alter the shape of the passband in such a way that the real and imaginary responses would not match.

A DC block is introduced into the real filter coefficients by subtracting a hamming weighted constant from the original filter coefficients

$$\tilde{W}_i = W_i - (\text{Hamming Window})_i * \kappa$$

where κ is chosen to match bandpasses above a -3 dB threshold. Figure 7.11.1 shows the impulse response for the real and imaginary FIR filters. The term $(\text{Hamming Window})_i * \kappa$ introduces a -5 dB gain in the real filter coefficients.

Using equation 7.11.1,

$$\text{White Noise Gain}_{\text{real}} = 1283 \text{ Pz} \sim 0.88 \times 2^{10.5}$$

$$\text{White Noise Gain}_{\text{imaginary}} = 1218 \text{ Pz} \sim 0.85 \times 2^{10.5}$$

resulting in an output signal with dynamic range of approximately $2^{10.5} + 2^4 = 2^{14.5}$. This is done to ensure that the output of the adder and subtracter, which have a $-/2$ white noise digital gain, will fill the full scale dynamic range of a fixed bit length dataword.

Figure 7.11.2 is the frequency response of the real and imaginary FIR filters. The bottom diagram is an increased resolution of the filter passbands, showing matching passbands above the -3 dB threshold.

7.11.3. Filter Implementations and Techniques. DSP applications in FPGAs are structured on one of two different architectures: serialism or parallelism. Whilst serialism has the advantage of comparatively much smaller area-time products, which grow linearly with their operands, it manages only 1 output per N clock cycles, where N is the bit length of the function operands. Parallelism manages up to a sample per clock cycle, yet takes up considerably more real estate within the FPGA and grows exponentially with the amount of functionality required. Most FPGA filter implementations adopt a hybrid combination of the two, with digit-serial and distributed arithmetic being popular.

More specialised DSP structures include the use of customised lookup tables to perform addition and multiplication [2], specialised DSP operator arrays [14] and augmented cells for lookup table based FPGAs, allowing tighter inter-block communication [136].

7.11.3.1. Digit Serial Approach. Given their pipelined nature, bit serial multipliers implemented in FPGAs typically possess excellent time-area products. A bit serial design typically requires $1/n^{\text{th}}$ of the real estate needed for the equivalent n bit parallel implementation [13]. The drawback, however, is the need for n clock cycles to process an n bit data word.

A common bit serial approach used is digit serial processing, a combination of a bit serial and parallel implementation in which a data word is split into 2 bit words and processed digitwise serially. This technique provides a 17% reduction in area-time product of the bit serial

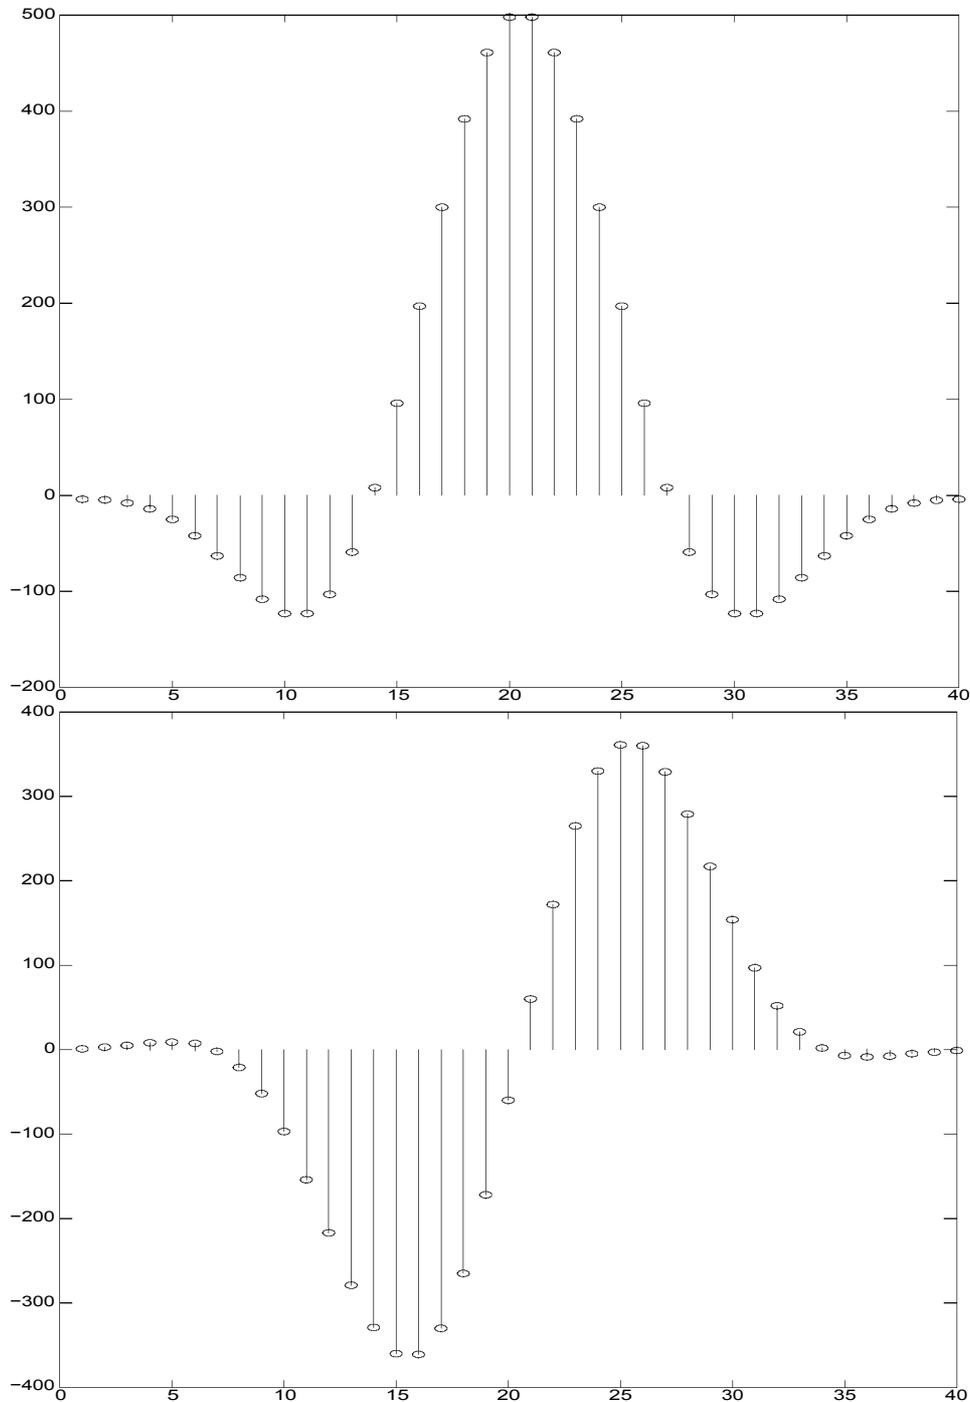


FIGURE 7.11.1. Filter coefficients for the real (top) and imaginary (bottom) FIR filters

implementation [97], offering a flexible tradeoff between data throughput and FPGA real estate size of operators.

Figure 7.11.4 is an unsigned digit serial multiplier. The multiplier requires a parallel input value, where all the bits are input simultaneously, and a serial value (split into N bit datawords), where all the bits are input sequentially. Different implementations of bit serial and digit serial structures are discussed by Lee and Sobelman (1997, 1998) [97, 98].

7.11.3.2. *Distributed Arithmetic Approach.* Distributed arithmetic delivers the performance of a parallel circuit within the space of a serial circuit, making it a highly efficient structure. It has been shown to produce efficient filter designs [128] with the successful implementation of complex digital filters for the purpose of radar signal processing [13, 129].

The array of configurable logic blocks (CLBs) within the FPGA is well suited to the use of distributed arithmetic. By viewing the CLB as a two bit functional block, they can be interconnected to perform any core DSP function.

Consider a linear, time-invariant system ie. an FIR filter. The response $y[n]$ of this system is

$$(7.11.3) \quad y[n] = \sum_{i=0}^{N-1} W_i x[n-i]$$

To obtain a single output $y[n]$, an accumulation of N products is required. This process would usually require a number of clock cycles in conventional, serial structured architectures.

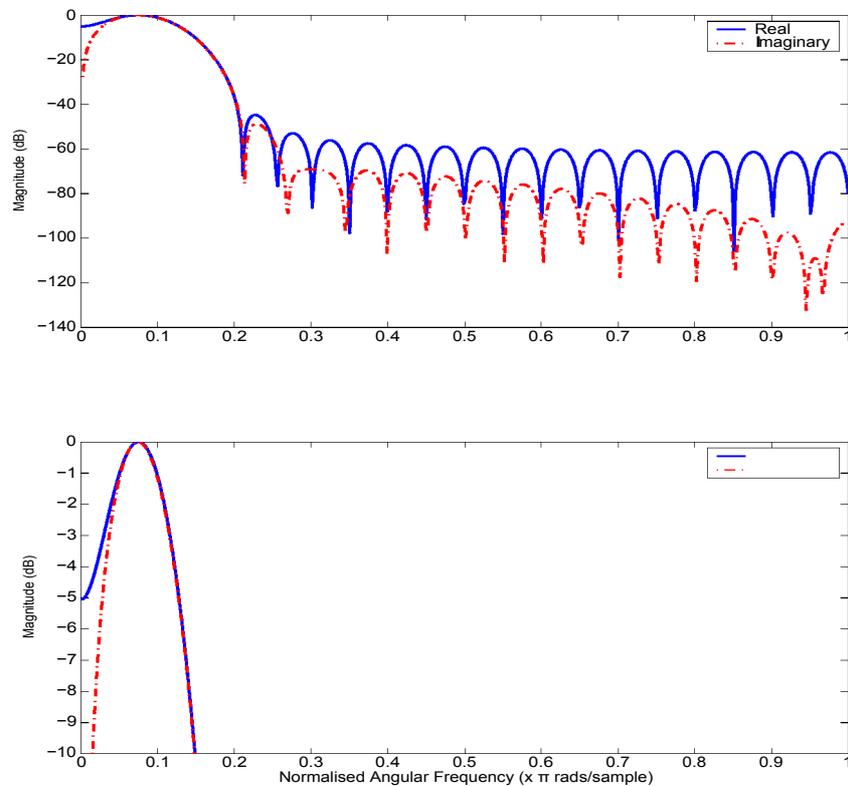


FIGURE 7.11.2. Frequency response of the real and imaginary FIR filters

In distributed arithmetic, this process is replaced by the use of a CLB look-up table (LUT) architecture. Dropping the variable n and rewriting $x[n - i] = x_i$, the binary representation of x_i is

$$x_i = -x_{i0} + \sum_{j=1}^M x_{ij} 2^{-j}$$

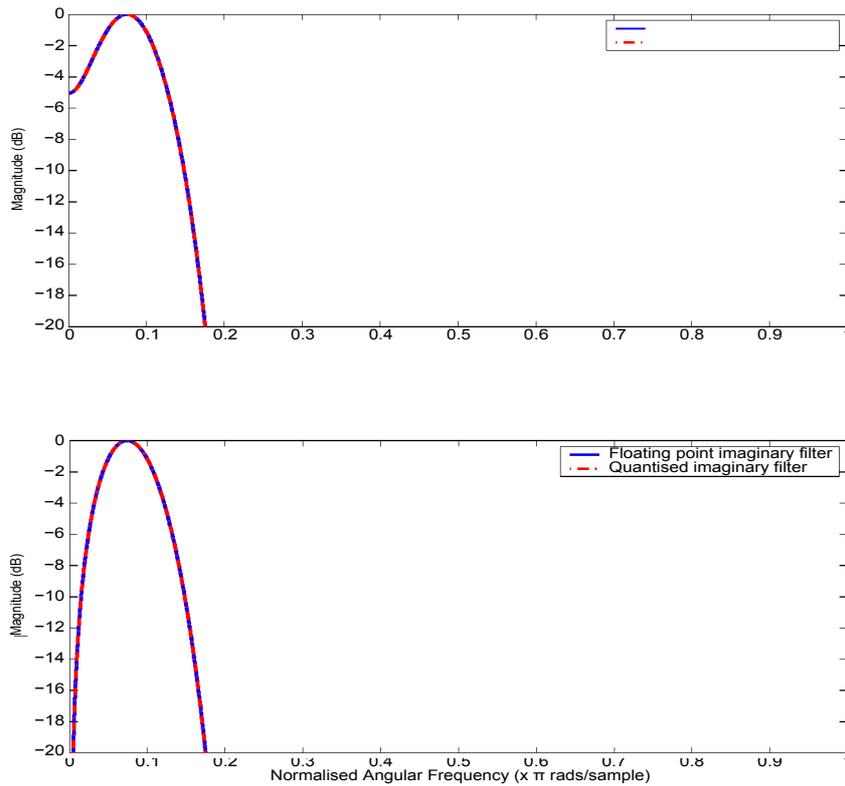


FIGURE 7.11.3. Comparison of frequency responses for floating point and 10 bit quantised filter coefficients.

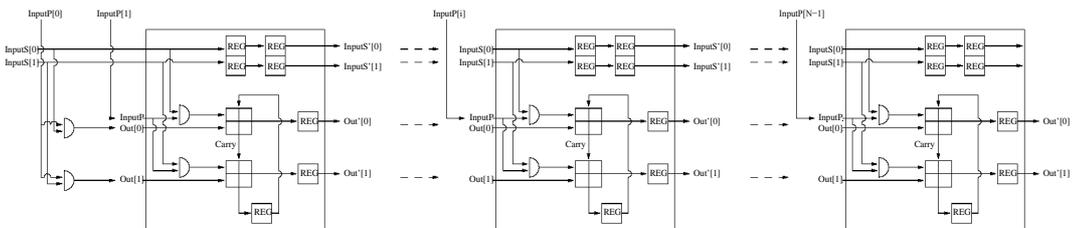


FIGURE 7.11.4. Unsigned digit-serial multiplier

where x_{ij} is a binary variable between 0 and 1, and $-x_{i0}$ represents a sign bit. Substituting into equation 7.11.3

$$(7.11.4) \quad y = \sum_{i=0}^N W_i [-x_{i0} + \sum_{j=1}^M x_{ij} 2^{-j}]$$

$$(7.11.5) \quad \sum_{i=0}^N W_i x_{i0} + \sum_{i=0}^N \sum_{j=1}^M W_i x_{ij} 2^{-j}$$

$$(7.11.6) \quad = -[W_0 x_{00} + W_1 x_{10} + \dots + W_{(N-1)x(N-1)0} + W_N x_{N0}] \\ + [W_0 x_{01} + W_1 x_{11} + \dots + W_{(N-1)x(N-1)1} + W_N x_{N1}] 2^{-1} \\ \cdot \\ \cdot \\ \cdot \\ + [W_0 x_{0M} + W_1 x_{1M} + \dots + W_{(N-1)x(N-1)M} + W_N x_{NM}] 2^{-M}$$

Equation 7.11.6 represents the distributed arithmetic algorithm. Each bit of the input sample is ANDed with all the bits of the particular scaled coefficients and then summed. By constructing LUTs addressed by the same scaled bit of the input samples, the sum of products within each pair of brackets can be obtained. The exponential factors of each bracket are scaling factors.

The LUTs can be implemented serially or in parallel. Highest speeds would be maintained in parallel, while maximum gate efficiency is achieved using a single LUT with parallel adder. It is possible to implement structures between these two extremes. Due to their reduced area of implementation, serial architectures have been attractive in the past. However, with the increase in gate density and size of reconfigurable logic, more parallel implementations of functional units has occurred to take advantage of data parallelism.

In the context of an FIR filter, equation 7.11.6 can be represented graphically as in Figure 7.11.5. The LUT is addressed by input samples a bit at a time. For a large number of taps, multiple LUTs can be used and their results added. In a complete parallel implementation, all inputs would be computed in parallel and then combined in a shifting adder tree [12].

7.11.4. FPGA Filter Implementation. The Xilinx Distributed Arithmetic FIR Filter V5.0 LogiCORE was used for the implementation of the FIR filters. This core is structured to produce area efficient FPGA implementations. For more information on the design structure of the DA FIR Filter V5.0, the datasheet can be downloaded from http://www.xilinx.com/ipcenter/catalog/logicore/dpcs/da_fir.pdf. The parallel structure reduces the 4 bit input samples to single bit data subwords, generating an 18 bit output sample every clock cycle. The use of symmetric and anti-symmetric filter coefficients (see Figure 7.11.1) reduces the number of hardware multiplications by one half. The initialisation files `fir_filter_0125_version2.coe` and `fir_hilbert_filter_0125_version2.coe` can be found in appendix L.

7.12. Sideband Separation

Upper and lower sidebands are obtained by adding or subtracting the complex components

$$\text{Upper sideband (USB)} = Y_r(\Omega) + \tilde{Y}_i(\Omega) = X_U(\Omega) + X_U^*(-\Omega)$$

$$\text{Lower sideband (LSB)} = Y_r(\Omega) - \tilde{Y}_i(\Omega) = X_L(\Omega) + X_L^*(-\Omega)$$

where $Y_r(\Omega)$ is the real filter output, $\tilde{Y}_i(\Omega)$ the imaginary filter output, and $X_U(\Omega) = X_U^*(-\Omega)$ (see appendix C.1 for mathematical detail). Following the implementation of a -5 dB block in the real filter, leakage between adjacent sidebands is reduced to -8 dB.

7.12.1. FPGA Sideband Separation Implementation. The adder and subtracter are 18 bit input, 18 bit output core generated modules, which separate the complex signal into an upper and lower sideband. With a white noise digital gain of 2^{15} , the output signals have a dynamic range of approximately 2^{15} . Adjusting for truncation errors, the signal is truncated to 6 bits using a VHDL (VHSIC Hardware Description Language) implemented module, the code for which can be found in 20BIT_CLIPPER.vhd in appendix L.

7.13. Detection and Post-Detection Integration

The radiometer equation

$$(7.13.1) \quad \Delta T_{\text{sys}} = \frac{\sqrt{T_{\text{sys}}}}{B \times \tau}$$

describes the statistics of the detection process of noise signals, where T_{sys} is the system temperature, B the receiver bandwidth and τ an integration time constant [57].

Each sideband signal is square law detected and the 6 bit output signal is accumulated for 1 bin period, the length of which is a user defined parameter.

7.13.1. FPGA Squarer Implementation. The squarer is implemented as a 6 bit input, 6 bit truncated output lookup table. The initialisation file square4_distrib_mem.coe can be found in appendix L.

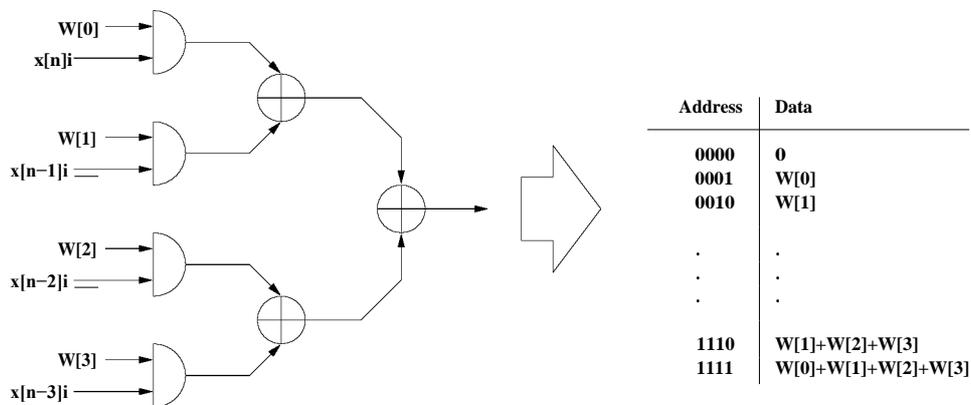


FIGURE 7.11.5. Adder tree implemented as a LUT addressed by input samples

7.13.2. FPGA Accumulator Implementation. The accumulator consists of a counter controlled by the microcontroller interface, accumulator and output interface (see appendix H.1 for schematic designs). The accumulated data is latched onto the output interface before the accumulator is cleared by the counter. This occurs when the counter matches a predefined number $N[0:13]$. The cleared signal initiates the read data routines on the microcontrollers.

7.14. Data Stacking

The data stacking performed within the microcontrollers requires knowledge of the exact period of the pulsar during the integration process.

By the Central Limit theorem [162], the amplitude distribution of the squared and integrated noise signal approaches a Gaussian distribution if sufficient samples are accumulated during a single bin period. The variance of the Gaussian distributed bin samples in the integrated pulse profile decreases with the number of integration periods N .

By the propagation of errors [22], for large N the standard deviation σ_μ of the averaged signal μ is

$$\sigma_\mu = \frac{\sigma_E}{\sqrt{N}}$$

where σ_E is the standard deviation of the additive noise signal [51, 142]. The SNR of the input signal is

$$\text{SNR}_{\text{in}} = \frac{x}{\sigma_E}$$

whilst for the averaged signal

$$\begin{aligned} \text{SNR}_{\text{avg}} &= \frac{x}{\sigma_\mu} \\ &= N \text{SNR}_{\text{in}} \end{aligned}$$

The gain in SNR for an averaged signal is therefore directly proportional to the square root of the number of averaged samples N .

1000 samples across a pulse profile results in 32 bits of dynamic range per bin. Using the full scale dynamic range of the FPGA implemented accumulator, a maximum value for the total number of samples per bin is 2^{12} . Given that the FPGA accumulates $\sqrt{6}$ bit values, the maximum value for $N = 2^{26}$, corresponding to a maximum increase in SNR of $N = 2^{13}$ per accumulated bin, or an increase of 2^7 for the integrated pulse profile.

7.14.1. Microcontroller Algorithm: ATmega128L (Data Capture). Figure 7.14.1 is a simplified flowchart indicating the upper sideband microcontroller algorithms. A more detailed flowchart is located in appendix G. Assembler code is located in `FPGADataReaderUSB.txt` in appendix L. The lower sideband microcontroller performs the same algorithm, but does not perform FPGA initialisation procedures ie. setting the number of samples per bin and clearing the FPGA. A detailed flowchart is located in appendix G. Assembler code is located in `FPGADataReaderLSB.txt` in appendix L.

The processor can exist in one of three states: ready, enabled or disabled. The ready state indicates when the processor can begin an observation at a predefined time. An observation will

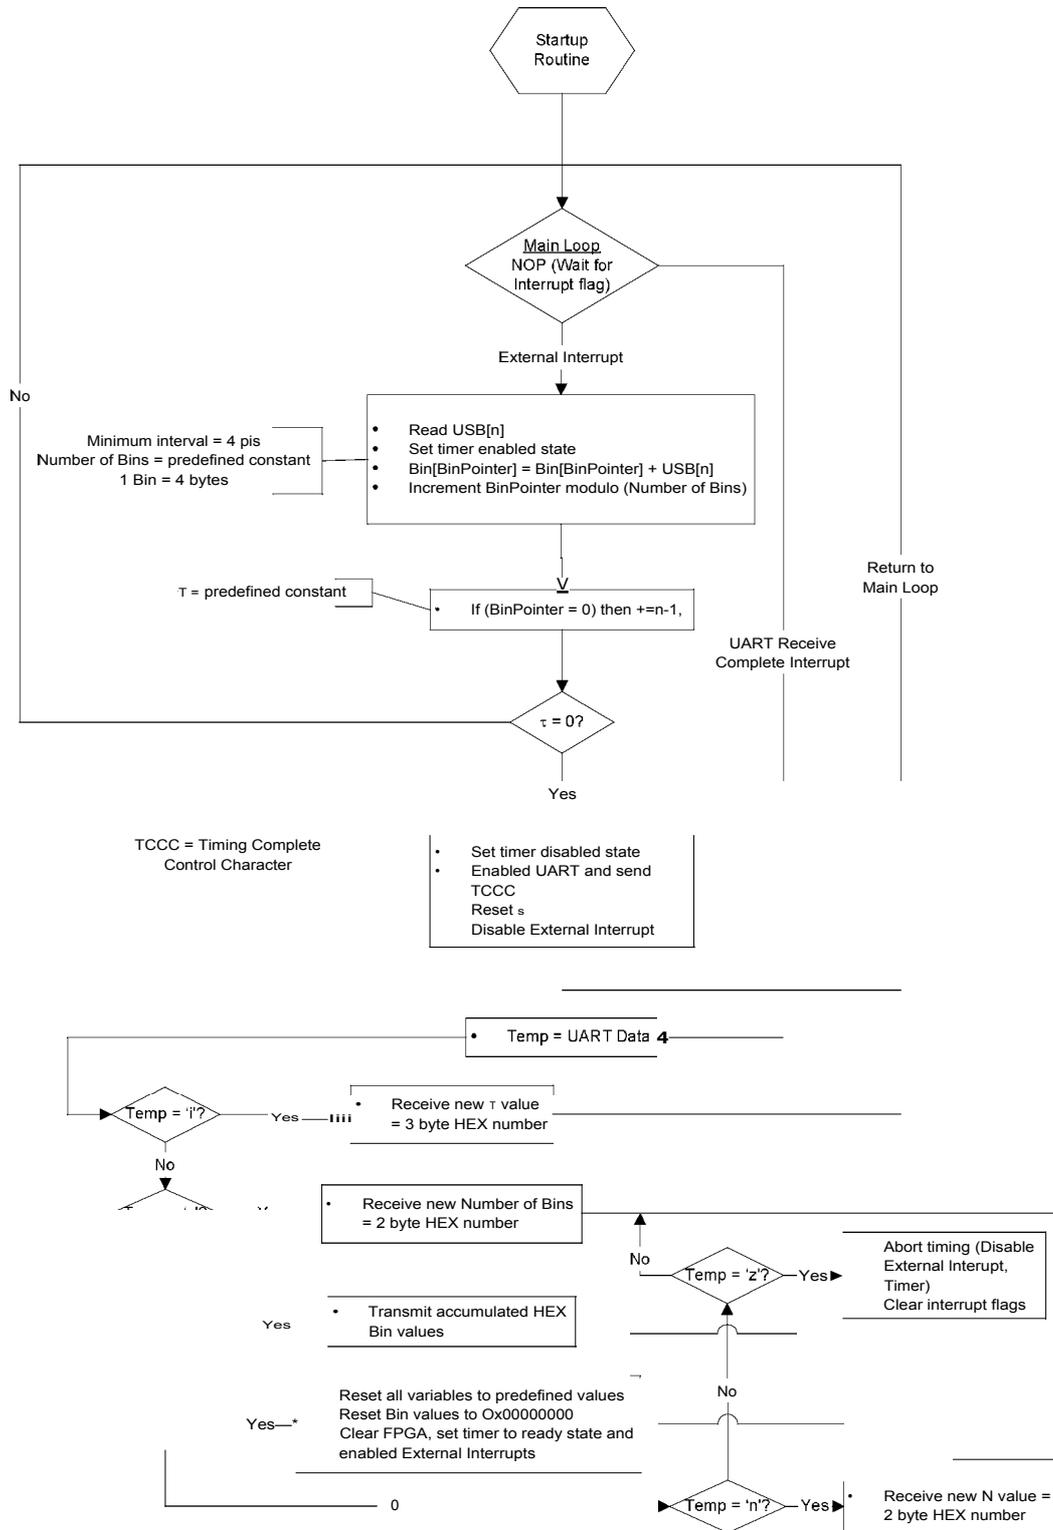


FIGURE 7.14.1. Simplified flowchart for upper sideband data capture, implemented on the ATmega128L

only begin in this state, during which it will be in an enabled state. All serial communication is ignored when the processor is enabled, apart from the timing abort signal. Upon completion of an observation, the processor is disabled. To place in a ready state, all variables, interrupts and data arrays are re-initialised.

