



UNIVERSITY COLLEGE LONDON

Faculty of Mathematics and Physical Sciences

Department of Physics & Astronomy

Technical development and scientific
preparation for the e-MERLIN Cygnus
OB2 radio survey

Thesis submitted for the Degree of Doctor of
Philosophy of the University of London

by
Luke William Peck

Supervisors:

Prof. Raman Prinja
Dr. Jeremy Yates

Examiners:

Prof. Ian Howarth
Dr. Neal Jackson

September 17, 2014

To my family and friends.

I, Luke William Peck, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

e-MERLIN is a recent upgrade to the MERLIN radio array. This enhanced facility utilises recent developments in wide bandwidth receivers, a new WIDAR correlator, and a new optical fibre network. This upgrade provides an increase in sensitivity and image fidelity, but also results in a significant increase in data volume. This thesis is motivated by the Cygnus OB2 Radio Survey (COBRaS), an e-MERLIN Legacy project observing the core region of the largest OB association in the northern hemisphere. COBRaS has been awarded ~ 300 hours observing time, resulting in a total Legacy dataset of tens of terabytes. It is not feasible to calibrate this amount of data manually, highlighting the necessity for automated procedures.

This thesis primarily contains technical development for e-MERLIN during the commissioning phase and early Legacy observations from COBRaS, which focuses on the creation of automated flagging and calibration pipelines. This includes an automated RFI-mitigation and reduction tool (SERPent), as well as a full calibration pipeline consisting of: phase calibration with fringe fitting, amplitude calibration with the flux calibrator 3C286, bandpass calibration with spectral index and curvature fitting, and automated self-calibration on combined or individual IFs. A program for extracting fluxes for resolved and unresolved sources from radio maps with a detection significance boosting module has also been developed.

In addition to the technical work, scientific preparations and initial results for COBRaS are also presented. A catalogue amalgamation routine for the Cyg OB2 association cross correlates previous surveys of Cyg OB2 into one definitive catalogue. Subsequent specific catalogues are compiled from this one catalogue to create an OB star catalogue and candidate catalogue. The predicted mass loss rates and radio fluxes from the winds

of O-type stars and early B-type supergiants are determined, and this includes predictions from smooth wind models as well as predictions including the effects of clumping in the winds. The inclusion of an X-ray variability study of the Chandra Cyg OB2 Legacy dataset, provides a multi-wavelength view of the population of Cyg OB2, which complements COBRaS. The first COBRaS 1.6 GHz and 5 GHz radio images of Cyg OB2 are presented with source and flux lists and some initial analysis.

The technical developments presented in this thesis are discussed in the context of COBRaS and of future interferometers such as the SKA and its associated pathfinders.

Acknowledgements

I would say many thank yous, but I'm not the grateful type... Only joking! Seriously though, this thesis would not exist if it were not for a number of people helping and supporting me over these last few years. Firstly the academic thank yous. A huge thank you to Raman who nurtured me for many years, firstly with my MSci project during my undergraduate, and all the way through this PhD. Your enthusiasm, humour and guidance has kept me going all the way. I can't thank you enough for the opportunity to work on this exciting project with you.

Secondly, many thanks to Jeremy for keeping tabs on me through the years, keeping me on track to finish, and introducing me to the *complex* (ahem) world of radio interferometry.

Thirdly, a massive thank you to Danielle, for being on the front-line helping me with matters on radio interferometry, AIPS, programming, life, and everything in between. I really couldn't have done it without you! You always made time to answer my questions (however silly) even when you were working from home (aka 'Pimms day').

A big thank you also goes out to all the people who have helped me in some way, shape or form with the work in this thesis. Thanks go to: Dugan for helping me in the early days with your general computer wizardry, Rob and Megan for testing SERPent on the machines in Manchester and ASTRON, Robert and Anita with their useful feedback on SERPent, Nick and Gregor for help with the Chandra X-ray study, and Jason and Boris for discussions on the scary Bayesian stuff.

Next, the social (-ly unacceptable) thank yous. These go to pretty much everyone in Group A (a catch-all statement as I can't name everyone!). Thanks go to: Ingo and Cris, for our non-Euclidean adventures, stimulating conversations at both 'ends' of the spectrum, and general drunkenness and fun times. To Nic for giving me a daily summary

of the internet, for our many bike rides, and general drunkeness and fun times. To Japh, Cam, Adam, Paul and Georgie for some amazing film nights and general drunkeness and fun times. To Stephanie and Harpreet for your support and general drunkeness and fun times. To Tomek, Patrick and Tomas for the musical entertainment, concerts and... general drunkeness and fun times. To the people who went on the Oktoberfest (particularly Dave for drinking 10 litres of beer in a day!), Amsterdam and cottage trips, these were some of the best weekends I've ever had!

Thanks again to all the people in Group A, too numerous to list fully, that I've met over the last few years. Thanks to my office mates in G14 and that 'popcorn office'. I have thoroughly enjoyed every minute and thank every one of you!

Thanks also to all the radio interferometry people I've met at the different conferences and workshops all over the world... (for the general drunkeness and fun times).

I would like to thank my loving family for all the support and encouragement you've provided throughout my life. I can't put into words how much I appreciate you all being there for me, through the good, the bad, and the ugly. I am so lucky to have such an amazing family and thank you for your support in everything I do.

Lastly, but most certainly not least, a enormous thank you to Sílvia. I can't think of enough superlatives to describe and thank all of the love and support you've given me. All the adventures we've been on, conversations we've had, and for being with me for all of the events and listings stated above... (and yes of course, general drunkeness and fun times). You make me so happy, and our time together during our PhDs have made these years the best years of my life. *T'estimo moltíssim.*

“Anything is possible, but not everything is probable.”

Contents

Table of Contents	7
List of Figures	13
List of Tables	19
1 Introduction	21
1.1 Introduction to Radio Interferometry	22
1.1.1 The Two Element Interferometer	24
1.1.2 Radio Antenna and Receivers	29
1.1.3 Complex Correlators	34
1.1.4 Visibilities and the Visibility Plane (Fourier Space)	36
1.1.5 Radio Interferometric Imaging via Deconvolution	38
1.1.6 Polarisation Effects	43
1.2 e-MERLIN	43
1.3 The e-MERLIN Cygnus OB2 Radio Survey	46
1.3.1 The Cygnus OB2 Association	46
1.3.2 The Role of Massive Star Winds in the Galaxy	48
1.3.3 COBRaS Scientific Analysis Goals	49
1.3.4 COBRaS Technical Analysis Goals	51
1.3.5 COBRaS Observations	52
1.3.6 Synergies with Multi-wavelength Studies of Cyg OB2	52
1.4 Thesis Aims	54

2 Radio Frequency Interference Mitigation with SERPent	56
2.1 RFI Mitigation Techniques	56
2.1.1 Pre-correlation RFI Mitigation	57
2.1.2 Post-correlation RFI Mitigation	58
2.1.3 Reduction and RFI Mitigation for e-MERLIN	58
2.2 SumThreshold Method	59
2.2.1 Statistical Variance Estimators	62
2.3 SERPent	63
2.3.1 Lovell Stationary Scan Removal	68
2.3.2 Zero-level Amplitude Dropouts Removal	70
2.3.3 RFI Mitigation Performance	72
2.3.4 Computational Performance	74
2.4 Conclusions and Discussions	81
2.4.1 Conclusions	81
2.4.2 Discussion	82
2.5 Future Work on SERPent	85
3 Radio Interferometric Calibration Pipeline for COBRaS	87
3.1 Calibration Coefficients	87
3.2 The Closure Relations	89
3.2.1 The Closure Phase	89
3.2.2 The Closure Amplitude	90
3.3 Outline of the Calibration Pipeline	91
3.4 Calibration Datasets	93
3.5 Preparation of Data for Calibration	93
3.6 Phase Calibration	94
3.7 Amplitude Calibration	100
3.8 Bandpass Calibration	106
3.9 Self-Calibration	110
3.10 Calibration Performance on COBRaS L-band Legacy Pointings	116
3.11 COBRaS Imaging and Mosaicing	119
3.11.1 Wide-band Imaging	121
3.11.2 Wide-field Imaging	124

3.11.3	Imaging Considerations with e-MERLIN	126
3.12	e-MERLIN Mosaicing Techniques	127
3.13	Discussion	128
4	Source Detection, Extraction and Classification	132
4.1	Building the Cyg OB2 Amalgamated Catalogue	132
4.1.1	Catalogue Amalgamation with Massey and Thompson (1991) Identifiers	135
4.1.2	Catalogue Amalgamation with Albacete Colombo et al. (2007b) Identifiers	136
4.1.3	Catalogue Amalgamation for Cyg OB2 Radial Velocity Survey	136
4.1.4	Catalogue Amalgamation with Sky Coordinates	136
4.1.5	The Cyg OB2 Amalgamated Catalogue	138
4.1.6	The OB Star Catalogue	138
4.1.7	The Candidate OB Star Catalogue	140
4.2	Source Detection Algorithms	143
4.2.1	Floodfill Algorithm	145
4.2.2	Modifications to the Floodfill Algorithm	145
4.3	Flux Extraction in the COBRaS Pipeline	147
4.3.1	Pixel-by-pixel Flux Extraction	147
4.3.2	Source Position Determination	148
4.3.3	Source Detection Performance	149
4.3.4	Flux Extraction Performance	153
4.3.5	Source Position Performance	161
4.4	Source Classification via Cross Referencing and Bayes' Theorem	163
4.4.1	Cross Referencing Maps and Catalogues	165
4.4.2	Significance Boosting with Bayes' Theorem	166
4.4.3	Empirical Simulations of the Significance Boost	168
4.5	Conclusions and Discussions	171
5	Radio Emission of Massive Stars in Cyg OB2	176
5.1	Massive Stars	176
5.2	The Winds of Massive Stars	178
5.2.1	Mass Loss	178

5.2.2	Radiative Line-Driven Wind Theory	180
5.3	Radio Emission from Massive Stars	184
5.3.1	Thermal Free-free Radio Emission Spectrum	184
5.3.2	Radio Emission from Single Massive Stars	187
5.3.3	Radio Emission from Massive Binaries	188
5.4	Clumping in the Winds of Massive Stars	191
5.4.1	The Effects of Clumping on Mass Loss Diagnostics	191
5.4.2	The Clumping Factor	194
5.4.3	Radio Free-Free Emission with the Clumping Factor	195
5.5	Predicted Smooth Mass Loss Rates and Radio Fluxes	195
5.5.1	Theoretical Mass Loss Relationships	195
5.5.2	Theoretical Stellar Parameters for Massive Stars	197
5.5.3	Smooth Mass Loss Rates and Radio Fluxes	201
5.5.4	Comparison of Smooth Mass Loss Rates and Fluxes with the Literature	207
5.6	Clumped Wind Radio Fluxes for Massive Stars in Cyg OB2	208
5.7	Current COBRaS L-band and C-band Maps	217
5.7.1	COBRaS C-band Commissioning Maps	217
5.7.2	COBRaS L-band Legacy Maps	218
5.7.3	Summary of Sources in COBRaS Maps	221
5.8	Conclusions and Discussions	223
5.9	Future Work with COBRaS	226
6	X-ray Variability of Massive stars in Cyg OB2	227
6.1	X-ray Emission from Massive Stars	228
6.1.1	X-ray Variability from Massive Binaries	228
6.1.2	Luminosity relations for O Stars in Cyg OB2	229
6.2	The Chandra Cygnus OB2 Legacy Survey	230
6.3	Introduction to Chandra	232
6.4	Chandra ACIS Data Analysis and Calibration	233
6.4.1	ACIS Photon Pileup	234
6.4.2	Chandra Observatory Response	234
6.4.3	Photometry with ACIS	235
6.5	Statistics and Variability Methods for X-ray Astronomy	236

6.5.1	Null Hypothesis	237
6.5.2	Chi-Squared Test	237
6.5.3	Kolmogorov-Smirnov Test	238
6.5.4	Kuiper's Test	240
6.5.5	Cash Statistic	240
6.5.6	Lomb-Scargle Periodograms	241
6.6	Chandra X-ray variability of the O star population of Cyg OB2	244
6.6.1	Variability in Previous Chandra Studies	244
6.6.2	Overview of Variable Candidates in the Chandra dataset	244
6.6.3	Inter-pointing X-ray Variability	247
6.6.4	Intra-pointing X-ray Variability	269
6.6.5	Summary of X-ray Variability	275
6.7	Conclusions and Discussions	277
6.8	Future Work	281
7	Summary, Conclusions and Future Work	282
7.1	Summary of Thesis	282
7.2	Thesis in the Context of COBRaS	284
7.3	Thesis in the Wider Context of Software Engineering	286
A	AIPS Nomenclature	288
B	Derivations	293
B.1	Derivation of Equation 1.8	293
B.2	Error Propagation Laws	294
B.3	Derivation of Mass Loss Rate Errors ($\Delta\dot{M}$)	295
B.4	Derivation of Flux Errors	296
Bibliography		299

This page was intentionally left almost blank

List of Figures

1.1	Simple two element interferometer	25
1.2	Signal path through a radio interferometer	26
1.3	Antenna optics systems	31
1.4	Correlator Architectures	35
1.5	u,v coverage of a single baseline from MERLIN and e-MERLIN	37
1.6	Map of the positions of the e-MERLIN antennas	44
1.7	MERLIN and e-MERLIN u, v coverage	45
1.8	Image of Cyg OB2 and the surrounding regions from Herschel	47
1.9	C-band and L-band mosaicing strategy for COBRaS	53
2.1	Example of time and frequency varying RFI from e-MERLIN commissioning data	60
2.2	A logic flow chart of the SERPent process	66
2.3	Before and after plots of visibility data with the Lovell and zero-level flagging passages in SERPent	69
2.4	Before and after plot of the zero-level flagging passage in SERPent	71
2.5	RFI in the COBRaS commissioning C-band (512 MHz bandwidth) data	73
2.6	SERPent flagging of RFI in the COBRaS commissioning C-band (512 MHz bandwidth) data	74
2.7	RFI in the COBRaS commissioning L-band (384 MHz bandwidth) data	75
2.8	SERPent flagging of RFI in the COBRaS commissioning L-band (384 MHz bandwidth) data	75
2.9	SERPent flagging of the time and frequency variant RFI in the e-MERLIN commissioning C-band data	76

2.10	Before and after example of the RFI flagging by SERPent	77
2.11	Parallelisation performance of SERPent	78
2.12	Memory performance of SERPent	80
3.1	Figure of the closure phase circuit	90
3.2	Logic diagram of the calibration procedure	92
3.3	Logic diagram of the phase calibration with AIPS tasks	96
3.4	Before and after plots of the phase calibration of J2007+4029	98
3.5	Another before and after plots of the phase calibration of J2007+4029	99
3.6	3C286 radio spectrum from 300 MHz to 50 GHz	100
3.7	Logic diagram of the amplitude calibration procedure	104
3.8	Before and after plots of the amplitude calibration of J2007+4029	105
3.9	Example of the uncalibrated e-MERLIN L-band bandpass response	106
3.10	Logic diagram of the bandpass calibration procedure	108
3.11	C-band Spectral Fit of OQ208	109
3.12	Spectrum of OQ208 from the MOJAVE survey	109
3.13	Before and after plots of the bandpass calibration of J2007+4029	111
3.14	Logic diagram of the self-calibration procedure	113
3.15	Before and after plots of the self-calibration of J2007+4029	114
3.16	C-band contour map of J2007+4029	115
3.17	Before and after plots of the L-band phase calibration of J2007+4029	117
3.18	Before and after plots of the L-band amplitude calibration of J2007+4029	118
3.19	L-band spectral fit of OQ208	119
3.20	Before and after plots of the L-band bandpass calibration of J2007+4029	120
3.21	L-band contour map of J2007+4029	121
3.22	Differences in u, v coverage and the dirty beam across the 512 MHz L-band	122
3.23	Diagram of the effects of wide-field imaging	126
3.24	Diagram of the method of faceting	127
3.25	e-MERLIN beam sensitivity areas	128
4.1	The Cyg OB2 Amalgamated Catalogue	139
4.2	The Known OB Catalogue	140
4.3	J–H versus H–K magnitude colour plot of star in Cyg OB2	141

4.4	J–H versus H–K magnitude colour plot of stars in the Cyg OB2 Amalgamated Catalogue and known OB catalogue	142
4.5	J–H versus H–K colour-colour plot of stars in the Cyg OB2 Amalgamated Catalogue, known OB catalogue and OB candidate catalogue	143
4.6	J–H versus H–K magnitude colour plot of stars in the 2MASS catalogue and the 2MASS OB candidate catalogue	144
4.7	Source detection results for point sources	152
4.8	Source detection results for resolved sources	153
4.9	14 point sources with $\text{SNR} = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100\}$	154
4.10	Performance of the pixel-by-pixel and JMFIT methods for extracting the integrated flux from simulated points sources	155
4.11	14 resolved sources with $\text{SNR} = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100\}$	157
4.12	Performance of the pixel-by-pixel and JMFIT methods for extracting the integrated flux from simulated resolved sources	158
4.13	Image of a simulated resolved source with no noise	159
4.14	Performance of the pixel-by-pixel and JMFIT methods for extracting the integrated flux from flux-corrected simulated resolved sources	160
4.15	Performance of PP methods and JMFIT positional offsets for point sources	162
4.16	Performance of PP methods and JMFIT positional offsets for resolved sources	164
4.17	Plot of two 2D Gaussian PDFs; one from the catalogue matching radius and one from source position extraction	166
4.18	Posterior probability as a function of the prior probability for three input SNRs	169
4.19	Figure of the prior probabilities needed to boost source detection sigma significance to the next integer level	171
4.20	Figure showing the prior probability required to boost a source detection sigma to the 5σ significance level	172
5.1	Hertzprung-Russell Diagram	178
5.2	Cartoon of radiative line-driving wind theory	181
5.3	Free-free emission (Bremsstrahlung) process with two oppositely charged particles	185

5.4	Massive star wind-wind collision region depicting the creation of non-thermal radiation	189
5.5	Discrepancies between UV P^{4+} and ρ^2 mass loss diagnostics	191
5.6	The formation regions of density squared free-free emission processes	193
5.7	Visual to radio fluxes demonstrating the clumping gradient in massive star winds	193
5.8	Smooth 1.5 GHz and 5 GHz (20 cm and 6 cm) fluxes of O stars and B Supergiants	202
5.9	Mass loss rates comparison with Muijres et al. (2012) methods	209
5.10	Radio flux comparison of nine Cyg OB2 members	211
5.11	Contour map of the first 300 μ Jy source in the C-band commissioning data field at $20^{\text{h}} 32^{\text{m}} 46.335^{\text{s}} +41^{\circ} 12' 51.070''$	218
5.12	Contour map of a 1.8 mJy source in the L-band pointing C field at $20^{\text{h}} 33^{\text{m}} 58.31452^{\text{s}} +41^{\circ} 09' 14.5002''$	219
5.13	Contour map of Cyg OB2 No. 9 in the L-band pointing D field at $20^{\text{h}} 33^{\text{m}} 10.735^{\text{s}} +41^{\circ} 15' 08.353''$	221
5.14	Contour map of a strong source in the L-band pointing D field centred at $20^{\text{h}} 32^{\text{m}} 56.795^{\text{s}} +41^{\circ} 08' 53.487''$	222
5.15	High resolution contour map of a strong source in the L-band pointing D field centred at $20^{\text{h}} 32^{\text{m}} 56.798^{\text{s}} +41^{\circ} 08' 53.392''$	223
6.1	The Chandra Cygnus OB2 survey's deep mosaic coverage of the Cygnus OB2 region	230
6.2	X-ray source distribution over the field of view of the Chandra Cyg OB2 Legacy survey	231
6.3	Image of the Chandra X-ray Observatory	233
6.4	A cumulative distribution plot of the individual photon events from the observation of MT91 534, epoch 10960	239
6.5	Binned light curve of Schulte 3	248
6.6	Binned light curve of Schulte 4	249
6.7	Binned light curve of Schulte 5	250
6.8	Lomb-Scargle periodogram of Schulte 5	251
6.9	Cumulative histogram of the MC simulations for Schulte 5	252

6.10	Histogram of the strongest periods found from each MC simulation for Schulte 5	253
6.11	Light curve of Schulte 5 folded onto the period of 6.64 days	254
6.12	Binned light curve of MT91 267	255
6.13	Lomb-Scargle periodogram of MT91 267	256
6.14	Cumulative histogram of the MC simulations for MT91 267	257
6.15	Histogram of the strongest periods found from each MC simulation for MT91 267	258
6.16	Light curve of MT91 267 folded onto the period of 15.95 days	258
6.17	Binned light curve of Schulte 22	259
6.18	Lomb-Scargle periodogram of Schulte 22	260
6.19	Cumulative histogram of the MC simulations for Schulte 22	261
6.20	Light curve of Schulte 22 folded onto the reported period of 4.61 days	261
6.21	Binned light curve of Schulte 9	262
6.22	Binned light curve of Schulte 8A	263
6.23	Binned light curve of MT91 516	264
6.24	Binned light curve of Schulte 11	265
6.25	Binned light curve of Schulte 75	266
6.26	Binned light curve of Schulte 73	267
6.27	Binned light curve of MT91 771	268
6.28	Binned light curve of WR 146	269
6.29	Binned light curve of Schulte 3	270
6.30	Binned light curve of MT91 267; epoch 4511	271
6.31	Binned light curve of WR 145	272
6.32	Binned light curve of WR 145 with regressional fit	273
6.33	Binned light curve of MT91 534 showing a 16 ks flare	274
6.34	Binned light curve of MT91 534 with sigma levels	275
6.35	600s binned light curves for 20 sources from Albacete Colombo et al. (2007a)	276

This page was intentionally left almost blank

List of Tables

1.1	Technical Capabilities of e-MERLIN	46
1.2	Physical Properties of Cyg OB2	48
2.1	SERPent Performance Test Datasets	67
2.2	Computer Systems	76
3.1	Calibration Datasets	93
3.2	Approximate Coherence Times for e-MERLIN	95
3.3	Coefficients for the Spectral Flux Density of 3C286	101
3.4	Shortest e-MERLIN Baselines	102
4.1	Table of existing Cyg OB2 catalogues	134
4.2	Table of Source Detection Results from Simulations	153
4.3	Table of boosting threshold probabilities	170
5.1	Stellar Parameters for O stars from Martins et al. (2005)	198
5.2	Stellar Parameters for B Supergiants from Searle et al. (2008)	199
5.3	Terminal Velocities for O stars and B Supergiants from Prinja et al. (1990)	200
5.4	Full Table of Smooth Mass Loss Rates and Radio Fluxes for O Supergiants and Giants	204
5.5	Full Table of Smooth Mass Loss Rates and Radio Fluxes for O Dwarfs	205
5.6	Full Table of Smooth Mass Loss Rates and Radio Fluxes for B Supergiants	206
5.7	Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for O Supergiants and Giants	212
5.8	Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for O Dwarfs	213

5.9	Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for B Supergiants	214
5.10	Sources in the COBRA-S Maps	222
6.1	Table of X-ray Variable O-type and WR Candidates	246
6.2	Table of Epoch Timescales for Variable Candidates	247
6.3	Summary of the X-ray Variability in this Study	277

Chapter 1

Introduction

“In the beginning the Universe was created. This has made a lot of people very angry and has been widely regarded as a bad move.”

Douglas Adams

This thesis presents the technical development for the Cygnus OB2 Radio Survey (CO-BRaS; P.I. Prof Raman Prinja, UCL), the deepest and most sensitive radio survey of an OB association to date. As one of the Legacy projects with the expanded Multi Elemental Radio Linked Interferometry Network (e-MERLIN), COBRaS exploits the recent upgrades in receiver technology (bandwidth) to produce a substantial dataset on the Cyg OB2 region. The key aim of this thesis is to process this large volume of data, automatically editing and calibrating the interferometric datasets, and analysing the resulting radio maps. These technical goals are essential in order to achieve the full scientific goals of CO-BRaS. The primary subject of COBRaS is to analyse the radio properties of massive stars within the core region of Cyg OB2, with particular emphasis on mass loss studies and the effects of clumping structure throughout the winds of massive stars. In addition to these technical challenges, this thesis contains a substantial amount of scientific preparation in the form of the amalgamation of previous surveys into one definitive archive catalogue and producing OB candidate catalogues. The predicted mass loss rates and radio fluxes for smooth and clumped winds for O-type stars and B supergiants are presented. These

are compared with observed radio fluxes of massive stars in Cyg OB2, including the first radio maps from COBRaS. This thesis also explores the X-ray variability properties of O stars with the Chandra Cyg OB2 Legacy survey, giving a multi-wavelength perspective into the massive star population of Cyg OB2.

This Chapter introduces the theory of radio interferometry (Section 1.1), a description of the e-MERLIN radio telescope (Section 1.2), an overview of the COBRaS project (Section 1.3) and finally a restatement of the thesis aims with a summary of the structure of the thesis (Section 1.4).

1.1 Introduction to Radio Interferometry

Scientific advancement is more than just the pursuit of knowledge. More often than not it is limited by technology, and therefore drives the need for technological innovations, which can have applications outside the original research field. In the case of radio astronomy, interferometry (the method of combining signal from multiple telescopes) has furthered our understanding of the radio universe by dramatically improving the available angular resolution of any observations, previously limited by the diameter of single radio dishes. The first radio observations were conducted by Karl Jansky (after whom the units for radio flux are named) in the 1930s, who discovered a radio source fixed on the sidereal day (Jansky 1933). This source was later designated Sagittarius A, one of the brightest sources of radio waves in the sky, emitted by electrons moving in a strong magnetic field.

The first radio interferometric observations came later in 1946 in Australia, using an old World War II radar antenna on a sea cliff. By observing the direct radiation from the source and the reflections from the sea, an extra path length, or phase change was introduced into the system resulting in clear maxima and minima interference fringes (McCready et al. 1947). The observations showed daily intensity variations over six months correlating with the incidence of sunspots and also rapid fluctuations on the order of seconds to minutes in duration.

The technique of *phase switching* was developed by Sir Martin Ryle in 1952 and was a way to combine signals for early interferometers. If the phase of a signal is periodically switched then the difference between these signals is proportional to the time average of the cross-correlation between two antennas (Ryle 1952). This is equivalent to the combination of the sine and cosine correlators in modern interferometers, see Section 1.1.3 for details.

The technique of aperture synthesis, pivotal to radio interferometry imaging with multiple antennas, was developed and applied to the Cambridge one-mile telescope (Elsmore et al. 1966) and signalled a new era of interferometry for both connected interferometers such as the Very Large Array (VLA), the Multi Elemental Radio Linked Interferometer Network (MERLIN), and for non-connected Very Long Baseline Interferometry (VLBI) such as the Very Long Baseline Array (VLBA) and the European VLBI Network (EVN). Deconvolution methods deployed in aperture synthesis such as the CLEAN algorithm (Högbom 1974; Clark 1980) and the Maximum Entropy Method (MEM) (Narayan and Nityananda 1986) have been applicable to medical imaging techniques (Simpson et al. 1995; Jannetta et al. 2004; Jackson and Jannetta 2006).

Interferometry allows milli-arcsecond astronomy at radio frequencies and allows even more accurate astrometric measurements. It is possible to combine the observations from multiple interferometric arrays together to increase the resolution and sensitivity on different physical scales by covering different baseline lengths with the different arrays.

The current generation of interferometers is benefiting from numerous upgrades to receivers, electronics, correlators and optical fibre networks, increasing the available frequency bandwidth and thus sensitivity. Recent upgrades to the VLA, now named the Jansky Very Large Array (JVLA), MERLIN → e-MERLIN and new interferometers such as the Low Frequency Array (LOFAR) and the Atacama Large Millimeter/submillimeter Array (ALMA) to name a few, now provide the most sensitive windows into the radio spectrum (from MHz to THz frequencies) pushing the sensitivity to μJy levels.

Along with new hardware such as wide bandwidth receivers, software development is necessary to exploit these technological advancements, either in the form of modifications to existing programs or entirely new software packages or pipelines. This is particularly true for systems managing large datasets, i.e. several TBs, where the creation of automated algorithms is a necessity and large amounts of manual intervention is unfeasible.

The future will contain ever bigger and more sensitive arrays with the Square Kilometre Array (SKA), being arguably the biggest experiment ever undertaken in sheer size, computational requirements and science implications. It has been confirmed that the SKA will be split between South Africa and Australia, with the mid-frequency dishes (500 MHz - 3 GHz) planned for the South African site and the low frequency (\sim 50 - 300 MHz) phased arrays situated at the Australian site (Taylor 2013). The SKA will provide an increase in sensitivity by many orders of magnitude, observing the radio universe at the \sim nJy level.

An introduction to radio interferometric theory is presented here and some of the current features of the current crop of interferometers and their recent upgrades are discussed. Firstly, a simplified example using the two element interferometer providing a brief overview of the theory is given. This will be followed by more detailed sections on the separate parts and methods of the interferometer (Antenna and Receiver Systems, Complex Correlators), the visibility plane, radio interferometric imaging via deconvolution theory and polarisation effects.

1.1.1 The Two Element Interferometer

Interferometry is the process of coherently correlating two or more electric fields measured at different locations. To achieve this, a collecting area or ‘sensor’ (antenna) is used at each location to collect the signal and convert it to a voltage, which is then correlated with other voltages from other antennas. The quantities which the interferometer retains are the amplitude and the phase part of the electromagnetic wave.

The astronomical source is assumed to be located at an infinite distance from the interferometer, thus the wavefront will be in the Fraunhofer regime, i.e. flat. This plane parallel approximation is a fair assumption given the vast, albeit finite distances to sources. Thus the electric field component of a source emitting a monochromatic electromagnetic planar wave is given by the real part of

$$E(t) = \text{Re} \{ A(t) e^{2\pi i \nu t} \}, \quad (1.1)$$

i.e. using Euler’s formula $e^{i\phi} = \cos(\phi) + i \sin(\phi)$, where the real part is the cosine term and the imaginary part is the sine term, this reduces to

$$E(t) = A(t) \cos(2\pi\nu t). \quad (1.2)$$

This wave arrives at two elements separated by a baseline \mathbf{b} , from the direction of the source which is denoted by the unit vector \mathbf{s} . The geometry of this simple interferometer is shown in Figure 1.1. Each individual antenna within an array measures the incoming wave as a signal voltage. For any given wavefront there are two voltages (V_i and V_j), a time of arrival t at antenna j , giving a distance $d = c\tau$, where c is the speed of electromagnetic radiation in a vacuum, and τ is the geometric delay for a wavefront between the two

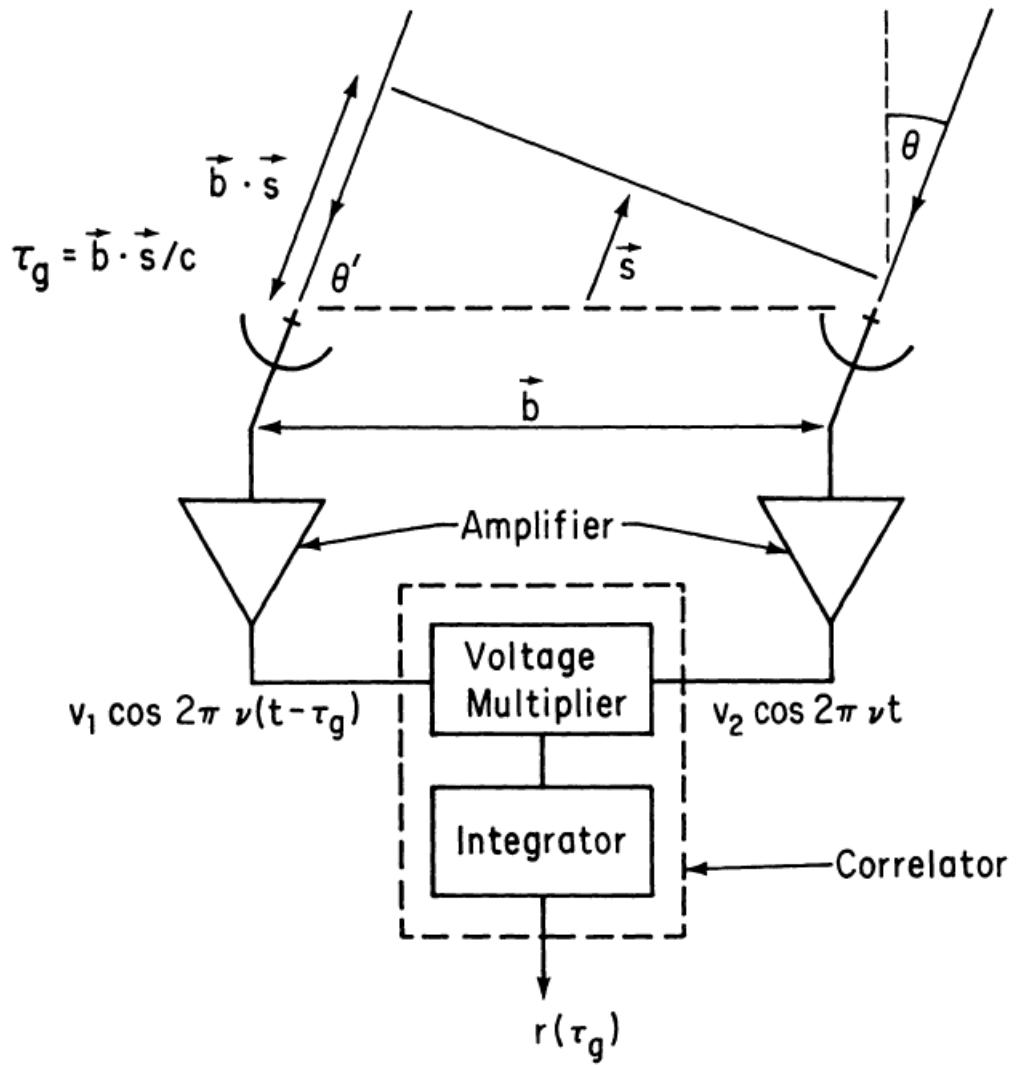


Figure 1.1: Simple two element interferometer diagram, including baseline geometry with respect to the incoming wave and a simplified electronics and receiver systems. Figure from Thompson (1999).

antennas. From inspection we find that $d = \mathbf{b} \cdot \mathbf{s}$ (or $d = \mathbf{b} \cos \theta'$, in trigonometric form where $\theta' = \pi/2 - \theta$). After some substitution we arrive at:

$$\tau_g = \frac{\mathbf{b} \cdot \mathbf{s}}{c} \quad (1.3)$$

which shows the geometric time delay and gives the position on the sky.

The recorded voltage contains both the amplitude and phase information of the wave, and is passed through a series of systems as depicted in Figure 1.2. This modifies the signal in ways which will be discussed in greater detail later (see Antenna and Receiver Systems, Section 1.1.2, and Complex Correlators, Section 1.1.3).

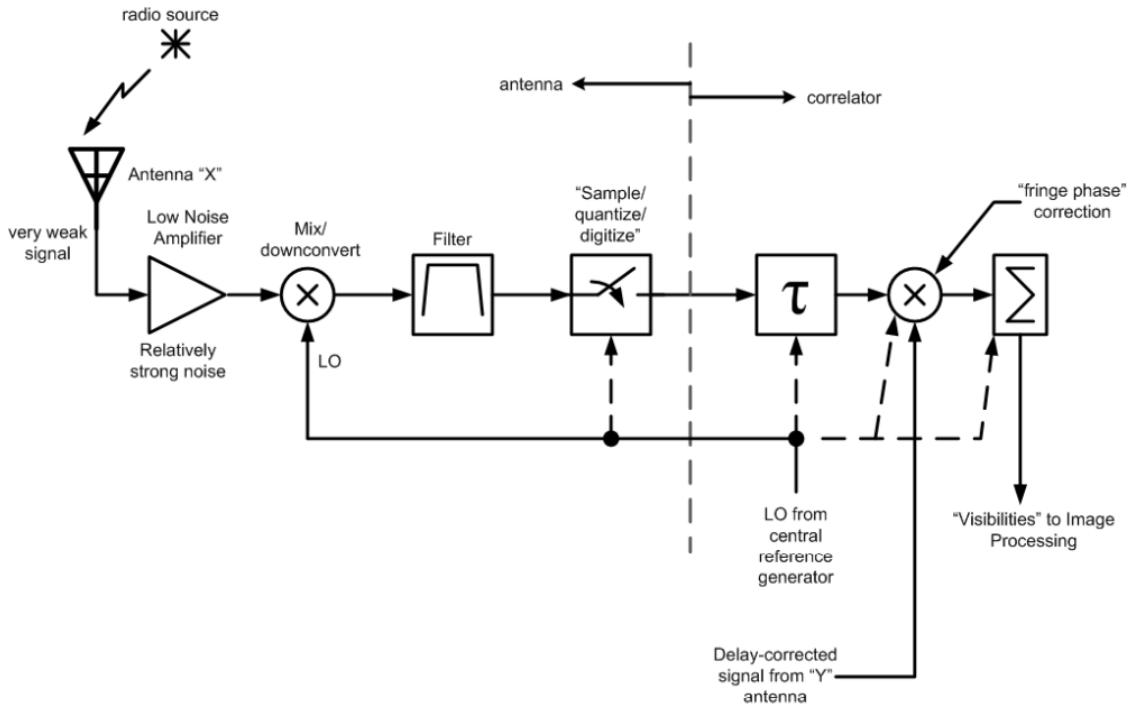


Figure 1.2: Signal path from source to visibilities via the antenna, correlator and their subsystems. Figure taken from Carlson (2012).

Essentially there are two voltages from each antenna in this system:

$$V_i(t) = v_i \cos(2\pi\nu(t - \tau_g)) \text{ and } V_j(t) = v_j \cos(2\pi\nu t) \text{ which arise from Equation 1.2.}$$

These pass through a correlator which multiplies and time averages (or integrates; denoted by the angled brackets) the two signals to produce the correlator output $r_c(\tau_g)$:

$$\begin{aligned}
 r_c(\tau_g) &= \langle V_i(t) V_j(t) \rangle \\
 &= \langle v_i v_j \cos(2\pi\nu(t - \tau_g)) \cos(2\pi\nu t) \rangle \\
 &= \left\langle \frac{v_i v_j}{2} [\cos(2\pi\nu(t - \tau_g) - 2\pi\nu t) + \cos(2\pi\nu(t - \tau_g) + 2\pi\nu t)] \right\rangle \\
 &= \left\langle \frac{v_i v_j}{2} [\cos(2\pi\nu\tau_g) + \cos(4\pi\nu t + 2\pi\nu\tau_g)] \right\rangle \\
 &= v_i v_j \cos(2\pi\nu\tau_g), \tag{1.4}
 \end{aligned}$$

where the product-sum cosine rule: $\cos \theta \cos \phi = \cos(\theta - \phi) + \cos(\theta + \phi)/2$ is used, and the rapidly varying term $\cos(4\pi\nu t + 2\pi\nu\tau_g)$ reduces to zero.

The geometric delay varies as the Earth rotates, thus the wave arrival orientation with respect to the different baselines will also change due to the cosine term in Equation 1.4. Since the interferometer measures the interference fringe pattern of the source, the array tracks the source as it passes through each fringe, where each fringe size is defined by

$\sim \lambda/\mathbf{b}$. The cosine term represents the fringe phase and the $v_i v_j$ term represents the fringe amplitude which is proportional to the received power¹.

The sky brightness $B(\nu, \mathbf{s})$ is defined as the intensity in the direction of the source (i.e. the unit vector \mathbf{s}) for some frequency ν . Since in this example, monochromatic radiation is assumed, the frequency dependence of the sky brightness is neglected, i.e. the sky brightness is a function of intensity in the direction of the source $B(\mathbf{s})$ only. The power received by the correlator over the entire sky surface (4π steradians) is found by integrating the correlator response over the solid angle $d\Omega$. Thus the output of the correlator $r(\mathbf{b})$ in terms of the sky brightness and baseline geometry is given by

$$r(\mathbf{b}) = \int_{\Omega} B(\mathbf{s}) \cos \frac{2\pi\nu\mathbf{b} \cdot \mathbf{s}}{c} d\Omega \quad (1.5)$$

where the sky brightness represents the power received i.e. $v_i v_j$ in Equation 1.4, and the cosine term has been converted as a function of baseline position by substituting Equation 1.3 for τ_g .

As can be seen in Equation 1.4; the correlator's response follows a cosinusoidal pattern which is only sensitive to the even part of the sky brightness. To reveal the odd part, the correlator must be responsive to fringes with a phase difference of $\pi/2$. This is achieved by adding in a phase shift to the signal from one of the antennas in a baseline before cross correlation occurs. This changes the signal's phase but leaves the amplitude unaffected. Thus the modified signal at antenna j would be

$$V_j(t) = v_j \cos(2\pi\nu t - (\pi/2)). \quad (1.6)$$

This modifies the input signal into the correlator to $V_j(t) = v_j \sin(2\pi\nu t)$, as $\cos(\phi - \pi/2) =$

¹This depends on the bandwidth of the observation as well as the efficiency of the antenna elements, affecting subsequent equations and will be explored in Section 1.1.2.

$\sin(\phi)$. Thus Equation 1.7 shows the output of the sine correlator:

$$\begin{aligned}
 r_s(\tau_g) &= \langle V_i(t) V_j(t) \rangle \\
 &= \langle v_i v_j \cos(2\pi\nu(t - \tau_g)) \sin(2\pi\nu t) \rangle \\
 &= \left\langle \frac{v_i v_j}{2} [\sin(2\pi\nu t + 2\pi\nu(t - \tau_g)) - \sin(2\pi\nu t - 2\pi\nu(t - \tau_g))] \right\rangle \\
 &= \left\langle \frac{v_i v_j}{2} [\sin(4\pi\nu t - 2\pi\nu\tau_g) - \sin(2\pi\nu\tau_g)] \right\rangle \\
 &= v_i v_j \sin(2\pi\nu\tau_g)
 \end{aligned} \tag{1.7}$$

where the product-sum rule: $\sin \theta \cos \phi = \sin(\theta + \phi) - \sin(\theta - \phi)/2$ is used, and the rapidly varying term $\sin(4\pi\nu t - 2\pi\nu\tau_g)$ reduces to zero. Equation 1.7 has the same response as Equation 1.4, except that the cosine term is switched for a sine term.

A correlator which produces both the even and odd part of the sky brightness response is called a *complex* correlator. These terms can be combined to produce the full response of the interferometer, which we define as the complex visibility $V(\mathbf{b})$, and after some manipulation², becomes:

$$V(u, v, w) = \int_l \int_m \int_n B(l, m, n) e^{-2\pi i(lu + mv + nw)} \frac{dl dm}{dn} \tag{1.8}$$

where u, v, w are baseline coordinates, l, m, n are the corresponding coordinate system for the sky and are the directional cosines with respect to the baseline vectors, and $B(l, m, n)$ is the sky brightness.

It is important to stress that this is not the whole picture as we have assumed monochromatic radiation, and in practice a number of coefficients are included to account for effects from the antenna response and bandwidth which will modify the interferometer's response (see Chapter 3 on calibration). Furthermore, the full response of the interferometer given by Equation 1.8 describes the raw measurements of the interferometer and requires further manipulation to calibrate issues such as phase delays from atmospheric distortion, gain (amplitude and phase) variations in time, and a non-flat bandpass response causing complex gain variations in frequency. Chapter 3 discusses all of these issues and how they are solved in detail. However, for demonstration purposes they have been neglected for this simplified example.

²See Appendix B for details.

1.1.2 Radio Antenna and Receivers

The main advantage of interferometry compared to single dish radio observations is the vast increase in angular resolution. The angular resolution θ of single dish radio astronomy is diffraction limited by the diameter D of the antenna for a wavelength λ ; $\theta \approx \frac{\lambda}{D}$. Interferometry is a system of two or more telescopes separated by a distance B , or ‘baseline’, which observe the same source. The length of a baseline is measured in units of wavelengths of the central frequency of the observation. Thus the angular resolution at a certain wavelength λ is now limited by the longest baseline of an interferometric array; $\theta \approx \frac{\lambda}{B}$.

Before any interferometric properties are presented or discussed, we must treat each component of the array as an individual system, i.e. for a correlated interferometer each antenna is, to begin with, separate. Firstly the type of antenna and receiver is to be considered, and the defining factor on the shape or nature of the radio antenna is the frequency of the observation. There is a general division between two main types of antenna design (although there are some overlaps). For observations with wavelengths shorter than 1 m (high frequency > 300 MHz), a parabolic reflector dish is favoured, because a mesh reflector will not have sufficient spacing to reflect the wave. This occurs when the spacing exceeds a fraction of the wave, typically on the order of $\lambda/20$. For observations with wavelengths longer than 1 m (low frequency < 300 MHz) the antennas can utilize a cheaper wire mesh design with ‘gaps’ in the collecting surface because the required surface accuracy is less relative to higher frequency requirements.

e-MERLIN and other higher frequency (> 300 MHz) arrays use reflector antennas which can vary in diameter from metres to 100 m for fully steerable dishes. There are two considerations to be made when choosing a certain reflector design: the type of mount and the optics system.

There are two types of mounts; equatorial mounts and elevation over azimuth (alt-azimuth) mounts. The alt-azimuth mount is simpler to calibrate for as the effects of gravity (which affects the aperture shape) are constant throughout the observation. However, since the aperture rotates with respect to the source, the beam profile will rotate. If the beam shape is elliptical this may cause the brightness distribution to smear over time.

Equatorial mounts have the advantage of being aligned to the polar axis, parallel to the rotation of the Earth and thus are able to track the target source. Hence, equatorial

mounts do not suffer from beam rotation. The big disadvantage is that gravity does not act in the same observation plane, and calibrating for these effects can be complicated.

Figure 1.3 shows a number of optical systems with different orientations of the telescope feed and receivers, each with advantages and disadvantages. The prime focus system ((a) in Figure 1.3) is implemented at Westerbork Radio Synthesis Telescope (WRST) and the Lovell, Mark II and Defford antennas in e-MERLIN (Garrington et al. 2004). This arrangement allows full frequency coverage over the primary reflector whereas the inclusion of a secondary reflector in the other systems limits the coverage by the practical size of the secondary. The disadvantages of this setup are the restricted space and access to the feed and receiver and the noise reflected off the ground.

The other e-MERLIN antennas, Darnhall, Pickmere, Knockin (referred to as the E-systems) and Cambridge are Cassegrain in design with an off-axis feed carousel which rotates the feeds into the focal plane. The advantage of the Cassegrain design is that the feeds and receivers are easy to access, there is no noise from ground spillover and the reflectors can be shaped to provide more illumination (Napier 1999). From here onwards, the term ‘antenna’ refers to reflector telescopes associated with high frequency (> 300 MHz) observations as applicable with the various e-MERLIN antennas.

The sensitivity of the individual antenna is defined by its effective collecting area $A(\nu, \theta, \phi)$ m² where ν the frequency, and θ and ϕ are the directional coordinates. The overall efficiency of this aperture depends on a number of factors. Imperfections in the antenna surface cause phase errors at the feed and thus a decrease in power, represented by the surface efficiency η_s . Any secondary reflector and its supports physically blocking radiation from reaching the primary reflector or feed are described by the blockage efficiency η_b . The feed spillover efficiency η_{spill} accounts for power lost when radiation is directed from the reflector (primary or secondary depending on antenna design) to the feed. The illumination taper efficiency η_t represents the discrepancy in illumination between the outer and inner parts of the reflector (Napier 1999).

The antenna performance is thus modified to account for these aperture inefficiencies to produce the relation between on-axis response (response at the centre of the main lobe of $A(\nu, \theta, \phi)$) A_0 , the physical area of the aperture A and the aperture efficiency η (where $\eta = \eta_s \eta_b \eta_{\text{spill}} \eta_t$):

$$A_0 = \eta A. \quad (1.9)$$

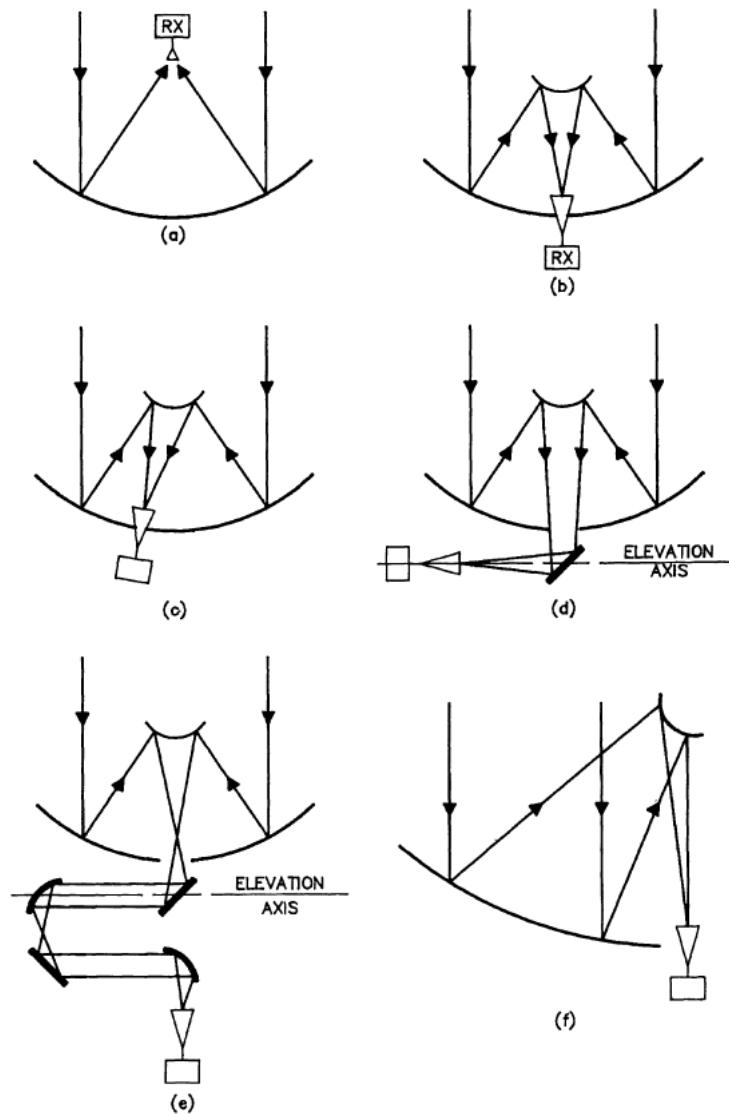


Figure 1.3: Antenna optics systems: (a) Prime focus, (b) Cassegrain, (c) Off-axis Cassegrain, (d) Naysmith, (e) Beam waveguide, (f) Offset Cassegrain. Figure from Napier (1999).

The power received from a source follows a diffraction pattern where the maximum response is at the pointing centre (peak of the beam). If a source were positioned at an angle θ from the pointing centre, the power is reduced as a function of this angle, $A(\theta)$ of the main lobe. The response at the pointing centre is $A(\theta = 0) = 1$ and almost no radiation is received at $\theta \gg \lambda/D$, where D is the diameter of the antenna. Combining the effective collecting area $A(\nu, \theta, \phi)$ with the response at the pointing centre A_0 gives an equation for the normalized pattern or primary beam

$$A_{\text{primary}}(\nu, \theta, \phi) = \frac{A(\nu, \theta, \phi)}{A_0} \quad (1.10)$$

where $A_{\text{primary}}(\nu, \theta, \phi)$ is known as the primary beam and is essentially the point spread function (PSF) for the antenna. Thus the true power received by the antenna includes this PSF or primary beam term in the form

$$P(\nu, \theta, \phi) = A_{\text{primary}}(\nu, \theta, \phi) B(\nu, \theta, \phi) \Delta\nu \Delta\Omega \quad (1.11)$$

where $P(\nu, \theta, \phi)$ is the power received (in Watts), $B(\nu, \theta, \phi)$ is the sky brightness, $\Delta\nu$ is the bandwidth and $\Delta\Omega$ is the solid angle on the sky. The primary beam response correction can be applied as another calibration component which is discussed in Chapter 3.

The waves focused from either the primary reflector or sub reflector (depending on antenna design) are directed to the antenna feed, where a splitter separates the two orthogonal polarisation components which are special cases of elliptical polarisation. Polarisation effects in interferometry are discussed in Section 1.1.6.

After the electromagnetic radiation has passed through the antenna element and the necessary feed for the observed wavelength, the signal is very weak. A low noise amplifier increases the signal strength before any subsequent electronics or processes add additional noise to the signal. The signal voltage is then fed into an intermediate frequency (IF) mix converter where it is multiplied with a local oscillator (LO) frequency by synthesizers to a lower frequency range (also known as ‘down-conversion’)³. The amplitude modulated signal is symmetrical about this local oscillator frequency, displaying a mirror signal on

³This frequency range can in some cases be higher than the observed frequency range. It will be translated to a sufficiently high frequency as to retain a wide bandwidth and so the modulation has no negative frequency components, but low enough for the electronic signals to be easier to handle.

either side. For single-sideband mixers where the observed frequencies are $\nu < 100$ GHz, one of the sidebands is removed. This is to ensure that the signals from two different antennas are at the same frequency before cross correlation, as well as increasing the output power of the signal. See Napier et al. (1983) for a more detailed description of the amplifiers and converters used at the VLA.

At this stage the signal is transferred to a quantizer (also known as the digitizer) and is sampled using Nyquist sampling theorem i.e. sampled at twice the signal bandwidth to obtain all the information (Carlson 2012). The Nyquist theorem states that a function with no frequencies above ν can be determined by sampling at $N_{\text{samples}} < 1/2\nu$, or twice a frequency cycle. As the signal is quantized there is a loss in sensitivity due to the digital representation of the signal deviating from the true signal shape. Increasing the sampling to twice the Nyquist rate will result in greater sensitivity, however it may be more prudent to represent each sample by *2-bits*⁴. A *2-bit* quantisation scheme allocates one of the *bits* to represent the sign of the voltage (0 for positive and 1 for negative), and the other bit is assigned 0, if the input voltage is between two chosen levels $\pm V_0$ and 1 if the voltage is outside the transition levels (Cooper 1970). Every eventuality is designated a state with a corresponding weighting factor which is saved for that sample of the signal.

Following quantisation, the signal is passed to the correlator, where the voltages from two antennas are multiplied together to produce a range of *bit* products shown in Table 2 in Cooper (1970). The rest of the process (shown in Figure 1.2) is performed by the cross correlator and will be covered in Section 1.1.3 on correlators.

So far only the antenna response (efficiency) has been discussed, and described by Equation 1.9, which defines the sensitivity of the antenna. However, this does not include the noise contributes from the electronics and other parts of the receiver system. Thermal noise present in wires or cables and the amplifier etc, will lower the sensitivity of the receiver system by decreasing the signal-to-noise ratio, thus the system is cryogenically cooled to reduce this effect. The final sensitivity of the antenna and receiver system (commonly known as the Radiometer equation) is defined by

$$\Delta T = \frac{T_{\text{sys}}}{\sqrt{B \tau}} = \frac{T_{\text{source}} + T_{\text{antenna}} + T_{\text{Rx}}}{\sqrt{B \tau}} \quad (1.12)$$

where ΔT is the sensitivity of the receiver, T_{source} is the astronomical noise from the

⁴Modern correlators e.g. WIDAR at e-MERLIN and JVLA and the correlator at ALMA, have *3-bit* to *8-bit* correlator modes available.

source, T_{antenna} is the noise contribution from the antenna, e.g. antenna shape, spillover etc (described previously), T_{Rx} is the noise contribution from the receivers, B is the bandwidth of the observation, and τ is the integration time (total time on source). The sensitivity of any radio map σ_s is also a function of the antenna efficiency A_0 (Equation 1.9) and the number of antennas in the array N , and is given by

$$\begin{aligned}\sigma_s &= \frac{\sqrt{2} k \Delta T}{A_0 \sqrt{\frac{N(N-1)}{2}}} \\ &= \frac{\sqrt{2} k T_{\text{sys}}}{A \eta \sqrt{\frac{N(N-1)}{2}} B \tau},\end{aligned}\quad (1.13)$$

where k is the Boltzmann constant. Typical values for the variables in Equation 1.13 for e-MERLIN are: $\eta = 0.7$, $T_{\text{sys}} = 50$ Kelvin (K), $N = 7$, and $A = \pi (r/2)^2$ where r is the antenna diameter. For e-MERLIN this area is the average antenna area for the array $\sim 1113 \text{ m}^2$.

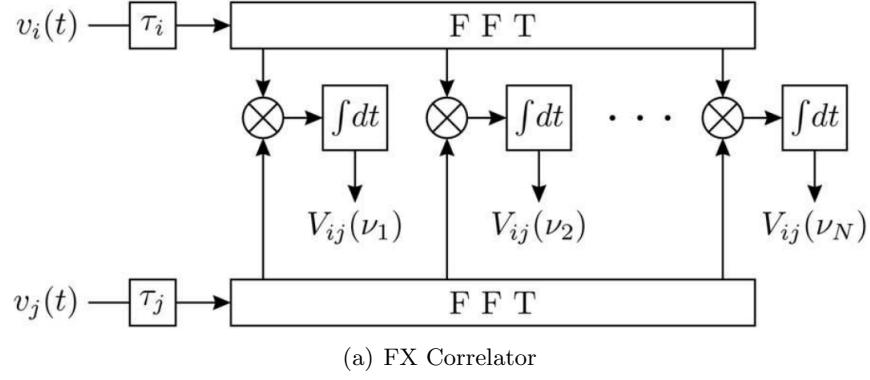
1.1.3 Complex Correlators

The principle task of the correlator is to combine the input voltage signals from different antennas in an interferometric array, for all possible permutations, to produce a coherent interference fringe or *complex visibility function*. This is broken down into multiple steps, where the correlator may include parts of the receiver system such as the quantizer. Since this operation acts purely on the voltage from a single antenna, it is discussed separately in Section 1.1.2. This section will focus on correlation of the output sampled signals.

For connected interferometers the voltage from all antennas is routed directly to the correlator, whereas the signal from non-connected (VLBI) arrays is recorded onto external harddrives (or tapes in times of old) and transported to the correlator site of that particular VLBI network. However, some VLBI arrays such as e-VLBI and e-MERLIN are connected via an optical fibre network. The exact processes the correlator performs depends on the correlator design, architecture and setup which will vary with each system.

The first correction is conducted on the antenna-dependent wavefront errors, i.e. the tracking delay as described previously. A digital delay is introduced to the voltage signal, along with the LO reference frequency used at the IF mixer. Because the LO reference frequency differs from the frequency detected at the antenna, also causing differences to the digital delay, a fringe phase correction is introduced to compensate (Carlson 2012). At

this fringe phase correction, the signals from other antennas are cross-correlated together twice, once for the even part of the sky and once for the odd part, or previously described as the cosine and sine correlations. The signals are then summed to produce the visibility function or interferometric equation (Equation 1.8).



(a) FX Correlator

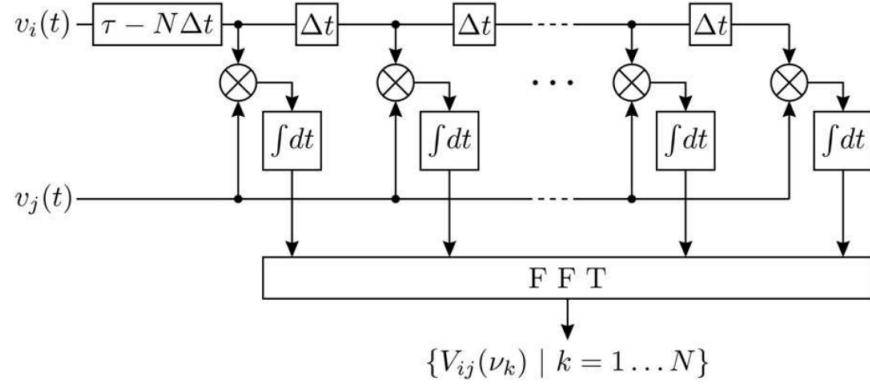


Figure 1.4: FX correlator (a) and XF correlator (b) flow diagrams. See text for full description of both correlator architectures. Figures taken from Brisken (2004).

The complex correlators have a number of architectures, or methods to create visibilities from the individual antenna signals. The two most common are FX and XF correlators, where F denotes the Fast Fourier Transform (FFT) and X denotes cross-correlation. FX correlators perform a FFT on the signal after the digital delay and before the cross correlation. A spectrum is available before integration (summation). The XF correlator cross correlates before with different delays for different frequencies (and antennas), integrates and then performs a FFT. Both architectures are shown in Figure 1.4.

The Wideband Interferometric Digital Architecture (WIDAR) correlator implemented at JVLA, ALMA and e-MERLIN, is a hybrid of the two architectures described above (FXF). This correlator contains a filter which first divides the signal into sub-bands or IFs

(F) and then performs the same procedure as the XF correlator architecture.

1.1.4 Visibilities and the Visibility Plane (Fourier Space)

The complex visibility is the final output from the correlator, and represents the interferometer's response to the sky brightness. There is a visibility for every baseline configuration, frequency channel in the bandwidth grouped into intermediate frequencies (IFs) or sub-bands, and integration time (the smallest sample rate of the correlator), which is usually on the order of \sim one second. Each visibility is approximately equal to one Fourier component of the observed sky brightness, a fact demonstrated by Equation 1.8. The baseline vectors (u, v, w) comprise the coordinate system of the interferometer and represent directions to the East, the North and the phase tracking centre (where the antenna are pointing directly at) respectively. The corresponding coordinate system for the sky is (l, m, n) , these are the directional cosines with respect to the baseline vectors (u, v, w) .

Visualisation of the baseline geometry is possible by plotting the visibilities as a function of their baseline vectors \mathbf{u} and \mathbf{v} to form what is commonly known as the u, v plane or u, v coverage. A visibility exists for every integration time (from the correlator) and for every channel in the bandwidth receiver (i.e. the smallest frequency increment of the observation). Figure 1.5 shows the u, v coverage of a single baseline from MERLIN (a) and the u, v coverage of a single baseline from e-MERLIN (b). The short periodical gaps in the ellipse arc are due to the interferometer observing a phase calibrator source for a technique called *phase referencing* which will be discussed in Chapter 3.

The rotation of the Earth continually changes the orientation of every baseline, and thus the visibility and the corresponding Fourier component change in time. The baseline projection traces out an ellipse in the u, v plane, as the source crosses the sky and the baseline geometry changes. This essentially maps out a different part of the source structure and thus fills in a different part of the u, v plane which is the Fourier transform of the sky brightness. The ellipse arc length is then simply the total integration time of the observation.

For the assumption concerning monochromatic radiation, the arc thickness (or visibilities in the radial direction from the origin) would be one visibility thick. Interferometers have a bandwidth of finite size divided into IFs which are comprised of a number of channels also of finite frequency. Therefore for every integration time a visibility exists for each channel within the full bandwidth. As the u, v plane is in units of wavelength, and

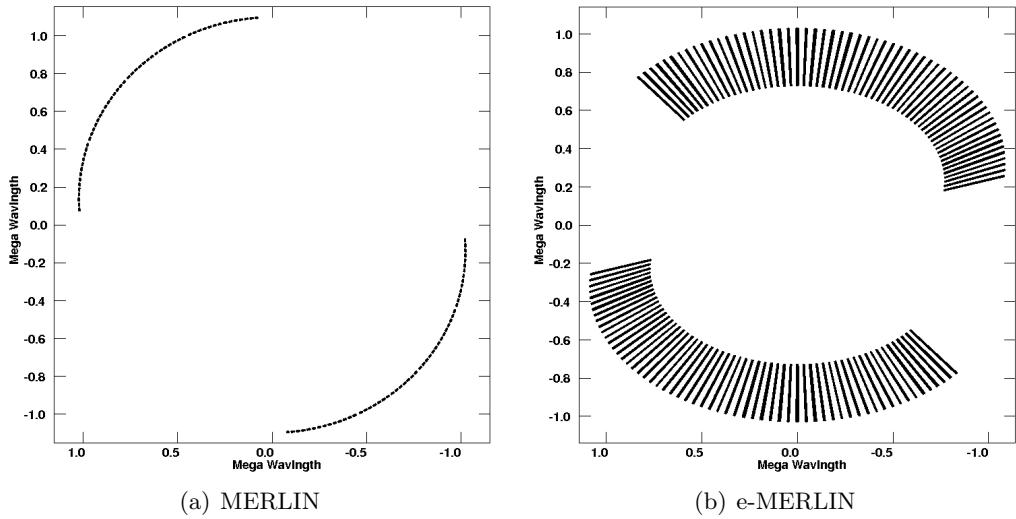


Figure 1.5: (a) MERLIN single baseline u, v coverage of M82 and (b) e-MERLIN single baseline u, v coverage of 2007+4029. The elliptical morphology arises from the baseline geometry change from the Earth's rotation.

a range of wavelengths are observed by the bandwidth, the visibility from each channel is displaced from the central frequency (centre of the ellipse arc, and band) by a factor of the wavelength difference.

The u, v coverage can be thought of as equivalent to the aperture of an optical telescope, where a larger aperture equates to a greater collecting area and thus greater sensitivity. For interferometry, the collecting area is not always complete and therefore it is a fuller u, v coverage (i.e. more visibilities) which produces more Fourier components for image reconstruction via deconvolution which leads to greater sensitivity.

The baseline lengths and position in this plane determine what part of the source structure is seen. Small baselines are sensitive to large scale structures and large baselines detect small scale structure. Alternatively, this can be described in Fourier space as; large u, v scales (large baselines) correspond to high frequency sinusoids of the Fourier transform, which can detect small scale structure from the source, and conversely; small u, v scales (small baselines) are represented by low frequency sinusoids which transform to large scale structure.

There is a second arc (Figure 1.5) which is a reflection through the origin of the u, v plane of the first arc track. This is a consequence of the sky brightness being a real function and thus its Fourier transform (visibility) $V(u, v)$ is Hermitian, i.e. $V(u, v) = V(-u, -v)^*$.

1.1.5 Radio Interferometric Imaging via Deconvolution

Before the imaging process, the data will have undergone calibration, modifying the interferometric equation to account for a number of issues, explained briefly in Section 1.1.1 and in more detail later in Chapter 3 (Radio Interferometric Calibration Pipeline for CO-BRaS). Now a description on imaging with deconvolution is given.

The interferometric equation (Equation 1.8) from Section 1.1.1 is three dimensional, but can be approximated to two dimensions when the $2\pi i w n$ component is much less than unity. This is the coplanar baseline assumption, and is valid when the baseline vectors trace out a concentric circle with the Earth's rotation axis (Thompson 1999). This is applicable to East-West interferometers. However, one-dimensional East-West interferometers suffer a lack of visibilities in the v direction. For two-dimensional interferometers, certain conditions will allow for a two dimensional representation of the visibility equation. If the $|l|$ and $|m|$ terms in Equation 1.8 are small enough i.e. a small field of view (as l and m describe the source structure), the n term can be treated as approximately zero (Equation 1.14 is from Thompson 1999)

$$\left(\sqrt{1 - l^2 - m^2} - 1 \right) w \approx -\frac{1}{2} (l^2 + m^2) w \approx 0. \quad (1.14)$$

Ignoring the w term can induce errors from aberration into the observations, similar to field curvature in an optical telescope (Anita Richards, MERLIN user guide). Removing the w dependence reduces the measurement equation to

$$V_\nu(u, v) = \int_l \int_m B_\nu(l, m) e^{-2\pi i (lu + mv)} dl dm. \quad (1.15)$$

where $V_\nu(u, v)$ and $B_\nu(l, m)$ are now functions of frequency, i.e. not assuming monochromatic radiation as in Equation 1.8 for the simplified two element interferometer example.

Reducing the interferometric equation to 2D has a significant benefit in aperture synthesis. The computational cost of the 2D Fourier transform is more feasible than the transform of the 3D relation. This is applicable for observations with a small field of view, or where the source of interest is located at the phase centre. However, for COBRaS this is not true, with wide-field imaging necessary for the science goals (see Sections 1.3.3 and 1.3.4), the problems and solutions to wide-field and wide-band imaging are discussed later in Chapter 3, but involve the 2D relation.

The u, v coverage is the response of the interferometer, whose Fourier transform is the complete sky brightness only if the u, v coverage is completely filled. In practice this is not true as is evident in Figure 1.5. Therefore a sampling function $S(u, v)$ is introduced with the visibility function in Equation 1.15 where $S(u, v)$ is zero where no data exists in the u, v plane. Directly inverting Equation 1.15 to find the sky brightness, or as is the new case, the dirty image $B^D(l, m)$, produces

$$B_\nu^D(l, m) = \int_u \int_v V_\nu(u, v) S(u, v) e^{2\pi i (ul + vm)} du dv. \quad (1.16)$$

It is possible to apply convolution theory to the variables in Equation 1.16. Convolution theory states that for the Fourier transform of two functions f and g ,

$$f * g = \mathcal{F}^{-1} \{ \mathcal{F}(f) \cdot \mathcal{F}(g) \}, \quad (1.17)$$

where $*$ denotes convolution and \mathcal{F} represents the Fourier transform. This relation is used for the terms on the right hand side of Equation 1.16 giving

$$\begin{aligned} B_\nu^D &= \mathcal{F}^{-1} \{ \mathcal{F}(V_\nu) \cdot \mathcal{F}(S) \} \\ B_\nu^D &= B_\nu * P, \end{aligned} \quad (1.18)$$

where the inverse Fourier transform of Equation 1.15 is used for B_ν and

$$P(l, m) = \int_u \int_v S(u, v) e^{2\pi i (ul + vm)} du dv \quad (1.19)$$

is the point spread function (PSF), also known as the synthesised beam or dirty beam.

To calculate the Fourier transform for Equation 1.18, Fast Fourier Transforms (FFT) are used with the data distributed over a grid. This is preferred to the direct Fourier transform which in \mathcal{O} formalism (denoting processing time or number of operations i.e. performance) is $\mathcal{O}(N^4)$ compared to $\mathcal{O}(N^2)$ for the FFT (Briggs et al. 1999). Therefore considering FFTs and the incomplete u, v coverage, gridding methods or non-linear methods are necessary for imaging in radio interferometry.

The gridding of the visibility data for the FFT enables the synthesised beam shape to be manipulated by different weighting schemes. These weighting schemes extrapolate over the grid positions where no visibility data exists in different ways, maximising different

properties in the final image.

Natural weighting is one of the density weighting schemes which gives equal weight to all visibilities, thereby emphasising grid cells containing many visibilities. This increases sensitivity at the cost of producing a synthesised beam with large side-lobes, decreasing the resolution.

Uniform weighting is another density weighting scheme which gives equal weight to all grid cells, regardless of visibility distribution. This maximises the effect of long baselines on the beam shape, creating a narrow beam profile with lower side-lobe levels thereby increasing resolution at the cost of sensitivity.

Robust weighting (Briggs 1995) is a density weighting scheme which is a hybrid of natural and uniform weighting. Robust weighting attempts to find a balance between the two schemes by creating a PSF that smoothly varies from one scheme to the other depending on a single tunable parameter.

Tapering is a weighting scheme which multiplies the weights by a Gaussian. Tapering can be combined with any other weighting scheme to compromise between sensitivity and resolution. Gaussian tapering is the optimum weighting for detecting Gaussian sources and increases the detectability of an extended source (SMA user guide).

CLEAN Algorithm

The most common traditional imaging algorithms used in aperture synthesis are the CLEAN and Maximum Entropy Method (MEM) algorithms. More recent and sophisticated imaging algorithms derived from these two are discussed in the wide-field and wide-band sections in Chapter 3.

The original CLEAN algorithm devised by Högbom (1974) solves the convolution equation (Equation 1.16) by representing radio sources as a number of point sources. The peak strengths and positions of these point sources is found iteratively and the final image is the sum of these components convolved with the CLEAN beam, usually represented by a Gaussian (Cornwell et al. 1999). The algorithm proceeds as follows:

1. Search for the strength and position of the highest intensity peak in the dirty image B^D .
2. Subtract the dirty beam P from the dirty image at the position of the point source found in step 1. The beam P is multiplied by a loop gain $\gamma \leq 1$.

3. The position and strength of the point source subtracted is recorded in a model.
4. Go back to step 1, unless all remaining peaks are below a specified user level. The remaining dirty image (minus subtracted peaks) is the residual image.
5. Convolve the model of point sources with a CLEAN beam, which is usually an elliptical Gaussian fitted to the central lobe of the dirty beam.
6. Add the residual image to the CLEANed image.

A FFT-based CLEAN algorithm by Clark (1980) performs in a similar fashion to Högbom's algorithm, but finds the positions and strengths of point sources by only using a fraction of the dirty beam profile. The Clark algorithm operates in two cycles, the major and minor cycles and proceeds as follows:

1. **Minor Cycle** - A segment of the beam is selected with the highest exterior side-lobe i.e. the central portion of the dirty beam.
2. **Minor Cycle** - Peaks are selected from the dirty image if the strength of the source is greater than the highest side-lobe of the beam from step 1.
3. **Minor Cycle** - A Högbom CLEAN is performed on all of the selected points from step 2 using the segment of the beam from step 1. This continues for this list of points until all sources selected are weaker than the side-lobe.
4. **Major Cycle** - The model created in step 3 is then transformed via a FFT, multiplied by the sampling function (the inverse transform of the PSF) transformed back and subtracted from the dirty image. Also any errors in the residual images from previous minor cycles are corrected by subsequent minor cycles.

This algorithm is sufficient to find CLEAN components for dirty beams which have fairly good side-lobe patterns. One further modification to the CLEAN algorithm is the Cotton-Schwab algorithm (Schwab 1984) where the major cycle subtracts CLEAN components from the un-gridded visibility data. This enables the removal of aliasing noise and gridding errors if the inverse Fourier transform of these components to each u, v sample is accurate enough (Cornwell et al. 1999). The algorithm also decides whether to use the direct Fourier transform for a small number of CLEAN components to increase accuracy, or to use the FFT for large numbers of CLEAN components.

A major advantage of the Cotton-Schwab algorithm is its ability to CLEAN many separate fields independently in the minor cycles and remove all of the CLEAN components together in the major cycle. Calculating the residual for each small field allows for the w -term in the full 3D measurement equation (Equation 1.8) to be determined. This corrects for non-coplanar baselines which is a problem for large arrays and wide-field imaging (see Chapter 3 for discussions). This is the algorithm used in the AIPS task IMAGR⁵.

Maximum Entropy Algorithm

Another algorithm is the MEM, a form of information theory, which selects a probability distribution that best fits the data within the noise level from all the possible distributions and also has maximum entropy. Maximum entropy is defined as a positive image with a compressed range of pixel values forcing the image to be “smooth” (Cornwell et al. 1999). Entropy takes the general form:

$$\mathcal{H} = - \sum_k I_k \ln \frac{I_k}{M_k} \quad (1.20)$$

where I_k is the reconstructed image and M_k is the expected or a priori image.

Each visibility cannot be exactly fitted by the probability distribution to produce a positive value and therefore the data is constrained by a χ^2 fit of the probability distribution to the observed image being equal to the expected value

$$\chi^2 = \sum_k \frac{|V(u_k, v_k) - \widehat{V}(u_k, v_k)|^2}{\sigma_{V(u_k, v_k)}^2}, \quad (1.21)$$

where $V(u_k, v_k)$ is the probability distribution, $\widehat{V}(u_k, v_k)$ is the observed image and $\sigma_{V(u_k, v_k)}^2$ is the variance of the image, i.e. Gaussian noise (Cornwell et al. 1999).

Images with around one million pixels take a similar amount of time for MEM and CLEAN, with MEM being faster for larger images and CLEAN faster for smaller images. MEM can also be faster than CLEAN when the image is filled with emission, but CLEAN is faster for sources which are well represented by a small number of point sources and for images with high dynamic ranges (Cornwell and Evans 1985).

⁵A full list of the AIPS Nomenclature is given in Appendix A.

1.1.6 Polarisation Effects

This section briefly discusses some of the general effects of polarisation for interferometers. Both the electric and magnetic field vectors of an electromagnetic wave trace an ellipse, the orientation of which defines the polarisation state of the wave. Only the electric field of the wave can be measured by the antenna feeds. The ellipse defining the electric wave is comprised of two orthogonal components which sum to produce either linear or circular polarisations. Linear polarisations have two cartesian components X and Y, and circular polarisations have left and right hand circular components (simply L and R).

The relation of the intensity of the wave and the polarisation ellipse are defined by Stokes parameters named after George Stokes who devised them in the 1850's. The Stokes parameters are intensity I, linear polarisation Q and U, and circular polarisation V. The relation of the Stokes parameters and circular polarisations is given by

$$\begin{aligned}
 E_R &= \text{Re}\{A_R e^{2\pi i \nu t}\} \\
 E_L &= \text{Re}\{A_L e^{i\delta_{RL}} e^{2\pi i \nu t}\} \\
 I &= \langle A_R^2 \rangle + \langle A_L^2 \rangle &= \langle E_R E_R^* \rangle + \langle E_L E_L^* \rangle \\
 Q &= \langle 2 A_R A_L \cos \delta_{RL} \rangle &= \langle E_R E_L^* \rangle + \langle E_L E_R^* \rangle \\
 U &= \langle 2 A_R A_L \sin \delta_{RL} \rangle &= i (\langle E_R E_L^* \rangle - \langle E_L E_R^* \rangle) \\
 V &= \langle A_R^2 \rangle - \langle A_L^2 \rangle &= \langle E_R E_R^* \rangle - \langle E_L E_L^* \rangle.
 \end{aligned} \tag{1.22}$$

Interferometers have dual polarisation feeds which produce four correlation states, RR, LL, RL and LR for circular and XX, YY, XY, YX for linear polarisations. These four states can be utilised with the relations in Equation 1.22 to create intensity and polarisation maps. Numerous considerations involving the calibration for polarisation effects from the atmosphere and system electronics are beyond the scope of this thesis, because the initial goal of COBRaS is an intensity study, utilising only parallel polarisations.

1.2 e-MERLIN

e-MERLIN⁶ is a UK National Facility operated by The University of Manchester on behalf of the Science and Technology Facilities Council (STFC). It is an upgrade to the MERLIN (Multi-Element Radio Linked Interferometer Network) array, consisting of seven radio

⁶e-MERLIN: <http://www.e-merlin.ac.uk/>

telescopes. Figure 1.6 shows the distribution of telescopes spanning across the UK.



Figure 1.6: Positions of the seven radio telescopes of e-MERLIN across the United Kingdom. Clockwise from top; Lovell, Mark II, Cambridge, Defford, Knockin, Darnhall and Pickmere.