The startup routine initialises interrupts, variables and input and output ports. The system is interrupt driven, and so executes a no operation procedure until one of two interrupt flags change state: the external interrupt, which is triggered by the FPGA to indicate latched data, or the Universal Asynchronous Receiver Transmitter (UART) Receive Complete interrupt.

The external interrupt algorithm has a maximum runtime of 4 μs , limiting the resolution across a pulse profile to a minimum of 4 μs . During the enabled state, the external interrupt is triggered by the FPGA, indicating that a new 20 bit accumulated bin value has been latched. The data is read and added to the correct bin number. A counter \mathbf{T} indicates the number of pulse period integrations to be performed. Once $\mathbf{T} = 0$, the processor moves into the disabled state and a Timing Complete Control Character is sent to the master controller. \mathbf{T} is reset and the external interrupt is disabled to ensure no more data is read.

The UART Receive Complete Interrupt is triggered when data is received. The characters sent instruct the microcontroller to perform one of six different procedures. These are: change total number of pulse period integrations (\mathbf{T}), change number of bins (BinCount), instruct the microcontroller to send the data in each bin, reset the processor and FPGA, abort the current observation and change the number of samples per bin to be accumulated (N). Aborting the observation will put the processor into a disabled state and require a reset of the processor to begin an observation.

7.15. Starter Module

The starter module is a standalone unit that initiates pulsar observing experiments, commensurate with a precision 1 pulse per second (1 PPS) station signal.

7.15.1. Microcontroller Algorithm: AT90S8535 (Starter Module). Figure 7.15.1 is a simplified flowchart of the starter module algorithm. Assembler code is located in `StarterModule.txt` in appendix L. A user defined start time for an observation is sent to the starter module by the master controller. Immediately prior to the start of an observation, the starter module enables an external flip-flop. The output of the flip-flop, connected to the ENABLE FPGA connector, is toggled by the 1 PPS signal and the observation begins. Once the processor is enabled, the flip-flop is disabled. The ENABLE FPGA signal is cleared by the user once the processor is in the disabled state.

7.16. Digital Backend Assembly

7.16.1. Schematic Design. To reduce development cost of the prototype, the starter module and digital backend module were implemented on the same printed circuit board.

Figure 7.16.1 is the schematic design for the digital backend and starter modules, showing labelled components and separate functional blocks. Both upper and lower sideband microcontrollers read data from the FPGA, whilst the upper sideband also performs initialisation of the FPGA via the communications bus. The STARTER MODULE uses an AT90S8535 microcontroller. The station signals that are input into the starter module are the 1 PPS, NASA time code and an emergency reset button.

7.16.2. Digital Testpoints. The AD7303 was used at digital signal testpoints. It is a serial input, dual voltage output 8 bit digital to analogue converter. The full product datasheet can be downloaded from http://www.analog.com/UploadedFiles/Data_Sheets/48853020AD7303_0.pdf.

7.16.3. Manufacture. The printed circuit board design guidelines provided by S.A. Signal Electronics [168] were followed in the manufacture of the digital backend module. Standard

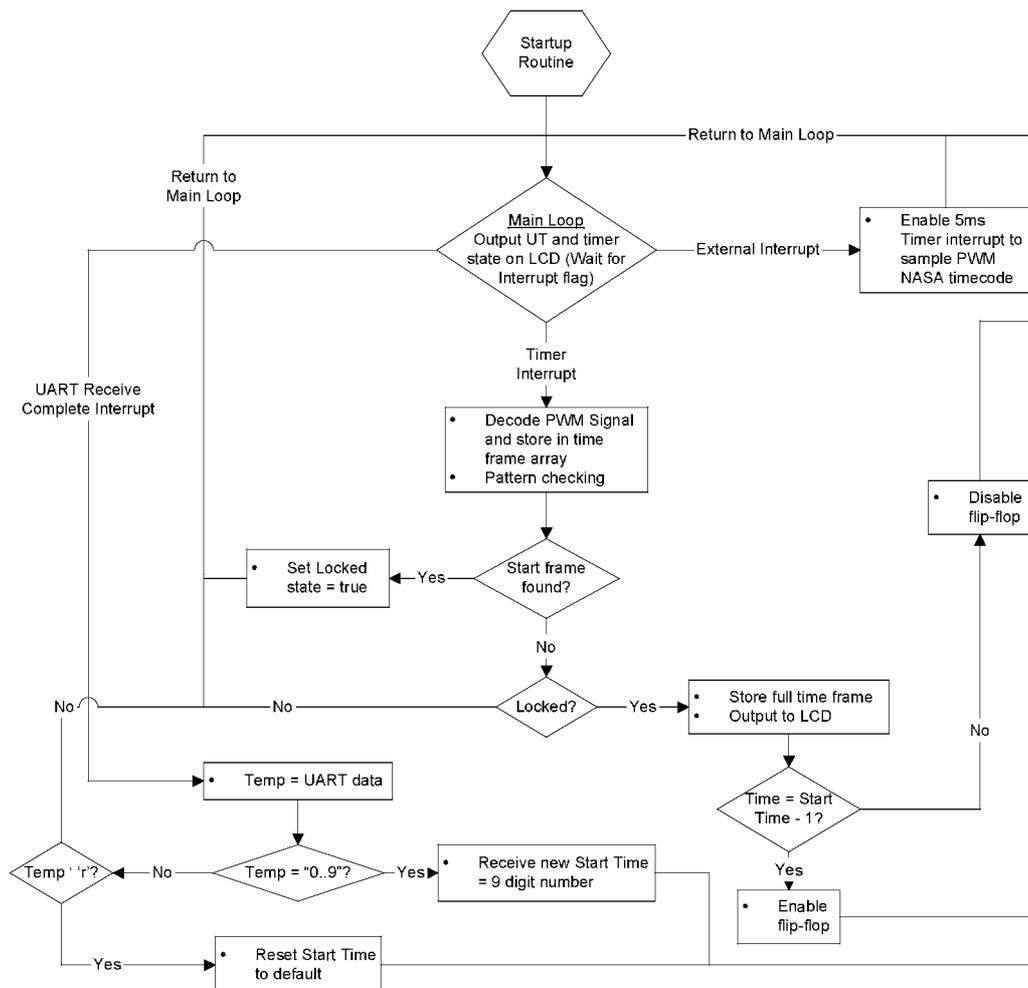


FIGURE 7.15.1. Simplified flowchart for starter module, implemented on the AT90S8535

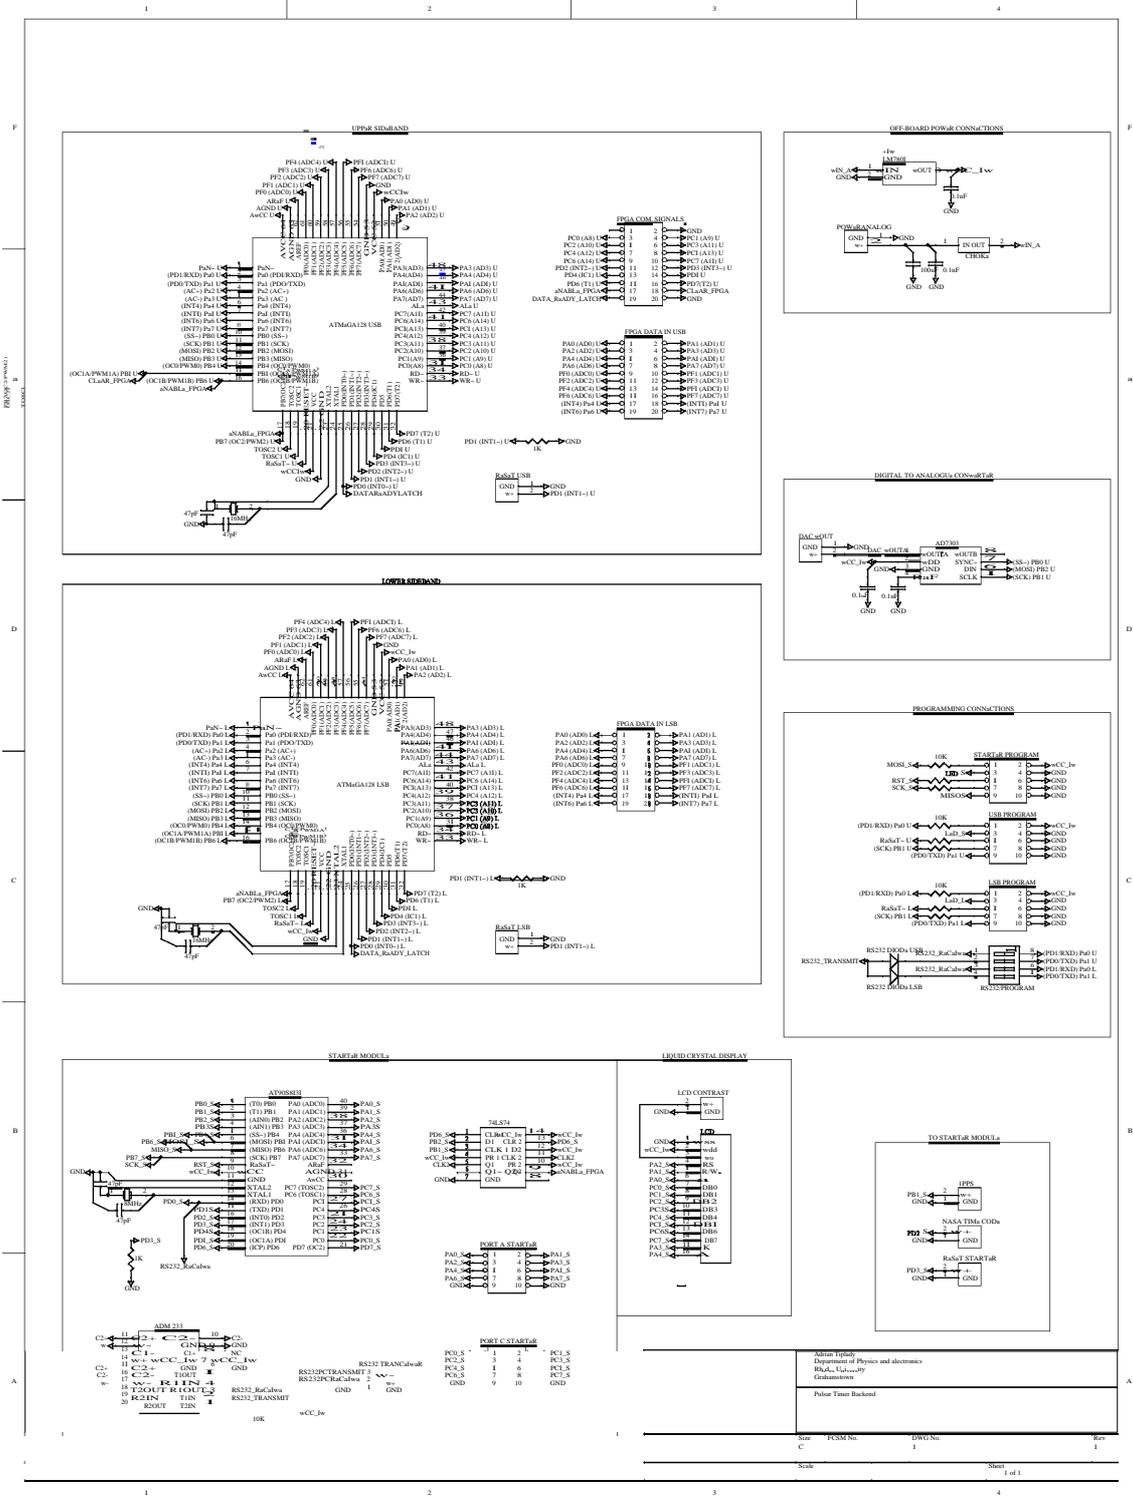


FIGURE 7.16.1. Schematic design for digital backend

practice was followed to maintain signal integrity. Printed circuit board silkscreens for the digital backend are in appendix F.

7.17. Pulsar Processor Assembly

7.17.1. **Prototype Specification.** The prototype was assembled in a standalone rack with single analogue frontend module, digital backend module and an XC2S200 evaluation board. The front panel connections included inputs for unregulated analogue and digital power supplies, regulated power outputs, a vernier potentiometer for analogue gain control, clock inputs and outputs, local oscillator input, quadrature mixer outputs (for anti-aliasing filters), ADC inputs, NASA time code input, 1 PPS input, digital to analogue output, an RS232 serial interface, signal LEDs and an LCD display. The digital to analogue output was used as an indicator to determine the level of saturation at the input of the FPGA.

7.17.2. **Eurocard Specification.** The final implementation of the pulsar processor uses the eurocard specification, incorporating the DIN 41612 and IEC 603-2 connector standards, IEEE 1101 PC board standards and the DIN 41494 and IEC 297-3 rack standards. The implementation of the eurocard specification makes use of a motherboard into which digital backend modules are slotted. The hardware implementation has been redesigned for the use of both polarisations. Schematics for the analogue frontend, digital backend, motherboard and starter modules are shown in appendix J.1, J.2, J.3 and J.4 respectively. Silkscreens for the printed circuit boards are shown in appendix K.

7.18. Operation

A general outline of proposed operation of the pulsar processor is given.

- (1) An observation is initiated by the master controller. The user input start time is sent to the starter module, which, while not observing, holds the FPGA and microcontroller accumulators in a disabled state.
- (2) The master controller initiates a setup control sequence. This sequence clears the accumulators in the FPGA and microcontrollers, and enables microcontroller interrupts connected to data ready latches on the FPGA. The latches initiate an interrupt routine at the end of each bin period, causing the microcontrollers to read the accumulated data from the FPGA and store in SRAM.
- (3) The master controller is given three user defined parameters: the number of bins per pulse period, the number of accumulated samples per bin and the length of integration. To ensure that the period of the total number of accumulated samples is an integer multiple of the pulse period, the programmable sampling clock is adjusted about a

nominal 32 MHz. The algorithm performed is:

$$\begin{aligned} \text{Nominal sampling period } P &= \frac{1}{32 \times 10^6} = 3.125 \times 10^{-8} \\ \text{Pulse period per bin} &= \frac{\text{Pulsar period}}{\text{Number of bins}} \\ \text{Nominal total samples per bin} &= \frac{\text{Pulse period per bin}}{\text{Nominal sampling period } P} \\ \text{Total samples per bin} &= \text{round}(\text{Nominal total samples per bin}) \\ \Rightarrow \text{Sampling clock frequency} &= \frac{\text{Total samples per bin} \times \text{Number of bins}}{\text{Pulsar period}} \end{aligned}$$

The sampling clock is set by the master controller. The microcontrollers, after obtaining observation parameters, will set the FPGA accumulator to the correct number of samples per bin.

- (4) Using the NASA timecode, the start of observation is commensurate with the 1 PPS station signal. The starter module enables the FPGA and microcontroller accumulators.
- (5) The FPGA accumulates N samples, then flags the microcontrollers to read the accumulated value, which is latched whilst the FPGA continues to accumulate the next bin. The microcontrollers have 1 bin period to read and store the accumulated value.
- (6) After a set number of pulse period integrations, the microcontrollers disable themselves until they have been reset by the master controller. Once identified and requested to do so, the microcontrollers will send data to the master controller.

CHAPTER 8

Diagnostic Tests and Field Trials

Testing of the pulsar processor was performed in a laboratory, using an artificial noise source, before final installation of the prototype instrument at HartRAO. Test results, together with ‘first light’ pulsar observations, are presented in this chapter.

8.1. Statistical Analysis and Diagnostic Tests

The statistical distribution of an input noise source at each stage of the FPGA signal flow was determined and measured. A standard Elsat satellite television receiver was used, bandpass filtered ($B = 50\text{--}100$ MHz) and mixed to baseband using an 80 MHz (160 MHz) local oscillator. Datastreams were sampled using an Agilent LogicWave E9340A logic analyser. With a maximum clock rate of 250 MHz, this instrument performs 32 bit sampling with a 32768 sample buffer size.

8.1.1. Signal Distribution Statistics. The empirical characterisation of a signal in terms of a random process model falls within a branch of statistics known as estimation theory.

The gaussian probability mass function, used to describe the distribution of a discrete white noise signal, is

$$(8.1.1) \quad P(\mathbf{x}) = \sqrt{\frac{1}{2\pi\sigma_x}} \exp\left[-\frac{(x-\mu_x)^2}{2\sigma_x^2}\right]$$

where

$$\sigma_x = \sqrt{\frac{1}{N} \sum_{i=1}^N x_i^2}$$

is the standard deviation of a sample population of N samples, and

$$\mu_x = \frac{1}{N} \sum_{i=1}^N x_i$$

is the population mean ([116], pp. 476-484).

A white noise model of the quantisation error can be reasonably assumed for a small quantisation step ([138], pp. 415). However, for a small number of bits this approximation is crude and departures can be expected from the theoretically predicted statistics (σ_x and μ_x), caused by convolution of a Gaussian quantised signal with a non-Gaussian quantisation error signal.

The error between the means of a sampled and unsampled stationary stochastic process is

$$(8.1.2) \quad \varepsilon_\mu = \sqrt{\frac{\sigma_x}{N}}$$

For large N , $\varepsilon_\mu \rightarrow 0$. The theoretical values of σ_x and μ_x , calculated for each stage of the signal flow, are determined for $N = \infty$.

The cumulative distribution function is

$$4b(m) = \int_{-\infty}^m x_j dx$$

For a normalised distribution, $4b(\infty) = 1$. The histogram, theoretical and measured values for $4b(m)$ are referred to as $4b(m)_h$, $4b(m)_{Th}$ and $4b(m)_M$ respectively. Dropping the subscript x , the theoretical and measured values for the standard deviation and mean are defined as σ_{Th} , σ_M , μ_{Th} and μ_M respectively.

8.1.2. Artificial Input Noise Source. Figure 8.1.1 is the histogram, theoretical and measured probability mass distribution of the 6 bit noisy input signal, which is dominated by the Gaussian distribution of a white noise source. For large $N = 32768$, the histogram

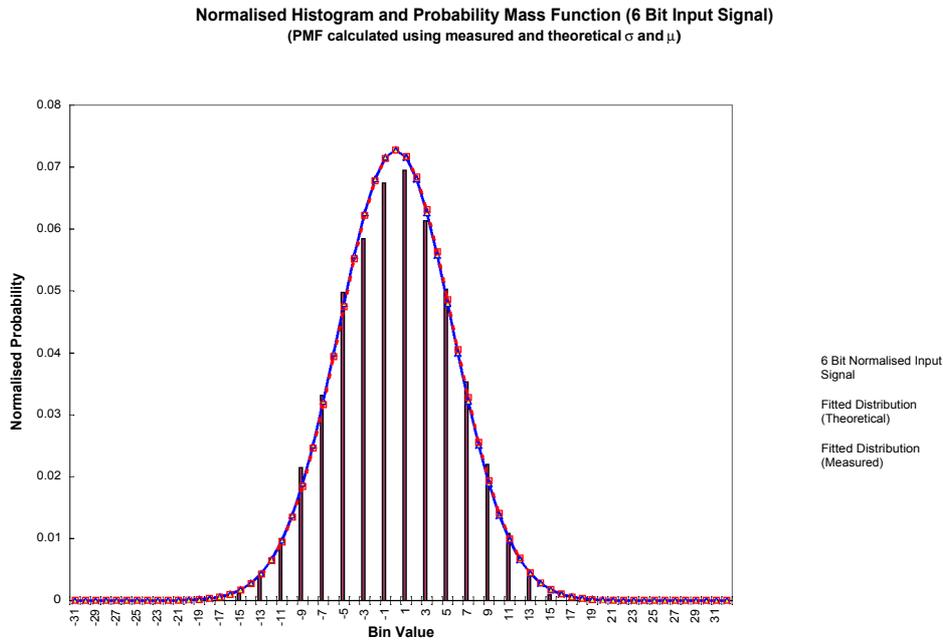


FIGURE 8.1.1. Normalised histogram and probability mass distributions of the 6 bit noisy input signal $x[n]$. $\sigma_{Th} = \sigma_M = 5.48$, $\mu_{Th} = 0$ and $\mu_M = 0.076$

closely matches the theoretical probability mass distribution. Normalisation of the histogram is performed by

$$y_i = \frac{z_i}{N \times dx}$$

where y_i is the normalised probability, z_i is the original summation, N the total number of samples and dx the bin width.

The choice of bin width seems natural here, with only 63 possible values for a 6 bit input. In later stages of the signal flow, however, this choice is based on a compromise between high

statistical noise, resulting from too small a value for dx , and low x axis resolution, as a result of large dx .

The error in μ_M is

$$\varepsilon_\mu = \frac{\sigma_{Th}}{\sqrt{N}} = 0.0303$$

(see equation 8.1.2), which accounts for 40% of the calculated error $\mu_M - \mu_x$. Instrumental errors may account for further discrepancies between μ_M and μ_{Th} .

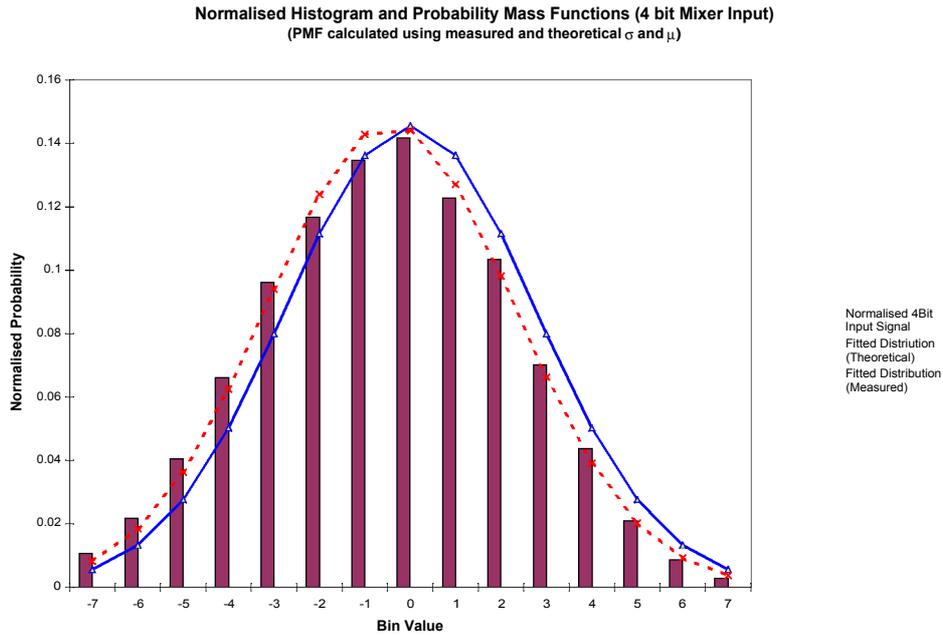


FIGURE 8.1.2. Normalised histogram and probability mass distributions of the 4 bit mixer input signal. $\sigma_{Th} = 2.740$, $\sigma_M = 2.735$, $\mu_{Th} = 0$ and $\mu_M = -0.437$

8.1.3. Input Interface. Figures 8.1.2 and 8.1.3 are the probability mass distributions and cumulative distributions of the 4 bit input signal to the complex mixer. Differences between measured and theoretical distributions result from a truncation error, which introduces a constant offset of $\mathbf{f}^{2l} \text{LSB} = \mathbf{f}0.5 \approx \mu_M = -0.437 \mathbf{f} \varepsilon_\mu$.

σ_{Th} is calculated as

$$\sigma_{Th} = \sigma_{Th}^0 \times G$$

where G is the digital gain of the relevant stage and σ_{Th}^0 the standard deviation of the input signal to the relevant stage. For 1 bit truncation, $G = \frac{1}{2}$. Hence

$$\sigma_{Th} = 5.480 \times \frac{1}{2} = 2.740$$

8.1.4. Digital Complex Mixer. Figure 8.1.4 is the histogram and probability mass distributions of the 4 bit mixer output signal. With a digital gain of unity, there is little change in the signal distribution relative to the input signal.

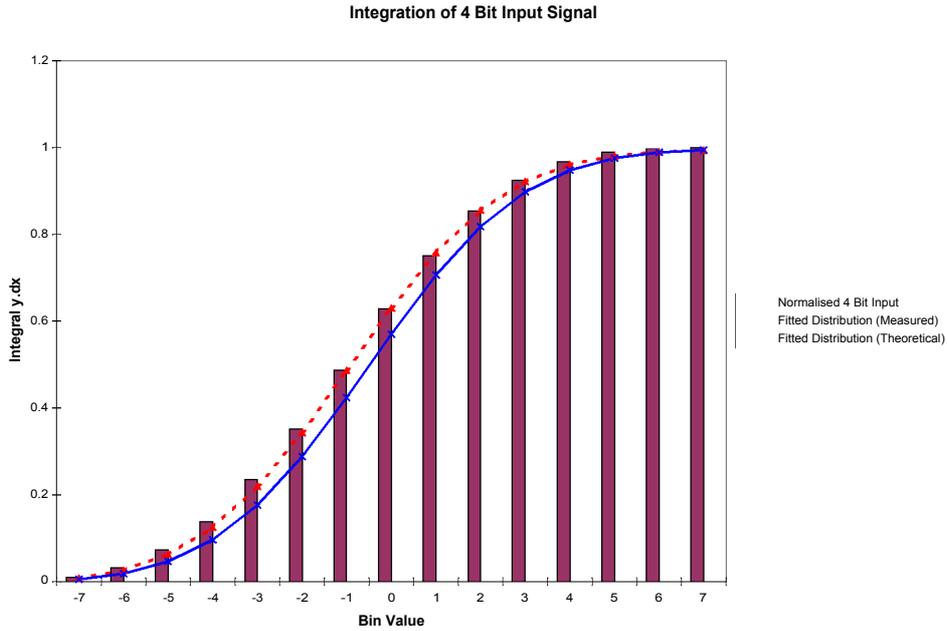


FIGURE 8.1.3. Cumulative distribution of 4 bit mixer input signal. $(D(\infty))_h = 1$, $(D(\infty))_{Th} = 0.9940$ and $(D(\infty))_M = 0.9935$

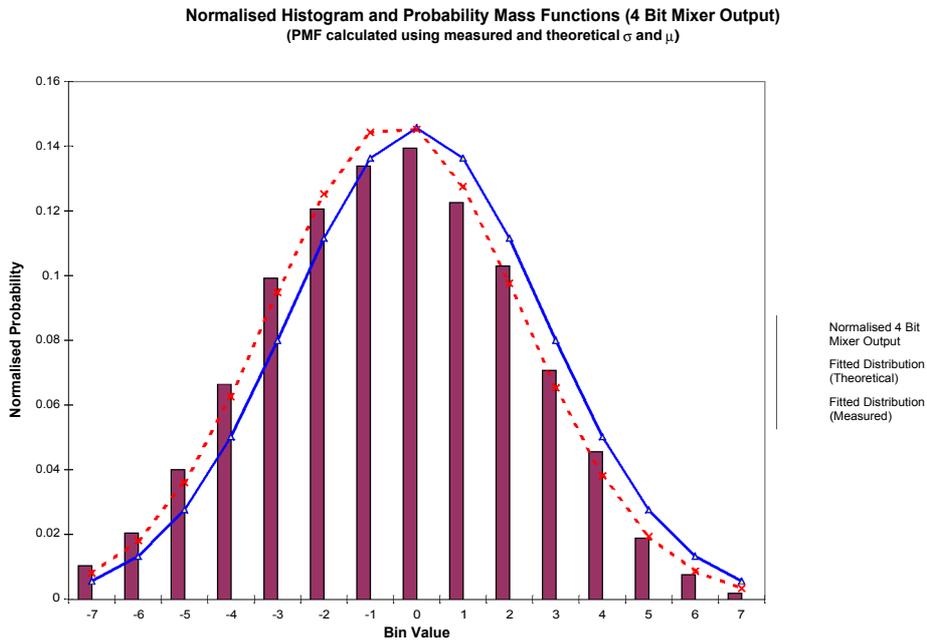


FIGURE 8.1.4. Normalised histogram and probability mass distributions of the 4 bit mixer output signal. $\sigma_{Th} = 2.740$, $\sigma_M = 2.709$, $\mu_{Th} = 0$ and $\mu_M = -0.457$

8.1.5. FIR Filters. The digital gain of the real filter for a white noise input is

$$G = \sqrt{\sum_{i=1}^M W_i^2} = 1282$$

$$\sim \sigma_{Th} = G \times \sigma_{Th} = 1282 \times 2.740 = 3513$$

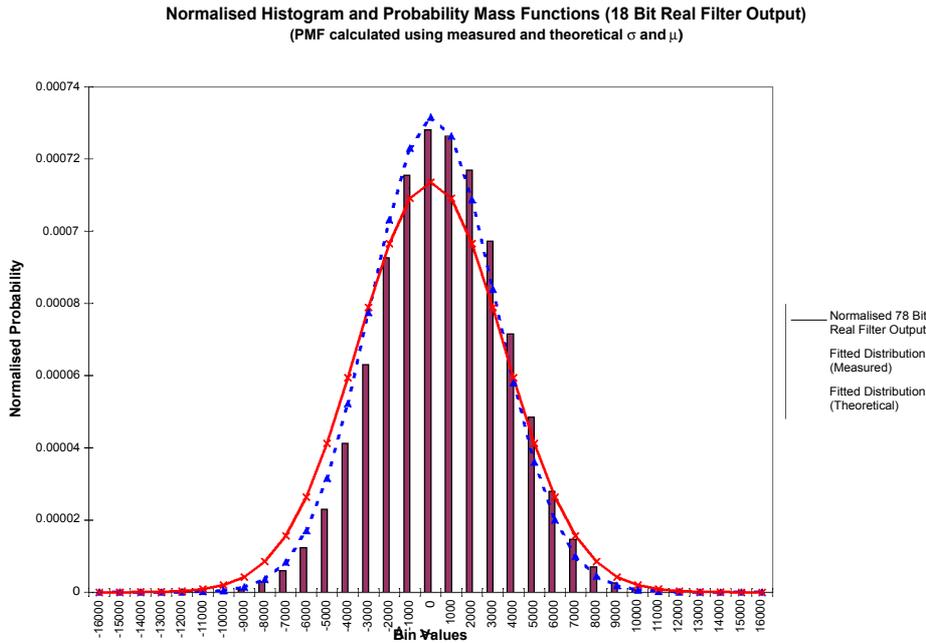


FIGURE 8.1.5. Normalised histogram and probability mass distributions for the 18 bit real filter output signal. $\sigma_{Th} = 3513$, $\sigma_M = 3030$, $\mu_{Th} = 0$ and $\mu_M = 121$

Figure 8.1.5 is the histogram and probability mass distributions of the 18 bit real filter output signal. The filter output is a narrowband signal and correlated due to the filter operation $y[n] = \sum_{i=0}^{M-1} W_i x[n-i]$, where $y[n]$ is the filter output. It is therefore not strictly a Gaussian process, and its probability distribution law will not be preserved through the linear system, resulting in larger discrepancies between theoretical and measured distributions.

Choosing $dx = 1000$ allows sufficient x axis resolution and a reduction in statistical noise by a factor 31.6. The limiting values for $4D(oc) \rightarrow 1$, with the slope of $4D(m)_M$ steeper than $4D(m)_{Th}$ since to $\sigma_{Th} > \sigma_M$.

The power spectrum of the real filter output, shown in figure 8.1.6, is obtained by the summation of 8×4096 point Discrete Time Fourier Transforms (DTFT). Smoothing is performed using an eight point moving average. Using a flat frequency spectrum for the noisy input signal, the filter response

$$H(\Omega) = \frac{Y(\Omega)}{X(\Omega)} = Y(\Omega)$$

confirms the theoretical filter response in figure 7.11.2.

The white noise digital gain of the imaginary filter is $G = 1218$. Therefore

$$\sigma_{Th} = G \times \sigma_{Th}^0 = 1218 \times 2.740 = 3337$$

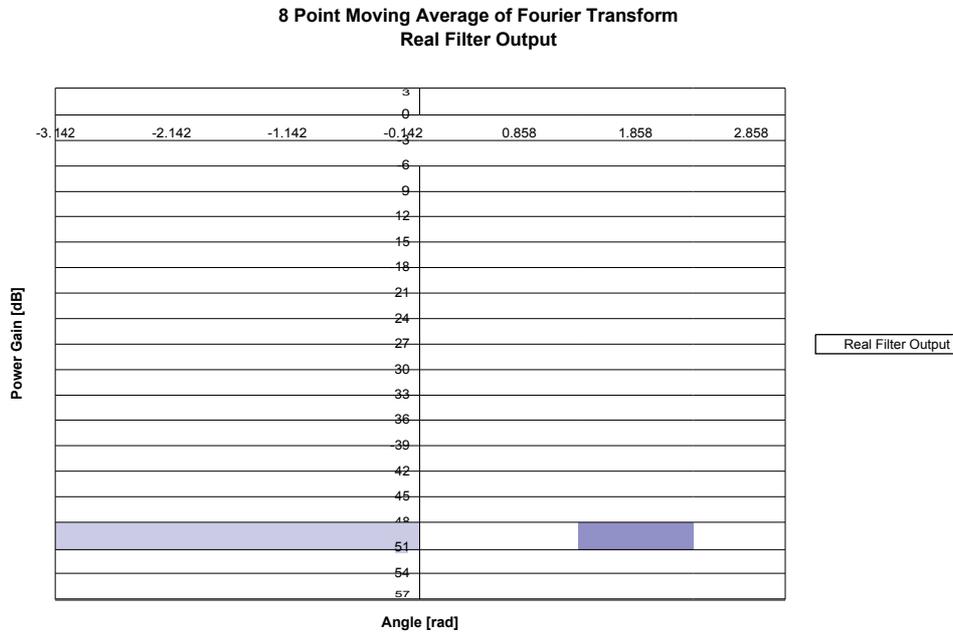


FIGURE 8.1.6. Smoothed 8×4096 point DTFT of real filter output

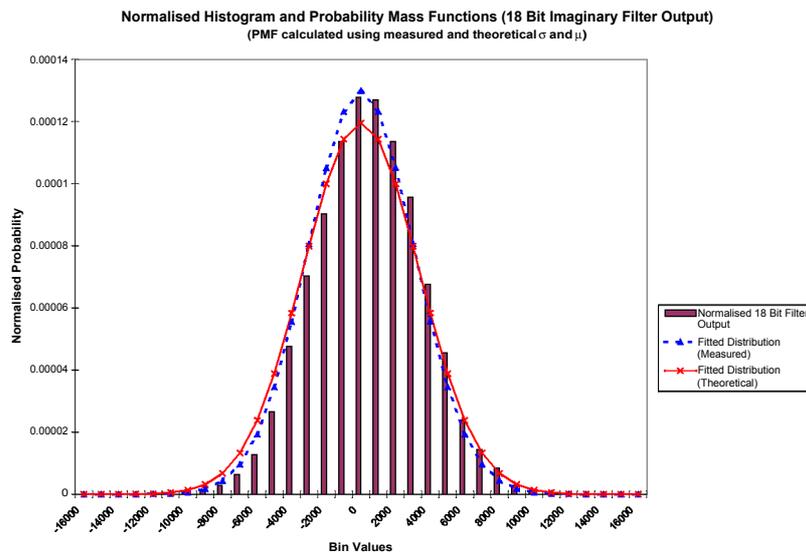


FIGURE 8.1.7. Normalised histogram and probability mass distributions of 18 bit imaginary filter output signal. $\sigma_{Th} = 3337$, $\sigma_M = 3072$, $\mu_{Th} = 0$ and $\mu_M = 1.77$

Figure 8.1.7 is the histogram and probability mass distributions of the imaginary filter output signal. Figure 8.1.8 is the smoothed 8×4096 point DTFT of the imaginary filter output $Y(\Omega)$, which is used to confirm the theoretical response of the imaginary filter in figure 7.11.2.

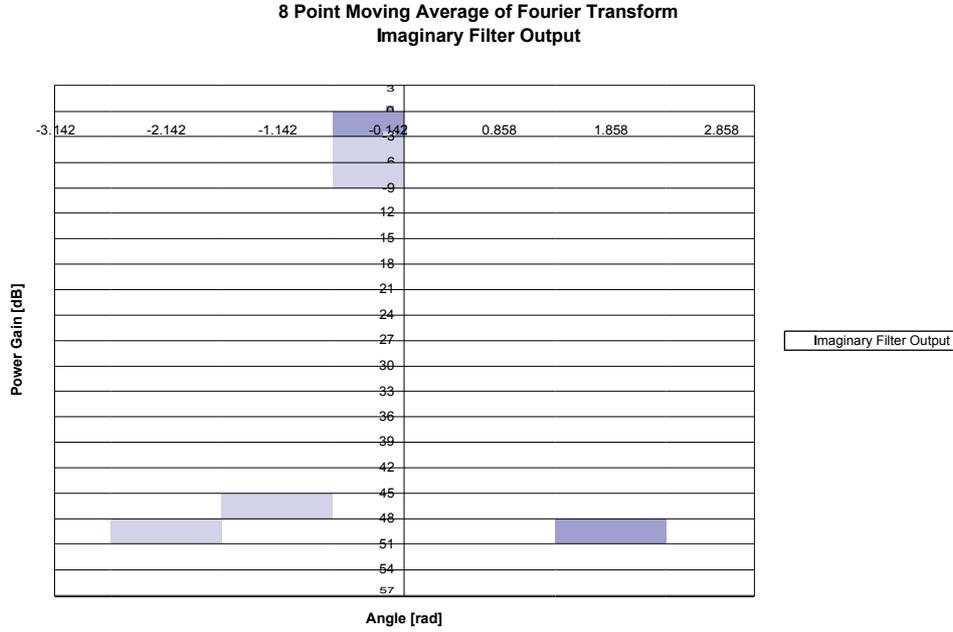


FIGURE 8.1.8. Smoothed 8×4096 point DTFT of imaginary filter output

8.1.6. Sideband Separation. The sideband separation (adder and subtractor) has a digital gain $G = \sqrt{2}$. Therefore

$$\sigma_{Th} = G \times \sigma_{avg}^0 = \sqrt{2} \times 3425 = 4844$$

where σ_{avg}^0 is the linear average of σ_{Th} for the real and imaginary filters. The linear average is used since there is no loss of total signal power ie. where two signals undergo total construction in the adder, they will undergo total destruction in the subtractor. Hence

$$\begin{aligned} \text{Power}_{in} &= \text{Power}_{out} \\ \Rightarrow \sigma_{Total_{in}} &= \sigma_{Total_{out}} \end{aligned}$$

Figure 8.1.9 is the histogram and probability mass distributions of the 18 bit adder output. The upper sideband separation is shown in the smoothed 8×4096 point DTFT of the adder output in figure 8.1.10. There is a -8 dB noise rejection between adjacent channels as a result of the 3 dB increase in peak power in the adder and -5 dB DC block in the real filter.

Figure 8.1.11 is the histogram and probability mass distributions of the 18 bit subtractor output signal. The 5% difference in σ_M between the subtractor and adder stage, and approximate anti-symmetry in values for μ_M , is the result of a small DC offset in the analogue real input signal. The lower sideband separation is shown in figure 8.1.12.

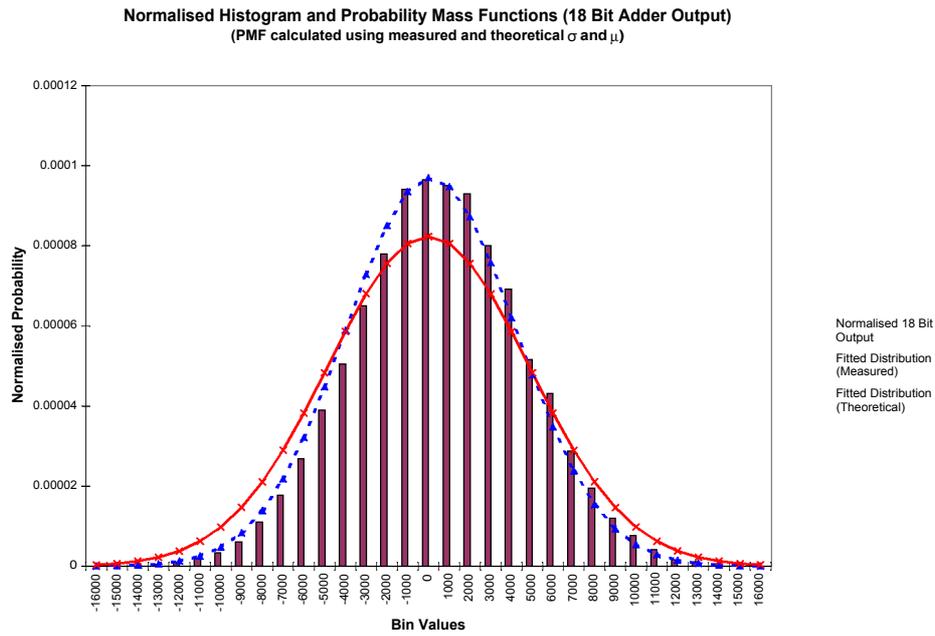


FIGURE 8.1.9. Normalised histogram and probability mass distributions of 18 bit adder output signal. $\sigma_{Th} = 4844$, $\sigma_M = 4111$, $\mu_{Th} = 0$ and $\mu_M = 110$

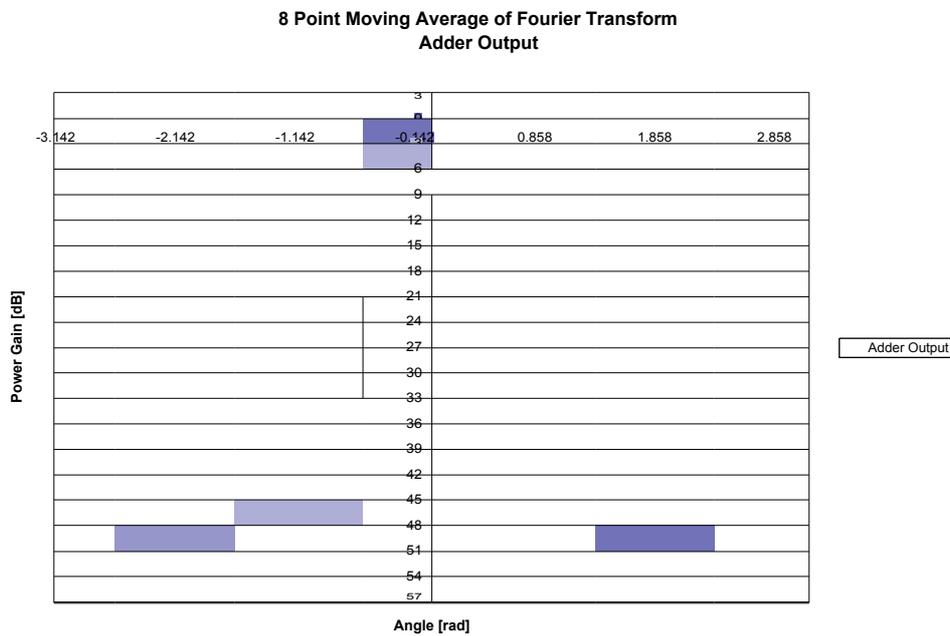


FIGURE 8.1.10. Smoothed 8×4096 point DTFT of adder output

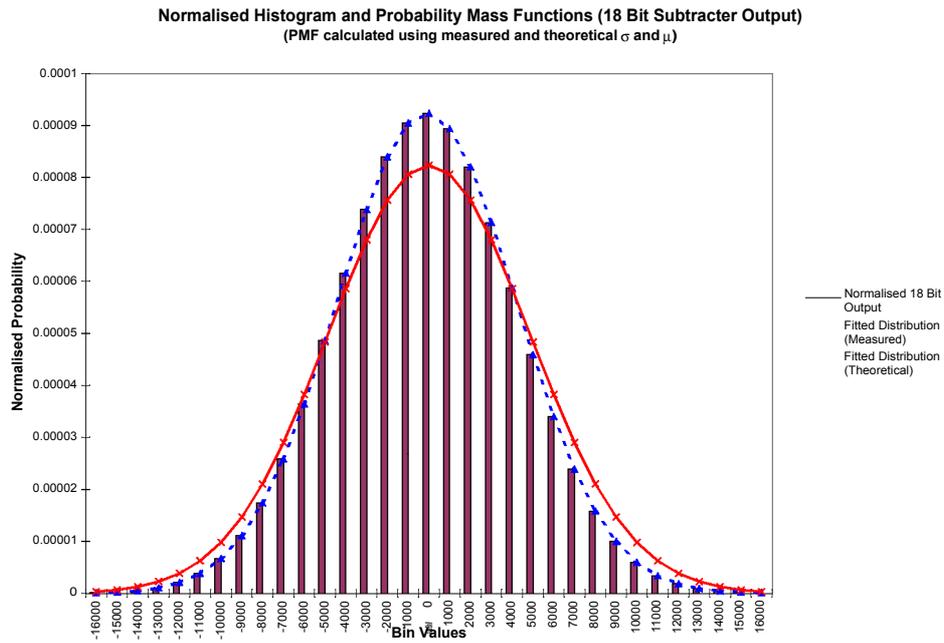


FIGURE 8.1.11. Normalised histogram and probability mass distributions of 18 bit subtractor output signal. $\sigma_{Th} = 4844$, $\sigma_M = 4319$, $\mu_{Th} = 0$ and $\mu_M = -109$

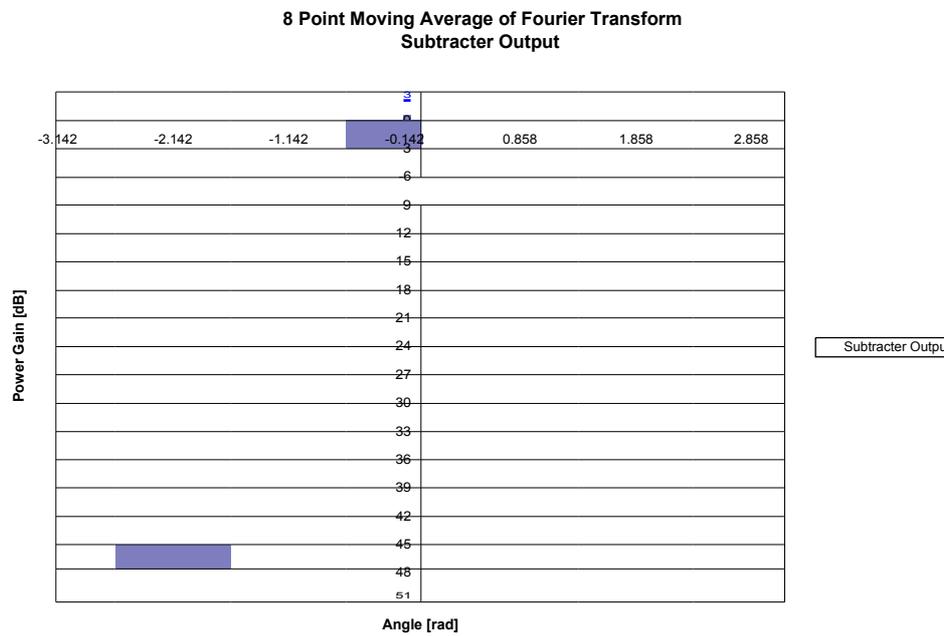


FIGURE 8.1.12. Smoothed 8×4096 point DTFT of subtractor output

8.1.7. Detection. Bits [9:14] of the upper and lower sidebands are used as input to the squarer, resulting in $G = \frac{1}{2^9}$

$$\Rightarrow \sigma_{Th} = G \times \sigma_{Th}^I = \frac{1}{2^9} \times 4111.08 = 8.03$$

where σ_{Th} is the theoretical standard deviation of the input signal to the squarer. σ_M^I of the adder output is used as opposed to σ_{Th}^I in order to adjust for the cascaded accumulation of offset errors.

The squared signal distribution is determined by making the substitution $y = x^2$ in the distribution

$$P(x) = \sqrt{\frac{1}{2\pi\sigma_x}} \exp^{-\frac{x^2}{2\sigma_x^2}}$$

So

$$P^I(y) = \frac{2P(x)}{\frac{dy}{dx}}$$

$$\Rightarrow P^I(y) = \sqrt{\frac{1}{2\pi y \sigma_x}} \exp^{-\frac{y}{2\sigma_x^2}}$$

which describes a Chi squared distribution with $\sigma_y = \sigma_{Th} = \sqrt{2\sigma_x^2}$ [132].