The upgrade consists of a new optical fibre network connecting each telescope to the Jodrell Bank Observatory, where the new WIDAR correlator developed by the Dominion Radio Astrophysical Observatory (DRAO) resides. New bandwidth receivers increase the useable bandwidth by two orders of magnitude, resulting in a continuum sensitivity increase of a factor of 10 or more compared to the old MERLIN array. To demonstrate this sensitivity increase, Figure 1.7 shows the u, v coverage of the old MERLIN array displayed alongside the u, v coverage of e-MERLIN. The MERLIN u, v coverage is a 12 hour observation of M82 with declination: $+69^\circ 40' 47''$ (data courtesy of Danielle Fenech; private communication), and the e-MERLIN u, v coverage is a COBRA-S L-band Legacy observation with declination: $+41^\circ 22' 48''$. Despite the lower declination and shorter total observation time, the e-MERLIN u, v plot shows a dramatic increase in coverage in the radial direction due to the increase in L-band bandwidth from 16 MHz to 512 MHz.

There are three observing bands for e-MERLIN. L-band operates at 1.3 - 1.8 GHz, C-band at 4 - 8 GHz and K-band at 22 - 24 GHz, with the available maximum bandwidths of 512 MHz for L-band and 2048 MHz for C and K-bands per polarisation (circular). Throughout this thesis, observations with frequencies between 1.3 - 1.8 GHz will be referred to as L-band as the new bandwidth encompasses all of these frequencies. Observations

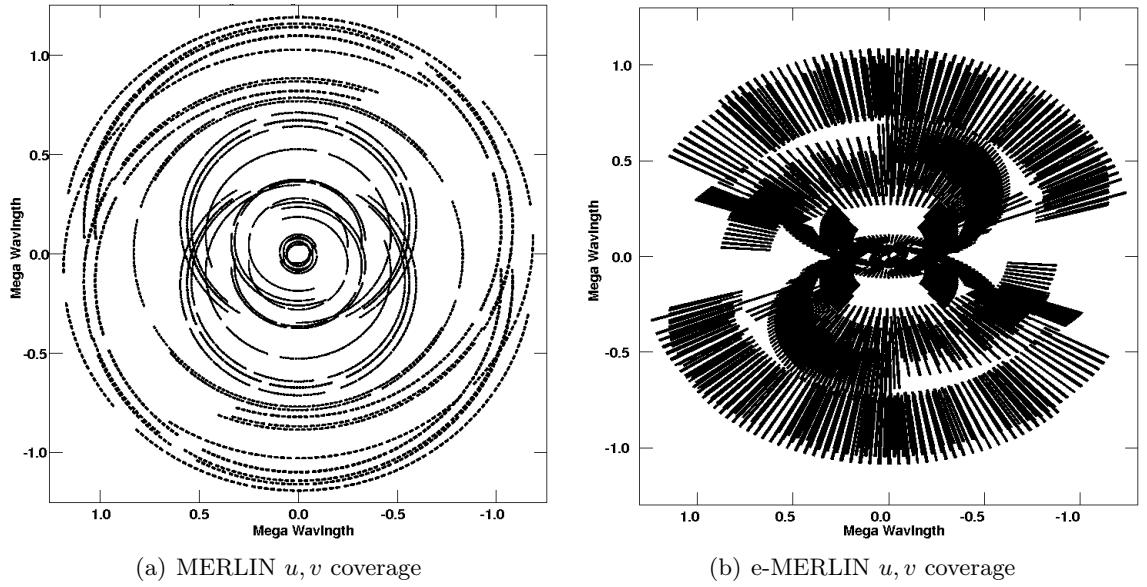


Figure 1.7: (a) MERLIN u, v coverage of M82 observation at 1.6 GHz with 16 MHz bandwidth and (b) e-MERLIN u, v coverage of L-band Legacy Pointing C (including calibration sources) at 1.5 GHz with 512 MHz bandwidth.

with frequencies between 4 - 8 GHz will be referred to as C-band. All bands are comprised of smaller sub-bands or intermediate frequencies (IFs; in the AIPS nomenclature), which segregate the total bandwidth into groups of channels.

The new correlator at Jodrell Bank Observatory is a smaller version of the WIDAR correlator at the JVLA. A range of correlator capabilities is available for both continuum and spectral-line observations, and the reader is referred to the relevant literature for details (Garrington et al. 2004; <http://www.e-merlin.ac.uk/tech/>).

The shortest baseline of e-MERLIN is the Lovell - Mark II baseline of 400 m. However, the large difference in uv-spacing sampled between this baseline and the next shortest baseline of ~ 11 km (Mark II - Pickmere), means the Lovell - Mark II baseline is not used. This is because there are inadequate data to fully recover any diffuse structures on large spatial scales seen on this very short baseline and connect the smaller spatial scales detected by the other baselines in the array during the imaging process (Rob Beswick, private communication). The smallest useable baseline is therefore Mark II - Pickmere (11 km) and the largest baseline is Lovell - Cambridge (217 km).

This provides e-MERLIN with resolutions of ~ 150 , 40 and 12 mas for L-, C- and K-band observations respectively (with small changes over the frequency range of each band). Table 1.1 gives the expected technical capabilities of a fully commissioned e-MERLIN array.

Table 1.1: Technical Capabilities of e-MERLIN

	1.5 GHz (L-band)	5 GHz (C-band)	22 GHz (K-band)
Resolution (mas)	150	40	12
Field of View (arcmin)	30	7	2.0
Bandwidth (GHz)	0.5	2	2
Freq. Range (GHz)	1.3 - 1.8	4 - 8	22 - 24
Sensitivity (μ Jy/bm) in full imaging run	5 - 6	1.8 - 2.3	\sim 15
Surface brightness sensitivity (K)	\sim 190	\sim 70	\sim 530
Astrometric performance (mas)	\sim 2	\sim 1	\sim 2
Amplitude calibration	2%	1%	10%

General capabilities of the full e-MERLIN array. The sensitivity and surface brightness numbers include e-MERLIN and the Lovell telescope. The field of view decreases with inclusion of the Lovell telescope by approximately $20/(\nu/ 1.4\text{GHz})$ arcmin, where ν is the observed frequency. This table is taken from the e-MERLIN website: <http://www.e-merlin.ac.uk/tech/>.

1.3 The e-MERLIN Cygnus OB2 Radio Survey

The Cygnus OB2 Radio Survey (COBRaS; P.I. Prof. Raman Prinja, UCL) is an e-MERLIN Legacy project on the Cygnus OB2 association (Cyg OB2), the largest OB association in the northern hemisphere and one of five OB associations in the Cygnus X region. COBRaS will provide the deepest, most sensitive radio survey of the region using observations at L and C-band, with \sim 300 hours of observing time on the core region of the association. This substantial dataset will enable a range of diverse astrophysical topics to be explored, and will have direct links to other wavelength surveys, generating a multi-wavelength window into one of the most spectacular massive stellar clusters in the Galaxy.

1.3.1 The Cygnus OB2 Association

The Cyg OB2 association is a young massive cluster in the heart of the Cygnus X region of the Galaxy, residing behind a large molecular cloud called the Great Cygnus Rift. Cyg OB2 was first observed by Münch and Morgan (1953), who noticed a number of blue giants in the region, with subsequent photometric and spectroscopic observations by Johnson and Morgan (1954). Further studies were made by Morgan et al. (1954), Schulte (1956b), Schulte (1956a) and UBV photometric investigations by Lawrence and Reddish (1965) and Reddish et al. (1967). The first general census of the massive star

population in Cyg OB2 using UBV photometry and spectroscopy for spectral classification was conducted by Massey and Thompson (1991) with subsequent surveys in near-IR with the Two Micron All Sky Survey (2MASS; Knöldlseder 2000; Comerón et al. 2002), and in X-rays with Chandra (Albacete Colombo et al. 2007b; Drake 2009).

All surveys have indicated a rich diverse population of stars within the Cyg OB2 region. Knöldlseder (2000) showed Cyg OB2 to have a total cluster mass of $(4 - 10) \times 10^4 M_{\odot}$, with 120 ± 20 O-type and 2600 ± 400 OB type stars. In addition to these OB stars, Cyg OB2 is home to Be stars, young stellar objects (YSOs), two known Wolf-Rayet (WR) stars (WR 145, WR 146), a candidate luminous blue variable (LBV) (G79.29+0.46), HII regions (DR 15, DR 17, DR 18, DR 20 etc.) and a γ -ray source (TeV J2032+4130). Figure 1.8 shows a Herschel image at 70, 160, and 250 μ m of the Cyg OB2 region with labels of the surrounding HII regions. Most of the population of Cyg OB2 is obscured by the molecular cloud.

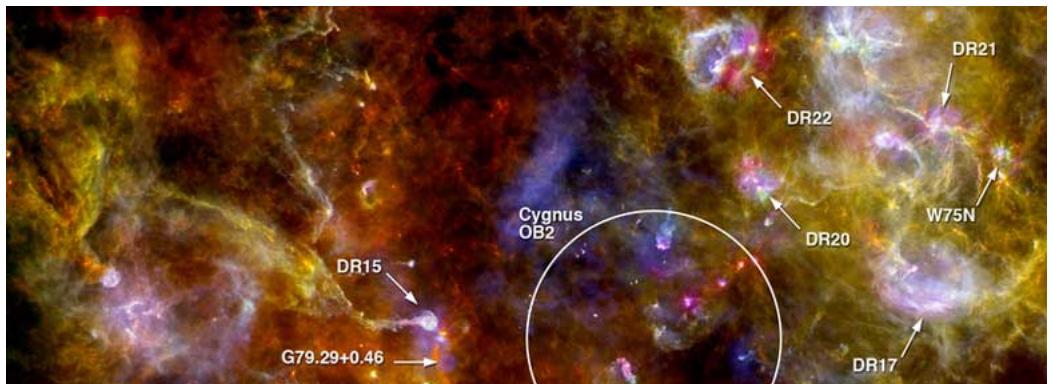


Figure 1.8: Image of Cyg OB2 and the surrounding regions from Herschel observed at 70 (blue), 160 (green), and 250 (red) μ m. North is situated towards the right, and East towards the top of the image. Figure taken from the ESA website.

The Cyg OB2 association extends across ~ 2 degrees of sky, which corresponds to ~ 60 parsecs at an assumed distance of 1.7 kpc (Massey and Thompson 1991). However, this distance is uncertain with values ranging over 1.2 - 2.1 kpc in the literature (Hanson 2003; Reddish et al. 1966), with the uncertainty originating from the variable visual extinction ($4^m - 10^m$; Knöldlseder 2000) in the Cygnus X region. Recent parallax observations of 6.7 GHz methanol and 22 GHz water masers have constrained the distance of Cyg OB2 to 1.40 ± 0.08 kpc (Rygl et al. 2012; Wright et al. in prep.). The estimated age of Cyg OB2 is $\sim 2 - 3$ Myr (Albacete Colombo et al. 2007b), in keeping with that of young massive clusters.

Table 1.2: Physical Properties of Cyg OB2

Physical Parameter		Reference
Centre (J2000)	$\alpha = 20^{\text{h}} 33^{\text{m}} 10^{\text{s}}$, $\delta = +41^{\circ} 12'$	
Total Stellar Mass	$(4 - 10) \times 10^4 M_{\odot}$	1
O Star Members	120 ± 20	1
OB Star Members	2600 ± 400	1
Association Diameter	$\sim 2^{\circ}$	1
Distance	1.40 ± 0.08 kpc	2
Visual Extinction	4^{m} to 20^{m}	1
Age	$2 - 3$ Myr	3

References: 1. Knödlseder (2000), 2. Rygl et al. (2012), 3. Albacete Colombo et al. (2007b).

The Cyg OB2 association thus provides a good laboratory to study a diverse range of stellar and cluster astrophysics. Radio observations from e-MERLIN and COBRaS have the benefit of piercing through the giant molecular cloud avoiding problems with variable extinction. A summary of all the adopted physical properties of Cyg OB2 is given in Table 1.2.

1.3.2 The Role of Massive Star Winds in the Galaxy

Massive stars are a rare but important population in the Galaxy. The powerful winds of massive stars influence the dynamics of the clusters in which the majority of them reside, expelling mass and ionising the matter in the surrounding regions. In particular the intense UV photons from OB stars can cause evaporation and compression of the surrounding ISM, forming comet shaped clouds which induce the formation of low and intermediate mass stars (Lee and Chen 2007). Such triggered star formation is possibly occurring in Cyg OB2 where the collective winds of OB stars in the core of Cyg OB2 are triggering star formation in the preperh HII region DR 15 (Vink et al. 2008).

Massive stars are one of the most important objects to transfer chemically enriched material into the ISM from internal nuclear processes throughout their lifetimes. Helium is produced via the CNO cycle during the main sequence phase, and also during the WR phase and then subsequently, carbon oxygen, silicon and iron (and nickel, but often not referred to in the literature because of its 10% abundance compared to iron), with the resulting energy emitted from the eventual supernova, intense enough to create further heavier elements.

Moreover, massive stars are the progenitors of supernova events, making the whole

lifetime of a massive star, from birth to death, the most important component in galactic chemical evolution. This includes further star and planet formation from the expelled material, as well as metallicity fractions for the ISM and future stellar systems.

1.3.3 COBRaS Scientific Analysis Goals

The COBRaS Legacy project aims to advance the scientific knowledge and understanding of a number of inter-related themes in stellar astrophysics with three major scientific areas.

1. The mass loss and the evolution process in massive stars.

There is currently a major discrepancy between the mass loss rates of OB stars determined from observed H α and radio free-free processes (so called density-squared methods; see Chapter 5) and those from ultra-violet (UV) resonance lines of P $^{4+}$, by an order of magnitude or more (Prinja et al. 2005; Fullerton et al. 2006; Puls et al. 2006). This disparity has significant consequences for the evolution and final stages of a massive star’s life (which is strongly driven by mass loss) and the interaction of massive stars and the instellar medium, where feedback mechanisms of the stellar winds from massive stars drives the galactic chemical evolution. The proposed cause of this discordance is believed to arise from *clumped* and/or *porous* stellar winds as evidenced by observations of the aforementioned P $^{4+}$ discrepancy (Prinja and Massa 2010; Prinja and Massa 2013; and references therein), X-ray spectroscopy (Crowther 2007; Puls et al. 2008), and theoretical hydrodynamical models, which predict instabilities from the radiation driving mechanism (Owocki and Rybicki 1984; Dessart and Owocki 2005). COBRaS provides a robust way to measure accurate radio fluxes for OB stars (at 6 cm and 20 cm), which is free from uncertain wind parameters such as the terminal velocity⁷ (see Chapter 5 for details). This in turn will provide a foundation to investigate the discrepancy between observed radio fluxes and fluxes predicted from smooth wind models on a large scale, and also to determine the clumping factors and the effects of wind structure far out in the winds of massive stars.

2. The massive binary frequency and particle acceleration in colliding-wind binaries.

The frequency of massive binaries in stellar clusters is currently uncertain, with

⁷However, mass loss rates are dependent on terminal velocity.

estimates ranging from 20-60% of all O stars as members of binary systems (Sana et al. 2008, NGC 6231 cluster; De Becker et al. 2006, IC 1805). Binary fractions, information about mass distribution among binaries and orbital parameters etc. are important for evolution and population modeling (Dionne and Robert 2006) and for understanding stellar populations and cluster dynamics (Eldridge et al. 2008).

The accurate measurement of radio fluxes with the spectral index information of massive stellar winds and colliding wind binaries from COBRaS will provide a good laboratory to investigate particle acceleration via the Fermi mechanism in real physical scenarios (Pittard and Dougherty 2006).

3. Triggered and on-going star formation.

Vink et al. (2008) conducted a survey using the INT Photometric H-alpha Survey of the Northern Galactic Plane (IPHAS) towards Cyg OB2 and the H II region DR 15, detecting 50 new pre-main-sequence (PMS) candidates via their strong H α emissions. A prominent arc of T Tauri stars was found close to DR 15 (Figure 1 in Vink et al. 2008), suggesting that these strong emission-line objects could just be the tip of a much larger low mass PMS population yet to be discovered. The authors also suggest the possible scenario, where the central OB star winds are triggering star formation in the southern periphery region of DR 15. No confirmed instance of triggered star formation from OB star winds has ever been observed because of the difficulty in making this connection. Whilst DR 15 is not in the COBRaS survey region, a complementary dataset *could* arise from a recently accepted JVLA (Priority C) proposal ‘Triggered star formation in DR 15: the Cyg OB2 HII region’ (P.I. Dr Danielle Fenech) observing at L-band and C-band.

For on-going star formation in the main core of Cyg OB2, radio observations could reveal so-called weak T-Tauri stars (WTTS) which may be the dominant population of pre-main sequence (PMS) stars. These stars differ from classic T-Tauri stars (CTTS) by their weak H α emission line ($W_\lambda < 10 \text{ \AA}$), making these objects difficult to identify. However, WTTS are non-thermal radio emitters whereas CTTS and Herbig Ae/Be stars are thermal (O’Neal et al. 1990; Skinner et al. 1993). The emphasis from COBRaS on providing spectral index information with accurate radio fluxes will therefore help in distinguishing the population of PMS stars in Cyg OB2.

1.3.4 COBRaS Technical Analysis Goals

In order to achieve the full scientific goals stated in Section 1.3.3, a number of technical areas are to be addressed. Automated software is required where possible, due to the large Legacy data volume. These technical areas are:

Data Processing

The development of data reduction processes to initially assess the data quality and create four main tiers of data output: calibrated visibility data, individual pointing images, full mosaiced images of the core at L-band and C-band, and source and flux lists with spectral information.

Calibration

The creation of a large pipeline to process the observations from raw data (AIPS UVFITS or CASA measurement sets) to fully calibrated datasets ready for imaging. This includes automated data editing due to the infeasibility of manual invention.

Imaging

Considerations for wide-band and wide-field imaging for e-MERLIN in the data analysis of COBRaS datasets, containing faceted imaging and removal of aliasing sources for each pointing. Developments for achieving optimal mosaicing using a heterogeneous array are also required.

Analysis

Pipelines for source location, flux and position extraction and determining spectral index information are necessary for the science goals of COBRaS. This provides the first step in source classification in determining between thermal and non-thermal emission, the second step is to cross correlate the source lists generated with existing catalogues. A definitive catalogue of previous surveys is thus also required.

Data Archiving

Outputs from the project are to be made available to the community via the Virtual Observatory (VO) library. The fully calibrated pointings will also be available through VO links to an external location as direct access to AIPS UVFITS or CASA measurement sets may not be feasible.

All pipelines are to be written in Python/ ParselTongue languages which are commonly used in connection with the AIPS software package.

1.3.5 COBRaS Observations

COBRaS has been awarded ~ 300 hours of e-MERLIN observing time, including the Lovell telescope, which is split between L-band (42 hours) and C-band (252) observations. In total there are 7 L-band pointings and 42 C-band pointings centred on the core region of Cyg OB2. Figure 1.9 shows a close up of the C-band (a) and L-band (b) mosaicing strategies of the core region of Cyg OB2.

The e-MERLIN observations involve *phase referencing* (see Chapter 3 on calibration), which consists of alternating between observing a phase calibrator and the target field. For COBRaS Legacy observations, the observed target field cycles through all of the different pointings which make up the mosaic to optimise the u, v coverage of each field and the observing time required to achieve enough sensitivity per pointing.

1.3.6 Synergies with Multi-wavelength Studies of Cyg OB2

The COBRaS project has direct links with other surveys at different wavelengths, providing a unified view across the electromagnetic spectrum of Cyg OB2. In the X-rays, the Chandra Cyg OB2 Survey (Drake 2009; Wright 2011) contains a 100 ks $16' \times 16'$ observation of the core region of Cyg OB2. The sensitivity of this survey is such that the dataset is expected to be complete to 95% for stellar masses down to one solar mass. The Chandra dataset provides probes into the X-ray behaviour of the shocked winds of OB stars and colliding wind interaction regions of massive binaries.

In the optical, the Isaac Newton Telescope Galactic Plane H α Survey (IPHAS; Drew et al. 2005) has obtained deep images of Cyg OB2 in the H α , r' and i' bands. In addition to this photometry, IPHAS has MMT HectoSpec spectroscopy of ~ 1000 sources in the Cyg OB2 region. Another study in the optical (which also includes near-IR spectroscopy) is the H α Radial Velocity Survey of Cyg OB2 (Kiminki et al. 2007). This survey presents a long term study on massive binaries in Cyg OB2, and in conjunction with COBRaS will enable the properties of wind clumping to be investigated from the mass loss diagnostics from different wavelengths, which probes different regions in the stellar wind (See Chapter 5).

In the near-IR, 2MASS (Skrutskie et al. 2006), and mid-IR Spitzer (Werner et al. 2004)

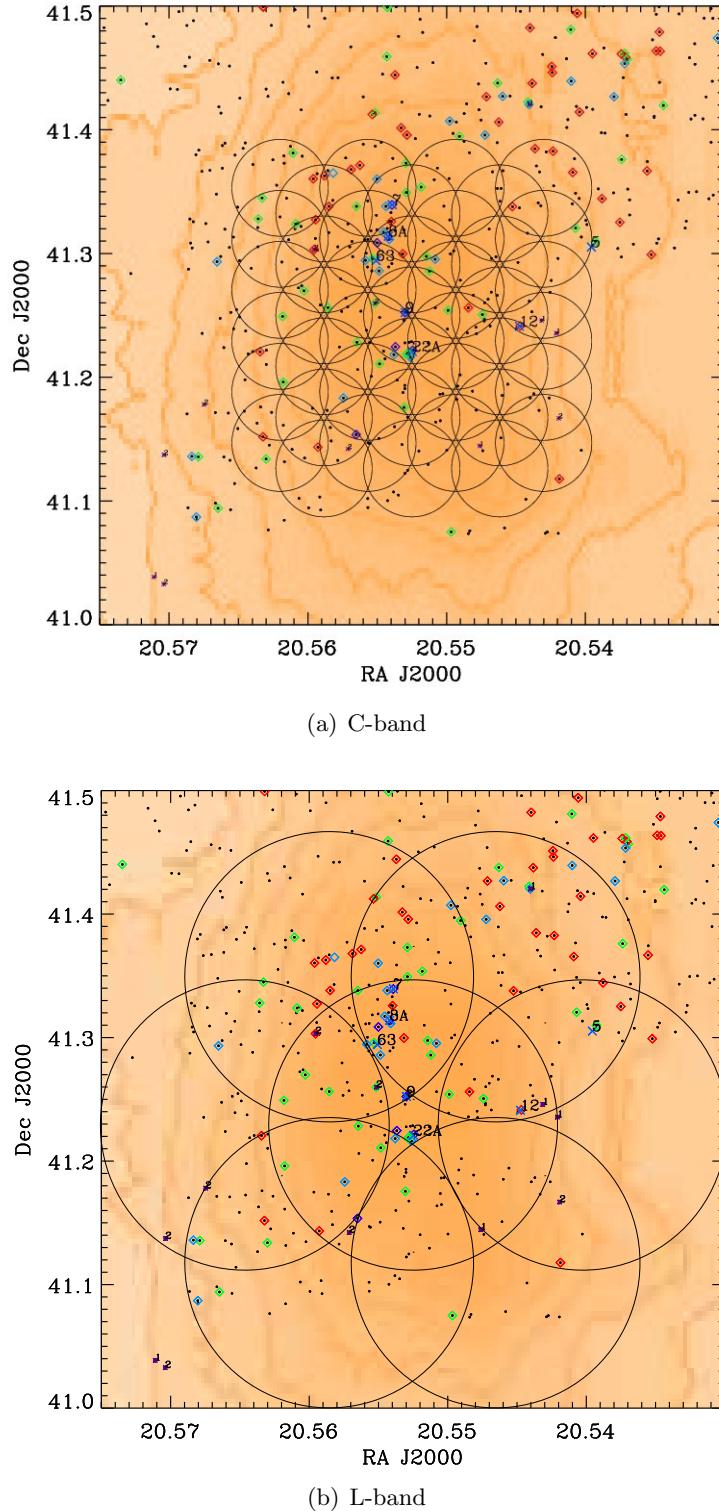


Figure 1.9: (a) C-band and (b) L-band mosaicing strategy for COBRaS observations on the core region of Cyg OB2. The black points are X-ray sources from the Chandra Cyg OB2 Legacy survey (Drake 2009), and known radio OB sources in blue.

IRAC and MIPS (Fazio et al. 2004; Rieke et al. 2004) surveys are essential in helping complete the picture of star formation in the surrounding regions. Radio continuum measurements from COBRaS with IR photometry will lead to new information on the HII regions and whether the winds of massive stars induce triggered star formation.

1.4 Thesis Aims

This thesis addresses many areas of the data analysis issues presented in Section 1.3.4 from the COBRaS proposal. This includes the technical development of pipelines and algorithms to solve these technical problems (excluding imaging) as well as investigating the first results from the COBRaS L-band Legacy and C-band commissioning observations and the Chandra X-ray Legacy dataset. The layout of the thesis is as follows.

Chapter 2 describes a new reduction and RFI-mitigation program for e-MERLIN, SERPent, which takes a raw interferometric AIPS UVFITS dataset and prepares it for calibration by editing and flagging bad, unusable data.

Chapter 3 introduces the automated COBRaS calibration pipeline which consists of phase, amplitude, bandpass calibration and self-calibration. The pipeline is tested on C-band e-MERLIN commissioning data and COBRaS L-band Legacy data. The chapter also includes discussions on wide-band and wide-field imaging techniques and mosaicing strategies with e-MERLIN for the COBRaS project.

Chapter 4 contains catalogue amalgamation routines to compile a complete catalogue of previous surveys of the Cyg OB2 region. The chapter also contains analysis packages for source detection and flux extraction on a pixel-by-pixel basis. Monte Carlo simulations with somewhat idealised Gaussian noise are conducted to test the performance of the methods against the standard AIPS flux extraction method JMFIT. In addition to the flux extraction, a significance boosting module is created to cross-reference the sources in the maps with the aforementioned catalogues and increase the detection sigma significance using Bayes' Theorem.

Chapter 5 contains the predictions of mass loss rates and radio fluxes from free-free emission of OB stars in Cyg OB2 using theoretical stellar parameters from the literature. Radio fluxes are determined for O-type stars and B supergiants for smooth winds and clumped winds. The first L-band Legacy maps and C-band commissioning maps are presented here with results in the form of source and flux lists and some analysis.

Chapter 6 is a study of X-ray variability of O-type stars from the Chandra Legacy survey of Cyg OB2. Finally, Chapter 7 contains a summary and conclusions from this thesis, and discussions on future work in the field of software development for radio interferometry.

Chapter 2

Radio Frequency Interference Mitigation with SERPent

A common mistake that people make when trying to design something completely foolproof is to underestimate the ingenuity of complete fools.

Douglas Adams

The first part of the automated radio interferometric pipeline is the preparation of the raw dataset by editing bad or unusable data. The reduction and removal of Radio Frequency Interference (RFI) is essential to enable accurate calibration of e-MERLIN data. This chapter presents the Scripted E-merlin Rfi-mitigation PipelinE for iNTerferometry (SERPent) algorithm applied to commissioning e-MERLIN and COBRaS Legacy datasets. Some details from this chapter are also presented in Peck and Fenech (2013).

2.1 RFI Mitigation Techniques

Modern interferometers are becoming increasingly more sensitive and powerful, with resulting datasets becoming ever bigger. Therefore, the need for automation of certain procedures in reduction and calibration of interferometric data is vital. The manual removal of radio-frequency interference (RFI) and other bad unusable data by the user presents a major ‘bottleneck’ in this reduction and calibration procedure.

Until recently, the manual flagging of typical datasets took a reasonable amount of time, with data sizes being on the order of Megabytes (MB). Improvements in receivers, electronics, correlators and optical fibre networks, result in observations which span a wide frequency range, into bands which are not protected for radio astronomy, thereby increasing the incidence of RFI. With future emphasis on multi-observation and full sky surveys (in the case of the SKA), data sizes will be on the order of Terabytes (TB) for the current crop of interferometers (including e-MERLIN) and even larger for future instruments (e.g. SKA), making manual flagging unfeasible. It is clear that automation of this process is necessary for the current generation of interferometers such as e-MERLIN, JVLA, ALMA, LOFAR and for future interferometers (MeerKAT, ASKAP, SKA).

One of the toughest challenges in RFI mitigation is accounting for its variable intensity, morphology and unpredictable nature. RFI can arise from many sources such as radio stations, microwaves, lightning, aeroplanes, mobile phones, CCTV etc. Some of these occur at specific frequencies (radio stations, mobile phones) and may only be problematic for certain arrays. The individual array characteristics and the corresponding RFI environment, needs to be considered to achieve optimal RFI reduction. Therefore, creating robust methods to mitigate RFI to a certain level to enable the highest sensitivity mapping possible, is essential.

Mitigation can be applied at two stages in the interferometric data reduction process: pre-correlation and post-correlation of the antenna signals. Both of these methods can be complementary to one another, as they will remove different kinds of RFI.

2.1.1 Pre-correlation RFI Mitigation

Pre-correlation is a very powerful option for RFI mitigation because the observational data are still at their highest time resolution (sub-integration time) (Offringa et al. 2010a), although executing the processes on small sections of the entire observation at the station in real time is challenging. Niamsuwan et al. (2005) test a strategy termed *asynchronous pulse blanking* (APB) on simulated data to constrain the inherent parameters associated with the algorithm. The method maintains a selection of the real-time data within the computer's memory to enable deletion of contamination when an RFI pulse is detected. Only when a sample exits the buffer is the blanking operation performed. The amount of time before and after the RFI spike to be removed depends on one of the input parameters, including an aggressiveness parameter (β) which controls the detection level threshold.

The Cumulative Sum (CUSUM) threshold method of RFI in the time-frequency domain has been conducted by Baan et al. (2004) at the Westerbork Synthesis Radio Telescope (WSRT) as part of the RFI mitigation subsystem (RFIMS) before the correlation of the signal. Another thresholding method using χ^2 statistics is proposed by Weber et al. (1997) and applied to spectral line observations.

2.1.2 Post-correlation RFI Mitigation

One advantage of post-correlation RFI mitigation is an inherent mitigation routine through the cross correlation of visibilities from different antennas. An RFI incident at any individual antenna location will not exist after correlation, because the corresponding RFI signal is not present at other antenna locations. The appearance of correlated RFI in VLBI is therefore unlikely, but interference can be correlated on ‘short’ baselines (Roshi and Perley 2003). Whilst technically e-MERLIN conducts VLBI in the sense that e-MERLIN observations require fringe-fitting (See Section 3.6), it still suffers from relatively short baselines compared to those comprising the EVN, VLBA and Global VLBI array configurations. Therefore RFI is seen in e-MERLIN observations.

Post-correlation is the final stage to remove RFI before calibration procedures commence. Methods include the use of an independent RFI reference signal from an additional receiver horn pointing in the direction of the RFI source to subtract the RFI from the data (Briggs et al. 2000), and fringe-fitting for spatially and temporally constant RFI (Athreya 2009), where the RFI causes the observed visibilities to deviate from the *true* visibilities when plotted in the *real* and *imaginary* visibility plane (see Figure 1 in Athreya 2009).

Thresholding methods can also be applied in post-correlation RFI mitigation, because the cross-correlated visibility amplitudes increase when RFI is present. Offringa et al. (2010a) analyse a number of threshold methods with simulated and real data from LO-FAR and WSRT, and demonstrate that the SumThreshold method (explained in Section 2.2) performs better by successfully characterising more RFI than the other rival methods. These include the Cumulative Sum method (stated earlier, see Baan et al. 2004), VarThreshold and Singular Value Decomposition (SVD).

2.1.3 Reduction and RFI Mitigation for e-MERLIN

Every interferometer around the world is unique, having a different baseline distribution, location, observed frequency band, RFI environment etc. Therefore, the method of mit-

igation needs consideration and the implementation and any parameters used may need to be optimised to suit any individual array characteristics. For example, WRST is a large and sparse interferometer where the RFI is sometimes partially coherent. For this reason, post-correlation spatial processing algorithms are not always effective. Baan et al. (2004) conclude that real-time, pre-correlation time-frequency analysis conducted at each antenna would be more effective than any post-correlation method.

Automated flaggers are compared on accuracy, computational performance, robustness and any technical requirements they impose (Offringa et al. 2010b). These criteria and the needs of the interferometer, will define which method is the most practical for that particular array.

For e-MERLIN, there is no hardware in place at Jodrell Bank to enable the implementation of pre-correlation RFI mitigation methods such as those described in Section 2.1.1. Commissioning data from e-MERLIN revealed the incidence of RFI varying simultaneously over time and frequency, an example is given in Figure 2.1. Evaluating the number of post-correlation methods available, and consulting the recent (at the time of inception of SERPent) study by Offringa et al. (2010a), the SumThreshold methodology offers a strong strategy to form the basis of an e-MERLIN specific reduction and RFI mitigation algorithm, SERPent (Peck and Fenech 2013).

Additional specific e-MERLIN issues arose during the creation of SERPent, including the Lovell stationary scan (Section 2.3.1), and the zero-level amplitude dropouts (Section 2.3.2). These problems are described and discussed later, but highlight the need for a new software package which reduces and mitigates RFI for the e-MERLIN system.

2.2 SumThreshold Method

The most effective thresholding method currently available is demonstrated by Offringa et al. (2010a) to be the SumThreshold and this is the adopted RFI detection algorithm for SERPent. An overview of the method is given here, for a more in depth analysis of the process see Offringa et al. (2010a).

Threshold methods work on the basis that RFI increases visibility amplitudes for the times and frequencies in which they are present. Therefore there is a considerable difference compared to other RFI-free visibility amplitudes, making these RFI contaminated visibilities statistical outliers. If the RFI amplitudes fulfil a certain threshold condition,

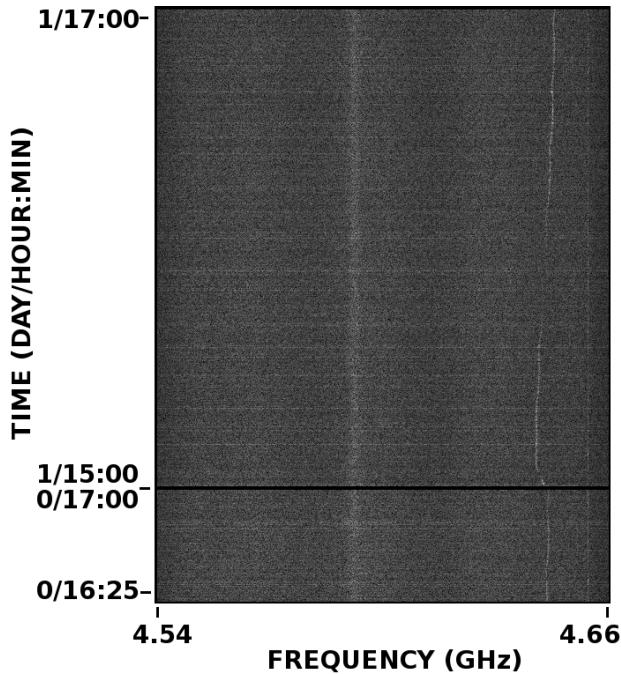


Figure 2.1: Time-frequency plot of the visibilities of the source 0555+398 from dataset number 3 (see Table 2.1). A single IF and RR polarisation is shown with a frequency range from 4.54 to 4.66 GHz from the baseline Knockin-Pickmere (5 – 7). RFI is seen to vary both in time (vertical axes) and frequency (horizontal axes) at around 4.64 GHz.

they are detected and flagged. The threshold level is dictated by the statistics of the relevant visibility subset, which can be the entire observation (all time scans, frequency channels, baselines etc.) or a smaller portion, for example: separate baselines, IFs and polarisations. This has the advantage of increasing the reliability of the statistics, because RFI may be independent of baseline and the distribution between IFs may differ. This is particularly relevant for L-band (1.3 - 1.8 GHz) observations where the RFI is more problematic.

The visibility data within AIPS is sorted by time and then baseline (TB format). Within each time-baseline data sample the data is further divided by IF, channels, stoke parameters and the real, imaginary and weight of the visibility (re: Section 1.1.1). The SumThreshold method applied in SERPent works on visibility data which is separated by baselines and polarisations and arranged in a 2D array, with the individual time scans and frequency channels comprising the array axes i.e. time-frequency space. The frequency axis is further split by IFs. The idea is that peak RFI and broadband RFI will be easily detectable when the visibility amplitudes are arranged in time-frequency space.

When appending visibilities in the time-frequency space, if the visibility weight is

greater than 0.0 i.e. if data exists for that time and frequency, then the magnitude of the real and complex part of the visibility is taken to constitute the amplitude. The visibility weight is present from cross correlation, where it operates as a noise scaling factor during correlation to account for the different antenna sensitivities in the e-MERLIN array. If the weight is 0.0 or less, i.e. no data exists for this time-frequency position on this baseline, then the amplitude is set to ‘NaN’. This datum has no effect on the sample statistics or threshold value, but acts as a structural substitute for that elemental position within the array, which both AIPS and SERPent require to retain the correct time-frequency information. The Python module NumPy is employed to create and manipulate the 2D arrays, as the module is implemented with performance-optimised Fortran code¹.

There are two concepts associated with the SumThreshold method: the threshold (χ) and the subset size (N). The threshold levels are discussed below in Section 2.2.1. A subset is defined as N number of elements in one direction (i.e. the window is one dimensional) of the 2D array (time or frequency) which is to be tested. Within this subset window, the amplitudes are averaged and if the average amplitude exceeds the threshold level $\chi(N)$, the elements within the subset window are flagged in a separate flag array (which is a float array in NumPy with the same size and structure as the data array). A 0.0, in this flag array denotes a normal visibility, 1.0 signifies RFI in the time direction, 2.0 for the frequency direction and higher values for any subsequent runs of the flagger function.

Once all permutations of the subset size N in the time direction has finished, all flagged elements are then set to the next threshold level. This is a unique feature of the SumThreshold method which differs from normal thresholding methods. The algorithm then proceeds to the next subset size in a specified series and repeats. This subset series increases as $S_N = [1, 2, 4, 8, 16, 32, 64]$, which provides a good balance between flagging performance and computational performance (Offringa et al. 2010a).

Once all the subset sizes in the series have checked for RFI in the time direction, the process is repeated in the frequency direction in exactly the same manner. Running the SumThreshold method in both time and frequency direction constitutes one full run of the algorithm. Subsequent full runs of the algorithm check the flag array for flagged amplitudes which are then set to the next threshold level before the algorithm commences. This removes previously found RFI and helps the algorithm to search for any remaining weaker RFI.

¹It should be noted here that how this module is compiled can have a significant effect on performance.

In addition to the SumThreshold methodology, certain clauses have been added to prevent the algorithm from over-flagging the dataset. If any threshold level reaches the $mean + \alpha \cdot variance\ estimate$, where $0.0 < \alpha < 5.0$, the flagging run for that direction (time or frequency) stops. The default kickout clause added in SERPent is set at $\alpha = 3.0$, with smaller values allowing more flagging and higher values restricting the algorithm.

The flagging process can run multiple times at the cost of computational time, and by default an initial run of subset $N = 1$ only, is included to remove extremely high amplitude RFI. This is followed by two full runs of the algorithm (as described above); the first with a subset size N up to 32 and the second size N up to 256. The subset size N , can be manually selected by the user for optimisation purposes.

The execution of these full two runs is conditional on two factors: that the maximum value within the array after each run is a certain factor of the median, and flags exist from the previous run. On each subsequent cycle, all flagged visibilities from the previous run are set to the next threshold in the visibility array so they don't skew the subsequent statistics and any weaker RFI which may remain can be found. This is necessary because some RFI in the e-MERLIN commissioning data are found to exist over a range of amplitude levels, even as high as 10,000 times the astronomical signal.

2.2.1 Statistical Variance Estimators

The variance of a sample is an important estimator of statistical outliers which represent RFI. Some statistical methods are sensitive to extreme values whereas others are robust against them. A study into a range of methods and various estimators is described and tested by Fridman (2008). The median absolute deviation (MAD) and median of pairwise averaged squares are the most effective estimators that remove outliers, although Fridman (2008) comments that both are not as efficient, (i.e. needs a larger sample population) as other methods. Since the sample size in any given observation from e-MERLIN will be sufficiently large, with over 500 channels per IF and over 100 time intervals for one scan giving $> 50,000$ values, this is not an issue. The breakdown point for MAD is also very high (0.5), i.e. almost half the data may be contaminated by outliers (Fridman 2008). MAD is adopted for this algorithm as an initial statistical estimator of the visibility population because of these robust properties, reducing the bias of RFI on the sample. Again, Fridman (2008) stresses that the type and intensity of RFI, type of observation and the method of implementation are important factors when deciding what estimate to

use for any given interferometer.

The MAD is the variance estimator employed in the SERPent algorithm and is defined by Equation 2.1, where $\text{median}_i(x_i)$ is the median of the original population. This median is then subtracted from every element in the population, creating a new modified sample of the same size as the original. The median of this new population is then calculated and multiplied by a constant scale factor 1.4286 to make this estimation consistent with that of an expected Gaussian distribution (Rousseeuw and Croux 1993; Fridman 2008).

$$MAD = 1.4286 \text{ median}_j \{ |x_j - \text{median}_i(x_i)| \} \quad (2.1)$$

The first threshold level $\chi(1)$ (i.e. when the size of the scanning window $N = 1$) is determined by the median of the sample ($\text{median}(x_i)$), the variance estimator (MAD) and an aggressiveness parameter β as shown in Equation 2.2 (Niamsuwan et al. 2005). Since the median is less sensitive to outliers, it is preferred to the traditional mean in this equation and the MAD to the traditional standard deviation for similar reasons. If the data is Gaussian in nature then the MAD value will be similar to the standard deviation (and the median to the mean). A range of values for β has been tested for multiple observations and frequencies and a stable value of around $\beta = 25$ was empirically found to be used as the default in the algorithm. Increasing the value of β reduces the aggressiveness of the threshold and decreasing the value increases the aggressiveness.

$$\chi(1) = \text{median}_i(x_i) + \beta MAD \quad (2.2)$$

The subsequent threshold levels (i.e. window sizes $N > 1$) are determined by Equation 2.3 where N is the subset value, and $\rho = 1.5$, empirically works well for the SumThreshold method (Offringa et al. (2010a)) and defines how coarse the difference in threshold levels is.

$$\chi(N) = \frac{\chi(N)}{\rho^{\log_2 N}} \quad (2.3)$$

2.3 SERPent

The aim of SERPent (Peck and Fenech 2013) is to provide a script which flags bad radio interferometric visibilities and that is quick and simple to set up and run. SERPent is

written in the Parseltongue programming language, because it is a Python based language which enables AIPS tasks to be imported as modules into the script. The direct application to e-MERLIN and COBRaS, has defined SERPent to interact with AIPS at the front-end and back-end of the algorithm, i.e. to read data into SERPent from AIPS and write the flag extension table (FG table) back into AIPS. Therefore, the Parseltongue programming language is an obvious choice for SERPent. It is also a popular choice for pipelines and is used extensively for European Very Long Baseline Interferometry (VLBI) Network (EVN) calibration.². Another advantage of using a Python based language is that it is independent of compilers, which enables quick easy execution to fulfill the aforementioned requirements. The SERPent software has been made publically available to download³ and has been tested on a number of systems.

SERPent is comprised of two text files; the first is the full body of code which executes all the modules within SERPent, and the second is an input file. This input file is where the user chooses which data to run SERPent on, and the different parameters within SERPent (although defaults are made for the flagging routines). These include: AIPS information variables, number of central processing units (NCPUs) to use, which baselines to flag, the phase calibrator name and flagging parameters (see Section 2.2.1). To demonstrate the ease of input selection, a section of the input file is shown below.

```
# Data Information:

AIPS_user_number = 101      # The AIPS user number the data is on.
Name = '1436+6336'         # The uvdata name of the catalogue.
Klass = 'SPLIT'             # The uvdata klass.
Disk = 1                    # The uvdata disk number.
Seq = 1                     # The uvdata sequence number.

# Parallelization Information:

NCPUs = 4                  # Define here the number of CPUs you want to use.
                             # Parallelization will distribute the jobs amongst
                             # the number specified here.

# Directory Information:

path2folder = '/home/lwp/aips_test/'      # Directory where SERPent
                                           # outputs are written to.
```

²EVN Parseltongue pipeline:

<http://www.jive.nl/wiki/doku.php?id=parsetongue:grimoire>

³http://www.ucl.ac.uk/star/research/stars_galaxies/cobras/technical/rfi

The main flagging parameters in the input file are the aggressiveness (β) and subset size (N) for the first and second full flagging runs and ρ (see Section 2.2.1), and the kickout sigma level where the flagging stops if the threshold reaches the median plus the designated factor of the MAD.

Figure 2.2 shows a logical flow diagram of the SERPent process to explain the many functions within the algorithm. The interaction with other programs such as AIPS, Python pickle files (storage files for Python information) and the input file are highlighted in this figure.

SERPent has been tested on a range of datasets and by multiple institutions around the world. To demonstrate the performance of the several passages within SERPent, a group of test datasets is given in Table 2.1. For ease of reference, each dataset has been assigned a number which will be used throughout this chapter.

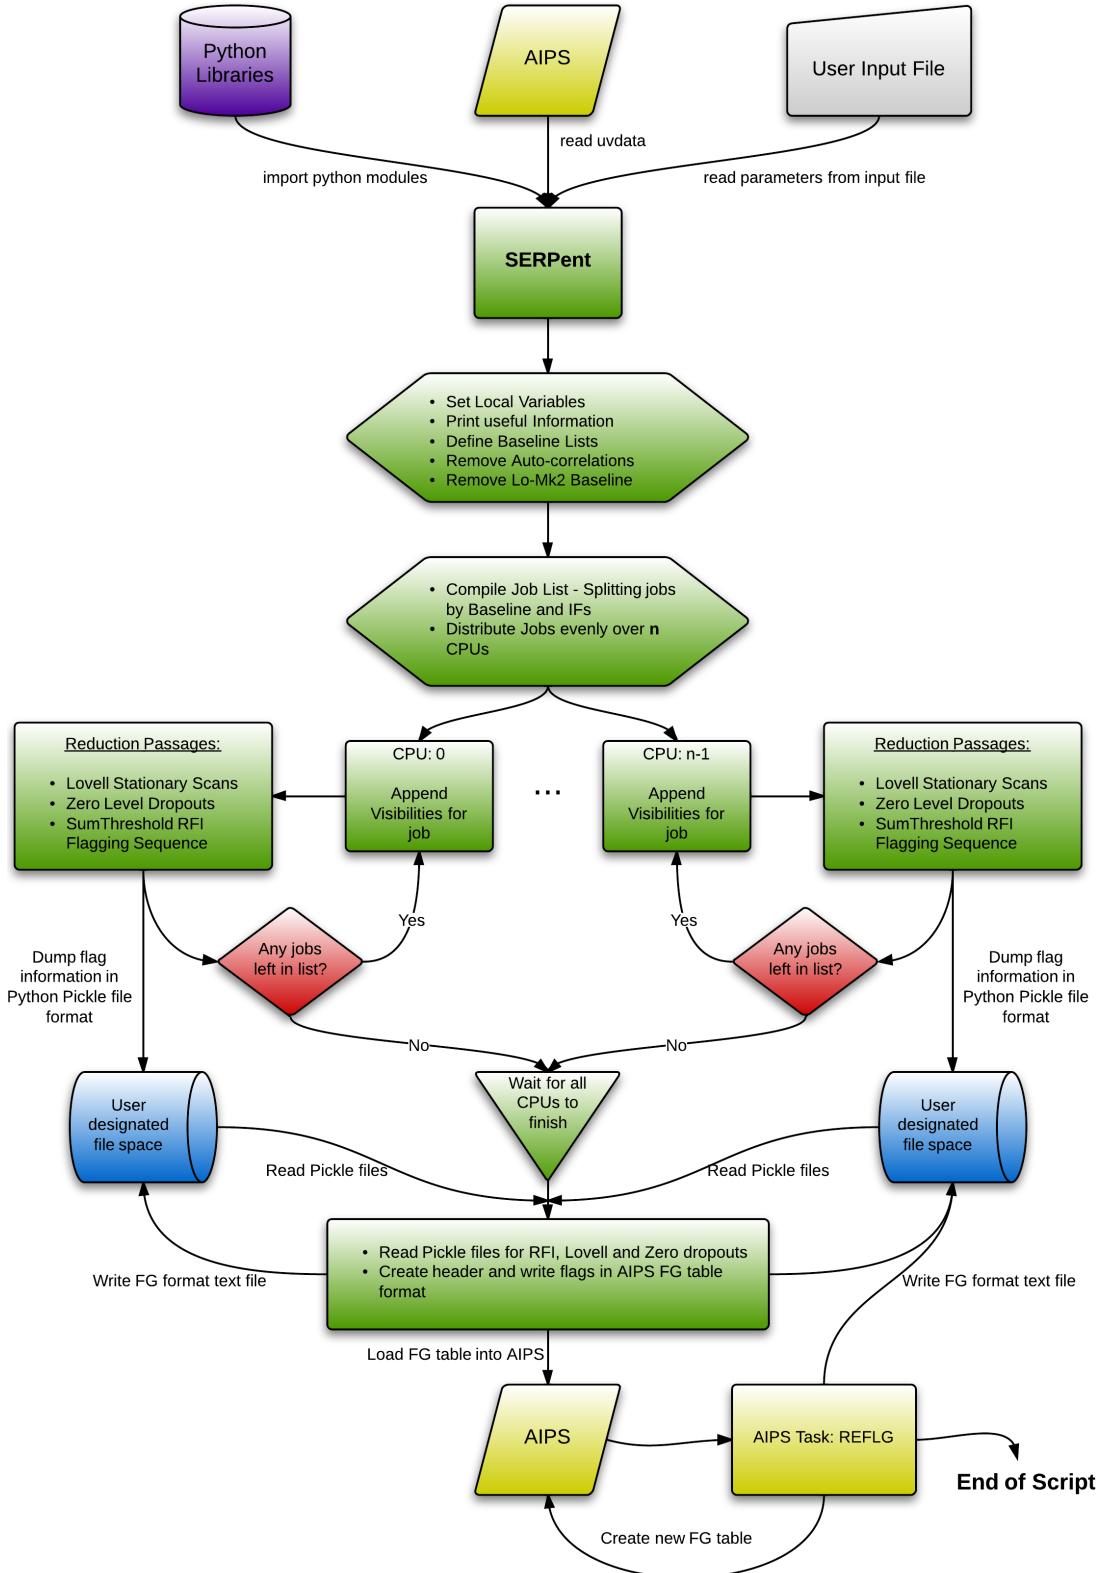


Figure 2.2: A logic flow chart of the SERPENT process.

Table 2.1: SERPent Performance Test Datasets

Dataset Number	Dataset Name	Size (GB)	Frequency Range (GHz)	Band	Number of Time-samples	Duration (Hours)	Baselines	Bandwidth (MHz)	IFs	Channels per IF
RFI Test Data:										
1	1436+6336	1.6	1.32 - 1.70	L	5812	0.17	10	384	12	512
2	1407+284	432 MB	1.63 - 1.69	L	73430	7	1	64	1	128
3	COBRaS W1 2011: 0555+398	2.3	4.41 - 4.92	C	99149	3	10	512	4	128
4	COBRaS 20th April 2012: 2033+4113	27.4	1.36 - 1.74	L	389839	20	21	384	12	128
5	COBRaS 18th July 2012: 0555+398	97	5.49 - 6.00	C	1033940	26	21	512	4	512
6	10.5	5.49 - 6.00	C	112631	2	21	512	4	512	
7	1331+305	3.2	5.49 - 6.00	C	34124	0.5	21	512	4	512
8	1407+284	3.4	5.49 - 6.00	C	35920	0.5	21	512	4	512
9	2007+404	23.7	5.49 - 6.00	C	252745	12	21	512	4	512
9	2032+411	56.1	5.49 - 6.00	C	598520	11	21	512	4	512

Every dataset was observed with e-MERLIN with full circular polarisations (RR, LL, RL, LR). A list of associated sources has been provided here for each dataset.

2.3.1 Lovell Stationary Scan Removal

A bright phase calibrator (phase-cal) is observed for the technique of phase referencing, which is necessary for Very Long Baseline Interferometry (VLBI), in order to provide complex (amplitude and phase) solutions during calibration. This is achieved by alternating scans of the target and phase-cal source.

The Lovell telescope has a slow slew speed in comparison to the other telescopes within the array. This presents a unique problem to the e-MERLIN array. When phase-referencing it only participates in every alternative phase-cal scan, remaining stationary on the target for the other scans. This results in baselines containing the Lovell telescope to have two different amplitude levels for the phase calibrator.

In most cases the phase-cal will be brighter than the target source, thus when the Lovell is observing the phase-cal, the received flux will be greater than when the Lovell does not participate in the phase-cal scan and remains on the target source.

Figure 2.3 shows the visibilities of the phase-cal for the Lovell-Knockin baseline, plotted in amplitude-time. The three windows display; top - before any flagging, middle - after flagging using the Lovell passage, and bottom - after flagging including the zero-level passage. There are two distinct amplitude levels, the highest is where the Lovell antenna contributes to the observation and the lowest is where the Lovell does not contribute.

SERPent detects whether the baseline contains the Lovell antenna and then executes the Lovell stationary scan passage appropriately. It defines each scan by checking whether the time duration between each scan is a factor larger than the integration time. If the average amplitude of all the visibilities within the scan is consistent with being a Stationary scan, it is flagged. This passage is essential for Lovell baselines. If the stationary scans (which make up 50% of the total data) remain, the good phase-cal data would be treated as RFI in the flagging sequence and therefore flagged.

In Figure 2.3 an additional effect can be seen which contributes to the zero-level amplitudes (see Section 2.3.2 for details). However, a careful inspection reveals that this additional zero-level contribution is part of the ‘on-target’ Lovell scan, and not part of the Lovell stationary scan. The antenna has started to receive signal before the antenna has been properly aligned causing in-scan zero-level amplitudes to be observed. These are dealt with in another SERPent passage (see Section 2.3.2 on zero-level dropouts).

SERPent’s Lovell Stationary Scan passage removes the time intervals involved in the

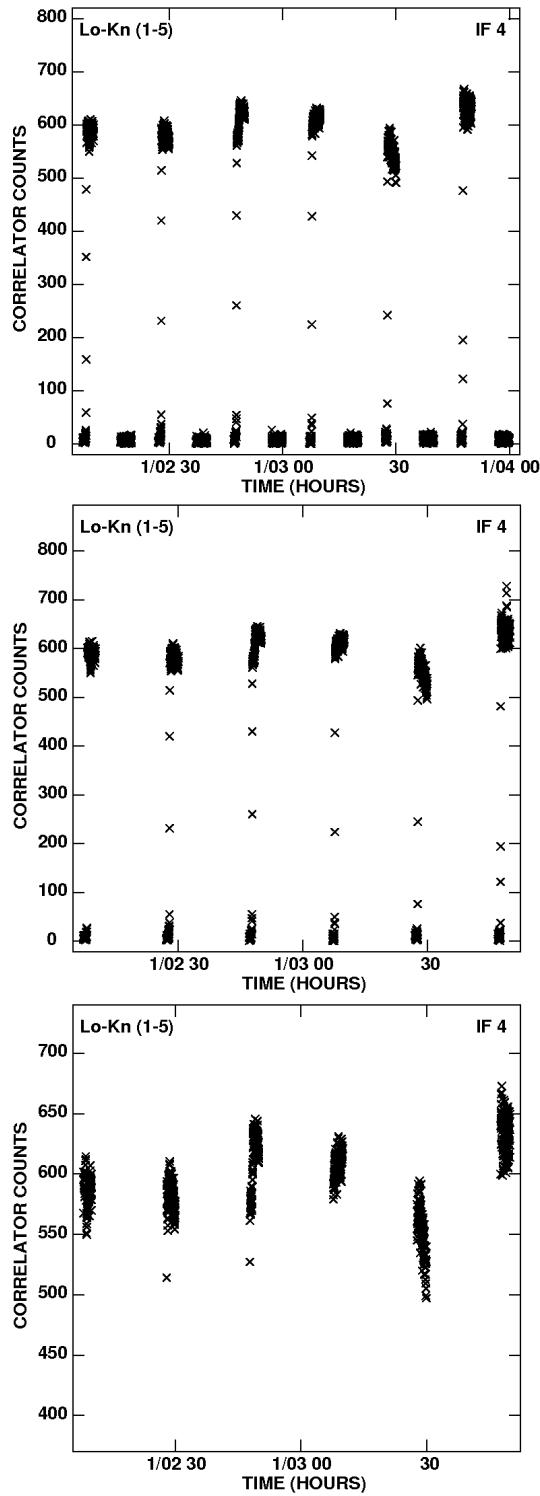


Figure 2.3: Amplitude-time plot with correlator counts on the y-axis and time on the x-axis, displaying a single IF and polarisation for the phase-cal source: 2007+404, baseline 1-5 (Lovell - Knockin) from dataset number 8. The top figure shows the visibilities before any flagging is done. The two distinct amplitude levels can be seen and the Lovell and zero-level dropouts are present. The middle figure shows the same visibilities after the Lovell stationary scan passage. The bottom figure shows the same visibilities after the Lovell and zero-level dropout passages. Both types of dropouts have been successfully flagged.

stationary scans for all channels within the tested IF from the NumPy visibility and flag arrays. A separate Lovell-only flag text file is created, as well as a combined master flag text file. This is done by dumping the flag information into a Python Pickle file, which is later read and combined with other files from other baselines and IFs. This combined flag text file is read into AIPS and attached to the input data as an AIPS flag extension table (FG), at the end of the script.

2.3.2 Zero-level Amplitude Dropouts Removal

Early COBRaS commissioning data revealed bad visibilities in the form of zero-level (visibility amplitudes \sim zero correlator counts) in-scan amplitudes (example: Figures 2.3 and 2.4), possibly a result of a system failure, telescope slew errors or the recording of data before the telescope was actually ‘on-source’. The zero-level amplitudes reside within scans containing good data and therefore need their own passage within SERPent to be flagged because these issues can arise on any baseline. This zero-level passage considers any visibility within all scans and it therefore does not matter where these zero-level amplitudes occur. It is expected that this effect will most likely occur either at the beginning or end of the scan.

The zero-level dropout separates the data in the same manner as the SumThreshold algorithm i.e. by baseline, IF and polarisation, and calculates the median of this sample. Then a series of tests are made on each visibility to determine whether it lies within a threshold i.e. beneath the median level minus a factor of $3 \times$ MAD, or beneath 0.0 amplitude plus a factor of $2 \times$ MAD etc. The breakdown limit of this process is when more than 50% of the data is at the zero-level amplitude.

In addition to the zero-level passage described above, it is useful to trim the very edges of every scan, because SERPent can miss a few visibilities which are in transition between the zero-level dropout and on-source amplitude levels. The AIPS task QUACK can be implemented for this job for a very short section of the scan ($\sim 5s$), and has the option to be implemented in the COBRaS calibration pipeline (see Section 3.5) after a full run of SERPent.

Figure 2.3 shows the visibilities in an amplitude-time plot, after the Lovell (middle) and zero-level (bottom) passages have been performed. It can be seen that both have removed low-level off-source amplitudes which would have affected calibration and RFI mitigation.

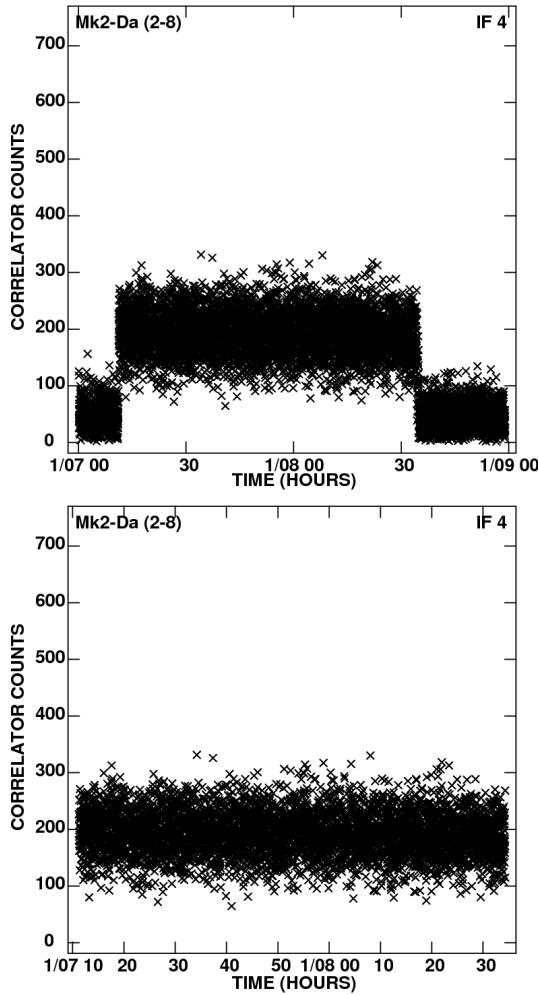


Figure 2.4: Amplitude-time plot with correlator counts on the y-axis and time in hours on the x-axis, of the source: 0555+398, baseline Mark II - Darnhall (2-8) from a single IF and polarisation from dataset number 5. The top figure shows the visibilities before the zero-level passage with two distinct amplitude levels at the beginning and end of the observation. The bottom figure shows the visibilities after the zero-level passage has been executed. The previous zero-level dropouts have been successfully removed.

To demonstrate the power of the zero-level dropouts passage, and its ability to remove dropouts from anywhere within the scan, dataset number 5 contains zero-level dropouts at the beginning and end of the scan and also contains a few minutes of the previous source scan. This reinforces the idea that the zero-level dropouts result from telescope slews or from the correlation. Figure 2.4 demonstrates that SERPent's zero-level passage can deal with dropouts at the beginning or end of the scan after the successful flagging of these low amplitudes.

2.3.3 RFI Mitigation Performance

As discussed in Section 2.1, RFI originates from a variety of sources. Some of the origins of RFI for e-MERLIN are known e.g. CCTV interference in L-band (1376 MHz), but others can be unpredictable, and neither are mitigated at the antenna or at the correlator level before data processing.

SERPent has been tested on both L-band (1.3 - 1.8 GHz) and C-band (4 - 8 GHz) observations (see Table 2.1 for datasets) which contain different amounts and types of RFI. L-band is typically more noisy with both broadband and narrowband RFI common in observations, whereas C-band is generally RFI quiet with only some narrow RFI present in observations (although broadband RFI has been seen).

The edges of the IFs often contain noise as a result of the reduced response of the bandpass. SERPent can detect and flag this because it behaves in the same way RFI does. We now present a series of before and after figures which depict SERPent's flagging ability on a range of e-MERLIN datasets.

Figure 2.5 displays the dataset number 9 (centred on 5.75 GHz), with the visibilities sorted in time along the y-axis and channels in frequency along the x-axis, with all four IFs side by side. There is some weak narrowband and broadband RFI in the central channels and some noise present at the edges of IFs 1 and 4. Figure 2.6 shows the same data after SERPent flagging. All of the narrowband and broadband RFI and IF edge noise has been detected and successfully flagged. This level of RFI detection and flagging is more accurate and delicate than what can be achieved with visual, manual flagging.

Dataset number 4 (centred on 1.56 GHz, with 12 IFs) provides a greater test of SERPent's flagging capabilities because of the increased incidence of RFI at L band. Once again the presence of narrowband and broadband RFI can be seen in Figure 2.7. There is in fact more RFI present at lower levels, but this can not be seen in the spectral window before flagging. Note, IF 9 (1.61 - 1.64 GHz) has been automatically flagged by the correlator, before any processing of the data has been done.

Figure 2.8 shows the L-band data following flagging by SERPent, again demonstrating the intricate nature of RFI detection by finding strong and weak RFI, as well as RFI which encompasses both large and small areas in the time-frequency space. Flagging to this level of accuracy on large datasets by hand would take an unfeasible amount of time.

There are examples of more complex RFI in the commissioning datasets from e-

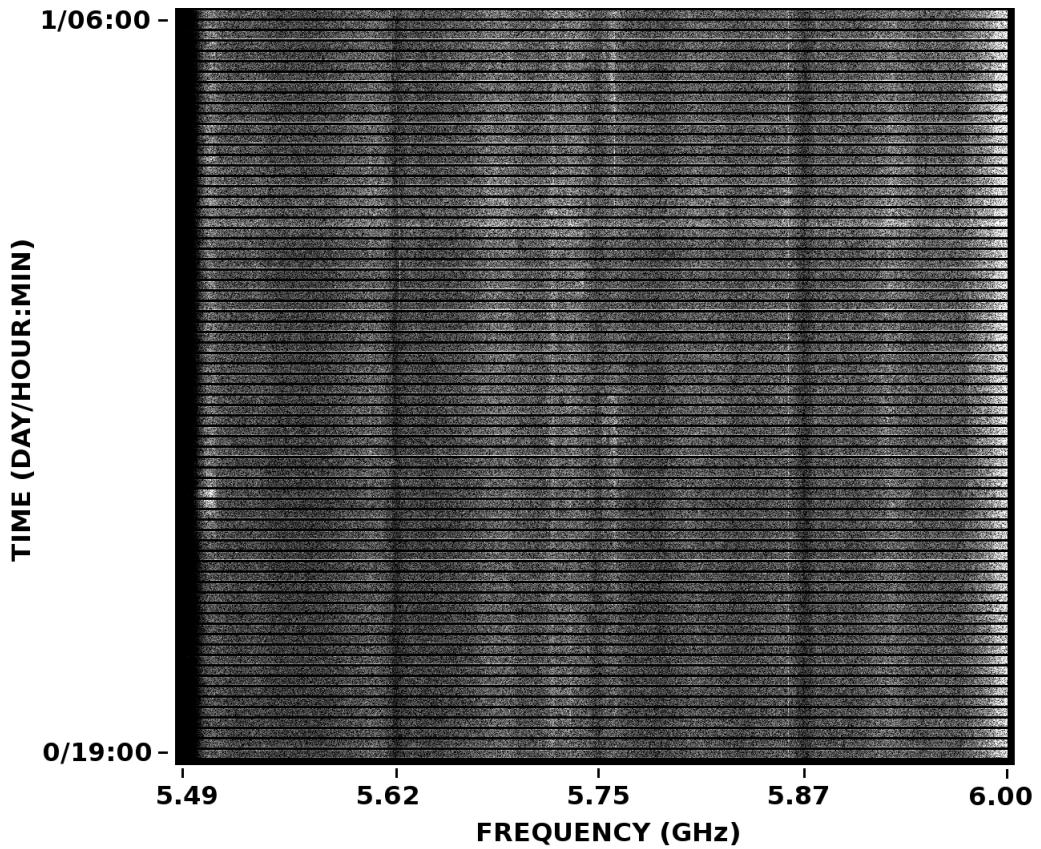


Figure 2.5: Time-frequency plot of the visibilities of the target field: 2032+411 from dataset number 9. Four IFs are plotted together with the bandwidth of 512 MHz from 5.49 to 6.00 GHz, from the Defford-Cambridge baseline. Weak narrowband and broadband RFI are present and noise in the edges of some IFs can also be seen.

MERLIN, varying in time and frequency simultaneously. The noisy COBRaS 2011 dataset at frequency 4.412 GHz and source: 0555+398 shown in Figure 2.1, demonstrates some RFI which varied over time and frequency. As stated before, thresholding methods are the most robust way to detect these unusual types of RFI, and Figure 2.9 displays how SERPent can deal with RFI of this nature.

One further example of some peculiar multiple RFI found in e-MERLIN commissioning datasets can be seen from the source 1407+284, on the baseline 1-8 (Lovell - Darnhall) in Figure 2.10. This RFI, of unknown origin, seems to drift in frequency over time and not necessarily in a constant direction. The before and after time-frequency plot in Figure 2.10 shows the complex shape of this RFI and how SERPent again has successfully flagged all of it.

These are only a small selection of examples from the commissioning e-MERLIN archives, but demonstrate the unpredictable nature of RFI and how thresholding detection

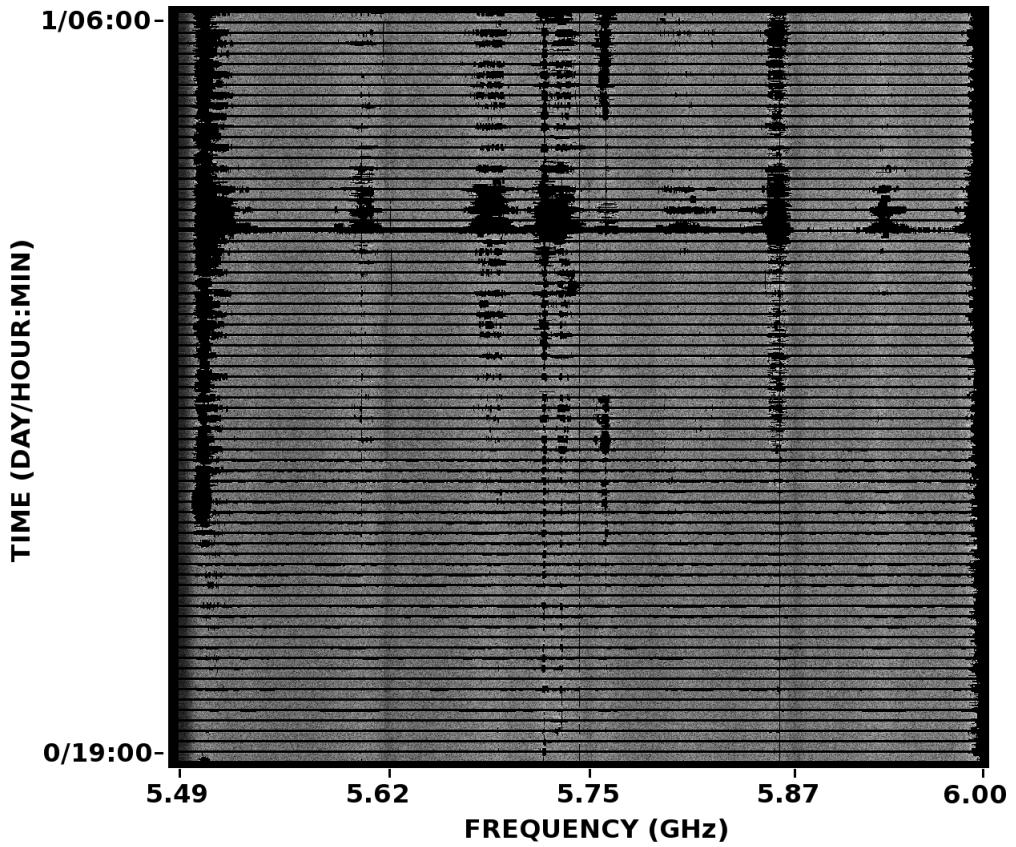


Figure 2.6: Time-frequency plot of the visibilities of the target field: 2032+411 from dataset number 9. Four IFs are plotted together with the bandwidth of 512 MHz from 5.49 to 6.00 GHz, from the Defford-Cambridge baseline. All of the visible narrowband and broadband RFI and the noise in the edges of IFs 1 and 4 has been flagged by SERPent.

methods can find RFI of any morphology. SERPent can easily convert this information into a readable AIPS FG table which is automatically appended to the input data in AIPS as part of the script.

2.3.4 Computational Performance

One important criteria for automated flaggers is computational performance. We have analysed the computational performance of SERPent on a number of computer systems, the details of which are given in Table 2.2. The difference in number of processors, Central Processing Units (CPUs) per processor and memory size covers a range of modest specifications available to institutions across the world (please refer to Table 2.2 for details).

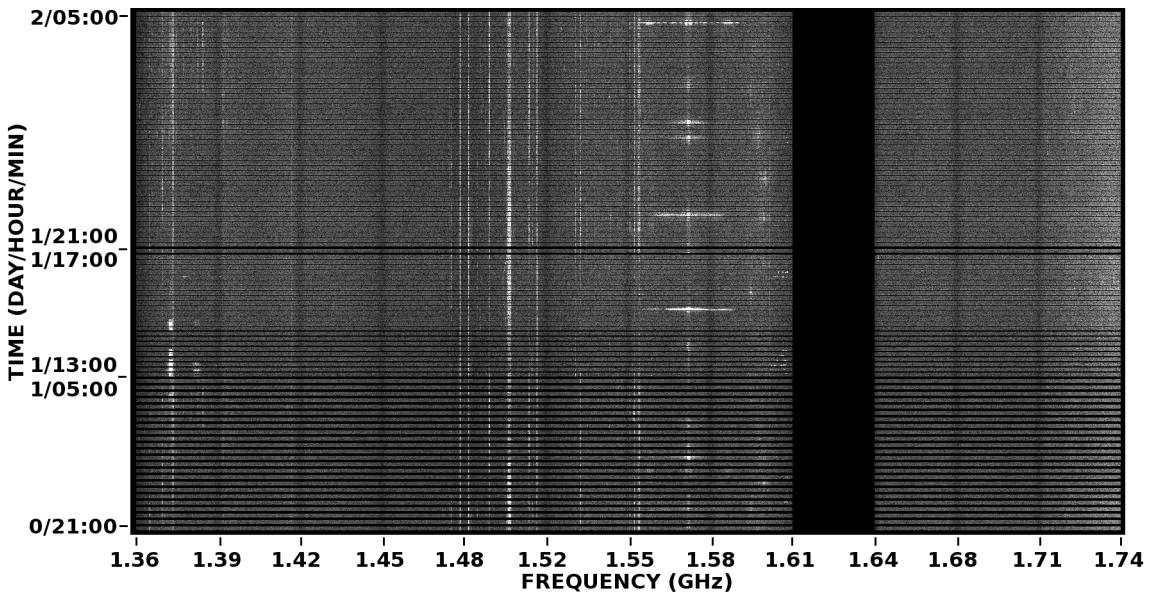


Figure 2.7: Time-frequency plot of the visibilities of the target field: 2033+411 from dataset number 4. Twelve IFs are plotted together with the bandwidth of 384 MHz from 1.36 to 1.74 GHz, from the Defford-Darnhall baseline. A variety of narrowband and broadband RFI can be seen, and many more weaker RFI are present but are below the current contrast levels, once the stronger visible RFI is removed, the weaker RFI is revealed. Note: IF 9 has been flagged by the online correlator before post-correlation reduction and processing.

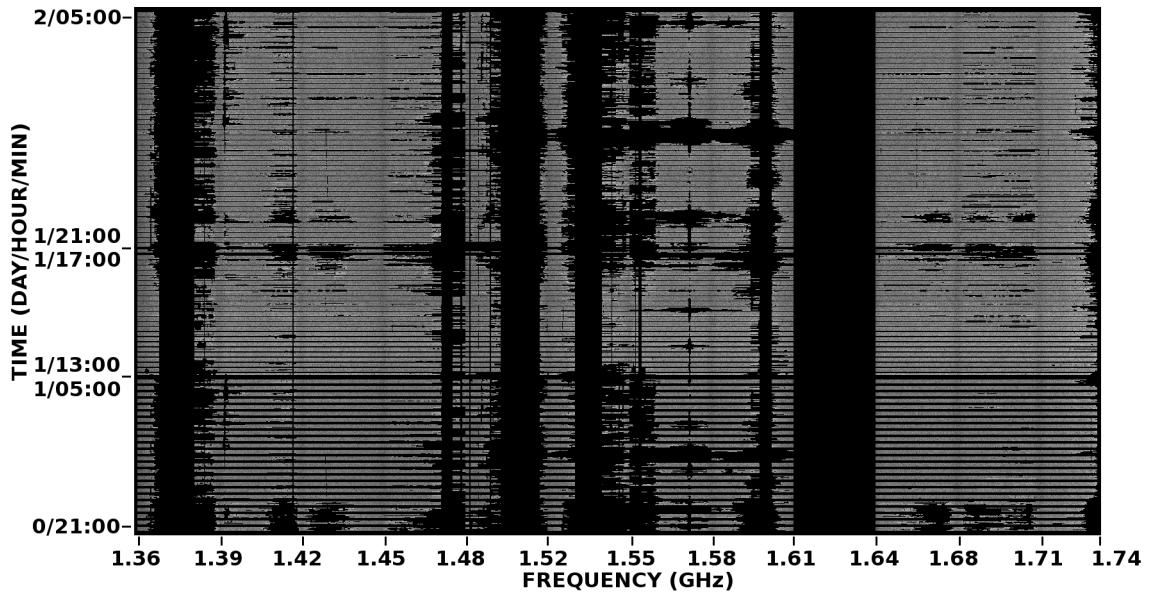


Figure 2.8: Time-frequency plot of the visibilities of the target field: 2033+411 from dataset number 4. Twelve IFs are plotted together with the bandwidth of 384 MHz from 1.36 to 1.74 GHz, from the Defford-Darnhall baseline. A lot of strong narrowband and broadband RFI has successfully been flagged, along with weaker RFI which was not visible in Figure 2.7.

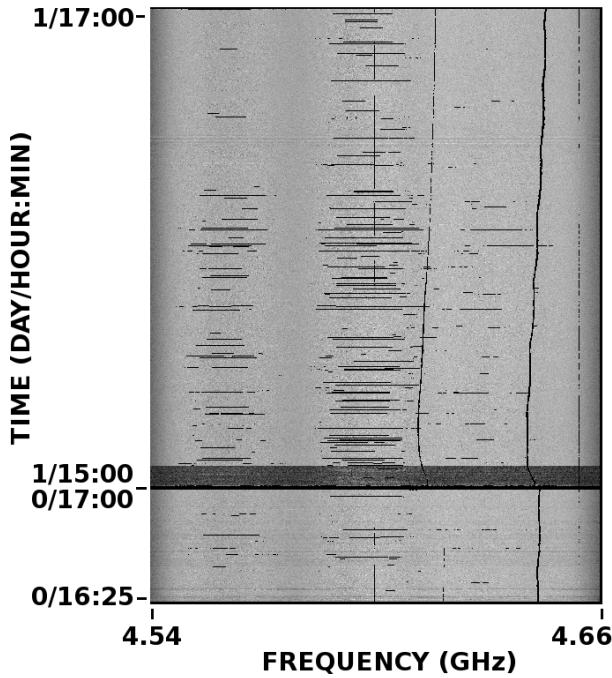


Figure 2.9: Time-frequency plot of the visibilities of the source: 0555+398 from dataset number 3. A single IF and RR polarisation is shown with a frequency range from 4.54 to 4.66 GHz from the baseline Knockin-Pickmere. The before image can be seen in Figure 2.1. After a run of SERPent, the time and frequency variant RFI has been flagged successfully.

Table 2.2: Computer Systems

Computer Name	Memory (GB)	Processor (GHz)	NCPUs
Leviathan (1 node) ¹	100	3.20	16
Kria ²	40	2.93	24
Cornish1 ²	16	3.20	8
Megan ³	48	2.40	16

Systems at: 1: University College London, UK. 2: University of Manchester, UK. 3: Netherlands Institute for Radio Astronomy (ASTRON), Netherlands.

To increase computational performance, SERPent is parallelised by splitting the data into ‘jobs’ which are evenly distributed across a number of CPUs. SERPent is parallelised in both baselines and IFs to maximize the even spread across CPUs and uses a user-designated number of CPUs specified in the input file. Initial tests on modest data sizes revealed a significant increase in performance. These tests also show that the processing

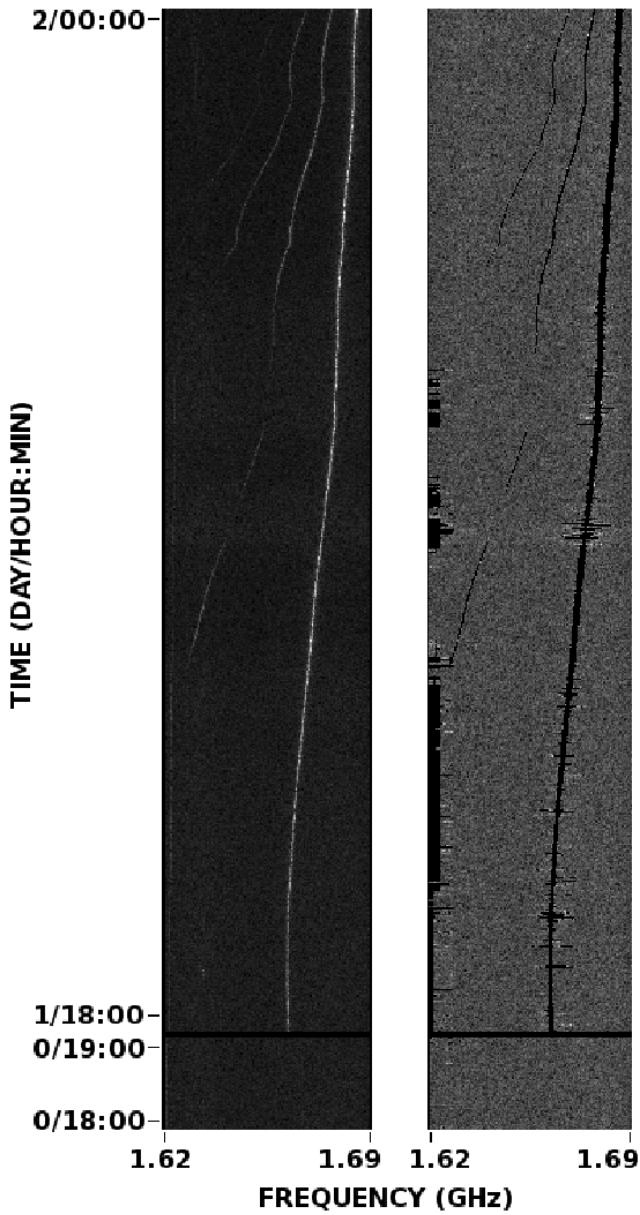


Figure 2.10: Time-frequency plot of the visibilities of the source: 1407+284 (dataset number 2). A single IF and RR polarisation is shown with a frequency range from 1.62 to 1.69 GHz from the baseline Lovell-Darnhall. Left; is the before image where the RFI varying in amplitude over time and frequency can be clearly seen, Right; is the clean, post-SERPent flagging image. Note that the contrast levels of the normal (unaffected) visibilities are different in each plot due to the influence of the RFI skewing the contrast levels.

time scales linearly with the data volume.

The effects of memory and the number of CPUs (NCPUs) used on the computational performance is assessed here, by testing SERPent on dataset number 1, 1.63 GB (see Table 2.1 for details). Figure 2.11 portrays the relative performance ratio to a single

CPU on the same system. All systems have linear relations with a peak CPU efficiency achieved at around 8 CPUs. At this point adding more CPUs still increases performance but at a slower rate. We can infer that using 10 CPUs on this dataset has increased the performance by a factor ~ 7 compared to using only 1 CPU on the same system. Runs on other datasets gave similar performance results.

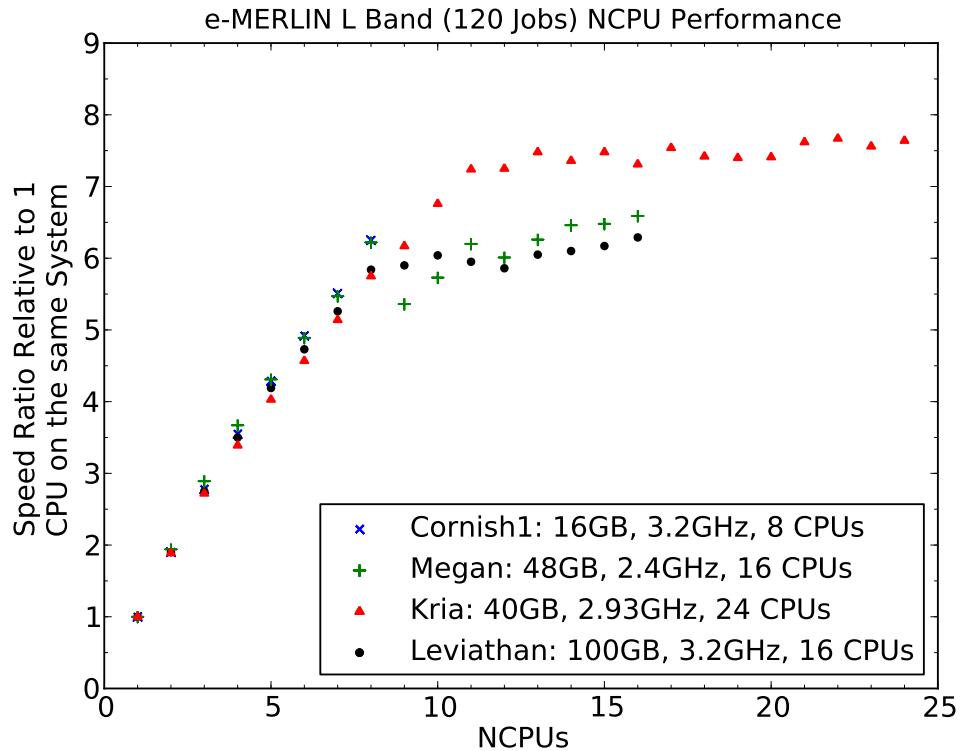


Figure 2.11: The speed relations of running SERPent on multiple CPUs on a range of computers relative to the performance of a single CPU on the same system. Even at high number of CPUs there are significant gains in performance which should increase further beyond 24 CPUs with datasets with a greater number of jobs available.

The tail-off in Figure 2.11 may result from this particular dataset, where a few IF-baseline combinations suffer from severe RFI (L-band observations). These jobs take more time to process, and as the number of jobs per CPU decreases, the portion of total time taken becomes biased towards the time taken by these ‘heavy’ jobs. This is because all other CPUs have finished processing and are waiting idle whilst the CPU with the ‘heavy’ job is still processing. Therefore, the performance relative to 1 CPU is affected by these jobs. However, the performance relative to 1 CPU is expected to increase further for a dataset with a higher number of total jobs (where the data has been distributed in a larger set of smaller segments), because the influence on the total computational time by

any lengthy jobs is minimised. This, and the (random) distribution of jobs, is the reason a turnover after 8 CPUs is seen. This could also explain the fluctuations in performance on certain systems after 8 CPUs.

Moreover, the limiting factor in the overall runtime of SERPent is calculating the averages of each window subset, which has been optimised for each subset N^4 , but is performed a large number of times. This number is estimated as follows. The number of polarisations, IFs and baselines, for Legacy data are 4, 16 (C-band), and 20 respectively. With two full runs with S_N going to $N = 32$ for the first run and $N = 256$ for the second run, resulting in 30 S_N cycles per observed source. Assuming 512 channels per IF (reasonable assumption for COBRaS) this results in the averaging function being called $\sim 2 \times 10^7 \times$ number of integration times of observation, per source.

Increasing the amount of memory of each CPU also increased the computational performance, albeit by a smaller factor than the parallelisation. Comparing computers with the same processing speed (Leviathan and Cornish1 both have 3.2 GHz processors), Leviathan has 6.25 GB memory per CPU, and Cornish1 has 2 GB memory per CPU. Figure 2.12 shows that the amount of memory per CPU decreases in significance as the number of CPUs increases from a factor of 1.22 for 1 CPU to 1.14 when running on 8 CPUs. This is because the effect of parallelisation on performance is greater than the benefit of having extra memory per CPU. This shows that the limiting factor of running SERPent on interferometric datasets is the shear volume of data that needs processing over a number of CPUs and not the result of a lack of memory.

The raw (unaveraged) COBRaS July 2012 C-band data (datasets 5 - 9) (97 GB) takes 20 hours to process with SERPent, yielding a flagging rate of ~ 110 GB/day. When the same dataset is averaged to 25 GB, SERPent takes ~ 6 hours with Leviathan (100 GB Memory and 16 CPUs), which is consistent with linear scaling in data size and time. This approximately results in a processing rate of 6.9 GB CPU $^{-1}$ day $^{-1}$. These extrapolations may vary in actual performance due to other factors such as the number of jobs SERPent creates, which is dependent on the number of baselines and IFs in any observation. The amount of RFI will also affect performance, as less RFI means SERPent can skip flagging runs due to kickout clauses in the flagging sequence etc. However, these remain reasonable

⁴For $N = 1$ the array element is simply tested against the threshold i.e. no averaging function is used. For $N > 2$ the Python `.sum()` function is found to have a flat performance, i.e. the time to calculate the summation of a slice (or window subset N) from the NumPy array is independent of the size of the slice. This function is found to be faster than all other summations and average functions.

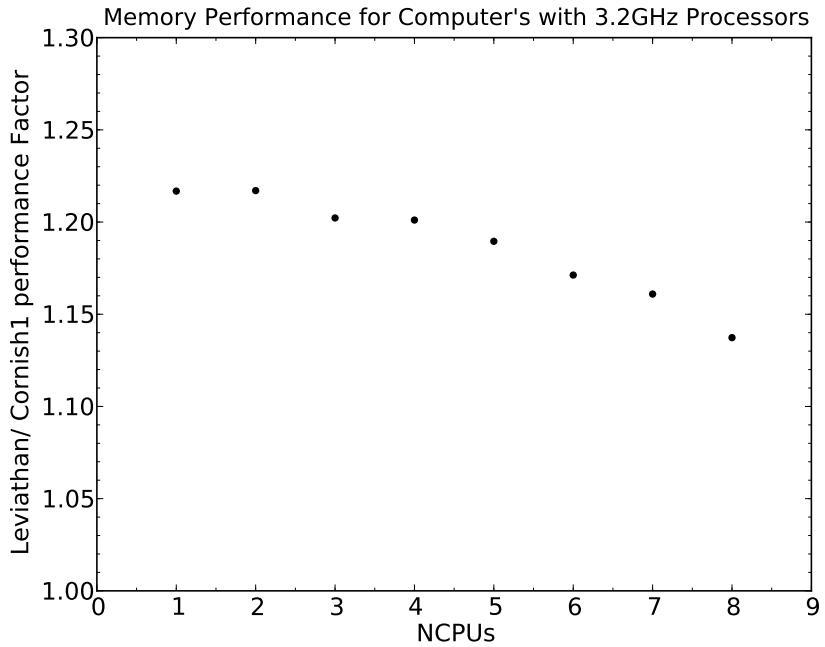


Figure 2.12: The performance relations of Leviathan/ Cornish1, both with 3.2GHz processors. As the number of CPUs increases the impact of Leviathan’s larger memory per CPU on computational performance decreases.

estimates for predicted performances.

To process a 1 TB dataset in a day, the user will require ~ 145 CPUs processing at $6.9 \text{ GB CPU}^{-1} \text{ day}^{-1}$ (rate taken from processing COBRaS July 2012 C-band dataset; dataset numbers 5 - 9). As discussed before, further increase in the number of CPUs won’t result in an increase in performance because the job/CPU factor is the limitation in the parallelisation for the COBRaS July 2012 C-band dataset. However, an increase in the number of jobs because of an increase in the number of IFs for full e-MERLIN Legacy data will provide an increase on the factor 7 seen in Figure 2.11 from the parallelisation.

It would be simple to parallelise even further in polarisations, as currently every polarisation for each baseline and IF are contained within the same ‘job’ but processed separately by the flagging sequence. This would potentially increase the number of jobs by a factor of 4 (for full polarisation studies). However, only computers with a high number of CPUs ($\text{NCPUs} > 100$) would predominantly benefit from this, in addition to the increase in the number of jobs resulting from the increase in bandwidth. There is scope to change the parallelisation method from using Python fork processes to utilising Python’s multiprocessing module to provide a more even distribution of jobs, and this is discussed in Section 2.5.

2.4 Conclusions and Discussions

2.4.1 Conclusions

SERPent automates the reduction, flagging and preparation procedures of post-correlated radio interferometric datasets, specifically those from e-MERLIN. SERPent is in the process of being tested on EVN and Global VLBI datasets, showing good early results. SERPent is written in Parseltongue, a common scripting language utilised prominently with the EVN, so that the user could start flagging data which has been loaded within the AIPS environment with relative ease. SERPent can be easily added to existing and future pipelines.

The entire SERPent program consists of only two text files. The first is the main SERPent code to be executed, and the second is a user input file designed so the user does not have to interact with the main body of code. The input file also has the benefit of making the input parameters obvious and therefore intuitive to set. This gives the freedom to the user to pursue their own flagging philosophy, i.e. whether they want to be aggressive or conservative with the flagging, but also includes a set of default inputs which will perform well on most datasets.

SERPent is designed to be run on ‘high-end desktop’ computer systems. The examples in this thesis used a system with 16 CPUs and 100 GB of Memory (Leviathan) and was flagging at ~ 110 GB day $^{-1}$. This throughput will increase with full e-MERLIN Legacy data as the number of ‘jobs’ will increase with full bandwidth, providing a higher throughput with a higher number of CPUs. It is unlikely that one will be able to process full e-MERLIN *legacy* data on a modest desktop computer. Although obvious advantages in increased computer facilities and real world limitations on smaller systems are apparent, SERPent can be used by institutions without access to super computer clusters.

Section 2.3 has demonstrated that SERPent can reduce and flag current e-MERLIN commissioning data, which will have many more complications than a stable fully commissioned e-MERLIN including the Legacy datasets. The benefit of using real data instead of simulated data is obvious, and SERPent is now part of the official pipeline for e-MERLIN, used at Jodrell Bank and other international groups.

In the wider context of this thesis and indeed the astrophysics community, more general conclusions can be gleaned from the experience of creating and developing SERPent. When testing and verifying the performance of any piece of software, real data is imperative. No

matter how many times some software is run on simulated data or even real observations, a new real dataset will almost certainly be different. This may lead to the discovery of bugs or a change in performance. This statement is magnified in the case of e-MERLIN where much of the upgrade process was conducted without any modelling of antenna responses etc. during the commissioning phase.

During the first release of SERPent to the international community, the e-MERLIN correlator's behaviour changed and the introduction of previously absent NaN's or empty visibilities from the correlator affected the performance on all of SERPent's passages. The only way to minimise the possibility of this is through the testing on as many real datasets as possible, to make the software robust and reduce the chance of failure on any given future run. In other words, you can never anticipate everything a real dataset (or a small selection of datasets) will throw at you.

Another conclusion from developing software like SERPent is to get people to use any piece of new software it has to have one or more of the following characteristics: it must be simple to use and run, contain a method which is more advanced than current methods/functions, be robustly tested to minimise failures, be modular or integrated with current packages and formats so users don't have to learn or modify existing code.

SERPent ticks most, if not all of these points, with the only limitation being the testing issue above, as SERPent is stable and working well, but is still a relatively new piece of software. Otherwise, being an algorithm written in Python/ Parseltongue with only two files, it is free of compilers and easy to use and interact with. It also benefits from Python's modularised nature and is part of COBRaS (see Chapter 3), and e-MERLIN's calibration pipeline. Finally, SERPent uses the most advanced existing method of RFI mitigation via thresholding (studies by Offringa et al. 2010a), along with passages dealing with e-MERLIN specific issues.

2.4.2 Discussion

When constructing an automated flagging script, the flagging philosophy has to be considered and decided. Whilst flagging all of the RFI and flagging none of the data is the idealistic scenario, even with implementing the SumThreshold Method with an extremely low false-positive detection percentage, either some RFI will remain or some good data will be flagged. This is the reality of working with real datasets from imperfect instruments and environments.

Following discussions at the e-MERLIN early science meeting (Manchester, 11th - 12th April 2014), some of the strong narrowband RFI in the L-band observations contaminate neighbouring channels via ringing. This has been seen to affect the entire IF at certain times with RFI. Suggestions of applying Hanning smoothing to the channels before running SERPent may reduce the impact of ringing and increase the performance of SERPent and calibration (Simon Garrington; private communication).

There is a philosophy which states ‘no data is better than bad data’ (a comment made on data editing at the 13th Synthesis Imaging Workshop at Socorro, May 2012), promoting aggressive flagging, while others who would rather flag 80-90% of RFI and have some of the weaker, lesser RFI remain (Rob Beswick; commenting on SERPent at the RadioNet Advanced Radio Astronomy workshop in Manchester, November 2012). Obviously both strategies can not be accommodated in total automation, therefore SERPent has the option for the user to decide some of the flagging parameters. These parameters include the aggressiveness (β), subset sizes (N) and kickout thresholds. The AIPS REFLG task has also been seen to over-flag at times, although it is necessary to condense the number of rows in the AIPS FG table. For further discussion on REFLG, see Section 2.5.

The computational performance of SERPent is probably the area which requires most improvement and future plans are in place to improve this (see Section 2.5). It currently flags ~ 110 GB/day with 16 CPUs, which is reasonable for commissioning e-MERLIN datasets. However, for e-MERLIN Legacy data this will be slow. It is obvious that including more CPUs could solve this problem, as 16 CPUs is still very modest in modern computing terms, however this is merely shifting the problem onto hardware (and isn’t very constructive). The flagging sequence makes two full passes through the SumThreshold method (the original AOflagger; Offringa et al. 2010b makes 5 passes) in order to maximise RFI detections, and skips these passes if the threshold level is low enough. This is currently the limiting factor in terms of performance. Reducing this to one full pass would speed SERPent up considerably at the expense of RFI mitigation performance. Note that the amount of RFI also affects computational performance, because more RFI means more full runs completed within SERPent, and less RFI means more cycles are skipped due to the invoked kickout clauses implemented in the flagging sequence to stop over flagging and increase speed performance.

Comparing SERPent with flagging implementations on the JVLA and LOFAR, the data volume per processing time appears to be slower. In the case with LOFAR, the

AOflagger has been written in a high-level language (C++) and includes specific compiler settings to achieve the optimal performance (Offringa 2012). In addition, the AOflagger is heavily parallelised over multiple cores and nodes on a super cluster, vectorised, and is part of the LOFAR pipeline which fully reduces and calibrates observations for users. This is different in the case of e-MERLIN, where the data will still be in a raw format when presented to the user, who will not have access to the same computing facilities as LOFAR. There is work currently being conducted on a general e-MERLIN pipeline, and SERPent is the flagging software implemented for the reduction passage. However, this is only a general pipeline and does not account for the many calibration techniques and methods needed for the many diverse projects e-MERLIN will observe for.

In the case of the JVLA, there is no implementation that is as sophisticated in mitigating RFI as the AOflagger or SERPent methods. The CASA software package is the main choice for the JVLA, and all developments are focused to this package. On the contrary, e-MERLIN currently favours AIPS because the ability of fringe fitting (the calculation of delay, rates and phase offset solutions of each antenna because residual errors remain after cross correlation from the correlator) exists within the program and is required to calibrate e-MERLIN data.

According to feedback received from users, SERPent can be rather aggressive at times. Whilst differing flagging philosophies can account for these views, it should also be considered that e-MERLIN is not a completely settled system, with noticeable improvements in data quality output from month to month. For example, there have been filter issues with some of the COBRaS April 2012 L-band datasets which have since been resolved (May 2013 L-band dataset), but caused amplitude level issues which then affected RFI mitigation performance. e-MERLIN is a heterogeneous array whose antennas have other responsibilities outside of e-MERLIN (Lovell and Cambridge partake in EVN observations). Compared to other, dedicated arrays such as the VLA/ JVLA and ALMA, both homogeneous (ALMA has 2 types of antennas) arrays which have been modelled extensively before commissioning. This provides a much smoother transition from the commissioning to fully-commissioned phases for the JVLA and ALMA. These factors should not be over looked with respect to e-MERLIN commissioning and early Legacy datasets, because both hardware and software changes make maintaining external software such as SERPent difficult.

The amount of RFI in the test datasets is variable but usually contains a high amount

of RFI. The kickout clause was added to ensure that good data was protected from the flagger as more aggressive flagging was required. Inadvertently this variable can set the overall aggressiveness by allowing the longer flagging runs to flag close to the median of each sample. The aggressiveness parameter β has stayed constant since the inclusion of the kickout clause. β may need to be automatically set depending on spread statistics (i.e. how many different levels of RFI there are) rather than empirically set, to achieve optimal flagging.

If β is large enough lowering the aggressiveness, then the kickout clause should not be envoked, but in the regime of aggressive flagging with a low β value it can have an effect on how much is flagged. In essence, β sets the first threshold level, ρ determines the difference in the subsequent threshold levels $\chi(N)$ (i.e. a function of the number of subsets N) and the kickout clause determines when to stop the algorithm if $\chi(N)$ gets too close to the median.

Furthermore the tweaking of SERPent flagging parameters may still yet yield the most optimised settings for both flagging and speed performances. The best time to conduct and hone these settings will be once e-MERLIN has settled and finished its commissioning phase (see Future Work).

2.5 Future Work on SERPent

Software can always be improved, and SERPent is no exception. One of the weaknesses of SERPent, highlighted in Section 2.4.2, is the computational performance.

Parallelisation has made massive improvements in the performance. However, the adopted method of parallelisation is not the most efficient. The Python fork process parallelisation, pre-allocates jobs to CPUs before reducing them. This results in a plateau in performance at higher NCPUs, because of the ‘heavy’ jobs issue (Figure 2.11). An alternative exists in Python 2.7, which contains a multiprocessing module. A recent study on Python parallelisation approaches was conducted by Singh et al. (2013), who showed that the Python Process/ Queue method from the multiprocessing module performed better than other methods from the same module and the Parallel Python module. The Process/ Queue method creates a queue of jobs from which CPUs can access and take as and when they’ve finished processing a previous job which is added to a second queue containing finished jobs. Once all jobs are finished the second queue can be read for post-

processing. This will create an even performance increase by removing the performance bias from CPUs with a heavy jobs.

Another weak link in SERPent is the behaviour of the AIPS task REFLG, which has not only been too aggressive at times, but has recently (November 2013) been seen to create more flag entries than the input table (reported by a number of people; Anita Richards; Rob Beswick; Megan Argo; Danielle Fenech etc.). This is in complete contradiction to one of its primary functions which is to condense the FG row entries. Investigations are ongoing to understand the reason(s) for this behaviour. Despite this, AIPS has an internal limit to how many FG rows can be read in (10 million), which is not always sufficient. A passage which does this job outside of AIPS (i.e. within Python/ Parseltongue) has been considered as a possible solution.

Following further discussion at the e-MERLIN early science meeting, the L-band RFI may be time dependent. Low declination observations (such as COBRaS $\sim 40^\circ$) will suffer from heightened levels of RFI from horizon sources and another big contributor to RFI has been identified as Manchester airport, just North of Jodrell Bank. Investigations will start in the future to determine and indentify different sources of RFI at L-band, which is the first step in order to mitigate the level of RFI during any particular observing run.

Empirical studies are currently underway at UCL and the University of Manchester, UK, to find the optimal flagging parameters for SERPent, for each antenna or baseline. This is an important investigation which will affect both the flagging and computational performance of SERPent, as some antennas are more stable than others within the array and differ in the amount of RFI in the local environment.

An e-MERLIN L-band RFI mask has also been created (Megan Argo, private communication), in the form of an FG table, which is applied to the dataset before being processed by SERPent. Whilst this can be added in the calibration pipeline, efforts may be made to have the option to introduce an FG mask into SERPent. The obvious benefits to this is that any known RFI will be removed, improving SERPent's ability to remove the rest. This RFI mask has been implemented on recent L-band datasets presented in Chapter 3.

Chapter 3

Radio Interferometric Calibration Pipeline for COBRaS

The only true wisdom is in knowing that you know nothing.

Socrates

This chapter addresses the interferometric calibration procedures for the COBRaS pipeline, which follows from the SERPent module (Chapter 2) for the initial reduction and flagging. This chapter also includes discussions on wide-field, wide-bandwidth imaging with e-MERLIN and mosaicing techniques with a heterogeneous array.

3.1 Calibration Coefficients

This chapter contains the methods which are used to calculate the relations between the measured interferometric visibilities and the true sky brightness described in Chapter 1. The interferometric equation (the two-dimensional simplified interferometric equation) is reproduced here as Equation 3.1 for reference,

$$V(u, v) = \int_l \int_m B(l, m) e^{-2\pi i (lu + mv)} dl dm, \quad (3.1)$$

where V are the visibilities measured by the interferometer, B is the sky brightness, u, v are the coordinate system for the interferometer and l, m are the corresponding sky coordinates.

The correlator outputs visibilities which contain errors representing effects within the interferometric system. These can range from atmospheric effects distorting the wavefront (which affect the phases at different antennas), differences in system electronics over the observing time, frequency dependent antenna gains etc., all affecting the polarised electromagnetic wavefront. To a good approximation, interferometers behave as linear systems (Fomalont and Perley 1999). The visibilities are sampled at discrete times and therefore the formula relating the measured visibilities V to the sky brightness B for antennas i and j (baseline $i j$) can be written linearly as

$$V_{i j}(t) = \int_l \int_m G_{i j}(t) B_\nu(l, m) e^{-2\pi i (u_{i j}(t)l + v_{i j}(t)m)} dl dm. \quad (3.2)$$

Equation 3.2 can be abbreviated to

$$V_{i j}(t) = G_{i j}(t) B_{i j}(t) + \epsilon_{i j}(t) + \eta_{i j}(t), \quad (3.3)$$

where t is the time of the observation, $G_{i j}(t)$ is the baseline based complex gain, $\epsilon_{i j}(t)$ is the baseline based complex offset and $\eta_{i j}(t)$ is the stochastic complex noise. The complex offset and complex noise derive from the use of two complex correlators (Sine and Cosine correlators).

Initial calibration is conducted on an antenna basis where the baseline based complex gain $G_{i j}(t)$ is the product of two associated antenna based complex gains $g_i(t)$ and $g_j(t)$:

$$\begin{aligned} G_{i j}(t) &= g_i(t) g_j^*(t) \\ &= a_i(t) a_j(t) e^{i(\phi_i(t) - \phi_j(t))}, \end{aligned} \quad (3.4)$$

where $a_i(t)$ is the antenna based amplitude correction and $\phi_i(t)$ is the antenna based phase correction for antenna i .

For e-MERLIN, there are two effects which need calibration, the complex electronic gain coefficients for phase and amplitude calibration ($G_{i j} = g_i g_j^*$) and the bandpass gain coefficients for bandpass calibration ($P_{i j} = p_i p_j^*$). These coefficients account for the dominant errors within the e-MERLIN system and are explained in Sections 3.6, 3.7,

3.8 and 3.9.

3.2 The Closure Relations

The closure relations (sometimes referred to as *self-calibration*), define the effects of a closed circuit of antennas on the visibilities which are independent of the individual antenna based errors. More explicitly, the net change in phase and amplitude in a closed circuit of antennas for a point source is zero and unity respectively. This condition provides a powerful technique to correct phase and amplitude errors inherent in the measured visibilities from an interferometer. The property of closure phases and amplitudes was first recognised by Jennison (1958).

The term self-calibration refers to the fact that a model of a source is used to create complex gain solutions, which are then applied to the same source, i.e. calibrating on itself. In this thesis, to avoid confusion, the term ‘*self-calibration*’ will refer to the cyclic process of imaging a point source calibrator with IMAGR and using CALIB to solve for the solutions (see Section 3.9). This is instead of self-calibration being a reference to utilising the closure relationships within phase and amplitude calibration which the FRING and CALIB algorithms use to solve the complex gain variations for phase and amplitudes (see Sections 3.6 and 3.7). All three calibration procedures: phase calibration, amplitude calibration and self-calibration, exploit the closure relations.

3.2.1 The Closure Phase

The closure phase (triple product) requires a minimum of three antennas to close the circuit and determine the phase component of the complex gain solutions. To demonstrate the closure phases in a three antenna system, Figure 3.1 shows antennas i , j and k in a circuit. If a phase delay is introduced above one of the antennas e.g. j (this delay could arise from structures in the atmosphere above antenna j), this causes a phase shift in the fringe for the baseline $i - j$, an equal but opposite delay for baseline $j - k$ and no effect for baseline $k - i$ because it is independent of antenna j . The sum of the three fringe phases is independent of the phase delay above antenna j , because the net change in phase around the closed triangle will equal zero (as you end up where you started). Therefore the closure phases are independent to all antenna-induced phase errors. The more antennas present in an array the more closure triangles can be made and thus more phase information of

the source can be obtained.

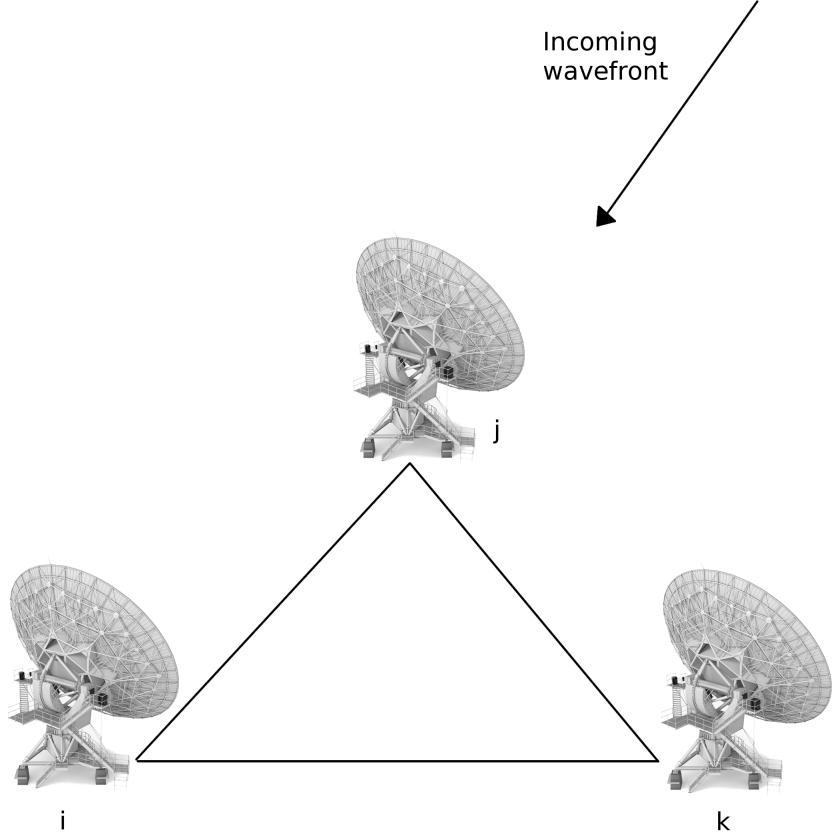


Figure 3.1: Three antennas i , j and k , depicting the closure phase circuit.

Quantitatively this is explained by defining the measured phase for baseline $i - j$ as ϕ_{ij} , equal to the true phase ϕ_{ij}^{true} plus errors $(\theta_i - \theta_j)$ from antennas i and j . The closure phase \mathcal{O} is defined as the sum of the three measured phases from the three baselines

$$\begin{aligned}
 \mathcal{O} &= \phi_{ij} + \phi_{jk} + \phi_{ki} \\
 &= \phi_{ij}^{\text{true}} + (\theta_i - \theta_j) + \phi_{jk}^{\text{true}} + (\theta_j - \theta_k) + \phi_{ki}^{\text{true}} + (\theta_k - \theta_i) \\
 &= \phi_{ij}^{\text{true}} + \phi_{jk}^{\text{true}} + \phi_{ki}^{\text{true}},
 \end{aligned} \tag{3.5}$$

where for a point source $\mathcal{O} = 0$.

3.2.2 The Closure Amplitude

The amplitude closure relationship is similar to the closure phase but is multiplicative, because it is the real part of a complex function (i.e. complex gain) rather than the imaginary exponential component. Therefore to reduce to unity the amplitude closure

requires four antennas. The amplitude closure relation for antennas i, j, k and l is

$$\frac{|r_{ij}| |r_{jk}|}{|r_{kl}| |r_{li}|} = \frac{|V_{ij}| |V_{jk}|}{|V_{kl}| |V_{li}|}, \quad (3.6)$$

where $r_{ij} = (g_i g_j^*) B_{ij}$ (Thompson et al. (1986)). The coefficients in Equation 3.6 cancel out to leave the closure amplitude equal to unity for a point source.

Exploiting these closure relation properties on a point source will produce phase and amplitude solutions equal to zero and unity respectively. Any deviations from these solutions are the associated residual errors in the calibration, which modify the complex gain coefficients in Equation 3.4.

3.3 Outline of the Calibration Pipeline

The COBRaS pipeline contains many individual modules which execute specific tasks within the calibration procedure. These include:

Editing section - Execute SERPent reduction and flagging module and AIPS QUACK tasks.

Phase calibration - Perform an initial delay correction and a full global fringe fit.

Amplitude calibration - Bootstrap refined VLA 3C286 fluxes to the point source OQ208 whilst applying corrections for the increased resolving power of e-MERLIN. Then determine the fluxes for all the other calibration sources.

Bandpass calibration - Spectral index and curvature corrections performed on a point source (OQ208, 0555+398 or phase calibrator) to calibrate the amplitude variations across the band.

Self-calibration - Iterative complex gain calibration refinement using a point source model.

A visual overview of the calibration pipeline is shown in Figure 3.2, depicting each stage of the pipeline from the raw dataset to fully calibrated visibilities, ready for imaging.

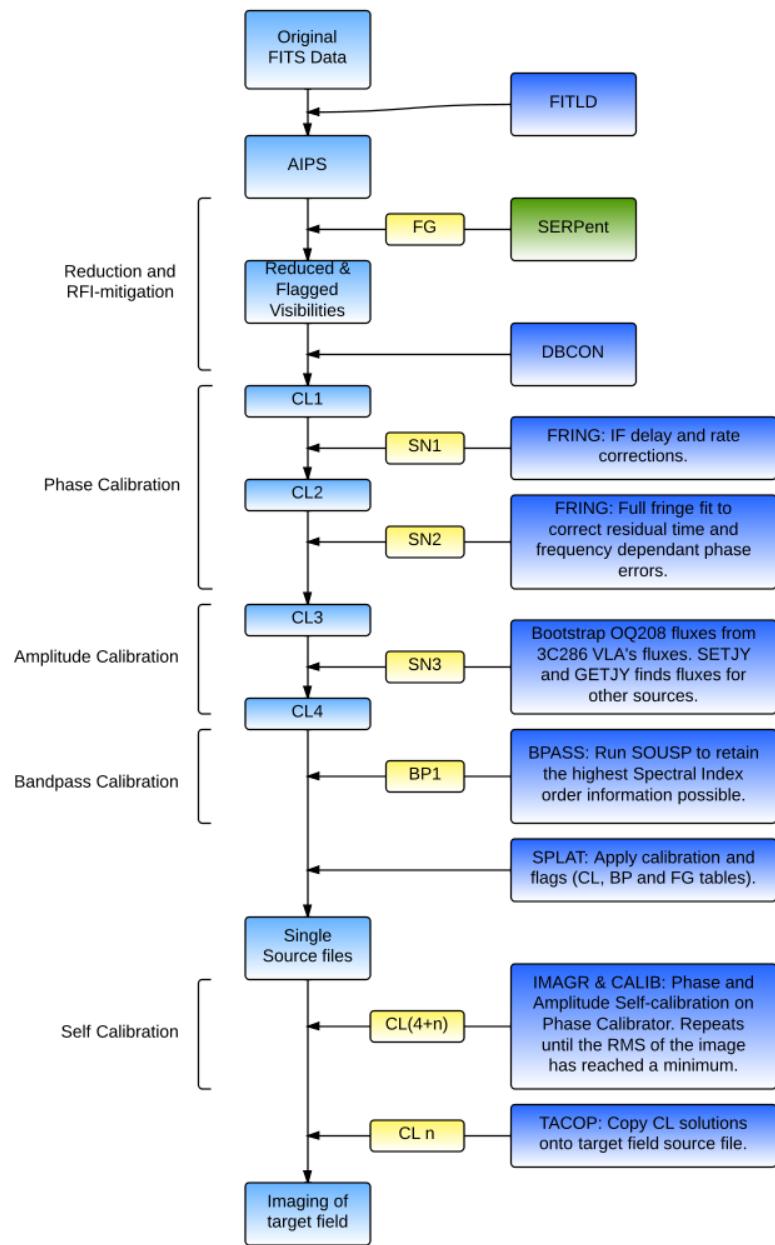


Figure 3.2: Logic diagram of the calibration procedure. The dark blue boxes; represent AIPS tasks, yellow boxes; the output extension table from the task (FG = flag table, SN = solutions table, CL = calibration table, BP = bandpass solution table), light blue boxes; the progression of the pipeline, and the green box represents SERPent.

Table 3.1: Calibration Datasets

Dataset Name	Frequency Range (GHz)	Total Integration Time (Hours)	IFs	Calibration Task
18th July 2012				
0555+398	5.49 - 6.00	2	4	Amplitude / Bandpass
1331+305	5.49 - 6.00	0.5	4	VLA Flux
1407+284	5.49 - 6.00	0.5	4	Amplitude / Bandpass
2007+4029	5.49 - 6.00	3	4	Phase / Self-cal
2032+411	5.49 - 6.00	8	4	Target field
4th May 2013				
0555+398	1.25 - 1.76	0.5	8	Amplitude / Bandpass
1331+305	1.25 - 1.76	0.5	8	VLA Flux
1407+284	1.25 - 1.76	2	8	Amplitude / Bandpass
2007+4029	1.25 - 1.76	3	8	Phase / Self-cal
Legacy L-band Pointing C	1.25 - 1.76	1.4	8	Target field
20 ^h 33 ^m 57.804 ^s + 41 [◦] 13' 12.000"				
Legacy L-band Pointing D	1.25 - 1.76	1.4	8	Target field
20 ^h 33 ^m 10.800 ^s + 41 [◦] 13' 12.000"				

COBRaS commissioning (18th July 2012) and Legacy (4th May 2013) datasets, all with full circular polarisations. The target fields are 2032+411 and Legacy L-band Pointings C and D.

3.4 Calibration Datasets

The calibration pipeline is tested against two available e-MERLIN commissioning and COBRaS Legacy datasets, the details of which are presented in Table 3.1. The pipeline performance results from the C-band commissioning dataset are included at the end of each calibration procedure, and the results from the L-band Legacy dataset are presented in Section 3.10.

3.5 Preparation of Data for Calibration

The beginning of the pipeline includes a range of user inputs for the calibration such as the solution interval for the global fringe fit, whether to average the data (or not) etc. Useful global variables are initialised here from reading the uvdata (visibility data) header to represent quantities associated with the observation.

The input raw dataset for the pipeline consists of single source files (via SPLIT/ SPLAT AIPS tasks) of each calibration source and target field and is set up within AIPS. The first procedure in the pipeline averages the data at the request of the user in time and frequency. For the initial COBRaS observations, no averaging is done in order to maintain the highest resolution possible for the initial calibration run. This comes at a cost of a large computational processing time for both SERPent (Chapter 2) and the rest of the calibration process.

Following this setup phase, the SERPent algorithm is run on each source in turn, appending an FG table to each when completed. An optional QUACK run can be made here if necessary to remove any remaining dropout visibilities at the beginning of each scan. The single source files are concatenated together via the DBCON task to create a multi-source file comprising of all the sources. Now the formal calibration procedures start.

3.6 Phase Calibration

The output visibilities from the correlator will almost certainly contain errors arising from a number of origins inherent within the interferometric system. Timing errors associated with the clocks at each antenna, errors in the position of antennas within the array, and atmospheric distortion of the radio wavefront will all cause the visibilities to contain delay residuals. If one were to integrate the visibilities to obtain the sky brightness (Equation 3.1) without ensuring that the visibility phases are coherent, it would result in a loss in signal-to-noise of the visibility amplitudes. It is therefore imperative to calculate these delay residuals over the observation.

This is conducted on a bright unresolved source in the vicinity of the target field which is called the *phase calibrator*. In Fourier space, an unresolved source or point source is constant. Therefore any corrections will be to ensure that the solutions represent a constant source in the Fourier domain. These are removed from the interferometer response, meaning any remaining phase structure represents the atmospheric distortion.

Whilst phase calibration can be conducted on resolved sources, this introduces additional errors from the source structure. Furthermore, the source will have to be modelled via self-calibration to create a map which can be used as the input model for calibration. This creates an additional step in the calibration procedure. Therefore the benefits of an unresolved phase calibrator are obvious. In the case of COBRaS, the phase calibrator is J2007+4029, a quasar classified as having a two-sided structure with a hot spot on one side (Kharb et al. 2010). For e-MERLIN resolutions, this source is slightly resolved at L-band and C-band, but not significantly and will not introduce major phase errors. However, self-calibration (Section 3.9) is part of the pipeline and should account for any phase offsets produced as a function of the resolved nature of J2007+4029.

The reason proximity to the target field is important when calibrating the phases is to

Table 3.2: Approximate Coherence Times for e-MERLIN

Frequency (GHz)	Coherence Time (minutes)	
	Good Conditions	Bad Conditions
1.6	40	1
5	40	5
22	10	30 seconds

Values for Coherence times taken from the MERLIN User Guide for the longest e-MERLIN baseline (Mark II - Cambridge).

ensure the phase solutions represent the same pathway through the atmosphere as a radio wave from the target field would experience. Hence any distortions should be mutually experienced by the target field and phase calibrator, assuming isoplanarity.

The atmosphere can be considered as having a refractive index, and therefore, turbulent structures within the atmosphere produce a variable refractive index over the path length (i.e. the speed of electromagnetic radiation varies). Hence the coherence time length of the atmosphere depends on the frequency ν of the observation. For $\nu < 1$ GHz, the ionosphere affects the wavefront, with free electrons contributing at $\nu < 300$ MHz. Furthermore, the activity of the solar cycle plays a significant role in coherence times, as the Sun affects the shape of the ionosphere. For higher frequencies, $\nu > 1$ GHz, it is the troposphere which dominates the turbulence of the wavefront.

The phase calibrator is observed periodically between target source observations, with the switching times not exceeding the atmospheric coherence time for the observed frequency. Table 3.2 gives the approximate coherence times for e-MERLIN frequencies.

To calibrate the phases, the AIPS task FRING is employed to perform an initial delay correction before also conducting a full global fringe fit. FRING uses a least squares (LS) algorithm given a source model (point source assumed in the case of J2007+4029) to estimate the antenna based components of the fringe rate and delays (Schwab and Cotton 1983). To aid the LS algorithm, a Fourier Transform method provides the starting guesses, as the LS method needs initial parameters to converge on the true global minimisation function instead of the local minima.

The phase calibration process utilising the AIPS tasks is shown as a logic diagram in Figure 3.3 as a visual aid.

Within the pipeline, the phase calibrator is defined as a required user input (although searches for variations are made) and a list of calibrators is also searched for and compiled.

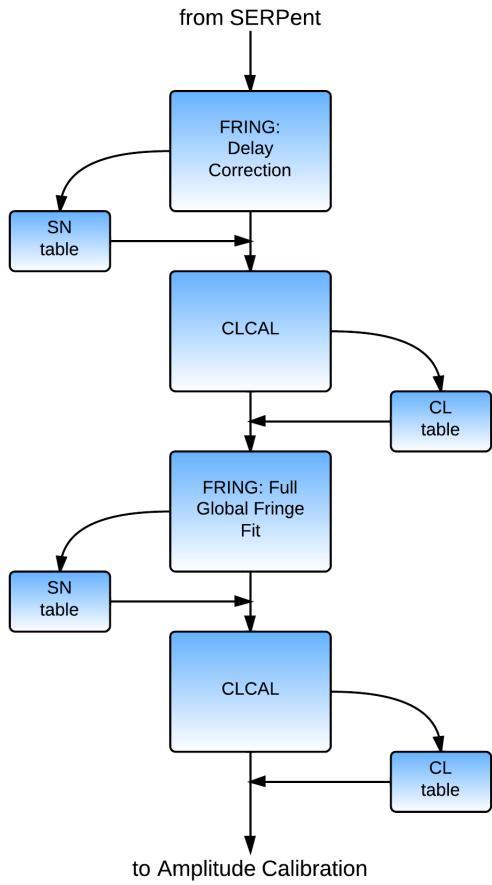


Figure 3.3: Logic diagram of the phase calibration with AIPS tasks.

These calibrators are also standard for COBRaS observations. During the FRING task, the phase offsets are referred to a reference antenna, where typically an antenna at the centre of the array is most suitable. The pipeline automatically searches and chooses either the Mark II antenna or Pickmere (if Mark II is not available).

The pipeline then searches through the visibilities to find a search window, e.g. ~ 2 minutes in length, where a fraction of the total number of baselines are available. Smaller time windows and fewer baselines are successively chosen if no window is found. Then the first instance of the AIPS task FRING is used to search for the single band delays only, in a delay window of 450 nanoseconds. This creates one solution for each IF, which is added to an output SN table.

The geometric delay is frequency independent and therefore calculating the single band delays for each individual IF *should* not make any difference than creating one solution for all IFs. However, these delays are used in the full fringe fit with atmospheric effects having a frequency dependence on the delays and phases. Moreover, there is no disadvantage for

not having a solution for each IF. Therefore the pipeline creates separate delays solutions for each IF and baseline.

The AIPS task CLCAL, takes an input SN solution table and creates the corresponding CL calibration table, which is used to apply the corrections to the visibility data. CLCAL can be used at any point in the calibration procedure where an SN table is created. As the calibration proceeds, new versions of CL tables are created, containing the previous and new calibrations to the visibilities. The newest table is always used in the next step of the calibration (see Figure 3.2).

Following the single-band delay correction a full global fringe fit is performed as described by Schwab and Cotton (1983). As with the single band delays, a reference antenna is required along with a designated solution interval (AIPS adverb SOLINT). The solution interval needs to be large enough to produce an adequate signal-to-noise ratio to find solutions, as well as short enough to maintain coherence and is a matter of balance. The user can once again define their own solution interval at the beginning of the pipeline script, however, the defaults provide the worst case scenario (bad conditions) limits given in Table 3.2. For L-band this is 1 minute, and for C-band this is 5 minutes. The fringe fit is performed for the entire length of the observation.

Another CLCAL run is made to convert the SN solutions output from the global fringe fit to CL tables. Figure 3.4 shows a before (a) and after (b) plots of phases as a function of frequency for OQ208, plotted with the AIPS task POSSM from the C-band commissioning dataset. This baseline (Lovell-Darnhall) suffers from *phase wrapping* due to the long baseline length and the observing configuration. After phase calibration on J2007+4029, solutions are converted to calibration weights for all sources.

Generally the phases for each antenna are good, flat and stable. Another demonstration of the phase calibration on longer baselines is shown in Figure 3.5 where the before (a) and after (b) phases as a function of frequency are plotted. Phase wrapping (phases looping through $\phi = -180^\circ$ to 180° etc.) is present before calibration along with other issues. These are corrected after the fringe fitting demonstrating good interpolation of the phase solutions.

The phase calibration corrections are tied up in the complex gain coefficients ($g_i(t)$ $g_j^*(t)$; Equation 3.3). Because of the complex nature of the gain coefficient (Equation 3.4), it will contain corrections from both phase ($e^{i(\phi_i(t) - \phi_j(t))}$) and amplitude ($a_i(t)$ $a_j(t)$) calibration procedures.

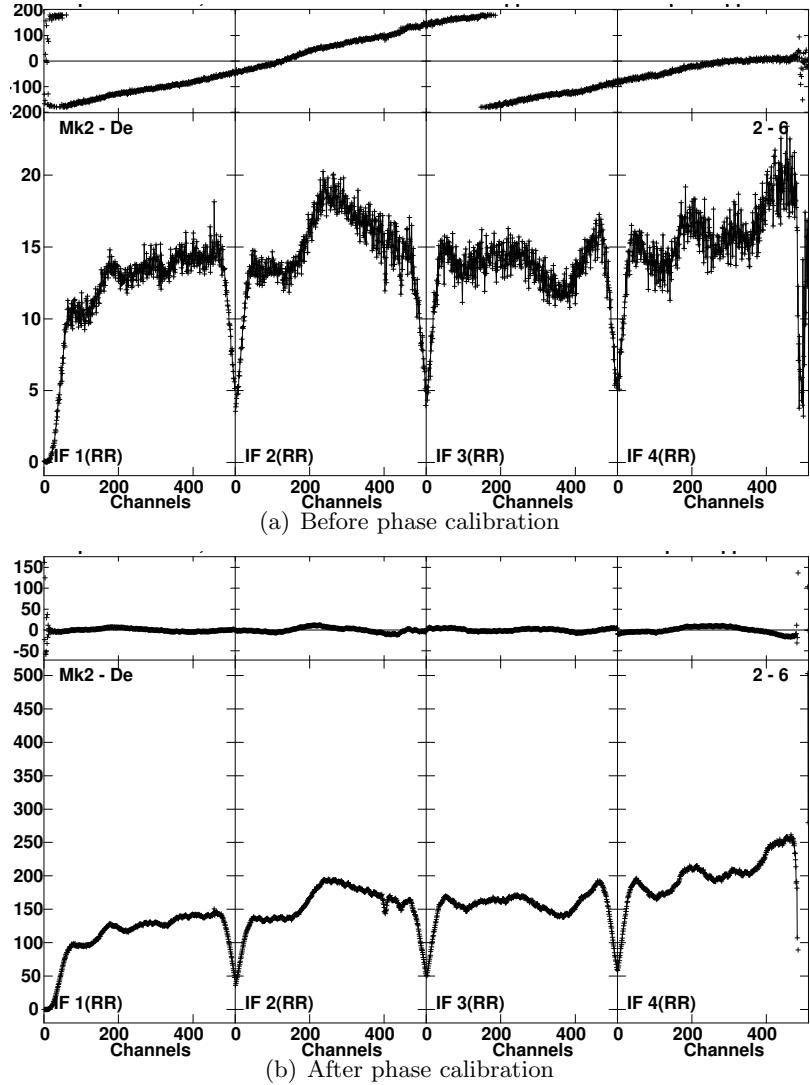


Figure 3.4: Phase calibration J2007+4029, C-band e-MERLIN commissioning data on the Mk II-Defford baseline. Top of each plot are the phases and the bottom are the amplitudes. (a) shows the phases as a function of frequency before calibration, and (b) after delay and phase calibration. Phase calibration has smoothed out the phases as a function of frequency.

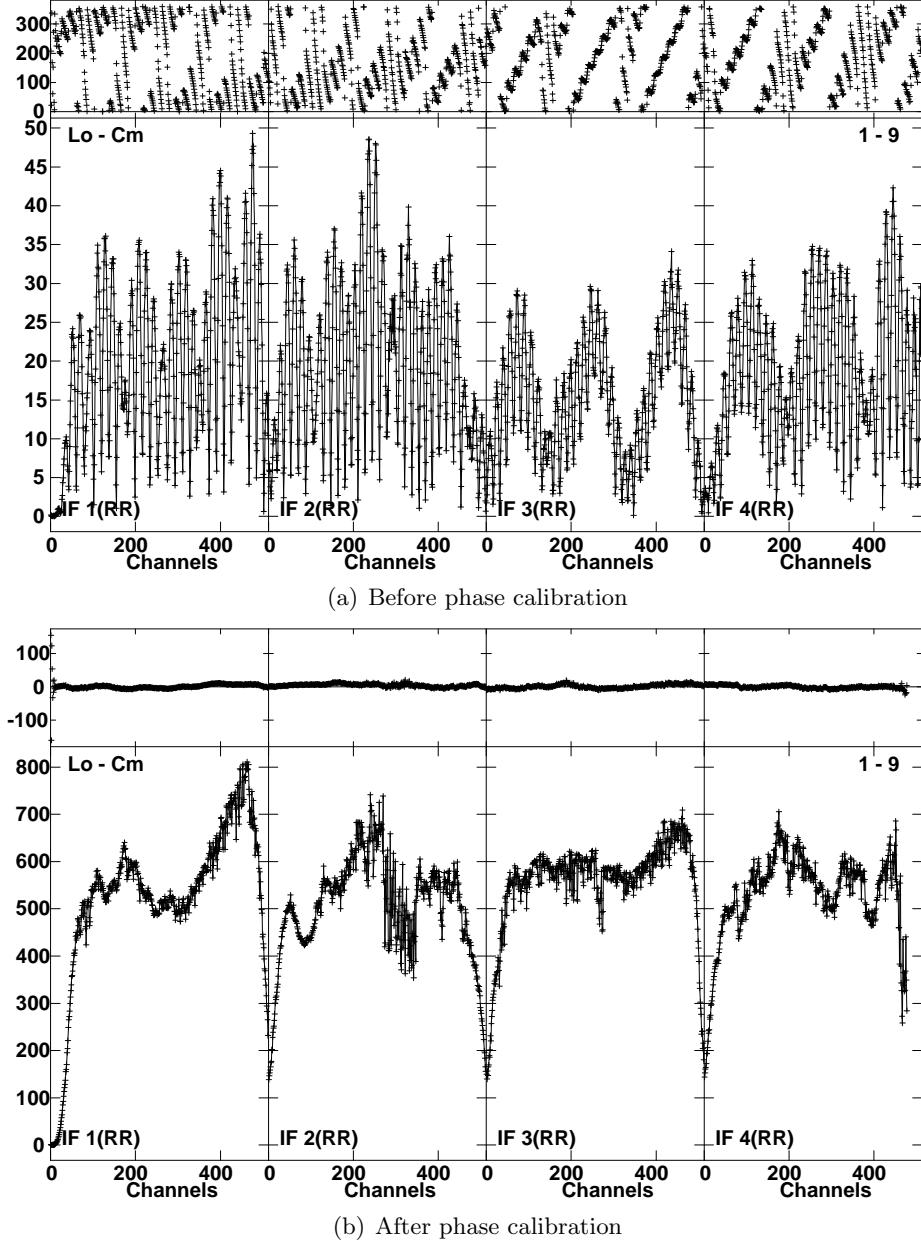


Figure 3.5: Phase calibration J2007+4029, C-band e-MERLIN commissioning data on the Lovell-Cambridge baseline. Top of each plot are the phases and the bottom are the amplitudes. (a) shows the phases as a function of frequency before calibration, and (b) after delay and phase calibration. Phase calibration has smoothed out the phases as a function of frequency.

3.7 Amplitude Calibration

After phase calibration, the pipeline performs amplitude calibration on all of the calibration sources. The aim of amplitude calibration is to set the absolute flux level in Janskys (Jy) for each source. The amplitudes of each source outputted from the correlator is only an arbitrary count level which has no physical meaning. To scale these correlator counts to a common flux scale, a source whose flux is well-known needs to be observed.

The method utilised for e-MERLIN involves the flux calibrator source 3C286, which has been observed extensively over the last few decades by the VLA following an initial study of the flux density scale by Baars et al. (1977). 3C286 is a compact steep-spectrum quasar at redshift $z = 0.846$, with a small resolved component $\sim 2''.5$ to the South-West (Perley and Butler 2013). Despite being partially resolved, 3C286 has a flux density which has been stable over a time span of decades, unlike many other VLA flux calibrators. Therefore the flux of 3C286 can be used with confidence as a stable flux calibrator to bootstrap its known flux to a less resolved calibrator for amplitude calibration. The benefits of using a more unresolved source for amplitude calibration are the same as for phase calibration (in Section 3.6), namely a point source is better modelled when conducting Fourier analysis and LS methods, leading to more accurate results. This assumes the point source is stable and the data is of sufficient quality to perform the extrapolation.

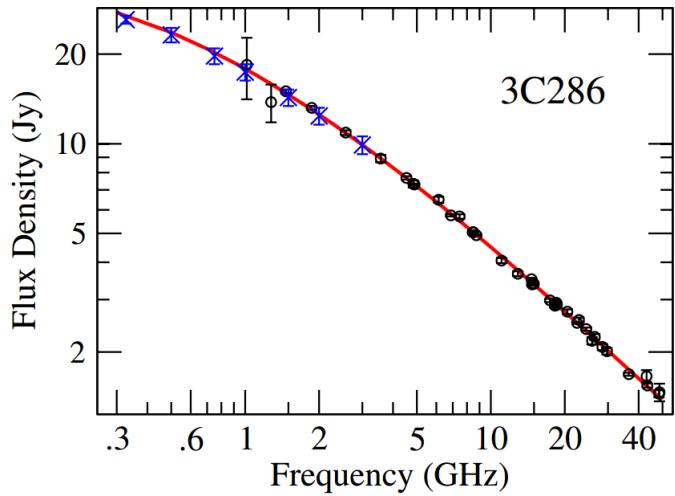


Figure 3.6: 3C286 radio spectrum from 300 MHz to 50 GHz, figure taken from Perley and Butler (2013).

Perley and Butler (2013) present a recent study on the common VLA flux calibrators including 3C286 with the upgraded JVLA. The authors observe the targets at all available

JVLA frequencies, from 300 MHz to 50 GHz (Figure 3.6) to produce accurate spectral flux densities, which are conveniently described as polynominal expressions. The spectral flux density S in Jy is given by

$$\log(S) = a_0 + a_1 \log(\nu_G) + a_2 [\log(\nu_G)]^2 + a_3 [\log(\nu_G)]^3, \quad (3.7)$$

where ν_G is the frequency in GHz. The coefficients a_n for 3C286 are given in Table 3.3.

Table 3.3: Coefficients for the Spectral Flux Density of 3C286

Source	a_0	a_1	a_2	a_3
3C286	1.2515 ± 0.0048	-0.4605 ± 0.0163	-0.1715 ± 0.0208	0.0336 ± 0.0082

Coefficients taken from Perley and Butler (2013), table 10.

This polynominal expression cannot be utilised directly for e-MERLIN observations because of the higher resolving power of e-MERLIN relative to the JVLA. In addition to Equation 3.7, the baseline length B and frequency ν are a function of the resolving power $\theta_R(\nu, B)$, and the reduction in central flux density $S(\theta)$ of a Gaussian PSF representing 3C286 (θ_S) is given by (Danielle Fenech, Peter Thomasson, Anita Richards; private communication):

$$\begin{aligned} \frac{S'(\theta)}{S(\theta)} &= \frac{\frac{1}{2\pi(\theta_R(\nu, B)^2 + \theta_S^2)}}{\frac{1}{2\pi\theta_R(\nu, B)^2}} \\ &= \frac{1}{1 + \frac{\theta_S^2}{\theta_R(\nu, B)^2}}. \end{aligned} \quad (3.8)$$

By designating $\rho = [\theta_S/\theta_R(\nu, B)]^2$, ρ has been calculated for a specific frequency (5 GHz) and baseline (Mark II - Pickmere) to produce a reference ratio to scale the resolving power to any frequency. Using $\theta_R(\nu, B) = k/\nu B$ gives

$$\rho(\nu, B) = \left(\frac{\nu B}{\nu_{\text{ref}} B_{\text{ref}}} \right)^2 \rho_{\text{ref}}, \quad (3.9)$$

where ρ_{ref} is fixed at 0.04 for reference frequency $\nu_{\text{ref}} = 5.00$ GHz and reference Mark II

- Pickmere baseline length $B_{\text{ref}} = 11236.79$ metres. This gives the final e-MERLIN flux $S_{\text{e-MERLIN}}$:

$$S_{\text{e-MERLIN}} = \frac{S_{\text{VLA}}}{1 + \rho(\nu, B)} \quad (3.10)$$

where S_{VLA} is the JVLA flux given by Equation 3.7 with the coefficients from Table 3.3. This e-MERLIN flux can then be applied to the 3C286 observation and used for flux and amplitude calibration, or scaled to an unresolved point source calibrator such as OQ208 for a theoretically optimal flux and amplitude calibration. This is dependent on data quality for each source, and therefore the pipeline has two options: to use the resolved 3C286 to determine the fluxes of the other calibrator sources, or the OQ208 fluxes scaled from 3C286 to determine the fluxes of the other calibrator sources.

The pipeline starts the amplitude calibration with a search for the primary flux calibrator 3C286 and the point source calibrator OQ208. The central frequencies of each IF are calculated from the uvdata headers, which will be used in the calibration of the VLA 3C286 fluxes. With the wide observing bandwidths of the new e-MERLIN receivers, there is a significant change in flux over both L-band and C-band observations. The L-band Legacy datasets have 512 MHz bandwidth with 8 IFs (64 MHz per IF) and COBRA-S C-band Legacy datasets will have 2 GHz bandwidth with 16 IFs (128 MHz per IF). Solving the flux for the centre of each IF should be sufficient for accurate amplitude calibration.

The shortest available baseline is then computed from two of three antennas; Mark II, Pickmere and Darnhall. Table 3.4 shows the three shortest baselines and the resolving percentages of 3C286 at L-band and C-band frequencies.

Table 3.4: Shortest e-MERLIN Baselines

Baseline	Baseline Length (in metres)	Resolving Percentage (%)	
		1.6 GHz	5 GHz
Mark II - Pickmere	11236.79	0.41	3.85
Darnhall - Pickmere	15923.55	0.82	7.44
Mark II - Darnhall	17737.45	1.01	9.06

The three shortest e-MERLIN baselines with the resolving percentages of 3C286 at 1.6 GHz (L-band) and 5 GHz (C-band).

The e-MERLIN resolved fluxes of 3C286 at the central frequencies of each IF are now

calculated using Equations 3.8, 3.9 and 3.10. If no scaling is to be done to a more point-like source, these 3C286 fluxes are directly used in SETJY.

However if the data quality is sufficient (using direct output from the correlator, i.e. correlator counts) then scaling these 3C286 fluxes to an unresolved source may yield more accurate calibration. To scale this against the correlator counts, the middle half of the time-scan of the 3C286 observation (i.e. centred in the middle of the timescan with a quarter of the scan either side of this centre included) is taken to reduce the chances of early slew errors and increase the probability that all the antennas have settled on the source. The middle quarter of each IF is considered to mitigate the effects of the bandpass edge response (which have not been corrected yet). This should average to the central frequency of the IF which corresponds to the central frequency flux for each IF as calculated by Equation 3.10. The correlator counts for parallel polarisations are considered only, because the cross-hand polarisations are generally noiser and more problematic. This is valid for COBRaS, because the science goals of COBRaS only require Stokes I (parallel polarisations).

The correlator counts from this slice (in time and frequency) of the observations is averaged to a singular value which corresponds directly to the calculations from Equation 3.10 for 3C286. The same process is made for OQ208 to obtain the average correlator counts for the centre of each IF. A direct scaling between the correlator counts of 3C286 and OQ208 is multiplied by the resolved e-MERLIN fluxes for each IF, to result in true e-MERLIN fluxes for the centre of each IF for OQ208.

The rest of the amplitude calibration is conducted with AIPS tasks in the standard manner. A logic diagram is presented in Figure 3.7 to depict the rest of the amplitude calibration process.

The first step involves setting the Stokes I flux in the source file (SU table) for OQ208 (or 3C286 depending on which method is used) with SETJY. The other Stokes parameters are ignored. The pipeline contains two possibilities at this point: run one instance of CALIB on all of the baselines over the full uv-range, or include a restricted CALIB run on the shortest baselines on 3C286 before running CALIB again on all calibrator sources over all baselines. This double run is adopted by the e-MERLIN pipeline and has shown to provide better calibration solutions (Rob Beswick, Megan Argo; private communication). The idea is that restricting to shorter baselines reduces the possibility of resolving 3C286 and therefore calculating more accurate solutions initially for the shorter baselines before

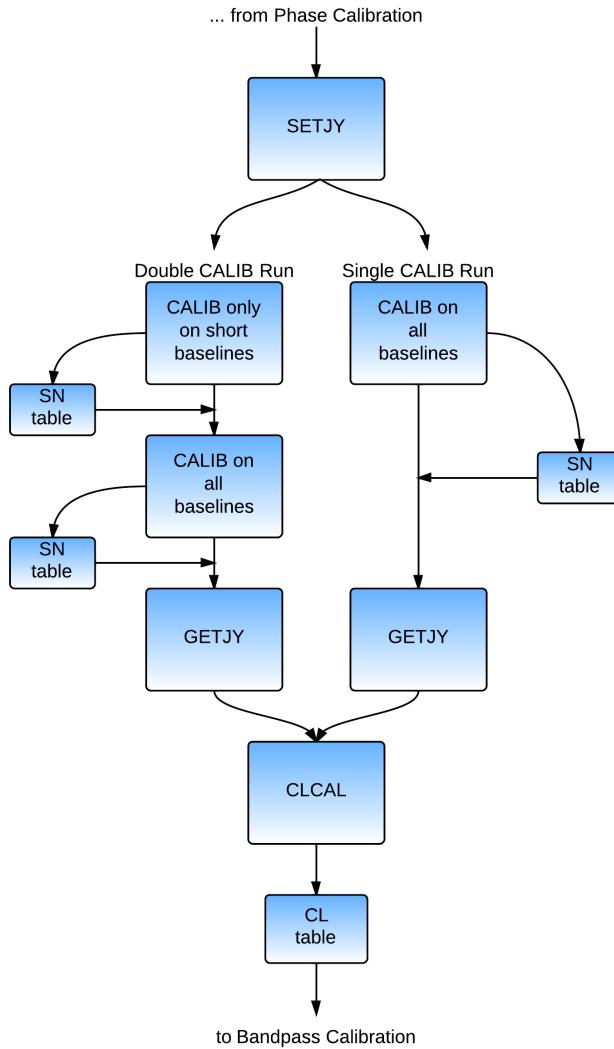


Figure 3.7: Logic diagram of the amplitude calibration procedure.

including the longer baselines into the solutions. AIPS does inherently contain input models for 3C286, but these can not be used in the calibration because the models are for the JVLA and do not consider the higher resolving power of e-MERLIN. There are however future plans for equivalent models for e-MERLIN.

Once the solutions have been found and written to the SN table, the GETJY task is used to determine the fluxes for individual sources and written into the source table (SU table). Then a run of CLCAL is made to convert these solutions to the calibration CL table to include the amplitude component of the electronic gain coefficients ($g_i g_j^*$). These will be further refined during self-calibration (Section 3.9) but will now be applied to the

interferometric equation:

$$V(u, v) = (g_i g_j^*) \int_l \int_m B(l, m) e^{-2\pi i (lu + mv)} dl dm. \quad (3.11)$$

Figure 3.8 shows the before (a) and after (b) amplitude calibration plots for the point source OQ208 from the C-band commissioning dataset. The before scaling is in correlator counts and the after scaling is in Jy.

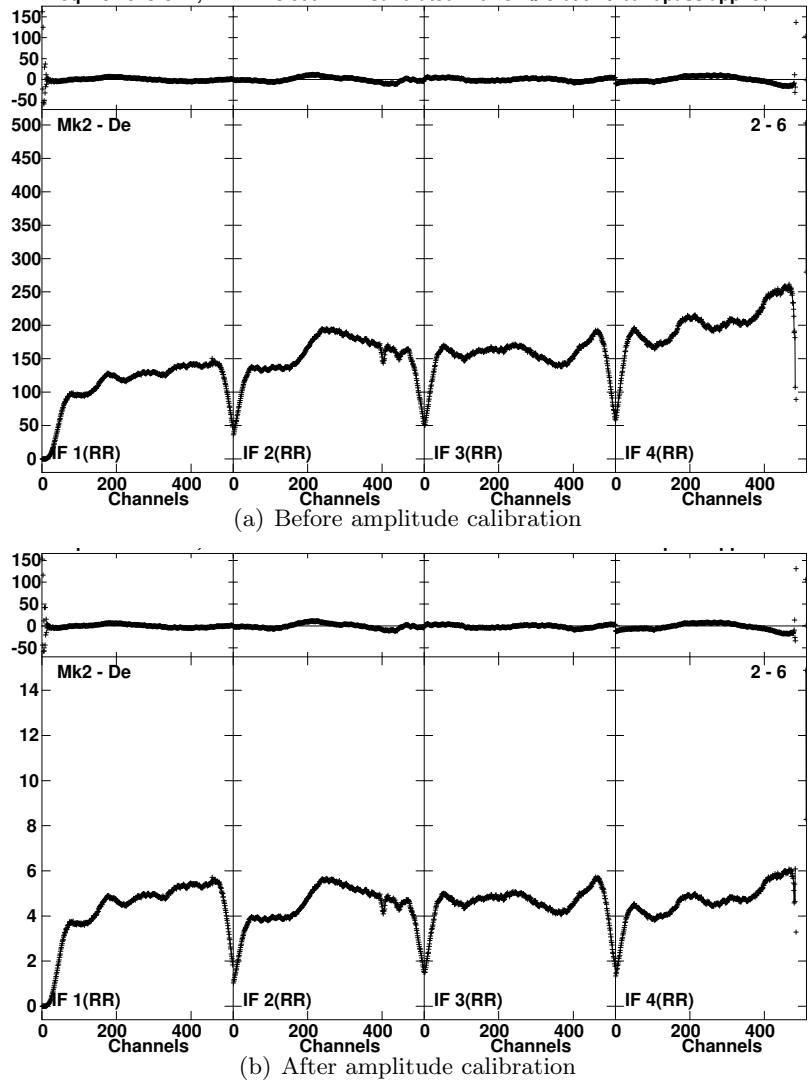


Figure 3.8: Amplitude calibration showing before (a) and after (b) plots of J2007+4029 from the C-band commissioning dataset, MkII-Defford baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency. Note the amplitude scales differ in each plot due to the calibration; in the before figure the scale is in correlator counts, and in the after figure the scale is in Jy.

3.8 Bandpass Calibration

After amplitude calibration, the response of the receiver across the band needs to be calibrated to correct the antenna based complex amplitude and phase gain variations as a function of frequency. Bandpass calibration is important to ensure the accurate measurement of spectral features e.g. spectral index and curvature of sources. In the case of COBRaS, spectral indices are required to detect and classify massive binary interaction regions, which are non-thermal and therefore have a negative spectral index (for details see Section 5.3.3).

An ideal bandpass would have a flat response at the maximum transmission available across the whole bandwidth. In reality, the response is not flat across the bandwidth and the edge frequencies suffer from effects attributed to the bandpass filters. To demonstrate the response of the bandpass of e-MERLIN, Figure 3.9 shows the L-band bandpass of 3C286 during the L-band Legacy observations.

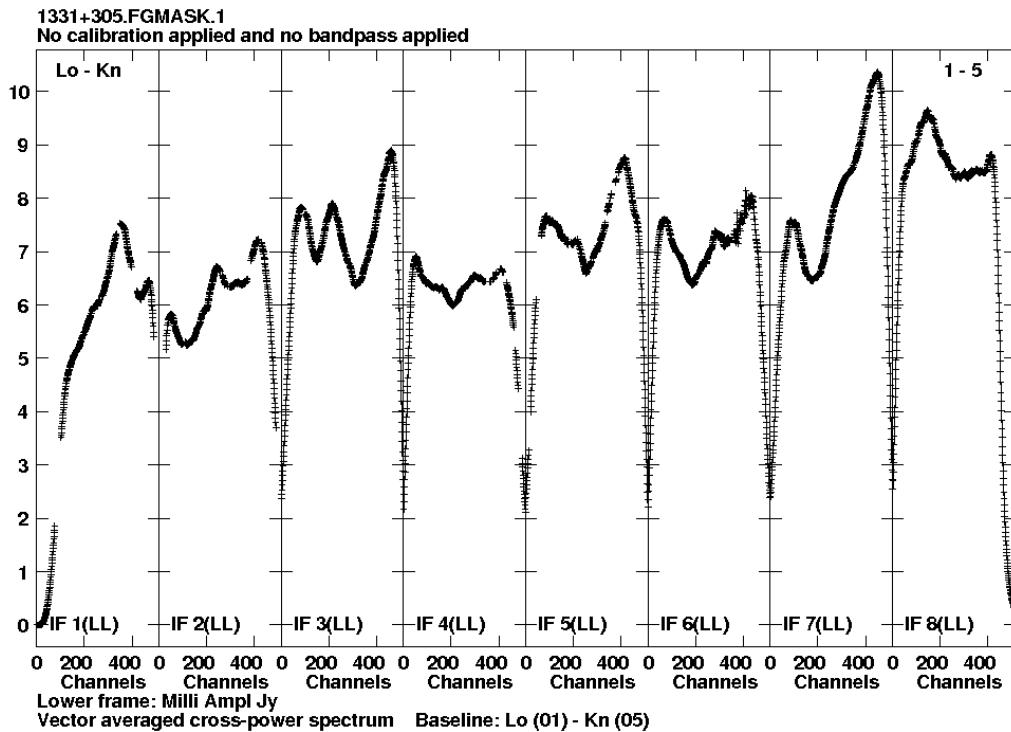


Figure 3.9: Example of the uncalibrated e-MERLIN L-band bandpass response, with channels on the x axis and correlator counts on the y axis. The 512 MHz bandwidth is split into 8 IFs, each 64 MHz in width, with an individual channel width of 125 kHz.

In addition to the effects of the baseband filters, the spectral index and curvature need accounting for during the bandpass calibration. This becomes important for new wide-

band receivers where the flux of a source will vary over the wide bandwidth, i.e. spectral index properties are observed.

Bandpass calibration uses the same methodology as amplitude and phase calibration, where a specified bandpass calibrator is used to determine the solutions. This need not be an unresolved source, as the AIPS bandpass calibration task BPASS has the option for a source model to be included and the visibility data can be divided by the source model prior to calibration. Although there is a frequency dependence with a resolved source which complicates calibration.

However, as stated before, to ease the automation a point-like source is used and assumed in the calibration procedure, because of the availability of such a source in the COBRaS observations. The ideal bandpass calibrator is bright and has a flat spectrum across the band. The bandpass calibrator for COBRaS is OQ208 (with 0555+398 as a backup/ additional calibrator). The phase calibrator J2007+4029 is also bright and point-like and therefore can be used for spectral fitting and bandpass calibration.

Figure 3.10 shows the logic diagram of the bandpass procedure, with the AIPS task BPASS and a spectral fit used to rectify the bandpass responses. The SOUSP task can fit a polynomial to the fluxes to determine the coefficients which describe the spectral index and spectral curvature

$$\begin{aligned} \log(f) = & \log(f_0) + S \log(f) + C(1) [\log(f)]^2 \\ & + C(2) [\log(f)]^3 + C(3) [\log(f)]^4, \end{aligned} \quad (3.12)$$

where f is the flux, f_0 is the flux at 1 GHz, S is the spectral index, C is the spectral curvature and all logs are base 10. It is unlikely that SOUSP will calculate reliable values for $C(3)$, with SETJY only utilising coefficients up to $C(2)$ for flux calibration.

However, SOUSP considers the fluxes from all IFs equally, and makes no judgment on whether individual fluxes are reasonable whilst fitting a polynomial. One bad flux from poor quality data or sub-optimal calibration can skew the spectral fit dramatically. Because of this flaw in the SOUSP task, a spectral fit module is created to remove any spurious fluxes and then fit a polynomial to the remaining fluxes to produce the spectral coefficients required for Equation 3.12 and BPASS. Figure 3.11 shows the before and after spectral fit and demonstrates the removal of outliers from the main trend. This is achieved by using a running median, where each flux is compared to the closest two fluxes in

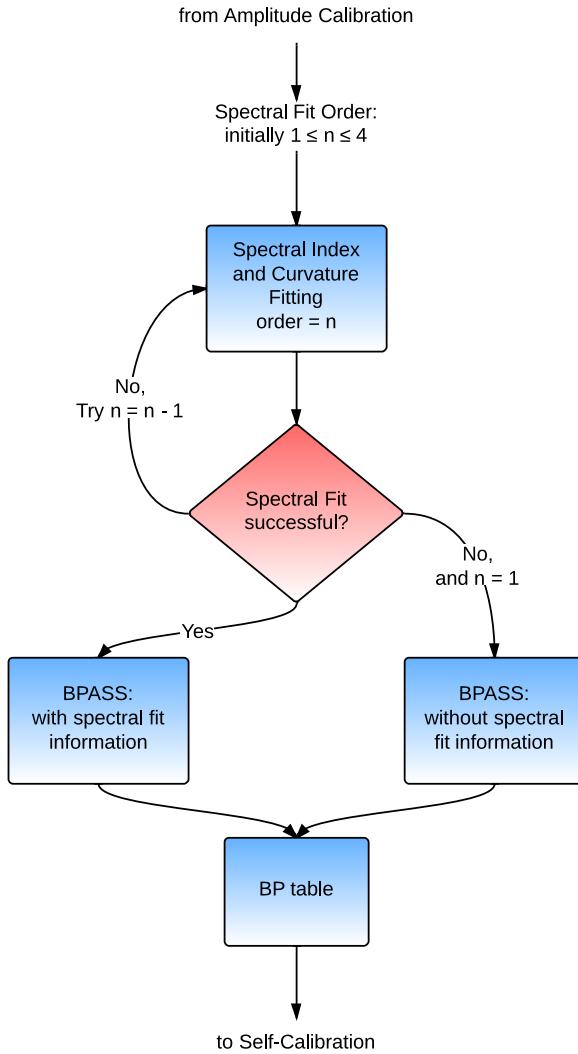


Figure 3.10: Logic diagram of the bandpass calibration procedure.

frequency and if the flux is beyond a threshold it is removed from the fit. The module raises a warning if half of the fluxes are removed.

The spectral fit module determines the spectral coefficients, which are expected to be linear for OQ208 for L-band and possibly quadratic for C-band as the spectrum turns over at ~ 6 GHz as seen in Figure 3.12. The outputs are read into BPASS as the adverbs `SPECINDEX` and `SPECURVE`, representing the spectral index and curvature respectively.