The squarer implements a linear scaling of 1 bit due to the unsigned output. Therefore $G = \frac{1}{2^5} \times 2$, and so

$$\sigma_{Th} = \sqrt{2} \times (G \times \sigma_{Th}^I)^2$$

Figure 8.1.13 is the histogram and probability mass distributions of the squarer output where $\sigma_{Th}^I = \sigma_M^I$ and $\mu_{Th} = \mu_M$.

8.1.8. Post-Detection Integration. Writing the accumulated output as $z = \sum_{i=1}^N x_i$, then by the Central Limit Theorem, for $N \gg 1$

$$P(z) = \sqrt{\frac{1}{2\pi\sigma_{Th(z)}}} \exp^{-\frac{(z - \mu_{Th(z)})^2}{2\sigma_{Th(z)}^2}}$$

where

$$\mu_{Th(z)} = N \times \sigma_{Th(x)}^2$$

and

$$\sigma_{Th(z)} = \sqrt{2N} \times \alpha \sigma_{Th(x)}^2$$

The constant α is determined by the number of channels in the total bandwidth. $\sigma_{Th(x)} = G \times \sigma_{Th(x)}^I$ where $\sigma_{Th(x)}^I = 8.029$ is the standard deviation of the squarer input signal. Accumulation for $N = 2^{14} - 9$, with $\alpha = 8$, gives

$$\sigma_{Th(z)} = \sqrt{2N} \times \alpha \sigma_{Th(x)}^2 = 206 \times 10$$

and

$$\mu_{Th(z)} = 660 \times 10^2$$

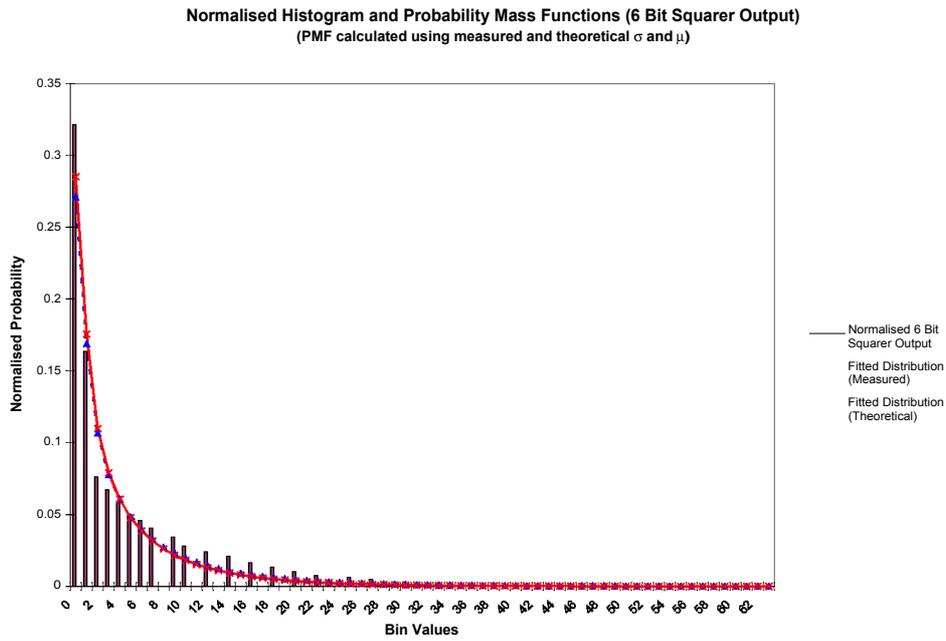


FIGURE 8.1.13. Normalised histogram and probability mass distributions of the 6 bit squarer output signal. $\sigma_{Th} = 5.70$, $\sigma_M = 6.13$, $\mu_{Th} = 4.03$ and $\mu_M = 4.24$

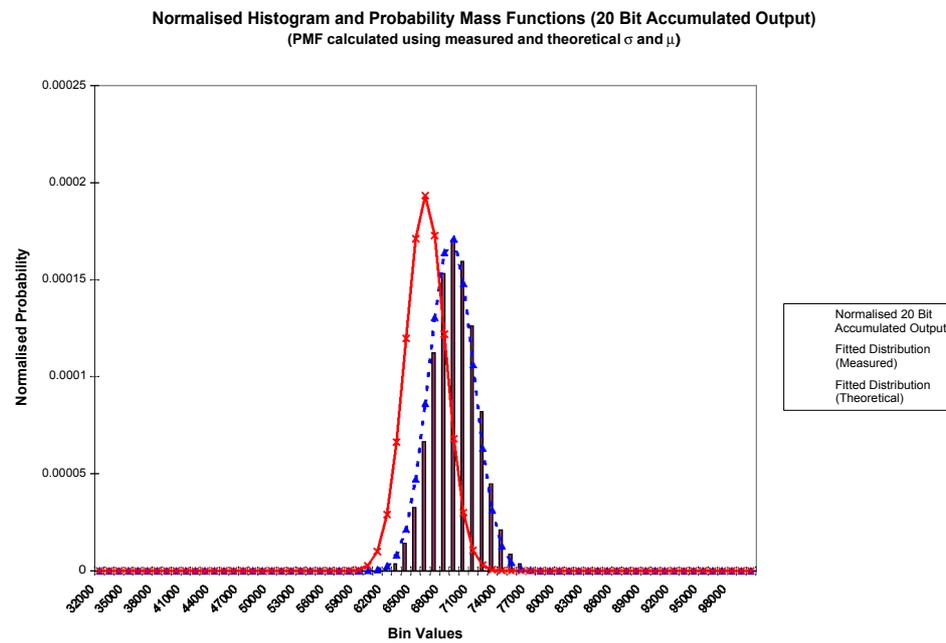


FIGURE 8.1.14. Normalised histogram and probability mass distributions of the accumulated square output signal. $\sigma_{Th(z)} = 206 \times 10$, $\sigma_{M(z)} = 232 \times 10$, $\mu_{Th(z)} = 660 \times 10$ and $\mu_{M(z)} = 687 \times 10^2$

Figure 8.1.14 is the histogram and probability mass distributions of the accumulated squarer output signal. The 10% difference between $a_{Th(z)}$ and $a_{m(z)}$ is a result of non-Gaussian processes in the quantisation error signal, a narrowband digital filter response and uneven frequency response of the bandpass filter used at IF. Values for $p_{m(z)}$ do not extend into bit 17 of the accumulator, allowing for a 3 bit overflow.

8.2. Bandpass Separation

Period integrations were performed using an injected CW signal from a sweep oscillator (HP 8690B) to measure the bandpass and sideband separation of the processor. Triggering the sweep oscillator at the beginning of each period, the CW signal is swept through the IF signal bandwidth at a rate $\frac{B}{t}$, where B is the signal bandwidth and t the pulse period. The resultant integration produces the frequency response of the associated sideband channel as each bin is linearly mapped with frequency

$$v = \text{Bin} \frac{B}{\text{Bin}_{\text{Total}}} \times \text{Bin} - \text{DC}_{,,}$$

where B is the total bandwidth swept through, $\text{Bin}_{\text{Total}}$ is the total number of bins, Bin is the bin number to be mapped and $\text{DC}_{,,}$ is the offset frequency.

Using an integration period of 10 seconds, with 200 bins across a pulse period of 100 ms, the upper and lower sidebands were measured for a 0 radian complex mixer, shown in figure 8.2.1. The bin offset is due to an incorrect integration start time.

Setting $k = -3, -2, -1, 0, 1, 2, 3$ for the $\frac{k\pi}{4}$ complex mixer (see section 7.10.1), repetition of the test allowed for the measurement of upper and lower sidebands of all narrowband channels. The composite passband is shown in figure 8.2.2, artificially rotated in order to shift 0 radians to bin 100. The narrowband channels do not spread across the entire composite bandpass since a π radian digital mixer was not implemented, due to the use of the two outer channels as guard bands. There is sufficient separation between the upper and lower sidebands, approaching the theoretical -8 dB limit for the $\frac{3\pi}{4}$ radian mixer. The inner sidebands fall short of this theoretical limit by approximately 2 dB. This could be due to measurement errors and the effect of the guard bands.

Separation between narrowband channels is approximately 3 dB. The inconsistency in separation between channels is a result of an uneven frequency response in the IF bandpass filter. Furthermore, the sweep oscillator was in need of service, with irregular amplitude modulations of the CW signal occurring during measurement.

8.3. Field Trials

The prototype processor with a 2×2 MHz bandwidth, consisting of a starter module, analogue frontend and digital backend, was installed at HartRAO for a period of two weeks for final adjustments and observations.

De-noising of the received pulse integration is performed using wavelet analysis, where thresholds are set using a heuristic variant of Stein's Unbiased Risk principle [59]. Each sideband is

normalised before a full passband is compiled. The parameter T corresponds to the number of pulse integrations.

8.3.1. 'First Light' Observations. Figures 8.3.1 to 8.3.11 are observations of a sub-population of pulsars regularly monitored at HartRAO. Figure 8.3.1 shows the upper and lower

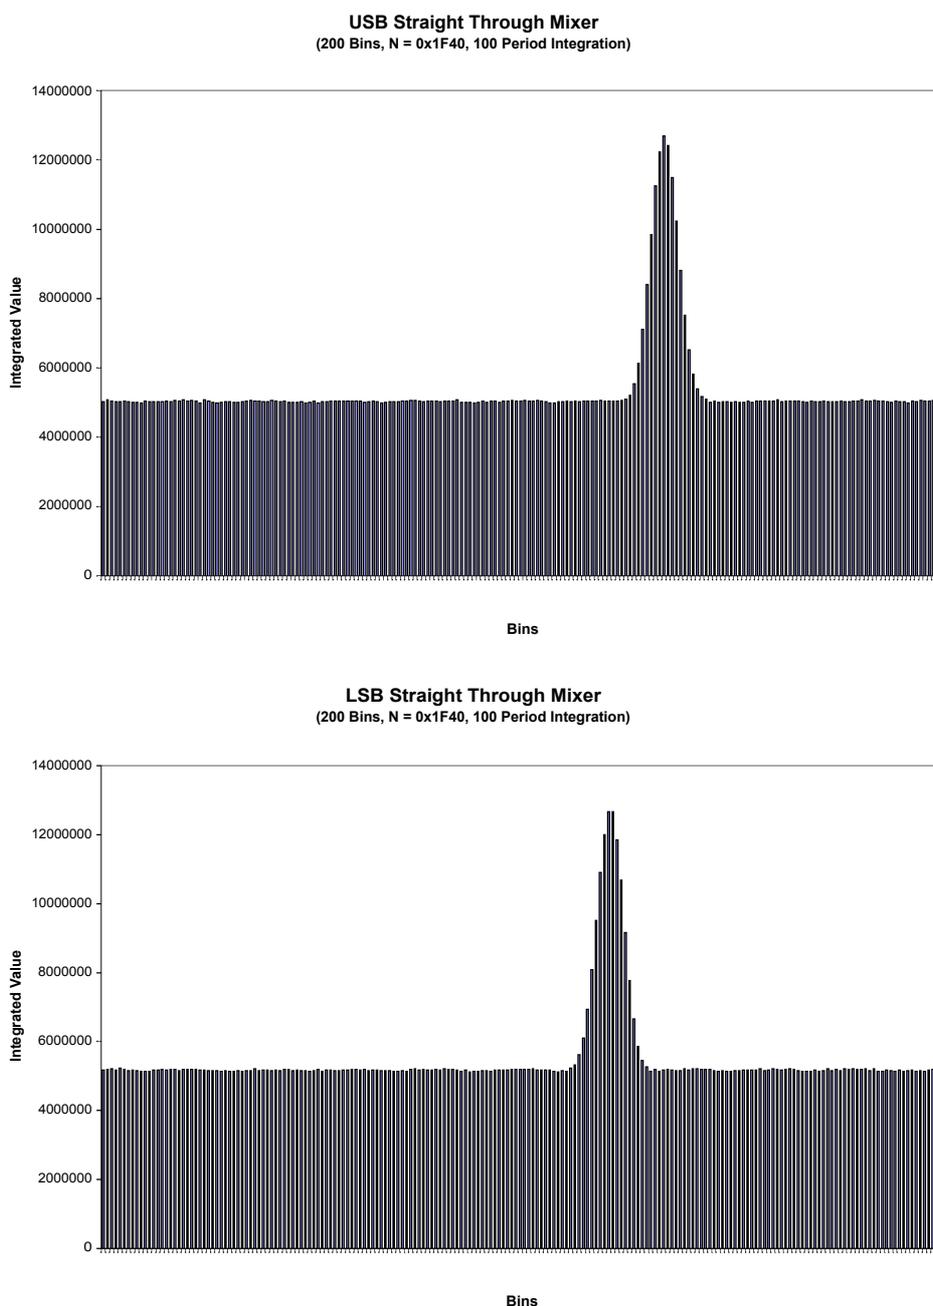


FIGURE 8.2.1. Frequency response of upper and lower sidebands, obtained using the 0 degree mixer and swept oscillator. Integration period = 10 seconds, oscillator sweep period = 0.1 second, 200 bins.

sidebands before integration into a single profile. The rest of the figures show integrated profiles only.

Since the pulsars PSR B0740-28 and PSR B1054-62 are at the limit of the test instrument’s sensitivity, their integrated pulse profiles, shown in Figures 8.3.2 and 8.3.4, are severely affected by stochastic fluctuations.

8.3.2. Calibration Measurements. Multiple observations of PSR 0833-45 and PSR 1641-45 were made for noise performance and individual channel bandwidth tests.

8.3.2.1. *Noise Performance.* Figure 8.3.12 shows integrated pulse profiles of PSR 0833-45 for different integration periods. Increases in the SNR are within 15% of those predicted by the radiometer equation.

8.3.2.2. *Individual Channel Bandwidths.* Figure 8.3.13 shows individual channel bandwidth tests, based on observations of PSR 1641-45, using different complex mixers. The integrated pulse profile for the $\frac{3\pi}{4}$ mixer shows a decrease in SNR, as well as a double peak. This is due to an erroneous update of the sampling frequency during the observation. Time constraints did not allow for this observation to be repeated.

8.4. Conclusion

Statistical analysis of each stage of the pulsar processor, using an artificial input noise source, seemed to agree with simulator predictions. This has allowed for accurate predictions of bit growth in the signal processing chain, thereby ensuring that the system will not saturate under normal operation.

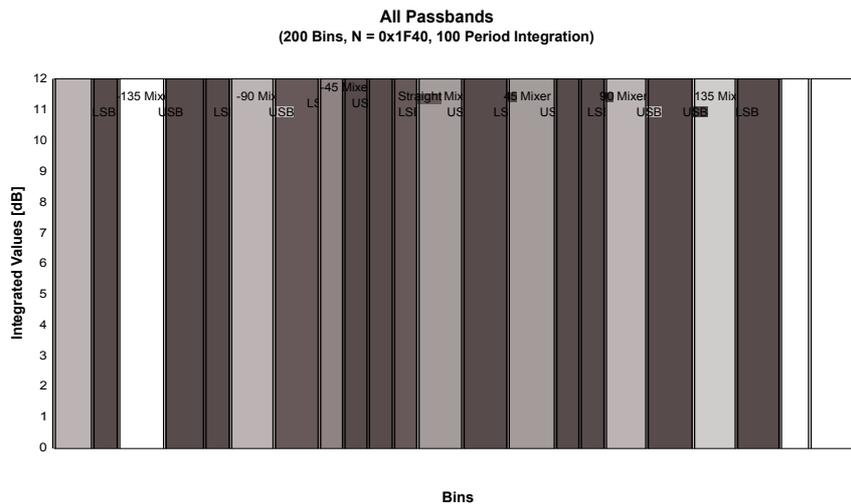


FIGURE 8.2.2. Composite bandpass of pulsar processor. 0 radians has been artificially shifted to Bin 100.

Complex digital mixers were shown to efficiently shift the receiver passband by units of $\frac{\pi}{4}$ with little or no adverse effects.

Field trials, using a prototype pulsar processor with $2 \times 2\text{MHz}$ channels, took place at HartRAO over a period of two weeks. While a number of pulsars were at the sensitivity limit of

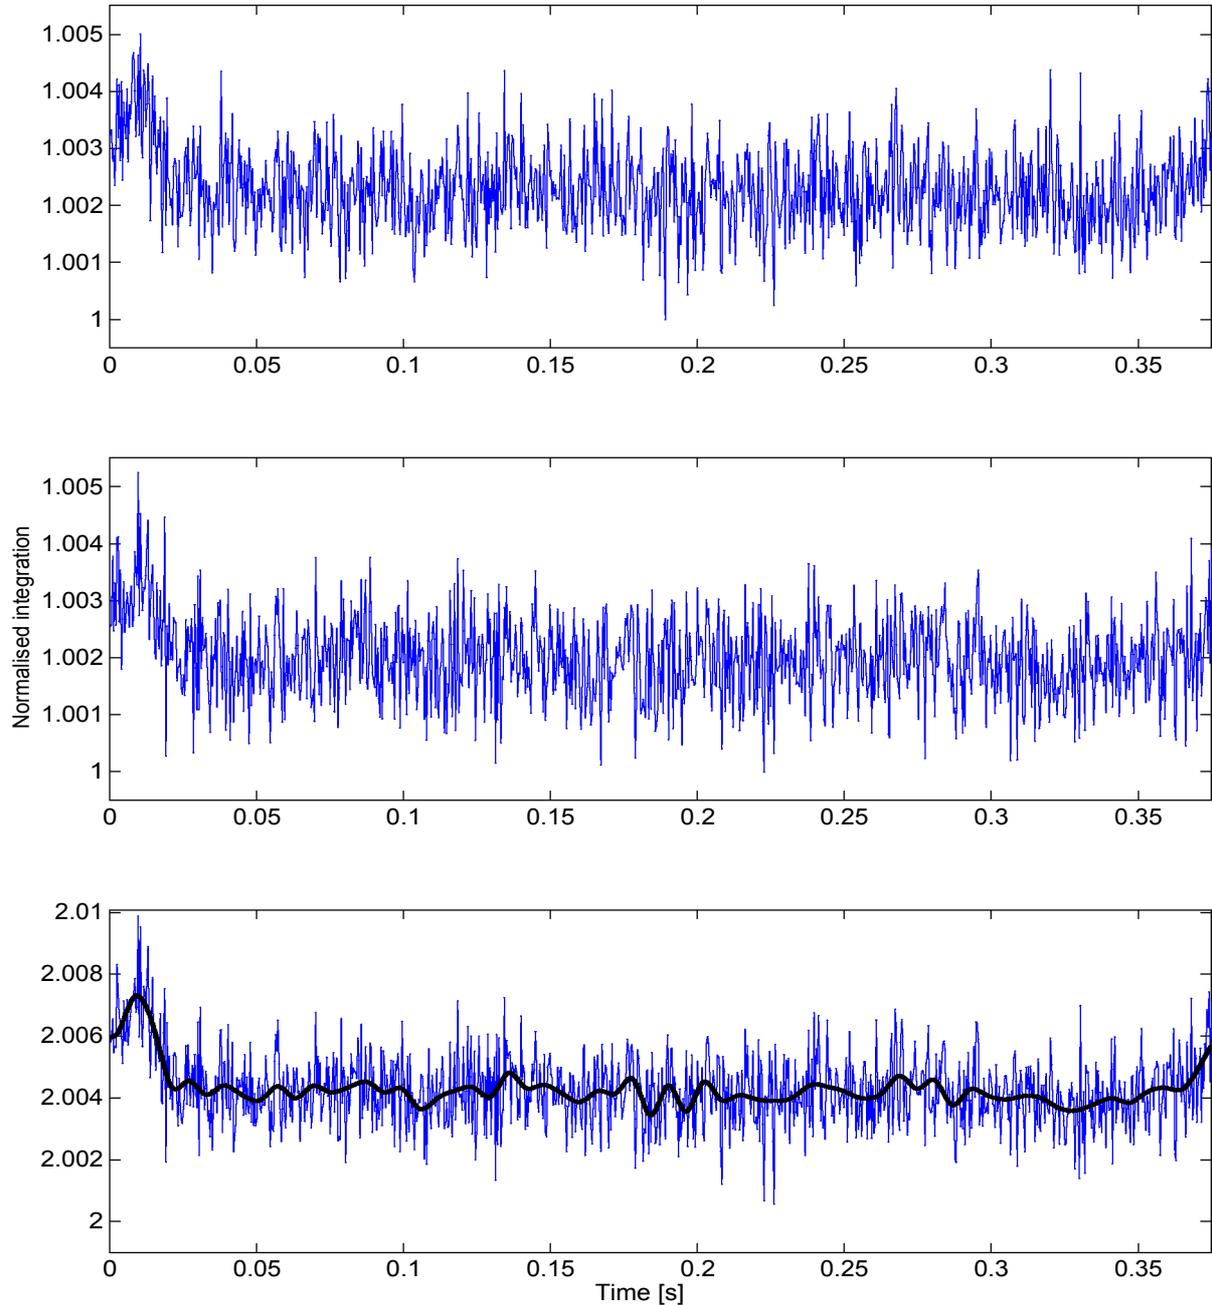


FIGURE 8.3.1. Compiled profile of PSR B0736-40 integration with $\mathbf{T} = 3000$, Bins = 993 and SNR = 0.471 dB. Top figure is the upper sideband, the middle figure is the lower sideband, and the bottom figure is an integrated pulse profile. Thick line is the de-noised integrated profile.

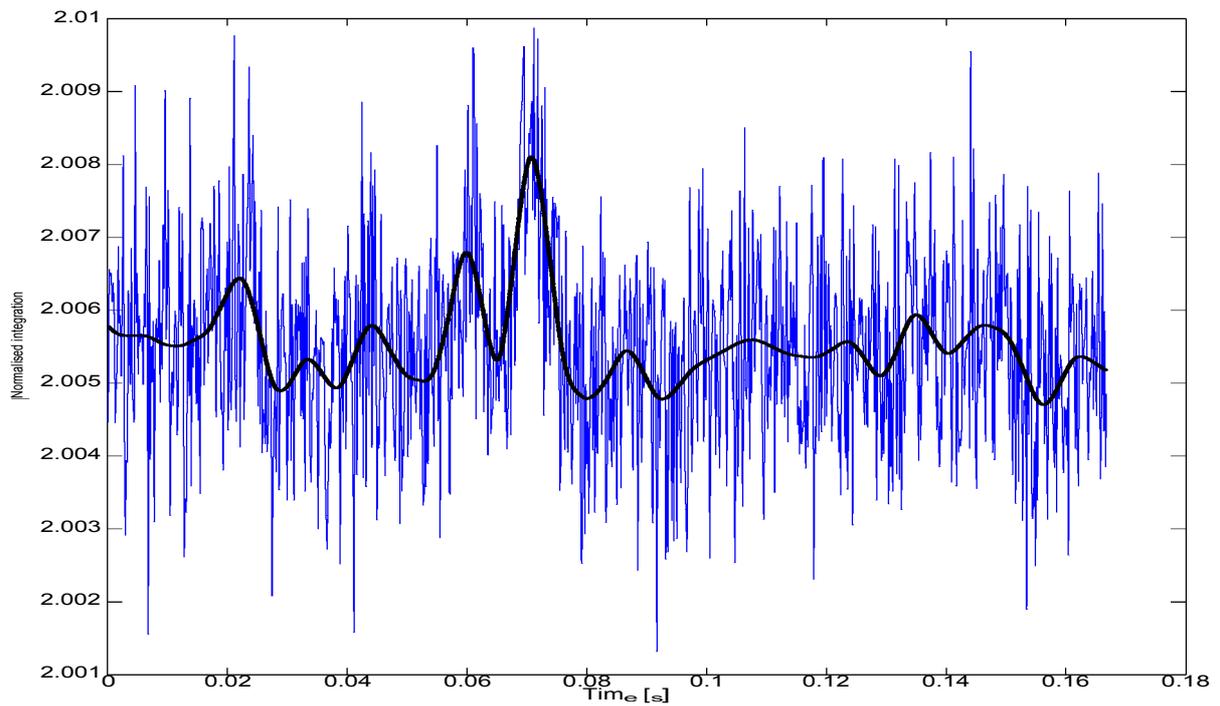


FIGURE 8.3.2. Integrated profile of PSR B0740-28. $\tau = 5000$, Bins = 1000 and SNR = 0.321 dB. Thick line is the de-noised integrated profile.