There are two methods to calibrate the bandpass with BPASS: (1); use a least squares method to calculate the cross-correlated spectra into antenna based complex functions, (2); use the auto-correlations to calculate the real part (amplitude) of the bandpass response by setting the phase part of the bandpass response to zero. Since the auto-correlations have been flagged (contain uncorrelated RFI) and the phases have already been calibrated,

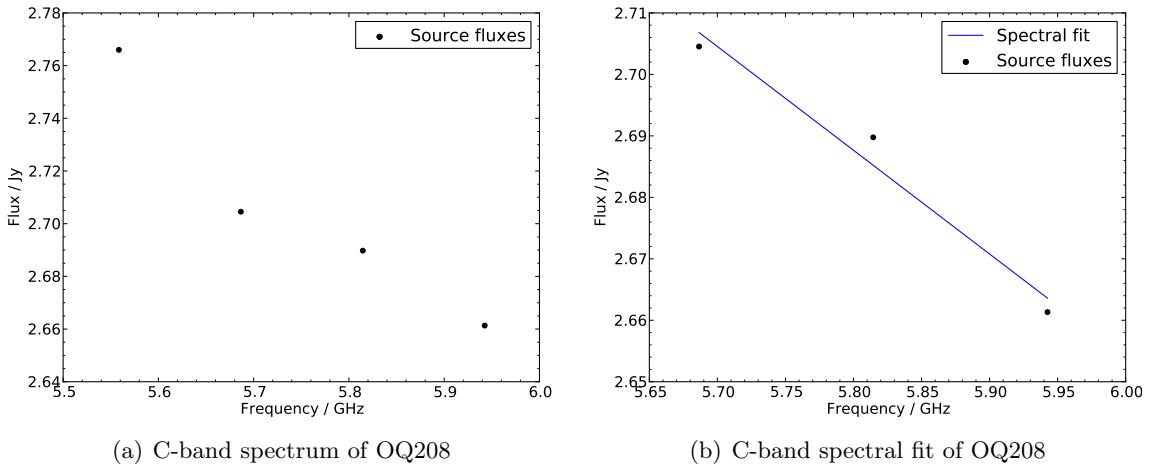


Figure 3.11: C-band spectral fit of OQ208. (a) is the spectrum of OQ208 from amplitude calibration and (b) is the spectral fit of OQ208. Note that the flux at 5.56 GHz (corresponding to IF 1) has been removed from the fit.

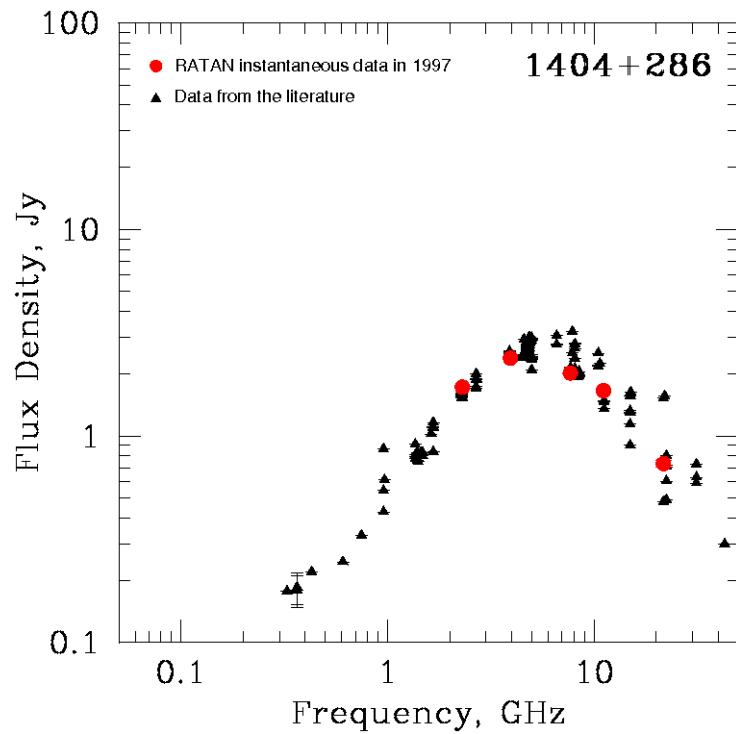


Figure 3.12: Spectrum of OQ208 taken from the MOJAVE survey (Lister et al. 2009).

decomposing the baseline based functions to antenna based functions via method (1) is adopted in the pipeline.

The bandpass calibrator (OQ208) is divided by a point source model and the scans averaged in time. Hanning smoothing is applied before the calibration and the centre 75% of the band is averaged to the so called ‘channel 0’ because channel-dependent flagging has occurred with SERPent (this is the recommended action in the AIPS cookbook).

Bandpass calibration with spectral correction with the spectral fit module on OQ208 (C-band) can be seen in Figure 3.13. The first plot (a) shows before bandpass calibration and the second plot (b) shows after, with only the central 75% of the bandwidth shown. Before calibration, the amplitudes varied between 4 - 6 Jy over the inner 75% of the bandwidth, with particularly strong variation in IF 1. Following bandpass calibration, the amplitudes are flatter across each IF with variations limited to 0.20 Jy for each IF. A slight increase in amplitude present at higher frequencies is probably a result of the spectral index of the source.

BPASS writes the solutions to a bandpass (BP) table, which include the amplitude and phase bandpass coefficients as a function of frequency ($p_i p_j^*$). These modify the interferometric equation as follows

$$V(u, v) = (p_i p_j^*) (g_i g_j^*) \int_l \int_m B(l, m) e^{-2\pi i (lu + mv)} dl dm. \quad (3.13)$$

3.9 Self-Calibration

After phase, amplitude and bandpass calibration, residual errors will usually remain within the visibility data with refinements to the complex gain coefficients possible through self-calibration. These errors are usually a result of using calibrators to correct for amplitude and phase errors which were observed at different times and/ or positions in the sky. This can lead to subtle differences in the electronics over the time span, and/ or differences in the atmosphere etc. which can leave errors in the calibration gains, i.e. self-calibration modifies the complex gain coefficient (Equation 3.11).

As discussed in Section 3.2, self-calibration is the process of imaging a point source calibrator and using the CLEAN components (CC table) from the image to help determine solutions for the complex gain coefficient via the CALIB algorithm. Self-calibration does not require an unresolved source (although the advantages have been discussed previously),

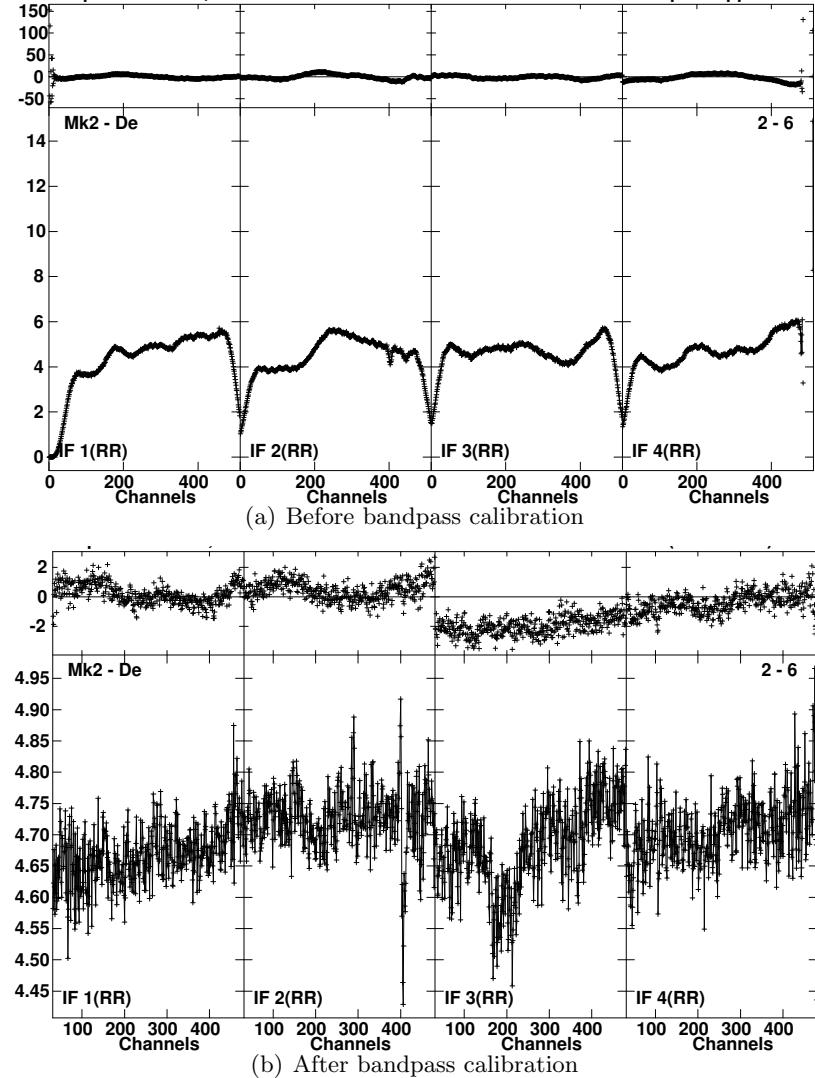


Figure 3.13: Bandpass calibration showing before (a) and after (b) plots of J2007+4029 from the C-band commissioning dataset, MkII-Defford baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency in Jy. The left image has a different x axis scale showing all the channels in each IF, whereas the right image is limited to the central 75% channels. After calibration the amplitudes across the bandwidth are flatter.

because the visibilities are normalised to the input model before the complex gain solutions are determined. The source model must be bright to produce a high signal-to-noise ratio to determine accurate solutions.

For the COBRaS pipeline the input source for self-calibration is the phase calibrator source J2007+4029. This source is bright (~ 2 Jy at L-band and ~ 4.5 Jy at C-band), and only slightly resolved at all e-MERLIN baselines and lies close to the target fields of the COBRaS project, therefore giving the most accurate phase solutions possible from self-calibration. Furthermore, COBRaS is a semi-blind survey, with no known other bright point-like sources in every L-band and C-band field available for more accurate phase solutions from self-calibration.

The procedure for self-calibration is shown in Figure 3.14 and is as follows. First an image of the source model is made with the AIPS task IMAGR, applying the most recent complex gain and bandpass solutions (highest CL and BP tables) from the calibration. The initial image must not be convolved with too many iterations to avoid introducing errors early on in the self-calibration. This model is then used in CALIB in self-calibration mode to calculate the phase gain solutions only. The output SN table solutions are then converted to CL table calibrations and then the self-calibration source is imaged again using the updated calibration solutions in the new CL table. The process is repeated until the root-mean-square (RMS) noise levels-off. Once the residual phase errors have been corrected, the same self-calibration process is conducted again, but making CALIB solve for both amplitude and phase solutions (amplitude self-calibration).

The pipeline contains two methods of self-calibration. The first, images the phase calibrator and creates solutions with all IFs simultaneously, while the second, images and solves IFs individually. Both methods have advantages and disadvantages which are considered before selecting one or the other. Combining the IFs in self-calibration will increase the signal-to-noise ratio and therefore increase the likelihood of finding good solutions in CALIB. The downside however, is this will introduce radial features or infidelities in the image as a result of frequency errors from a single solution applied to all channels in all IFs. This method can also introduce increasing errors away from the phase centre because of the multi-frequency beams.

Applying self-calibration on individual IFs has the advantage of reducing frequency-based errors by determining solutions for each IF across the bandwidth. It also retains spectral index information (a flux solution for each IF) as CALIB does not deal with the

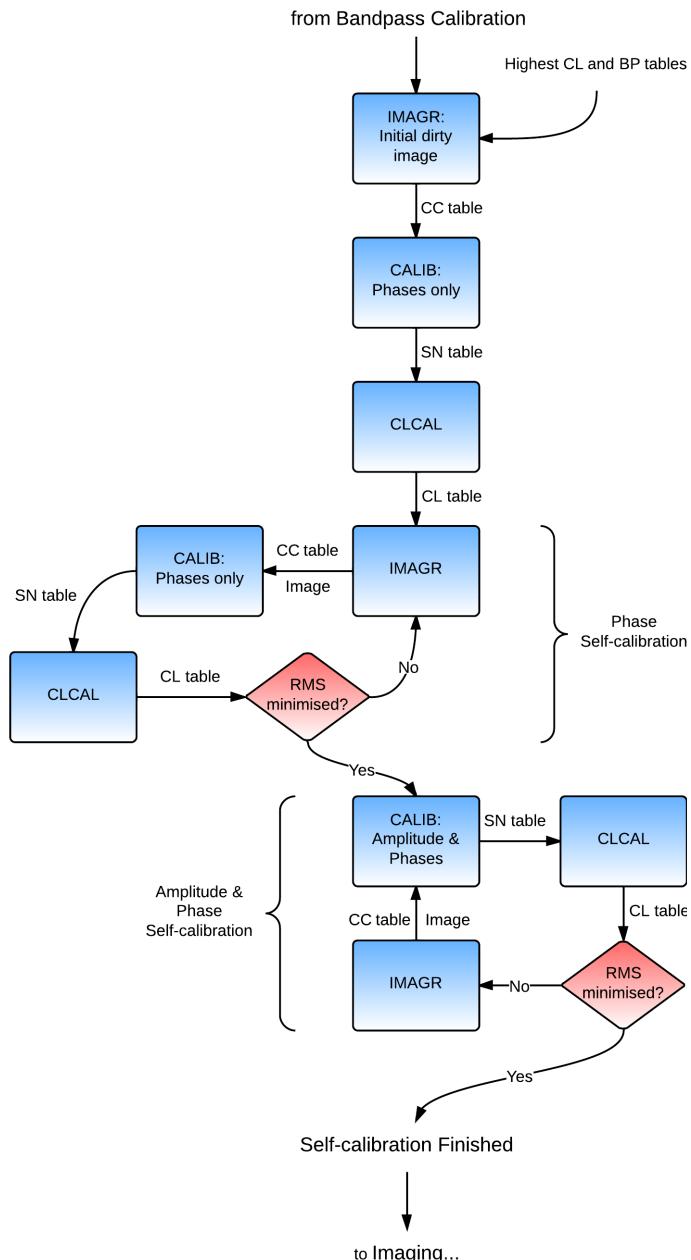


Figure 3.14: Logic diagram of the self-calibration procedure.

spectral index¹. The disadvantage of this method is the lower signal-to-noise ratio of the data to find solutions. However, due to the brightness of the phase calibrator (J2007+4029) used for self-calibration, this resolves the issue of signal-to-noise, allowing self-calibration to be performed successfully on individual IFs.

Figure 3.15 shows the before (a) and after (b) corrections from self-calibration on J2007+4029. Subtle corrections have been made to the data in this example such as the small delay offset between IFs 2 and 3 along with the amplitude levels.

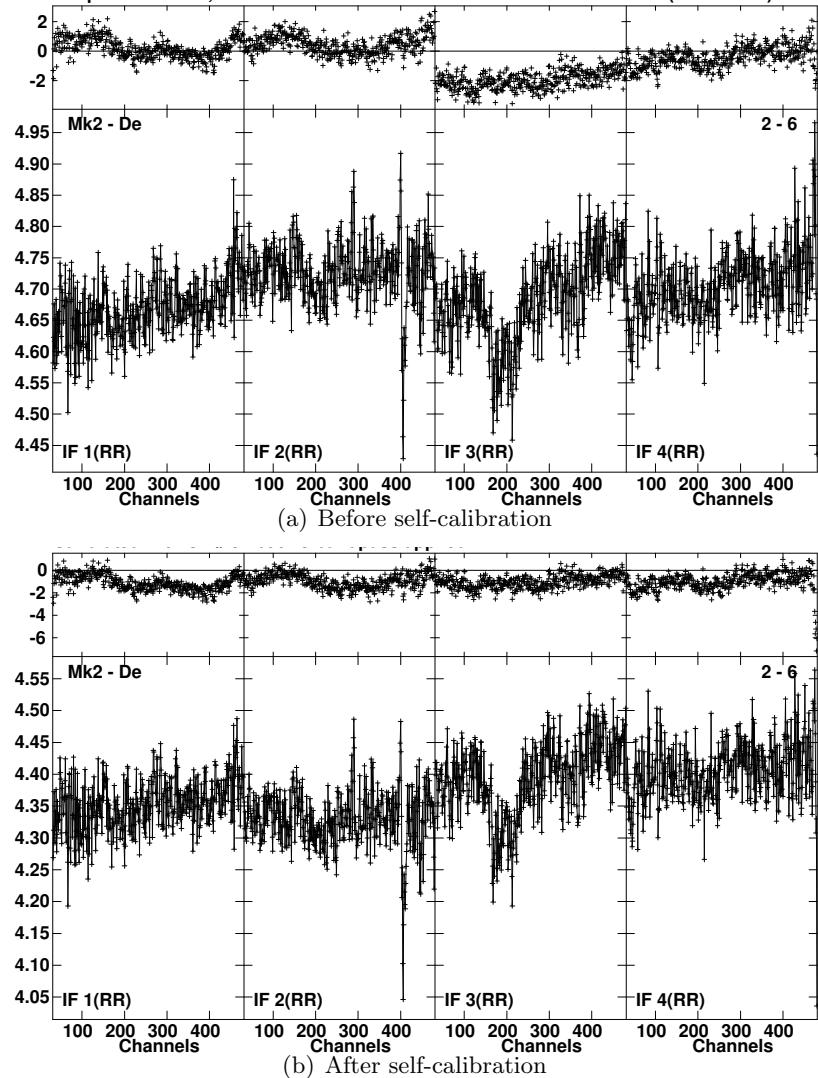


Figure 3.15: Self-calibration showing before (a) and after (b) plots of J2007+4029 from the C-band commissioning dataset, MkII-Defford baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency in Jy.

Finally an image of the phase calibrator J2007+4029 at C-band is made to demonstrate

¹ At the current date of this thesis.

the calibration performance. Figure 3.16 shows a contour map of J2007+4029 with the base contour level set at 8 mJy with intervals at -1, 1, 2, 4, 8, 16, 32, 64, 128 \times 8 mJy. The image RMS from the histogram is 4.37 mJy and the image has a peak flux density of 4.49 Jy/ beam with a dynamic range of just over 1000:1. It is clear that J2007+4029 is slightly resolved with features on the North and South side of the main component, confirming the two-sided structure description by Kharb et al. (2010).

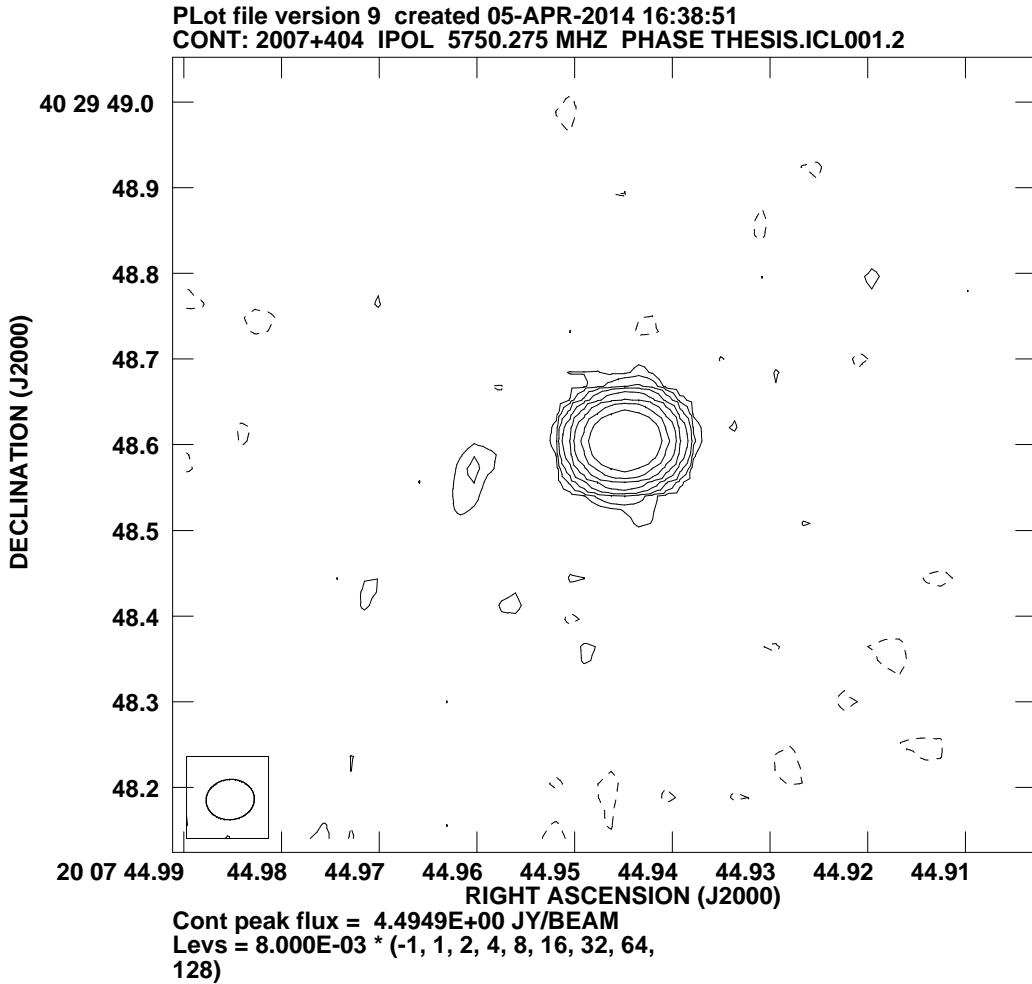


Figure 3.16: C-band contour map of J2007+4029. The base contour level is set at 8 mJy with intervals at -1, 1, 2, 4, 8, 16, 32, 64, 128 \times 8 mJy. The synthesised beam is shown in the bottom left corner.

With the calibration complete for the phase calibrator, the solutions for the complex gains and bandpass are copied over to the target fields and then one final run of CLCAL applies the solutions from the SN table in a CL table for the field. The target field can now be imaged.

3.10 Calibration Performance on COBRaS L-band Legacy Pointings

So far only the results from the C-band commissioning data calibration have been presented. C-band generally has less RFI than L-band and therefore passes through the pipeline within a higher ratio of good solutions at each calibration step. L-band is more problematic, with more RFI increasing the probability that any remaining bad data will produce a lower fraction of good solutions at each stage of the calibration. In essence, L-band is harder to calibrate than C-band data.

Comparing the relative fractions of good solutions of the C-band commissioning data and the L-band Legacy data, the commissioning dataset obtained better solutions than the Legacy data at every stage of the calibration. This further emphasises the difficulty of lower frequency calibration.

All the results presented in this section are of the phase calibrator J2007+4029 from the L-band 4th May 2013 dataset averaged over all scans. Figure 3.17 shows the data before any flagging and calibration has taken place (a), and after the initial delay and phase calibration (b). Some additional editing of the phase solutions was required to remove erroneous solutions using the AIPS task SNEDT. The phase solutions are smoothed out and calibrated as a function of frequency.

After phase calibration, flux scaling is performed using 3C286 as described in Section 3.7 and amplitude calibration to determine the real part of the complex gain correction. Figure 3.18 shows the calibration as a function of frequency before and after flux and amplitude calibration. Again manual editing of the SN table were made to remove spurious gain solutions.

The flux in IF 1 (Figure 3.18) is significantly higher than all of the other IFs. This is likely to be a result of heightened RFI at these frequencies. This is not the strong or narrowband RFI (which is mitigated by SERPent) but very weak and/ or broadband RFI which can cover all of the channels in the IF². This raises the flux level of every visibility in IF 1 and is impossible to mitigate without flagging the entire IF.

The next step is bandpass calibration including corrections for spectral index, using the spectral fit module as described in Section 3.8. The spectral fit is displayed in Figure 3.19. The spectral index of OQ208 is then passed to the AIPS task BPASS and the bandpass

²Such RFI arises from ringing effects which have been discussed in Section 2.4.2.

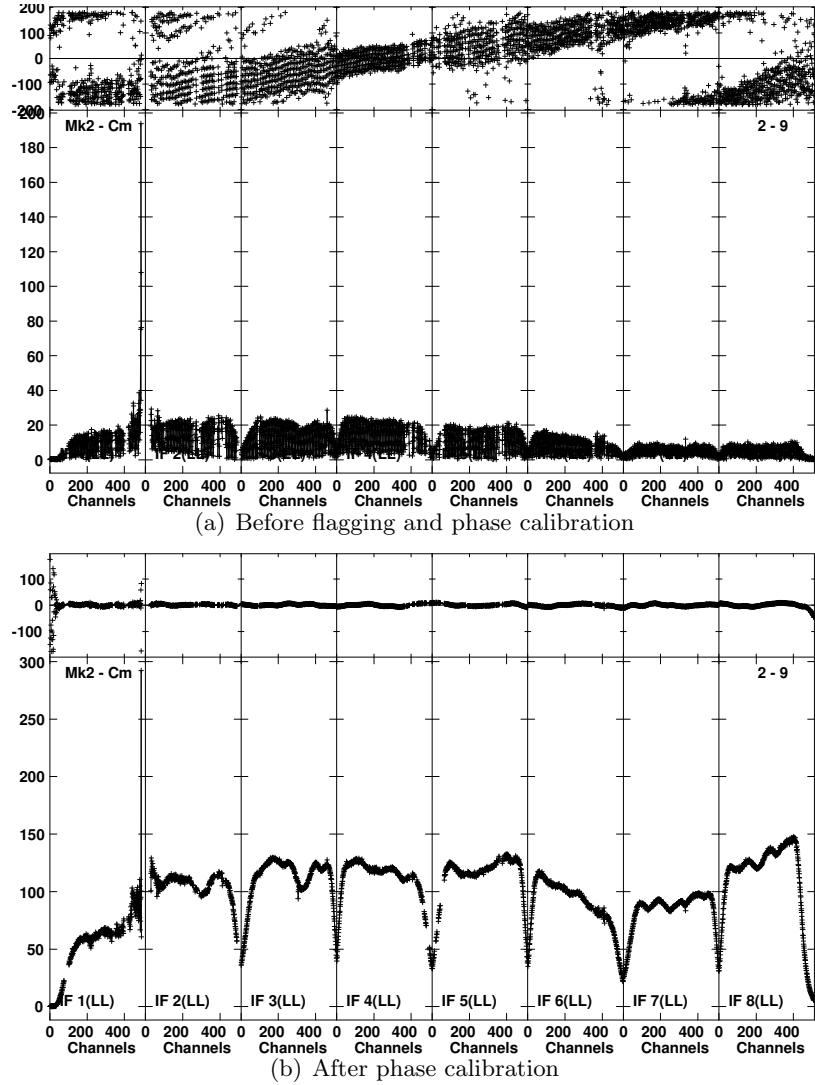


Figure 3.17: Phase calibration showing before (a) and after (b) plots of J2007+4029 from the L-band Legacy dataset, MkII-Cambridge baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency. The y axis is in correlator counts and the x axis is the channel number.

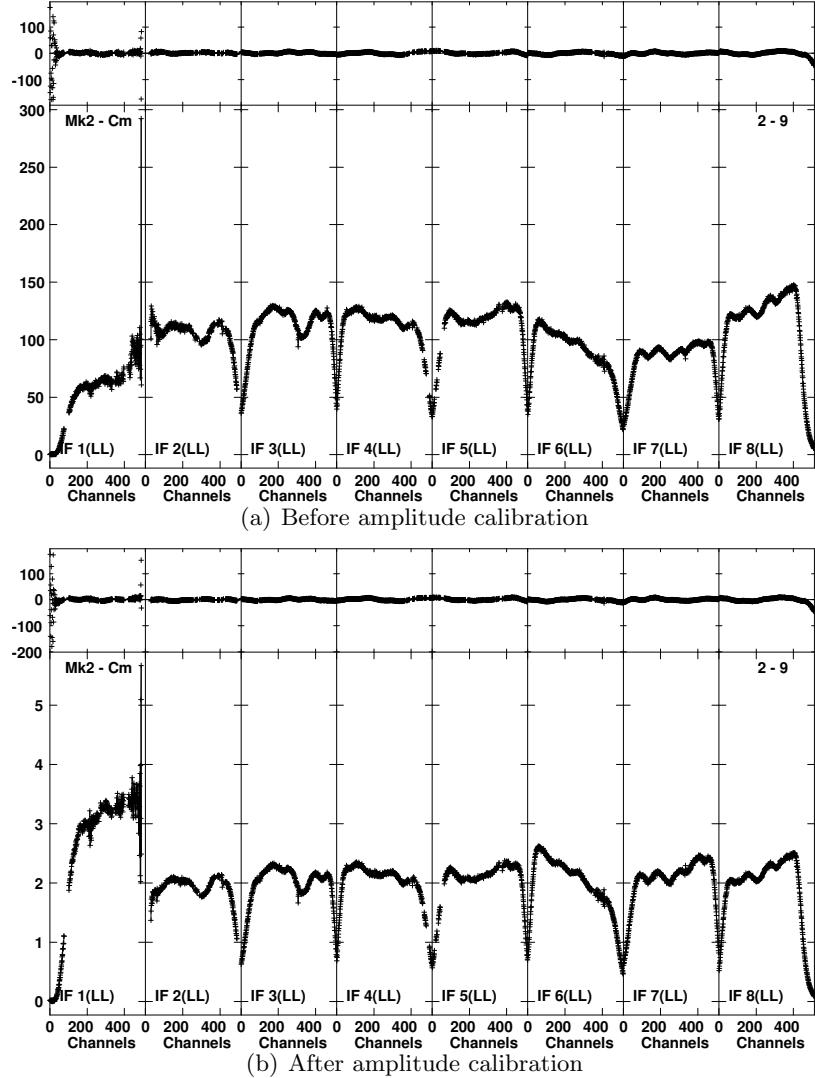


Figure 3.18: Amplitude (and flux) calibration showing before (a) and after (b) plots of J2007+4029 from the L-band Legacy dataset, MkII-Cambridge baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency. The y axis is in correlator counts (a) and Jy (b) and the x axis is the channel number.

response over frequency is flattened. Figure 3.20 shows a before and after plot of the bandpass calibration which includes the spectral index correction.

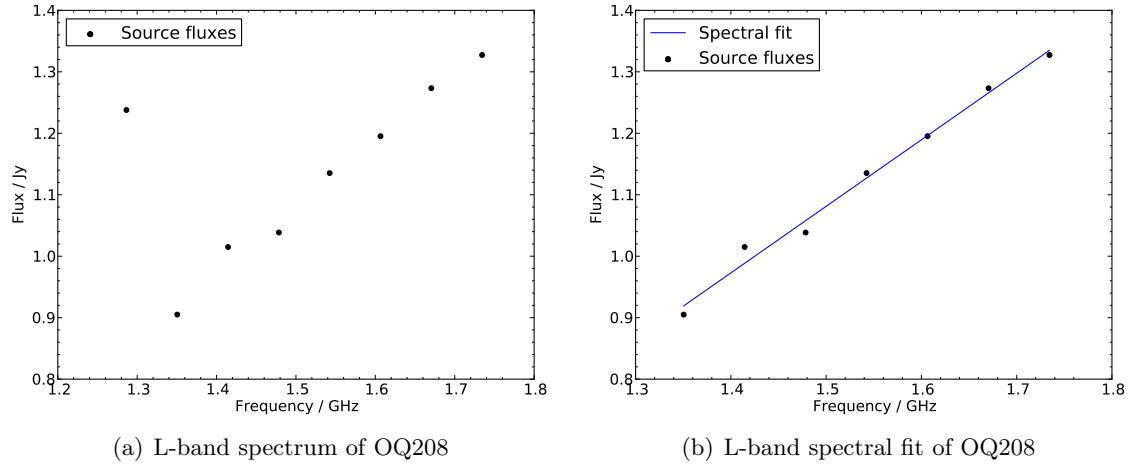


Figure 3.19: L-band spectral fit of OQ208. (a) is the spectrum of OQ208 from amplitude calibration and (b) is the spectral fit of OQ208. Note that the flux at 1.29 GHz (corresponding to IF 1) has been removed from the fit.

The amplitude level of IF 1 in Figure 3.20 is significantly higher than the other IFs, due to remaining RFI. IFs 1 and 2 produce the most problems during calibration because of the large amount of low level broadband RFI at these frequencies, which at times can consume over half the channels within those IFs. This is also why the flux value for IF 1 has been removed during the spectral fit by the spectral module.

Finally an image of the phase calibrator J2007+4029 at L-band is made to demonstrate the calibration performance. Figure 3.21 shows a contour map of J2007+4029 with the base contour level set at 15 mJy with intervals at -1, 1, 2, 4, 8, 16, 32, 64, 128 \times 15 mJy. The image RMS from the histogram is 4.80 mJy and the image has a peak flux density of 2.19 Jy/ beam with a dynamic range of 456:1. The limiting factor in the dynamic range is probably a result of residual weak RFI remaining in the data. It is clear that J2007+4029 is slightly resolved with features on the North and South side of the main component also seen in the C-band image (Figure 3.16).

3.11 COBRaS Imaging and Mosaicing

The next sections contain discussions on wide-field and wide-band imaging and mosaicing techniques with e-MERLIN and COBRaS.

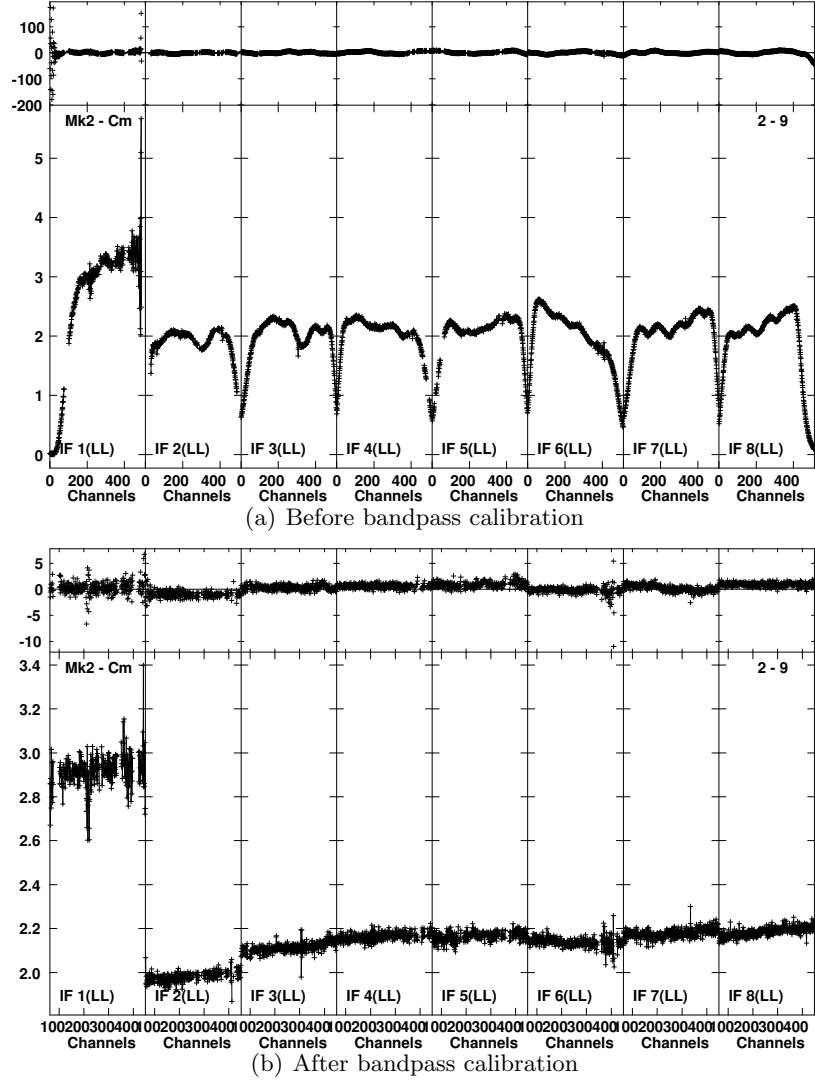


Figure 3.20: Bandpass calibration showing before (a) and after (b) plots of J2007+4029 from the L-band Legacy dataset, MkII-Cambridge baseline. Within each plot; (top) the phases against frequency, (bottom) the amplitudes against frequency. The y axis is in Jy and the x axis is the channel number.

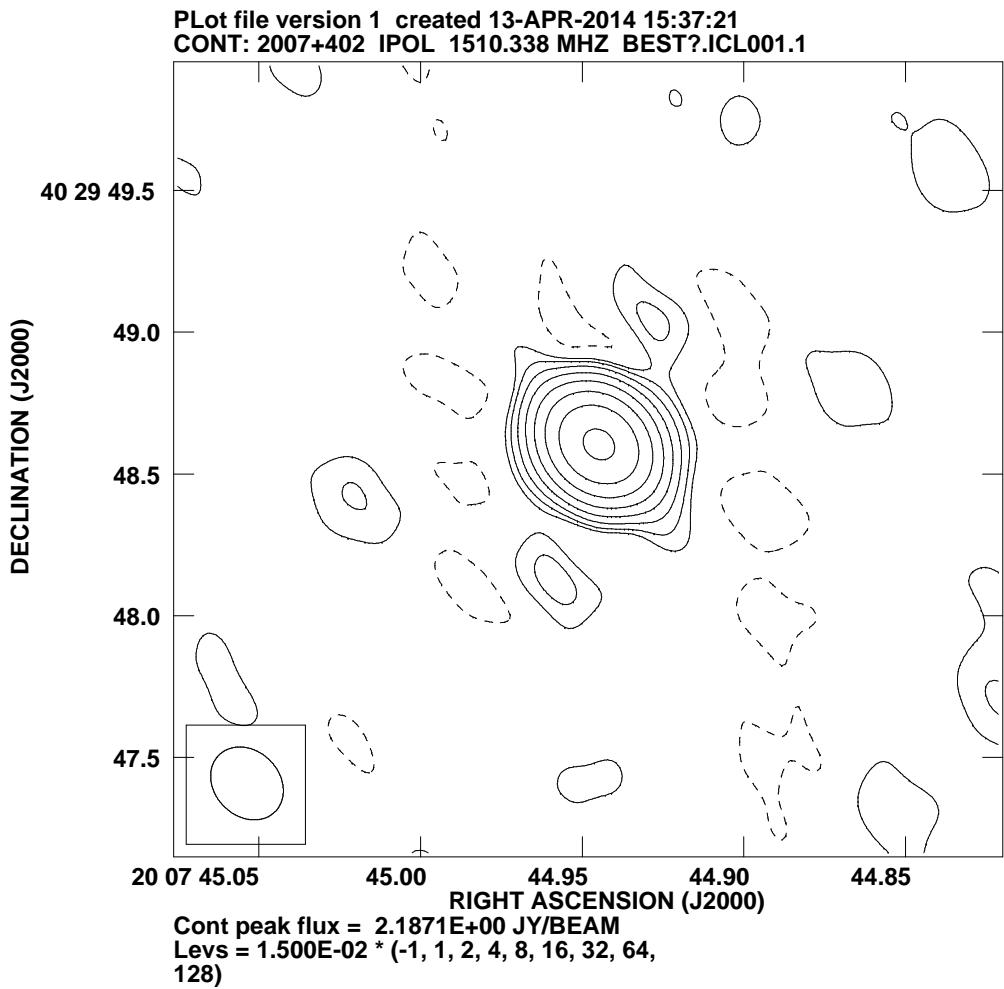


Figure 3.21: L-band contour map of J2007+4029. The base contour level is set at 15 mJy with intervals at -1, 1, 2, 4, 8, 16, 32, 64, 128 \times 15 mJy. The synthesised beam is shown in the bottom left corner.

3.11.1 Wide-band Imaging

The increase in receiver bandwidths leads to an increase in continuum sensitivity in images because of the increased u, v coverage, i.e. interferometer response. This can be easily demonstrated by comparing the dirty beam (the FT of the interferometer response) for an individual IF. Figure 3.22 shows the dirty beams at each end of the e-MERLIN L-band; IFs 1 (central frequency of 1.2865 GHz) and IF 8 (central frequency of 1.7345 GHz). When incorporating the whole bandwidth, the fidelity of the dirty beam is significantly improved as a consequence of lower sidelobe levels.

However, the increased bandwidth introduces additional issues into the imaging process. The multiple frequencies will cause the spectral index to vary significantly over the bandwidth, which is particularly problematic for bright sources. There are no expected

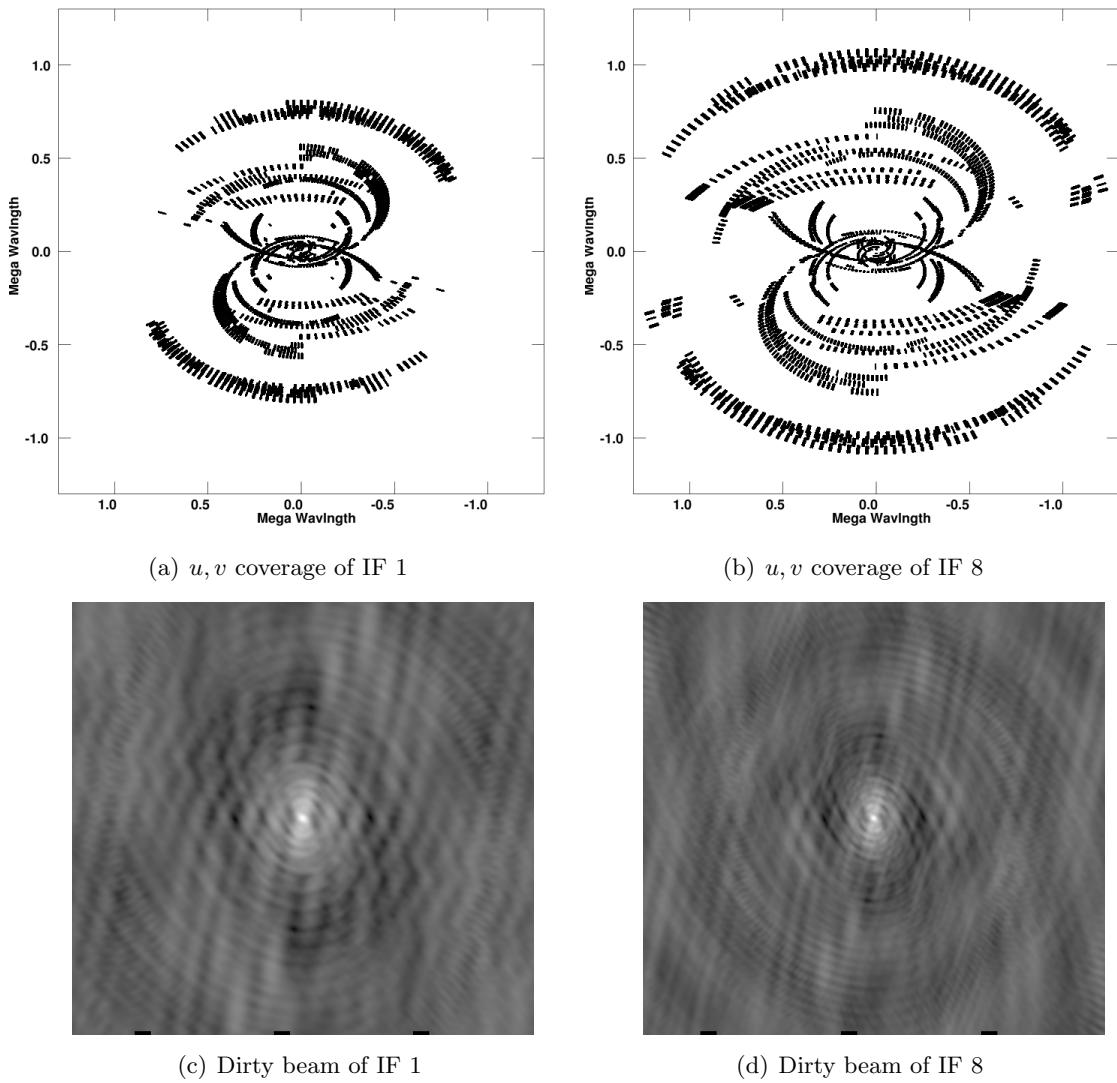


Figure 3.22: The differences in u, v coverage of IF 1 (a; central frequency 1.2865 GHz) and IF 8 (b; central frequency 1.7345 GHz) each with a 64 MHz bandwidth, and the respective dirty beams (c and d).

extremely strong sources in the COBRaS field of view, however, spectral index is a key component to some of COBRaS science goals. This is solved by multi-frequency synthesis which is explained later.

The primary beam is a function of frequency and therefore, over a wide bandwidth will change considerably. At the phase centre, the primary beams at all frequencies are at unity, but further away from the phase centre, radial errors become more pronounced. This can be accounted for by including a primary beam spectrum i.e. a primary beam model in the multi-frequency synthesis and by applying wide-field imaging techniques such as facets, which is explained in Section 3.11.2.

Furthermore, the angular resolution also scales with frequency, resulting in different possible resolutions across the bandwidth. A simple solution is to image at the lowest resolution i.e. at the lowest frequency. However, multi-frequency synthesis allows not only for imaging at the highest possible resolution (highest frequency) but also retains the correct spectral index which is affected by the resolution (or strictly, it is limited to the highest resolution where the signal is detected; Rau and Cornwell 2011).

Conway et al. (1990) produced the first multi-frequency synthesis method ‘*double deconvolution*’, where two successive CLEAN deconvolutions determine two components of the dirty beam; the normal dirty beam $B_0(l, m)$ and the spectral dirty beam $B_1(l, m)$ which estimates the spectral index α . This is essentially expressing the spectral function (flux as a function of frequency $I(\nu_0)$) of the dirty beam as the first two coefficients of a Taylor expansion as follows:

$$I_D(l, m) = I(\nu_0) B_0(l, m) + \alpha I(\nu_0) B_1(l, m). \quad (3.14)$$

Sault and Wieringa (1994) describe a multi-frequency CLEAN algorithm which calculates the Taylor coefficients via least squares minimisation. More recently, Rau and Cornwell (2011) presented an algorithm which combines multi-scale, deconvolution with different spatial scales, and multi-frequency synthesis (MS-MFS) which can be used in conjunction with other wide-field imaging algorithms. These later algorithms also determine the spectral curvature (3^{rd} Taylor coefficient) and higher orders of spectral information. MS-MFS is implemented in the reduction package CASA.

In the case of COBRaS imaging, the AIPS task IMAGR is an all-purpose cleaning task which utilises corrections and enhancements for wide-bandwidth (and wide-field) imaging. Addressing the issues stated above, a number of image parameter adverbs (IMAGRPRM) are available to allow corrections during deconvolution. The frequency dependent primary beam effects can be corrected via IMAGRPRM(1) = D, where the primary beam is assumed to be a uniformly illuminated disk of diameter D metres³.

A simple correction to the spectral index is available in IMAGR to allow a single correction to be made. Considering that the COBRaS maps will contain multiple point sources, each possibly having a different spectral index, one correction will not suffice. Instead spectral information will be gleaned from obtaining fluxes from imaging individual

³AIPS IMAGR help file.

IFs (or a couple of IFs) for objects bright enough to be detected with the temporally smaller bandwidth and from obtaining fluxes at both L-band and C-band. For maximising the absolute detection of sources, the spectral information can be neglected.

3.11.2 Wide-field Imaging

Wide-field imaging is described as an image with a large number of resolution elements and samples far out into the primary beam, which introduces distortions in the image arising from a number of origins.

Bandwidth Smearing

Bandwidth smearing (chromatic aberration) produces radial smearing which becomes more severe further away from the phase centre. If averaging has occurred then the gains as a function of frequency may not have been accurately determined before averaging. Otherwise it is an effect of the wide bandwidth with different frequencies probing different spatial scales.

The fractional reduction in amplitude of a point source due to bandwidth smearing $R_{\Delta\nu}$ is the ratio of intensity I to the peak response I_0 at $\Delta\nu = 0$, given by (Bridle and Schwab 1999; MERLIN user guide):

$$R_{\Delta\nu} = \frac{I}{I_0} = \sqrt{\frac{\pi}{4 \ln 2 \beta^2}} \operatorname{erf}(\sqrt{\ln 2} \beta) \quad (3.15)$$

where β is the fractional bandwidth \times the radial distance in arcsec from the phase centre in half power beam widths (HPBW) and is given by:

$$\beta = \frac{\Delta\nu}{\nu_0} \frac{\Theta_0}{\Theta_{\text{HPBW}}}. \quad (3.16)$$

Time-average Smearing

Time-averaging smearing (de-coherence) produces tangential smearing. This arises from the poor phase gain solutions for the observation which has then been averaged in time, losing coherence over the averaged time. Accurate solutions may have been determined for a point source at the phase centre, but further away, residual phase errors may remain.

The average fractional reduction in amplitude of a point source due to time-average

smearing $\overline{\langle R_\tau \rangle}$ is given by (Bridle and Schwab 1999; MERLIN user guide):

$$\overline{\langle R_\tau \rangle} = 1 - 1.22 \times 10^{-9} \left(\frac{\Theta}{\Theta_{\text{HPBW}}} \right)^2 \tau_a^2 \quad (3.17)$$

where Θ is the radial distance from the phase centre in compressed coordinates ($\Theta = \sqrt{l^2 + m^2 \sin^2 \delta}$), and τ_a is the averaging time. This relation is strictly only correct for a circumpolar point source i.e. $\delta = 90^\circ$.

Primary Beam Response

The response of the primary beam also affects wide-field imaging, because further out into the field the interferometer's response decreases. The AIPS task PBCOR can be used on images to correct the frequency dependent effects from the primary beam. This requires modelling of the primary beam response for each baseline for a heterogeneous array such as e-MERLIN. Another way is to use the previously stated IMAGR adverb IMAGRPRM(1), which makes a simple frequency dependent correction during imaging.

Non-coplanar Baselines

The last effect from wide-field imaging to consider is the assumption of coplanar arrays, which is valid for East-West interferometers i.e. antennas that lie on the same latitude, and for small fields of view, to neglect curvature of the celestial sphere. Non-coplanar effects can be considered either as the 3D visibility function $V(u, v, w)$ or the Sky Brightness $B(l, m, n)$ (which is the Fourier transform of the visibility function). Essentially, the third dimensional direction cosine in either term has to be accounted for. Doing a full 3D Fourier transform is computationally expensive and therefore corrections are made to project sources from the 3D sky to a tangential 2D plane. Figure 3.23 shows the effects of a wide field image on the third dimensional term and how sources in the wide field are projected to a flat plane.

There are two methods to correct the non-coplanar effects, the first is to break up the field into many facets, or smaller fields, hence approximating to many 2D fields where the sky curvature is negligible, or applying corrections to the w -term of the visibility function by reprojection to or from the $V(u, v, w)$ space from or to the $V(u, v, w = 0)$ plane (Cornwell et al. 2008). The two terms are related by the convolution of a known function

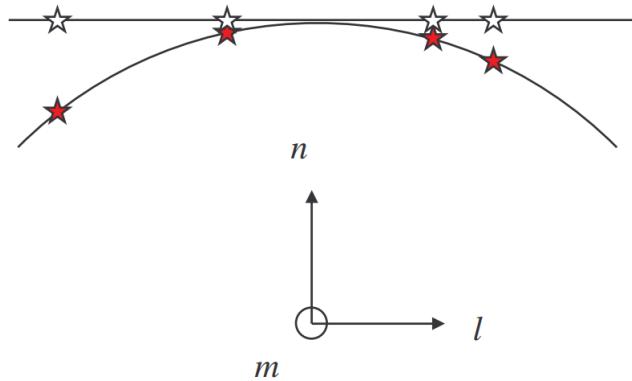


Figure 3.23: Diagram of the effects of wide-field imaging. The red stars are the locations of sources within the 3D sky and the white stars are the same sources projected onto a tangential 2D plane. Figure taken from Muxlow (2007).

$\tilde{B}(u, v, w)$ (the Fourier transform of $B(u, v, w)$) giving

$$V(u, v, w) = \tilde{B}(u, v, w) \star V(u, v, w = 0). \quad (3.18)$$

The w-projection algorithm is implemented in CASA, with no equivalent in AIPS currently.

The method of splitting the imaging field into facets solves a number of the issues above. Bandwidth and time-averaging smearing and non-coplanar effects can all be solved using many facets. For each facet, the specific gain solutions, frequency-dependent solutions and three dimensional positional information are supplied and then solved with the 2D Fourier transform. Facets can be placed at any point in the primary beam, on known sources of interest or overlapping to cover the entire primary beam. The latter is required for COBRaS. Figure 3.24 shows how the method of faceting solves the three dimensional problem for wide-field imaging.

IMAGR images and CLEANS the multiple facets with the adverb OVERLAP = 2, which subtracts the clean components from the current facet (and all other facets) before imaging the next strongest facet. The AIPS task FLATN is then used to regrid the facets from IMAGR onto a single map.

3.11.3 Imaging Considerations with e-MERLIN

e-MERLIN is an heterogeneous array with each antenna having a different primary beam. Not only will the response across each primary beam differ due to different antenna shapes and surfaces, but there are also three different antenna sizes: 25m, 32m and 76m. This results in the interferometer beam shape being baseline dependent. Currently, there are

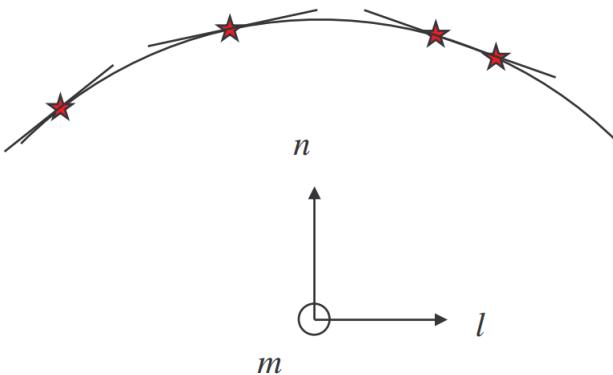


Figure 3.24: Diagram of the method of faceting. The example has split the original wide-field image into three separate facets which approximates the local fields as tangential 2D planes. Figure taken from Muxlow (2007).

on-going investigations into the e-MERLIN combined beam responses, which will consider the frequency dependent effects of the wide bandwidth as well as the varying beam shapes from the differing antenna sizes (Wrigley et al. in preparation).

3.12 e-MERLIN Mosaicing Techniques

COBRaS is comprised of 7 L-band pointings and 42 C-band pointings centred on the core region of Cyg OB2. Each pointing overlaps with one or more adjacent pointings to give continual coverage. In addition, the overlapped regions also increase the sensitivity due to the increase in u, v coverage.

e-MERLIN requires a unique mosaicing technique resulting from its heterogeneous nature. As mentioned in Section 3.11.3, e-MERLIN contains three different antenna diameters: 25m, 32m, and 76m. Baselines containing the Lovell telescope will benefit from an increase in sensitivity for the inner 50% of the 25m - 25m interferometer beam by a factor of ~ 2 . Figure 3.25 demonstrates these properties.

Not including the Lovell telescope during imaging will result in a dramatic increase in the field of view. If the original mosaic strategy considers the overlap regions assuming the inclusion of the Lovell telescope, then additional imaging without the Lovell will produce pointings which will overlap a greater area into adjacent pointings, increasing the sensitivity in those areas. Essentially there will be two maps generated, one with the Lovell and one without. Combining these together will result in the edge regions of the overlapped pointings to increase in sensitivity.

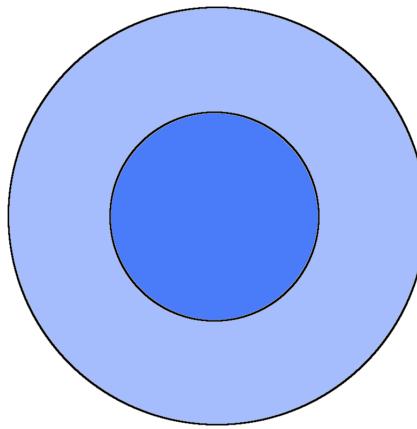


Figure 3.25: Sensitivity areas for the Lovell baselines, 76m - 25m interferometer beam (darker shade blue region), and the 25m - 25m interferometer beam (lighter shade blue region). The Lovell region has a greater sensitivity around a factor of 2 compared to the same region within the 25m - 25m beam.

3.13 Discussion

The COBRaS pipeline reduces, and calibrates L-band or C-band continuum data with passages for amplitude, phase, bandpass with spectral index corrections and self-calibration. It has considerations for K-band, such as defaults for the atmospheric coherence time. It does not apply any corrections for polarisations and other polarisation-dependent effects such as leakage etc. It also includes methods to complete simple manual tasks automatically, for example; selecting a small window for delay correction during phase calibration and selecting an IMAGR window for self-calibration. These types of tasks arise from the calibration process previously being a manual process and therefore all the AIPS tasks requiring manual input and choices.

The pipeline only calculates antenna based corrections and does not consider baseline based corrections of any kind (closure offset corrections). This is because baseline based corrections are usually required for images with high dynamic ranges $> 10,000:1$ (i.e. highest signal-to-noise ratio in the image). The expected sources in the COBRaS field of view are on the order of mJy or μ Jy with the best possible rms of 3 μ Jy and 7.5 μ Jy for 5 GHz and 1.6 GHz respectively (COBRaS proposal) and realistic noise levels of 12 -14 μ Jy for current e-MERLIN Legacy L-band datasets (Tom Muxlow; private communication). Therefore the expected highest flux from a source (background galaxy or a massive binary interaction region) in the COBRaS field is on the order of 10's mJy, which corresponds to a dynamic range of $\sim 1000:1$.

This pipeline also assumes isoplanicity between the phase calibrator and the target fields, because there are no known bright point-like sources in any of the fields. Therefore the assumption that applying the calibration made on the phase calibrator (J2007+4029) to the target fields is made. Whilst an entirely plausible and valid assumption, it comes with the caveat that the phase solutions will be limited to a certain degree, not only by the different atmospheric path between the phase calibrator and fields but also due to time-average smearing (discussed in Section 3.11.2).

Ideally a phase calibrator would exist in each of the target fields, and should a sufficiently suitable source be found after the initial COBRaS analysis, then further self-calibration may be possible to refine phase solutions within the target field. This is more likely for L-band where the field of view is significantly larger than at C-band.

Another improvement in the calibration procedure would arise from incorporating spectral indices in the CLEAN components during self-calibration (IMAGR). Currently IMAGR does not interpret the spectral index, and therefore the pipeline attempts to include this information by performing self-calibration on individual IFs. This is not an optimal solution to this issue, but it is the easiest to implement using AIPS and IMAGR.

In this chapter, the pipeline is tested on C-band commissioning data and L-band legacy data. Despite the commissioning nature of the C-band dataset, the data passed through the pipeline without issue, with a good image and dynamic range of the phase calibrator map. The L-band data, however, were more problematic, requiring the solutions from phase and amplitude calibrations to be manually edited with the AIPS task SNEDT. This may result from insufficient data quality, but is likely to be a result of very weak RFI remaining in the dataset. Avenues of pursuit to help SERPent and increase successful RFI-mitigation have been discussed previously (Section 2.4.2). This will help the calibration pipeline achieve good solutions and output maps, something it has proven with the RFI-quiet C-band observations. However, what is necessary is an automated passage to assess and edit the solutions from phase and amplitude calibration. Such a passage has recently been added to the e-MERLIN pipeline and because of the modulised nature of the e-MERLIN and COBRaS pipelines, it should be easy to integrate the SN editing module into the COBRaS pipeline.

As stated above, improvements to the pipeline will occur with more tests with the rest of the L-band Legacy datasets to increase robustness. Such additional tests will optimise the pipeline in preparation for the C-band COBRaS dataset, where effects from

a bandwidth of 2 GHz will produce even larger datasets and therefore greater frequency dependent errors. There are six times the number of individual pointings for C-band than L-band, making the stability of the pipeline imperative.

Creating an automated calibration pipeline is not trivial, and this is especially the case with e-MERLIN, due to the heterogeneous nature of the array. This is because a number of issues can arise from the data which will propagate through the pipeline, and accounting for every eventuality is difficult. It is not possible to create a pipeline which caters to every need. A general pipeline can be created to produce a simple map to give an idea of what is present in the field, or a specialised pipeline tailored to a single requirement.

With the abundance of sophisticated imaging software, CASA would provide a convincing argument as a choice of package to calibrate and image data. However, e-MERLIN requires fringe-fitting to correct the phases and no corresponding task currently exists in CASA. Converting terabytes worth of AIPS UVFITS data to CASA measurement sets is computationally expensive and with adequate equivalent algorithms in AIPS for imaging and mosaicing, seems unnecessary.

New calibration software such as MeqTrees (Noordam and Smirnov 2010) and imaging algorithms exploiting compressed sensing (Wiaux et al. 2009; McEwen and Wiaux 2011) are becoming available for 3rd generation interferometers. The MeqTrees package is designed and optimised for LOFAR and the SKA, in the same way AIPS is for the VLA/ JVLA and CASA for ALMA and the JVLA. That is not to say that this package is not applicable to other arrays, but to the author’s knowledge there is no example of this package being utilised with e-MERLIN.

Compressed sensing algorithms have shown great promise on simulated data (McEwen and Wiaux 2011) incorporating modern interferometric issues such as wide-field imaging. However, it remains unknown how such algorithms fair with real observations and all the problems and issues that come with them. The authors state that the *spread spectrum phenomenon* (termed w -projection in this thesis) has been idealised to a constant w , which will not be the case for a real instrument where values range from $w = 0$ to $w = u_{\max}$. Once again highlighting the importance of real data on testing algorithms, a conclusion also made from Chapter 2. Moreover, compressed sensing algorithms do not currently retain spectral information from the images.

It is clear that extensive research is needed on compressed sensing to meet the demands of other existing and operational algorithms. However, the potential and power of

compressed sensing should not be understated, and it is likely to be a popular route of investigation for imaging with the SKA, given the speed and performance of the algorithm.

Discussions and research are now starting with existing packages for calibration and imaging, comparing all the virtues and vices of these with simulated and real observations to decide whether the SKA can utilise existing software (and if so which package), or whether new SKA-specific programs will need to be created (Danielle Fenech; private communication).

Chapter 4

Source Detection, Extraction and Classification

All that glisters is not gold.

William Shakespeare - The Merchant of Venice

After reduction, calibration and imaging, the final task for the COBRaS pipeline is to analyse the resulting radio maps of the Cygnus OB2 association. This is done in three steps: source detection, source extraction and source classification. The classification step relies on existing knowledge of the association from the catalogues of previous studies. This chapter discusses techniques for each of these and presents the developed tools designed to accomplish the science goals set by COBRaS.

4.1 Building the Cyg OB2 Amalgamated Catalogue

The COBRaS Legacy survey expects to detect a number of different sources within the core region of Cyg OB2. These range from single OB stars, OB binaries, pre-main sequence stars, and background radio galaxies. Many OB stars will not have been either formally classified or even detected by previous studies over a number of observed wavelengths before the COBRaS observations. Therefore, COBRaS can be considered as a *semi-blind*

survey, where known positions of candidate sources will help the source detection and classification but is also a survey which will find new radio sources and should expect no external contribution.

It is therefore instructive to collect all known information on the area of interest, i.e. the Cyg OB2 association. Thus a definitive catalogue, amalgamated from all previous studies in any wavelength regime will be a powerful tool for COBRaS, and indeed, any large survey.

Catalogue amalgamation is no easy process, with the main issue being (or there not being) standardisation in cataloging practices. Positional information obviously varies due to different levels of accuracy and resolution for every survey. The same is the case for any quantitative information (fluxes, magnitudes etc.), but other fundamental information such as identifiers or names of sources, can, and should be standardised for specific regions. This makes archiving much simpler.

To aid this process, the Virtual Observatory (VO) library is used to find existing catalogues in its vast database, and the TOPCAT/ STILTS (Taylor 2005; Taylor 2006) programs with the Java scripting language are used to concatenate and manipulate these catalogues. TOPCAT/ STILTS uses the xmatcher algorithm to match catalogues by their sky coordinates, and has options to cross correlate any information given in the catalogue columns, e.g. names or identifiers.

A Jython (an implementation of the Python programming language written in Java) script is created and contains passages to cross correlate the input catalogues, create sub catalogues such as the OB catalogue (Section 4.1.6), and amend additional columns for near-infrared colours (e.g. J-H etc.), stellar parameters for OB stars (Prinja et al. 1990; Martins et al. 2005; Searle et al. 2008), mass loss rates as a function of metallicity (Vink et al. 2001) and empirical radio flux predictions using radio free-free emission (Wright and Barlow 1975). These fundamental parameters are discussed and investigated in Chapter 5.

The ALADIN interactive sky atlas (Bonnarel et al. 2000) is also employed in this chapter to visually represent the catalogues onto archival images such as those from the Sloan Digital Sky Survey (SDSS).

By placing restrictions on the VO search cone of 24 arcminutes around the central coordinates of the Cyg OB2 core region, right ascension (RA): $20^{\text{h}} 33^{\text{m}} 10.8^{\text{s}}$, and declination (DEC): $41^{\circ} 13' 12''$, a list of catalogues with sources in the Cyg OB2 region is found.

Table 4.1 shows all of the catalogues used in creating the Cyg OB2 Amalgamated Catalogue. For the Radio Master Catalogue a wider search radius of two degrees is adopted to consider any future radio observations of the region.

Table 4.1: Table of existing Cyg OB2 catalogues

Reference	Common Identifiers	Number of entries
Massey and Thompson (1991)	MT91	801
Condon et al. (1998)		72
Pigulski and Kołaczkowski (1998)	MT91	288
Comerón and Torra (2001)		320
Comerón et al. (2002)	Various	85
Setia Gunawan et al. (2003)		31
Wolff et al. (2007)	MT91	13
Albacete Colombo et al. (2007b)	AFM2007 & 2MASS	1003
Albacete Colombo et al. (2007a)	AFM2007 & 2MASS	147
Martí et al. (2007)		153
Kiminki et al. (2007)	MT91	303 (from 3 tables)
Kiminki et al. (2008)	Various	17
Vink et al. (2008)	Various	71
Kiminki et al. (2009)	Various	22
Wright and Drake (2009)	AFM2007	1696
Skiff (2010)	Various	150
Kobulnicky et al. (2010)	Various	17
Kiminki et al. (2012a)	Various	21
Kiminki and Kobulnicky (2012)	Various	46
Kobulnicky et al. (2012)	Various	28
Comerón and Pasquali (2012)	Various	240
Chandra Point Source Catalogue	AFM2007	1003
Radio Master Catalogue	Various	2850
Galactic O star Catalogue		15
Simbad Database	Various	2077

Chandra Point Source Catalogue, contains the same sources as Albacete Colombo et al. (2007b) catalogue but with different information.

Radio Master Catalogue, from two earlier catalogues: Dixon (1970) and Kuehr et al. (1979).

Galactic O star Catalogue (Maíz-Apellániz et al. 2004).

Simbad Database, VO search result with search radius of 24 arcminutes.

Table 4.1 shows a list of constituent catalogues which form the Cyg OB2 Amalgamated Catalogue. Most catalogues have a number of different identifiers (denoted as Various in Table 4.1), or multiple naming categories within the same ID column. Cross matching these is difficult, because any difference in the strings (e.g. whitespace) will be recognised as a false match. Furthermore, Java and STILTS do not allow cross matching of string-types. However, two common identifiers are present; MT91 (Massey and Thompson 1991) and AFM2007 (Albacete Colombo et al. 2007b). These form the base of two separate cat-

alogues before being co-added. All other smaller catalogues are added with sky positional information with a matching error radius. The error radius can be chosen to be small or large, and since TOPCAT can pick the best match out of all the candidate matches, a large radius (e.g. 2 arcseconds) can be chosen with confidence.

Very large surveys such as 2MASS (Skrutskie et al. 2006), are not included because the source density is too high and cross correlation will result in many false positive matches. Therefore the 2MASS survey is used as an independent catalogue for further studies, e.g. the JHK colour plots for OB candidate selection (see Section 4.1.7).

4.1.1 Catalogue Amalgamation with Massey and Thompson (1991) Identifiers

The first major census of Cyg OB2 with modern CCDs was conducted by Massey and Thompson (1991). Over 800 members within the core region of the Cyg OB2 association are catalogued in this UBV photometric and spectroscopic study (IDs are denoted by MT91) which, as stated previously, provides a good foundation to build a comprehensive catalogue of the Cyg OB2 region. Three other catalogues have the MT91 identifier information and are combined with the Massey and Thompson (1991) catalogue using the STILTS command *tmatch2*(‘*matcher=1d*’), matching the MT91 ID numbers of each entry catalogue.

Firstly, three tables from Kiminki et al. (2007) (Tables 2, 3, and 5 from Kiminki et al. 2007) are combined with the Massey and Thompson (1991) catalogue to form the first part of the Cyg OB2 Amalgamated Catalogue. This is a six year (1999-2005) radial velocity survey on 146 OB stars within the Cyg OB2 association, to collect evidence of binarity of the massive star population. All 303 entries have MT91 IDs. The next catalogue to be added is Wolff et al. (2007), a rotational velocity study of B0 - B3 stars in young clusters. All 13 stars from this study of the Cyg OB2 field have MT91 IDs. The last catalogue with MT91 IDs is Pigulski and Kołaczkowski (1998), a photometric study of variable stars in the Cyg OB2 central region. There are 96 entries which have the same MT91 IDs as those currently in the Cyg OB2 Amalgamated Catalogue, 182 entries have extended MT91 IDs (numbers which are not in the original Massey and Thompson (1991) catalogue), and 10 entries have no MT91 IDs. These latter 10 are amalgamated using their sky positions as explained in Section (4.1.4).

4.1.2 Catalogue Amalgamation with Albacete Colombo et al. (2007b) Identifiers

The second common identifier from Table 4.1 comes from Albacete Colombo et al. (2007b), a Chandra X-ray study of the Cyg OB2 population. Three more catalogues have the AFM2007 identifier and are again combined using the STILTS *tmatch2* command.

The Chandra point source catalogue contains exactly the same entries as the Albacete Colombo et al. (2007b) catalogue, but contains additional information about the sources. The Albacete Colombo et al. (2007a) is a subset of the original catalogue with detailed information on the flaring nature of some of the X-ray sources. The final catalogue with the AFM2007 ID is Wright and Drake (2009) which also contains 992 out of 1003 sources in the Albacete Colombo et al. (2007b) catalogue and another Chandra observation of Cyg OB2. Entries which do not have the AFM2007 ID information are added by their sky positions.

4.1.3 Catalogue Amalgamation for Cyg OB2 Radial Velocity Survey

The Cyg OB2 radial velocity survey (see original paper; Kiminki et al. 2007) is an on-going survey to map the massive binary content of the Cyg OB2 association. There are a number of subsequent studies: Kiminki et al. (2008), Kiminki et al. (2009), Kobulnicky et al. (2010), Kobulnicky et al. (2012), Kiminki et al. (2012a), and Kiminki and Kobulnicky (2012), all containing catalogues which have corresponding columns of information which are also in the same format, because these originate from the same authors. Therefore, these are amalgamated together prior to the inclusion into the full Cyg OB2 Amalgamated Catalogue for the best matching results. Note that Kobulnicky et al. (2010) is a study on the bow shocks in the Cygnus X region, but has the same catalogue style as those from the massive binary radial velocity study and is therefore included in this concatenation subset.

4.1.4 Catalogue Amalgamation with Sky Coordinates

All remaining catalogues either do not have ID information, or have complex names involving string-types which may include multiple different identifiers, are therefore not standardised and can not be cross matched using STILTS. These catalogues and the two base catalogues (from IDs MT91 and AFM2007) are cross correlated with one another

using the STILTS *tmatch2* command and the RAJ2000 and DEJ2000 (RA and DEC at ephemeris J2000.0) columns. The matching error radius is set to 2 arcseconds, with the option of only selecting the best match. A brief overview of the catalogues collected into the Cyg OB2 Amalgamated Catalogue by the sky coordinates is given below.

Condon et al. (1998). The NRAO VLA Sky Survey (NVSS) of the sky northwards of J2000.0 $\delta = -40^\circ$ at the radio frequency of 1.4 GHz. The complete catalogue contains 2×10^6 sources stronger than $S \approx 2.5$ mJy.

Comerón and Torra (2001). A near-infrared imaging survey of compact HII regions in the Cygnus X complex with JHK filters.

Comerón et al. (2002). A near-infrared spectroscopic survey of the Cyg OB2 association, aiming to list all of the O-type members of the association.

Setia Gunawan et al. (2003). A WSRT radio continuum survey at 1400 and 350 MHz centred on the centre of Cyg OB2, covering 2 square degrees with angular resolutions of 13 and 55 arcminutes respectively.

Martí et al. (2007). A GMRT and VLA survey at 610 MHz and 1400 MHz centred on the TeV J2032+4130 source position.

Vink et al. (2008). The INT Photometric H α Survey of the Northern Galactic Plane (IPHAS), observing towards the Cyg OB2 association and the HII region DR 15.

Skiff (2010). A catalogue of Stellar Spectral Classifications from the VizieR¹ online data catalogue.

Comerón and Pasquali (2012). A near-infrared census study (BJHK filters) using the USNO-B and 2MASS catalogues on the Cyg OB2 region.

Radio Master Catalogue. An amalgamation of Dixon (1970) and Kuehr et al. (1979) catalogues which is regularly maintained.²

Galactic O star Catalogues. A catalogue of Galactic O stars (Maíz-Apellániz et al. 2004).
Simbad Database.

¹<http://vizier.cfa.harvard.edu/viz-bin/VizieR?-source=B/mk>

²<http://heasarc.nasa.gov/W3Browse/all/radio.html>

4.1.5 The Cyg OB2 Amalgamated Catalogue

The Cyg OB2 Amalgamated Catalogue has 25 constituent catalogues with a total of 11469 input entries. This reduces to 6679 entries after the cross correlations by identifiers and sky positions. As certain sources have multiple components, with the catalogues not necessarily denoting these consistently from one to another, duplicates will exist after the matching process. Therefore a tidying algorithm is included to match any duplicates in the Cyg OB2 Amalgamated Catalogue. The catalogue is sorted by RA and a first run is made for identical positional matches, and a second run is made for those within an error radius of 2 arcseconds i.e. the same level of accuracy as the cross correlation. The tidying algorithm outputs a catalogue with 6500 entries.

Ambiguous catalogue entries with more than one component (e.g. Cyg OB2 No. 8) are not dealt with in any special manner. If each component does not have a unique identifier or sky position, then the matched component from any subsequent correlation is somewhat serendipitous in nature. However, when the cross correlation scripts (Section 4.4) are deployed, more than one positive result may be returned. Manual inspection will validate the final classification of the object, where catalogues with a high source densities will require more than just position coincidence to justify the origin of any radio flux. Analysis of the physical properties of each source and of the radio flux will determine the most probable origin.

To illustrate the coverage of this catalogue, Figure 4.1 shows the central square degree of the Cyg OB2 association with the Cyg OB2 Amalgamated Catalogue plotted in red. The core region, and region surrounding the TeV source of Cyg OB2 are clearly identifiable by the catalogue distribution at the centre and off centre of the image respectively.

4.1.6 The OB Star Catalogue

The Cyg OB2 Amalgamated Catalogue contains a variety of objects of interest from massive stars to late-type stars and pre-main sequence stars. For this thesis and the UCL led part of the COBRaS science, the OB stars are of particular interest. The Cyg OB2 Amalgamated Catalogue has multiple columns from many catalogues with spectral classifications for certain stars.

Within the Jython catalogue amalgamation script spectral columns are searched for and combined into one column for reference. If more than one spectral classification

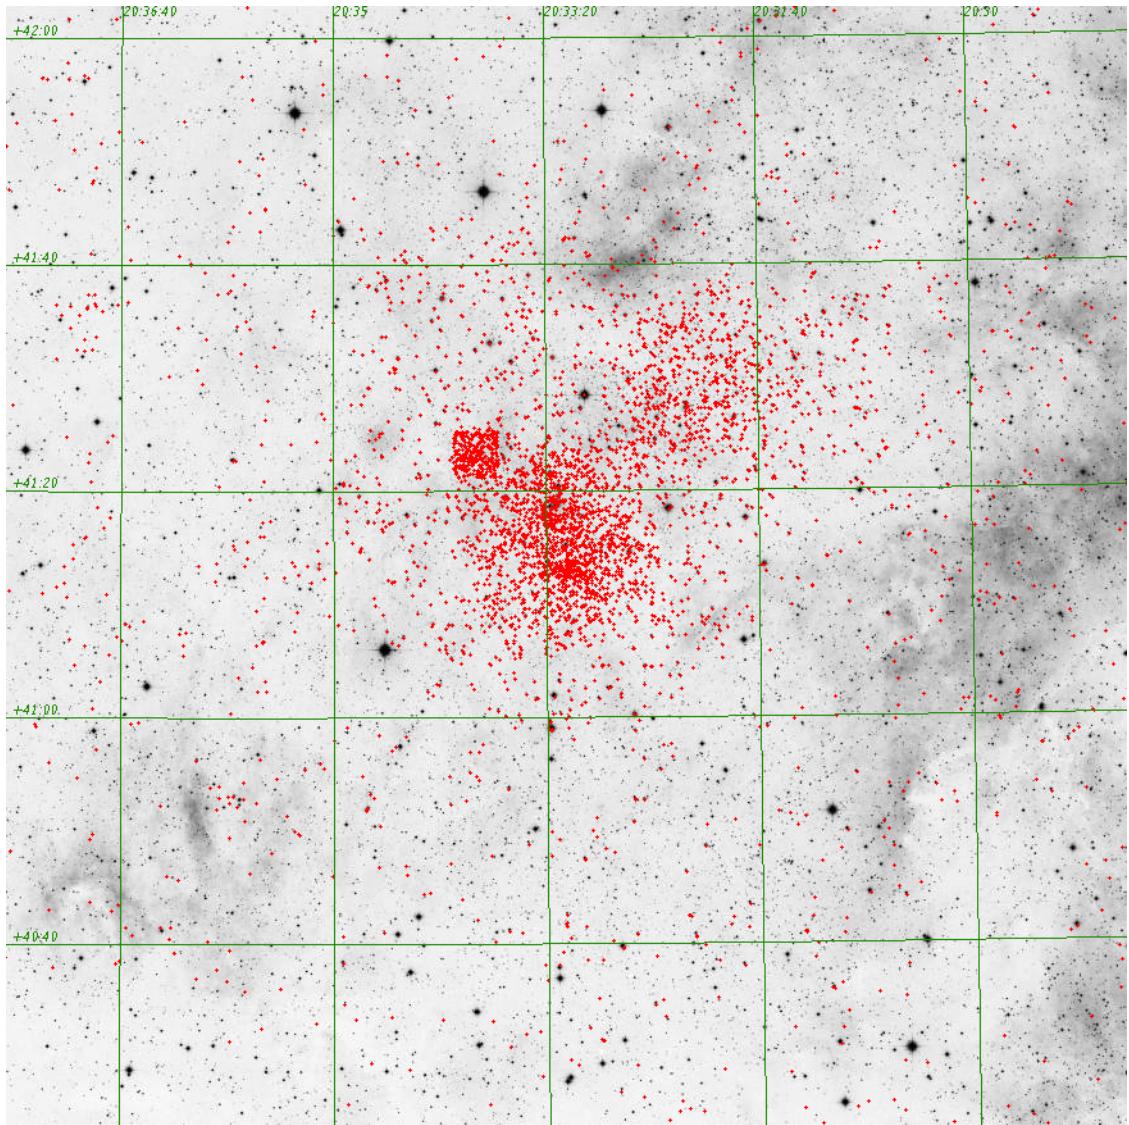


Figure 4.1: The Cyg OB2 Amalgamated Catalogue (red) over laid on a Palomar image at $0.645 \mu\text{m}$ over $1.7^\circ \times 1.7^\circ$.

exists for any source then the first spectral type found in the catalogue is used for this reference column. All other classifications remain in the catalogue in their respective original columns. Then a simple search for strings containing ‘O’ or ‘B’ in the spectral classification are selected to create a known OB catalogue. This known OB catalogue contains 290 stars, and is a valuable tool for not only source classification, but also for source detection (see Section 4.4.1 and 4.4.2).

Figure 4.2 shows the known OB catalogue (blue circles) plotted on the same Palomar image at $0.645 \mu\text{m}$ as in Figure 4.1. The distribution of OB stars is clustered around the core region of Cyg OB2, and extends across the association.

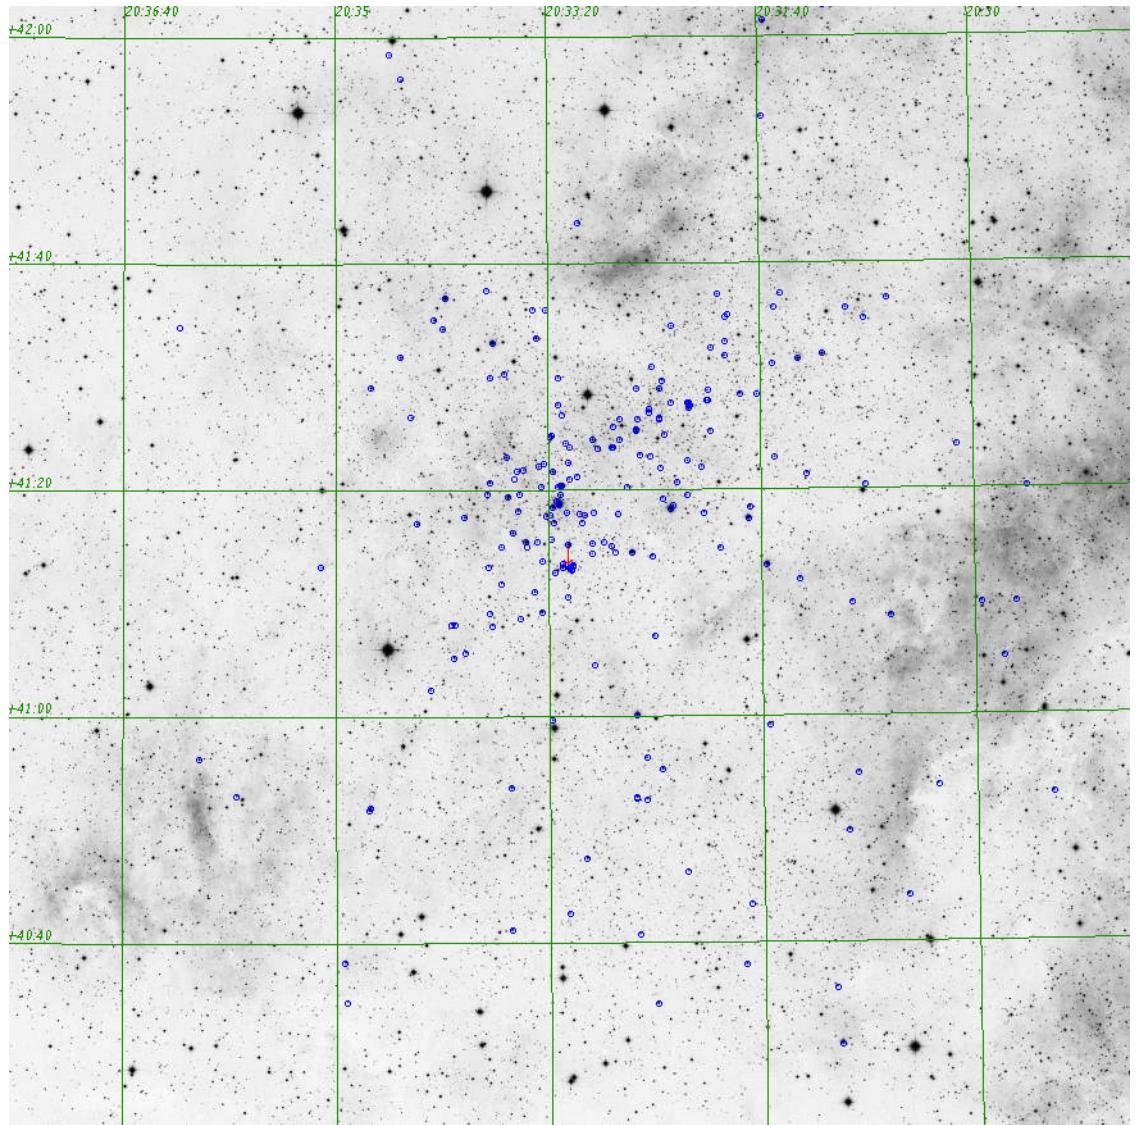


Figure 4.2: The Known OB Catalogue (blue circles) over laid on a Palomar image at $0.645 \mu\text{m}$ over $1.7^\circ \times 1.7^\circ$.

4.1.7 The Candidate OB Star Catalogue

A large number of sources (2584, $\sim 40\%$ of total sources) in the Cyg OB2 Amalgamated Catalogue have magnitudes for near-infrared JHK filters. When plotted together, the J–H and H–K colours can reveal possible spectral information about the population. Comerón et al. (2002) conducted a study using JHK colours from the 2MASS catalogue for the Cyg OB2 region. After following up with spectroscopic observations, the authors discovered that early-type stars resided together in a ‘blue group’, offset from the primary colour branch consisting of unreddened late-type giants. This is visible in Figure 4.3, taken from Comerón et al. (2002) as the cluster of circle points at $(\text{H} - \text{K}) < 0.5$. The authors comment

that the number of non-OB stars contaminating this blue group is very small.

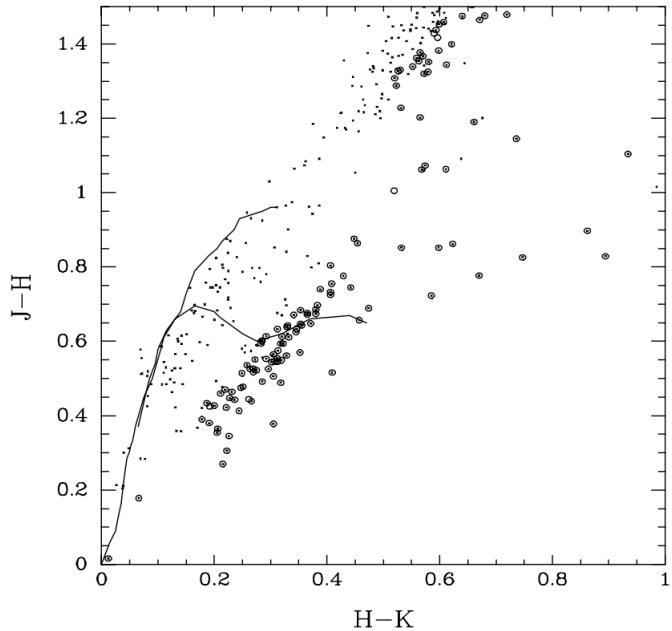


Figure 4.3: J–H versus H–K magnitude colour plot focusing on $(H-K) < 0.5$. The circles denoted sources with spectroscopic observations to verify spectral classification. Figure taken from Comerón et al. (2002).

Applying this technique to sources with JHK magnitudes from the Cyg OB2 Amalgamated Catalogue with the known OB stars with JHK magnitudes from the OB catalogue (see Figure 4.4), it is evident that the ‘blue group’ exists in the Cyg OB2 Amalgamated Catalogue. Moreover, a number of sources from the main Cyg OB2 Amalgamated Catalogue (red) which do not have any spectral classification are also coincident with members of the early-type blue group.

Therefore by using the known OB catalogue to plot a regression line through the blue group with a locus of ± 0.1 in (J–H) and ≤ 0.5 in (H–K), and applying this locus to the super catalogue, a candidate OB catalogue is created from the Cyg OB2 Amalgamated Catalogue. Only sources with $(H-K) < 0.5$ are considered as this is the limit stated by Comerón et al. (2002). Subtracting the known OB catalogue from this newly created candidate catalogue results in 350 additional unclassified but potential early-type stars within the Cyg OB2 association. Figure 4.5 shows the Cyg OB2 Amalgamated Catalogue (red), Known OB Catalogue (blue) and the OB Candidate Catalogue (green).

Until now, the 2MASS point source catalogue has not been considered during the concatenation process of the Cyg OB2 Amalgamated Catalogue. This is because of the

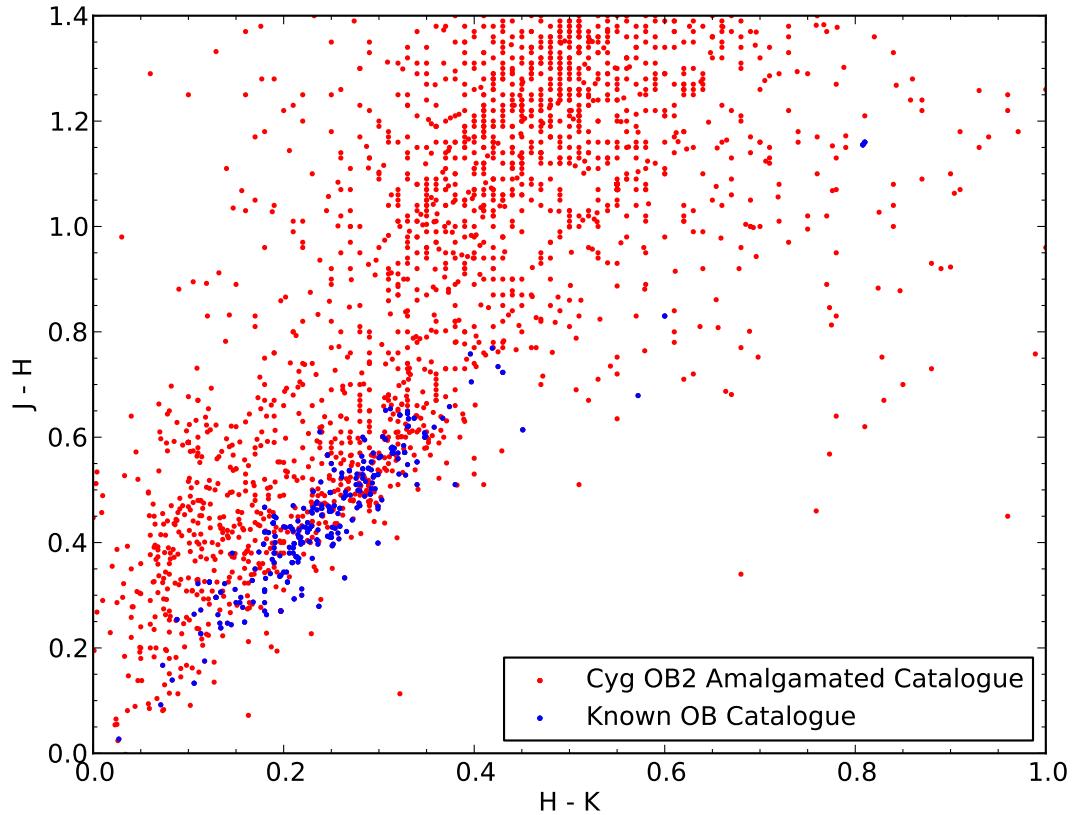


Figure 4.4: J-H versus H-K magnitude colour plot with the Cyg OB2 Amalgamated Catalogue (red) and the known OB catalogue (blue). The majority of the known OB stars reside within the early-type blue group. A number of unclassified (red data points) sources appear to belong to the blue group.

high density of sources from the 2MASS catalogue producing confused source matches during the cross correlation. For a search region of 24 arcminutes centred on the core Cyg OB2 coordinates, the 2MASS catalogue produces 27090 sources.

However, using the same blue group loci produced from the known OB catalogue, but this time with the 2MASS catalogue for the Cyg OB2 region as the input catalogue, produces a larger OB candidate catalogue with 1930 candidates. Figure 4.6 shows the J-H, H-K colour plots for the 2MASS point source catalogue with the 1930 2MASS OB candidates plotted on top.

At the end of the catalogue amalgamation routine, there are four main catalogues for the COBRaS project: the Cyg OB2 Amalgamated Catalogue containing everything from the constituent catalogues, the known OB catalogue containing classified early-type stars, the Candidate OB catalogue containing early-type candidates from their JHK colour

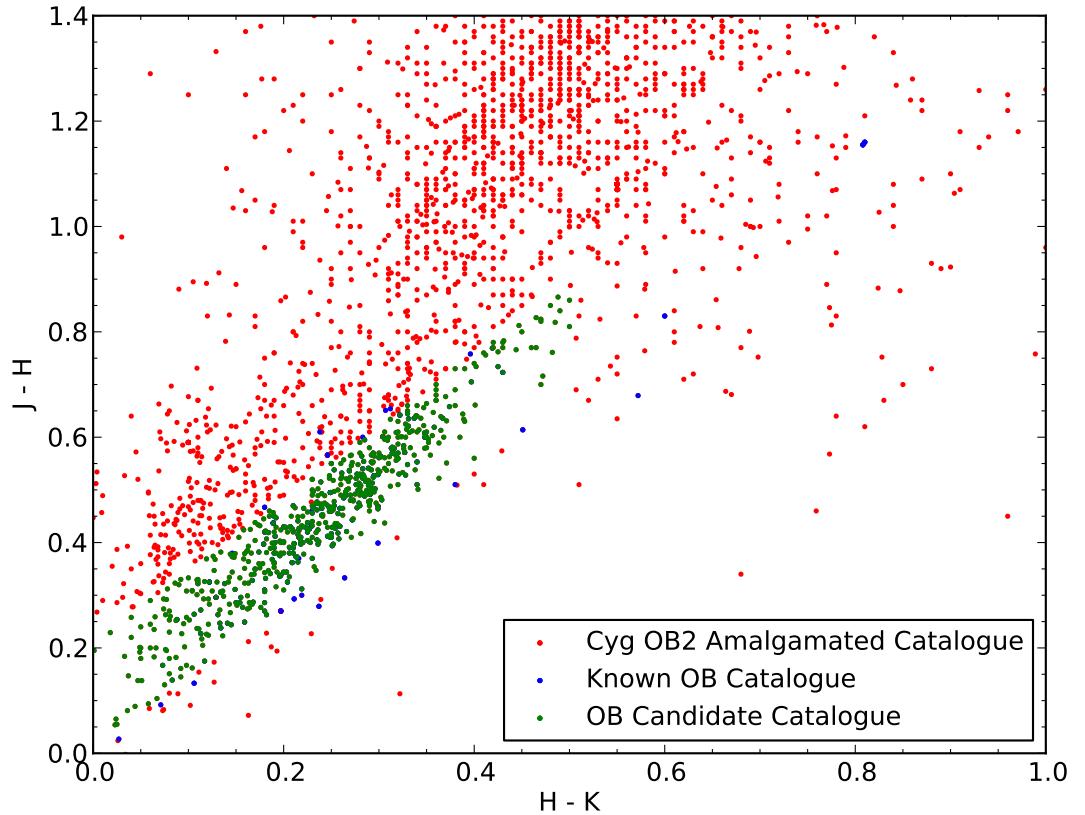


Figure 4.5: J-H versus H-K magnitude colour plot with the Cyg OB2 Amalgamated Catalogue (red), the known OB catalogue (blue) and the OB Candidate Catalogue (green).

excesses, and the 2MASS OB candidate catalogue. These are valuable resources for a large survey such as COBRaS, and are all contained within the Jython catalogue amalgamation script which can be easily updated with new catalogues as and when they are published in the literature.

4.2 Source Detection Algorithms

Continuing the theme of automated procedures for COBRaS (and indeed other radio surveys with exceptional data volumes), the detection of sources in the radio maps begins the final stage of the pipeline. The ideal source finder is one which is complete, finding all the sources in the image, and one with a low false-positive ratio i.e. finding only real sources. Due to the importance of source detection, a number of algorithms have been developed over the years using a multitude of techniques.

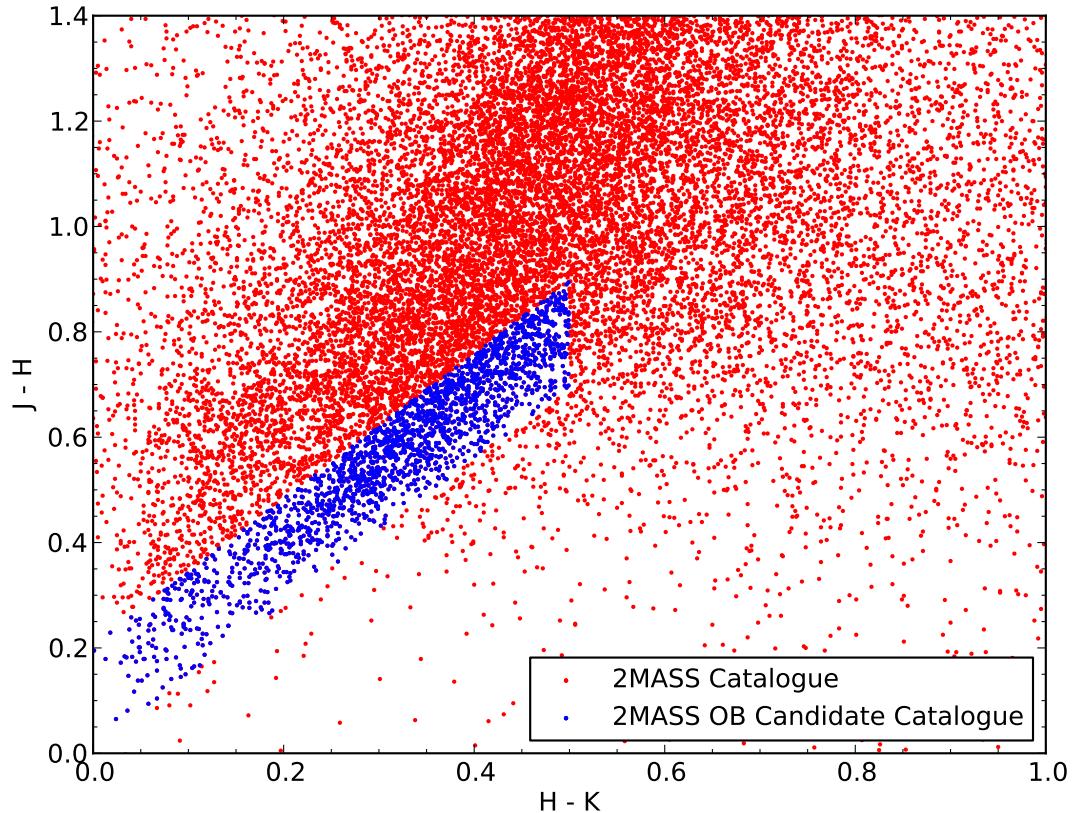


Figure 4.6: Same plot as Figure 4.4, J-H versus H-K magnitude colour plot with the 2MASS catalogue (red) and the 2MASS OB candidate catalogue (blue). The same regression line and loci from the OB catalogue is used to determine the 2MASS OB candidate catalogue.

A recent review of source detection methods was conducted by Masias et al. (2012), with discussions on basic image transformation, Bayesian inference, matched filtering, multi-scale approaches and wavelet transformations. These cover a range of objectives, such as source detection, faint source detection, point source detection, extended source detection and a range of observing bands from X-ray to radio. Recent developments in radio source detection algorithms for the Australian SKA Pathfinder (ASKAP) Evolutionary Map of the Universe (EMU) project have produced two codes: AEGEAN (Hancock et al. 2012), and BLOBCAT (Hales et al. 2012).

AEGEAN is a radio point source detection algorithm which adopts the floodfill algorithm to detect sources and fit Gaussian point spread functions to the sources, otherwise known as island of pixels. It assumes all sources are unresolved or point-like and convolves the image with a Lagrangian matrix to produce a covariance map. This breaks the islands

into ‘summits’ which signify the peak presence of a point source. Multiple summits within an island represent multiple blended point sources, and a Gaussian is fitted to each. The consequent flux extraction will account for multiple blended Gaussian sources.

BLOBCAT is a radio source detection algorithm which also adopts the floodfill algorithm to detect sources (Hales et al. 2012). It however does not assume the source is unresolved and the flux extraction performs both Gaussian and non-Gaussian routines to calculate the flux. For resolved sources, it performs a number of corrections for the peak surface brightness bias, where a Gaussian may over or under estimate the peak flux depending on the nature of the source (see Figure 3 of Hales et al. 2012), and the integrated surface brightness bias, where the floodfill algorithm does not fill the entire source volume before the cut off threshold.

Both algorithms use the same floodfill detection method (discussed below) and demonstrate a high completeness; 93.87% at 5σ , and reliability; 98.69% at 5σ (Hancock et al. 2012). Since COBRaS is also primarily a point source radio survey, this algorithm is desirable for its proven robustness and automated nature.

4.2.1 Floodfill Algorithm

The floodfill algorithm operates on an image with two inputs; the seed or detection threshold (T_s) and the flood threshold (T_f). If a pixel in the image is above the seed threshold T_s , then an island is grown around the seed pixel by testing whether adjacent pixels are above the flood threshold T_f . All pixels that fit the flood threshold criteria are added to the island and the next adjacent pixels are considered. This iterative procedure continues until all adjacent pixels to the island are below the flood threshold.

In the AEGEAN algorithm (Hancock et al. 2012) the seed threshold T_s is set to 5σ (5 times the noise level) and the flood threshold $T_f = 4\sigma$. This level of flooding is sufficient to characterise the islands and fit Gaussians to the components. BLOBCAT (Hales et al. 2012) is also seeded at $T_s = 5$, but the flood threshold $T_f = 2.6$, because the authors found that this enabled true source pixels to be flooded whilst avoiding flooding non-source pixels and perform a volume bias correction to account for any unfilled source pixels.

4.2.2 Modifications to the Floodfill Algorithm

It is evident that the floodfill algorithm is an effective source detection method. However, these algorithms have only been tested for 5σ sources, embedded in Gaussian noise. To

arbitrarily lower the seed and flood thresholds further without considering preventative ways to mitigate false detections will decrease the reliability.

Other source detection algorithms such as SExtractor (Bertin and Arnouts 1996) have a minimum island size (chosen by the user) to mitigate false detections. This is a useful variable to employ, but it is not automated. However, due to the nature of radio interferometric images during deconvolution, the image is convolved with a Gaussian PSF representation of the synthesised beam. Therefore, an infinitesimally small point source will assume the shape and size of the CLEANed beam, which in turn is comprised of a finite number of pixels. No real detectable source can be smaller than the beam with the addition of noise and therefore no island can have fewer numbers of pixels than the number of pixels comprising the beam with noise. Exploiting this condition inherent to radio interferometric images to the source detection mitigates the possibility of false detections, even in very noisy images, or sources $< 5\sigma$. AEGEAN does consider the synthesised beam when fitting Gaussian components to the summits in the curvature maps for flux extraction, but not during the island flooding, i.e. source detection stage.

For the COBRaS source detection pipeline, it is necessary to push the limit of source detection below 5σ in order to allow the detection of as many OB stellar winds as possible and to achieve a high completeness of OB stars in Cyg OB2. This condition is an important modification to the current floodfill algorithm utilised by AEGEAN and BLOBCAT.

The other modification to the COBRaS source detection is another additional run of the floodfill algorithm with a lower flood threshold to enable more source pixels to be added to the island. This is because the COBRaS algorithm will either fit Gaussians via the AIPS task JMFIT (least squares method to fit a Gaussian given a small search region) or calculate the flux directly from the island pixels (see Section 4.3.1). If the fluxes are derived from JMFIT, then only one run of the floodfill algorithm is required to obtain the max pixel position and pixel flux inputs for JMFIT.

This method has the benefits of the restrictive conditions above by having a stricter flood threshold when identifying sources, and then relaxing this flood threshold when the source islands are known. It is necessary to flood the island as much as possible for accurate flux determination from individual island pixels.

Another important aspect to the COBRaS source detection pipeline is that all the pixels within an island are assumed to be from a single source. This is because the flux determination is conducted on a pixel-by-pixel basis, summed and corrected for the

background noise to produce a flux independent from Gaussian fits. Certain expected sources in the COBRaS field are intrinsically non-Gaussian in shape. For example, binary interaction regions assume a crescent shape as the more powerful wind of one of the components of the binary distorts the adiabatic shock region. The absolute shape of the region depends on the orbital positions of the binary components. To achieve an accurate flux measurement, non-Gaussian determinations are required. Previously, the fluxes from these regions were determined by Gaussian fits (typically JMFIT, e.g. Watson et al. 2002) with manual TVSTAT checks (Dougherty et al. 2005), and as stated above is an option within the pipeline.

4.3 Flux Extraction in the COBRaS Pipeline

The following sections will analyse the pixel-by-pixel flux extraction method with that of the Gaussian JMFIT on a variety of simulated point sources with a range of signal to noise ratios (SNR) and the source positions from the weighted mean of the pixels in an island

4.3.1 Pixel-by-pixel Flux Extraction

After the source detection, a list of island pixels is examined assuming that each island is a single source. If the flux determination is chosen to be calculated from fitting Gaussians (via JMFIT) then information on the maximum island pixel position and flux and search box parameters (typically a few percent of the image size centred on the max pixel) are given to JMFIT to help with the least squares fitting routine. The peak fluxes and/ or integrated fluxes and the source position are returned with errors.

If flux analysis is conducted pixel-by-pixel (PP), then the beam area Ω_{area} is calculated using information on the resolving beam from the image header as

$$\Omega_{\text{area}} = \frac{\Theta_{\text{maj}} \Theta_{\text{min}} \pi}{4 \ln 2}, \quad (4.1)$$

where Θ_{maj} and Θ_{min} are the major and minor full width half maximum (FWHM) of the resolving beam. Dividing the beam area by the pixel area of the beam A produces the size of the beam Ω_{size} in pixels

$$\Omega_{\text{size}} = \frac{\Omega_{\text{area}}}{A}, \quad (4.2)$$

which is used to convert the fluxes into Janskys rather than Jansky Beam⁻¹. Next, the flux F of N pixels at positions i and j in the island are summed and divided by Ω_{size} to give the uncorrected flux F_u of the source:

$$F_u = \frac{\sum_{i=1}^N \sum_{j=1}^N F_{ij}}{\Omega_{\text{size}}}. \quad (4.3)$$

The global background noise level (B) of the image is calculated from the histogram of the pixels within AIPS (IMEAN function) and is multiplied by the number of pixels in the island N and divided by the beam size. This is subtracted from the uncorrected flux value determined in Equation 4.3 to give the corrected flux F :

$$F = F_u - \frac{B N}{\Omega_{\text{size}}}. \quad (4.4)$$

If the background noise level is chosen to be determined locally, a small annulus is created around the island to enclose the local background pixels adjacent to the island. To avoid contamination from other local sources, only non-island pixels are considered. The RMS is calculated using median statistics within the annulus to prevent bias from spurious pixels. This is more a consideration for real images with artefacts in the image rather than the simulations presented here.

The error on the flux (σ_F) is given by

$$\sigma_F = \frac{\text{RMS}}{\sqrt{N/\Omega_{\text{size}}}}, \quad (4.5)$$

where the RMS is calculated from the histogram of the pixels within AIPS (STD function) and N is the number of pixels in the island. Equations 4.4 and 4.5 give the integrated flux and error on the flux from the island of pixels.

4.3.2 Source Position Determination

The source positions can be determined from JMFIT, or from the PP analysis. From the pixel analysis, the source position can be extracted from the pixel with the maximum flux value within the island, or from the weighted mean of the pixel fluxes. Here the flux for each pixel acts as a weight for that pixel position, i.e. favouring the peak pixel more than the edge of the island pixels where the flux will be close to the noise level. Employing

this method ensures the most accurate measurement for source position for non-Gaussian sources as well as unresolved sources.

The weighted mean pixel position of the source (x_w, y_w) is given by

$$(x_w, y_w) = \frac{\sum_{i,j=1}^N x_i F(x_i, y_j), y_j F(x_i, y_j)}{\sum_{i,j=1}^N F(x_i, y_j)}, \quad (4.6)$$

where N is the number of pixels in the island at positions i and j , x_i and y_j are the pixel values at i and j respectively, and $F(x_i, y_j)$ is the flux at the corresponding pixels. The error on the weighted position (σ_x, σ_y) is given by

$$(\sigma_x, \sigma_y) = \sqrt{\sum_{i,j=1}^N \left(\frac{x_i, y_j}{F(x_i, y_j)} \right)^2}. \quad (4.7)$$

These pixel values for the position can be converted into RA and DEC positions simply by calculating the pixel offset position from the centre of the image (information from the image header) and multiplying by the RA and DEC increment for each pixel. For the errors, (σ_x, σ_y) are in the form of a fraction of a pixel, and can be converted by multiplying by the RA and DEC increment for a pixel.

4.3.3 Source Detection Performance

To test the performance of the source detection with the PP method, a range of Monte Carlo (MC) simulations were conducted on synthesised datasets. The obvious benefit of simulations is the ability to control all inputs and to obtain reliable performance statistics. The first datasets consist of point sources generated by the AIPS task UVCON randomly distributed over multiple images of 2048×2048 pixels. These are then convolved with IMAGR to create the sources in a map. Ten sources are created per image and a random variable drawn from a Gaussian distribution is added to each pixel to simulate Gaussian noise. This random variable has a mean equal to zero and a varying standard deviation to produce maps of differing signal to noise ratios (SNRs). For each SNR bin, a total of 1000 sources are generated and the median of the fluxes and the first and third quartiles for the errors are determined, in line with the similar study by Hales et al. (2012). Another dataset consisting of resolved Gaussian sources is also generated in the same manner.

The UVCON reconstructions do not contain any intentional simulated errors such as instrumental system errors or phase errors. Only idealistic Gaussian noise is added after the imaging process. However, as in real images it is a real possibility that residual phase errors from the CLEAN process may exist. The effect this will have on the source extraction is inaccurate source positional information and even decoherence of the source, particularly for weaker sources.

Another effect to consider during the simulations is ‘clean bias’. Clean bias is the effect of poor $u - v$ coverage, producing noisy sidelobes in the synthesised beam. This has the effect of creating clean components during CLEAN from the sidelobes which will reduce the flux at the true source position (Condon et al. 1998). Resolved sources are particularly prone to this where the an extended source may appear brighter at these sidelobes than at the true peak, and flux from the true position is shifted onto these sidelobes. However, for these simulations and the expected COBRaS field, no resolved sources are expected to have dimensions several times that of the synthesised beamwidth.

Clean bias has obvious effects for weak sources where the peak flux is already close to the RMS noise level. There are ways to reduce clean bias which include: using windows around the sources to restrict the area of cleaning. To use a cleaner synthesised beam, uniform weighting creates a narrow central profile but noisy sidelobes whereas natural weighting generally has lower sidelobe levels. Do not over clean as this will increase the chance of cleaning sidelobes and therefore reducing the fluxes at the true source positions as described above. Whilst windows are not used during the simulations, the sources are not ‘weak’ during the CLEAN part of the image reconstruction because noise is added after the deconvolution.

Unlike previous studies with the floodfill algorithm (Hales et al. 2012; Hancock et al. 2012), source detection and flux extraction performances are tested for sources with SNRs less than 5 with the total tested subset of SNRs = {2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100}. Therefore the two variable inputs for source detection, the seed threshold T_s and flood threshold T_f , differ from the values used in those studies. A range of values was empirically tested for SNRs less than 5 and for detecting sources with $\text{SNR} \geq 2$, $T_s = 1.9$ RMS and $T_f = 1/4 T_s$ for the first run of the floodfill algorithm and $T_f = 0.5$ RMS for the second run (T_s remains the same). For SNR greater than 10, the seed threshold performed better (fewer false positives from artefacts in the images) with $T_s = 5$ RMS for both runs and the flood threshold remaining the same as before. Note that overflooding is not a

problem for the PP flux extraction, because any background pixel added to the island will be normalised by the background subtraction routine, resulting in a negligible positive or negative contribution to the integrated flux value. However, a suitable cut-off is required to avoid endlessly adding pixels to the island and stop the program getting stuck in a loop.

If two (or more) sources are generated at the same location (i.e. the islands of the sources join), they are treated by the algorithm as a single source. For these simulations there is no penalty for the inability to distinguish blended sources, as this is a key assumption that islands are assumed as a single source. However, the results below do consider the ability to determine the flux contribution from all of the blended sources as part of the detection.

Point Sources

The detection of point sources in the simulations for the PP method with global and local noise estimations and JMFIT is shown in Figure 4.7. The PP method shows similar results for both noise estimations, with around 90% successful detection at $\text{SNR} = 2$, and 100% detection at $\text{SNR} = 3$ to 50. At $\text{SNR} = 100$, false positives are produced for both noise regimes. This is likely to be a result of artefacts from the image creation by IMAGR, where the deconvolution via CLEAN scatters a small fraction of the flux in the image. It becomes apparent at this SNR because of the low noise level, however there are only 2 instances for the global and 8 for the local noise estimations in 1000 simulated sources.

JMFIT shows clear false positives at $\text{SNR} = 2$ and 3, where the fitting of multiple Gaussians to single sources seemingly produces the best χ^2 minimisation. At higher SNRs, JMFIT has a detection level $\gtrsim 99.0\%$, despite the floodfill algorithm (as demonstrated by the PP method results) feeding 100% correct peak positions and peak fluxes to JMFIT for these SNRs. At $\text{SNR} = 30$, JMFIT has a 100% detection performance.

Resolved Sources

The results of the resolved source detections with the PP method (global and local noise estimations) and JMFIT is shown in Figure 4.8. The PP methods achieve similar performances for resolved sources as for the point sources, with the biggest difference occurring at low SNRs. At $\text{SNR} = 2$, only a third of total sources are detected and for $\text{SNR} = 3$, the performance is $\sim 99.0\%$. Only three false positives for $\text{SNR} = 100$ are detected for the global noise estimation and one for the local noise. Again this is likely to be resultant

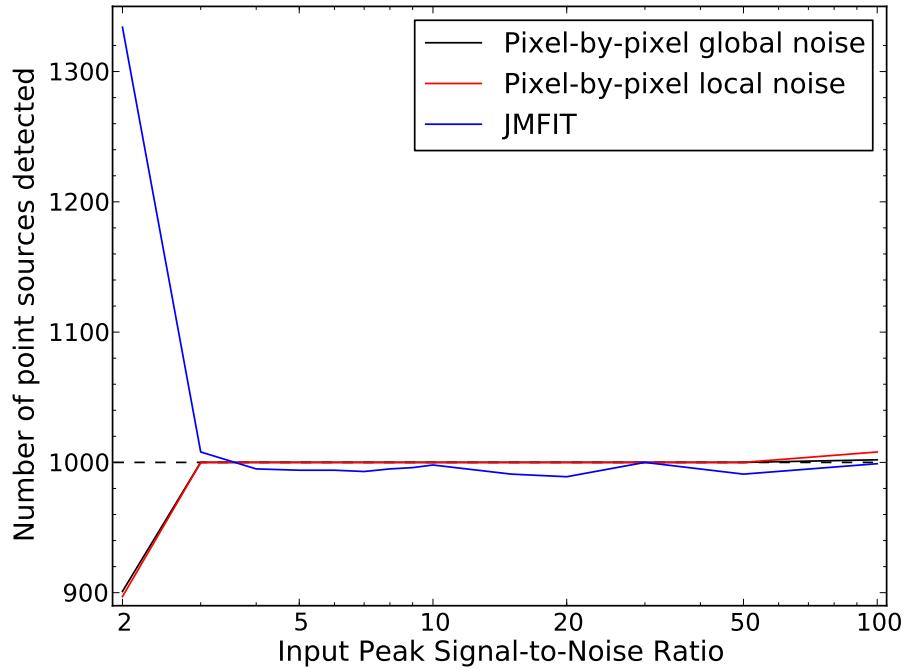


Figure 4.7: Source detection results for point sources. PP method with global noise in black, local noise in red and JMFIT in blue. Target number of sources per SNR bin is 1000.

from artefacts from IMAGR scattering flux in the image (and is seen later in Figure 4.13).

JMFIT also struggles at $\text{SNR} \leq 5$, with false positives arising from the χ^2 minimisation procedure fitting multiple Gaussian components. Furthermore, with only ~ 360 real sources detected by the PP method for $\text{SNR} = 2$ passed on to JMFIT, which finds > 1400 sources, results in JMFIT on average fitting four Gaussians per source in an attempt to recover all of the available flux. At $\text{SNRs} > 5$, JMFIT has a source detection performance $\sim 99.0\%$.

Summary

The source detection results for the PP method and JMFIT are summarised in Table 4.2 for all SNRs.

The precision of these results is determined by the number of sources simulated for each SNR, e.g. for 1000 sources, a precision of 0.1% can be obtained. Despite a number of 100% detection performances above, it is never certain that the same precision will be achieved with subsequent runs. In this particular scenario, the performance is deemed to be $> 99.9\%$, or accurate to within one decimal place.

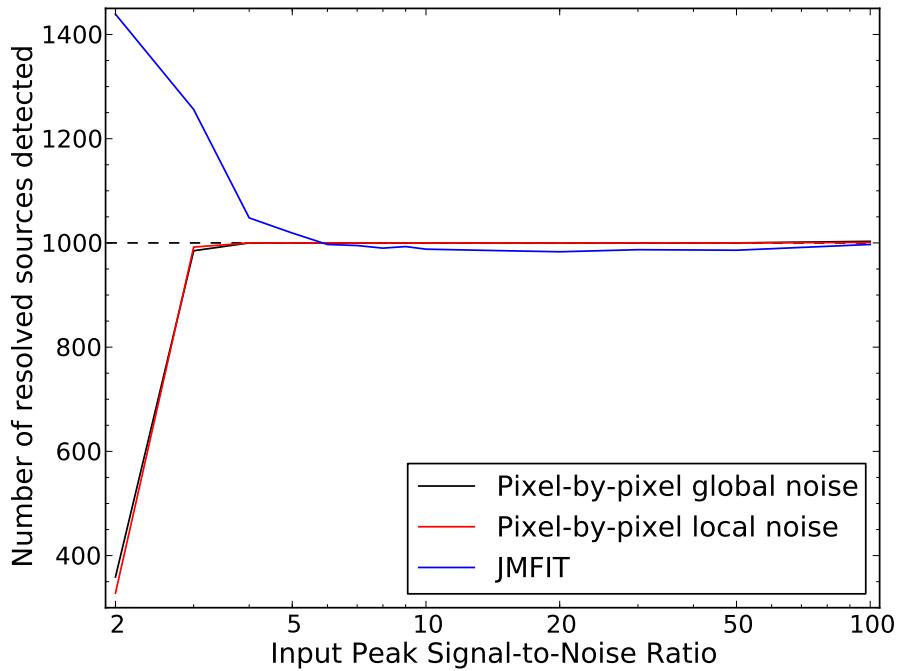


Figure 4.8: Source detection results for resolved sources. PP method with global noise in black, local noise in red and JMFIT in blue. Target number of sources per SNR bin is 1000.

Table 4.2: Table of Source Detection Results from Simulations

Signal-to-noise Ratio	Point Sources			Resolved Sources		
	Global*	Local*	JMFIT	Global*	Local*	JMFIT
2	901	897	1334	359	328	1439
3	1000	1000	1008	985	992	1256
4	1000	1000	995	1000	1000	1048
5	1000	1000	994	1000	1000	1019
6	1000	1000	994	1000	1000	997
7	1000	1000	993	1000	1000	995
8	1000	1000	995	1000	1000	990
9	1000	1000	996	1000	1000	993
10	1000	1000	998	1000	1000	988
15	1000	1000	991	1000	1000	985
20	1000	1000	989	1000	1000	983
30	1000	1000	1000	1000	1000	987
50	1000	1000	991	1000	1000	986
100	1002	1008	999	1003	1001	997

* The noise estimation regimes for the PP method.

4.3.4 Flux Extraction Performance

The flux extraction performances for point sources and resolved sources generated as described previously, are now presented for the two PP methods (global and local noise)

and JMFIT for the $\text{SNR} = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100\}$.

Point Sources

Figure 4.9 shows 14 point sources embedded with varying levels of random Gaussian noise to replicate the subset of tested SNRs.

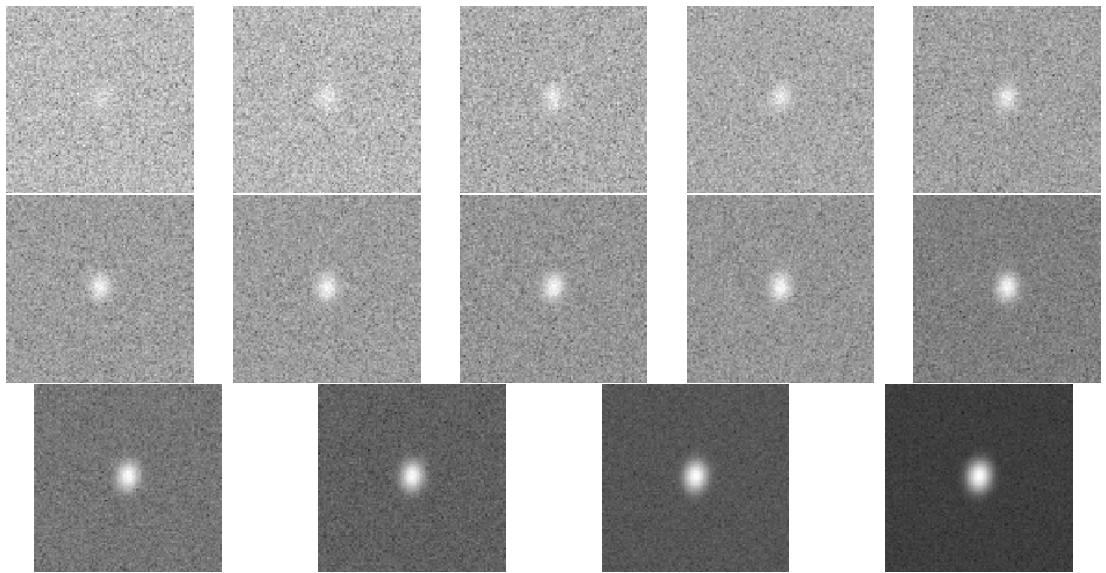


Figure 4.9: 14 point sources created with UVCON with added Gaussian noise, $\mu = 0$, $\sigma = 1 / \text{SNR}$. From the top going left to right, $\text{SNR} = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100\}$.

The PP flux extraction method is tested along side the common Gaussian fitting procedure JMFIT within AIPS. The input parameters for JMFIT are left unconstrained to ensure an optimal fit. Both global and local background noise determinations for the PP flux extraction are presented here to determine the performance of each method. The results for the point source (unresolved) integrated flux extraction can be seen in Figure 4.10, for the PP flux extraction with global (a) and local (b) noise determination and JMFIT (both figures).

The performance of both methods are generally within the dotted curves, defined by the random errors, represented by the MAD of the flux density sample from each SNR bin. JMFIT only overestimates the flux at $\text{SNR} = 2$, with large errors, whilst underestimating the flux at all other SNRs. The PP method always underestimates the flux with global background subtraction, but is more accurate than JMFIT for $\text{SNR} \geq 5$. For $\text{SNR} < 5$, the PP method performance suffers as a result of an intrinsic property of the method, in that it is based on a pure PP analysis to determine the noise level and, furthermore, the flux.

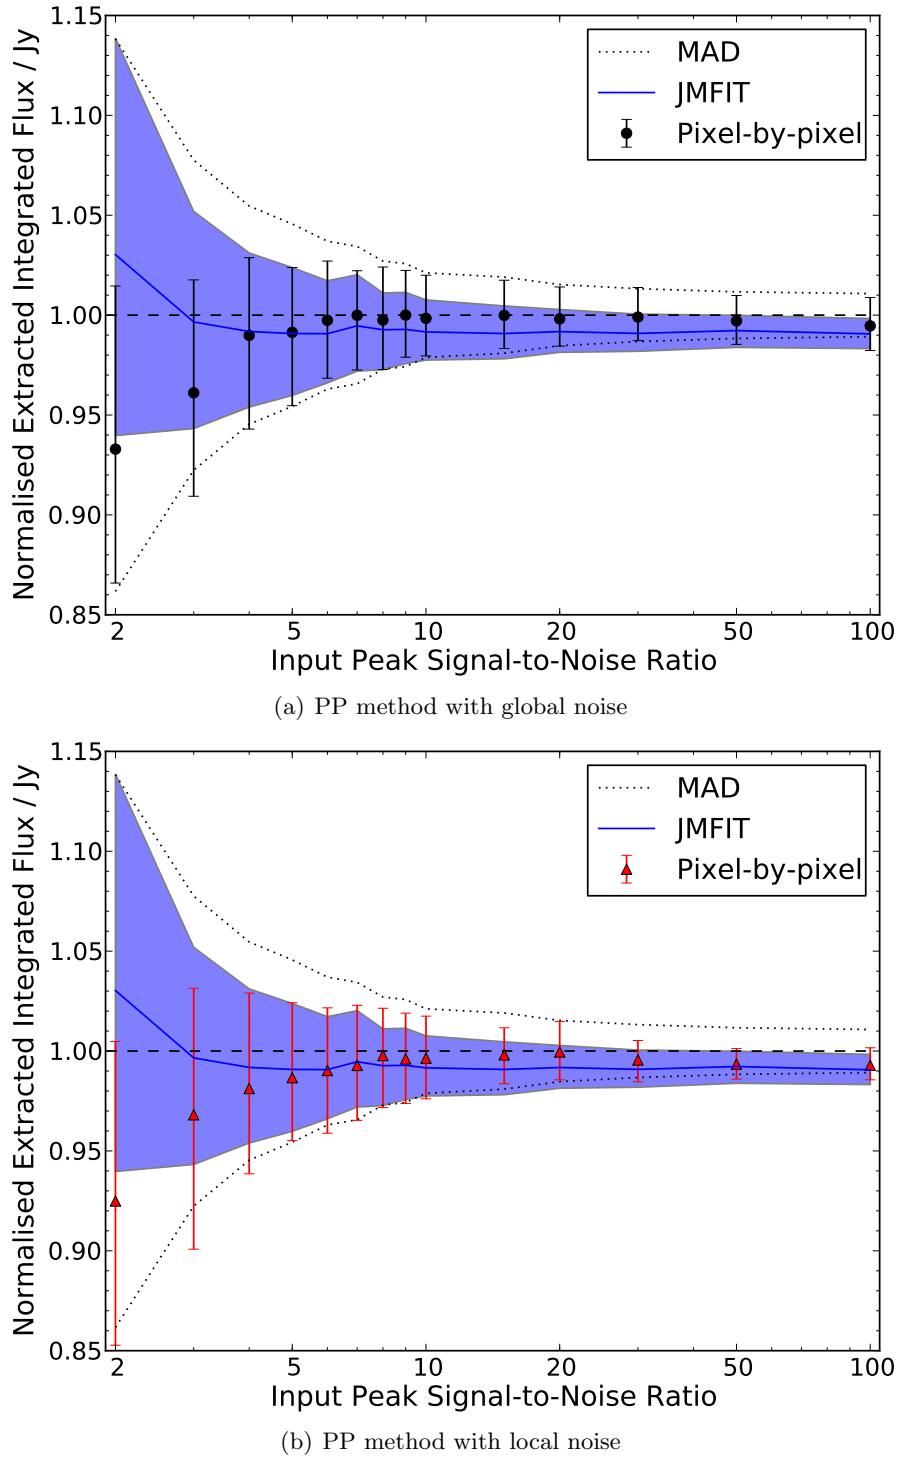


Figure 4.10: Performance of the pixel-by-pixel (black points) and JMFIT (blue line) integrated flux extractions of point sources. The integrated fluxes are summarised by the median of the 1000 sources and the first and third quartiles for the errors (error bars and shading). For reference, expected random errors are indicated by the MAD (dotted lines).

The greater the noise and thus the uncertainty, the worse the performance. One merit of this method is that the PP method will rarely overestimate the flux, even in the situation where too many pixels are added to the island than exists in the real source structure. This is because each additional pixel will be normalised by the background level during noise subtraction.

The reason JMFIT does not suffer from the same issue is because it fits Gaussians via a least squares algorithm to the pixels within a designated region around the source. The flux value with the smallest χ^2 is the preferred solution. However, a problem with fitting Gaussians to sources with low SNR is that multiple Gaussian solutions may be found, indicating the presence of multiple false sources. This has been the case during these simulations for $\text{SNR} = 2$ and 3 with a number of false positive sources (> 300 for $\text{SNR} = 2$, and 8 for $\text{SNR} = 3$). Therefore any individual fluxes determined for these SNRs should be treated with caution, because the median statistics mask a number of incorrect solutions.

The local noise fluxes have similar behaviour to the global noise fluxes over all SNRs. This is expected because of the uniform Gaussian noise determination, resulting in similar local and global noise properties. For a non-uniform noise distribution (which can be expected in real maps), the local noise estimates may be more accurate than globally determined noise estimates.

Results from JMFIT are dependent on the initial starting values given to the task. From the PP analysis of point sources, the positional accuracy of source detection is $\sim 90\%$ which drops to $\sim 30\%$ for resolved sources. It is therefore expected that most positions of the point sources passed onto the least squares algorithm used by JMFIT are valid, but the same statement is not true for resolved sources.

Because median statistics are considered for the demonstration of source detection for each SNR bin, the incorrect characterisation of islands by JMFIT for point sources is determined more by the inability of the task to correctly fit the islands, as opposed to incorrect starting values for the least square algorithm giving by the source detection. However, for resolved sources this statement does not hold as only $\sim 30\%$ of the total source positions were correctly given to JMFIT.

Resolved Sources

The flux extraction methods are also tested on resolved Gaussian sources generated by UVCON in the same manner as for the point sources. The sources in this simulation have a generic structure which is highly resolved in one axis. Figure 4.11 shows 14 resolved sources embedded with varying levels of random Gaussian noise to replicate the subset of tested SNRs.

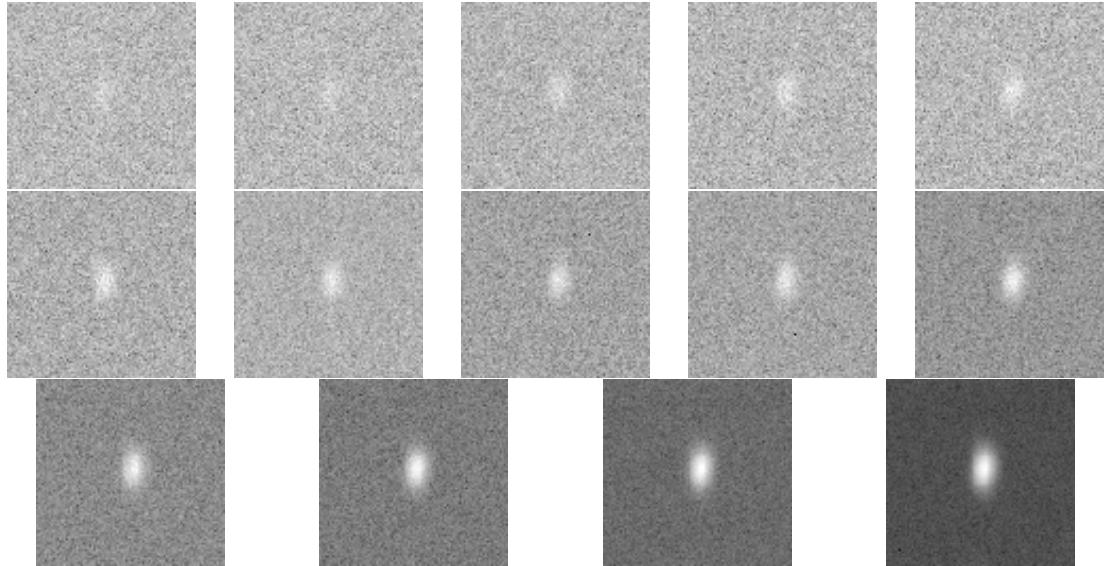


Figure 4.11: 14 resolved sources created with UVCON with added Gaussian noise, $\mu = 0$, $\sigma = 1 / SNR$. From the top going left to right, $SNR = \{2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 30, 50, 100\}$.

Both global and local background noise determinations for the PP flux extraction and JMFIT flux extraction are again presented here to determine the performance of each method on resolved sources. The results for the resolved source integrated flux extraction can be seen in Figure 4.12, for the PP flux extraction with global (a) and local (b) noise determination and JMFIT (both figures).

The apparent systematic offset for both PP methods and JMFIT arises from IMAGR, with CLEAN's attempt at deconvolving resolved sources. During this process some of the flux is scattered throughout the image due to CLEAN's inability to correctly deconvolve a resolved source. This affects all source detection and flux extraction methods and results in the main source component having a lower input flux than the assumed 1.00 Jy input flux which was entered into UVCON. To demonstrate this, Figure 4.13 is the resultant field of a resolved source deconvolved with IMAGR with no added Gaussian noise. Obvious

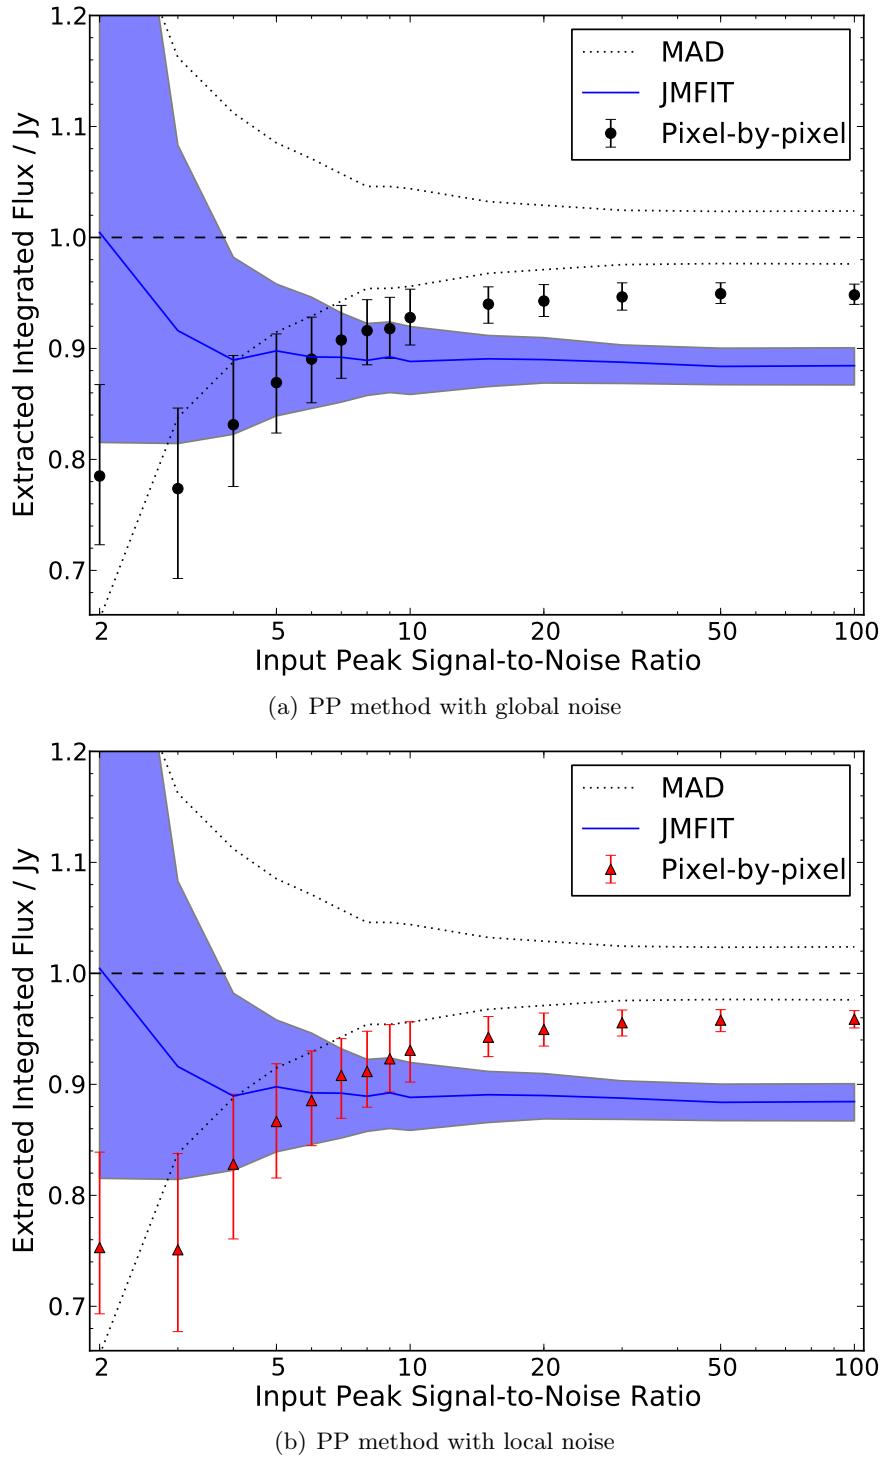


Figure 4.12: Performance of the pixel-by-pixel (black points) and JMFIT (blue line) integrated flux extractions of resolved sources. The integrated fluxes are summarised by the median of the 1000 sources and the first and third quartiles for the errors (error bars and shading). For reference, expected random errors are indicated by the MAD (dotted lines).

artefacts appear in the field, meaning the main source component does not contain the total 1.00 Jy input flux from UVCON. To determine the true flux of the main component, a box is drawn around the source and every pixel within is summed using the AIPS task IMEAN. To reduce the effects of random errors, 1000 input resolved sources with no noise are generated and the simulations showed the median flux of these 1000 resolved sources to be 0.94028 Jy with the median absolute deviation (MAD) equal to 0.00915.

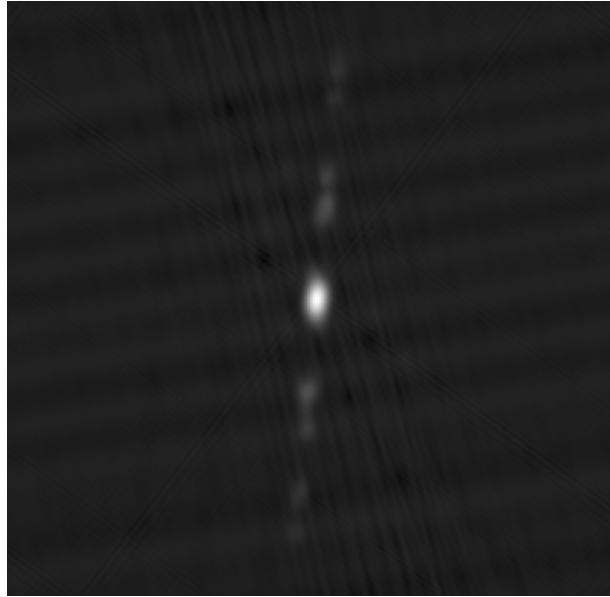


Figure 4.13: Image of a UVCON and IMAGR simulated resolved source with no noise. The main source component is at the centre of the image with a fraction of the flux scattered throughout the image.

Redoing Figure 4.12 with the true input flux, reveals the true performance of these algorithms without the CLEAN deconvolution bias. Figure 4.14 shows integrated flux extraction of a resolved source for the PP method using global noise (a) and local noise (b) estimations, and JMFIT (both figures). Note that the y-axis is no longer the normalised integrated flux (normalised to 1.00 Jy), but the actual integrated flux, with the horizontal dashed line representing the median of 1000 simulated flux-corrected resolved sources. The MAD (horizontal dash-dot lines) represents the uncertainty on this median input value.

The true input fluxes in Figure 4.14, show an improved performance for all methods compared to Figure 4.12. The PP method once again suffers at $\text{SNR} < 5$ due to the intrinsic noise problems discussed earlier. At around $\text{SNR} = 7$ the flux values are correct within errors, improving with higher SNRs. Although the fluxes appear to be slightly overestimated at higher SNRs, the majority of the values are within the third quartile

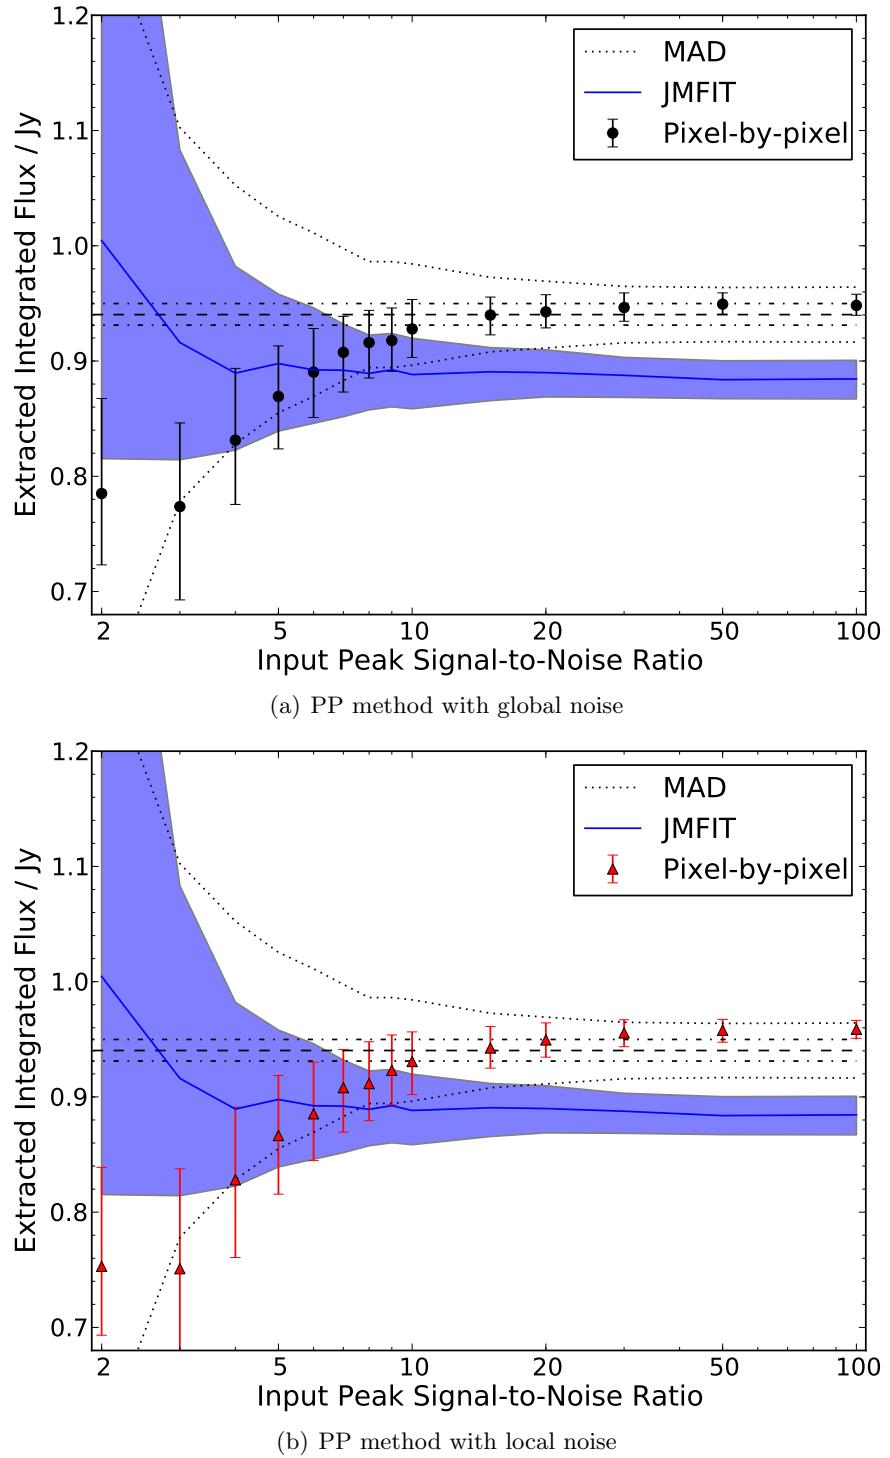


Figure 4.14: Performance of the pixel-by-pixel (black and red symbols) and JMFIT (blue line) integrated flux extractions of flux-corrected resolved sources. The integrated fluxes are summarised by the median of the 1000 sources and the first and third quartiles for the errors (error bars and shading). For reference, expected random errors are indicated by the MAD (dotted lines). The methods are compared to the median of the true resolved input fluxes from 1000 simulations, shown by the dashed line, with the MAD of the correct flux input given by dash-dot lines.

limit, and therefore a result of the nature of the flux-correction implemented above. This can be interpreted as being a function of random error associated with the flux-correction rather than an actual overestimate of the flux by the PP method.

The high SNR resolved source fluxes determined from local noise estimation are overestimated compared to the corrected input flux. This may result from the localised noise determination containing negative bowls (regions of negative flux) around the main source component generated during deconvolution (an effect which can arise from incorrect amplitude calibration as well as incomplete apertures in real sources). This could have the effect of increasing the resultant flux by subtracting less noise. This would explain why the global noise estimates are within the corrected input flux but the local noise estimates are not. Despite this, the flux values are within that expected from random errors given by the MAD.

The performance of JMFIT appears to suffer from the resolved nature of the sources even after the flux-correction. JMFIT consistently underestimates the flux by $\sim 4\%$ at higher SNRs and only correctly estimates the flux at $\text{SNR} < 7$ due to the large associated errors. Furthermore, it was shown earlier that JMFIT produces a number of false positives (see Section 4.3.3) for $\text{SNR} \leq 5$, and therefore individual fluxes for these SNRs should be treated with caution.

4.3.5 Source Position Performance

Similar MC simulations also test the performance of the weighted mean position determination from the PP method and JMFIT's position determination of the sources. The pixel offsets are calculated for a smaller subset sample of 100 sources³ and the median taken to represent that SNR bin.

Point Sources

For the point sources, Figure 4.15 shows the PP method with global and local background subtraction (black and red symbols) and JMFIT (blue lines) positional offsets in pixels.

The positional offsets determined by the PP methods (for both global (a) and local (b) noise estimates) have a sub-pixel positional accuracy for all SNRs. The accuracy increases with higher SNR, and starts to plateau at ~ 0.05 pixel offset. Any differences between the two different noise estimations are within the associated errors. The offsets from JMFIT

³Due to time restrictions, only 100 simulations were performed for each SNR and algorithm.

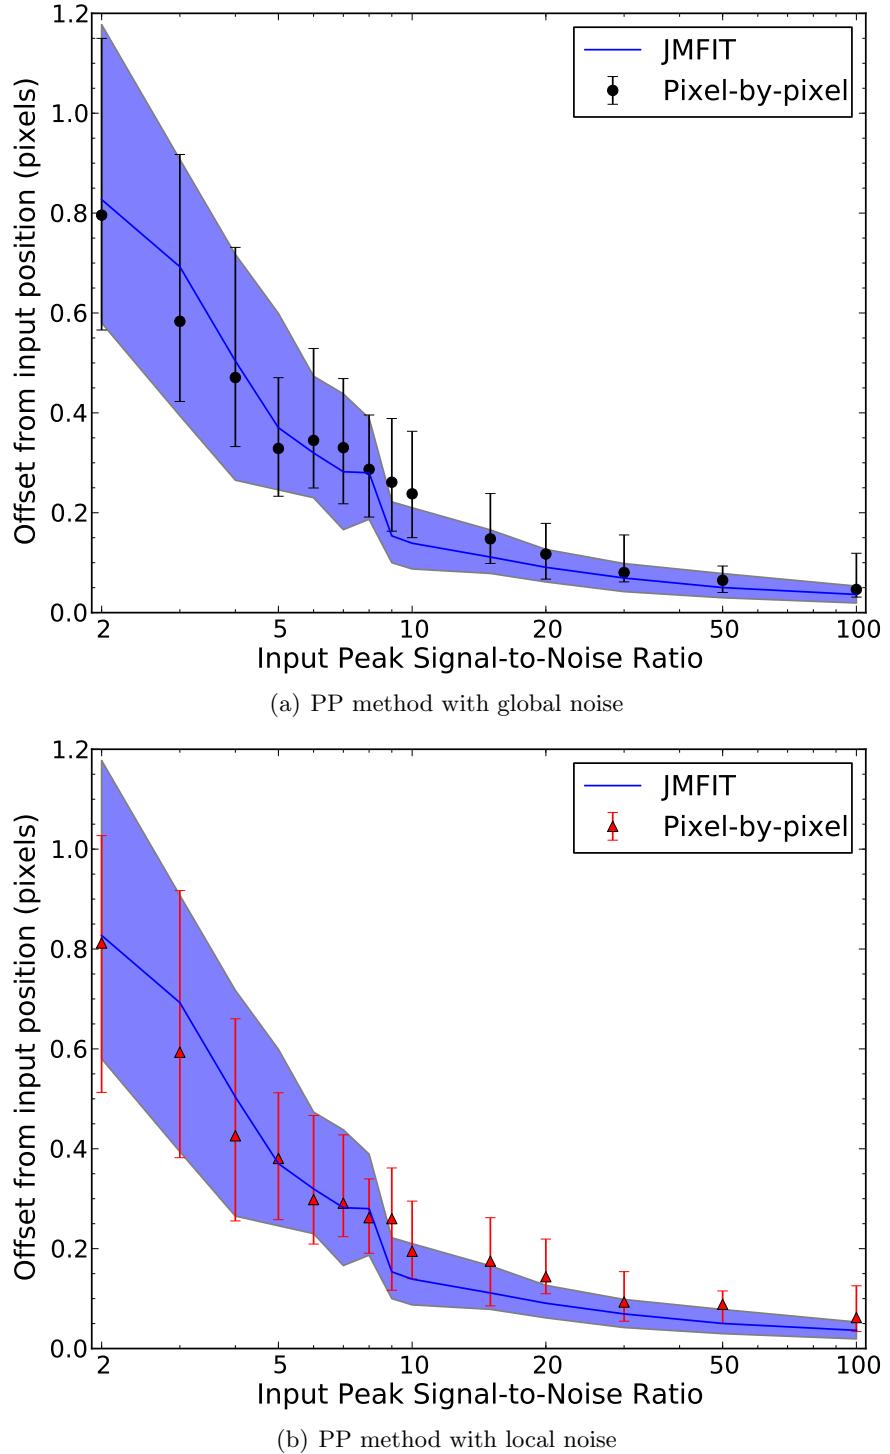


Figure 4.15: Performance of PP methods (black and red symbols) and JMFIT (blue line) positional offsets for point sources. Sub-figure (a) shows the PP method with global noise estimation and JMFIT fit and sub-figure (b) shows the PP method with local noise estimation (red triangles). The positional offsets are summarised by the median of the 100 sources and the first and third quartiles for the errors (error bars and shading). For reference, expected random errors are indicated by the MAD (dotted lines).

follow a similar trend to the PP methods, with noticeable increase in performance at $\text{SNR} > 10$ which also converges around the same pixel offset. Any visual differences between the PP method and JMFIT are accounted for by the associated errors given by the first and third quartiles.

Resolved Sources

For the resolved sources, Figure 4.16 shows the PP method with global and local background subtraction (black and red symbols) and JMFIT (blue lines) positional offsets in pixels.

The positional offsets determined by the PP methods (for both global and local noise estimates) have a sub-pixel positional accuracy for all SNRs but $\text{SNR} = 2$ for both and $\text{SNR} = 3$ for local noise estimates. Again the accuracy increases with increasing SNR, and plateaus around 0.10 pixel offset. The offsets from JMFIT follow a similar trend to those from the PP method with a marginally better performance for higher SNRs, but consistent within errors.

It is obvious that the PP methods and JMFIT achieve higher accuracy for the point sources than for the resolved sources. For JMFIT, the lower peak flux of the source makes determining the centre of the Gaussian profile more difficult. However, the performance of both increases with higher source SNR, with comparable positional offsets for point sources and resolved sources at $\text{SNR} \geq 100$.

4.4 Source Classification via Cross Referencing and Bayes' Theorem

The final part of the COBRaS pipeline focuses on source classification and significance boosting for low SNR sources (< 5 SNR) utilising the amalgamated catalogues from Section (4.1). Firstly the outputs from the source detection and flux extraction are cross referenced with an input catalogue within a designated search radius (similar to the ‘tidying’ algorithm in the catalogue concatenation), to create a list of potential matching candidates for each source. This matching list is then explored in detail to reveal the probabilities that the sources and candidates are equivalent.

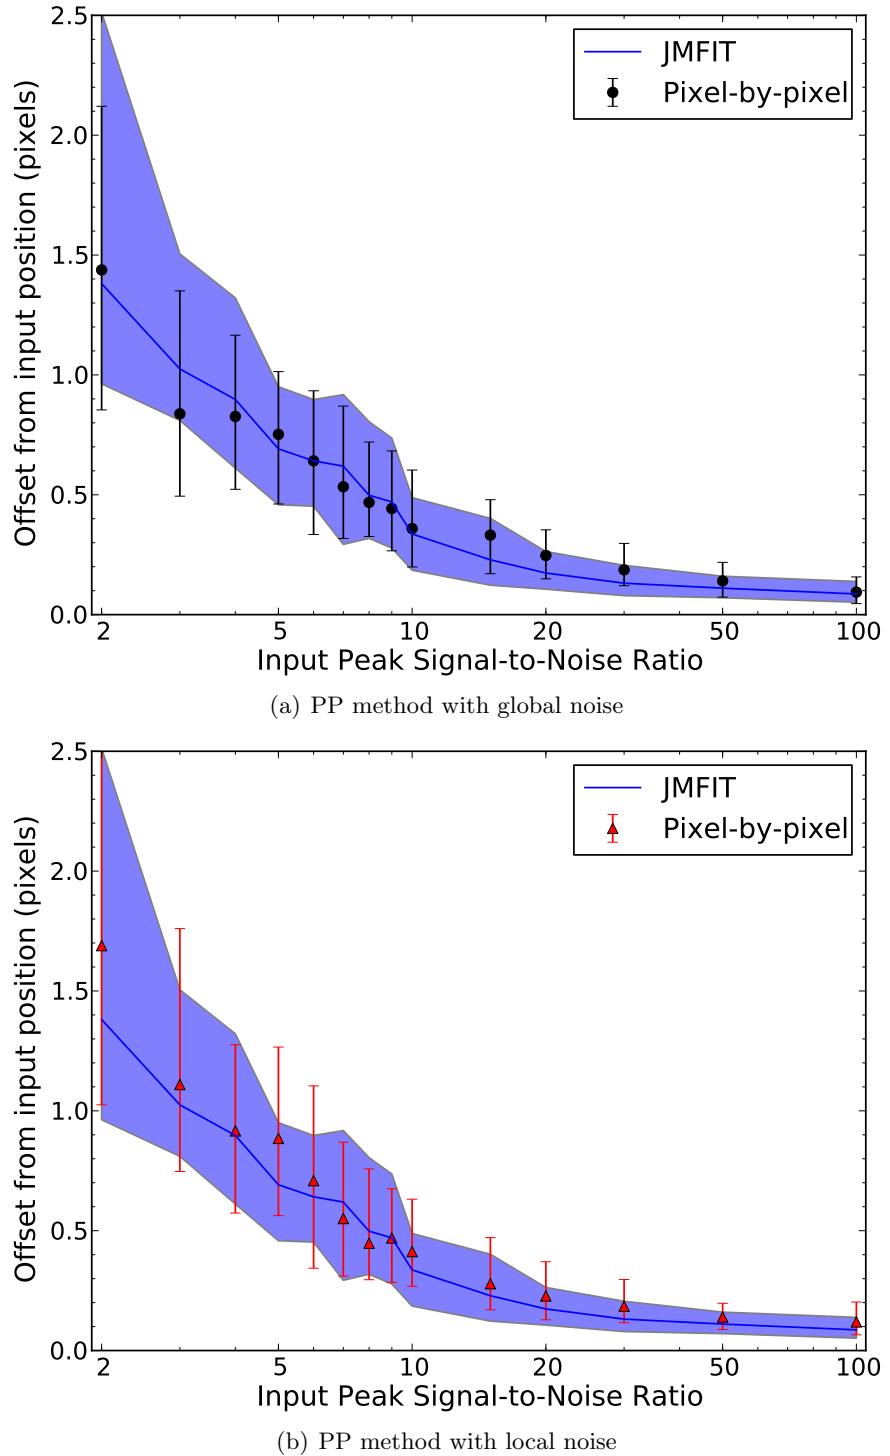


Figure 4.16: Performance of PP methods (black and red symbols) and JMFIT (blue line) positional offsets for resolved sources. Sub-figure (a) shows the PP method with global noise estimation and JMFIT fit and sub-figure (b) shows the PP method with local noise estimation (red triangles). The positional offsets are summarised by the median of the 100 sources and the first and third quartiles for the errors (error bars and shading). For reference, expected random errors are indicated by the MAD (dotted lines).

4.4.1 Cross Referencing Maps and Catalogues

The probability that the catalogue and map source are coincident can be determined by the product of the Gaussian distributions of each source. The majority of the OB stellar winds residing within the COBRaS maps are expected to be point sources, or can be described by Gaussians and therefore the positional probability can be well described by a Gaussian probability density function (PDF) with the position equal to the mean and error to the standard deviation of the distribution. The matching source is reasonably described by a Gaussian PDF, with the position equal to the mean of the PDF and the standard deviation equal to the matching radius. This produces a wide distribution, but still only significantly favours a source which coincides with the same position and insignificantly for a source on the periphery of the distribution.

Given the mean position of the source (x_{\circ} , y_{\circ}) and its associated errors σ_x and σ_y , the two dimensional (2D) PDF of the map source $f(x, y)$ at (x, y) , is given by

$$f(x, y) = \frac{1}{\sqrt{2\pi}\sigma_x\sigma_y} \exp \left[- \left(\frac{(x - x_{\circ})^2}{2\sigma_x^2} + \frac{(y - y_{\circ})^2}{2\sigma_y^2} \right) \right]. \quad (4.8)$$

The expression for the 2D PDF of the catalogue source ($g(i, y)$) is similar:

$$g(i, j) = \frac{1}{\sqrt{2\pi}\sigma_i\sigma_j} \exp \left[- \left(\frac{(i - i_{\circ})^2}{2\sigma_i^2} + \frac{(j - j_{\circ})^2}{2\sigma_j^2} \right) \right] \quad (4.9)$$

where the mean position given by the indices (i_{\circ}, j_{\circ}) , with errors σ_i and σ_j for the positions (i, j) .