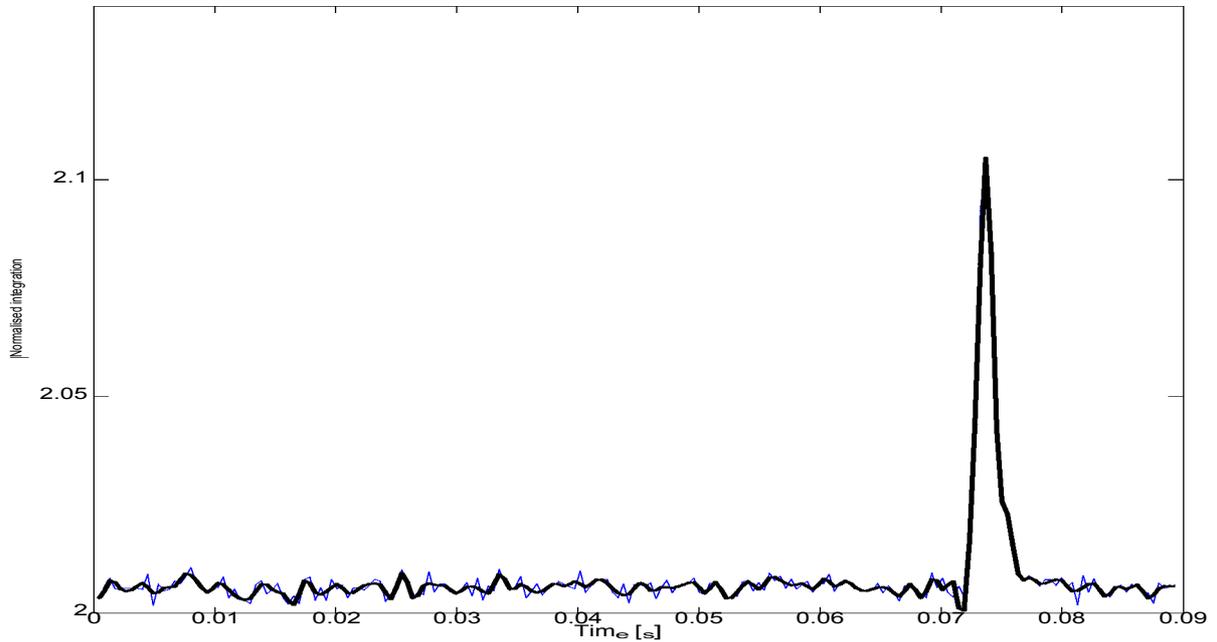


FIGURE 8.3.3. Integrated profile of PSR B0833-45. $\tau = 1000$, Bins = 200 and SNR = 1.738 dB. Thick line is the de-noised integrated profile.

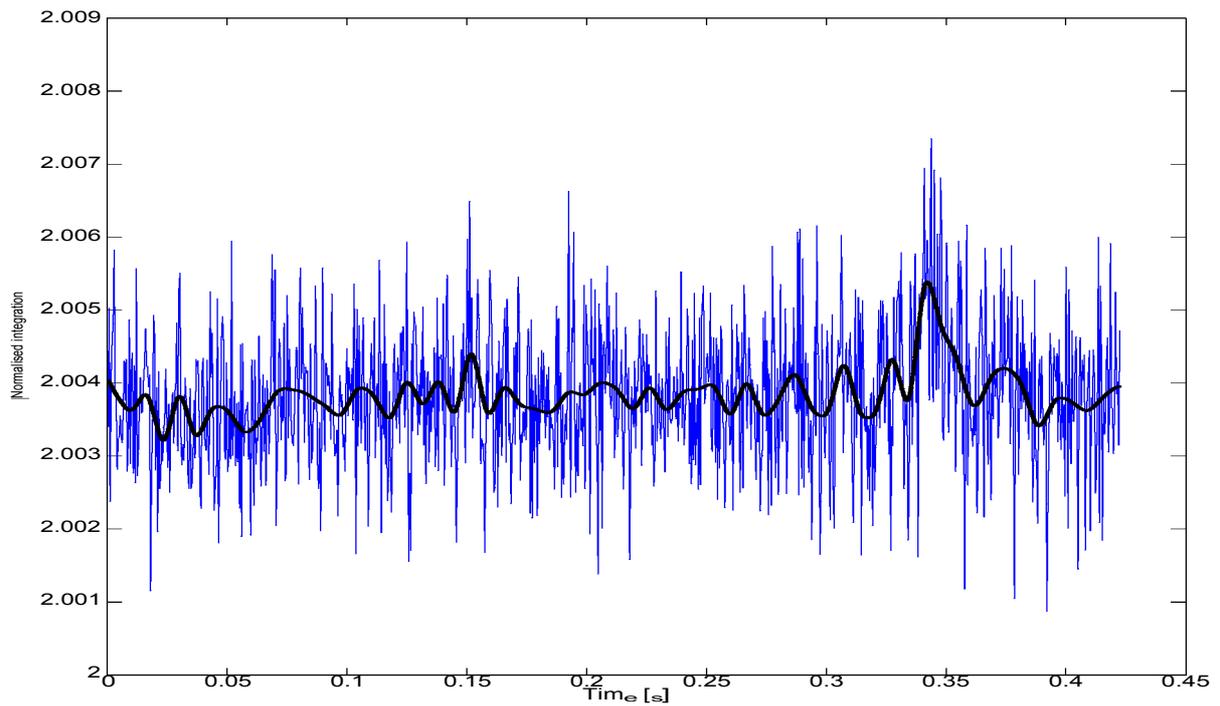


FIGURE 8.3.4. Integrated profile of PSR B1054-62. $\tau = 3500$, Bins = 1000 and SNR = 0.234 dB. Thick line is the de-noised integrated profile.

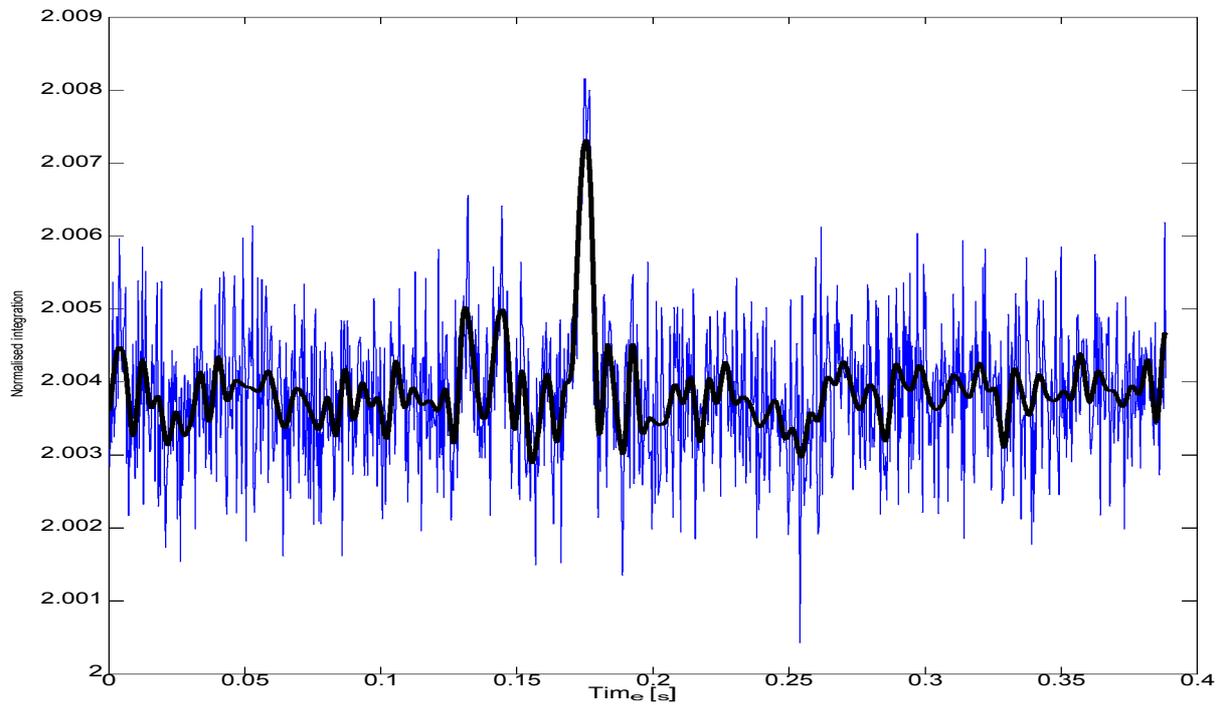


FIGURE 8.3.5. Integrated profile of PSR B1240-64. $\tau = 4000$, Bins = 1000 and SNR = 0.622 dB. Thick line is the de-noised integrated profile.

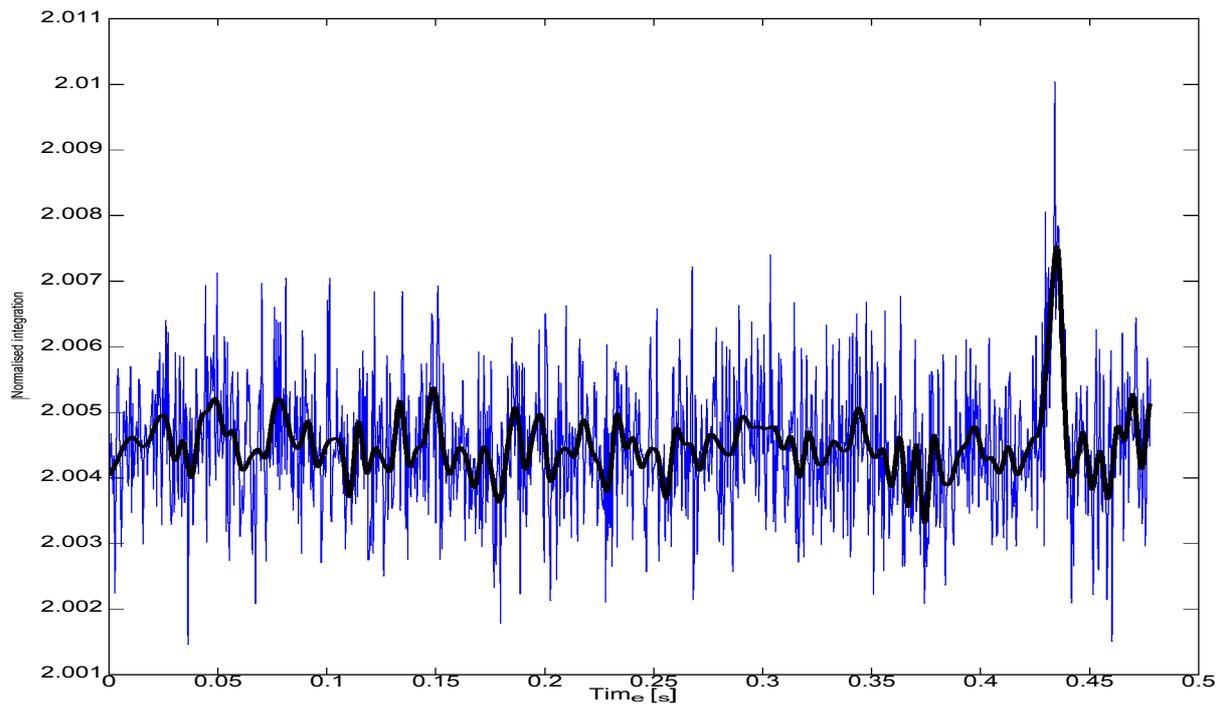


FIGURE 8.3.6. Integrated profile of PSR B1323-58. $\tau = 3000$, Bins = 1000 and SNR = 0.520 dB. Thick line is the de-noised integrated profile.

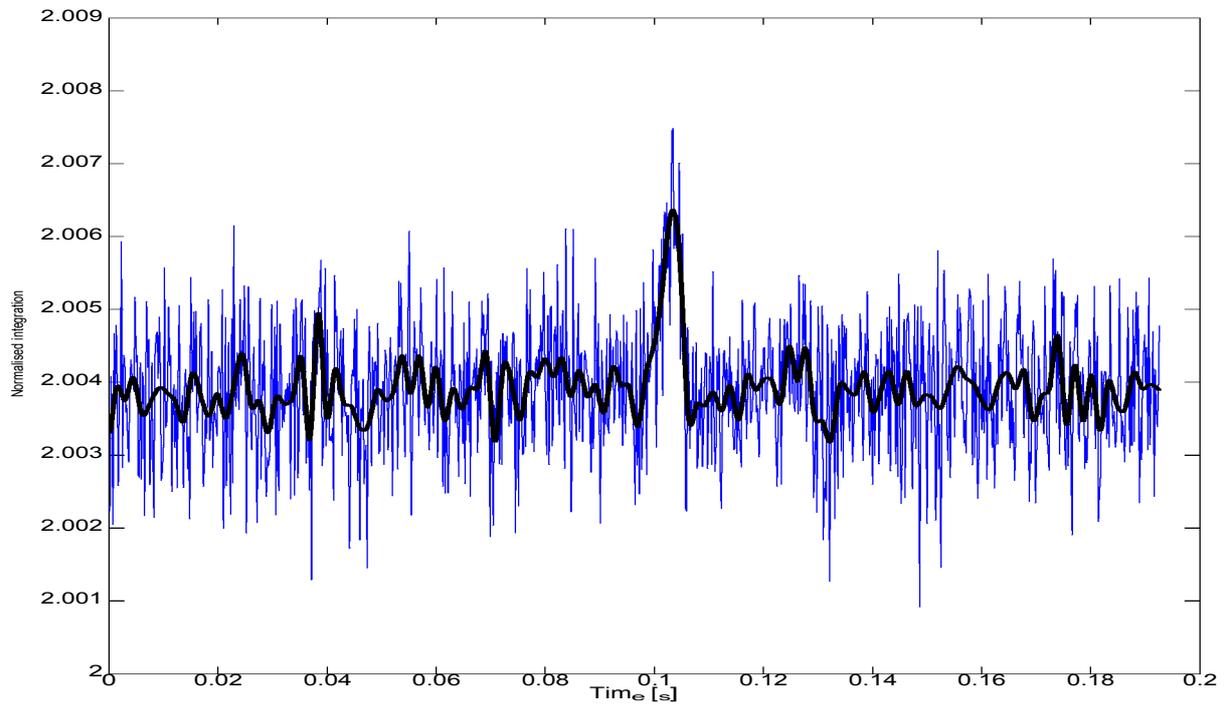


FIGURE 8.3.7. Integrated profile of PSR B1557-50. $\tau = 8000$, Bins = 1000 and SNR = 0.486 dB. Thick line is the de-noised integrated profile.

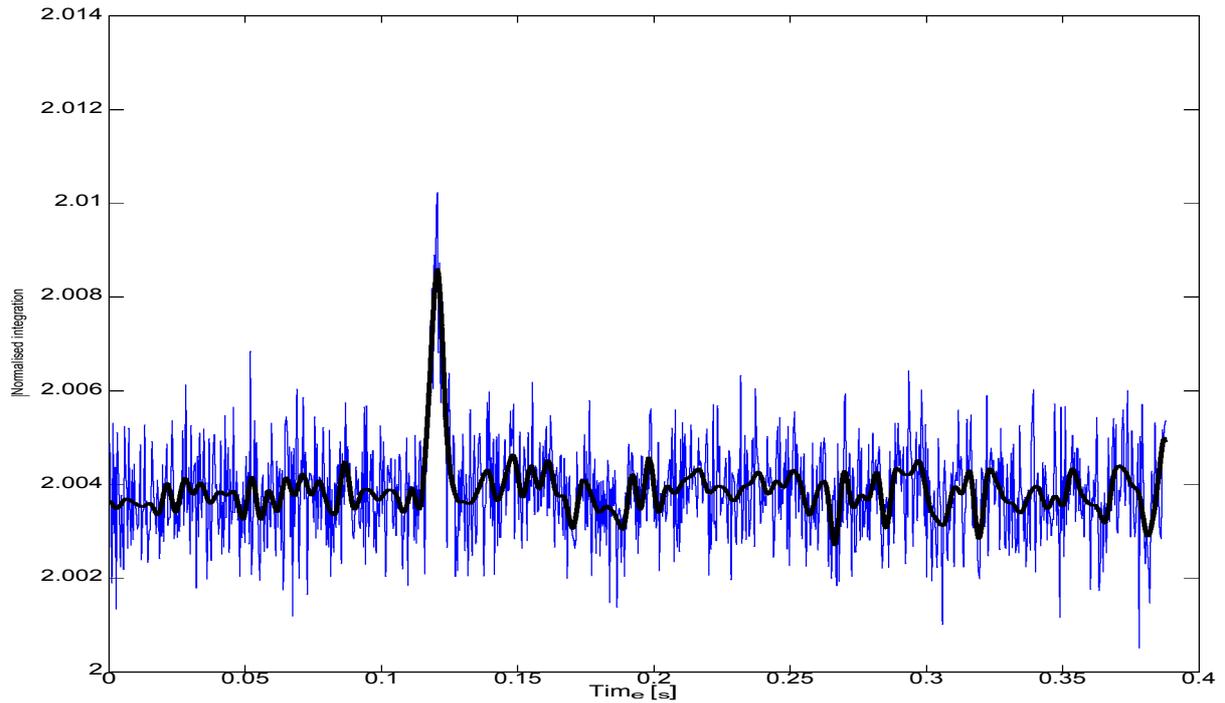


FIGURE 8.3.8. Integrated profile of PSR B1642-03. $\tau = 3600$, Bins = 1000 and SNR = 0.722 dB. Thick line is the de-noised integrated profile.

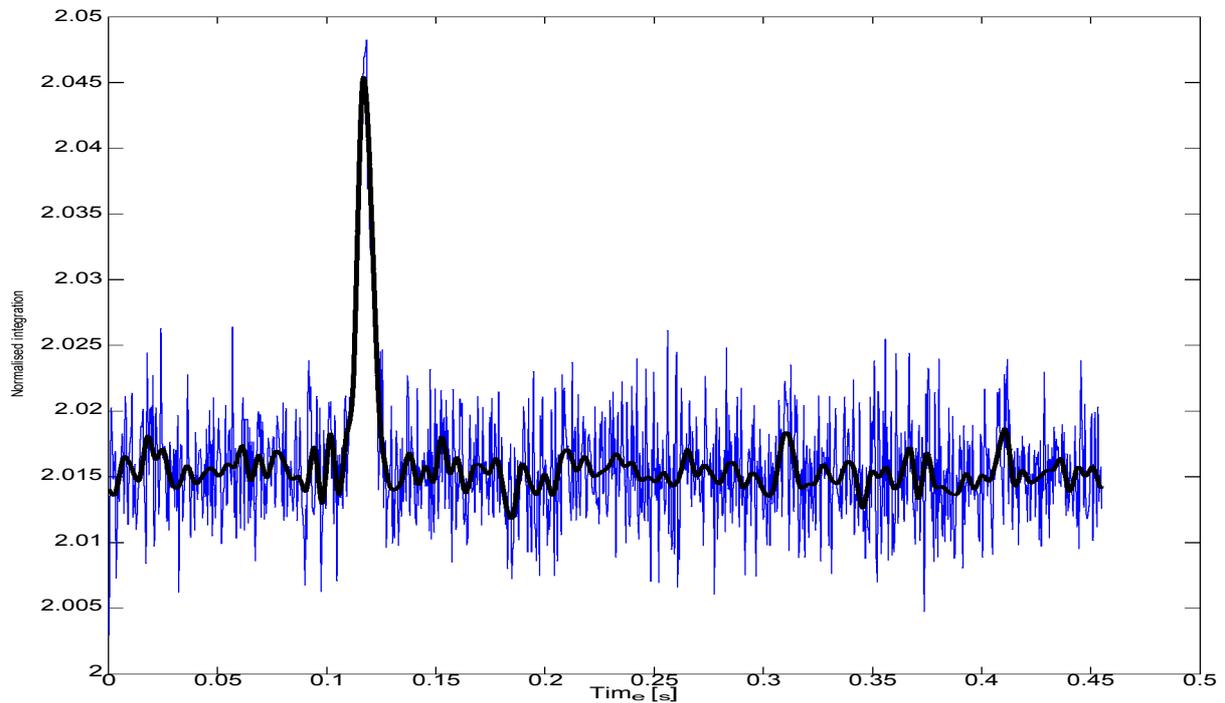


FIGURE 8.3.9. Integrated profile of PSR B1641-45. $\tau = 200$, Bins = 1000 and SNR = 0.925 dB. Thick line is the de-noised integrated profile.

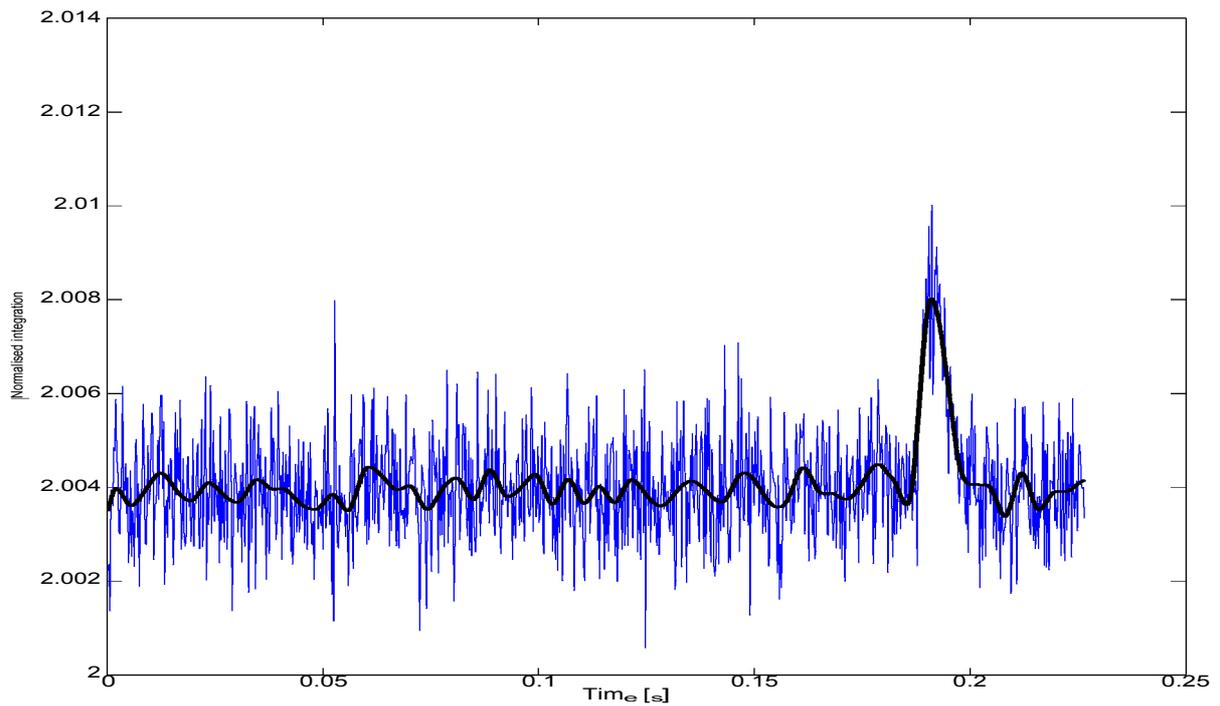


FIGURE 8.3.10. Integrated profile of PSR B1929+10. $r = 5000$, Bins = 1000 and SNR = 0.598 dB. Thick line is the de-noised integrated profile.

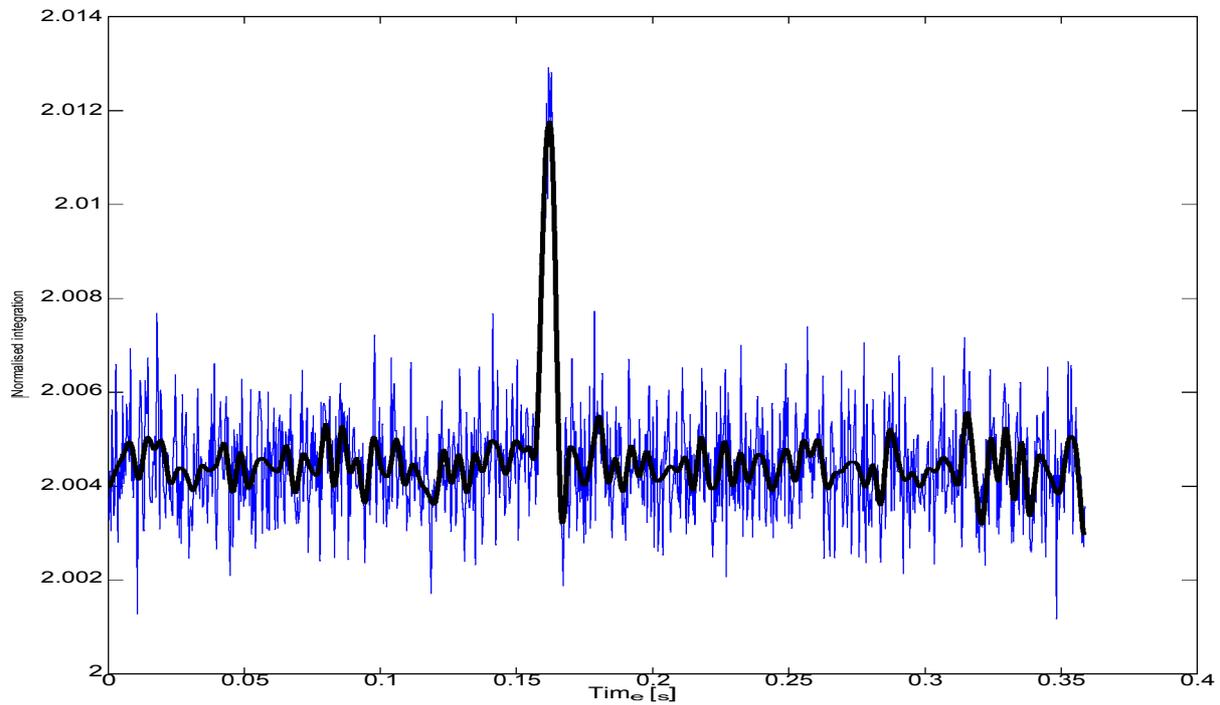


FIGURE 8.3.11. Integrated profile of PSR B1939+16. $r = 3200$, Bins = 1000 and SNR = 0.891 dB. Thick line is the de-noised integrated profile.