The product of two Gaussian PDFs is also a Gaussian PDF, therefore the product of Equations 4.8 and 4.9 reveals the PDF that the source in the map and catalogue match $m(x, y, i, j)$:

$$m(x, y, i, j) = \frac{1}{2\pi\sigma_x\sigma_y\sigma_i\sigma_j} \times \exp \left[- \left(\frac{(x - x_{\circ})^2}{2\sigma_x^2} + \frac{(y - y_{\circ})^2}{2\sigma_y^2} + \frac{(i - i_{\circ})^2}{2\sigma_i^2} + \frac{(j - j_{\circ})^2}{2\sigma_j^2} \right) \right] \quad (4.10)$$

This expression indicates the probability that the two sources are coincident. Generally, the source has a narrow PDF profile whereas the catalogue source, with its larger uncertainty has a broad profile. Figure 4.17 shows the two 2D PDFs of a simulated

example of these two profiles.

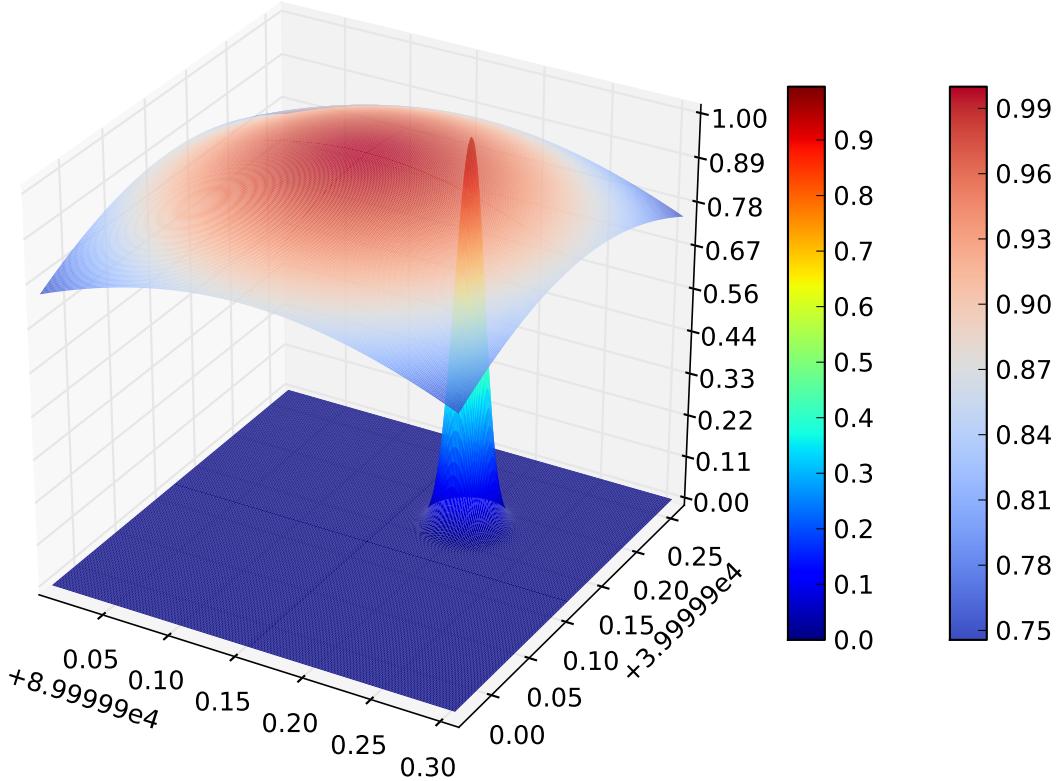


Figure 4.17: Plot of two 2D Gaussian PDFs. The broad Gaussian represents the catalogue source with its wide matching radius and the narrow Gaussian profile represents the source in the map. In this example there is a small offset between the source in the map and the source from the catalogue.

4.4.2 Significance Boosting with Bayes' Theorem

Bayesian statistics has been applied in detection algorithms (Hobson and McLachlan 2003; Savage and Oliver 2007; Carvalho et al. 2009) as an alternative method for source extraction utilising either Monte Carlo Markov Chains (MCMC) or the PowellSnakes algorithm for parameter estimation. One specific quality of Bayesian analysis is the effect of the *prior knowledge* on the posterior distribution, which can be the source positions determined from another survey or observing band (Savage and Oliver 2007). A brief introduction to Bayes' Theorem is given here, for a more formal and detailed description, the reader is encouraged to refer to Sivia and Skilling (2006).

Bayes' Theorem links the probability of hypothesis H being true, with that of the observed data D via three terms:

$$P(H | D) = \frac{P(D | H) \times P(H)}{P(D)} \quad (4.11)$$

where $P(H | D)$ is the posterior distribution (what we want to determine), $P(D | H)$ is the likelihood function of D given H, $P(H)$ is the prior knowledge of H, and $P(D)$ is the evidence term which acts as a normalisation constant and for many analysis problems can be omitted (Sivia and Skilling 2006). If the evidence term is removed then Equation 4.11 reduces to the proportionality:

$$P(H | D) \propto P(D | H) \times P(H). \quad (4.12)$$

The application of Bayes' theorem to cross correlate the source and catalogue positions using the source SNR is shown by the following. The posterior distribution $P(S | G_1, G_2)$ is described as the probability of there being a source S given the source in the map G_1 and the source in the catalogue G_2 . The likelihood function $P(S | G_1)$ is given as the probability of the source S given the source in the map G_1 and the prior knowledge $P(S, G_1 | G_2)$ is the probability of the source and source in the map given the source in the catalogue. Substituting these terms into Equation 4.12 gives

$$P(S | G_1, G_2) \propto P(S | G_1) \times P(S, G_1 | G_2). \quad (4.13)$$

The prior knowledge term in Equation 4.13 is the product of the probabilities of the source and catalogue positions, both described by a 2D Gaussian PDF, which is the same as Equation 4.10. The likelihood term in Equation 4.13 can be described by the probability of a tailed Gaussian distribution⁴ using the SNR as the standardised critical z value for calculating the normal distribution for $P(Z > z)$. The error function of a Gaussian distribution does precisely this, with an approximation from Bryc (2002) given by:

$$P(Z > z) \approx \frac{z^2 + 5.575192695z + 12.77436324}{\sqrt{2\pi}z^3 + 14.38718147z^2 + 31.53531977z + 2 \times 12.77436324} e^{-z^2/2} \quad (4.14)$$

⁴This is a reasonable assumption, despite the fact the pixel-by-pixel extraction does not assume Gaussianity of the source structure.

Multiplying the probability outputs from Equations 4.14 and 4.10 for the likelihood and prior knowledge terms respectively for Equation 4.13 does not increase the significance of the source. Instead the inverse of each probability is required, or the probability of there being no source $P(\text{no } S | G_1, G_2)$, with corresponding likelihood $P(\text{no } S | G_1)$ and prior probability $P(\text{no } S, G_1 | G_2)$. Each term is equivalent to 1 minus the corresponding term in Equation 4.13, resulting in the boosted significance $P(S | G_1, G_2)$:

$$\begin{aligned} P(\text{no } S | G_1, G_2) &\propto P(\text{no } S | G_1) \times P(\text{no } S, G_1 | G_2) \\ P(S | G_1, G_2) &\propto 1 - [(1 - P(S | G_1)) \times (1 - P(S, G_1 | G_2))]. \end{aligned} \quad (4.15)$$

4.4.3 Empirical Simulations of the Significance Boost

Equation 4.15 results in a boosted significance for any source in the vicinity of a known catalogue source, which is dependent on the SNR, or the level of significance of source detection. The cross correlation with a known position follows a 2D Gaussian PDF and increases the significance by larger amounts for more coincident sources than sources on the periphery of the distribution.

To empirically demonstrate the level of boosting, multiple simulations over a range of SNRs and over the range of the prior knowledge probability (product PDF of the source and catalogue position match given by Equation 4.10) are conducted. Figure 4.18 shows three plots of the posterior probability as a function of the prior probability for three input SNRs: 2, 3, and 4. When the prior probability is zero (i.e. when no corresponding catalogue source is matched with a map source), the posterior converges to the tailed distribution probability for the input SNR.

For example; $\text{SNR} = 2$, prior probability = 0.0, posterior probability ≈ 0.9545 , i.e. the same 2σ significance result as before. However, when the sources do coincide e.g. with a prior probability of 0.94, then the posterior is boosted to ≈ 0.9973 , i.e. a 3σ significance result.

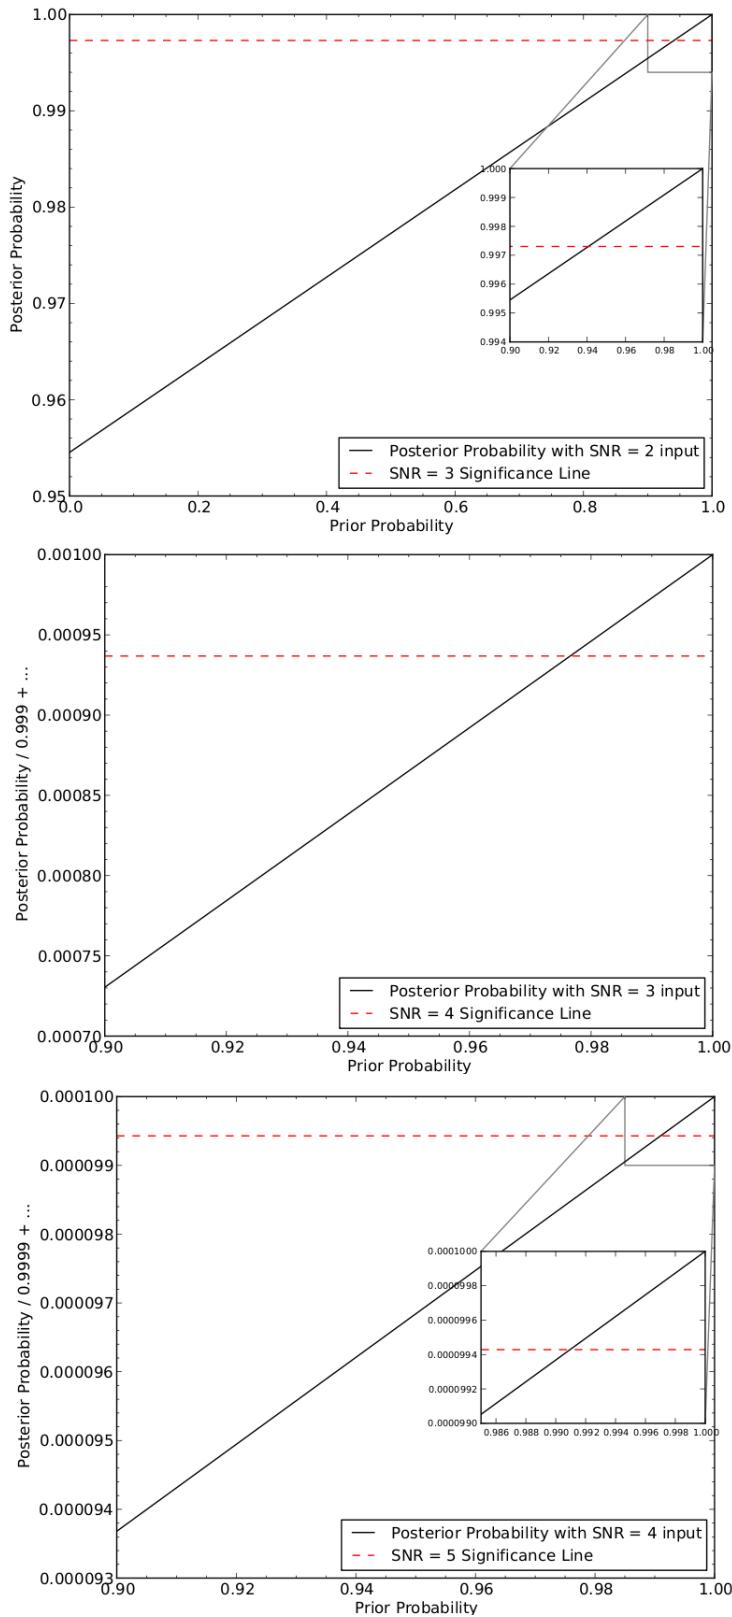


Figure 4.18: Three plots displaying the posterior probability as a function of the prior probability for input SNRs: 2 (top), 3 (middle) and 4 (bottom). The next significance level line has been plotted on each figure, with insets for clarity at the injunction of the posterior and next significance lines. Note the axes scales differ between each figure.

Table 4.3: Table of boosting threshold probabilities

Input SNR Level	Probability at input SNR	Next Significance Level	Prior Probability needed for
			next Significance Level
2	0.95452799	3	0.9408
3	0.99730431	4	0.9766
4	0.99993681	5	0.9910
5	0.99999942	-	-

Table of probabilities for input SNRs and the needed prior probabilities to boost to the next sigma significance level.

The threshold levels to boost up one significance level are given in Table 4.3. The posterior converges to a probability of 1.0 when the prior probability is also 1.0. Quantitatively, this is misleading as you can never completely reject the *null hypothesis* (please consult Section 6.5.1), but the fundamentals of probability theory state that the probabilities of something being true plus the probabilities of something being false equal 1.0 (Sivia and Skilling 2006). Qualitatively, this is seen as a *perfect* match between the two sources, which is unlikely, given different levels of precision and the constituent errors.

To demonstrate the requirements to boost over a continuous range of SNRs, Figure 4.19 shows the input SNRs for a range of non-discrete values (to 2 decimal places, i.e. 2.50, 3.95 etc...) and the required prior probability required to boost to the next significance level. It is evident that for each significance level, the trend is similar but with a greater exponential gradient when the prior probability $\rightarrow 1.0$ for the higher significance levels.

Boosting beyond the 5σ significance is unlikely to produce meaningful practical results for the applications of source detection. The level of precision in the prior probability needed to boost to higher levels of significance, far exceeds that of the expected precision on source position measurements, even at radio interferometry astrometric measurements. Moreover, a 5σ significance detection is considered *the* detection threshold in many physics and astrophysics communities.

For completeness, the values needed for the prior probability to boost to the 5σ significance level are shown in Figure 4.20, where the prior probability is calculated with a resolution to 4 decimal places, which equates to the minimum possible SNR to be boosted to 5σ being $\text{SNR} = 2.77$ (prior = 0.9999 rounded). To boost a 2.00 SNR source to the 5σ level, requires a prior probability of 0.999988 (rounded to 6 decimal places). This

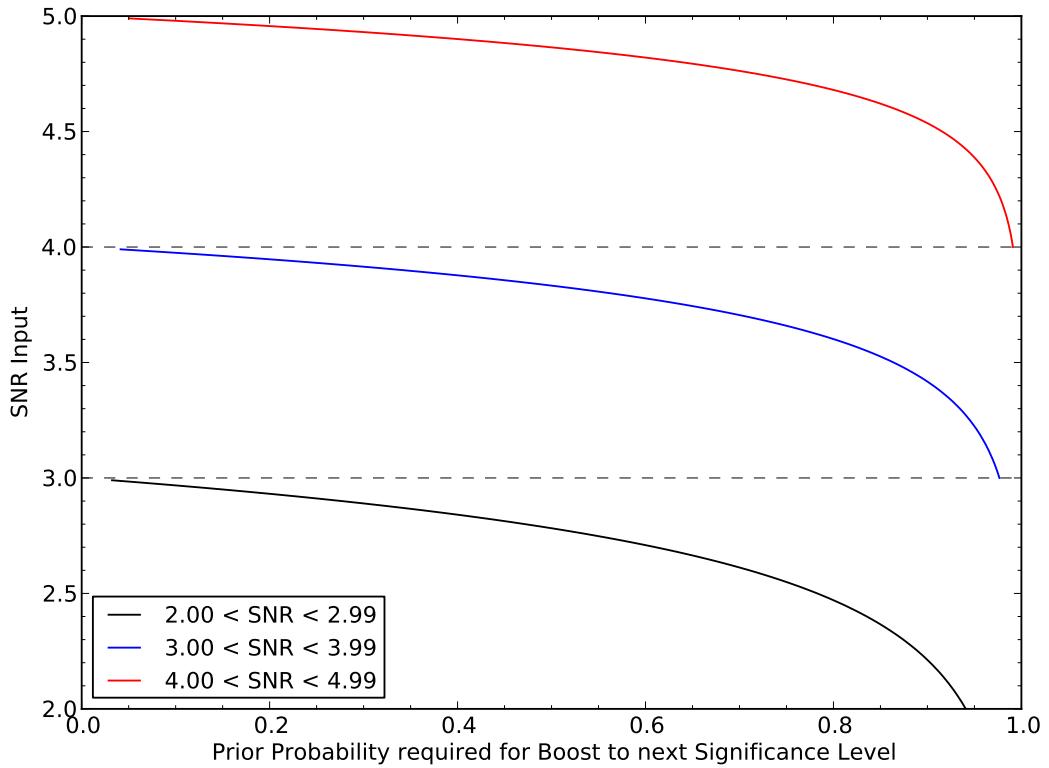


Figure 4.19: Plot showing the continuous range of $2.00 \leq \text{SNR} \leq 4.99$ prior probabilities needed to boost the sigma significance to the next integer level. i.e. $2.78 \rightarrow 3\sigma$ level etc. The dashed lines denoting the integer significance levels are shown for clarity.

again raises the issues of the level of precision needed for the source positions, and also for the resolution of the product of the two 2D Gaussian PDFs in Equation 4.10 (which is $1/10^{th}$ of the smallest of the two positional errors) in order to boost over multiple integer significance levels.

The resolution of the product PDF (Equation 4.10) can be set to any incremental level at the expense of computational time for higher resolutions. However, the amount of time necessary to perform these calculations is by no means a limiting factor, but is stated here as a consideration.

4.5 Conclusions and Discussions

Section 4.1 focuses on the catalogue amalgamation routine for the Cyg OB2 Amalgamated Catalogue and its constituent catalogues. The Java module STILTS is imported within the Jython environment to enable scripting of the concatenation of meta table data from the VO and other databases. The exact procedure is described here and the problems of

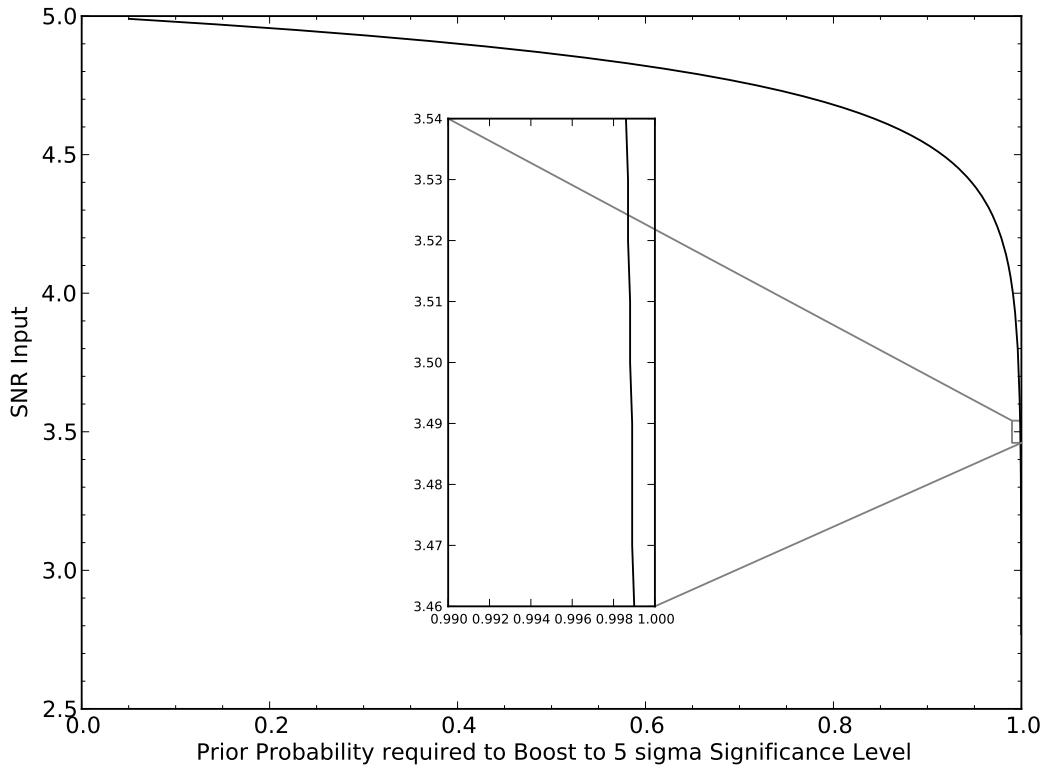


Figure 4.20: Figure showing the prior probability required to boost a SNR to the 5σ significance level. The figure is restricted to a prior resolution of 4 decimal places where the limit for boosting to 5σ is $\text{SNR} = 2.77$ (prior required = 0.9999). The inset shows a zoomed plot around an SNR of 3.5 where the prior needed to boost is ≈ 0.999 .

catalogue standardisation are also discussed. The main Cyg OB2 catalogue contains 6500 non-duplicate entries and the known OB catalogue contains 290 sources.

Near infra-red photometric methods such as those employed by Comerón et al. (2002) are also presented and identify 350 potential OB candidates. The same methods are applied to the 2MASS catalogue to create an additional 1930 OB candidates from the JHK colours. All of these catalogue will be pivotal in aiding the classification of radio sources in COBRaS maps.

The second part of this chapter contains a simple modification to the Floodfill algorithm utilised by Hancock et al. (2012) and Hales et al. (2012), exploiting a fundamental property of radio interferometric constructed images. By using the synthesised beam size as a condition for source detection (island size) results in the lowering of false positive detections from a Gaussian simulation. Only at high SNRs (50, 100) were there any signs of false positives (1 in every 500 sources for point sources) from the source detection, and this is most likely attributed to artefacts from IMAGR scattering flux during source creation.

Accounting for these effects, the reliability of the source detection is $> 99.9\%$ for point sources with an $\text{SNR} \geq 3$. The same modification allows fainter sources to be detected down to an $\text{SNR} = 2$, with $\sim 90\%$ detection for point sources, dropping to $\sim 36\%$ for resolved sources.

A new flux extraction method is presented in Section 4.3, which is independent of the source structure. The pixel-by-pixel (PP) flux extraction method is shown to be as accurate as existing methods while not assuming source Gaussianity. It is also tested on faint sources ($\text{SNR} < 5$), with all determined fluxes within the associated errors for point sources. For resolved sources, the fluxes coincide with input fluxes at $\text{SNR} \geq 7$ within errors, after accounting for the scattered flux from IMAGR. The PP method is comparable or performs better in flux extraction than JMFIT for $\text{SNR} \geq 5$. For lower SNRs the method suffers from its intrinsic property of determining fluxes from higher noise. The sources in the simulations presented here are Gaussian and resolved Gaussians, because a standard set of non-Gaussian sources with known signal properties does not appear in the literature (Hales et al. 2012).

The final section of the chapter describes a novel method to statistically boost the sigma significance level of a faint source with $\text{SNR} < 5.0$. Using Bayes' Theorem, the posterior probability is given as the tailed Gaussian distribution with the SNR substituting the critical z value (likelihood) and the product of two 2D Gaussian distributions of the map source and catalogue source (prior). Empirical simulations are made to find the theoretical prior probability required to boost a source SNR to a higher sigma significance level. The precision of the prior probability needed to boost $2.00 \leq \text{SNR} < 5.00$ to the accepted detection level of 5σ is also discussed.

The Bayesian boosting module uses prior information in the form of catalogues amalgamated from previous surveys. The boost requires the positions of source in the COBRaS maps to coincide with those from the catalogues. Two potential problems arise from this set up. The first is that the resolution of the COBRaS survey and any previous surveys will almost certainly differ, with the source positions from COBRaS more likely to have a higher precision from the high resolution of e-MERLIN. The other consideration is if COBRaS is compared to another high resolution observation with a significant elapse in time between the two observations, proper motion effects may influence the absolute positions from both surveys. These points need to be considered when deploying the Bayesian module to sources.

This chapter raises a new issue concerning standardisation, particularly with catalogues from surveys. Amalgamating catalogues with known identifiers enables accurate cross correlation to build complete and large catalogues over a common region in space. Concatenating catalogues with source sky positions is limited by the positional accuracy and precision of the constituent catalogues. While this offers the only way to compare every known source in the sky, because it is the only common property of any source, localised regions of space (and more practically useful, regions within the Galaxy) can and should be treated differently. For a definitive catalogue of Cyg OB2 for example, a common identifier is necessary for all objects belonging to the association for future reference and studies, as opposed to the various identifiers (e.g. MT91, AFM2007) currently available. It is possibly within the future scope of the COBRaS project to provide this, building on the back of the most detailed and deep radio map of the core region of Cyg OB2.

For the source and flux extraction, the programming language Python is again deployed here because of its link with Parseltongue and therefore the AIPS package, and because it is free of compilers and easy to run and modify (syntax friendly). The flux extraction code presented in this thesis has been set up to accept image files directly from AIPS, however it has been run on generic fits images, using the *pyfits* module to read the image into *NumPy* arrays. For the latter, the user will have to designate the beam size of the image to utilise the full capabilities of the algorithm.

The unique feature of the pixel-by-pixel flux extraction method is that the algorithm does not assume the source to have any specific structure. For non-Gaussian sources such as massive binary interaction regions this will produce an accurate flux determination (for a good SNR e.g. ≥ 7), free of any constraining fit parameters used with Gaussian fits. An interesting application would be to test this method on diffuse regions such as HII regions (or ultra-compact HII regions), where one would expect that the algorithm would be limited by its ability to assess the noise of the image from the minority of pixels representing the background, if the diffuse region covered a significant portion of the image. This would affect the flux calculation in conjunction with the associated errors from the algorithm.

It is possible to search for sources in the visibility plane instead of the image domain. The advantages of model fitting in the visibility plane is that after deconvolution with CLEAN or MEM, the noise distribution of the image is poorly understood (Pearson 1999).

Model fitting is most useful when the sky brightness can be described by a model with a small number of parameters. An example is a field with a small number of unresolved sources where the parameters are the position and flux densities.

The limitations of model fitting are dependent on the model fitting the data well, the errors are Gaussian and are known, and that there are no calibration errors and any errors in the observed visibilities are uncorrelated (Pearson 1999). There are a few programs which conduct model fitting: the AIPS task UVFIT, SLIME, an add-on AIPS task with a graphical interface, and DIFMAP (Shepherd 1997) another graphical editor.

Visibility plane source detection could be a avenue of future investigation for big projects such as MeerGAL with MeerKat and other SKA surveys.

Chapter 5

Radio Emission of Massive Stars in Cyg OB2

If we knew what it was we were doing, it would not be called research.

Albert Einstein

The theme of this chapter involves the radio emission from massive stars in Cyg OB2. Firstly, an introduction on the background of massive stars and massive star winds is given. Secondly, descriptions of the mechanisms of radio emission excess from massive stars is presented. Thirdly, an investigation into the predicted mass loss rates and radio fluxes of single massive stars, with and without clumping effects. Finally, a first look at the initial COBRA-S maps accompanied with source and flux lists.

5.1 Massive Stars

Stars are comprised of the fourth state of matter, plasma, and are in hydrostatic equilibrium, the balance between self-gravity and thermal pressure. They are internally powered by the thermonuclear fusion of elements in the core (or shell), where energy and synthesised elements are created and propagate through to the photosphere via convection and radiation processes. At the photosphere they radiate energy as electromagnetic radiation

or photons, and material via a stellar wind. These processes change over the lifetime of a star as it evolves, changing also the composition of the star e.g. neutron-degeneration of neutron stars.

The most fundamental stellar parameter is *mass*, which dictates the star's role and future in the Galaxy (see Section 1.3.2). Using our the mass of our Sun as the standard unit (1 solar mass = M_{\odot}), stars have a range of stellar masses from the smallest; e.g. OTS 44, 0.013 M_{\odot} (Luhman et al. 2005) to the largest; a1 in the R136 cluster $\sim 300 M_{\odot}$ (Crowther et al. 2010). Photospheric (or effective) temperature (T_{eff}) broadly scales with stellar mass (more mass leads to higher pressure and densities, which in turn leads to higher temperatures), and can be determined through the spectroscopic study of lines arising at known specific, discrete wavelengths.

To separate the diverse population of stars, a spectral classification scheme called the MK system devised by Morgan et al. (1943) is used, with some modern additions accounting for new discoveries. The *Havard sequence* assigns a letter from the group - OBAFGKM - to stars depending on which spectral species appear in their spectra (which roughly depends on their effective temperature). O denotes the hottest stars, decreasing through the sequence to M, the coolest stars. Each group is further split into subgroups, designated with a number from 0; the hottest, to 9 the coolest within a group, with an exception for O stars which start from O2 as the hottest. Spectral precision is such that the subgroup can be in increments of 0.5, or finer increments around effective temperatures where complex stellar behaviour is present (as will be seen later).

Stars are further subdivided by their luminosity class, or by the strength of certain absorption lines and are represented by Roman numerals. I stands for supergiants, II, III and IV for bright, normal and sub- giants respectively, V for main sequence stars (also known as 'dwarfs'), and VI for very faint dwarfs and white dwarfs. The supergiants category can be split again into Ia, Iab and Ib from most to least luminous. Further additions to the spectral classification can be made to include peculiar features and specific emission lines by including lower case letters to the description. Figure 5.1 shows an example of an Hertzprung-Russell (HR) diagram depicting the stellar classification of a number of stars with differing spectral types and luminosity classes.

This thesis is concerned with massive stars, which are defined as having an initial stellar mass $> 8 M_{\odot}$ and will conclude their lives as spectacular supernova events. This encapsulates a number of stellar objects at different evolutionary stages, including O and

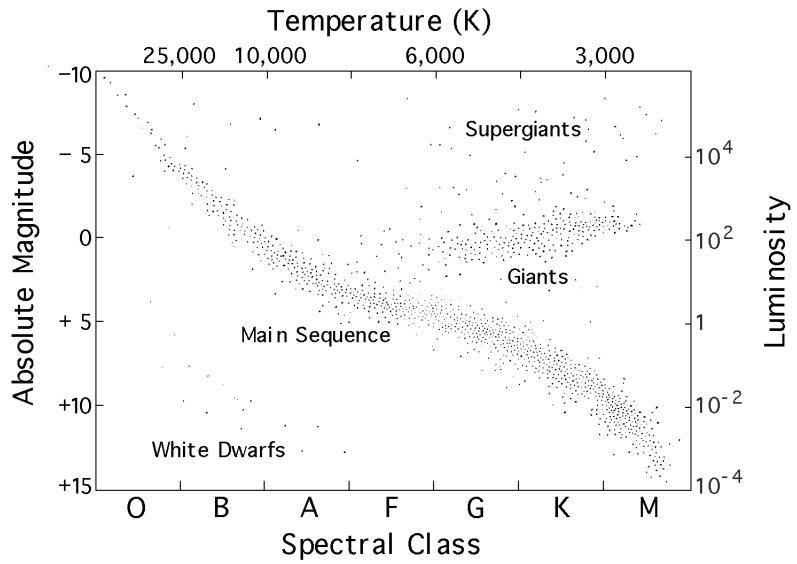


Figure 5.1: Hertzsprung-Russell Diagram showing the Havard sequence and luminosity classes.

early (the hottest) B stars, all Wolf-Rayet (WR) stars, Luminous Blue Variables (LBVs), and red and yellow supergiants.

5.2 The Winds of Massive Stars

The two quantities emitted from the photosphere are electromagnetic radiation and particles. Emission of the latter constitutes the stellar wind - a continuous outflow of material from the photosphere to the interstellar medium (ISM). This next section describes the nature of stellar winds from massive stars.

5.2.1 Mass Loss

For our Sun, the mass loss rate is on the order of $\sim 10^{-14} M_{\odot} \text{ yr}^{-1}$. Over the Sun's lifetime of ~ 10 billion years, it will lose $\sim 10^{-2}\%$ of its initial mass. This is feeble in comparison to the mass loss rates of massive stars which are on the order of $10^{-6} M_{\odot} \text{ yr}^{-1}$ (Puls et al. 2008), increasing to $10^{-3} M_{\odot} \text{ yr}^{-1}$ in the most extreme case of the LBV η Carina (Hillier et al. 2001). Therefore, massive stars will lose a substantial fraction of their initial masses over their lifetimes (typically a few million years), indicating the importance of mass loss on their evolution.

The single scattering limit for the mass loss rate (\dot{M}_{ss}), is driven by line radiation and assumes that so many spectral lines exist that all the photons leaving the star are absorbed

within the wind. In essence, the momentum of photons L/c is equal to the momentum in the wind $\dot{M}_{ss} v_\infty$, giving the mass loss rate

$$\dot{M}_{ss} = \frac{L}{v_\infty c}. \quad (5.1)$$

Inputting values for L and v_∞ for ζ Puppis: $L = 8 \times 10^5 \text{ L}_\odot$ and $v_\infty = 2200 \text{ km s}^{-1}$ (Lamers and Leitherer 1993) gives a single scattering mass loss rate $\dot{M}_{ss} = 10^{-5} \text{ M}_\odot \text{ yr}^{-1}$. This is a factor of 0.25 within the observed mass loss rate of $\dot{M}_{ss} = 2.4 \times 10^{-6} \text{ M}_\odot \text{ yr}^{-1}$ (Lamers and Leitherer 1993). There are also theoretical reasons why assuming that all of the momentum from photons is transferred to the wind material, and will become evident later.

To answer the question of how a massive star loses mass, one may first consider the radiative force only as a function of the frequency independent free electron scattering. The flux of radiative energy at any given radius r is $L/(4\pi r^2)$ and the corresponding radiative momentum flux is $L/4\pi r^2 c$. The electron scattering opacity is $\kappa_e = \sigma_e / \mu_e$, where σ_e is the Thomson scattering cross-section ($6.6 \times 10^{-25} \text{ cm}^2$), and μ_e is the mean atomic mass of the material in the atmosphere. In the majority of cases hydrogen makes up the bulk of the wind plasma, resulting in $\mu_e \sim 1$ and $\kappa_e = \sigma_e$. The product of the opacity of electron scattering and the radiative momentum flux produces the acceleration due to continuum opacity g_{cont}

$$g_{\text{cont}} = \frac{\sigma_e L}{4\pi r^2 c}. \quad (5.2)$$

The ratio of the acceleration of continuum opacity to that of the gravitational acceleration of the star given by $g_{\text{grav}} = GM/r^2$ is denoted by Γ_e and subsequently by

$$\begin{aligned} \Gamma_e &= \frac{g_{\text{cont}}}{g_{\text{grav}}} = \frac{\sigma_e L}{4\pi r^2 c} \cdot \frac{r^2}{GM} \\ &= \frac{\sigma_e L}{4\pi c GM}. \end{aligned} \quad (5.3)$$

This is known as the Eddington parameter, and for massive stars Γ_e approaches unity¹, signifying that matter only needs a small additional acceleration component to break free

¹This is a broad statement and neglects the effects of rotation which modifies to the $\Omega\Gamma$ limit (Maeder and Meynet 2000).

of the star's gravitational field and become part of the stellar wind. In practice if Γ_e reaches 0.639 or the so called break-up limit, then extremely high mass loss rates can occur such as those in LBVs and WR stars (Maeder and Meynet 2000). However, this addition acceleration component is gained through the opacity of metal lines, and is now discussed.

5.2.2 Radiative Line-Driven Wind Theory

It has been stated earlier that the ejection of particles from the outer atmosphere of massive stars makes up the stellar wind. However, it is the released photons, each carrying momentum $h\nu/c$ which drives and accelerates the wind through the opacity of metal ions (Milne 1926; Kudritzki and Puls 2000 and references within). Morton (1967) discovered from ultra-violet (UV) observations that if the acceleration from metal ions could be shared among the more abundant hydrogen and helium species, then significant mass loss from this radiative line-driving mechanism could be viable. Lucy and Solomon (1970) and Castor et al. (1975) (hereafter CAK) devised the first theories of radiative line-driven winds via the absorption in the UV resonance lines of a number of species.

There are two key concepts which explain the effective acceleration of the line-driving mechanism. The first is the momentum transfer from photons to ions through line scattering, where the photons originate with a range of angles in the stellar atmosphere and become more or less isotropic after absorption and re-emission (Puls et al. 2008). The change in direction angle leads to the radial net outwards transfer of momentum. The second concept concerns the Doppler effect. If the atmosphere were static but with strong line absorption, photons would be absorbed and scattered in the lower layers of the atmosphere and the upper layers would not receive direct radiation at the wavelength of the absorption line. If the atmosphere is moving radially outwards, a velocity gradient exists between the atoms in the outer layers and the seemingly red-shifted photons in the photosphere. Therefore any photons which are not absorbed in the lower layers can be absorbed in the upper layers at a different wavelength. In essence, the Doppler shift allows the absorption of unattenuated UV continuum photons, whose wavelength has shifted to the wavelength of the absorption ion(s).

The dominant transition lines that contribute the most to the acceleration are UV resonance lines of abundant elements, i.e. the absorption of photons from the ground state of C, N and O ions and meta-state levels of Fe. Additional contribution from less

abundant (but still significant) elements include P, Si, S ions. Lamers and Cassinelli (1999) show that it is possible for absorption to occur $\sim 10^9$ times per absorbing ion, which accelerates the wind to a terminal velocity v_∞ after 10^4 seconds. This amounts to the wind attaining terminal velocity within a few stellar radii. This result becomes highly significant later in this chapter (Section 5.3).

So far it has been stated that photons propagate radially outwards transferring momentum to ions in the process (Puls et al. 2008). In dense winds the transfer of momentum from the metal ions to the hydrogen and helium species within the wind is accomplished by Coulomb collisions (Puls et al. 2008). Via this mechanism it is possible to accelerate the entire bulk of the wind plasma. Figure 5.2 shows a simplified cartoon of the radiative line-driven theory setup. In this figure a layer of depth Δv lies a distance r away from the centre of a massive star, and has a mass Δm , density $\rho(r)$ and a local optical depth $\tau_s(r)$.

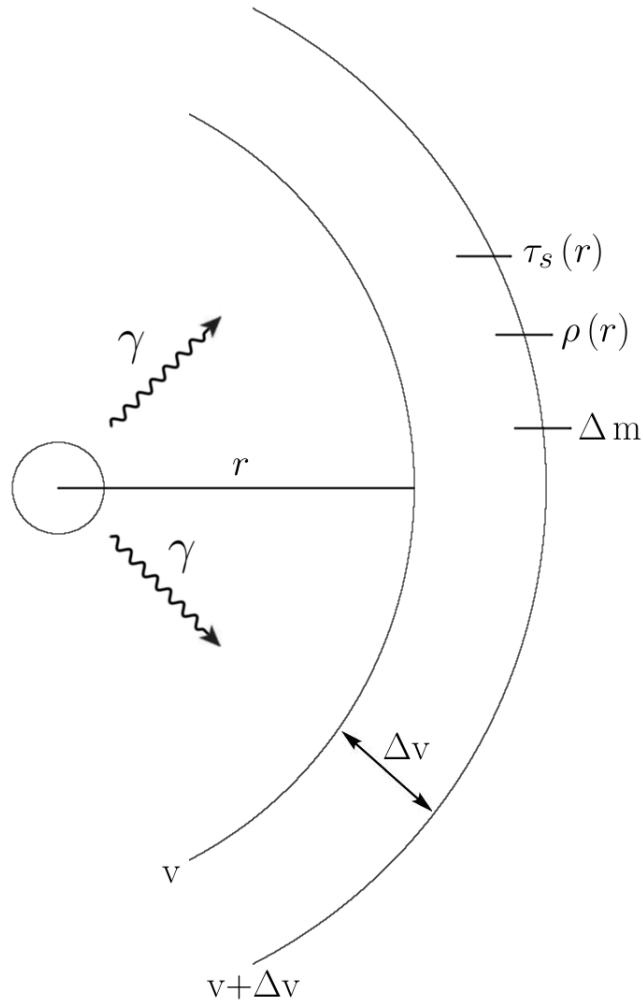


Figure 5.2: Cartoon of a simplified radiative line-driven wind theory.

A simplified derivation of how the spectral lines can drive a wind is now shown, for a full description the reader is referred to the work of CAK. First we utilise three simplifying assumptions:

1. The massive star is assumed to be a point source and only radially streaming photons interact with the layer.
2. Only strong lines drive the stellar wind with an optical depth $\tau_s \gg 1$, absorbing all the photons in a bandwidth $\Delta\nu_i$.
3. A large local velocity gradient $\Delta v = \frac{dv}{dr}\Delta r$ exists across the layer and a range of frequencies are absorbed by a line transition of rest frequency ν_i (ignoring intrinsic width).

The absorption bandwidth $\Delta\nu_i$ can be described in terms of velocity v as a consequence of the Doppler effect, which in turn can be described in terms of the velocity gradient

$$\Delta\nu_i = \frac{\Delta v}{c} \nu_i = \frac{dv}{dr} \Delta r \frac{\nu_i}{c}. \quad (5.4)$$

The acceleration due to the opacity of spectral lines g_{lines} is equal to the product of the total momentum provided by the star (momentum of photons) and the fraction absorbed by a strong line i of frequency ν_i and width $\Delta\nu_i$ (substituting Equation 5.4) divided by the change in mass ΔM

$$\begin{aligned} g_{\text{lines}} &= \frac{1}{\Delta M} \frac{L}{c} \cdot \frac{L_\nu}{L} \Delta\nu_i \\ &= \frac{L}{\Delta M c} \frac{L_\nu}{L} \frac{dv}{dr} \frac{\nu_i}{c} \Delta r. \end{aligned} \quad (5.5)$$

where the change in mass ΔM is defined by the mass continuity relation

$$\dot{M} = 4\pi r^2 \rho v \quad (5.6)$$

Equation 5.5 contains the contribution of only one spectral line. To include all species i ,

simply sum over i and substitute $\dot{M}/v = 4\pi r^2 \rho \Delta r$:

$$\begin{aligned}
 g_{\text{lines}} &= \frac{L}{c} \frac{1}{4\pi r^2 \rho \Delta r} \sum_i \frac{L_\nu}{L} \frac{dv}{dr} \frac{\nu_i}{c} \\
 &= \frac{L}{c^2} \frac{v}{\dot{M}} \frac{dv}{dr} \sum_i \frac{L_\nu \nu_i}{L} \\
 &= \frac{L}{c^2} \frac{v}{\dot{M}} \frac{dv}{dr} N_{\text{eff}}, \tag{5.7}
 \end{aligned}$$

where N_{eff} represents the number of effective optically thick absorbing lines (dimensionless). Now the total acceleration of the radiative force can be determined by combining the three acceleration components; line, continuum and gravity:

$$\begin{aligned}
 g_{\text{tot}} &= g_{\text{lines}} + g_{\text{cont}} - g_{\text{grav}} \\
 &= g_{\text{lines}} + \frac{\sigma_e L}{4\pi r^2 c} - \frac{GM}{r^2}, \tag{5.8}
 \end{aligned}$$

and substituting the Eddington parameter $\Gamma_e GM = \frac{\sigma_e L}{4\pi c}$,

$$\begin{aligned}
 g_{\text{tot}} &= g_{\text{lines}} + \frac{GM\Gamma_e}{r^2} - \frac{GM}{r^2} \\
 &= g_{\text{lines}} - \frac{GM}{r^2} (1 - \Gamma_e). \tag{5.9}
 \end{aligned}$$

Defining the effective mass M_{eff} as the stellar mass with the electron scattering correction Γ_e as $M_{\text{eff}} = M(1 - \Gamma_e)$,

$$g_{\text{tot}} = g_{\text{lines}} - \frac{GM_{\text{eff}}}{r^2}. \tag{5.10}$$

After a simple algebraic rearranging, the expression for the radiative acceleration required to generate a stellar wind is given by

$$\begin{aligned}
 g_{\text{lines}} &> \frac{GM_{\text{eff}}}{r^2} \\
 \frac{L}{c^2} N_{\text{eff}} \frac{v}{\dot{M}} \frac{dv}{dr} &> \frac{GM_{\text{eff}}}{r^2}. \tag{5.11}
 \end{aligned}$$

It can be shown using the single scattering limit for optically thick lines $\dot{M} \simeq N_{\text{eff}} L/c^2$,

and the typical parameters for massive stars; $L \simeq 10^5 L_\odot$, $\dot{M} = 4 \times 10^{-6} M_\odot \text{ yr}^{-1}$, that $N_{\text{eff}} \simeq 100$'s of spectral lines are required to drive the stellar wind of a massive star.

5.3 Radio Emission from Massive Stars

The winds of massive stars are comprised of hot plasma, mostly in the form of ionised hydrogen accelerated through radiative line-driving mechanisms explained in the previous section. This property provides a unique method of detection through the interaction of wind particles with one another when accelerated to high speeds. This next section explains the mechanisms where radio flux observations estimate the mass loss rate of massive stars.

The wind plasma from a massive star is able to produce emission from the interactions of ionised particles. The majority of these particles are protons and electrons from ionised hydrogen which makes up the bulk of the material in the wind. When two charged particles come in close proximity, the trajectory of one is deflected by the other, and the resulting deceleration causes the slowing particle to lose kinetic energy in the form of a photon. This is a consequence of the conservation of energy where the photon emitted can be at a frequency in the radio regime.

The physical mechanism is called *bremsstrahlung*, or ‘braking radiation’ because of the deceleration of one of the particles, or free-free radiation because of the interaction of two free particles. Figure 5.3 shows the Bremsstrahlung mechanism for two oppositely charged particles, in this instance a proton and electron. In the outer parts of the wind, the speed of the outflowing material has reached terminal velocity, increasing the incidence of Bremsstrahlung, assuming the wind is hot, maintaining ionised hydrogen throughout. With these conditions, the emission arises in the radio regime, where the intrinsically thermal process (positive spectral index) creates stronger fluxes at higher frequencies. This is typically observable at frequencies of a few GHz.

5.3.1 Thermal Free-free Radio Emission Spectrum

The free-free radio and IR excess emission spectrum were first derived independently by Wright and Barlow (1975) and Panagia and Felli (1975). For a full detailed explanation of this derivation, the reader is asked to consult these papers. A simplified proof is given here.

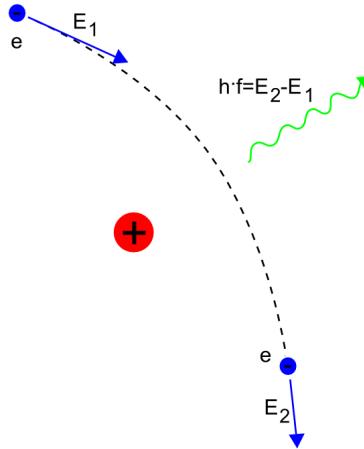


Figure 5.3: Free-free emission (Bremsstrahlung) process with two oppositely charged particles; one proton and one electron. Some of the emission from the winds of massive stars arises from the radio frequency part of the electromagnetic spectrum.

The free-free absorption coefficient κ_ν of an ionised gas with density ρ (Güdel 2002) is given by

$$\kappa_\nu \propto \nu^{-2} T^{-\frac{3}{2}} \rho^2, \quad (5.12)$$

where $\rho^2 = n_e n_i$, n_e and n_i are the electron and ion number densities in the wind, ν is the frequency of the excess radiation and T is the temperature of the wind. The equation of mass continuity (Equation 5.6) is substituted for ρ in Equation 5.12 giving

$$\begin{aligned} \rho^2 &\propto \left(\frac{\dot{M}}{v_\infty} \right)^2 \frac{1}{r^4} \\ \kappa_\nu &\propto \nu^{-2} T^{-\frac{3}{2}} \left(\frac{\dot{M}}{v_\infty} \right)^2 r^{-4}, \end{aligned} \quad (5.13)$$

where it is assumed the wind has reached terminal velocity v_∞ . To get the optical depth

τ_ν , Equation 5.13 is integrated from r to infinity

$$\begin{aligned}\tau_\nu &= \int_r^\infty \kappa_\nu dr \\ &= \nu^{-2} T^{-\frac{3}{2}} \left(\frac{\dot{M}}{v_\infty} \right)^2 \int_r^\infty r^{-4} dr \\ &\propto \nu^{-2} T^{-\frac{3}{2}} \left(\frac{\dot{M}}{v_\infty} \right)^2 r^{-3}.\end{aligned}\quad (5.14)$$

This equation shows that the optical depth $\tau_\nu = \tau_{\text{eff}} = 1$ is reached at a radius r ($\tau_\nu = \tau_{\text{eff}}$)

$$r(\tau_\nu = \tau_{\text{eff}}) = \nu^{-\frac{2}{3}} T^{-\frac{1}{2}} \left(\frac{\dot{M}}{v_\infty} \right)^{\frac{2}{3}}. \quad (5.15)$$

The luminosity of the wind can be represented as the zero order solution of radiative transfer in a stellar atmosphere with the Eddington-Barbier relation

$$L_\nu \propto r^2 (\tau_\nu = \tau_{\text{eff}}) B_\nu (T(\tau_\nu = \tau_{\text{eff}})). \quad (5.16)$$

Using the Rayleigh-Jeans approximation to the Planck function $B_\nu = T\nu^2$ and substituting Equation 5.15 for r , results in the wind luminosity taking the form

$$\begin{aligned}L_\nu &\propto \left(\nu^{-\frac{2}{3}} T^{-\frac{1}{2}} \left(\frac{\dot{M}}{v_\infty} \right)^{\frac{2}{3}} \right)^2 T\nu^2 \\ &\propto \nu^{\frac{2}{3}} \left(\frac{\dot{M}}{v_\infty} \right)^{\frac{4}{3}}.\end{aligned}\quad (5.17)$$

The observed radio flux received at Earth S_ν , is related to the luminosity by $L_\nu = 4\pi d^2 S_\nu \rightarrow L_\nu \propto d^2 S_\nu$, where d is the distance from Earth to the source. Substituting this relation into Equation 5.17 and solving for \dot{M} gives

$$\begin{aligned}d^2 S_\nu &\propto \nu^{\frac{2}{3}} \left(\frac{\dot{M}}{v_\infty} \right)^{\frac{4}{3}} \\ \dot{M} &\propto \nu^{-\frac{1}{2}} d^{\frac{3}{2}} S_\nu^{\frac{3}{4}} v_\infty.\end{aligned}\quad (5.18)$$

This relation requires a couple of modifications to compensate for the over-simplification of some of the assumptions made. Firstly a quantum mechanical correction called the Gaunt Factor g_ν , a function of the electron temperature T_e and frequency ν (or wavelength) is introduced with the consequence of flattening the spectrum $S_\nu \propto \nu^{2/3} \rightarrow S_\nu \propto \nu^{0.6}$ (Wright and Barlow 1975; Cox 2000):

$$g_\nu = 1.38 T_e^{-0.34} \nu_{\text{GHz}}^{-0.11} \quad (5.19)$$

which is valid for radio wavelengths (but not for mm or IR wavelengths). This spectrum is characteristic of an isothermal wind with a constant velocity (v_∞) and density $\rho \propto r^{-2}$.

The second modification allows a correction for the total gas number density relation $\rho^2 = n_e n_i$, which assumes the wind is entirely made up of hydrogen ions. Therefore a term containing the mean atomic mass of ions in the wind μ_i , the number of free electrons per ion γ_e (therefore $n_e = \gamma_e n_i$), and the charge state of the atoms Z . Both corrections modify Equation 5.18 to

$$\dot{M} = 0.095 (S_\nu d^2)^{\frac{3}{4}} \left(\frac{\mu_i^2}{Z^2 \gamma_e} \right)^{\frac{1}{2}} (\nu g_\nu)^{-\frac{1}{2}} v_\infty, \quad (5.20)$$

where ν is in the units of Hz, T is in K, v_∞ is in km s^{-1} , \dot{M} is in $\text{M}_\odot \text{ yr}^{-1}$, d is in kpc, and S_ν is in Jy ($10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$). This relation can be re-arranged to make S_ν the subject, as presented in Wright and Barlow (1975) and Panagia and Felli (1975):

$$\begin{aligned} S_\nu^{-\frac{3}{4}} &= 0.095 (d^2)^{\frac{3}{4}} \left(\frac{v_\infty}{\dot{M}} \right) \left(\frac{\mu_i^2}{Z^2 \gamma_e} \right)^{\frac{1}{2}} (\nu g_\nu)^{-\frac{1}{2}} \\ S_\nu &= 23.2 d^{-2} \left(\frac{\dot{M}}{v_\infty} \right)^{\frac{4}{3}} \left(\frac{\mu_i^2}{Z^2 \gamma_e} \right)^{-\frac{2}{3}} (\nu g_\nu)^{\frac{2}{3}} \quad [\text{Jy}]. \end{aligned} \quad (5.21)$$

Equation 5.21 is used later in this chapter to estimate the smooth wind radio fluxes of massive stars (Section 5.5.3).

5.3.2 Radio Emission from Single Massive Stars

It has been described in previous sections that radio emission can arise from the ionised winds of massive stars through free-free radiation. This process is dependent on the temperature of the wind, which needs to be hot enough to ionise hydrogen which constitutes

the bulk of the wind. Because of this dependence, free-free radio emission is termed *thermal*, where the radio spectrum steadily increases with increasing frequency. The relation between flux and frequency is called the spectral index (α) and has been shown for thermal emission to have a positive index ($\alpha \sim 0.6$).

The radio flux from Bremsstrahlung is weak with typical O stars located at 1 kpc having a flux ~ 0.1 mJy (1 to 15 GHz) (Puls et al. 2008), with the stronger emitters on the order of a few mJy at 5 GHz (6cm) (Abbott et al. 1980; Bieging et al. 1989; Howarth and Brown 1991). This only allows detectable radio emission from massive stars with winds which are completely ionised (Drake and Linsky 1989), which includes O stars, early B supergiants and WR stars.

Some OB stars have displayed non-thermal radio emission (Abbott et al. 1984; Bieging et al. 1989; Güdel 2002), and WR stars (Güdel 2002 and references therein) with spectral indexes approaching zero and even negative indexes. White (1985) proposed a mechanism in single stars where synchrotron emission from electrons accelerated by shocks propagating through the winds. These shocks are predicted to occur from radiative instabilities from the line-driving mechanism in the wind of a single star (Owocki and Rybicki 1984; Chen and White 1994). Some discussions are included at the end of Chen and White (1994) which reveal illuminating, critical comments concerning the possibility of shocks in the wind generating non-thermal emission.

5.3.3 Radio Emission from Massive Binaries

There is now considerable doubt as to whether single stars produce any non-thermal radio emission (Blomme et al. 2010). Observations of seemingly single stars (such as Cyg OB2 No. 8A and No. 9) where modelling of non-thermal processes appeared to explain the higher mass loss rates by including the non-thermal component, were later shown by other observations to be binaries (De Becker et al. 2004; Nazé et al. 2008). This is backed up by theoretical work on non-thermal emission from single stars by van Loo et al. (2006), who concludes that the non-thermal emission would be absorbed by the opacity of the stellar wind.

Whatever the origin, it is clear that the production of non-thermal radiation is a result of synchrotron emission, emitted from relativistic electrons which are accelerated via the first order Fermi mechanism (Fermi 1949; Bell 1978) occurring in the presence of hydrodynamical shocks (White 1985; Chen and White 1994; Güdel 2002; van Loo

et al. 2006; De Becker 2007; Puls et al. 2008 and references therein). In this mechanism, electrons are bounced across the shock, each time gaining energy from the shock and accelerating. A fraction of electrons make enough bounces to attain relativistic speeds (Blomme 2011). The electrons then spiral around magnetic field lines from the massive stars emitting synchrotron radiation which can be observed in the radio regime.

The currently favoured hypothesis of the origin of non-thermal radiation from massive stars is from the wind-wind collisions of massive binaries. In this model, the ram pressure of two radiatively driven winds generates large scale shocks where the two winds collide. The two shocked regions are divided by a contact discontinuity region and the bow shock front envelopes around the star with the weaker stellar wind as seen in Figure 5.4.

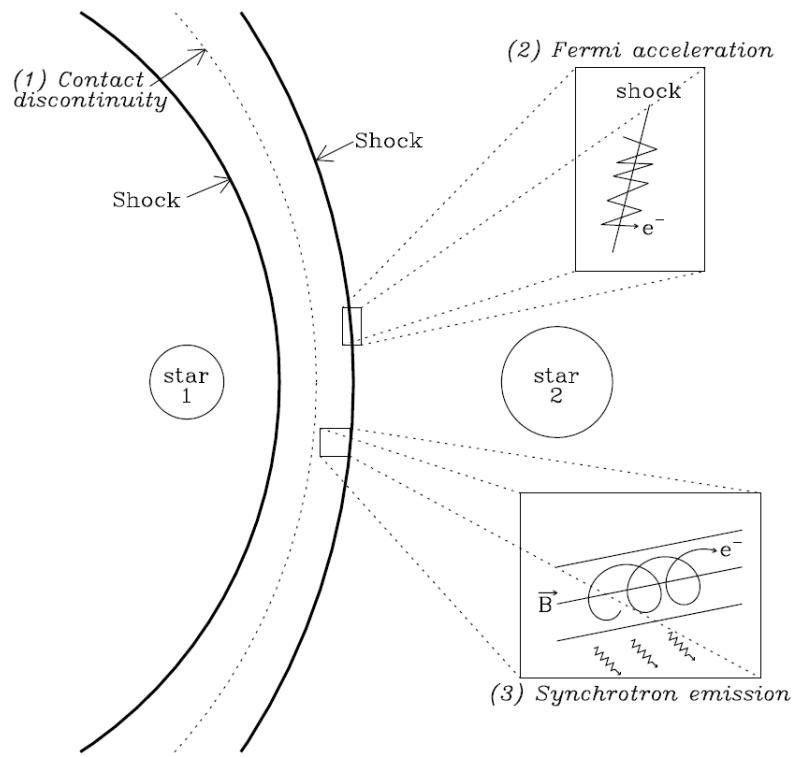


Figure 5.4: Massive star wind-wind collision region depicting the creation of non-thermal radiation from synchrotron emission. Two shock regions created from the collision of two radiatively driven winds, provide the nursing ground for electrons to be accelerated to relativistic speeds via the Fermi mechanism. Figure taken from Blomme (2011).

There is theoretical evidence supporting the physics behind the generation of non-thermal radiation from the wind-wind collision of massive binaries (Stevens et al. 1992; Eichler and Usov 1993; Dougherty et al. 2005; Pittard and Dougherty 2006). These include considerations such as: inverse Compton (IC) cooling, the process where the scattering of low energy photons by relativistic electrons transfers energy to photons from

the electrons i.e. electrons lose energy. IC cooling occurs at the high end of the energy spectrum. Coulomb cooling, the collisions of electrons results in a loss of energy for the electrons. Coulomb cooling occurs at the low end of the energy spectrum. Synchrotron self-absorption, where at low frequencies the synchrotron electrons are optically thick for the synchrotron radiation cause them to scatter of the electrons. Free-free absorption, as described in Section 5.3. The Razin effect, when relativistic particles are in the presence of a plasma, the beaming effects of synchrotron radiation are suppressed (?). Clumping, see Section 5.4 for details.

The inclusion of IC cooling particularly leads to a greater understanding of the electron spectrum and therefore the spectral index of the synchrotron radiation. Moreover, these models enable fits to the radio data to determine the spatial distribution of electrons without knowledge of the magnetic fields within which they lie (De Becker 2007).

There is also mounting evidence from observations that non-thermal radiation arises from massive binaries. For O stars, there are 16 known non-thermal radio emitters, 11 are confirmed binaries and two more are suspected binaries. For WR stars, there are 17 known non-thermal radio emitters, 12 are confirmed binaries and 2 more are suspected binaries (De Becker 2007).

Putting these numbers into context, Sana et al. (2008) determined the binary fraction of massive stars in NGC 6231 as 60% and De Becker et al. (2006) determined that the binary fraction of IC 1805 should range between 20% to 60%. The main issue with binary studies is observational bias which splits into two problems. (1) radial velocity observations typically last a few days with follow-ups on the order of months or years, this leads to the majority of detected spectroscopic binaries having periods of a few days. (2) the inclination angle of the orbital system, where radial velocity studies are biased by edge-on systems and do not include face-on systems.

Despite the high correlation between known binaries and non-thermal emission from massive stars, detecting a massive star with a negative spectral index (i.e. non-thermal radiation) requires follow-up observations from radial velocity studies before being irrefutably designated a binary system.

5.4 Clumping in the Winds of Massive Stars

5.4.1 The Effects of Clumping on Mass Loss Diagnostics

There are currently serious discrepancies between the theoretical radiation-driven mass loss estimations as described by the mass loss recipes of Vink et al. (2001), and a number of different mass loss diagnostics (Prinja et al. 2005; Fullerton et al. 2006; Puls et al. 2006). The differences between the theoretical models and observational mass loss rates is probably a consequence of modelling a *smooth* stellar wind (Vink et al. 2001) whereas there is observational evidence that the winds are not homogeneous, but contain structure (Crowther 2007; Puls et al. 2008 and references therein). Moreover, theoretical hydrodynamical models predict instabilities from the radiation driving mechanism (Owocki and Rybicki 1984; Dessart and Owocki 2005). In addition, there are discrepancies within different observational methods of mass loss rates, as demonstrated by Figure 5.5, which shows the difference between UV P^{4+} and $H\alpha$ / radio (ρ^2) methods.

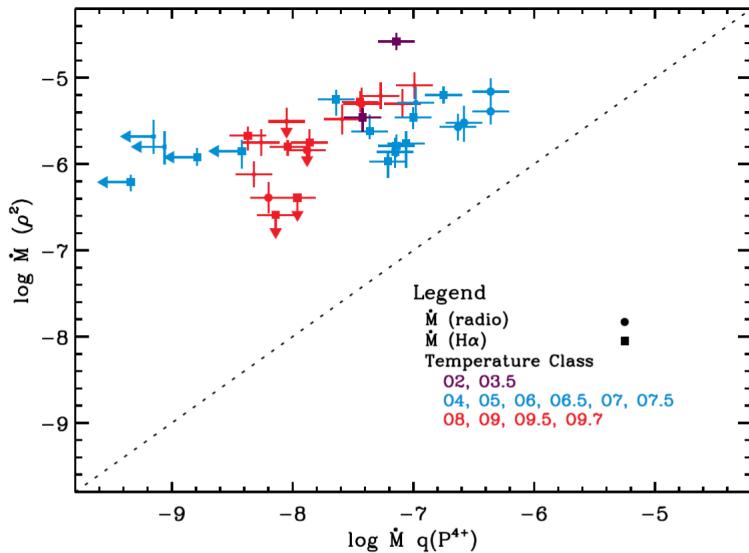


Figure 5.5: Discrepancies between two mass loss methods: UV P^{4+} and ρ^2 for a range of O stars. Figure from Fullerton et al. (2006).

The physical mechanisms of UV P^{4+} and the so called ρ^2 methods, provides insight into why such differences occur. The wind profile of UV resonance lines are insensitive to micro-clumping (see Section 5.4.2). This is because the analysis of P-Cygni profiles of any ion species involves the determination of the optical depth or *column density* of all the material associated with any particular ion (and are integral quantities), are therefore not sensitive to the distribution of material along the line of sight (Fullerton et al. 2006).

This holds true for optically thin clumps but becomes more complicated for optically thick clumps.

Optically thick clumps (large scale clumping) constitute porous winds, where ‘porosity’ is defined as a wind with localised clumps of increased density and large gaps or channels of low density material where photons are more readily transmitted (Owocki et al. 2004). For P Cygni profiles in porous winds, the clumps only cover a fraction of the solid angle surrounding the star affecting the optical depth of unsaturated UV resonance lines and thus underestimate the mass loss rate (Massa et al. 2003; Prinja and Massa 2013).

$H\alpha$ and radio free-free continuum mass loss rates are examples of ρ^2 diagnostics. The emission from these two methods arises from the interaction of two particles (recombination for $H\alpha$, Bremsstrahlung for radio, see Section 5.3 for details). Therefore the increase in local density will result in emission from these processes to be produced more strongly. Additional complications for $H\alpha$ measurements originate from the type of atmospheric model adopted (Fullerton et al. 2006), as well as the high frequency of variability in $H\alpha$ (Markova et al. 2005).

Considering these two methods gives rise to three possibilities. Either mass loss rates from UV P^{4+} resonance lines are systematically *underestimated*; mass loss rates from ρ^2 methods ($H\alpha$ and radio free-free processes) are systematically *overestimated*; or a combination of the two is affecting both measurements by different amounts. The last scenario is possible due to the fact that there are further discrepancies in the mass loss rates between $H\alpha$ and radio free-free emission (Blomme et al. 2003; Fullerton et al. 2006; Puls et al. 2006).

The combination of both scenarios is further explained by investigating the optical depth for free-free emission as a function of wavelengths. The dependence is proportional to λ^2 and results in observations at different wavelengths probing different regions of the stellar wind. Approximately, for an O star this equates to: $H\alpha$ within $1.5 R_*$, IR up to a few R_* , millimetre from $\sim 10 R_*$ and radio (cm) $\sim 100 R_*$. This dependence is illustrated in Figure 5.6.

It is unlikely that the amount of clumping is constant throughout the wind, but that a variable clumping gradient exists (Puls et al. 2006; Blomme 2011), which would explain the observed discrepancies in mass loss rates. Evidence for this gradient can be seen from flux observations of ϵ Ori from optical to radio wavelengths in Figure 5.7. There are deviations from the smooth model (dashed line in Figure 5.7) in the $H\alpha$ / IR regions which

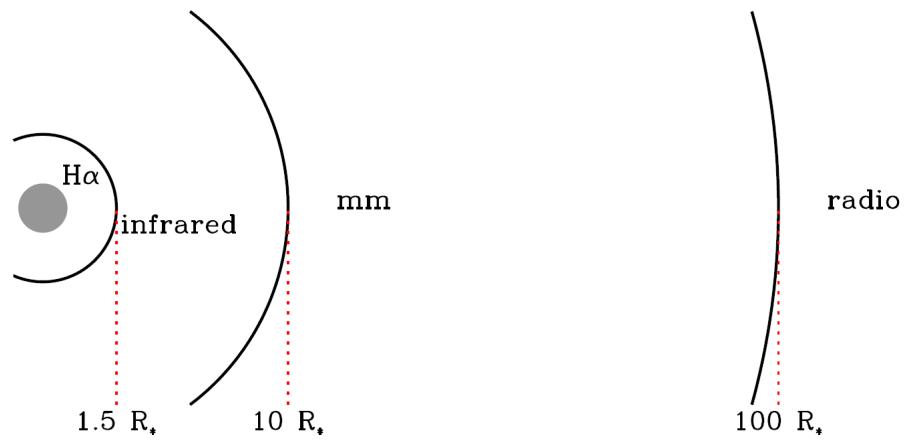


Figure 5.6: The formation regions of density squared free-free emission processes, for a typical O star. Figure taken from Blomme et al. (2002).

increases dramatically going into the millimetre regions and then declines into the radio cm wavelength region.

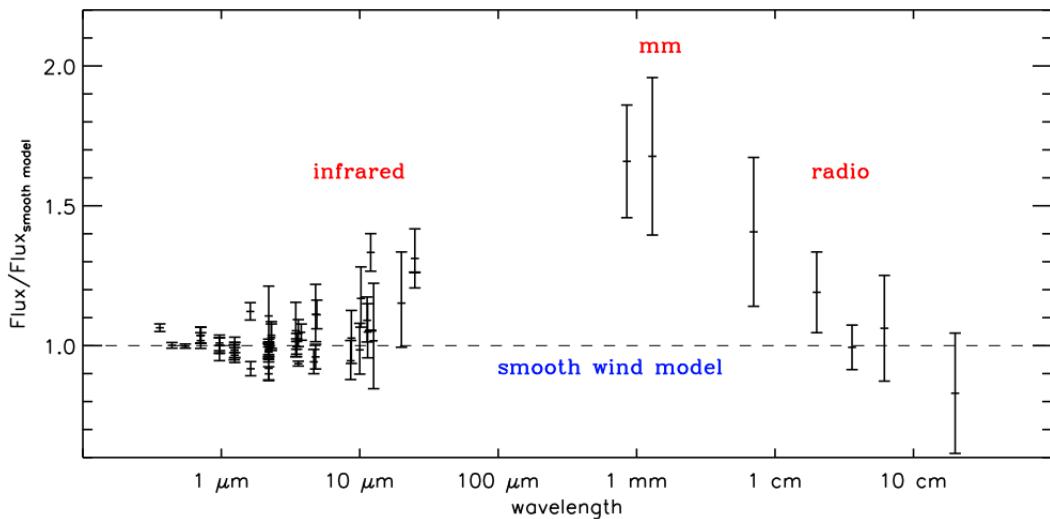


Figure 5.7: Visual to radio fluxes for ϵ Ori normalised to the expected fluxes with a smooth wind model. The greater flux excess, particularly at millimetre wavelengths indicates a higher amount of clumping in the region. Figure taken from Blomme (2011).

It is clear that clumping in the winds of *all* massive stars is present and that “stochastic wind clumping is a universal phenomenon in the radiation-driven, hot winds from all massive stars” (Lépine and Moffat 2008; Puls et al. 2008), and it significantly affects the mass loss rates of ρ^2 methods by varying degrees.

5.4.2 The Clumping Factor

One approach to quantifying the amount of clumping in the wind involves using a *clumping factor* f_{cl} , which may differ for different radii. A simple interpretation of the clumping scenario is to assume that the wind has two constituent parts; dense optically thin clumps and rarified (empty) interclump material (Puls et al. 2006; Blomme 2007; Puls et al. 2008). This definition gives rise to

$$f_{\text{cl}} = \frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2}, \quad (5.22)$$

where ρ is the mass density in the wind and the $\langle \rangle$ brackets denote an average over some small wind volume. The clumping factor is sometimes referred to as a volume filling factor f in the literature and has the relation

$$f = \frac{1}{f_{\text{cl}}}. \quad (5.23)$$

The clumping regime described above is also referred to as *micro-clumping*, due to the optically thin nature of the clumps at all wavelengths. In the scenario that the clumps are optically thick (clumping on large scales), or *macro-clumping*, the effects of porosity are somewhat different. The physical difference between the two regimes is the mean free path of interacting photons. Optically thin clumps are defined as clumps with sizes smaller than the photon mean free path of the matter-light interaction i.e. only interacting once before being scattered, whereas for optically thick clumps, photons interact many times before being destroyed or scattered into the interclump material (Puls et al. 2008).

Photons are able to pass through gaps in the wind unattenuated, and optically thick clumps can result in other material which lies behind these clumps becoming unimportant to the radiative transfer procedure, because they are hidden from incident photons (Austin 2011). The optical depth of spectral lines differs between species and therefore spatial porosity affects some lines by different amounts. Models including macro-clumping can match observed UV P⁴⁺ profiles without lowering the mass loss rate, whereas models without macro-clumping empirically underestimate mass loss rates (Oskinova et al. 2007). Macro-clumping does not affect H α empirical mass loss measurements where the line is optically thin.

In summary, *micro-clumping* reduces the empirical mass loss rates from ρ^2 diagnostics,

counteracting the *overestimated* rates from those methods. *Macro-clumping* increases the empirical mass loss rates from UV resonance lines (such as P^{4+}) counteracting the *underestimated* rates. A combination of micro- and macro-clumping can also resolve the discrepancies between resonance lines and ρ^2 methods (Oskinova et al. 2007; Hamann et al. 2008).

5.4.3 Radio Free-Free Emission with the Clumping Factor

The clumping factor f_{cl} modifies the mass loss rate \dot{M} by

$$\dot{M}_{\text{clumped}} = \dot{M} \sqrt{f_{\text{cl}}}, \quad (5.24)$$

for ρ^2 processes. For smooth winds (no clumping) $f_{\text{cl}} = 1$ and for clumped winds $f_{\text{cl}} > 1$, where larger clumping factors increases the emission for a given mass loss rate. Substituting Equation 5.24 into the radio free-free emission relation (Equation 5.21) gives

$$S_{\nu}^{\text{cl}} = 23.2 d^{-2} \left(\frac{\dot{M} \sqrt{f_{\text{cl}}}}{v_{\infty}} \right)^{\frac{4}{3}} \left(\frac{\mu_i^2}{Z^2 \gamma_e} \right)^{-\frac{2}{3}} (\nu g_{\nu})^{\frac{2}{3}} \quad [\text{Jy}]. \quad (5.25)$$

5.5 Predicted Smooth Mass Loss Rates and Radio Fluxes of Massive Stars in Cyg OB2

This section provides the ground work for predicting the radio fluxes (at L-band and C-band) of the massive stars in Cyg OB2. These fluxes are generated from models assuming a smooth wind, with the effects from different clumping factors presented in Section 5.6.

5.5.1 Theoretical Mass Loss Relationships

A grid of theoretical mass loss rates for Galactic OB stars was created by Vink et al. (2000), considering the effects of multiple scatterings, where a photon is able to be scattered more than once, in radiative line-driven winds. The authors generate mass loss formulae as a function of five stellar parameters; stellar mass, luminosity, effective temperature, terminal velocity and escape velocity, from the interpolation of multiple linear regression fits. They derive formulae for effective temperature ranges between 50,000 and 12,500 K with two distinct temperature groups; 50,000 to 27,500 K and 22,500 to 12,500 K.

For effective temperatures between these two regions (22,500 to 27,500 K) the authors state that either relation can be used. This region is where the bi-stability jump is located, where mass loss increases dramatically with decreasing effective temperature (Vink et al. 2000). This is a consequence of the ionisation balance of iron which dominates the line-driving at the base of the wind (Vink et al. 1999). For effective temperatures below 25,000 K, Fe^{3+} recombines to Fe^{2+} which is more efficient in accelerating the wind. This has been seen in UV observations of OB stars (Lamers et al. 1995).

An additional parameter considering the metallicity Z was introduced into the mass loss formulae by Vink et al. (2001) resulting in a total of six stellar parameters describing the mass loss rate. This last parameter Z is assumed to equal solar metallicity in the present work because many of the stellar parameters described below are assumed to be in solar units.

The mass loss rate for the effective temperature range $27,500 \leq T_{\text{eff}} \leq 50,000$ K, given by Vink et al. (2000) is:

$$\begin{aligned}
 \log \dot{M} = & -6.697 (\pm 0.061) \\
 & + 2.194 (\pm 0.021) \log (L_*/10^5) \\
 & - 1.313 (\pm 0.046) \log (M_*/30) \\
 & - 1.226 (\pm 0.037) \log \left(\frac{v_{\infty}/v_{\text{esc}}}{2.0} \right) \\
 & + 0.933 (\pm 0.064) \log (T_{\text{eff}}/40000) \\
 & - 10.92 (\pm 0.90) \{\log (T_{\text{eff}}/40000)\}^2 \\
 & + 0.85 (\pm 0.10) \log (Z/Z_{\odot}),
 \end{aligned} \tag{5.26}$$

where \dot{M} is in units of $M_{\odot} \text{ yr}^{-1}$, L_* is in solar units, M_* is the stellar mass without electron scattering correction also in solar units, and T_{eff} is in Kelvin K.

The mass loss rate for the effective temperature range $12,500 \leq T_{\text{eff}} \leq 22,500$ K, given

by Vink et al. (2000) is:

$$\begin{aligned}
 \log \dot{M} = & -6.688 (\pm 0.080) \\
 & + 2.210 (\pm 0.031) \log (L_*/10^5) \\
 & - 1.339 (\pm 0.068) \log (M_*/30) \\
 & - 1.601 (\pm 0.055) \log \left(\frac{v_\infty / v_{\text{esc}}}{2.0} \right) \\
 & + 1.07 (\pm 0.10) \log (T_{\text{eff}}/20000) \\
 & + 0.85 (\pm 0.10) \log (Z/Z_\odot),
 \end{aligned} \tag{5.27}$$

where the units are the same as for Equation 5.26.

For this study the mass loss rates for O and B0 stars are determined by the hotter relation (Equation 5.26), and stars with spectral types B0.2 - B5 are determined by the cooler relation (Equation 5.27). This rule designates a relation for all stars, including those whose effective temperature is approximately around where the bi-stability jump occurs.

5.5.2 Theoretical Stellar Parameters for Massive Stars

The five stellar parameters (M_* , L_* , T_{eff} , v_∞ and v_{esc}) in Equations 5.26 and 5.27 are selected according to the spectral classification of the star. Theoretical values for each are taken from a number of studies in the literature.

For O stars, Martins et al. (2005) provide the M_* , L_* , T_{eff} , R stellar parameters for supergiants, giants and dwarf luminosity classes (I, III, V). The authors compile a theoretical table of parameters from interpolating in a grid of non-local thermodynamic equilibrium (LTE) spherically extended line-blanketed models computed with the code CMFGEN (Hillier and Miller 1998). There are no parameters for bright giants (II) and sub giants (IV), therefore the calculations for bright-giants assume the parameters of supergiants and the sub-giants assume the parameters of giants. In addition, the stellar radius R is also provided for the determination of the escape velocity (v_{esc} ; see below). Table 5.1 shows the stellar parameters for O stars from Martins et al. (2005).

Table 5.1: Stellar Parameters for O stars from Martins et al. (2005)

Spectral Type	T _{eff} K	log L _* L _⊙	R R _⊙	M _{spec} M _⊙	T _{eff} K	log L _* L _⊙	R R _⊙	M _{spec} M _⊙	T _{eff} K	log L _* L _⊙	R R _⊙	M _{spec} M _⊙
Dwarfs (V)	Giants (III)	Supergiants (I)										
O3	44616	5.83	13.84	58.34	42942	5.92	16.57	58.62	42551	6.00	18.47	66.89
O4	43419	5.68	12.31	46.16	41486	5.82	15.83	48.80	40702	5.94	18.91	58.03
O5	41540	5.51	11.08	37.28	39507	5.70	15.26	41.48	38520	5.87	19.48	50.87
O5.5	40062	5.41	10.61	34.17	38003	5.63	15.13	38.92	37070	5.82	19.92	48.29
O6	38151	5.30	10.23	31.73	36673	5.56	14.97	36.38	35747	5.78	20.33	45.78
O6.5	36826	5.20	9.79	29.02	35644	5.49	14.74	33.68	34654	5.74	20.68	43.10
O7	35531	5.10	9.37	26.52	34638	5.43	14.51	31.17	33326	5.69	21.14	40.91
O7.5	34419	5.00	8.94	24.15	33487	5.36	14.34	29.06	31913	5.64	21.69	39.17
O8	33383	4.90	8.52	21.95	32573	5.30	14.11	26.89	31009	5.60	22.03	36.77
O8.5	32522	4.82	8.11	19.82	31689	5.24	13.88	24.84	30504	5.58	22.20	33.90
O9	31524	4.72	7.73	18.03	30737	5.17	13.69	23.07	29569	5.54	22.60	31.95
O9.5	30488	4.62	7.39	16.46	30231	5.12	13.37	21.04	28430	5.49	23.11	30.41

Notes. M_{spec} is the spectroscopic mass adopted from Martins et al. (2005). This table is an excerpt of their tables 1, 2, and 3.

The authors compare their results with a previous study by Vacca et al. (1996) and found that the effective temperature of all luminosity classes studied are cooler by 2000 to 8000 K, and luminosities for all luminosity classes are also reduced.

For the B supergiants (Ia, Ib), Searle et al. (2008) also utilise CMFGEN (Hillier and Miller 1998) and results from Crowther et al. (2006) to generate the fundamental stellar parameters (M_* , L_* , T_{eff} , R) for classes B0 to B5. They also find reduced luminosities compared to previous studies. Table 5.2 shows the stellar parameters for B supergiants from Searle et al. (2008).

Table 5.2: Stellar Parameters for B Supergiants from Searle et al. (2008)

Spectral Type	T_{eff} K	$\log L$ L_{\odot}	R R_{\odot}	M_* M_{\odot}
B0 Ia	28.1	5.60	26.9	25
B0 Ib	29.7	5.66	23.8	37
B0.2 Ia	26.7	5.62	30.4	36
B0.2 Ib	28.5	5.65	27.8	49
B0.5 Ia	24.7	5.58	33.8	33
B0.5 Ib	25.4	5.58	32.2	47
B0.7 Ia	23.6	5.53	35.1	23
B0.7 Ib	24.4	5.51	33.9	37
B1 Ia	22.0	5.44	36.5	12
B1 Ib	21.7	5.38	34.9	22
B1.5 Ia	19.9	5.44	44.5	18
B1.5 Ib	19.3	5.29	39.7	19
B2 Ia	18.3	5.41	51.0	19
B2 Ib	18.1	5.27	44.4	21
B2.5 Ia	17.2	5.39	56.5	19
B2.5 Ib	17.6	5.25	46.2	22
B3 Ia	16.4	5.37	60.4	19
B3 Ib	17.5	5.23	45.5	19
B4 Ia	15.8	5.34	63.5	16
B4 Ib	17.4	5.18	43.2	13
B5 Ia	15.7	5.33	63.0	15
B5 Ib	15.2	5.09	51.7	13

Notes. M_* is the spectroscopic mass adopted from Searle et al. (2008). This table is an excerpt of their table 5.

The terminal velocities (v_{∞}) of O stars and B supergiants are given by Prinja et al. (1990). The authors characterised 181 O stars, 70 early B supergiants and 35 WR stars using the central velocity of narrow absorption features and the violet limit of zero residual

intensity in saturated UV P Cygni profiles. They then generate the mean and range of v_∞ for O stars with luminosity classes I, III and V and BI. They ignore “f” qualifiers and extreme BI⁺ supergiants. Table 5.3 shows the terminal velocities for O stars and B supergiants from Prinja et al. (1990).

Table 5.3: Terminal Velocities for O stars and B Supergiants from Prinja et al. (1990)

Spectral Type	\bar{v}_∞ (km s ⁻¹)		
	Supergiants (I)	Giants (III)	Dwarfs (V)
O3	3150	3150 ^b	3190
O4	2325	2810 ^c	2950
O5	1885	2810	2875
O5.5	1885 ^a	2810 ^c	1960
O6	2300	2560	2570
O6.5	2180	2545	2455
O7	2055	2600	2295
O7.5	1980	2175	1975
O8	1530	2125	1755
O8.5	1955	2255	1970
O9	1990	1875	1500
O9.5	1765	1505	1500 ^e
O9.7	1735	1505 ^d	1500 ^e
B0	1535		
B0.2	1215		
B0.5	1405		
B0.7	1155		
B1	1065		
B1.5	750		
B2	790		
B2.5	490		
B3	590		

Notes. Superscripts denote no value in original table. ^a adopted O5I v_∞ ; ^b adopted O3I v_∞ ; ^c adopted O5III v_∞ ; ^d adopted O9.5III v_∞ ; ^e adopted O9V v_∞ .

For stellar classifications where no velocities are present, the nearest similar class velocity is chosen and shown in Table 5.3. For the bright-giant and sub-giant (II, IV) classes for O stars, supergiant and giant (I, III) velocities are used respectively. For B4 and B5 supergiants, the terminal velocity for B3 is adopted.

The final parameter required for the Vink et al. (2000) mass loss recipes (Equations

5.26 and 5.27) is the escape velocity, defined as

$$v_{\text{esc}} = \sqrt{\frac{2GM_*(1 - \Gamma_e)}{r}}, \quad (5.28)$$

where Γ_e is described by Equation 5.3, which reduces to:

$$\Gamma_e = 7.66 \times 10^{-5} \sigma_e \left(\frac{L}{L_\odot} \right) \left(\frac{M_\odot}{M_*} \right), \quad (5.29)$$

σ_e is the electron scattering coefficient and depends on the degree of ionisation and on the abundance of helium near the base of the wind (Lamers and Leitherer 1993). This is broken down into three regimes (in cgs units):

$$\begin{aligned} \sigma_e &= 0.34 \text{ cm}^2 \text{ g}^{-1} & \text{if } T_{\text{eff}} \geq 35,000 \text{ K}, \\ \sigma_e &= 0.32 \text{ cm}^2 \text{ g}^{-1} & \text{if } 30,000 \text{ K} \leq T_{\text{eff}} < 35,000 \text{ K}, \\ \sigma_e &= 0.31 \text{ cm}^2 \text{ g}^{-1} & \text{if } T_{\text{eff}} < 30,000 \text{ K}. \end{aligned} \quad (5.30)$$

As part of the catalogue building script in Chapter 4, Jython routines are made to calculate these five stellar parameters for each star depending on their spectral types from the known OB catalogue.

5.5.3 Smooth Mass Loss Rates and Radio Fluxes

The rest of the Jython script calculates the mass loss rates using Equations 5.26 and 5.27. To determine the smooth wind radio fluxes at 1.5 GHz and 5 GHz, Equation 5.21 is used with the following values for variables; \dot{M} from Equations 5.26 and 5.27 (Vink et al. 2000), v_∞ from Table 5.3 (Prinja et al. 1990), the distance d as 1.4 kpc (Rygl et al. 2012), the mean atomic mass $\mu_i = 1.3$ (90% hydrogen and 10% helium; Scuderi et al. 1998), number of free electrons $\gamma_e = 1.1$ (assuming both species completely ionised: $0.9(1/1) + 0.1(2/1)$), the charge state of atoms $Z = 1.13$, electron temperature $T_e = 0.85 \times T_{\text{eff}}$ (Scuderi et al. 1998), and ν is the central bandwidth frequency of the observation.

Figure 5.8 shows the 1.5 GHz and 5 GHz fluxes as a function of T_{eff} . Each luminosity class is identifiable with different symbols, with the O stars in red and B supergiants in blue.

These fluxes, as with the mass loss rates calculated above, are predictions from a

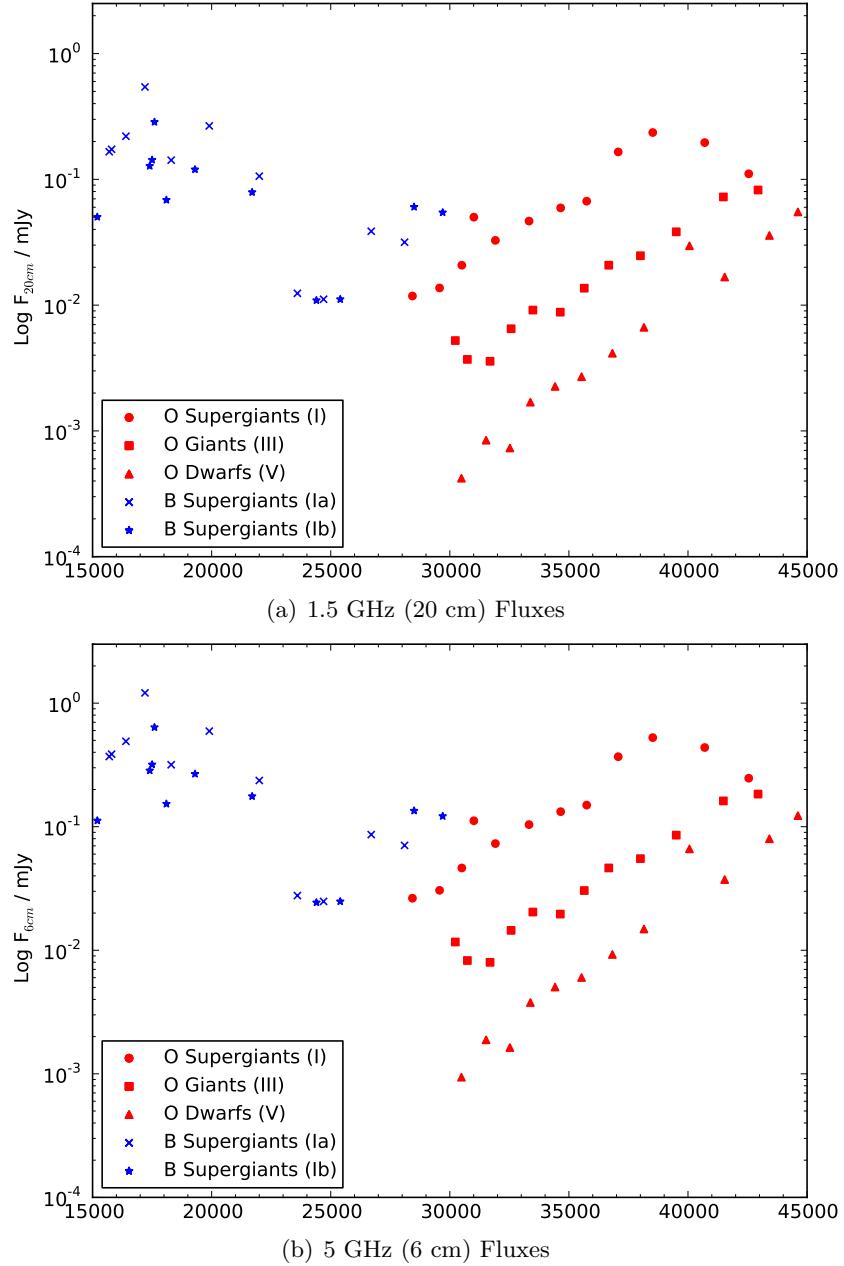


Figure 5.8: Smooth 1.5 GHz and 5 GHz (20 cm and 6 cm) fluxes of O stars and early B Supergiants. O stars are in red; circles (supergiants I), squares (giants III) and triangles (dwarfs V) and B supergiants are in blue; crosses (Ia) and stars (Ib).

smooth wind model. If clumping is present in the radio-opacity region of the winds, it is expected that the observed radio fluxes at these frequencies will be higher than the fluxes determined here. Furthermore, it is expected that the observed 5 GHz fluxes will be larger than the observed 1.5 GHz fluxes because of the thermal spectral index of massive star winds.

It is important to note that these fluxes are predictions for the wind component of single stars, and do not consider any influence from binary effects. Binaries emit both thermal radiation from the winds of the single star components, and non-thermal radiation from the interaction region between both stars. Depending on the line of sight and inclination of the binary, the amount of thermal radiation (which is calculated here) will change as parts of the wind may become obscured by the interaction region (and vice versa). Moreover, the nature of the interaction region will have a significant effect on the thermal emission. For example, the interaction region of two colliding clumpy winds, destroys the clumps which smooths out the structure in the interaction region (Pittard 2007). Therefore whilst it is possible to compare these fluxes with thermal fluxes from known binaries, the results must be interpreted with caution.

The predicted fluxes indicate that most OI and BI, half of the OIII and early OV are expected to be seen in 1.5 GHz COBRaS maps, assuming an image rms $\sim 11 \mu\text{Jy}$ (for ~ 5 hours on each target field). This increases for 5 GHz where the thermal nature of the winds increases the fluxes. For the 5 GHz maps roughly all of the OI, OIII, and BI stars and early OV stars are expected to be visible with an image rms $\sim 5 \mu\text{Jy}$ (for ~ 5 hours on each target field). These results are subject to large errors arising from the large uncertainties on some of the individual stellar parameters (e.g. $\Delta M = 50\%$; Martins et al. 2005) which dramatically changes the predicted flux.

Tables 5.4, 5.5, and 5.6 give the full set of parameters used to determine the 1.5 GHz and 5 GHz radio fluxes for O supergiants and giants, O dwarfs, and B supergiants respectively.

Table 5.4: Full Table of Smooth Mass Loss Rates and Radio Fluxes for O Supergiants and Giants

ST	T_{eff} K	L $\log L/L_{\odot}$	M M_{\odot}	R R_{\odot}	v_{∞} km s^{-1}	Γ_e	v_{esc} km s^{-1}	$\log \bar{M}$ $M_{\odot} \text{ yr}^{-1}$	\bar{M} $M_{\odot} \text{ yr}^{-1}$	$\Delta \bar{M}$ $M_{\odot} \text{ yr}^{-1}$	$S_{6\text{cm}}$ mJy	$\Delta S_{6\text{cm}}$ mJy	$S_{20\text{cm}}$ mJy	$\Delta S_{20\text{cm}}$ mJy
O3I	42551	6.0	66.89	18.47	3150	0.389	918	-5.230	5.888E-06	1.212E-06	0.2469	0.0685	0.1107	0.0307
O4I	40702	5.94	58.03	18.91	2325	0.391	844	-5.175	6.691E-06	1.394E-06	0.4373	0.1303	0.1960	0.0584
O5I	38820	5.87	50.87	19.48	1885	0.380	786	-5.204	6.252E-06	1.320E-06	0.5258	0.1683	0.2356	0.0754
O5.5I	37070	5.82	48.29	19.92	1885	0.356	771	-5.319	4.802E-06	1.016E-06	0.3686	0.1181	0.1652	0.0529
O6I	35747	5.78	45.78	20.33	2300	0.343	751	-5.525	2.987E-06	6.275E-07	0.1496	0.0449	0.0671	0.0201
O6.5I	34654	5.74	43.1	20.68	2180	0.313	739	-5.587	2.587E-06	5.459E-07	0.1323	0.0404	0.0593	0.0181
O7I	33326	5.69	40.91	21.14	2055	0.293	722	-5.690	2.041E-06	4.331E-07	0.1040	0.0324	0.0466	0.0145
O7.5I	31913	5.64	39.17	21.69	1980	0.273	707	-5.820	1.513E-06	3.228E-07	0.0730	0.0231	0.0327	0.0103
O8I	31009	5.6	36.77	22.03	1530	0.265	684	-5.793	1.611E-06	3.509E-07	0.1117	0.0394	0.0501	0.0176
O8.5I	30504	5.58	33.9	22.2	1955	0.275	650	-5.972	1.065E-06	2.284E-07	0.0463	0.0147	0.0208	0.0066
O9I	29569	5.54	31.95	22.6	1990	0.258	632	-6.100	7.952E-07	1.708E-07	0.0305	0.0097	0.0137	0.0043
O9.5I	28430	5.49	30.41	23.11	1765	0.241	617	-6.198	6.332E-07	1.373E-07	0.0264	0.0088	0.0118	0.0039
O3III	42942	5.92	58.62	16.57	3150	0.370	922	-5.327	4.711E-06	9.679E-07	0.1836	0.0508	0.0823	0.0228
O4III	41486	5.82	48.8	15.83	2810	0.353	872	-5.417	3.830E-06	7.904E-07	0.1617	0.0459	0.0725	0.0205
O5III	39507	5.7	41.48	15.26	2810	0.315	843	-5.623	2.381E-06	4.928E-07	0.0854	0.0243	0.0383	0.0109
O5.5III	38003	5.63	38.92	15.13	2810	0.285	837	-5.765	1.719E-06	3.565E-07	0.0551	0.0157	0.0247	0.0070
O6III	36673	5.56	36.38	14.97	2560	0.260	828	-5.861	1.378E-06	2.874E-07	0.0464	0.0135	0.0208	0.0060
O6.5III	35644	5.49	33.68	14.74	2545	0.239	814	-5.999	1.001E-06	2.092E-07	0.0304	0.0089	0.0136	0.0040
O7III	34638	5.43	31.17	14.51	2600	0.212	803	-6.132	7.375E-07	1.543E-07	0.0196	0.0057	0.0088	0.0026
O7.5III	33487	5.36	29.06	14.34	2175	0.193	789	-6.196	6.364E-07	1.344E-07	0.0204	0.0062	0.0091	0.0028
O8III	32573	5.3	26.89	14.11	2125	0.182	771	-6.317	4.820E-07	1.021E-07	0.0145	0.0045	0.0065	0.0020
O8.5III	31689	5.24	24.84	13.88	2255	0.171	752	-6.484	3.278E-07	6.942E-08	0.0080	0.0024	0.0036	0.0011
O9III	30737	5.17	23.07	13.69	1875	0.157	736	-6.553	2.802E-07	6.005E-08	0.0083	0.0027	0.0037	0.0012
O9.5III	30231	5.12	21.04	13.37	1505	0.154	713	-6.535	2.917E-07	6.365E-08	0.0117	0.0041	0.0052	0.0019

Table 5.5: Full Table of Smooth Mass Loss Rates and Radio Fluxes for O Dwarfs

ST	T _{eff} K	L Log L / L _⊙	M M _⊙	R R _⊙	v _∞ km s ⁻¹	Γ _e	v _{esc} km s ⁻¹	Log M̄	M̄ M _⊙ yr ⁻¹	ΔM̄ M _⊙ yr ⁻¹	S _{6cm} mJy	Δ S _{6cm} mJy	S _{20cm} mJy	Δ S _{20cm} mJy
O3V	44616	5.83	58.34	13.84	3190	0.302	1059	-5.453	3.522E-06	7.230E-07	0.1229	0.0339	0.0551	0.0152
O4V	43419	5.68	46.16	12.31	2950	0.270	1022	-5.627	2.362E-06	4.864E-07	0.0799	0.0224	0.0358	0.0100
O5V	41540	5.51	37.28	11.08	2875	0.226	996	-5.884	1.305E-06	2.695E-07	0.0373	0.0105	0.0167	0.0047
O5.5V	40062	5.41	34.17	10.61	1960	0.196	994	-5.863	1.370E-06	2.878E-07	0.0661	0.0208	0.0296	0.0093
O6V	38151	5.3	31.73	10.23	2570	0.164	994	-6.231	5.878E-07	1.224E-07	0.0149	0.0043	0.0067	0.0019
O6.5V	36826	5.2	29.02	9.79	2455	0.142	985	-6.404	3.945E-07	8.248E-08	0.0092	0.0027	0.0041	0.0012
O7V	35531	5.1	26.52	9.37	2295	0.124	972	-6.572	2.679E-07	5.629E-08	0.0060	0.0018	0.0027	0.0008
O7.5V	34419	5.0	24.15	8.94	1975	0.101	962	-6.694	2.021E-07	4.288E-08	0.0050	0.0016	0.0023	0.0007
O8V	33383	4.9	21.95	8.52	1755	0.089	946	-6.839	1.450E-07	3.107E-08	0.0038	0.0012	0.0017	0.0006
O8.5V	32522	4.82	19.82	8.11	1970	0.082	925	-7.061	8.693E-08	1.853E-08	0.0016	0.0005	0.0007	0.0002
O9V	31524	4.72	18.03	7.73	1500	0.071	909	-7.132	7.383E-08	1.609E-08	0.0019	0.0007	0.0008	0.0003
O9.5V	30488	4.62	16.46	7.39	1500	0.062	892	-7.357	4.391E-08	9.597E-09	0.0009	0.0003	0.0004	0.0001

Table 5.6: Full Table of Smooth Mass Loss Rates and Radio Fluxes for B Supergiants

ST	T _{eff} K	L $\log L/L_\odot$	M M_\odot	R R_\odot	v_∞ km s^{-1}	Γ_e	v_{esc} km s^{-1}	$\log \dot{M}$ $M_\odot \text{yr}^{-1}$	\dot{M} $M_\odot \text{yr}^{-1}$	$\Delta \dot{M}$ $M_\odot \text{yr}^{-1}$	$S_{6\text{cm}}$ mJy	$\Delta S_{6\text{cm}}$ mJy	$S_{20\text{cm}}$ mJy	$\Delta S_{20\text{cm}}$ mJy
B0Ia	28100	5.6	25.0	26.9	1535	0.378	469	-5.938	1.153E-06	2.612E-07	0.0706	0.0252	0.0316	0.0113
B0.2Ia	26700	5.62	36.0	30.4	1215	0.275	572	-5.973	1.065E-06	2.320E-07	0.0863	0.0342	0.0387	0.0153
B0.5Ia	24700	5.58	33.0	33.8	1405	0.274	520	-6.313	4.867E-07	1.064E-07	0.0249	0.0091	0.0111	0.0041
B0.7Ia	23600	5.53	23.0	35.1	1155	0.350	403	-6.361	4.354E-07	1.063E-07	0.0277	0.0117	0.0124	0.0052
B1Ia	22000	5.44	12.0	36.5	1065	0.545	238	-5.696	2.015E-06	9.064E-07	0.2366	0.1375	0.1060	0.0616
B1.5Ia	19900	5.44	18.0	44.5	750	0.363	313	-5.545	2.848E-06	9.135E-07	0.5935	0.3535	0.2660	0.1584
B2Ia	18300	5.41	19.0	51.0	790	0.321	310	-5.725	1.886E-06	5.715E-07	0.3171	0.1796	0.1421	0.0805
B2.5Ia	17200	5.39	19.0	56.5	490	0.307	298	-5.494	3.208E-06	1.174E-06	1.2109	1.0095	0.5427	0.4524
B3Ia	16400	5.37	19.0	60.4	590	0.293	291	-5.706	1.970E-06	6.600E-07	0.4912	0.3494	0.2201	0.1566
B4Ia	15800	5.34	16.0	63.5	590	0.325	254	-5.782	1.651E-06	6.350E-07	0.3868	0.2847	0.1733	0.1276
B5Ia	15700	5.33	15.0	63.0	590	0.338	245	-5.797	1.597E-06	6.495E-07	0.3699	0.2766	0.1658	0.1239
B0Ib	29700	5.66	37.0	23.8	1535	0.293	647	-5.762	1.729E-06	3.616E-07	0.1217	0.0422	0.0545	0.0189
B0.2Ib	28500	5.65	49.0	27.8	1215	0.216	726	-5.830	1.480E-06	3.107E-07	0.1346	0.0527	0.0603	0.0236
B0.5Ib	25400	5.58	47.0	32.2	1405	0.192	670	-6.314	4.855E-07	1.018E-07	0.0248	0.0090	0.0111	0.0040
B0.7Ib	24400	5.51	37.0	33.9	1155	0.208	574	-6.404	3.945E-07	8.722E-08	0.0244	0.0100	0.0109	0.0045
B1Ib	21700	5.38	22.0	34.9	1065	0.259	422	-5.791	1.617E-06	4.066E-07	0.1762	0.0786	0.0790	0.0352
B1.5Ib	19300	5.29	19.0	39.7	750	0.244	371	-5.804	1.570E-06	4.820E-07	0.2674	0.1573	0.1198	0.0705
B2Ib	18100	5.27	21.0	44.4	790	0.211	377	-5.962	1.092E-06	3.058E-07	0.1529	0.0848	0.0685	0.0380
B2.5Ib	17600	5.25	22.0	46.2	490	0.192	383	-5.703	1.979E-06	6.582E-07	0.6373	0.5223	0.2856	0.2341
B3Ib	17500	5.23	19.0	45.5	590	0.212	354	-5.849	1.417E-06	4.740E-07	0.3184	0.2264	0.1427	0.1015
B4Ib	17400	5.18	13.0	43.2	590	0.276	288	-5.885	1.304E-06	6.036E-07	0.2850	0.2222	0.1277	0.0996
B5Ib	15200	5.09	13.0	51.7	590	0.225	272	-6.185	6.534E-07	3.035E-07	0.1120	0.0875	0.0502	0.0392

Both the mass loss rates and fluxes have associated errors, propagated from the errors on the individual stellar parameters from their respective studies (Prinja et al. 1990; Vink et al. 2000; Martins et al. 2005; Searle et al. 2008). Full details on these errors is given in Appendix B.

5.5.4 Comparison of Smooth Mass Loss Rates and Fluxes with the Literature

The work on the smooth mass loss rate and flux predictions was completed in late 2010. Since that time, Muijres et al. (2012) have published similar studies on the predictions of mass loss rates and wind velocities of O stars (I, III and V luminosity classes). They also use stellar parameters from Martins et al. (2005) and compare the velocities of massive stars with those from Prinja et al. (1990). However, this thesis study differs from Muijres et al. (2012) in a number of ways. This investigation also includes early B supergiants with stellar parameters from Searle et al. (2008), and goes one step further by using the calculated mass loss rates from the Vink et al. (2000) recipes to predict the radio flux at 6 cm and 20 cm (5 GHz and 1.5 GHz respectively). This is then applied to known OB stars in Cyg OB2 derived from the catalogues created in Chapter 4.

Muijres et al. (2012) generate mass loss rates with two different methods. The first termed “best- β solution”, uses a line force representation of the radiative acceleration of lines (their Equation 6) to analytically find the best solution for the velocity law in the outer part of the wind. The second method is a hydrodynamical solution which calculates a numerical solution of the equation of motion (their Equation 11) which assesses whether the wind is isothermal and the electron scattering Γ_e is constant through the wind. These are used to estimate the terminal velocities and the v_∞/v_{esc} ratio of each O star.

These methods differ from this study where the v_∞/v_{esc} ratio is determined from observations in Prinja et al. (1990) and Equation 5.28. Vink et al. (2000) also differs by estimating the ratio as a step function. Stars which qualify for the hot regime (Equation 5.26) the ratio ~ 2.6 and for the cool regime (Equation 5.27) the ratio ~ 1.3 (Lamers et al. 1995). Note, the method of Lamers et al. (1995) and Vink et al. (2000) also uses the terminal velocities of Prinja et al. (1990) to determine the step function.

Figure 5.9 shows the comparison of the O star mass loss rates from this thesis study and that of Muijres et al. (2012) (method A; (a), method B; (b)). There is a clear systematic discrepancy between the two studies for both methods A and B, whereby the mass loss

rates of Muijres et al. (2012) are underestimated relative to this study. Furthermore, the rates for late O dwarfs also appear to deviate from the general trend. This is because the terminal velocities and mass loss rates are overestimated for the weak wind cases of late O dwarfs.

The cause of these discrepancies can be explained as follows. The systematic difference in mass loss rates is explained by the aforementioned difference in calculation involving the v_∞/v_{esc} ratio. The resulting ratio from methods A and B are all larger than those from observations, and therefore, because of the inverse proportionality, underestimates the mass loss. When the same velocity ratio is used for their methods and the method of Vink et al. (2000), the mass loss rates are comparable. Muijres et al. (2012) also state that the uncertainty of terminal velocities is widely quoted as 20%, but should realistically be increased to 30-40% for empirical values, due to the uncertainties in the measurement of v_∞ from UV resonance lines.

The trend of the late O dwarf mass loss rates is due to the ionisation of iron, particularly Fe^{4+} . Similar behaviour of Fe^{3+} and Fe^{2+} ions was mentioned in Section 5.5.1 with the bi-stability jump around B0 to B1 stars. In this case, Fe^{4+} is more efficient in absorbing stellar flux (and therefore accelerating the wind) than Fe^{3+} . This leads to the methods of Muijres et al. (2012) to not find physical solutions for method B, and method A mass loss rates heavily caveated for late O dwarfs. They further state that supergiants and giants (other than O7III) do not suffer from this problem. The authors conclude that another, unknown mechanism is contributing to the acceleration at the base of the wind for late O dwarfs, suggesting magnetic pressure, turbulence or pulsations as possible explanations.

5.6 Clumped Wind Radio Fluxes for Massive Stars in Cyg OB2

Clumping is likely differential in the winds of massive stars (see Figure 5.7) resulting in different clumping factors throughout the wind. This directly results in different clumping factors associated with different wavelength mass loss diagnostics such as $\text{H}\alpha$, mm/ sub-mm and radio (cm) observations. Hydrodynamical simulations have shown the amount of clumping rises in the inner region as the clumps are formed by line-driven instabilities and peaks in the $\text{mm} \sim 20 R_*$, where the line driving force no longer maintains the clumps which can survive to large distances $\sim 100 R_*$ (Runacres and Owocki 2002).

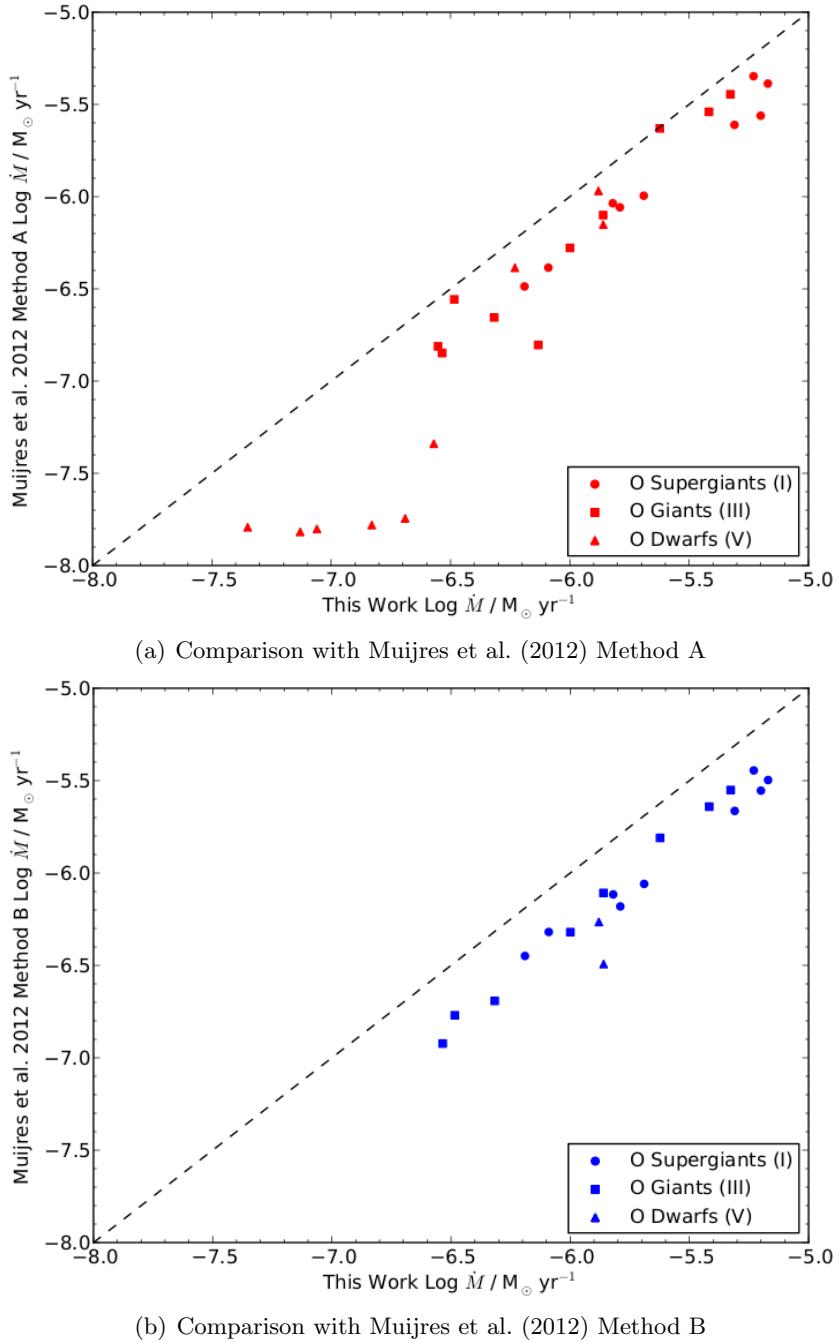


Figure 5.9: Mass loss rates comparison with Muijres et al. (2012) methods A and B for O stars with luminosity classes I, III, and V. Against both methods a systematic discrepancy exists, whereby Muijres et al. (2012) mass loss rates are underestimated in comparison to this study. This is particularly evident for late O dwarfs, where their solutions for O6.5V and later are considered non-physical solutions. This is explained further in the text.

Simpler models of clumping stratification using piece-wise linear curves throughout the wind were investigated by Blomme et al. (2002). Two models were made to fit the observations of ϵ Ori, one assumed no wind clumping in the radio regime and fits the fluxes in the mm regime well, and the other assumed clumping in the radio regime to coincide with results from Runacres and Owocki (2002). This second model required the mass loss rate to be decreased by a factor of two in order to fit the fluxes, but resulted in a clumping factor of 4 for the radio fluxes which is also consistent with those from hydrodynamical simulations (Runacres and Owocki 2002).

Other studies have investigated the clumping factors from $H\alpha$ diagnostics, whose mass loss rates are similar to those determined from radio excess (Figure 5.5). Repolust et al. (2004) detect an excess in mass loss rates from $H\alpha$ observations including two Cyg OB2 stars (No. 7 and No. 11), resulting in a clumping factor of 5.2 necessary to reduce the mass loss rate to match the wind luminosity relation (WLR). More $H\alpha$ studies find clumping factors of 5.7 (Puls et al. 2003), 10 (Hillier et al. 2003), 4.3 (Markova et al. 2004) and 4.1 ± 1.4 (Puls et al. 2006). Furthermore, Markova et al. (2004) found that for strong winds where the $H\alpha$ profile is in emission, the mass loss rates from $H\alpha$ is larger than radio determined rates by a factor of two, but for weaker winds where the $H\alpha$ profile is in absorption, the reverse is true or the rates are similar. This would suggest that the amount of clumping in the region close to the star ($H\alpha$) is more clumped than in the outer regions where radio emission occurs. However, further uncertainties exist from determining mass loss rates from $H\alpha$ absorption profiles for weak winds, meaning interpretation of these results should be met with caution.

The amount of predicted clumping from simulations of the $H\alpha$ and radio regions is similar (Runacres and Owocki 2002), and independent observations of $H\alpha$ and radio clumping factors also show similar factors (see references above). Whilst this is not always true for any given star, it is reasonable to assume for a first-order model that the clumping factors for $H\alpha$ and radio diagnostics are equivalent. A small grid of clumping factors can be drawn to encompass the common clumping factor of 5 from the majority of $H\alpha$ observations, with an upper limit of $f_{cl} = 10$ for highly clumped winds, and a lower limit of $f_{cl} = 2$ for slightly clumped winds. A clumped wind ($f_{cl} > 1$) emits more radio flux than a smooth wind ($f_{cl} = 1$) and the smooth wind fluxes are also included for reference and are equal to the radio fluxes determined in Section 5.5.3.

The 1.6 GHz (20 cm) and 5 GHz (6 cm) clumped radio fluxes determined by Equation

5.25 for the subset of clumping factors $f_{\text{cl}} = \{1, 2, 5, 10\}$ are shown in Tables 5.7, 5.8 and 5.9 for O supergiants and giants, O dwarfs, and B supergiants respectively. The errors on the clumped fluxes are determined by the errors on the smooth wind fluxes (ΔS_{ν} for $f_{\text{cl}} = 1$) multiplied by $f_{\text{cl}}^{2/3}$.

Bieging et al. (1989) conducted a VLA radio study on some OB stars in Cyg OB2 at 1.5 GHz and 5 GHz. Comparing these fluxes with the predicted radio fluxes with varying clumping factors provides early insight into whether these fluxes show signs of clumping. Figure 5.10 shows Cyg OB2 members (thermal components where possible) selected from Bieging et al. (1989), compared with the predicted fluxes of the same stars from the OB Catalogue. The dashed lines show the fluxes expected with different clumping factors.

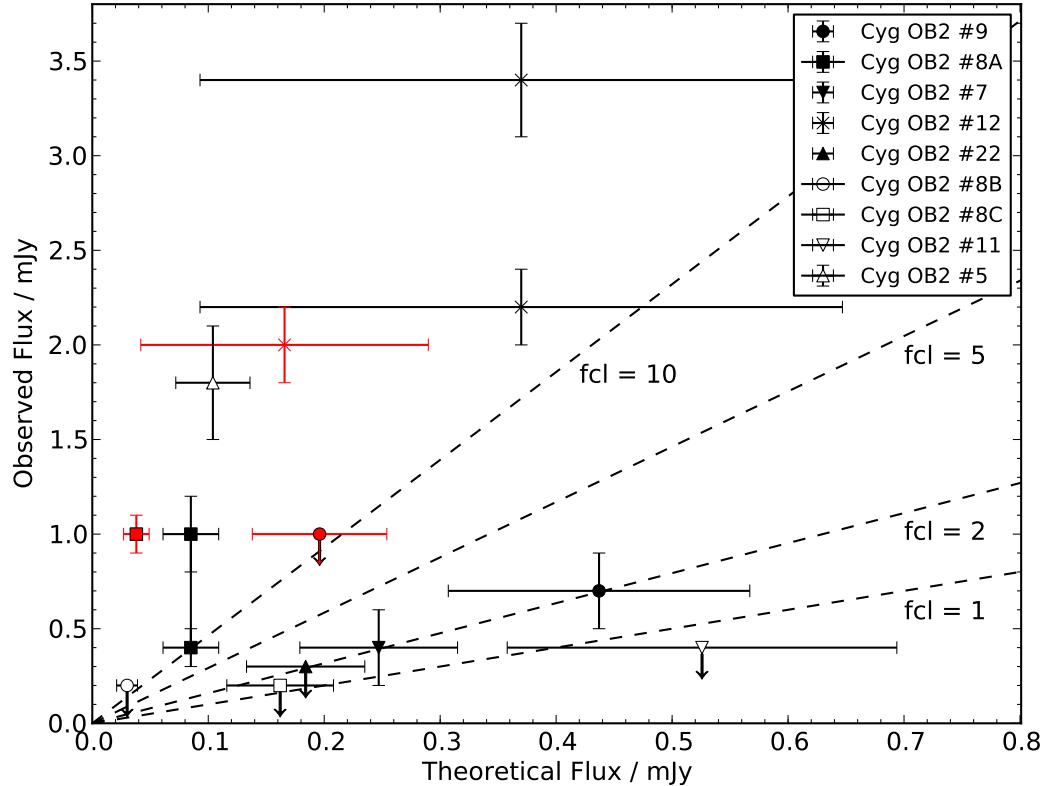


Figure 5.10: 1.5 GHz (red) and 5 GHz (black) radio flux comparisons between the observed flux from the VLA (Bieging et al. 1989), and the predicted fluxes determined in this chapter. The dashed lines show the expected fluxes for the clumping factors $f_{\text{cl}} = 1, 2, 5, 10$.

Table 5.7: Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for O Supergiants and Giants

ST	$S_{6\text{cm}}^{\text{cl}}$	$\Delta S_{6\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$								
	$f_{\text{cl}}=1$	$f_{\text{cl}}=1$	$f_{\text{cl}}=2$	$f_{\text{cl}}=2$	$f_{\text{cl}}=5$	$f_{\text{cl}}=5$	$f_{\text{cl}}=10$	$f_{\text{cl}}=10$	$f_{\text{cl}}=1$	$f_{\text{cl}}=1$	$f_{\text{cl}}=2$	$f_{\text{cl}}=2$	$f_{\text{cl}}=5$	$f_{\text{cl}}=5$	$f_{\text{cl}}=10$	
O3I	0.247	0.068	0.392	0.109	0.722	0.200	1.146	0.318	0.111	0.031	0.176	0.049	0.324	0.090	0.514	0.142
O4I	0.437	0.130	0.694	0.207	1.279	0.381	2.030	0.605	0.196	0.058	0.311	0.093	0.573	0.171	0.910	0.271
O5I	0.526	0.168	0.835	0.267	1.537	0.492	2.441	0.781	0.236	0.075	0.374	0.120	0.689	0.220	1.094	0.350
O5.5I	0.369	0.118	0.585	0.187	1.078	0.345	1.711	0.548	0.165	0.053	0.262	0.084	0.483	0.155	0.767	0.246
O6I	0.150	0.045	0.237	0.071	0.437	0.131	0.694	0.208	0.067	0.020	0.107	0.032	0.196	0.059	0.311	0.093
O6.5I	0.132	0.040	0.210	0.064	0.387	0.118	0.614	0.188	0.059	0.018	0.094	0.029	0.173	0.053	0.275	0.084
O7I	0.104	0.032	0.165	0.051	0.304	0.095	0.483	0.150	0.047	0.014	0.074	0.023	0.136	0.042	0.216	0.067
O7.5I	0.073	0.023	0.116	0.037	0.223	0.068	0.339	0.107	0.033	0.010	0.052	0.016	0.096	0.030	0.152	0.048
O8I	0.112	0.039	0.177	0.063	0.327	0.115	0.518	0.183	0.050	0.018	0.080	0.028	0.146	0.051	0.233	0.082
O8.5I	0.046	0.015	0.073	0.023	0.135	0.043	0.215	0.068	0.021	0.007	0.033	0.010	0.061	0.019	0.097	0.031
O9I	0.030	0.010	0.048	0.015	0.089	0.028	0.142	0.045	0.014	0.004	0.022	0.007	0.040	0.013	0.064	0.020
O9.5I	0.026	0.009	0.042	0.014	0.077	0.026	0.123	0.041	0.012	0.004	0.019	0.006	0.035	0.011	0.055	0.018
O3III	0.184	0.051	0.291	0.081	0.537	0.149	0.852	0.236	0.082	0.023	0.131	0.036	0.241	0.067	0.382	0.106
O4III	0.162	0.046	0.257	0.073	0.473	0.134	0.751	0.213	0.072	0.020	0.115	0.033	0.212	0.060	0.337	0.095
O5III	0.085	0.024	0.136	0.039	0.250	0.071	0.396	0.113	0.038	0.011	0.061	0.017	0.112	0.032	0.178	0.051
O5.5III	0.055	0.016	0.087	0.025	0.161	0.046	0.256	0.073	0.025	0.007	0.039	0.011	0.072	0.020	0.115	0.032
O6III	0.046	0.014	0.074	0.021	0.136	0.039	0.215	0.063	0.021	0.006	0.033	0.010	0.061	0.018	0.097	0.028
O6.5III	0.030	0.009	0.048	0.014	0.089	0.026	0.141	0.041	0.014	0.004	0.022	0.006	0.040	0.012	0.063	0.019
O7III	0.020	0.006	0.031	0.009	0.057	0.017	0.091	0.026	0.009	0.003	0.014	0.004	0.026	0.008	0.041	0.012
O7.5III	0.020	0.006	0.032	0.010	0.060	0.018	0.095	0.029	0.009	0.003	0.014	0.004	0.027	0.008	0.042	0.013
O8III	0.014	0.004	0.023	0.007	0.042	0.013	0.067	0.021	0.006	0.002	0.010	0.003	0.019	0.006	0.030	0.009
O8.5III	0.008	0.002	0.013	0.004	0.023	0.007	0.037	0.011	0.004	0.001	0.006	0.002	0.011	0.003	0.017	0.005
O9III	0.008	0.003	0.013	0.004	0.024	0.008	0.039	0.013	0.004	0.001	0.006	0.002	0.011	0.004	0.017	0.006
O9.5III	0.012	0.004	0.019	0.007	0.034	0.012	0.054	0.019	0.005	0.002	0.008	0.003	0.015	0.006	0.024	0.009

All flux values are in mJy.

Table 5.8: Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for O Dwarfs

ST	$S_{6\text{cm}}^{\text{cl}}$	$\Delta S_{6\text{cm}}^{\text{cl}}$	$S_{6\text{cm}}^{\text{cl}}$	$\Delta S_{6\text{cm}}^{\text{cl}}$	$S_{6\text{cm}}^{\text{cl}}$	$\Delta S_{6\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$	$S_{20\text{cm}}^{\text{cl}}$	$\Delta S_{20\text{cm}}^{\text{cl}}$					
	$f_{\text{cl}}=1$	$f_{\text{cl}}=1$	$f_{\text{cl}}=2$	$f_{\text{cl}}=2$	$f_{\text{cl}}=5$	$f_{\text{cl}}=5$	$f_{\text{cl}}=10$	$f_{\text{cl}}=10$	$f_{\text{cl}}=1$	$f_{\text{cl}}=2$	$f_{\text{cl}}=5$	$f_{\text{cl}}=10$	$f_{\text{cl}}=10$				
O3V	0.123	0.034	0.195	0.054	0.359	0.099	0.570	0.157	0.055	0.015	0.087	0.024	0.161	0.044	0.256	0.071	
O4V	0.080	0.022	0.127	0.036	0.234	0.065	0.371	0.104	0.036	0.010	0.057	0.016	0.105	0.029	0.166	0.046	
O5V	0.037	0.010	0.059	0.017	0.109	0.031	0.173	0.049	0.017	0.005	0.027	0.007	0.049	0.014	0.078	0.022	
O5.5V	0.066	0.021	0.105	0.033	0.193	0.061	0.307	0.097	0.030	0.009	0.047	0.015	0.087	0.027	0.137	0.043	
O6V	0.015	0.004	0.024	0.007	0.044	0.013	0.069	0.020	0.007	0.002	0.011	0.003	0.020	0.006	0.031	0.009	
O6.5V	0.009	0.003	0.015	0.004	0.027	0.008	0.043	0.013	0.004	0.001	0.007	0.002	0.012	0.004	0.019	0.006	
O7V	0.006	0.002	0.010	0.003	0.018	0.005	0.028	0.008	0.005	0.003	0.001	0.004	0.001	0.008	0.002	0.013	0.004
O7.5V	0.005	0.002	0.008	0.003	0.015	0.005	0.023	0.007	0.007	0.002	0.001	0.004	0.001	0.007	0.002	0.011	0.003
O8V	0.004	0.001	0.006	0.002	0.011	0.004	0.018	0.006	0.006	0.002	0.001	0.003	0.001	0.005	0.002	0.008	0.003
O8.5V	0.002	0.000	0.003	0.001	0.005	0.001	0.007	0.002	0.002	0.001	0.000	0.001	0.000	0.002	0.001	0.003	0.001
O9V	0.002	0.001	0.003	0.001	0.006	0.002	0.009	0.003	0.003	0.001	0.000	0.001	0.000	0.002	0.001	0.004	0.001
O9.5V	0.001	0.000	0.001	0.000	0.003	0.001	0.004	0.001	0.001	0.000	0.001	0.000	0.001	0.000	0.002	0.000	0.000

All flux values are in mJy.

Table 5.9: Full Table of 6 cm (5 GHz) and 20 cm (1.6 GHz) Clumped Fluxes for B Supergiants

ST	$S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=1$	$\Delta S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=1$	$S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=2$	$\Delta S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=2$	$S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=5$	$\Delta S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=5$	$S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=10$	$\Delta S_{6\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=10$	$S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=1$	$\Delta S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=1$	$S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=2$	$\Delta S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=2$	$S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=5$	$\Delta S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=5$	$S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=10$	$\Delta S_{20\text{cm}}^{\text{cl}}$ $f_{\text{cl}}=10$
B0Ia	0.071	0.025	0.112	0.040	0.206	0.074	0.328	0.117	0.032	0.011	0.050	0.018	0.092	0.033	0.147	0.052
B0.2Ia	0.086	0.034	0.137	0.054	0.252	0.100	0.401	0.159	0.039	0.015	0.061	0.024	0.113	0.045	0.180	0.071
B0.5Ia	0.025	0.009	0.040	0.014	0.073	0.027	0.116	0.042	0.011	0.004	0.018	0.007	0.032	0.012	0.052	0.019
B0.7Ia	0.028	0.012	0.044	0.019	0.081	0.034	0.129	0.054	0.012	0.005	0.020	0.008	0.036	0.015	0.058	0.024
B1Ia	0.237	0.138	0.376	0.218	0.692	0.402	1.098	0.638	0.106	0.062	0.168	0.098	0.310	0.180	0.492	0.286
B1.5Ia	0.594	0.354	0.942	0.561	1.735	1.034	2.755	1.641	0.266	0.158	0.422	0.251	0.778	0.463	1.235	0.735
B2Ia	0.317	0.180	0.503	0.285	0.927	0.525	1.472	0.834	0.142	0.080	0.226	0.128	0.416	0.235	0.660	0.374
B2.5Ia	1.211	1.010	1.922	1.602	3.541	2.952	5.620	4.686	0.543	0.452	0.861	0.718	1.587	1.323	2.519	2.100
B3Ia	0.491	0.349	0.780	0.555	1.436	1.022	2.280	1.622	0.220	0.157	0.349	0.249	0.644	0.458	1.022	0.727
B4Ia	0.387	0.285	0.614	0.452	1.131	0.832	1.795	1.321	0.173	0.128	0.275	0.203	0.507	0.373	0.804	0.592
B5Ia	0.370	0.277	0.587	0.439	1.082	0.809	1.717	1.284	0.166	0.124	0.263	0.197	0.485	0.362	0.770	0.575
B0Ib	0.122	0.042	0.193	0.067	0.356	0.123	0.565	0.196	0.054	0.019	0.087	0.030	0.159	0.055	0.253	0.088
B0.2Ib	0.135	0.053	0.214	0.084	0.394	0.154	0.625	0.245	0.060	0.024	0.096	0.037	0.176	0.069	0.280	0.110
B0.5Ib	0.025	0.009	0.039	0.014	0.073	0.026	0.115	0.042	0.011	0.004	0.018	0.006	0.032	0.012	0.052	0.019
B0.7Ib	0.024	0.010	0.039	0.016	0.071	0.029	0.113	0.046	0.011	0.004	0.017	0.007	0.032	0.013	0.051	0.021
B1Ib	0.176	0.079	0.280	0.125	0.515	0.230	0.818	0.365	0.079	0.035	0.125	0.056	0.231	0.103	0.367	0.163
B1.5Ib	0.267	0.157	0.424	0.250	0.782	0.460	1.241	0.730	0.120	0.070	0.190	0.112	0.350	0.206	0.556	0.327
B2Ib	0.153	0.085	0.243	0.135	0.447	0.248	0.710	0.394	0.068	0.038	0.109	0.060	0.200	0.111	0.318	0.176
B2.5Ib	0.637	0.522	1.012	0.829	1.863	1.527	2.958	2.424	0.286	0.234	0.453	0.372	0.835	0.685	1.326	1.087
B3Ib	0.318	0.226	0.505	0.359	0.931	0.662	1.478	1.051	0.143	0.102	0.227	0.161	0.417	0.297	0.662	0.471
B4Ib	0.285	0.222	0.452	0.353	0.833	0.650	1.323	1.031	0.128	0.100	0.203	0.158	0.373	0.291	0.593	0.462
B5Ib	0.112	0.088	0.178	0.139	0.327	0.256	0.520	0.406	0.050	0.039	0.080	0.062	0.147	0.115	0.233	0.182

All flux values are in mJy.

A brief description of the observation of each star from the Bieging et al. (1989) study is now given. Cyg OB2 No. 9 has multiple observations with radio fluxes ranging from 0.7 to 8.1 mJy at 5 GHz and 1.0 to 6.3 mJy at 1.5 GHz. However, the spectral index is either not known or negative for these observations, indicating a probable non-thermal origin. One observation did herald a positive spectral index (0.6 ± 0.3) at both 1.5 GHz and 5 GHz, and these fluxes are those displayed in Figure 5.10. The 1.5 GHz flux is only an upper limit value but the 5 GHz flux has a thermal spectral index and resides ontop of the $f_{\text{cl}} = 2$ line.

Cyg OB2 No. 8A also has a number of flux values, but half of these have a negative spectral index and others have none. The fluxes of 0.4 ± 0.1 mJy and 1.0 ± 0.2 mJy both at 5 GHz in Figure 5.10 should be treated with caution.

Cyg OB2 No. 7 has one observation deemed ‘probable free-free emission’ by Bieging et al. (1989), but with no quoted spectral index. The observed flux of 0.4 mJy at 5 GHz corresponds to a predicted flux with a clumping factor of $f_{\text{cl}} = 2$, although the large observed errors also encompasses predicted fluxes with $f_{\text{cl}} = 1$.

Cyg OB2 No. 12 has a few observations and these are shown to also be thermal emissions; 3.4 ± 0.2 mJy ($\alpha = +0.5 \pm 0.3$) and 2.2 ± 0.2 mJy ($\alpha = +0.1 \pm 0.1$) at 5 GHz, with an increase in observed fluxes with higher frequencies (6.0 ± 2.0 mJy at 15 GHz with $\alpha = +0.5 \pm 0.3$). There is also a 1.5 GHz flux of 2.0 ± 0.2 mJy ($\alpha = +0.1 \pm 0.1$). Furthermore, Bieging et al. (1989) label the flux from this star as ‘definitely’ free-free emission. The observed fluxes at both 1.5 GHz and 5 GHz are substantially larger than the predicted fluxes with $f_{\text{cl}} = 10$, signifying a significant amount of clumping. This star is classified as a B8Ia by Bieging et al. (1989), although the COBRaS OB catalogue (originally classified by Skiff (2010)), Massey and Thompson (1991) and Klochkova and Chentsov (2004) all define it as a B5Ia. For a mid B supergiant this is an extremely high radio flux at these frequencies, and Klochkova and Chentsov (2004) note that the wind is time variable.

Cyg OB2 No. 22AB has one radio flux with an upper limit which is consistent with free-free emission (Bieging et al. (1989)). The flux value of < 0.3 mJy at 5 GHz is situated on the $f_{\text{cl}} = 2$ line and only further radio observations from COBRaS will be able to constrain this value to confirm the amount of clumping from this star.

Cyg OB2 No. 8B has one radio flux with an upper limit which is consistent with free-free emission (Bieging et al. (1989)). The flux value of < 0.2 mJy at 5 GHz is situated just

above the $f_{\text{cl}} = 10$ line. However, being so close to the origin (at 0 mJy) and only being an upper limit, there is still some uncertainty to the amount of clumping in the wind of this star.

Cyg OB2 No. 8C has one radio flux with an upper limit which is consistent with free-free emission (Bieging et al. (1989)). The flux value of < 0.2 mJy at 5 GHz is situated just above the $f_{\text{cl}} = 1$ line. The upper limit means no clear conclusions can be made with respect to clumping.

Cyg OB2 No. 11 has one radio flux with an upper limit which is consistent with free-free emission (Bieging et al. (1989)). The flux value of < 0.4 mJy at 5 GHz lies below but within the large errors of the $f_{\text{cl}} = 1$ line. This result would suggest that no clumping exists in the outer winds of this star.

Cyg OB2 No. 5 has one radio observation in the Bieging et al. (1989) survey of 1.8 ± 0.3 mJy at 5 GHz but with an unknown spectral index. Other observations of Cyg OB2 No. 5 were made by Persi et al. (1985) with 5 GHz fluxes of 5.8 ± 0.7 mJy and 6.3 ± 0.3 mJy with a spectral index of $\alpha = 0.2 \pm 0.1$ for the latter observation. Additional radio observations of Cyg OB2 No. 5 were made by Ortiz-León et al. (2012) with the VLA and found flux values of 7.9 ± 0.2 mJy at 5 GHz and 6.8 ± 0.4 mJy at 1.5 GHz with a spectral index $\alpha = +0.12 \pm 0.05$. Cyg OB2 No. 5 is a quadruple system (Kennedy et al. 2010) with reported periods of 6.6 days (Persi et al. 1985; Rauw et al. 1999; this thesis, see Section 6.6.3) and 6.7 years (Kennedy et al. 2010; Cazorla et al. 2014). The known multiple components, variable flux and flat spectrum all point towards a non-thermal contribution to the flux and therefore the apparent heightened flux consistent with $f_{\text{cl}} > 10$ is not a result of clumping in the wind but of non-thermal emission associated with a multiple component system.

Considering all of the above, there is evidence of deviation from predicted fluxes with $f_{\text{cl}} = 1$ after including uncertainties with some of the spectral indexes, and more significantly the errors on the predicted fluxes. However, what is required is a statistically viable sample of stars to demonstrate differences between observed and theoretical fluxes and thus mass loss rates. COBRaS will ultimately provide an unbiased sample of massive stars in Cyg OB2, with the sensitivity to robustly examine the ratio between observed and predicted fluxes with a high level of accuracy.

5.7 Current COBRaS L-band and C-band Maps

The COBRaS C-band Legacy data observations will commence from mid 2014 to 2015, when the e-MERLIN C-band receivers have been upgraded to the full 2 GHz bandwidth. The C-band pointings will detect more of the massive star population than L-band will, due to the thermal nature of the stellar winds. Only when the C-band Legacy datasets are delivered, can the extensive radio clumping investigation begin. This Section contains only fragmented datasets from the commissioning period of e-MERLIN and the first fraction of L-band Legacy data.

At the time of writing this thesis (April 2014) the COBRaS dataset includes one usable C-band 18th July 2012 commissioning (512 MHz bandwidth) pointing, 8 hours on target, and one run through all of the L-band 4th May 2013 pointings with ~ 1.4 hours on each pointing with full 512 MHz bandwidth (see Table 3.1 in Chapter 3). Due to only having roughly a quarter of the required time on each target field for L-band, the expected RMS level of each map will be significantly higher than with the full allocated Legacy time. Using Equation 1.13 with the typical values as stated in Section 1.1.2, the expected RMS noise level for the C-band commissioning pointing is 17 μ Jy and for L-band pointings is 22 μ Jy (but this value is assuming no RFI is present).

5.7.1 COBRaS C-band Commissioning Maps

The map of the C-band commissioning pointing is noise dominated and the RMS level from the image histogram is 54 μ Jy, which is higher than the expected RMS. However, due to the commissioning nature of the observation, there may be additional system noise from the electronics which has not been considered in the calculation, or perhaps residual errors from phase calibration.

Running the source detection algorithm from Chapter 4, yields two sources over the 5σ detection limit. Neither of the two sources are found to correlate with any positions of known stars in the OB catalogue, COBRaS catalogue or 2MASS catalogue. The first source is shown in Figure 5.11 as the central contours in the field. JMFIT returns a flux of $343 \pm 98 \mu$ Jy and the PP method gives a flux of $327 \pm 35 \mu$ Jy resulting in a 6σ detection at the noise level of 54 μ Jy. The second source lies at the edge of the map where the noise is higher and is therefore not considered further.

Due to the source not correlating with any known source in the catalogues (and the

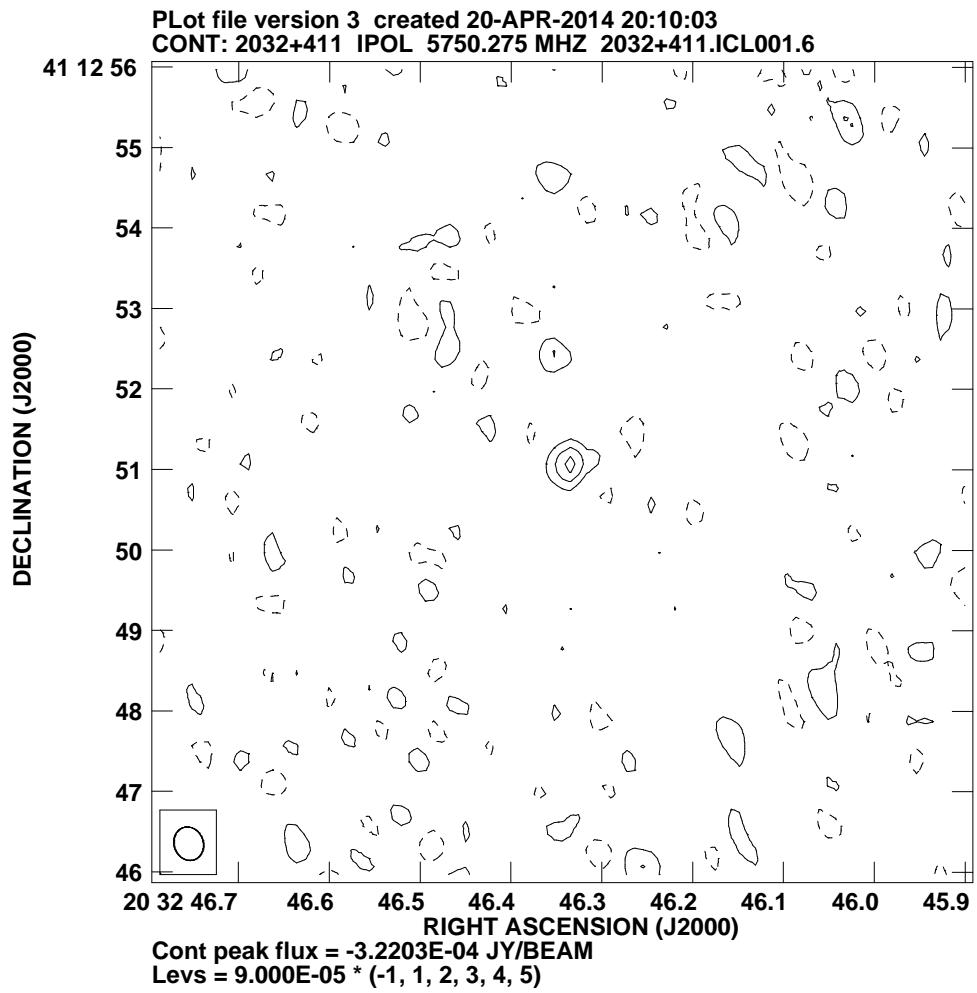


Figure 5.11: A contour map of a $\sim 300 \mu\text{Jy}$ source found in the C-band 18th July 2012 commissioning data at $20^{\text{h}} 32^{\text{m}} 46.335^{\text{s}}$ $+41^{\circ} 12' 51.070''$. Contour levels are -1, 1, 2, 3, 4, 5 $\times 90 \mu\text{Jy}$. The synthesised beam is given in the bottom left corner.

NASA/IPAC Extragalactic Database), it is impossible to classify any further. This source will be sought after with the full C-band Legacy dataset where the noise levels should be considerably lower than the C-band commissioning data, enabling maps with higher dynamic ranges. These future datasets will definitively show whether this source is real or not.

5.7.2 COBRaS L-band Legacy Maps

Wide-field maps for the L-band pointings C and D (see Table 3.1 for details) are created and have a RMS noise level of $113 \mu\text{Jy}$ and $111 \mu\text{Jy}$ respectively. These are the most sensitive L-band maps to date from COBRaS and only contain a fraction of the total allocated time. Despite this, sources have been found in both maps.

The first source is shown in Figure 5.12, and has a flux of 1.8 ± 0.2 mJy (from JMFIT). This source is coincident with a source in Setia Gunawan et al. (2003) and is within the positional errors of that study. Moreover, the flux at 1.4 GHz from that study is given as 2.2 ± 0.3 mJy with the authors stating that their amplitude calibration uncertainty reaches 10 - 20% for their 350 MHz observations (nothing stated for their 1.4 GHz observations). Given the proximity and similar fluxes, it is reasonable to infer that these sources are the same.

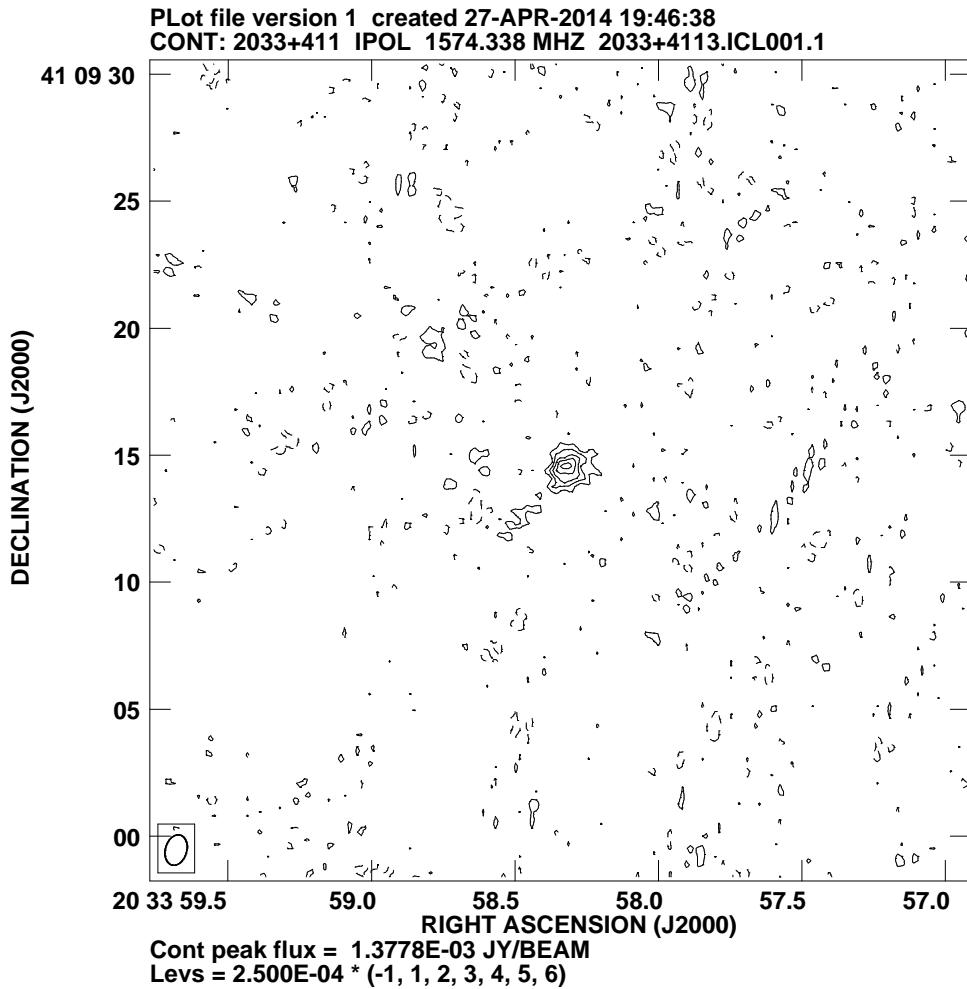


Figure 5.12: Contour map of 1.8 mJy source in the L-band pointing C field at $20^{\text{h}} 33^{\text{m}} 58.31452^{\text{s}}$ $+41^{\circ} 09' 14.5002''$. Contour levels are $-1, 1, 2, 3, 4, 5, 6 \times 0.25$ mJy. The synthesised beam is given in the bottom left corner. This source coincides with a source from Setia Gunawan et al. (2003).

Furthermore, Setia Gunawan et al. (2003) cite a spectral index of $>+0.37$, suggesting thermal emission. Therefore this source should be even brighter in the COBRaS C-band Legacy observations. This does however raise doubts on whether this source is a massive

star. It is around the brightness level expected of a massive binary at L-band frequencies, but does not show non-thermal emission according to the spectral index cited by Setia Gunawan et al. (2003). If it is indeed a massive star, the clumping factor for even the brightest O or B supergiant would be around $f_{\text{cl}} \geq 10$. Definitive classification can only occur with COBRaS determined spectral indices with the completion of the C-band Legacy dataset.

Cyg OB2 No. 9 is detected (confirmed by cross correlation with the OB catalogue) in the L-band pointing D and is shown in Figure 5.13. The detected flux is 3.8 ± 0.2 mJy (from JMFIT) which falls between the 1.0 - 6.3 mJy fluxes determined by previous observations (Bieging et al. 1989; Setia Gunawan et al. 2003 and references therein). Cyg OB2 No. 9 is a known binary system and the strong 1.5 GHz flux would strongly suggest that non-thermal emission is detected here, although 5 GHz fluxes from COBRaS are required to confirm this.

It may be possible to obtain fluxes from individual IFs from this incomplete dataset due to the strong flux from Cyg OB2 No. 9 to determine the spectral index. Furthermore this may be necessary because of the variable flux (from binary motion) and because the C-band Legacy data has not been observed at the same time. This will affect any obtained spectral index from comparing fluxes at two different frequencies, observed at different epochs.

Pointing D reveals another strong source. Figure 5.14 shows a source with a flux of 7.3 ± 0.2 mJy (JMFIT) located at $20^{\text{h}} 32^{\text{m}} 56.795^{\text{s}}$ $+41^{\circ} 08' 53.487''$. The position of the bright source correlates to a known source in the COBRaS catalogue, matching a 14.7 ± 0.8 mJy source from Setia Gunawan et al. (2003) and a 12.99 ± 1.9 mJy source from the Radio Master catalogue, originally from the VLA study by White et al. (2005).

The flux from the COBRaS L-band Legacy map differs drastically from the previous surveys of Setia Gunawan et al. (2003) and White et al. (2005). Setia Gunawan et al. (2003) determined the spectral index of the source to be $\alpha = -0.77$, indicating a steep spectrum non-thermal emission.

The image containing the source in Figure 5.14 was created to be a wide-field map to image across the primary beam. A higher resolution map (220 milli arcseconds) is created of the same source as is shown in Figure 5.15. A similar flux of 7.1 ± 0.3 mJy (JMFIT) is obtained, however, the PP method returns a slightly higher flux of 8.09 ± 0.01 mJy (although the uncertainty on the amplitude calibration will be significantly larger than

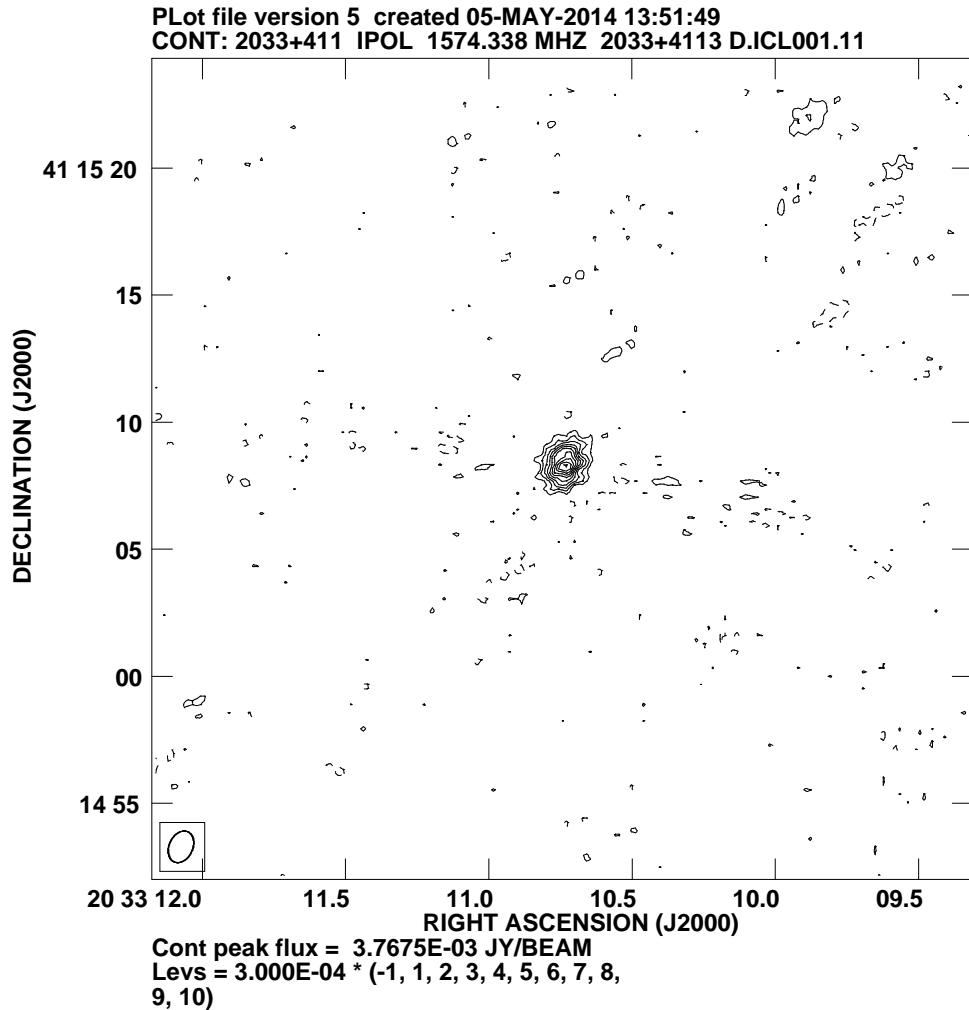


Figure 5.13: Contour map of Cyg OB2 No. 9 in the L-band pointing D field at $20^h 33^m 10.735^s +41^\circ 15' 08.353''$. Contour levels are $-1, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 \times 0.3$ mJy. The synthesised beam is given in the bottom left corner.

this value). This is because JMFIT attempted to fit the resolved source with one Gaussian, whereas the PP method determines the flux independent of source structure. Although it is unknown why the PP flux for the resolved source is higher than the unresolved source. The noise RMS of the high resolution map is $55 \mu\text{Jy}$.

5.7.3 Summary of Sources in COBRaS Maps

All the contour maps presented in this section are created using the AIPS task KNTR. Table 5.10 contains the sources found in the COBRaS early science maps.

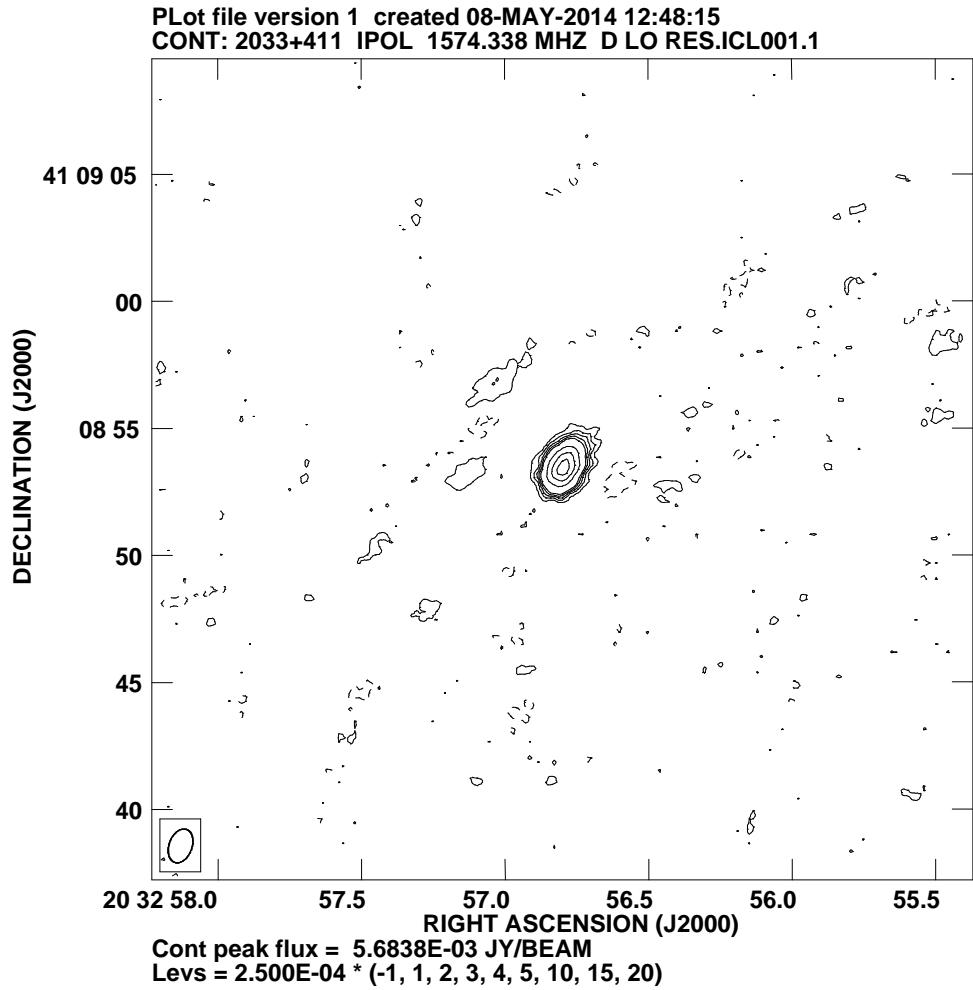


Figure 5.14: Contour map of a strong source in the L-band pointing D field centred at $20^{\text{h}} 32^{\text{m}} 56.795^{\text{s}}$ $+41^{\circ} 08' 53.487''$. Contour levels are $-1, 1, 2, 3, 4, 5, 10, 15, 20 \times 0.25$ mJy. The source has a flux of 7.3 ± 0.2 mJy (JMFIT) and an image rms $\sim 91 \mu\text{Jy}$. The synthesised beam is given in the bottom left corner.

Table 5.10: Sources in the COBRaS Maps

Right Ascension ^a	Declination ^a	Map	Flux / μJy		Known Source*
			JMFIT	PP	
20 32 46.335	41 12 51.070	C-band 18 th July 2012	343 ± 98	327 ± 35	No
20 32 56.795	41 08 53.487 [†]	L-band Pointing D	7325 ± 184	7299 ± 39	Yes
20 33 10.735	41 15 08.353	L-band Pointing D	3772 ± 234	3814 ± 55	Yes
20 33 58.315	41 09 14.500	L-band Pointing C	1790 ± 240	1690 ± 60	Yes

^a Right Ascension in units of hours minutes and seconds and Declination is in units of degrees arcminutes and arcseconds.

* i.e. in one of the catalogues created in Chapter 4.

[†] Information from the unresolved source.