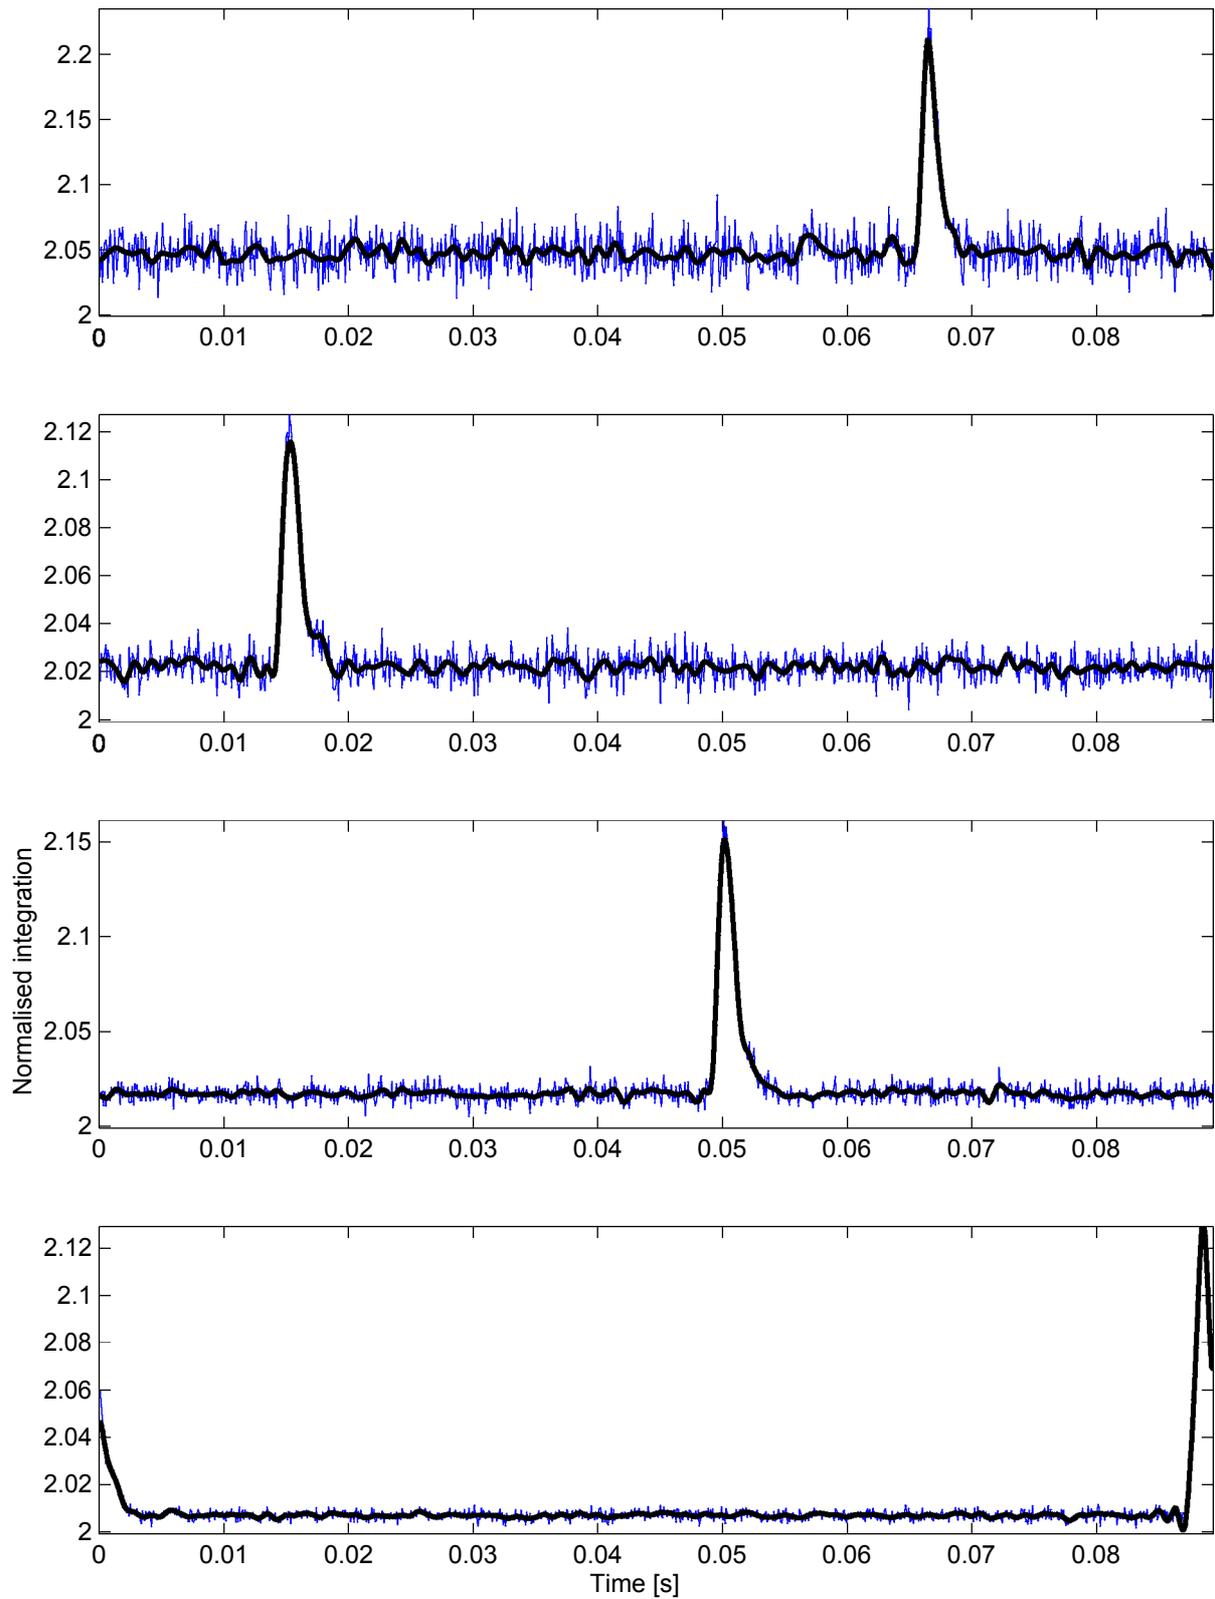


FIGURE 8.3.12. Noise performance test, based on integrated pulse profiles for PSR 0833-45. From top to bottom, $\tau = 100, 500, 1000$ and 5000 , Bins = 1000 and SNR = 1.139 dB, 1.266 dB, 1.560 dB and 1.802 dB respectively. The increases in SNR are within 15% of those predicted by the radiometer equation.

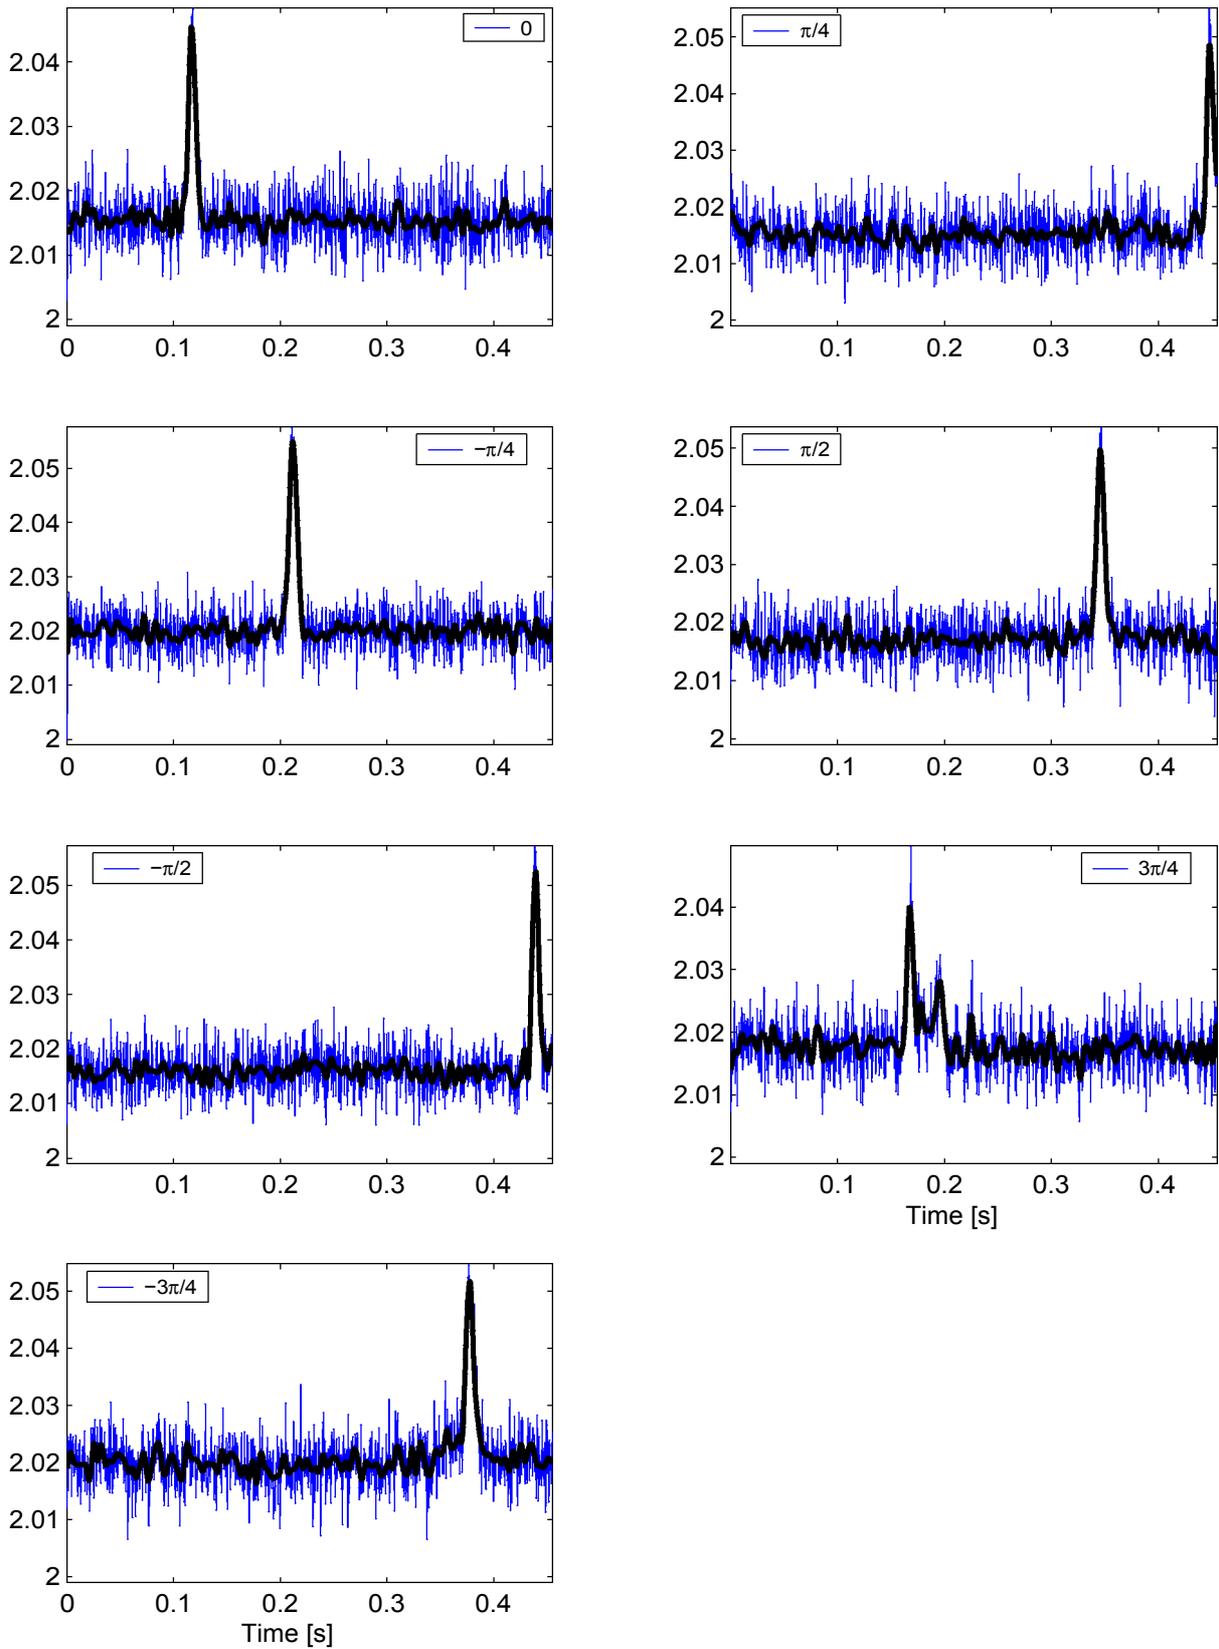


FIGURE 8.3.13. Individual channel bandwidth tests, based on multiple observations of PSR 1641-45 using different complex mixers. $\tau = 200$, $B_{\text{sig}} = 1000$. Mixers used are, from top to bottom, left to right: 0 , $\frac{\pi}{4}$, $-\frac{\pi}{4}$, $\frac{\pi}{2}$, $-\frac{\pi}{2}$, $\frac{3\pi}{4}$, $-\frac{3\pi}{4}$.

the prototype, namely PSRs B0740-28 and B1054-62, measurements of pulsars in the HartRAO catalogue confirmed and proved the correct operation of the new pulsar processor.

CHAPTER 9

Conclusion and Recommendations

This project aimed to consolidate and extend what we know about the dynamical behaviour of isolated pulsars. The long term pulsar monitoring program at HartRAO has produced timing residual data that has resulted in the selection and fitting of candidate pulsars with a torqued precession model. The extension of the monitoring program is used as motivation for the design and development of a new digital pulsar processor for HartRAO. The design methodology and implementation of the processor, using COTS components and reprogrammable logic to ensure sufficient generality in the instrument and to complement the HartRAO observing procedure, has resulted in a multi-purpose digital frontend receiver that could be used for future large scale array telescope designs.

9.1. Braking Pulsars

A literature survey found the current understanding of braking pulsars was incomplete, with various inconsistencies in the concept of a braking law being revealed. A distinction between an observable and theoretical braking index provides an understanding to the discrepancy between theoretically expected and measured values. Measurements of the braking index for pulsars monitored at HartRAO were found to be heavily contaminated with timing noise. As a result, measured braking indices ranged between large positive and negative values, with only two pulsars having physically realistic values in the range $0 < n_{obs} < 10$. The existence of a braking mechanism, in the form of an internal or external viscous torque, that could result in extreme values is improbable. The effects of a viscous torque may be better analysed by the introduction of a dynamic frictional coefficient in a precession model, which may lead to phenomenological changes in the observed timing residuals.

The dynamical behaviour of precessing torqued pulsars was analysed analytically and numerically. The characteristic effects of magnetic field strength, inclination of the magnetic axis, mass distribution, initial angular momentum and vacuum radius on the rotational behaviour of pulsars were determined. The introduction of torques due to a magnetic field adds a second point of attraction in the configuration phase space upon which the nature of precession, Eulerian or radiative or coupled, is dependent. Given the effect of coupled magnetic and crustal stresses on the rotational behaviour of pulsars, it was determined that the observed timing residual behaviour of precessing torqued pulsars would be characteristic of precessing asymmetric bodies.

Using the long dataspans available at HartRAO, previously identified cyclic behaviour in the observed timing residuals of certain pulsars led to the selection of candidate pulsars to be used to test a torqued precession model. It was found that, whilst the quality of timing data

is dependent on the accuracy of the fitted ephemeris, the observed timing residuals of PSRs B1642-03, B1323-58 and B1557-50 could be fitted by a torqued precession model. The physical parameters derived from a manual fitting of the model to the observed timing residuals are listed in table 9.1.1.

TABLE 9.1.1. Fitted pulsar parameters

| | PSR B1642-03 | PSR B1323-58 | PSR B1557-50 |
|--|-------------------|------------------------|--------------------|
| $\dot{P}_1 = \frac{I^{(3)} - I^{(2)}}{I^{(3)}} [10^{-10}]$ | -21.3674029845679 | 1.129041877345285 | 4.99815017433773 |
| $\dot{P}_2 = \frac{I^{(2)} - I^{(1)}}{I^{(2)}} [10^{-10}]$ | \dot{P}_1 | $100 \times \dot{P}_1$ | \dot{P}_1 |
| B [T] | 1.1×10^9 | 2.6×10^8 | 2.25×10^8 |
| χ [deg] | 80.1 | 89.89 | 89.84 |
| $\theta(0)$ [rad] | 0.937 | 0.022 | 0.010 |

PSR B1642-03 was modeled as a prolate rigid body. It is unlikely that this is the original configuration of the pulsar, since the crust is partially elastic at birth and would yield to reconstructions due to the minimisation of elastic energy. It is most likely to have resulted, after sufficient cooling, from a discrete structural event such as a starquake. PSRs B1323-58 and B1557-50 are both modeled as oblate rigid bodies, which is the generally expected configuration for self gravitating, rotating bodies. The magnitude of the fitted B fields for these two pulsars are, in general, the same order of magnitude as those values listed in the ATNF Pulsar Catalogue [81]. The fitted value of B for PSR B1642-03 is an order of magnitude greater than the catalogued value. Including the effect of the fitted values for χ in the model results in calculated pulse period derivatives that are comparable with catalogued values for all three pulsars (see table 2.6.1). The large wobble angle $\theta(0)$ for PSR B1642-03 would need to have resulted from either a massive structural event, such as a collision, or the original configuration of the pulsar at birth. Speculation on the large wobble angle is not likely to be confirmed any time soon.

It is probable that, with further monitoring, the nature of cyclic variations in the timing residuals may change due to more accurate ephemeris describing the spin down.

The implementation and fitting of the torqued model is limited by available computer processing power. Although only six parameters were adjustable for data fitting (magnetic field B, precession parameter P fp, magnetic inclination angle χ , initial wobble angle $\theta(0)$, oblateness pro and z), and two for a minimum Chi squared approximation ($[B - P \text{ fp}]$ and $[\chi - \theta(0)]$), the number of explicit variables is 10 for a torqued model with a braking index of 3. The use of a high performance computing cluster would allow for multi-variate optimisation. Using such a facility, in the future further braking mechanisms could be included into the model for the study of spin down behaviour.

9.2. Remarks on Modeling Precession

Nonlinear multi-variate optimisation is required for the fitting of timing residuals with a torqued precession model. The Levenberg-Marquadt algorithm is shown to be an effective strategy for a broad spectrum of optimisation problems [130]. With an automatic method

of parameter optimisation, the precession model provides a method for the identification of braking mechanisms. The standard dipole torque has been adopted for the purposes of this thesis. This could be adapted to accept different forms of the braking mechanism, including a dynamic frictional torque, which would result not only in changes in the spin down of the pulsar, but also the characteristic timing residual behaviour of precession.

Discontinuities in timing residuals, resulting from a discontinuous piecewise function calculated for $0_1 < \chi < 0_2$, could be misinterpreted as regular glitches in current observing programmes. Damped precession has previously been identified as a result of starquake induced glitches [65], resulting in inter-glitch precessional behaviour. However, for certain time intervals between discontinuities for $0_1 < \chi < 0_2$, the observed timing residuals may not show any evidence of precession. Because of the doubling of processing power required in order to fit a piecewise discontinuous function, this has not been explored in depth in this thesis. Furthermore, the detection of a change in precession mode (radiative to Eulerian), which would result in either discontinuous or long term changes in the pulsar period, is left for future work.

For the purposes of this thesis, an isolated neutron star was considered, rotating *in vacuo* with a fixed vacuum radius. This radius was chosen with limited knowledge and evidence of the actual length of interaction between the near field and far field radiation zones. Magnetospheric effects may strongly alter the precessional behaviour, including temporal distortions in the radiative torque, an effect that has not been taken into account. The effect on precession due to variations in the pulsar period derivative, although not significantly altering the observed precessional behaviour, was removed from the fitted timing residuals, but should be included in future modeling.

9.3. Future Pulsar Processor Implementation

The original pulsar processor described in chapter 6 is implemented as a hybrid filterbank receiver, with digital backend and analogue frontend. Digital backend modules perform an eight way digital division of each 32 MHz IF slice mixed to baseband by the analogue frontend modules. The overall sensitivity of the new processor is a factor 3 improvement of the current HartRAO pulsar processor, which results in a threefold increase in the detectable pulsar sub-population. A large fraction of the total capital outlay for the new processor is dependent on the number of digital backend modules used.

Cost effectiveness and efficiency of the pulsar processor improves almost sevenfold with the implementation of a polyphase filtering system in place of the eight FIR digital filters. However, implementation of the polyphase system was stalled by problems encountered with the Xilinx ISE 4.2i Development System, which was used for the initial development of the polyphase filter. Whilst simulation of the design performed to specification, including verification of the design with the Xilinx company, the automatic bit file generation, intended for uploading to the FPGA, contained errors and caused the system to be unusable. This section describes the design of a polyphase filter system that is to be implemented using the Xilinx ISE 6.3i

Development System. This software is free of the errors experienced while using the Xilinx ISE 4.2i Development System.

9.3.1. Polyphase Filtering System. A comprehensive analysis of polyphase and multi-rate filtering is given by Vaidyanathan (1992) [172].

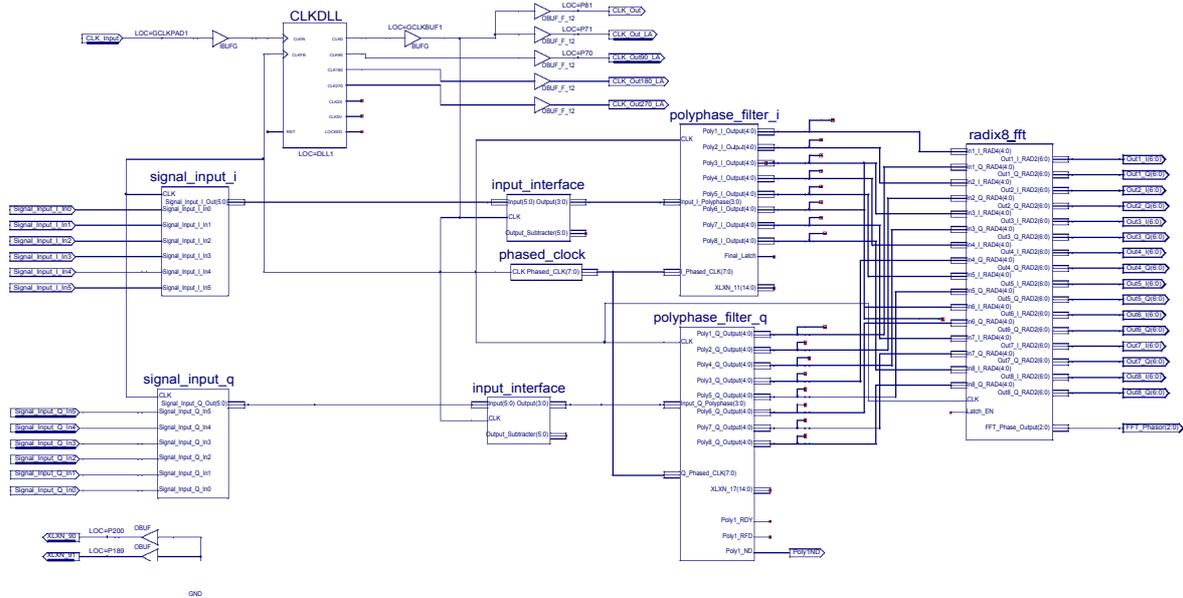


FIGURE 9.3.1. FPGA schematic design for top hierarchy of polyphase filtering system. The design is split into three functional blocks: input interface, polyphase filters and a Fast Fourier transform. The input interface removes DC offsets and outputs a 5 bit, two’s complement signal. The polyphase filter system includes eight subfilters, each with a 5 bit output signal. The radix8 FFT performs a $4 \times 2 = 8$ point FFT with eight complex output signals.

Figure 9.3.1 is the schematic design for a polyphase filter system and 8 point complex FFT, which produces the complex components of each narrowband channel in the receiver’s composite passband. The design is split into three functional blocks: an input interface, polyphase filter and a $4 \times 2 = 8$ point FFT.

The polyphase filter system, shown in figure 9.3.2, comprises eight serial distributed arithmetic FIR subfilters. These are area efficient filters, which need only produce an output word every eight clock cycles. The input signal is multiplexed across the filters by a 3 bit selector. The coefficients of each subfilter, $n = 1..8$, are

$$W_n^0[i] = W[8 \times i + n]$$

where $i = 0..4$, i an integer and W are the filter coefficients used in the original FIR low pass filter design (see figure 7.11.1).

The 8 point FFT is partitioned into a 4 point and 2 point FFT (see appendix I.1), based on the Cooley Tukey algorithm for $N = r_1 r_2 = 4 \times 2 = 8$ [30], with calculated weights listed in table I.2.1. Figure 9.3.3 is the schematic design for the 4 point FFT subsystem, which is composed of three functional blocks: the input interface, 4 point FFT weight multipliers and the output interface (see appendix I.1 for the 2 point FFT subsystem). The eight complex input signals are fanned out to four signal paths, each of which is multiplied by a 2 bit rotating phasor. Similarly, the 2 point FFT uses two signal paths with multiplication by a 3 bit phasor. Phasor weights for the FFT subsystems are listed in table 9.3.2.

Repetition of weights in each branch increases the scope for optimisation. Schematic designs of the weight multipliers for both the 4 point (sets 1 to 4) and 2 point (sets 1 and 2) FFT subsystems are shown in appendix I.1. Whilst the 4 point FFT phasor requires the swapping of complex components and simple addition and subtraction, the 2 point FFT makes use of lookup table based multiplication.

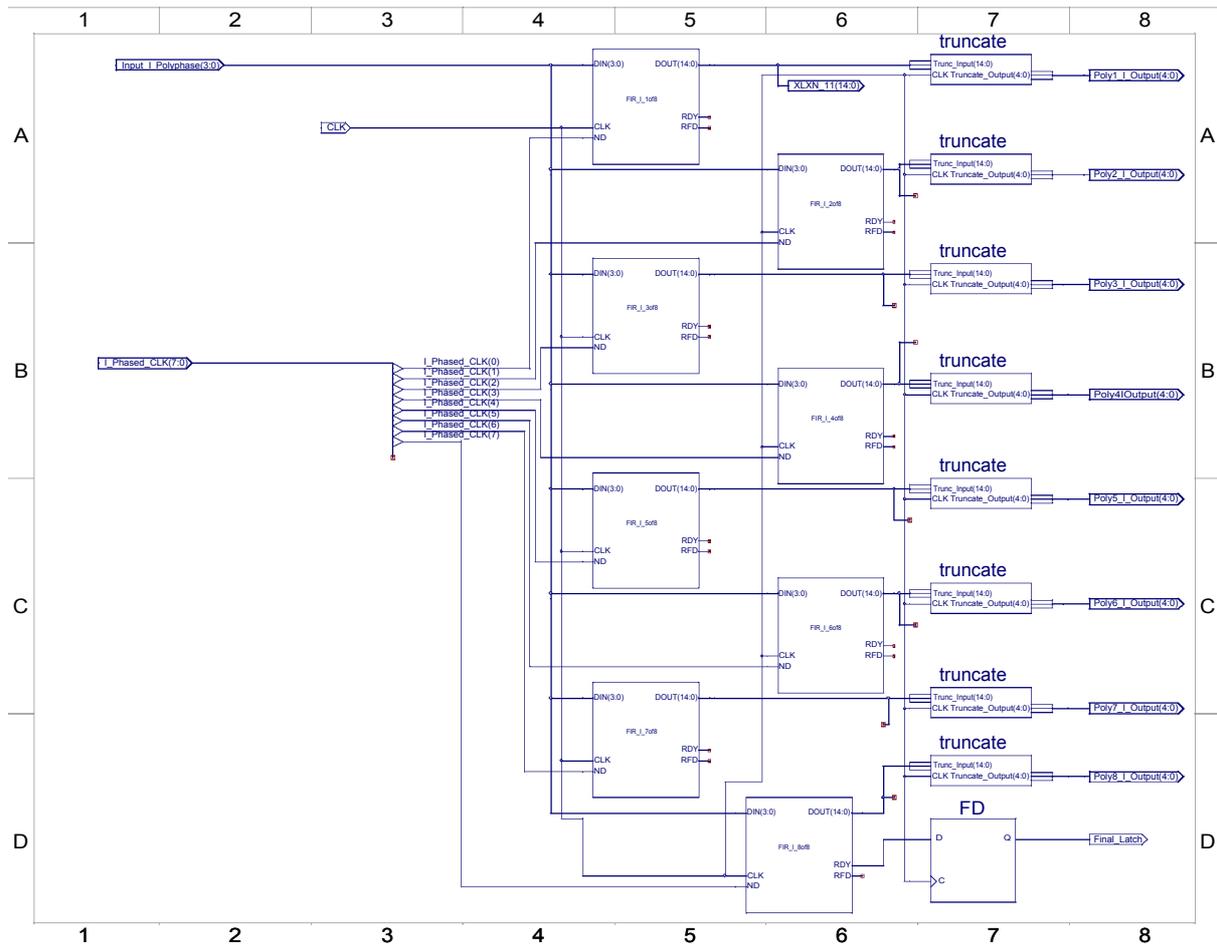


FIGURE 9.3.2. FPGA schematic design of polyphase filter. The phased clock signal, connected to the new data (ND) pins of each filter, effectively multiplexes the input signal across the polyphase subfilters. The 15 bit output signal is truncated to 5 bits and latched.

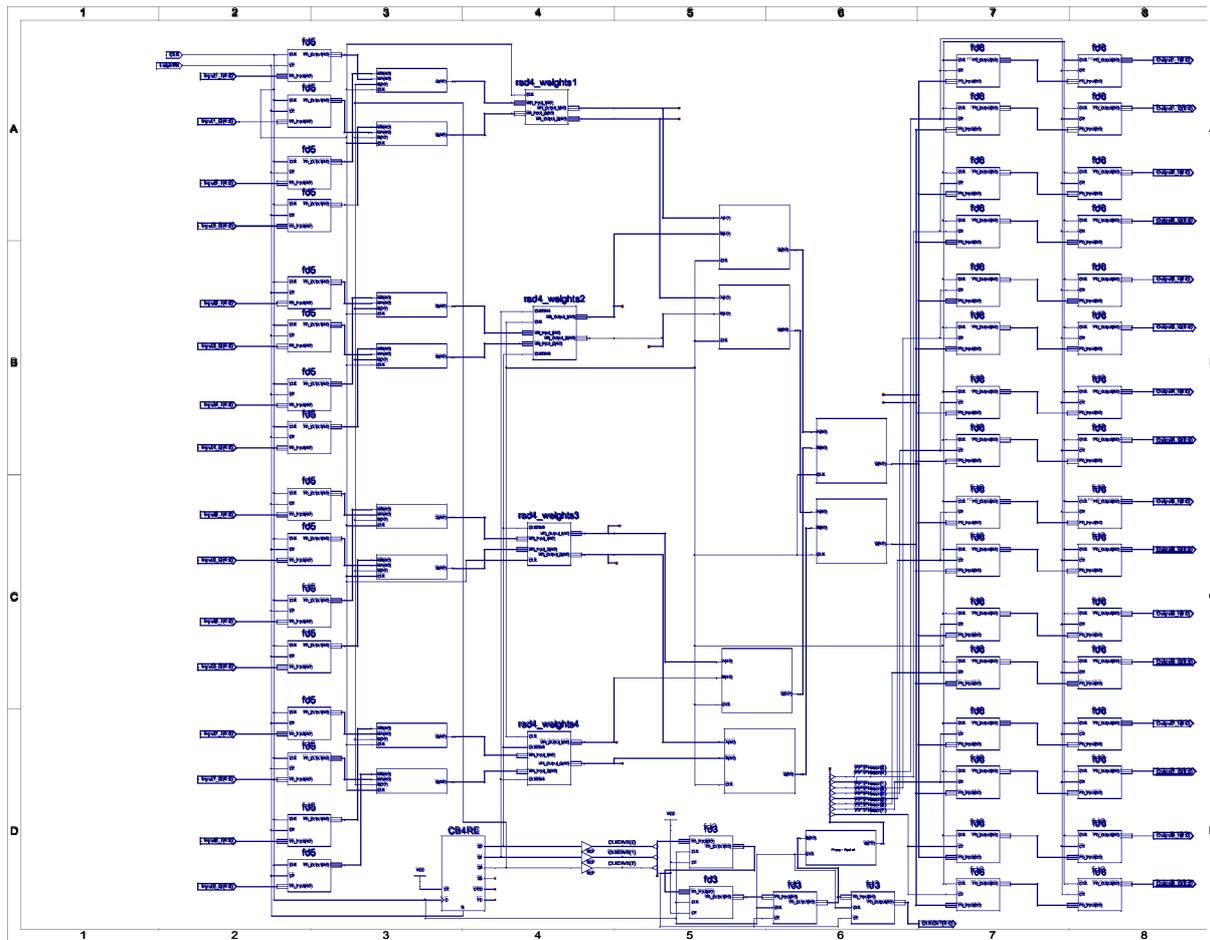


FIGURE 9.3.3. FPGA schematic design of 4 point FFT. The system is separated into three functional blocks: the input interface, the 4 point FFT weight multipliers and the output signal interface. The eight complex input signals are multiplexed between four signal paths. The signal is multiplied by a rotating weighted phasor, and demultiplexed by the output interface, where the signal is latched. Data flip-flops are used to insert serial delays on the clock phasor.

TABLE 9.3.2. Weights for each signal branch of 4 point and 2 point FFTs

| Phasor Values | |
|---------------|--|
| 4 Point FFT | |
| Branch 1 | +1,+1,+1,+1 |
| Branch 2 | +1,-j,-1,+j |
| Branch 3 | +1,-1,+1,-1 |
| Branch 4 | +1,+j,-1,-j |
| 2 Point FFT | |
| Branch1 | $+1,+1,+1,\frac{\sqrt{2}}{2} - j\frac{\sqrt{2}}{2},+1,-j,+1,-\frac{\sqrt{2}}{2} - j\frac{\sqrt{2}}{2}$ |
| Branch2 | $+1,-1,+1,-\frac{\sqrt{2}}{2} + j\frac{\sqrt{2}}{2},+1,+j,+1,\frac{\sqrt{2}}{2} + j\frac{\sqrt{2}}{2}$ |

9.3.2. Concluding Remarks. The future implementation of a polyphase based pulsar processor will place high emphasis on its performance as a multi-purpose radio astronomy digital radiometer. The ability to produce low cost scientific instrumentation is important to the construction of large scale array telescope designs such as the Square Kilometer Array (International, <http://www.skatelescope.org>) and associated concept designs, where multiple beam forming will require high speed, reprogrammable digital instrumentation.

The South African Square Kilometer Array team (<http://www.ska.ac.za>) have recently embarked on a program to design, develop and build an SKA science demonstrator, known as the Karoo Array Telescope (KAT), to be commissioned by the end of 2008. The polyphase based pulsar processor is to form the basis of a fully digital frontend receiver for KAT, with a total bandwidth of 500 MHz, maximum configurable frequency resolution of 125 KHz per channel (for up to 4000 channels) and time resolution of 0.1 μ s. Vivaldi elements will receive dual polarisation, from which full Stokes parameters will be produced by the frontend receiver. A costing analysis for the system forms part of a risk reduction project currently being defined.

APPENDIX A

Pulsar Mechanics

A.1. Determination of Euler Equations

The motion of a rigid body about a fixed point O is considered. The angular momentum in tensor form is written as

$$\mathbf{J} \sim = \mathcal{I}^{\mu\nu} \omega_\nu \mathbf{e}_\mu$$

where \mathbf{e}_μ is any set of basis vectors for a given reference frame. Using a body frame removes time dependencies in the moment of inertia $\mathcal{I}^{\mu\nu}$. The equation of motion is then found to be

$$\begin{aligned} \frac{d\mathbf{J} \sim}{dt} &= \mathcal{T} \sim \\ &= \frac{d}{dt} (\mathcal{I}^{\mu\nu} \omega_\nu \mathbf{e}_\mu) \\ &= \mathcal{I}^{\mu\nu} \dot{\omega}_\nu \mathbf{e}_\mu + \mathcal{I}^{\mu\nu} \omega_\nu \dot{\mathbf{e}}_\mu \\ &= \mathcal{I}^{\mu\nu} \dot{\omega}_\nu \mathbf{e}_\mu + \mathcal{I}^{\mu\nu} \omega_\nu \omega_\mu^\alpha \mathbf{e}_\alpha \\ &= \mathcal{I}^{\alpha\nu} \dot{\omega}_\nu \mathbf{e}_\alpha + \mathcal{I}^{\mu\nu} \omega_\nu \omega_\mu^\alpha \mathbf{e}_\alpha \\ \Rightarrow \mathcal{T} \sim^\alpha &= \mathcal{I}^{\mu\nu} \dot{\omega}_\nu + \mathcal{I}^{\mu\nu} \omega_\nu \omega_\mu^\alpha \end{aligned}$$

which is the tensor form of Euler's equations.

A.2. Co-ordinate Transformations

In an attempt to describe the motion of a rigid body, three independent parameters are required to describe the orientation of the body. The most commonly used are known as the Euler angles.

It is possible to move from a cartesian coordinate system to another, such as from an inertial space frame of reference to an body based frame of reference, by means of three sets of rotations in a given sequence. Each transformation matrix describes a rotation by a fixed angle, called an Euler angle. The first rotation consists of a counterclockwise rotation of the space axes, described by the basis \mathbf{E}_j , about \mathbf{E}_3 by a fixed angle φ . The new set of axes $\tilde{\mathbf{e}}_i$, is obtained by a transformation matrix \mathbf{a}_i^j , such that

$$\tilde{\mathbf{e}}_i = \mathbf{a}_i^j \mathbf{E}_j \quad \mathbf{a}_i^j = \begin{bmatrix} \cos \varphi & \sin \varphi & 0 \\ -\sin \varphi & \cos \varphi & 0 \\ 0 & \mathbf{0} & \mathbf{1} \end{bmatrix}$$

A second rotation consists of a counterclockwise rotation of \bar{e}_j about \bar{e}_1 by a fixed angle θ . The new set of axes n_j is known as the nodal axes which describe the nodal frame. Thus

$$n_i = b_{i \ a k}^j E_k \quad b_i^j = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & -\sin \theta & \cos \theta \end{bmatrix}$$

The third rotation consists of a counterclockwise rotation of the nodal axes n_j about n_3 by a fixed angle ψ . The new set of axes e_j describes the body frame.

$$e_i = c_i^j b_j^k a_k^l E_l \quad c_i^j = \begin{bmatrix} \cos \psi & \sin \psi & 0 \\ -\sin \psi & \cos \psi & 0 \\ 0 & \mathbf{0} & \mathbf{1} \end{bmatrix}$$

The matrix $c_i^j b_j^k a_k^l$ to transform the system from an inertial space from to a body frame is then

$$c_i^j b_j^k a_k^l = \begin{bmatrix} \cos \psi \cos \varphi - \sin \psi \sin \varphi \cos \theta & \cos \psi \sin \varphi + \sin \psi \cos \varphi \cos \theta & \sin \psi \sin \theta \\ -\sin \psi \cos \varphi - \cos \psi \sin \varphi \cos \theta & -\sin \psi \sin \varphi + \cos \psi \cos \varphi \cos \theta & \cos \psi \sin \theta \\ \sin \varphi \sin \theta & -\cos \varphi \sin \theta & \cos \theta \end{bmatrix}$$

APPENDIX B

Precession Model Fits

B.1. PSR B1642-03

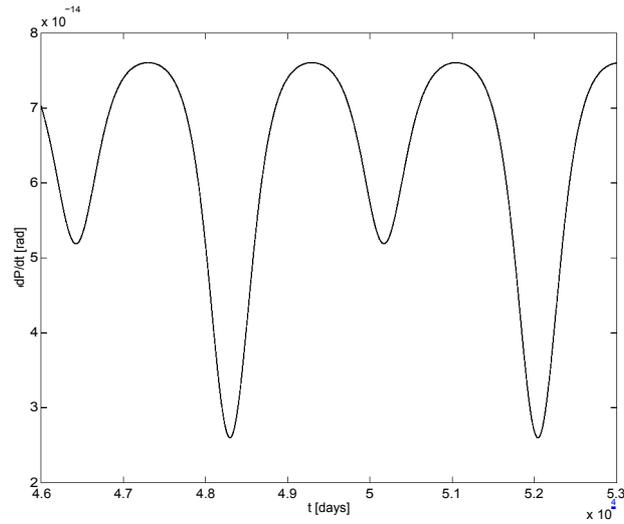


Figure B.1.1: Calculated period derivative of PSR B1642-03

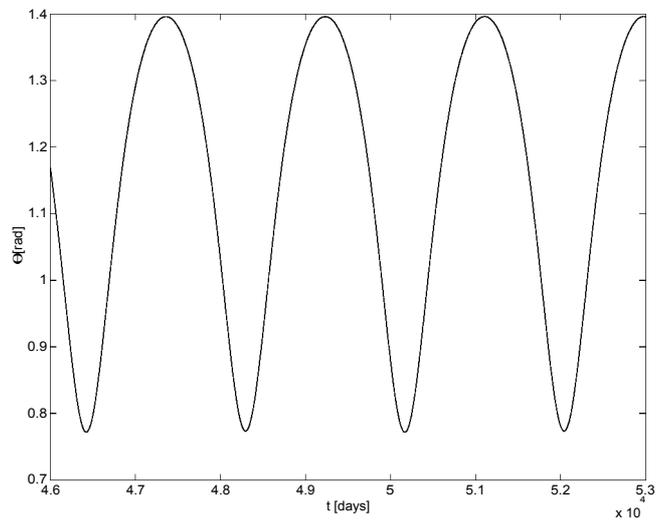


Figure B.1.2: Calculated $\theta(t)$ of PSR 1642-03

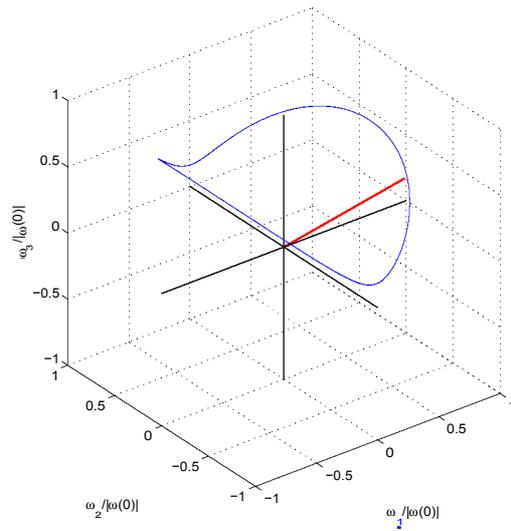


Figure B.1.3: Calculated path of $\boldsymbol{\omega}(t)$ for PSR B1642-03

B.2. PSR B1323-58

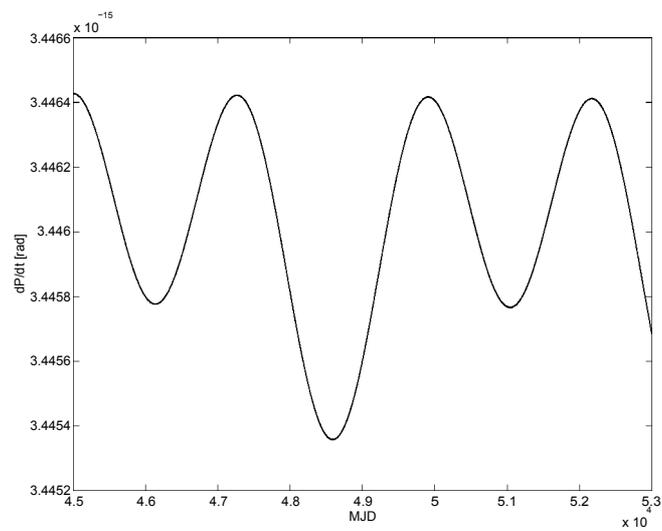


Figure B.2.1: Calculated period derivative of PSR B1323-58

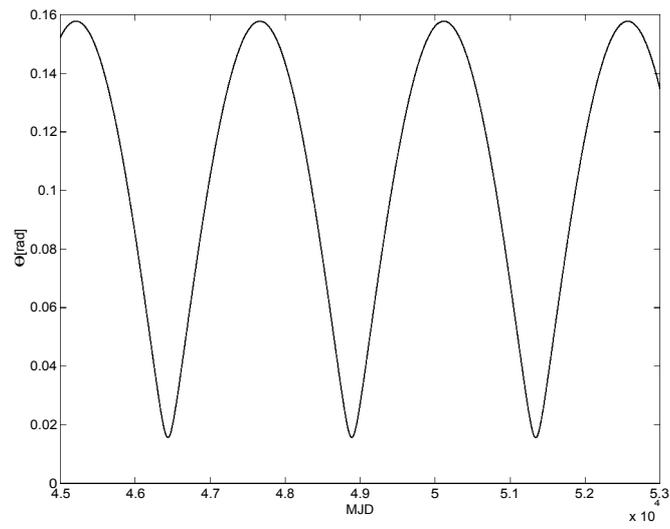


Figure B.2.2: Calculated $\theta(t)$ of PSR B1323-58

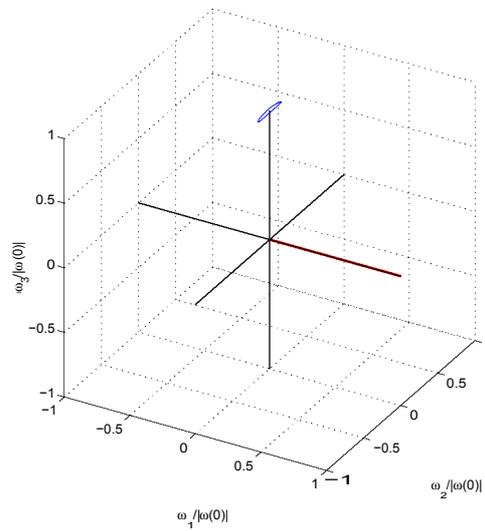


Figure B.2.3: Calculated path of $\omega(t)$ for PSR B1323-58

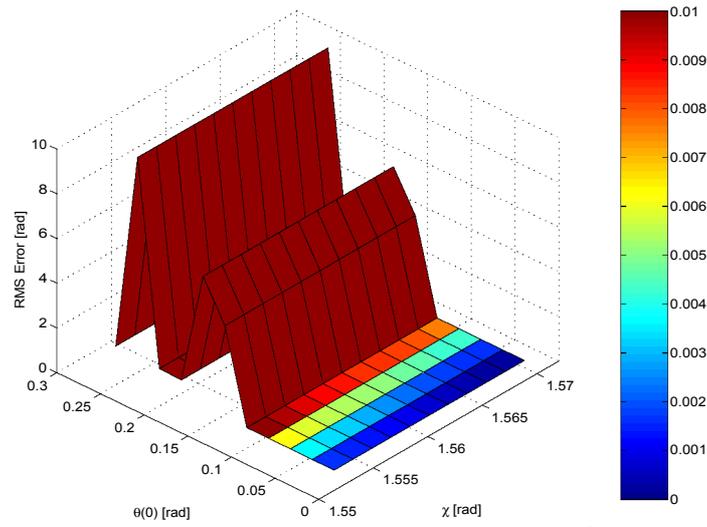


Figure B.2.4: Chi squared surface for constant $B = 2.6 \times 10^8$ [T], $Pfp = 49000$ [days]

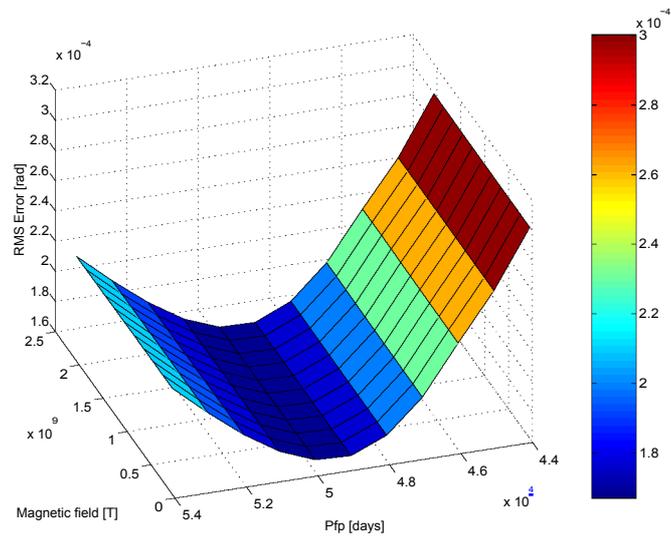


Figure B.2.5: Chi squared surface for constant $\chi = 89.89$ [deg] and $\theta(0) = 0.022$ [rad]

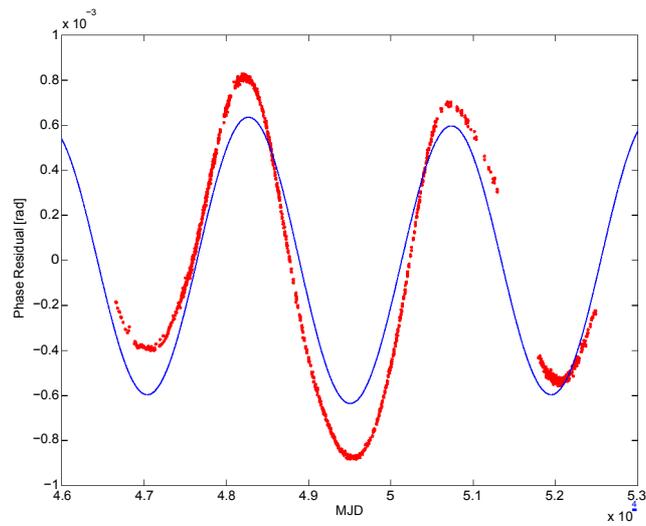
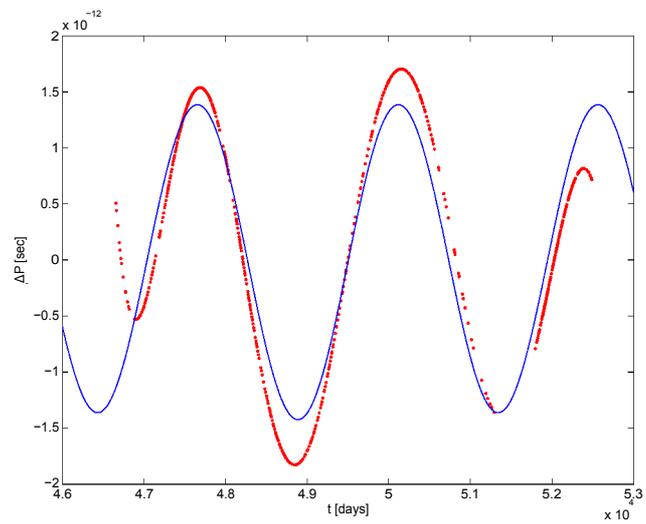


Figure B.2.6: Calculated phase residual for best fit of PSR B1323-58

Figure B.2.7: Calculated ΔP for best fit of PSR B1323-58 timing residuals

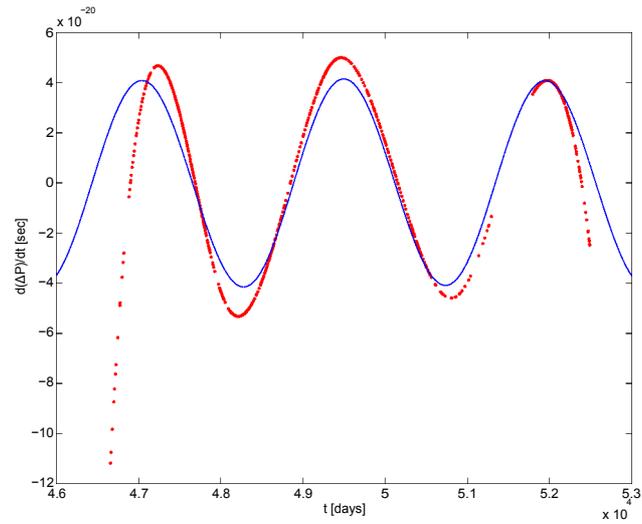


Figure B.2.8: Calculated $\frac{d(\Delta P)}{dt}$ for best fit of PSR B1323-58 timing residuals

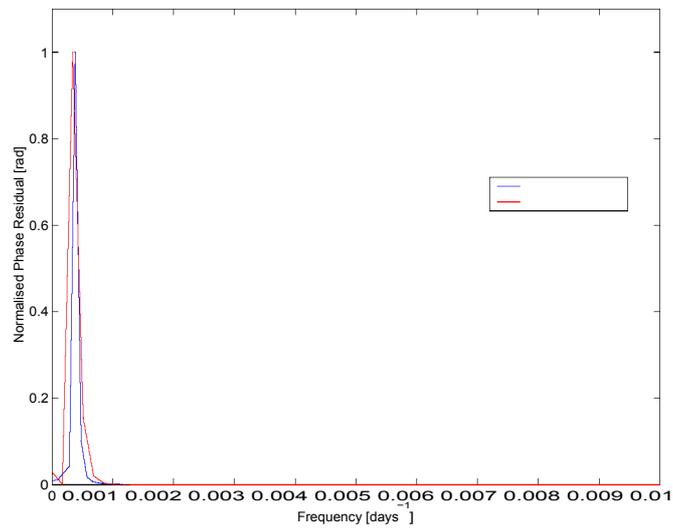


Figure B.2.9: Spectrum of phase residual calculated for best fit of PSR B1323-58

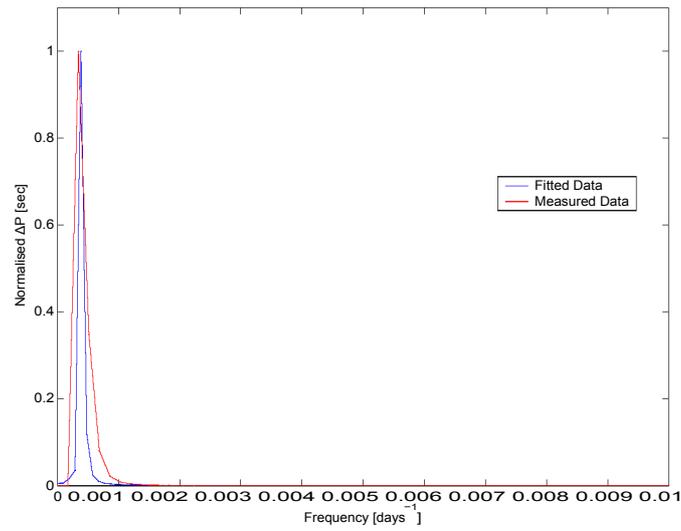


Figure B.2.10: Spectrum of ΔP calculated for best fit of PSR B1323-58 timing residuals

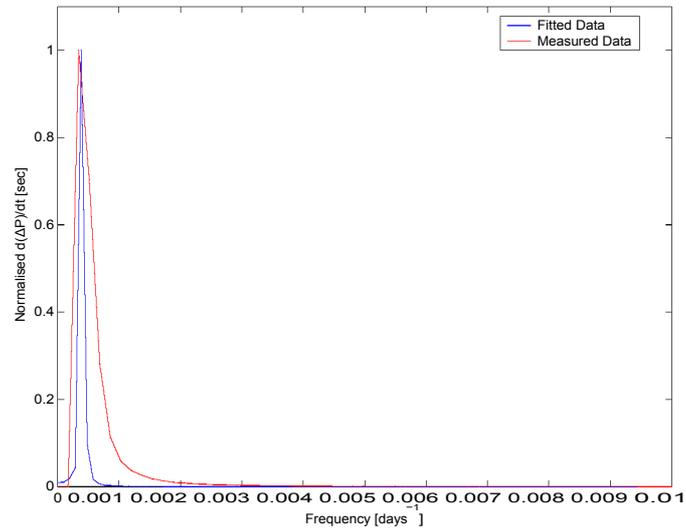


Figure B.2.11: Spectrum of $\frac{d(\Delta P)}{dt}$ calculated for best fit of PSR B1323-58 timing residuals

B.3. PSR B0736-40

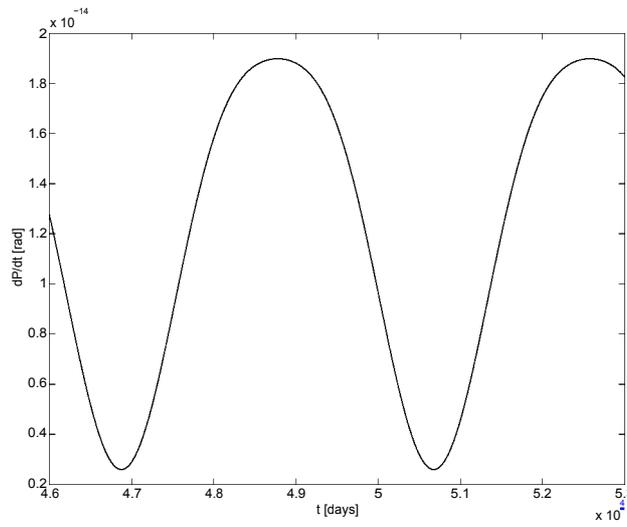


Figure B.2.11: Figure B.3.1: Calculated period derivative of PSR B0736-40

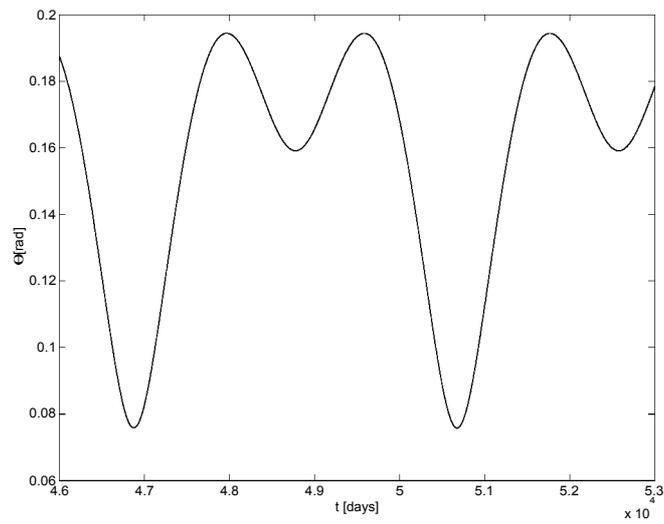


Figure B.3.2: Calculated $\theta(t)$ of PSR B0736-40

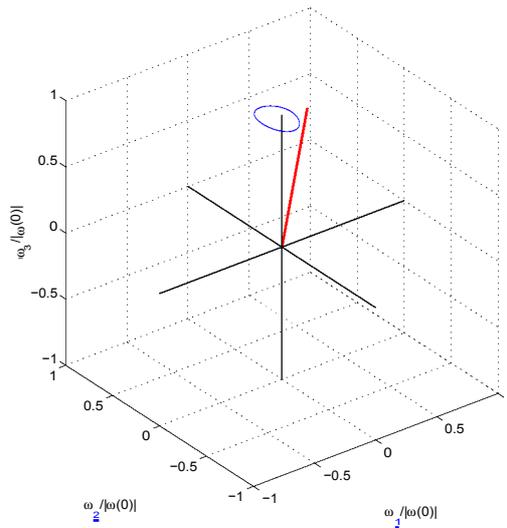


Figure B.3.3: Calculated path of $\boldsymbol{\omega}(t)$ for PSR B0736-40

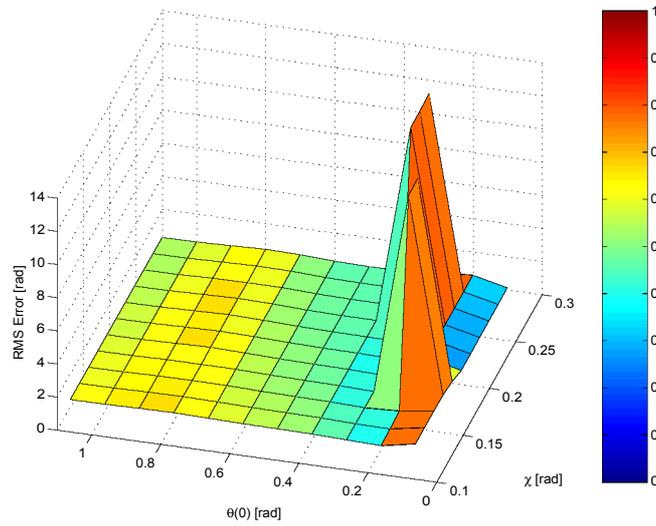


Figure B.3.4: Chi squared surface for fitting of PSR B0736-40, constant $B = 1.6 \times 10^9$ [T] and $Pfp = 3800$ [days]

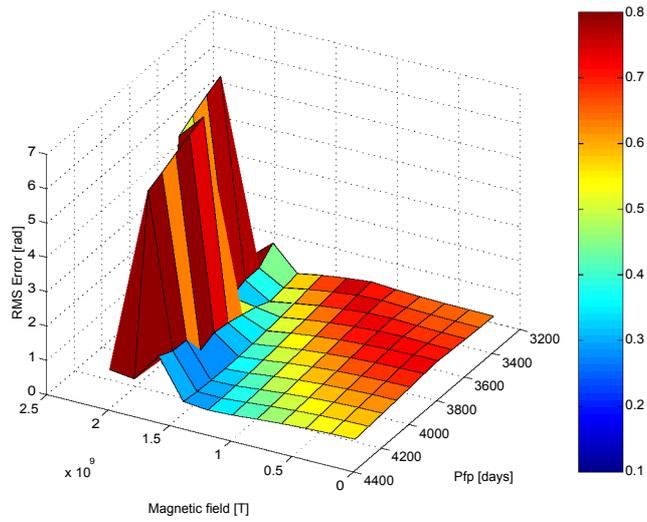


Figure B.3.5: Chi squared surface for fitting of PSR B0736-40, constant $\chi = 12.00$ [deg] and $\theta(0) = 0.096$ [rad]

B.4. PSR B1557-50

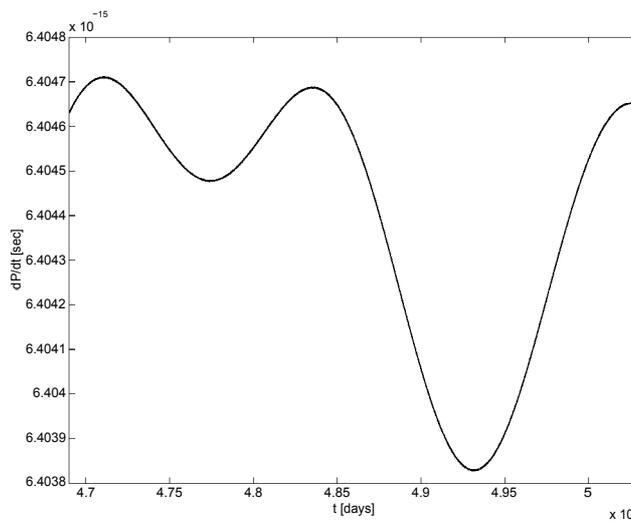


Figure B.4.1: Calculated period derivative of PSR B1557-50

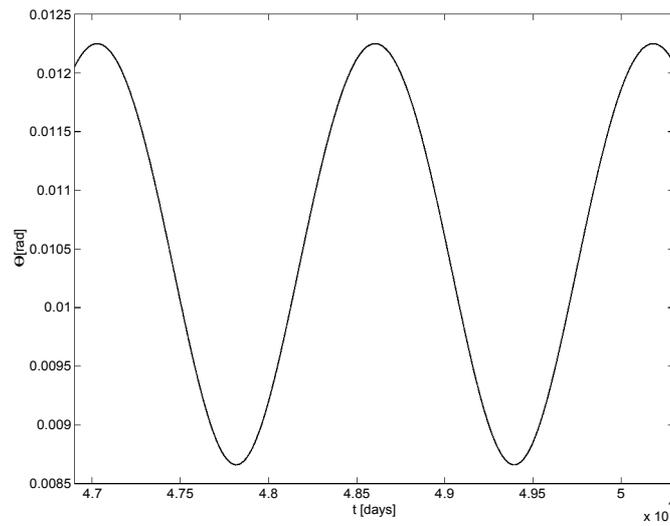


Figure B.4.2: Calculated $\theta(t)$ of PSR B1557-50

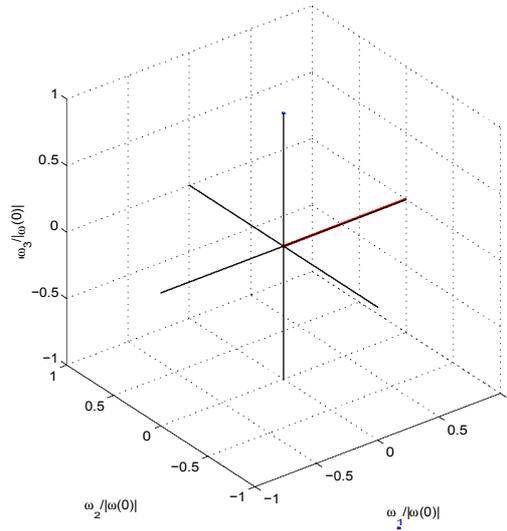


Figure B.4.3: Calculated path of $\omega(t)$ for PSR B1557-50

APPENDIX C

Signal Processing

C.1. Quadrature Mixing and Sideband Separation

The Fourier transform of $x(t)$ is

$$(C.1.1) \quad X(\omega) = X_L(\omega) + X_U(\omega)$$

where $X_L(\omega)$ and $X_U(\omega)$ are the lower and upper sidebands. The input signal is

$$(C.1.2) \quad \tilde{X}(\omega) = X(\omega + \omega_c) + X^*(-\omega - \omega_c)$$

in terms of standard signal theory.

The Fourier components for $\cos(\omega_c t)$ are

$$\delta(\omega + \omega_c) + \delta(\omega - \omega_c)$$

and for $\sin(\omega_c t)$ are

$$j[\delta(\omega + \omega_c) - \delta(\omega - \omega_c)]$$

The real and imaginary outputs are

$$Y_r(\omega) = \tilde{X}(\omega) * [\delta(\omega + \omega_c) + \delta(\omega - \omega_c)] \dots \text{convolution with } \cos(\omega_c t)$$

$$Y_i(\omega) = \tilde{X}(\omega) * j[\delta(\omega + \omega_c) - \delta(\omega - \omega_c)] \dots \text{convolution with } \sin(\omega_c t)$$

By substituting in equations C.1.1 and C.1.2, the real output is

$$(C.1.3) \quad Y_r(\omega) = [X(\omega + \omega_c) + X^*(-\omega - \omega_c)] * [\delta(\omega + \omega_c) + \delta(\omega - \omega_c)]$$

$$(C.1.4) \quad = X(\omega + 2\omega_c) + X^*(-\omega) + X(\omega) + X^*(-\omega - 2\omega_c)$$

$$(C.1.5) \quad = X(\omega) + X^*(-\omega) \text{ post anti aliasing filter}$$

$$(C.1.6) \quad = X_L(\omega) + X_U(\omega) + X_L^*(-\omega) + X_U^*(-\omega)$$

while for the imaginary output is

$$(C.1.7) \quad Y_i(\omega) = [X(\omega + \omega_c) + X^*(-\omega - \omega_c)] * j[\delta(\omega + \omega_c) - \delta(\omega - \omega_c)]$$

$$(C.1.8) \quad = j[X(\omega) - X^*(-\omega) - X(\omega + 2\omega_c) + X^*(-\omega - 2\omega_c)]$$

$$(C.1.9) \quad = j[X(\omega) - X^*(-\omega)] \text{ post anti aliasing filter}$$

$$(C.1.10) \quad = j[X_L(\omega) + X_U(\omega) - X_L^*(-\omega) - X_U^*(-\omega)]$$

The Hilbert transform of $Y_i(\omega)$ is

$$\begin{aligned}\tilde{Y}_i(\omega) &= j[(j)X_L(\omega) + (-j)X_U(\omega) - (-j)X_L^*(-\omega) - (j)X_U^*(-\omega)] \\ &= -X_L(\omega) + X_U(\omega) - X_L^*(-\omega) + X_U^*(-\omega)\end{aligned}$$

C.2. FIR Filter Coefficients

TABLE C.2.1. Quantised filter coefficients for real and imaginary filters

| Real Filter | Coefficient Index | Imaginary Filter |
|-------------|-------------------|------------------|
| -4 | 1 | 1 |
| -5 | 2 | 3 |
| -8 | 3 | 5 |
| -14 | 4 | 8 |
| -25 | 5 | 9 |
| -42 | 6 | 7 |
| -63 | 7 | -2 |
| -86 | 8 | -21 |
| -108 | 9 | -52 |
| -123 | 10 | -97 |
| -123 | 11 | -154 |
| -103 | 12 | -217 |
| -59 | 13 | -279 |
| 8 | 14 | -329 |
| 96 | 15 | -360 |
| 197 | 16 | -361 |
| 300 | 17 | -330 |
| 392 | 18 | -265 |
| 461 | 19 | -172 |
| 498 | 20 | -60 |
| 498 | 21 | 60 |
| 461 | 22 | 172 |
| 392 | 23 | 265 |
| 300 | 24 | 330 |
| 197 | 25 | 361 |
| 96 | 26 | 360 |
| 8 | 27 | 329 |
| -59 | 28 | 279 |
| -103 | 29 | 217 |
| -123 | 30 | 154 |
| -123 | 31 | 97 |
| -108 | 32 | 52 |
| -86 | 33 | 21 |
| -63 | 34 | 2 |
| -42 | 35 | -7 |
| -25 | 36 | -9 |
| -14 | 37 | -8 |
| -8 | 38 | -5 |
| -5 | 39 | -3 |
| -4 | 40 | -1 |

APPENDIX D

Low Frequency Concept Design

D.2. Flowcharts

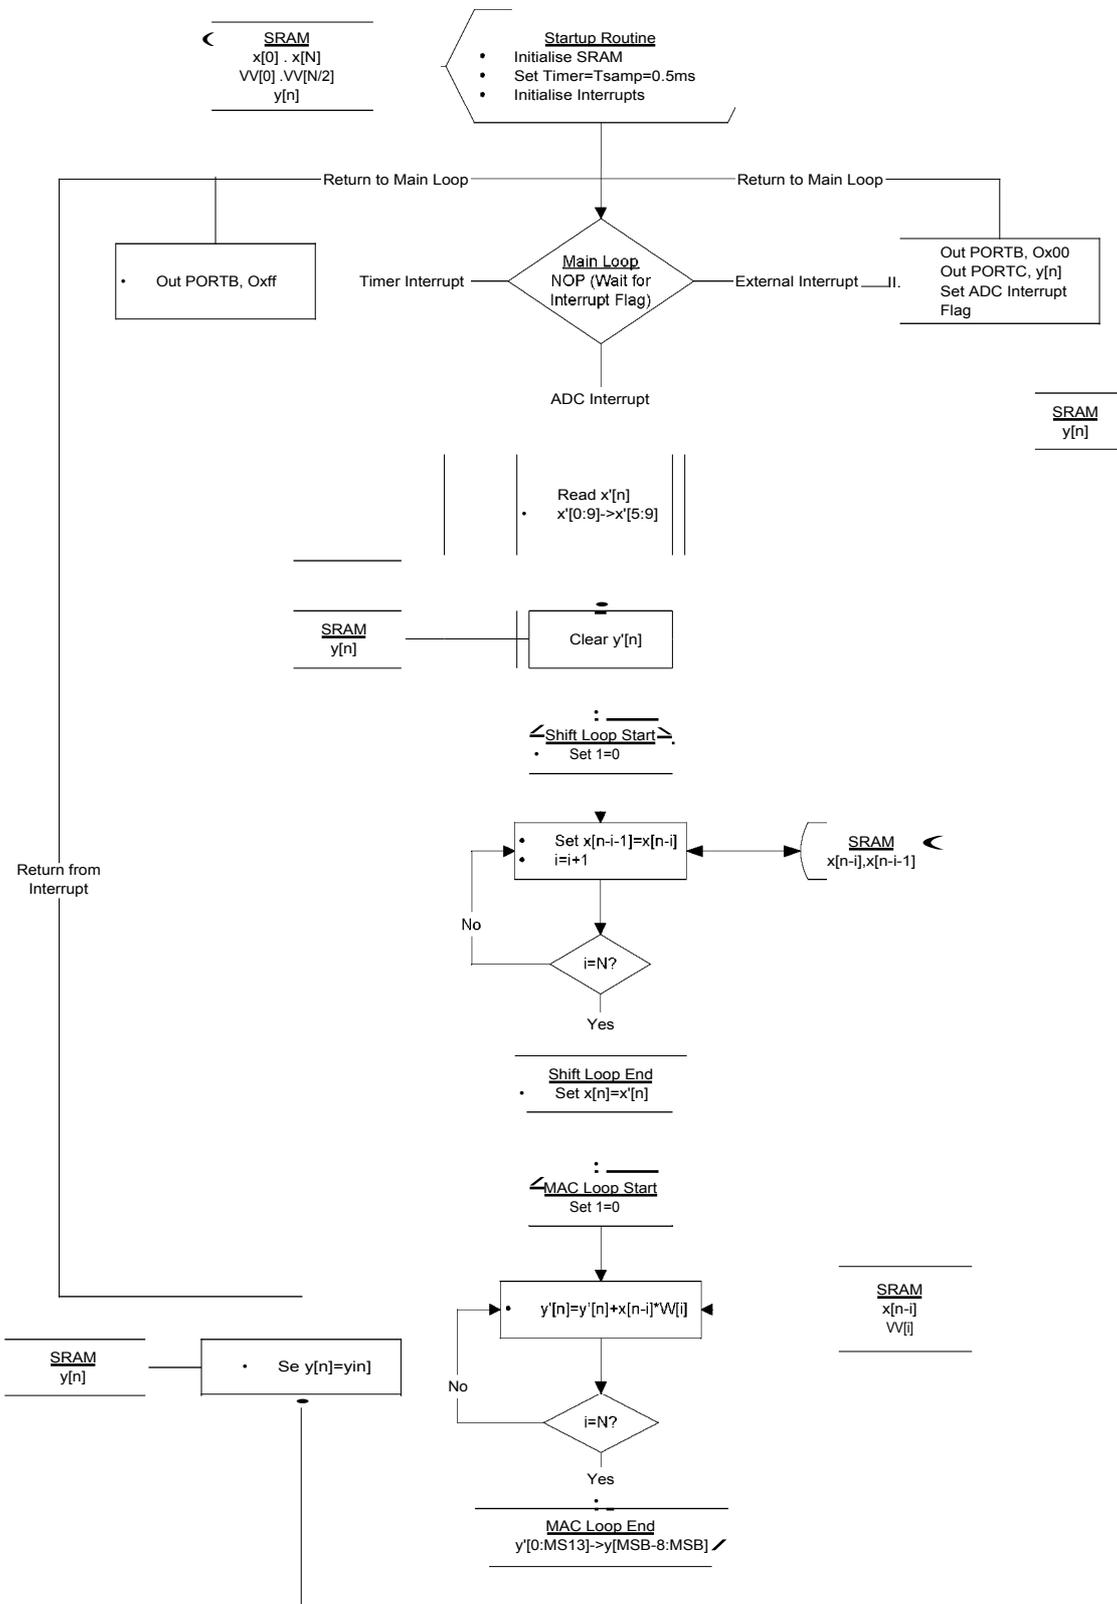


Figure D.2.1: Flowchart for real filter implemented on an AT90S8535 microcontroller