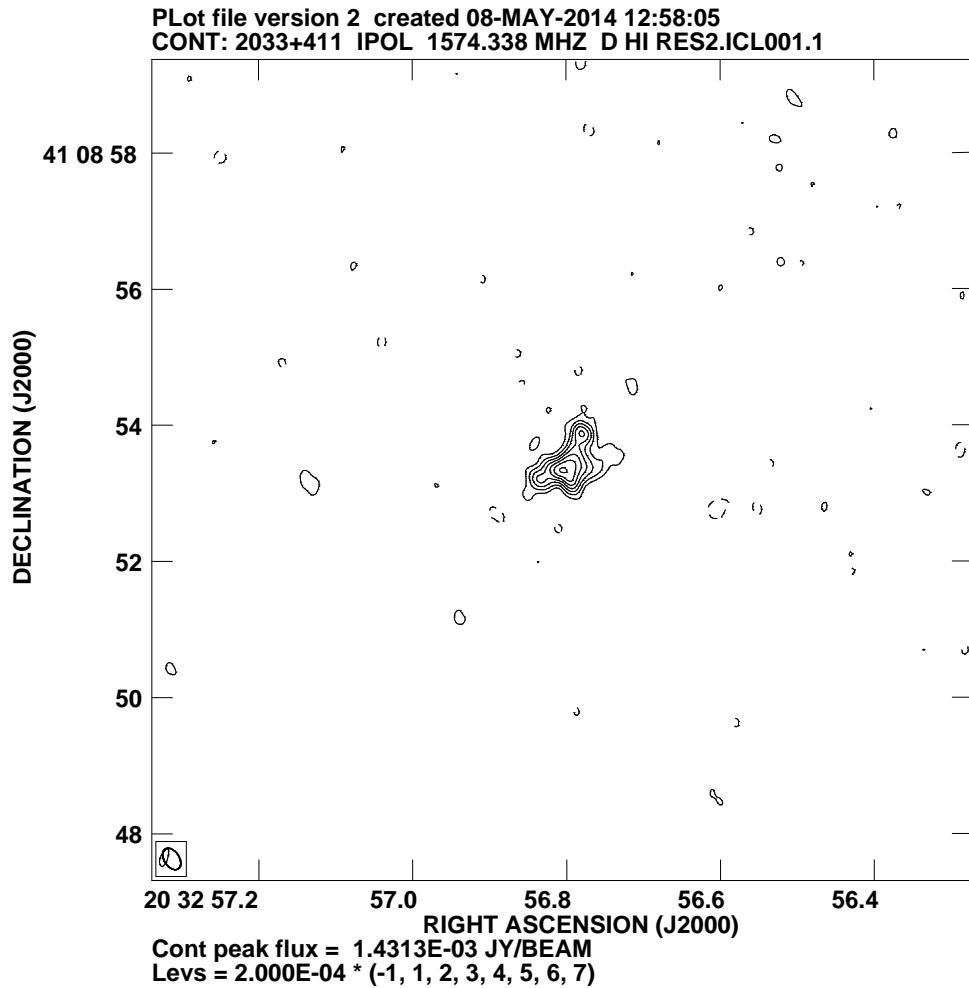


Figure 5.15: High resolution contour map of a strong source in the L-band pointing D field centred at $20^{\text{h}} 32^{\text{m}} 57.2^{\text{s}}$ $+41^{\circ} 08' 58''$. Contour levels are $-1, 1, 2, 3, 4, 5, 6, 7 \times 0.2$ mJy. The source has a flux of 7.3 ± 0.2 mJy (JMFIT) and an image rms $\sim 55 \mu\text{Jy}$. The synthesised beam is given in the bottom left corner.

5.8 Conclusions and Discussions

This Chapter contains an introduction to the winds from massive stars, radio emission from massive stars and massive binaries, as well as mass loss rates and the role of clumping in the winds. Using the mass loss recipes as described by Vink et al. (2000) and Vink et al. (2001) with the fundamental parameters for O stars (Martins et al. 2005), early B supergiants (Searle et al. 2008) and terminal velocities (Prinja et al. 1990), a grid of predicted smooth wind mass loss rates have been created. Applying this mass loss rate to the relation of excess free-free emission (Equation 5.21), the radio fluxes at 1.5 and 5 GHz for O stars and early B supergiants are determined. Quantifying the amount of clumping

using a clumping factor (Equation 5.22), adjustments to the radio fluxes from massive stars is made with a range of clumping factors $f_{\text{cl}} = \{1, 2, 5, 10\}$, consistent with those from the literature for H α emission and radio mass loss rates (see Section 5.6). Comparisons of the observed radio fluxes of Cyg OB2 stars from previous surveys is made with the grid of clumped radio fluxes to give an initial impression of the amount of clumping in the winds of Cyg OB2 massive stars. Finally, the first radio maps from the COBRaS project are presented with source and flux lists.

Some of the errors on the predicted radio fluxes are large due to the large uncertainties on some of the fundamental parameters. This also includes adopting one fundamental parameter for each spectral class when in reality these will vary for different stars of the same class. The terminal velocity is one parameter which varies considerably within a spectral class and also has a significant effect on the radio flux. Some spectral classes only had one reference star to determine the terminal velocity, thereby skewing the predicted flux trend. For example, the brightest predicted flux for an O supergiant from Table 5.7 is from an O5I star, not an O3I star. The terminal velocity for O5I is significantly lower than the neighbouring classes. Investigating in greater detail reveals that the terminal velocities for O3I and O5I stars from Prinja et al. (1990) only had one example for each class. Furthermore, the terminal velocity for O4I came from three stars with a range of $v_{\infty} = 1880 - 2605 \text{ km s}^{-1}$, demonstrating the fallacy of adopting a ‘one size fits all’ parameterisation from a limited sample. It should be noted that the rarity of early O supergiants does makes this task difficult.

Moreover, once COBRaS has identified and obtained radio fluxes for massive stars and candidates, follow-up multi-wavelength surveys will help with obtaining unique fundamental parameters for each target, thereby enabling precise determination of predicted radio fluxes. This is necessary for definitive analysis on clumping in the winds of massive stars, and this thesis provides the foundation for that analysis.

The source detection algorithm has been tested on real maps from COBRaS for the first time as opposed to idealised simulations and the 5σ seed threshold run finds a possible candidate source at C-band and more known sources at L-band. Additional runs with lower seed thresholds were made finding a large number of potential sources. Any faint sources from these searches will require cross correlation with catalogues to have any confidence of a source detection over an image artefact.

One test of realism on whether the C-band source is real is to determine the probability

that one pixel in an image of over 16 million pixels will randomly have a flux over 5σ . Then, what is the probability that two pixels situated next to one another have a flux over 5σ ? The source in Figure 5.11 has four pixels at the third contour level (6σ) which strengthens the possibility that it is indeed a real source and not a random noise feature.

The expected noise level for the C-band commissioning map is lower than the one obtained from the calibration pipeline. The pipeline may be the limiting factor in sensitivity of the maps, however, due to the commissioning state the assumed system temperature of 50 K may be incorrect or more likely additional noise from the unsettled system electronics may cause the actual RMS level to be higher than the expected level.

The L-band legacy pointings C and D, contain a number of interesting sources. The most relevant to this thesis is the strong detection of Cyg OB2 No. 9, a massive binary, and the first confirmed detection of a massive star from COBRaS. In creating the map of Cyg OB2 No. 9, all of the software presented in this thesis was required; SERPent to flag the RFI, the pipeline to calibrate the data, the source detection algorithm to find the source and determine the flux, and the catalogues and cross correlation scripts to identify the source. This result demonstrates the effectiveness and the importance of these programs when searching for sources. There are other interesting sources found in the L-band Legacy pointings, with two sources previously detected by Setia Gunawan et al. (2003) shown in Figures 5.12, and 5.14 and 5.15.

An important note to make here is the final noise RMS level of the wide-field L-band Legacy maps, pointings C and D have a noise RMS level of 113 and 111 μJy respectively. This only includes a fraction of the total allocated time for each field. However, even with 5 hours on target, this will not reduce the noise RMS to the expected RMS of 11 μJy at L-band for e-MERLIN. This is due to the amount of data removed, because RFI is fairly substantial at L-band, meaning the realistic noise RMS level for any COBRaS L-band field may be up to a few factors higher than expected. Moreover, IFs 1 and 2 were removed for imaging due to the substantial RFI in these IFs for this observation. Future e-MERLIN sensitivity focused projects at L-band and low declinations, may have to consider the final noise levels from COBRaS maps when planning observations.

Furthermore, the second Setia Gunawan et al. (2003) source was located near the edge of the wide-field pointing D map. The original source contained phase errors around this source, and therefore was re-imaged with a RA and Declination shift to the source's location. This solved the abberation problem from the wide-field imaging, and further

emphasises a point made in Chapter 3 about the possibility of imaging the primary beam as a number of facets to reduce wide-field effects. Moreover, the noise RMS level for this shifted map was $91 \mu\text{Jy}$ for the low resolution (900 milli arcseconds) map (Figure 5.14) and $55 \mu\text{Jy}$ for the high resolution (220 milli arcseconds) map (Figure 5.15). Therefore, in order to achieve the expected image RMS, mosaicing within individual primary beams may be necessary.

5.9 Future Work with COBRaS

COBRaS Legacy data has only just started to arrive, and it is expected that the remaining L-band data will be received throughout 2014. This will increase the sensitivity of the current pointings, enabling more than just a ‘first look’ at the core of Cyg OB2 at low frequencies. The expectation is that these maps will include a number of background radio galaxies, weak T-Tauri stars, massive-binary interaction regions and other non-thermal emitters, whilst including some of the thermal emission from massive stars. Only when the extensive C-band Legacy observations start, can the real census of massive stars in Cyg OB2 begin. A large sample of radio fluxes and radio mass loss rates can answer the question, ‘is there clumping in the outer winds of massive stars, from radio observations?’, the implications of which have already been discussed.

Chapter 6

X-ray Variability of Massive stars in Cyg OB2

Lies, damned lies, and statistics.

Benjamin Disraeli

One of the key science goals of COBRaS is synergy with other wavelength surveys, providing a full spectral window to the stellar population of Cyg OB2. One of these surveys is the Chandra Cygnus OB2 Legacy survey, observing the same core region of Cyg OB2 as COBRaS, and also the region around the TeV source (J2032+4130), in X-rays.

X-rays trace high-energy phenomena, which in the context of COBRaS, applies to flaring pre-main sequence stars and therefore star formation, and the collision region between two massive binaries. The former results in a dramatic increase in X-ray photon counts on a short timescale (\sim hours) in comparison to the non-flaring X-ray continuum level, whereas the latter exhibits a gradual X-ray variability over a longer period (days to weeks).

This chapter focuses on the X-ray variability from massive stars, with particular emphasis on known and potential massive binaries. This will contribute towards better determination of the binary fraction of massive stars within young stellar associations, which is one of the key science goals of COBRaS.

Firstly, an introduction to X-ray emission from massive stars is given, followed by an introduction to the Chandra Cyg OB2 Legacy survey and the Chandra X-ray observatory. Then, an introduction to the statistical methods used in X-ray astronomy is given. Finally, the X-ray variability study of the O-type and WR stars in the Chandra Cyg OB2 Legacy survey is presented.

6.1 X-ray Emission from Massive Stars

Theories on the origins of X-rays from single O-type stars are confined to the stellar winds of these objects. One explanation generally attributed for this emission from single O-type stars is the presence of many small hydrodynamic shocks arising from instabilities intrinsic to the radiatively-driven wind associated with massive stars (Feldmeier et al. 1997). These Line Deshadowing Instabilities (LDI) average out over the emitting volume, resulting in an X-ray flux which is similar to the observed fluxes. The authors used fundamental stellar parameters similar to those of the O supergiant ζ Ori for their simulations.

Another model applies to stars with strong magnetic fields, where in the presence of a strong dipole magnetic field, the radiatively-driven wind is channeled along the magnetic field lines from each hemisphere, meeting at the magnetic equator (Babel and Montmerle 1997; ud-Doula and Owocki 2002; Gagné et al. 2005). The collision between the two streams at high velocity produces very high shocked temperatures which emit harder X-rays than those expected from LDIs.

In the case of massive binaries, an additional source of X-rays can be found from the large shocks at the centre of the interaction region of the stellar winds (Stevens et al. 1992). As the winds meet, a contact discontinuity separates two shocks in the interaction region, assuming the winds have achieved terminal velocity. The shocks will produce X-rays and because the wind is at terminal velocity in these systems as opposed to an accelerating wind in the situation of LDIs, the spectrum is expected to be harder than the X-rays produced in the individual stellar winds.

6.1.1 X-ray Variability from Massive Binaries

The X-ray emission from massive binaries is expected to demonstrate similar variable behaviour to the non-thermal radio emission also associated with massive binaries, because of the origins of these emissions. Non-thermal radio emission arises from the electrons

spiralling along the magnetic field lines emitting synchrotron emission, initially accelerated by the shocks in the interaction region, whereas the X-rays arise directly from the shocks. Therefore the orientation of the binary can affect the amount of emission which is seen by the observer. This translates into a variable non-thermal radio and X-ray emission over time, as the two stars orbit one another, changing the line of sight to the shocked region. The subsequent X-ray variability should be modulated on the orbital period if the emission arises from the interaction region.

6.1.2 Luminosity relations for O Stars in Cyg OB2

From the very first observations of X-rays from massive stars, a distinct relation between the X-ray luminosity and the bolometric luminosity $L_X/L_{\text{bol}} \sim 10^{-7}$ has been observed from a number of studies on a handful of O-stars (see Feldmeier et al. 1997 and references therein). Larger surveys with ROSAT (Berghoefer et al. 1997) and XMM-Newton (Nazé 2009) have confirmed this relation for O-type stars. Interestingly, the same relation does not exist for Wolf-Rayet (WR) stars (Wessolowski 1996), and only the very earliest B-stars appear on the end of this luminosity relation.

The X-ray luminosities can be derived for a source of known distance (see Section 6.4.3). The distance to the Cygnus OB2 massive star association has historically been somewhat uncertain with values ranging from no less than 1.2 kpc (Hanson 2003) to 2.1 kpc (Reddish et al. 1966). However, recent parallax observations of 6.7 GHz methanol and 22 GHz water masers determine the distance of the whole Cygnus X complex to be 1.40 ± 0.08 kpc with the distances to the constituent HII regions consistent with this value (Rygl et al. 2012; Wright et al. in prep.).

Rauw et al. (2014) derives the relation (Equation 6.1) between the X-ray luminosity and bolometric luminosity by using the X-ray fluxes and bolometric fluxes which are independent of the distance of Cyg OB2. The authors compute the X-ray fluxes from the ACIS spectra and correct for interstellar absorption only. The bolometric fluxes are calculated from V magnitudes assuming visual extinction $R_V = 3.1$ and adopting the bolometric corrections of Martins and Plez (2006). If no B or V magnitudes are available, the authors relied on near-IR photometry of Negueruela et al. (2008) for bolometric corrections. This reveals the interstellar absorption corrected relation between the X-ray luminosity and bolometric luminosity, where the fluxes are evaluated between the 0.5 - 8 keV energy range.

Rauw et al. (2014) then discuss whether to adopt a scaling-law or a power-law between the X-ray flux and bolometric flux, and conclude that both methods are consistent within 2σ . This results in the following X-ray luminosity and bolometric luminosity relation:

$$\log L_X/L_{\text{bol}} = -7.21 \pm 0.24 \quad (6.1)$$

which is in agreement with the scaling relation found by Nazé et al. (2011) for the O stars in the Chandra Carina Complex Project.

6.2 The Chandra Cygnus OB2 Legacy Survey

The Chandra Cygnus OB2 Legacy survey is a 1.08 Ms, one square degree survey with Chandra centred on the core of the Cygnus OB2 massive star association (Drake et al. in prep.). The uniform deep exposure study detects stars down to the lower mass limit of 1 solar mass, providing the most comprehensive study of a massive, young star cluster in the X-ray regime. The survey took place between the dates of 25th January - 10th March 2010 using the ACIS-I camera on Chandra. Figure 6.1 shows the mosaics of 36 pointings of the Chandra survey from the design phase of the project (left) and the executed phase (right).

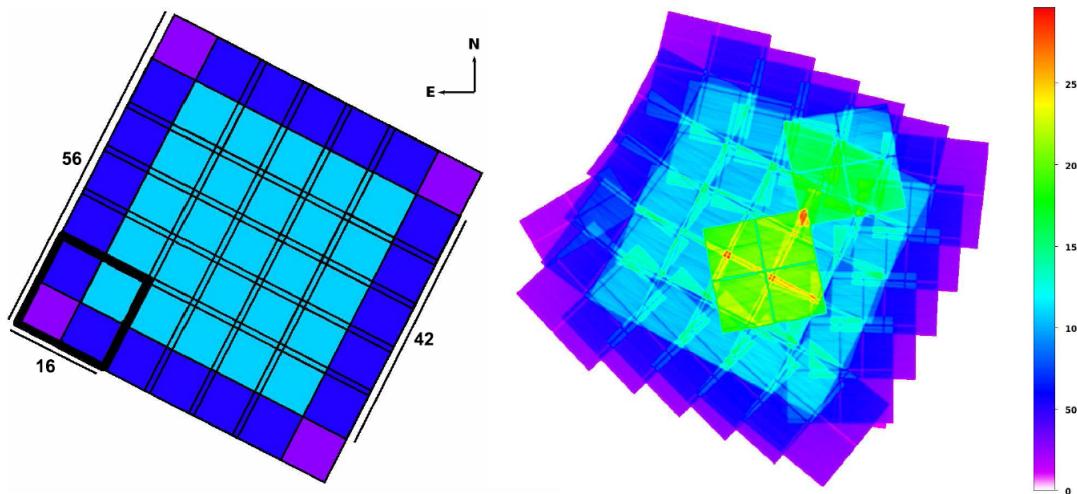


Figure 6.1: The Chandra Cygnus OB2 survey's deep mosaic coverage of the Cygnus OB2 region with a total of 36-pointings each with 30 ks integration time. The left image is the “designed” coverage and the right the “executed” coverage. The survey extends to one square degrees, centred on the core region of Cyg OB2, which is the green tiling in the middle of the right figure. This has a exposure time of ~ 214 ks and is also the survey region for COBRaS. This figure is taken from Wright et al. (in prep.).

The mosaic strategy employed in the survey was to observe the central 0.5 square degrees to an effective exposure of 120 ks, the outer 0.4 square degrees to 60 ks, and an extra 100 ks observation for the core of Cyg OB2 (reaching 220 ks total integration). This observed region is to be complemented with other surveys including COBRA-S with e-MERLIN, VLA (radio), Spitzer and Herschel (IR), XMM-Newton and INTEGRAL (X-ray) to produce a multi-wavelength study.

Chandra detected in total, 8352 point sources, and Figure 6.2 displays the distribution of sources within the Chandra survey field of view, with the applied methods of detection (see Wright et al. in prep. for details).

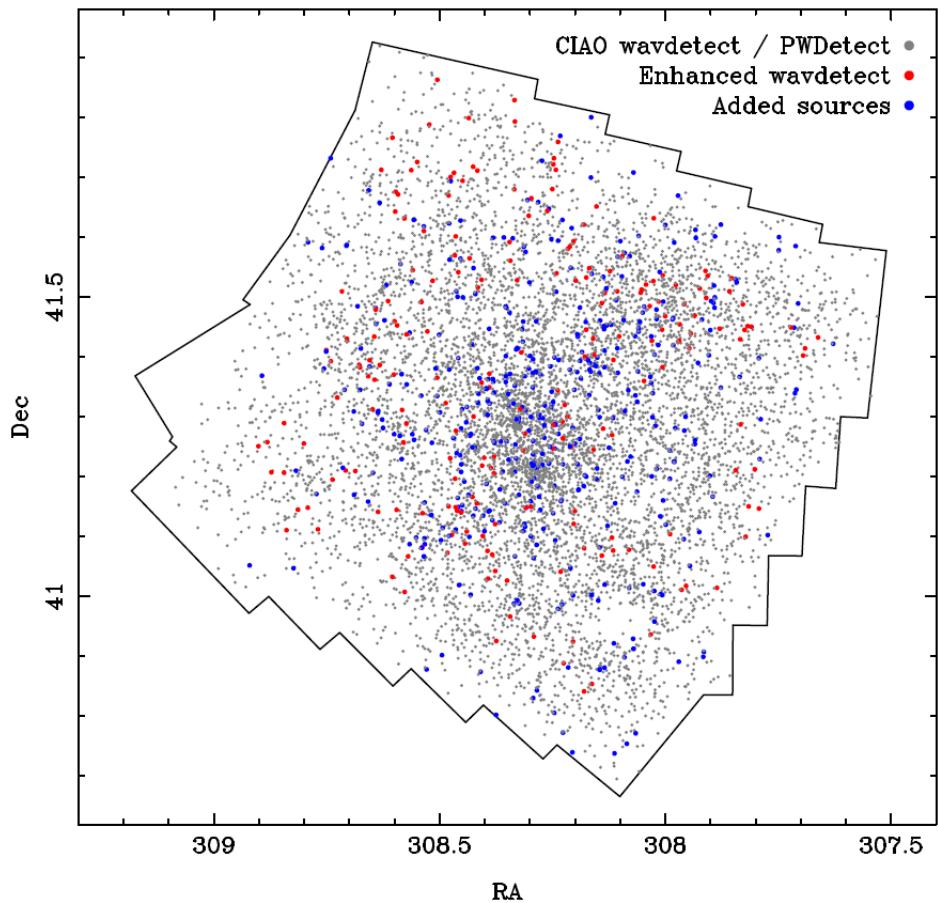


Figure 6.2: X-ray source distribution over the field of view of the Chandra Cyg OB2 Legacy survey. In total 8352 sources were detected by various methods. This figure is taken from Wright et al. (in prep.).

There are a number of scientific objectives for the Chandra Cyg OB2 Legacy survey, including the assessment on the size, mass and structure of the association, and the formation, evolution and dynamics of a massive star cluster. Investigations into the X-ray

properties of massive stars, intermediate stars and the nature of protoplanetary disks in a cluster environment are also objectives. Studies on Cygnus X-3 an X-ray binary is also a key objective. In relation to COBRaS, the X-ray variability of massive stars is important to characterise, confirm and detect massive binary systems, the analysis of which will contribute towards the effort of determining the binary fractions of massive stars in young stellar clusters.

The field of view of the Chandra survey detected a rich population of massive stars in Cyg OB2 with 100 B stars, 49 O stars and 3 WR stars (Wright et al. in prep.). The spectral types and luminosity classes range from O3I to O9.5V, with the most frequent spectral type being O8V (Rauw et al. 2014).

6.3 Introduction to Chandra

Satellite technology has revolutionised studies in the Ultra-violet (UV), X-ray and Gamma-ray regions of the electromagnetic spectrum. The Earth's atmosphere is opaque to radiation of these wavelengths, rendering observations in these bands impossible. Early attempts to circumvent these problems were made in 1948, adapting German V-2 rockets to launch detectors beyond the majority of Earth's atmosphere to detect the first astronomical X-rays emanating from the Sun (Keller 1995). Other methods included using a scintillation counter mounted on a hot-air balloon (Clark 1965), and combining hot air balloons and rockets (Kundu 1961).

The 1960s saw the launch of the first orbiting X-ray satellite SOLar RADiation satellite program (SOLRAD) amongst many Orbiting Geophysical Observatories (OGOs), and the first X-ray imaging telescope Einstein Observatory (HEAO-2) at the end of the 1970s. The Einstein Observatory adopted a near-circular orbit at an initial altitude of 500 km and operated in the energy range of 0.2 - 20 keV.

Around the same time, the Advanced X-ray Astrophysics Facility (AXAF; later renamed the Chandra X-ray Observatory) was first proposed. Throughout the years, many designs and changes were made until its final launch in 1999. These changes included reducing the number of mirrors (from 12 to 8) and only 4 of the initial 6 scientific instruments would be utilised. The orbit of Chandra was also reconsidered and an elliptical orbit reaching $\sim 140,000$ km apogee and $\sim 10,000$ km perigee was chosen as the majority (70%) of the orbit is above the Van Allen radiation belts (Weisskopf et al. 2000).

Figure 6.3 shows an illustration of the Chandra X-ray Observatory with the constituent parts of the spacecraft. Since the first light image of the supernova remnant Cassopeia A, which revealed a compact object at the centre (Pavlov et al. 2000), Chandra has made a host of scientific discoveries in a range of fields within astronomy. Chandra's superior sensitivity and angular resolution to any previous X-ray telescope has advanced the understanding of X-rays from astronomical sources.

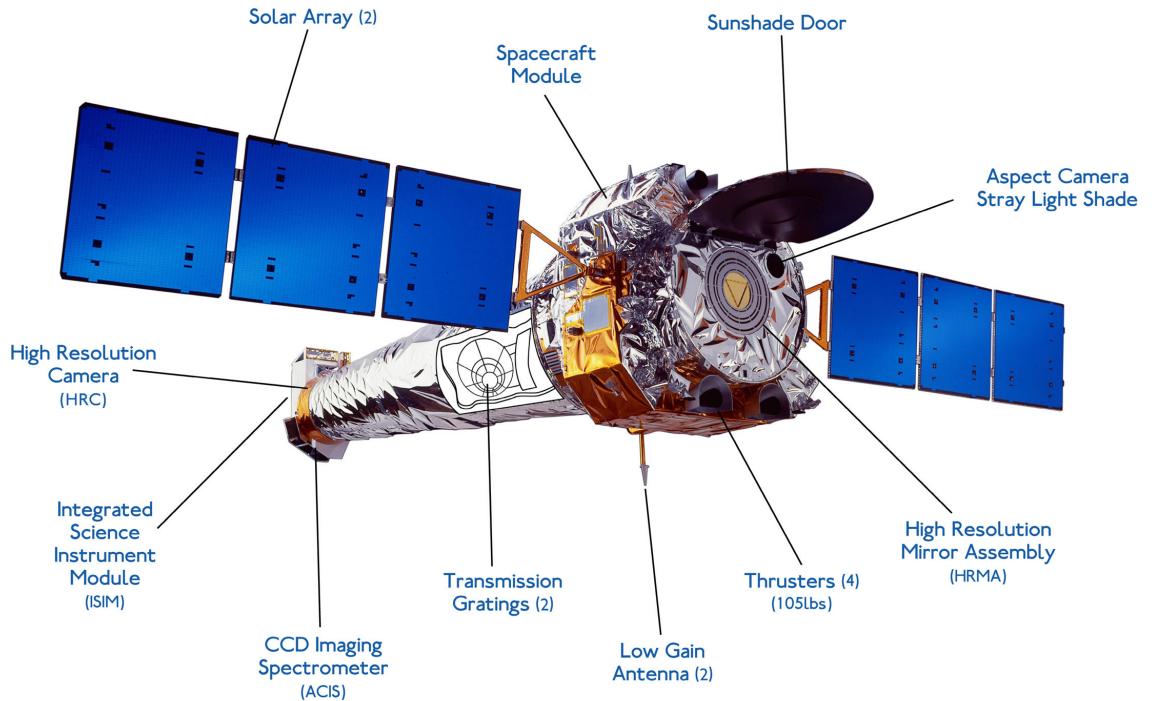


Figure 6.3: Image of the Chandra X-ray Observatory. Labelled are the main components of the spacecraft including the location of the X-ray subsystems.

6.4 Chandra ACIS Data Analysis and Calibration

The calibration procedure for Chandra is extensive and not in the scope of this thesis to describe in detail. The observations presented here are all taken with the ACIS CCD and are reduced via the method described by Broos et al. (2010)¹. The following sections (6.4.1, 6.4.2, 6.4.3) will describe ACIS specific considerations and procedures in the data analysis.

¹http://www2.astro.psu.edu/xray/acis/acis_analysis.html

6.4.1 ACIS Photon Pileup

The ACIS CCD is sensitive to X-ray photons and therefore behaves differently to optical CCDs. Optical CCDs work by releasing an electron into a potential well every time an optical photon hits a photoactive silicon layer. This charge builds up until it is read out by transferring the charge along the chip rows to the reader. However, the ACIS CCD only behaves as a photon counter, with a full-frame readout time of 3 seconds. If a photon is detected during this time window the CCD records the photon as an ‘event’ or ‘count’. If a compact bright X-ray source is observed, more than one photon may arrive at the same or neighbouring pixels during the 3 second readout time, but only one event will be recorded. This is known as ‘pileup’ as the X-ray photons pileup before readout.

For a count rate of 0.2 photons per pixel per frame, there is a significant pileup of $\sim 10\%$ based on Poisson statistics². That is, every 1 in 10 events will actually correspond to two or more photons. This will lead to an underestimation of the peak flux in the observed spectrum compared to the ‘true’ spectrum. Another effect is a shift to higher energies for any pileup-affected events, producing a high energy tail in the spectrum. To mitigate this effect, when observing a bright compact source, a short integration time is needed for each event. However, the total overhead to readout the entire frame is limited to 3 seconds, regardless of the integration time.

6.4.2 Chandra Observatory Response

The effective area of the Chandra observatory is a strong function of energy and therefore a point spread function (PSF) model is required. This model accounts for the point source light which falls outside of the aperture. This is necessary because the Chandra data assumes an infinitely large detector and extraction aperture (Broos et al. 2010). The fraction of power of the PSF that falls within the aperture is calculated at five monochromatic energies. Interpolations between these five energies are made to estimate the PSF fraction at every energy.

The quantum efficiency is the fraction of photons registered by the detector. An ideal detector would have a quantum efficiency of 100%. However, realistically this percentage is lower and varies as a function of energy and affects both the ACIS and HRC detectors. The quantum efficiency also deviates as a function of the position on the CCD.

²http://cxc.harvard.edu/newsletters/news_05/node12.html

All of these factors along with off-axis vignetting (where the edge of detector is less responsive compared to the centre), and the diffraction grating efficiency (also a function of energy) are tied up in an Auxiliary Response File (ARF). The ARF represents the effective area of the observatory and needs to be convolved with the observed spectrum to derive physical properties such as fluxes or luminosities.

6.4.3 Photometry with ACIS

To obtain fluxes and luminosities from standard aperture photometry, the calibration procedure begins by computing the counts for a range of energy bands, $S(E)$. A scaled background level is subtracted by calculating the counts in the neighbouring region to the source, using various masking techniques depending on the level of crowding, as explained in detail by Broos et al. (2010) (Section 5.4). The net counts, $S(E)$, is given by

$$S(E) = C^s(E) - \left(A^s / A^b \right) C^b(E) \quad (6.2)$$

where $C^s(E)$ is the number of counts in the source aperture for energy E , and $C^b(E)$ is the number of counts in the background aperture for energy band E . A^s and A^b are the source and background areas and are derived by integrating exposure maps (Broos et al. 2010). In this study, however, the background is not subtracted but monitored, and any observations with significant backgrounds are ignored. This is to simplify the calibration process. The observatory response (a function of multiple variables all tied up in the ARF) has not yet been considered in the calibration.

There are two methods to calculate the incident photon flux from the standard net counts, designated F^* and $F^\#$, both in units of photon $\text{cm}^{-2} \text{ s}^{-1}$. The first flux estimate (F^*) is determined by dividing the net counts, $S(E)$, by the observatory response (ARF) for every discrete energy band E , and then summed over a chosen energy range (E_{\min} and E_{\max}):

$$F_{\text{photon}}^* (E_{\min} < E < E_{\max}) = \sum_{E_{\min}}^{E_{\max}} \frac{S(E)}{\text{EXPOSURE} \times \text{ARF}(E)} \quad (6.3)$$

where $S(E)$ is the net counts (given by Equation 6.2), $\text{ARF}(E)$ is the observatory response at energy E , and EXPOSURE is the exposure time of the source. The energy range limits for Chandra are $E_{\min} = 0.5 \text{ keV}$ and $E_{\max} = 8 \text{ keV}$, with two commonly used bands being:

$E_{\min} = 0.5$ keV and $E_{\max} = 2$ keV for Chandra ‘soft band’ and $E_{\min} = 2$ keV and $E_{\max} = 8$ keV for Chandra ‘hard band’.

The second method in estimating the incident photon flux ($F^\#$) sums the net counts, $S(E)$, and the observatory response, $\text{ARF}(E)$, for the given energy range E_{\min} to E_{\max} , before averaging. This gives the photon flux $F^\#$ as

$$F_{\text{photon}}^\# (E_{\min} < E < E_{\max}) = \frac{\sum_{E_{\min}}^{E_{\max}} S(E)}{\text{EXPOSURE} \times \sum_{E_{\min}}^{E_{\max}} \text{ARF}(E)}. \quad (6.4)$$

The F^* estimator has limitations when the ARF has a small value because the corresponding source or background events have a large influence on the estimator. This leads to large Poisson errors and is therefore not recommended for weak sources (Broos et al. 2010). The $F^\#$ estimator averages across the observatory response (ARF) for all energies, which is only applicable for a flat incident spectrum.

Equations 6.3 and 6.4 refer to the photon flux, which has units of photons $\text{cm}^{-2} \text{ s}^{-1}$. This can be converted to an energy flux with units of $\text{erg cm}^{-2} \text{ s}^{-1}$, by multiplying by a conversion factor of 1.6×10^{-9} ergs/ 1 keV.

If information on the spectral shape and distance d to the target is known, it is possible to determine the apparent X-ray luminosity L_x for a broad energy band with either flux estimator and the background median corrected energy (E_{median}):

$$L_x = 4\pi d^2 [F^* \text{ or } F^\#] E_{\text{median}}. \quad (6.5)$$

The accuracy of Equation 6.5 for bright and faint sources is tested by Getman et al. (2010) who also use E_{median} to scale the observed luminosity (Equation 6.5) to the intrinsic luminosity corrected for absorption. Errors are larger for absorption-corrected luminosities and very large for heavily absorbed luminosities for the soft band regime (Broos et al. 2010). The authors conclude that for faint sources the flux estimator $F^\#$ is preferred and the absorption-corrected luminosity for the hard band (2 - 8 keV) over the luminosities in the whole Chandra band (0.5 - 8 keV).

6.5 Statistics and Variability Methods for X-ray Astronomy

Due to the limited number of photon counts from any observation, X-ray astronomy resides in the Poisson statistical regime. Traditional methods involving χ^2 or data obeying

Gaussian statistics are not applicable for these X-ray datasets, where sampling is often uneven with few counts per bin. Therefore a range of alternative statistics and methods are required to search for variability in the Chandra Legacy datasets.

6.5.1 Null Hypothesis

In statistics, the null hypothesis tests whether a relationship between two measured variables exists. It can only be disproved or rejected if the data or distribution reaches a set significance level, but can never be ‘proven’. Ronald Fisher first coined the term *null hypothesis* in *The Design of Experiments* (Fisher 1966) as an explanation to how a hypothesis could be disproven by a single experiment, but no finite number of experiments can prove it.

There are two types of null hypothesis; extrinsic null hypothesis, where the number of outcomes (n) predicted by the null hypothesis is known, and intrinsic null hypothesis, where estimations of one or more parameters from the data are made to test the hypothesis.

The number of degrees of freedom can be determined for the extrinsic null hypothesis as the number of final outcomes n of the variable, minus one. For the intrinsic null hypothesis, the number of degrees of freedom is expressed as the number of final outcomes n minus the number of estimation parameters, minus one.

As stated previously, the null hypothesis is rejected if the distribution meets a significance requirement (commonly denoted as p , or the p -value). The conventional criterion for statistical significance is set at the 5% level (0.05), or that the null hypothesis is rejected at the $p = 0.05$ level. In the case of Chandra, the significance levels of a source varying are described as possibly variable at $p = 0.05$, and definitely variable at the $p = 0.005$ level.

6.5.2 Chi-Squared Test

The Pearson’s chi-squared (χ^2) test assesses whether an observed frequency differs from an expected or theoretical frequency. It calculates the summed normalisations of the squared differences between observed and expected frequencies. Therefore the larger the difference between the observed and expected distributions, the larger the test statistic becomes.

Equation 6.6 shows the value of the test statistic χ^2

$$\chi^2 = \sum_{i=1}^n \frac{(O_i - E_i)^2}{E_i} \quad (6.6)$$

where O_i is the observed frequency, E_i is the expected or theoretical frequency, and n is the number of cells or number of final outcomes. The χ^2 test follows a discrete uniform distribution, where the possible results are finite and have an equal probability of being observed. The expected frequency therefore acts as a normalisation factor such as a mean. It is often equal to the total observed frequencies N_{total} divided by the number of outcomes n , i.e. $E_i = N_{\text{total}}/n$.

This test statistic is not a level of significance, moreover it needs to be compared with a χ^2 distribution value table which contains the number of degrees of freedom and the related p-values. The p-value can be extrapolated from such a table (or calculated) to give the statistical significance of the distribution and whether the null hypothesis can be rejected.

Alternatively the χ^2 of two distributions can be compared using the difference in χ^2 values $\delta\chi^2$ and the difference in degrees of freedom δdof . If $\delta\chi^2$ rejects the null hypothesis at a certain p-value, then the model with the lowest χ^2 is deemed to be the better test fit.

6.5.3 Kolmogorov-Smirnov Test

The Kolmogorov-Smirnov (KS) test assesses the difference between the empirical distribution function of a uni-variate sample and an expected or theoretical cumulative distribution function (for a one-sided KS test). It works on cumulative distributions and can be utilised as a goodness-of-fit test. The fit calculates the observed maximum absolute deviation from the theoretical cumulative distribution. A cumulative distribution function $F_N(x)$ with N ordered events can be described by

$$F_N(x) = \frac{1}{N} \sum_{i=1}^N I[X_i \leq x] \quad (6.7)$$

where $I[X_i \leq x]$ is an indicator function which sums all of the values X_i which lie below a given value x for N events. For an entire observation the test value x is the last event. Therefore, the one-sample K-S statistic of two different cumulative distribution functions

$F_{N1}(x)$ (observed) and $F_{N2}(x)$ (expected) is given by

$$D = \sqrt{N} \sup_x |F_{N1}(x) - F_{N2}(x)| \quad (6.8)$$

where sup is the supremum, described as the least upper bound value of a set of real numbers (or simply the maximum), of the absolute differences of the two cumulative distributions. This generates a p-value which can be tested against any given significance level to determine whether the null hypothesis can be rejected.

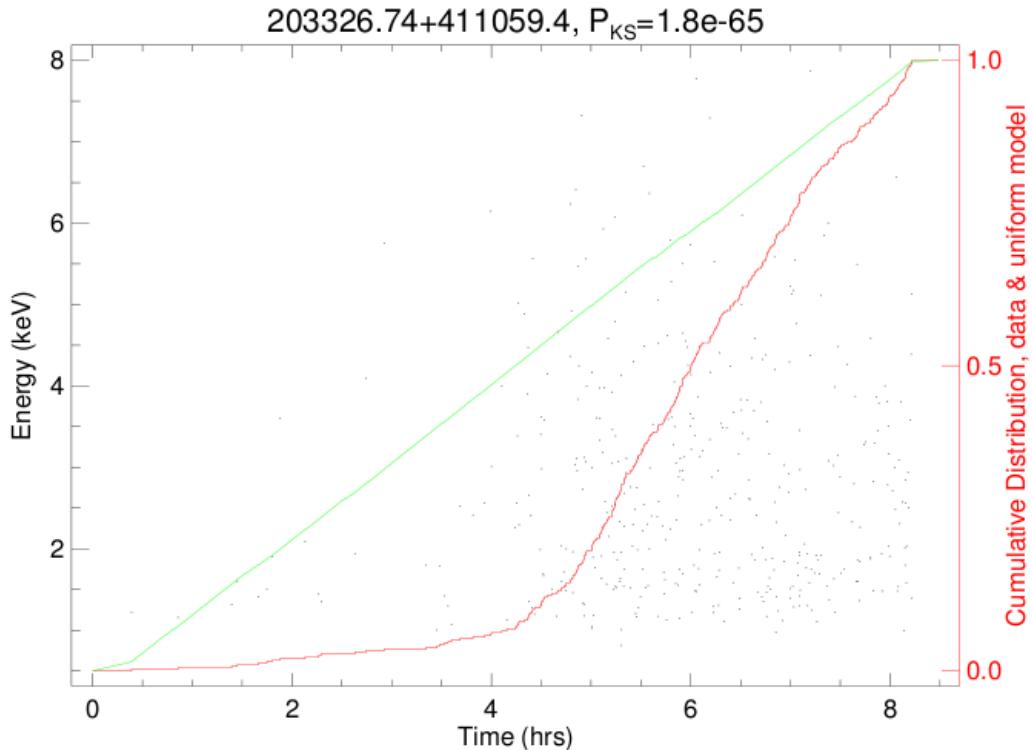


Figure 6.4: A cumulative distribution plot of the individual photon events from the observation of MT91 534, epoch 10960. The black data points refer to the individual photon events and correspond to the energy vertical axis. The green and red curves are the cumulative distribution functions of a constant source model and that of the real Chandra observation of MT91 534. The KS statistic is calculated as the biggest deviation between the real data curve and the model curve.

Figure 6.4 demonstrates the KS statistic (the biggest deviation between the two cumulative distributions) with the example star MT91 534. The maximum difference between the constant model (green curve) and Chandra data (red curve) occurs around 4.5 hours. This equates to a p-value of 1.8×10^{-65} which is highly significant.

6.5.4 Kuiper's Test

Kuiper's test (Kuiper 1960) is similar to the KS test as it evaluates the differences between two cumulative distributions using the same KS statistic as in Equation 6.8. Kuiper's statistic V , assesses the most positive (D^+) and most negative (D^-) deviations between the distributions

$$V = D^+ + D^- = \sqrt{N} \sup_x |F_{N1}(x) - F_{N2}(x)| + \sqrt{N} \sup_x |F_{N2}(x) - F_{N1}(x)| \quad (6.9)$$

Because Kuiper's statistic, V in Equation 6.9, is the sum of D^+ and D^- , it is invariant to any change in the starting point of either distribution. It is as sensitive to the tails of the distribution as it is around the median, whereas the KS statistic is only sensitive around the median. In physical terms, the KS statistic is insensitive to variations at the start and end points of an observation (Broos et al. 2010). Furthermore, the KS statistic is good at finding shifts in a distribution and Kuiper's statistic is powerful at finding spreads in a probability distribution.

For applications to Chandra X-ray datasets, the test of variance needs to access the uniformity of continuous distributions of a single variable which may contain a small number of samples. As with the KS test, the Kuiper test uses a uniform distribution which resembles a constant source as one of the distributions. Any deviations from this constant cumulative distribution at the statistical significance level will indicate a variable source candidate.

The Kuiper test has been used by Paltani (2004) to search for periodic behaviour in Röntgensatellit (ROSAT) X-ray data and the KS and Kuiper statistics are employed in the Chandra Legacy survey for their ability to determine these variations in the X-ray data.

6.5.5 Cash Statistic

The traditional χ^2 maximum likelihood analysis assumes the data is binned, Gaussian and contains at least 5 events per bin. In practical X-ray astronomy, none of these criteria are met and therefore another statistic is needed. Cash (1979) derived a likelihood function

from Poisson distributions by using the product of individual probabilities for each bin i

$$\mathcal{L} = \prod_i \frac{M_i^{D_i}}{D_i!} \exp(-M_i) \quad (6.10)$$

where M_i is the sum of the source and background model amplitudes, and D_i is the observed counts. Taking the logarithm of \mathcal{L} , inverting the signs, dropping the factorial term and multiplying by two, converts the maximum likelihood function into the Cash statistic (Sherpa Documentation for Chandra³):

$$C = 2 \sum_i (M_i - D_i \log M_i) \quad (6.11)$$

The factor of two arises so that when two distributions are compared, δC approximately scales with $\delta\chi^2$ when the number of counts in each bin exceeds 5. The δC can be used in comparison tests but does not evaluate a goodness-of-fit.

6.5.6 Lomb-Scargle Periodograms

One of the most commonly used techniques to find periodic signals in astronomical data is the *periodogram* (Scargle 1982). The classical periodogram takes a time series $\{t_j, j = 1, 2, \dots, N\}$ of a variable e.g. flux; $F(t)$ and transforms the series into frequency via a discrete Fourier transform:

$$\mathcal{F}_F(\omega) = \sum_{j=1}^N F(t_j) \exp(-i\omega t_j). \quad (6.12)$$

The periodogram P of frequency ω , takes the modulus-squared of the Fourier transform:

$$\begin{aligned} P(\omega) &= \frac{1}{N} |\mathcal{F}_F(\omega)|^2 \\ &= \frac{1}{N} \left| \sum_{j=1}^N F(t_j) \exp(-i\omega t_j) \right|^2 \\ &= \frac{1}{N} \left[\left(\sum_j F_j \cos(\omega t_j) \right)^2 + \left(\sum_j F_j \sin(\omega t_j) \right)^2 \right] \end{aligned} \quad (6.13)$$

as defined in the literature (Scargle 1982 and references therein). The periodogram

³<http://cxc.harvard.edu/sherpa4.4/statistics/index.html#cash>

works when the variable F contains a sinusoidal component of frequency ω_0 , which will maximise the summation in Equation 6.13 because the components will be in phase. At other frequencies the components will be randomly positive and negative making the summation relatively small. If the time sampling is evenly distributed, Equation 6.13 translates to

$$P_F(\omega) = \frac{1}{N} \left| \sum_{j=1}^N F_j \exp(-i\omega j) \right|^2 \quad (6.14)$$

where $t_j = j$ and $F(t_j) = F_j$.

The similarity of periodograms and least squares analysis of sine waves to the data was noted by Lomb (1976), who investigated least squares analysis on unevenly sampled data convolved in the frequency domain. Scargle (1982) “slightly modified” the classic periodogram so the new periodogram could perform analysis on unevenly sampled data, be equivalent to least squares analysis of sine waves and would reduce to the classic periodogram when the time samples are even. The new periodogram, which is referred to as the Lomb-Scargle (LS) periodogram because of the influence of both authors’ work on the subject, is defined as

$$P_F(\omega) = \frac{1}{2} \left\{ \frac{\left[\sum_j F_j \cos \omega (t_j - \tau) \right]^2}{\sum_j \cos^2 \omega (t_j - \tau)} + \frac{\left[\sum_j F_j \sin \omega (t_j - \tau) \right]^2}{\sum_j \sin^2 \omega (t_j - \tau)} \right\} \quad (6.15)$$

where τ is defined as

$$\tau = \frac{1}{2\omega} \arctan \left[\frac{\sum_j \sin 2\omega t_j}{\sum_j \cos 2\omega t_j} \right]. \quad (6.16)$$

The periodogram produces a number of peaks of varying power. To determine the significance of any peak, a certain power level must be exceeded. The false alarm probability (FAP) p_0 is a small number, typically the nominated p-value for any investigation. The

power level (z_0) of a peak is related to the FAP p_0 by

$$z_0 = -\ln \left[1 - (1 - p_0)^{1/N} \right]. \quad (6.17)$$

N represents the number of frequencies searched for in LS periodogram. For example: with $N = 30$ and $p_0 = 0.05$, a peak must have a LS power level greater than 6.37 to have a significance of 95%.

Scargle (1982) commented that when the periodogram is used on noisy data, “*surprisingly large spurious spectral peaks can occur and be erroneously taken to indicate the presence of a periodic signal. Hence it is important to critically analyse the statistical significance of a suspected spectral feature, by answering the question: what is the probability that this feature could have arisen from chance (noise) fluctuations?*”. In order to answer this question, Monte Carlo (MC) simulations can be performed to add random noise to the data and then evaluate whether the LS periodogram detects the same periods with the added noise. This method has been used by Hoffman et al. (2012) on XMM–Newton targets to search for coronal X-ray cycles. The simulations modify the data simply by

$$F_{MC}(t_j) = F_j + X_j \quad (6.18)$$

where F_j is the flux observed at time j , and X_j is a random variable from a Gaussian distribution sampled with a zero mean and a standard deviation equal to the uncertainty of the flux value (taken as one standard deviation of the total pointing sample). The LS periodogram is then performed on the MC simulated data and the FAP for each MC run is recorded. Over a large number of MC runs (1000 were used by Hoffman et al. 2012, 10,000 are used in this thesis study), if 68% of the MC runs have a $\text{FAP} < 0.30$ for a given period, then the dataset is said to be periodic (or warrants further investigation) (Hoffman et al. 2012).

In the context of this thesis study and the inter-pointing variability investigations, only targets with known periods will be tested. This eliminates any problems with aliasing frequencies from the LS periodogram, as spurious peaks are ignored and only significant peaks around the quoted period are of interest. Searches for peaks from LS periodogram studies of the MC simulations are conducted around the period $\pm 5\%$ (a small error in the period value). The frequency with the lowest FAP within this small period window is deemed to be the true period of the MC simulation found by the LS periodogram. This

is because of the nature of the unevenly sampled dataset, not necessarily representing the quoted period, which was most likely determined by other methods in other wavelength regimes.

6.6 Chandra X-ray variability of the O star population of Cyg OB2

6.6.1 Variability in Previous Chandra Studies

A previous 97 ks-long, $17' \times 17'$ Chandra Cyg OB2 study by Albacete Colombo et al. (2007a), investigated the variability of 135 of the 1003 detected sources. The authors found 85 X-ray sources ($\approx 8.5\%$ of the total) to have a KS p-value (P_{KS}) < 0.001 and 49 sources with $0.001 < P_{KS} < 0.01$. In the context of OB stars, 26 were detected in the study, 20 O stars and 6 B stars, with only two stars (both B-types) having a $P_{KS} < 0.01$.

Among the 20 detected O stars, two were discussed by the authors in detail. Cygnus OB2 No. 8 is an O6If + O5.5III(f) binary with an orbital period of 21.9 days and evidence of phase-locked X-ray variability (De Becker and Rauw 2005). The observations of Albacete Colombo et al. (2007) only covered 5% of the orbital period and therefore did not detect any X-ray variability.

Cygnus OB2 No. 9 is an O5I+O3.5III binary with an orbital period of 2.355 years (Nazé et al. 2010). Other studies have shown Cyg OB2 No. 9 to be the strongest and most variable non-thermal radio emitter in the Cyg OB2 association (Van Loo 2005). There is no reported X-ray variability from any previous studies using older X-ray observatories (Albacete Colombo et al. 2007a).

6.6.2 Overview of Variable Candidates in the Chandra dataset

The O star dataset from the Chandra Cyg OB2 Legacy survey contains 49 O-type and 3 Wolf-Rayet stars within the Chandra field of view. Twelve sources are found to have a KS statistic (either inter-pointing or intra-pointing KS) below the ‘definitely variable’ condition of $P_{KS} < 0.005$ and an additional 5 sources which do not meet this requirement but have a binary spectral classification and hence the possibility of inter-pointing variability.

The 17 X-ray variable candidates found are listed in Table 6.1. Of these 17 candidates, Schulte 27 and A38 are not included in any further analysis because the background photon

count is equal or greater than any substantial source X-ray photon count for all observed epochs. The rest are split into inter-pointing and intra-pointing variability studies.

Table 6.1: Table of X-ray Variable O-type and WR Candidates

Chandra Name	Common Catalogue Designation	Spectral Classification	Known Period (days)	Reference	KS Statistic
203137.50+411321.1	Schulte 3	O6IV+O9III	4.7464	1	4.223E-01
203213.84+412711.4	Schulte 4	O7III			6.781E-07
203222.42+411819.0	Schulte 5	O7I+O9Pe/WN9	6.6	2	1.358E-129
203231.54+411408.1	MT91 267	O7.5III	15.511	3	2.095E-07
203234.87+405617.0	A38	O8V			1.089E-16
203308.77+411318.7	Schulte 22	O3If+O6V	4.61	4	1.958E-01
203310.73+411508.2	Schulte 9 (MT91 431)	O5I+O3.5III	2.355 yr (1.22/5.6)	5, (6)	4.908E-11
203315.07+411850.5	Schulte 8A (MT91 465)	O6I+O5.5III	21.908	7	0.000E+00
203323.48+410912.6	MT91 516	O5.5V	few weeks	4	3.121E-11
203359.56+411735.5	Schulte 27	O9.5V+BOV			4.541E-01
203408.52+413659.3	Schulte 11 (MT91 734)	O5If+B0V	72.4	3	1.583E-06
203409.51+413413.9	Schulte 75	O9V			4.946E-17
203421.95+411701.5	Schulte 73	O8III+O8III			9.041E-05
203429.60+413145.3	MT91 771	O7V+O9V	2.8635	1	2.544E-01
203547.08+412244.7	WR 146	WC6+O8III	~ 300 yr	8	2.733E-02
203206.26+404829.6	WR 145	WN7o/CE+O7V((f))	22.54977	9	5.754E-07 ^a
203326.74+411059.4	MT91 534	O7.5V			1.803E-65 ^a

List of X-ray variable O stars from the Chandra survey. The first 15 sources demonstrate variability on the inter-pointing scales with $P_{\text{KS}} < 0.005$ over all pointings, or binary spectral classifications. The last two sources display intra-pointing variability with $P_{\text{KS}} < 0.005$ over a single pointing.

All KS statistics are ‘MERGED’ KS statistics (inter-pointing) unless otherwise stated. ^a KS statistics from the single epoch (intra-pointing).

Notes. Spectral Classifications are from the Chandra survey fits headers. Periods are given in days unless otherwise stated.

References. (1) Kiminki et al. (2012), (2) Rauw et al. (1999) (and many others), (3) Kobulnicky et al. (2012), (4) Rauw et al. (2014), (5) Nazé et al. (2010), (6) Pignksi and Kołaczkowski (1998), (7) De Becker et al. (2004), (8) Dougherty et al. (1996), (9) Muntean et al. (2009).

6.6.3 Inter-pointing X-ray Variability

Of the 15 O-type and WR stars which qualified as X-ray variable candidates, 13 show possibilities of variance over inter-pointing observations from either their spectral class and / or KS statistic over multiple epochs. Nine sources have the spectral classification of binaries and four of single stars. LS periodograms complemented with MC simulations are performed on each target to determine whether any significant periods are present. In the case of an established period from the LS periodogram and MC simulations, the light curves are folded on the determined period.

Table 6.2 shows the timescales between the first and last observations and the number of epochs available for those timescales. Early epochs are not included because of the large difference in time (~ 6 years) compared to the Chandra Cyg OB2 Legacy observations.

Table 6.2: Table of Epoch Timescales for Variable Candidates

Chandra Name	Common Catalogue Designation	Timescale of Epochs (days)	Number of Epochs
203137.50+411321.1	Schulte 3	27.51	6
203213.84+412711.4	Schulte 4	21.55	4
203222.42+411819.0	Schulte 5	11.43	6
203231.54+411408.1	MT91 267	21.88	5
203308.77+411318.7	Schulte 22	12.69	8
203310.73+411508.2	Schulte 9 (MT91 431)	2.61	4
203315.07+411850.5	Schulte 8A (MT91 465)	2.61	4
203323.48+410912.6	MT91 516	7.88	4
203359.56+411735.5	Schulte 27	10.73	4
203408.52+413659.3	Schulte 11 (MT91 734)	26.63	5
203409.51+413413.9	Schulte 75	25.36	4
203421.95+411701.5	Schulte 73	10.73	4
203429.60+413145.3	MT91 771	25.36	4
203547.08+412244.7	WR 146	13.91 hours	2
203206.26+404829.6	WR 145	8.72 hours	1
203326.74+411059.4	MT91 534	7.88	4

Note. Timescales and number of epochs do not include early epochs: 4358, 4501, 4511, 7426, which are generally six years earlier than the Chandra Cyg OB2 Legacy observations.

The inter-pointing, long term variability results for each of the O-type and WR stars are now presented.

203137.50+411321.1 - Schulte 3

Schulte 3 is an O6IV+O9III binary of type SB2/EB with an orbital period of 4.7464 days (Kiminki et al. 2012b). Figure 6.5 shows the binned Chandra pointings adjacent to one another, with relatively constant fluxes over all the pointings. This is consistent with the KS statistic of 4.223E-01, i.e. not significant (see Table 6.1). The reported pile-up rate for Schulte 3 is $\sim 3\%$ with errors on the fluxes of individual pointings between 6 and 15% (Rauw et al. 2014). There are hints of intra-pointing variability in the form of a few ‘spikes’ in the fluxes in the epoch 10953, which are discussed later (Section 6.6.4, page 270). The extreme peaks in epoch 10962 are most likely calibration errors.

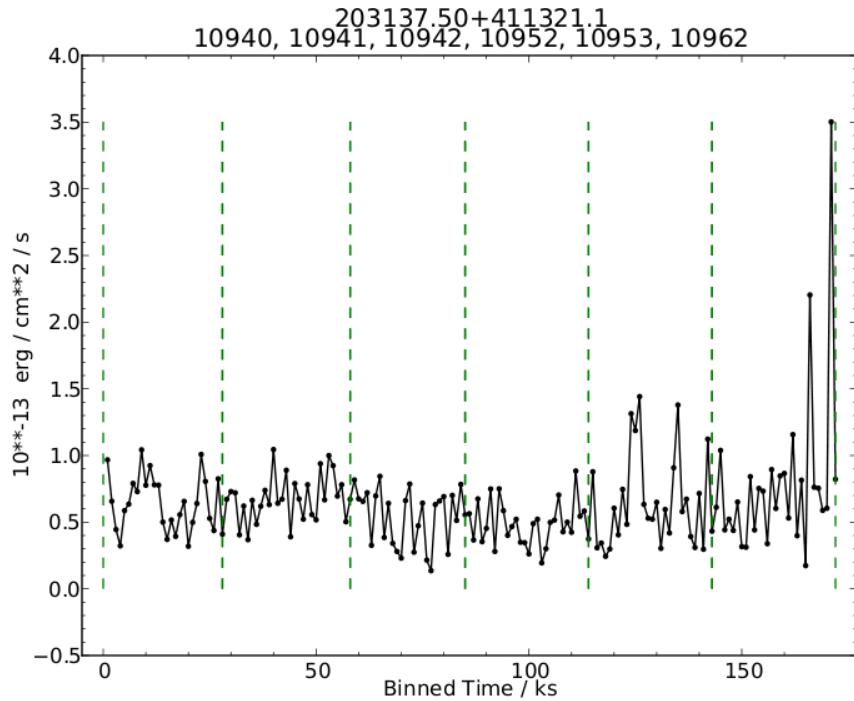


Figure 6.5: Binned light curve of Schulte 3 with the observed epochs: 10940, 10941, 10942, 10952, 10953, 10962. The y-axis is the energy flux in 10^{-13} erg $\text{cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The fluxes are relatively constant for all epochs, with a few spikes in epochs 10953 and 10962.

The LS periodogram deployed on the dataset found a singular significant peak around $4.5 \mu\text{Hz}$ (2.57 days), with the closest peak to the reported period still well below the FAP line.

Results show that only 1.77% of the MC simulations around the reported period ($\pm 5\%$) have a FAP < 0.30 , ruling out the possibility of detecting the orbital period through the variability of this Chandra dataset. Therefore, there is no evidence of inter-pointing

X-ray variability arising from the orbital motion of Schulte 3.

203213.84+412711.4 - Schulte 4

Schulte 4 is a single O7III star (Chandra catalogue classification), included in this study because of a KS statistic of 6.781E-07. Figure 6.6 shows the binned Chandra pointings adjacent to one another, with the majority of the points lying close to $0.00 - 0.01 \times 10^{-13}$ erg cm $^{-2}$ s $^{-1}$. Epoch 10962 is not included here because of the high background photon counts for that observation. The significant KS statistic for inter-pointing variability is a result of the background counts in epoch 10962, with the single epoch $P_{KS} = 3.211E-02$ ('possibly variable'). The single epoch P_{KS} for 4501 is 8.495E-01, suggesting that the prominent spike residing in that epoch is an erroneous calibration error.

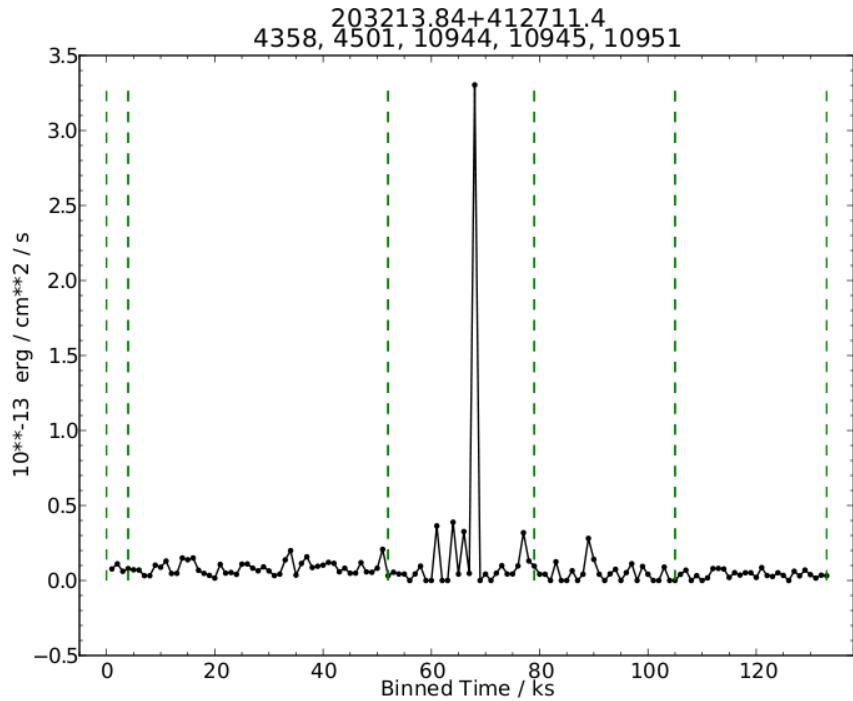


Figure 6.6: Binned light curve of Schulte 4 with the observed epochs: 4358, 4501, 10944, 10945, 10951. Epoch 10962 is not shown here because of high background noise. The y-axis is the flux in 10^{-13} erg cm $^{-2}$ s $^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The fluxes are relatively constant for all epochs, with a few spikes in epochs 4501 and 10944. The prominent spike in epoch 4501 is most likely an error in the calibration.

As a single star with no known associated period of any origin, showing constant inter-pointing fluxes as can be seen in Figure 6.6 lead to the conclusion that Schulte 4 is not variable on the inter-pointing timescale.

203222.42+411819.0 - Schulte 5

Schulte 5 is an O7I+Ofpe/WN9 binary with a period of 6.6 days (Rauw et al. 1999). Figure 6.7 shows the binned Chandra pointings adjacent to one another. Schulte 5 suffers from a high pile-up percentage of $\sim 34\%$ (Rauw et al. 2014) which affects the flux significantly. Moreover, it can account for much of the variability seen in Figure 6.7. Epoch 10956 has the highest mean flux value of around $0.90 \cdot 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ which corresponds to an error of $0.30 \cdot 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ at the 34% level. Epoch 10952 has a lower mean value ($0.50 \cdot 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$) beyond these errors, suggesting the possibility of some of the variability being real. Epochs 10951 and 10953 are omitted due to extreme scattering, possibly as a result of pile-up.

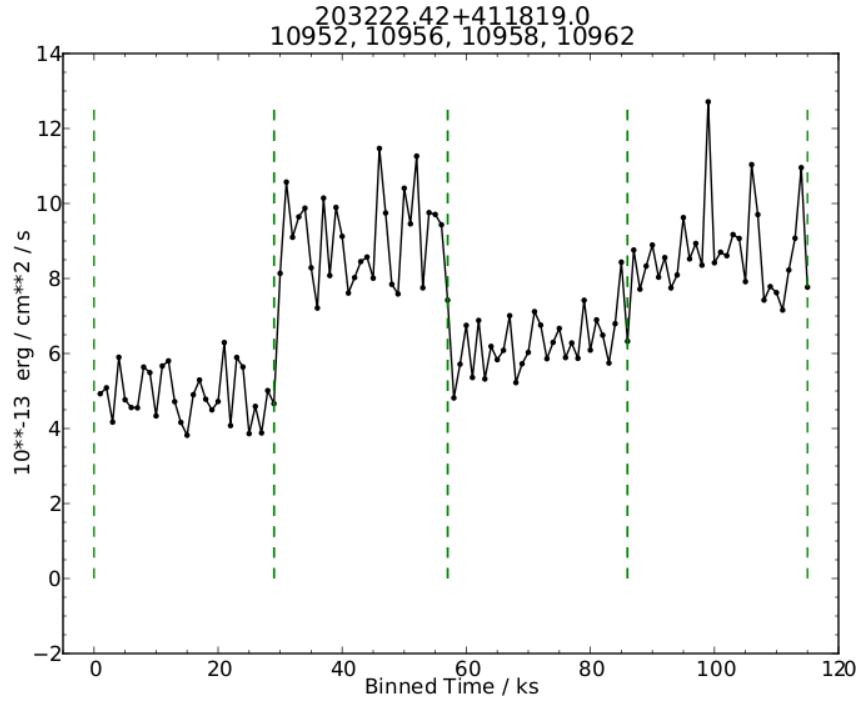


Figure 6.7: Binned light curve of Schulte 5 with the observed epochs: 10952, 10956, 10958 and 10962. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Clear variability can be seen over all epochs, however, with the high pile-up of $\sim 34\%$, most of this *can* be accounted for by the pile-up.

LS periodograms are again used to detect the cited period of 6.6 days (Rauw et al. 1999). Figure 6.8 shows the LS periodogram of the Schulte 5 Chandra dataset. The blue line traces the period of 6.6 days and the red dashed line shows the FAP = 0.005 confidence level. The strongest peak is only slightly offset from the true period line and is very significant with a Lomb power value of over 40. The other peaks of similar strength

are most likely aliases of the period peak, as the flux variability seen in Figure 6.7 is quite pronounced. The peak profile shapes are also similar leading to the same conclusion.

Because of the poor phase coverage of the Chandra pointings, the flux values (y-axis) are randomly shuffled and then a LS periodogram is deployed once more to determine whether the ‘aliases’ are due to sampling effects, purely related to the x-axis values. The resulting periodograms did not reveal peaks of similar strength or positions of the aliases seen in Figure 6.8, ruling out data sampling aliases.

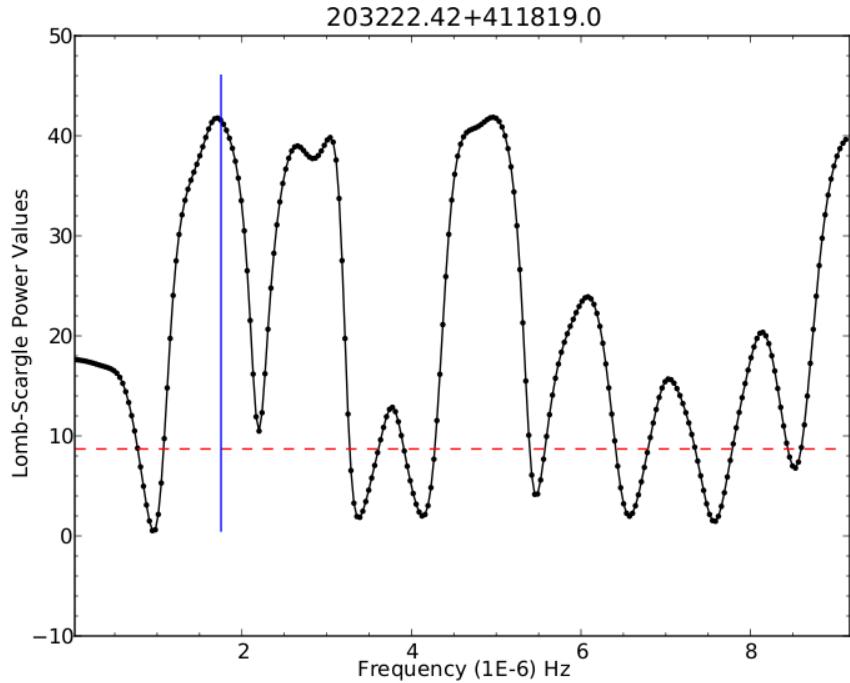


Figure 6.8: Lomb-Scargle periodogram of Schulte 5. The y-axis represents the Lomb power value (z_0) and the x-axis is the frequency in μHz . The blue line denotes the reported orbital period of 6.6 days, and the red dashed line shows the $\text{FAP} = 0.005$. The peak near the true period is the strongest in the periodogram and is above the FAP threshold.

MC simulations are conducted seeking out the highest Lomb power value between the reported period $\pm 5\%$ and the resulting histograms are presented in Figure 6.9. The strength of the peaks in the LS periodogram (Figure 6.8) are also obvious in the histogram with 99.80% of MC runs having a $\text{FAP} < 0.05$, and 99.96% having a $\text{FAP} < 0.30$ around the reported period of 6.6 days (Rauw et al. 1999).

There are a finite number of frequencies which the LS periodogram searches for, and because this study limited the MC simulations to search between the reported period $\pm 5\%$, the most frequent period within these limits can be extracted from the MC runs. Figure 6.10 shows a histogram of the most significant peaks from each MC run around

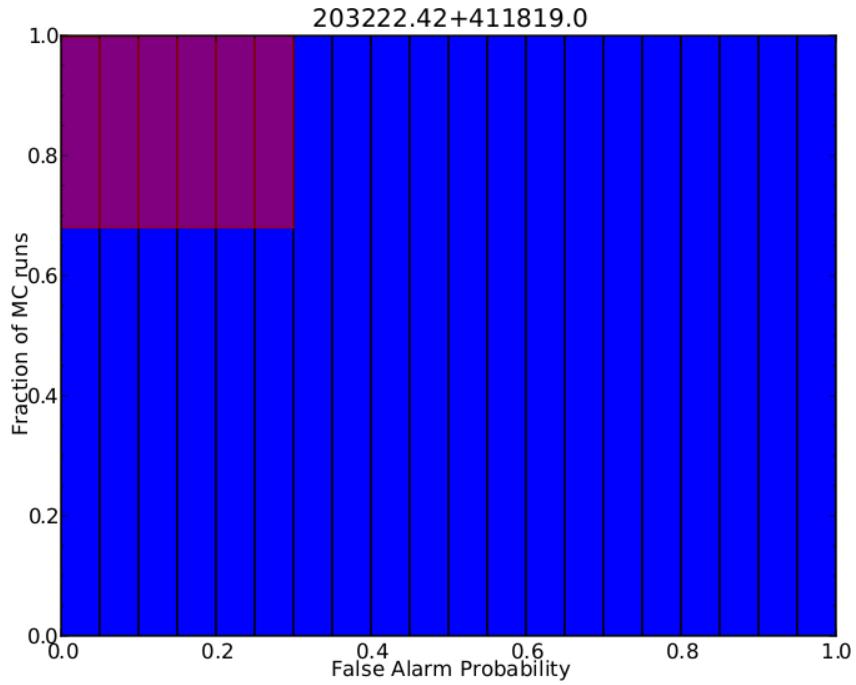


Figure 6.9: Cumulative histogram of the MC simulations for Schulte 5. The y-axis shows the fraction of the total (10,000) MC simulations and the x-axis shows the FAP. The x-axis is binned to 0.05 FAP per bin. The red box in the top left corner represents the area when 68% of the MC simulations have an $FAP < 0.30$, i.e. a positive result. The cumulative histogram of the MC runs breaks the red box marker for a positive result for variability for Schulte 5.

the cited period of 6.6 days. The majority of the MC runs (44%) are most significant at a period of 6.87 days.

However, Figure 6.10 only shows the fraction of MC simulations that occurred at each frequency, and does not account for the relative strength or corresponding FAP of each MC run. Almost all (99.80%) of the MC runs have a $FAP < 0.05$. Therefore, a weighted mean can still produce a truer representation of the period found by the LS periodograms and MC simulations, utilising the exact FAP found from each MC run (however small). The weighted mean is found to be 6.64 days, which is in agreement with the value for the period found by Rauw et al. (1999). The errors on this period are defined by the search window for the MC runs of $\pm 5\%$.

Finally, the light curve folded on the found period of 6.64 days is shown in Figure 6.11. Visually, the variability is clearly modulated on the period found by this study, with the major argument against being the high pile-up fraction quoted by Rauw et al. (2014). However, this pile-up fraction is for on-axis sources, and most if not all massive stars are observed off-axis, leading to lower pile-up estimations (Facundo Albacete-Colombo;

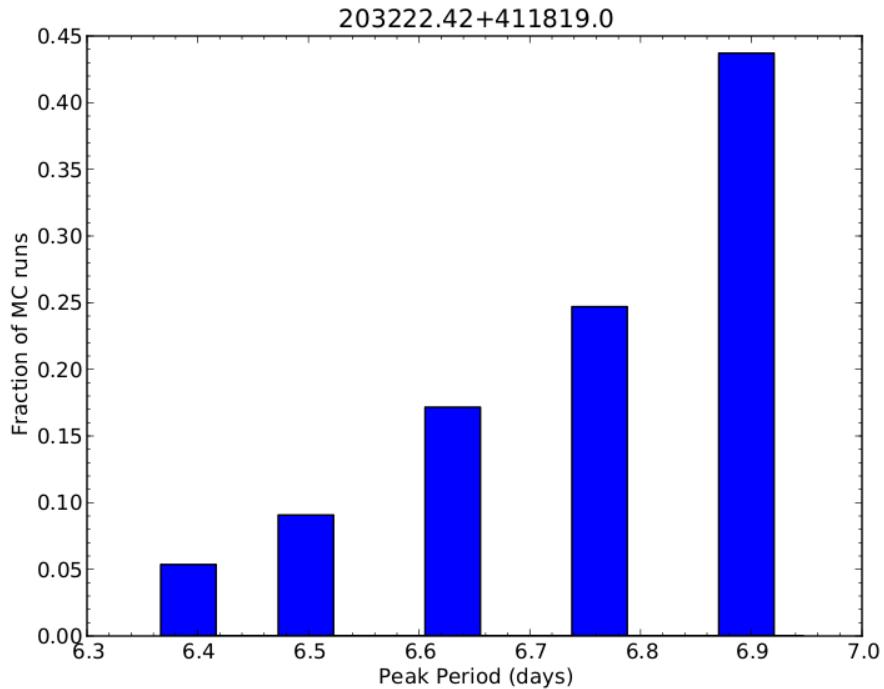


Figure 6.10: Histogram of the strongest periods found from each MC simulation for Schulte 5. The y-axis shows the fraction of the total (10,000) MC simulations and the x-axis shows the strongest frequencies as periods (units of days). The largest histogram bar is at 6.87 days.

private communication). The seemingly spurious epochs of 10951 and 10953 which are omitted from the periodic study may be a result of pile-up. However, the LS periodogram, MC simulations and visual folded light curve offer strong evidence that Schulte 5 shows signs of inter-pointing X-ray variability beyond the pile-up issues, and evidence that this is locked on the orbital phase of the binary.

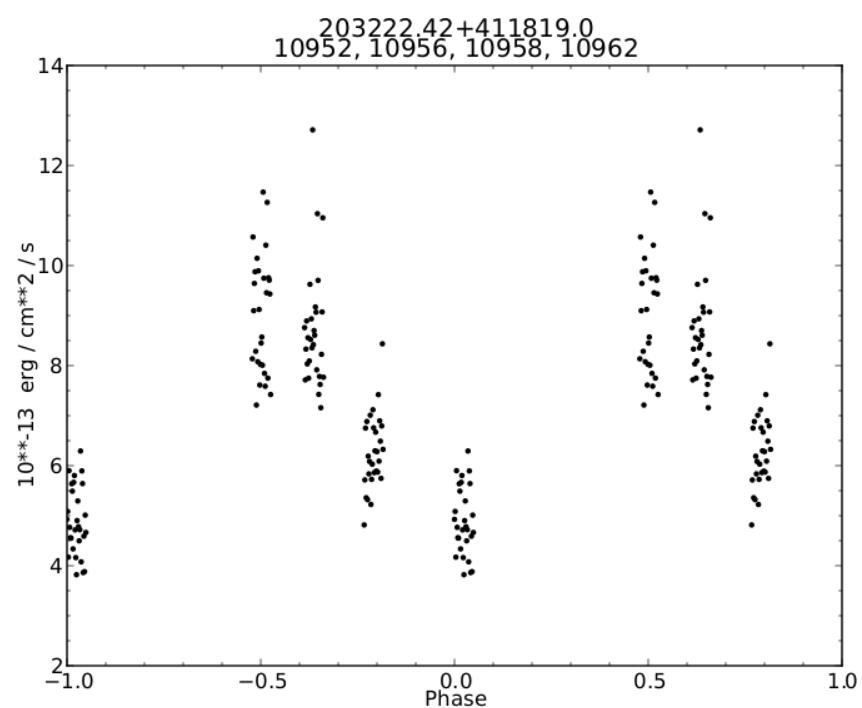


Figure 6.11: Light curve of Schulte 5 folded onto the period found by this study of 6.64 days.

203231.54+411408.1 - MT91 267

MT267 has a spectral classification of O7.5III (from the Chandra catalogue), but has a recently discovered period of 15.511 days (Kobulnicky et al. 2012). Figure 6.12 shows the binned Chandra pointings adjacent to one another, with variability on long and short timescales visible. The later epochs 10956 and 10958 have a lower mean flux than that of the earlier epochs. Spikes are seen in most epochs, the most significant of which occurs in the 4511 epoch around 60 ks. This peak lasts for 5 ks above the neighbouring flux levels, strengthening the argument that it is not an anomaly. This possible intra-pointing variability will be discussed later (Section 6.6.4).

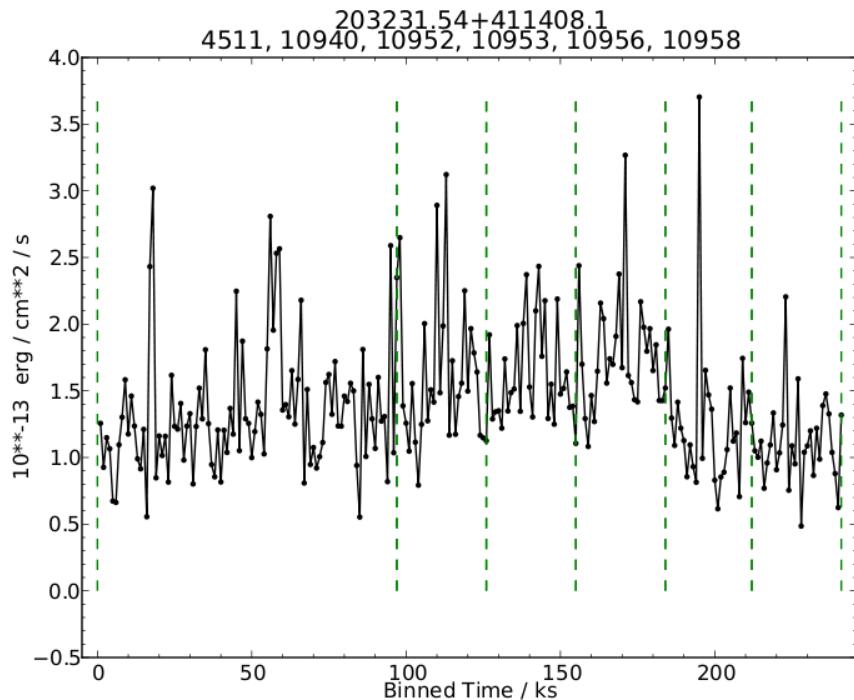


Figure 6.12: Binned light curve of MT91 267 with the observed epochs: 4511, 10940, 10952, 10953, 10956 and 10958. The y-axis is the flux in 10^{-13} erg cm^{-2} s^{-1} and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Variability can be seen on both long and short timescales. Epochs 10956 and 10958 show a lower general flux and a number of spikes are present in many epochs, notably in epoch 4511.

Periodical analysis using the LS periodograms is used to investigate whether the cited period can be recovered by Fourier techniques. Figure 6.13 shows the LS periodogram of MT91 267 with the blue line representing the period of 15.511 days and the red dashed line showing the FAP = 0.005 confidence level. The reported period of 15.511 days (Kobulnicky et al. 2012) lies between two local maxima, both above the FAP line, with the longer period

being closer to the period. Once again as with Schulte 5, many aliases of similar strength are present. A couple of aliases have the same morphology as the peak near the cited period, with a distinct double peak profile.

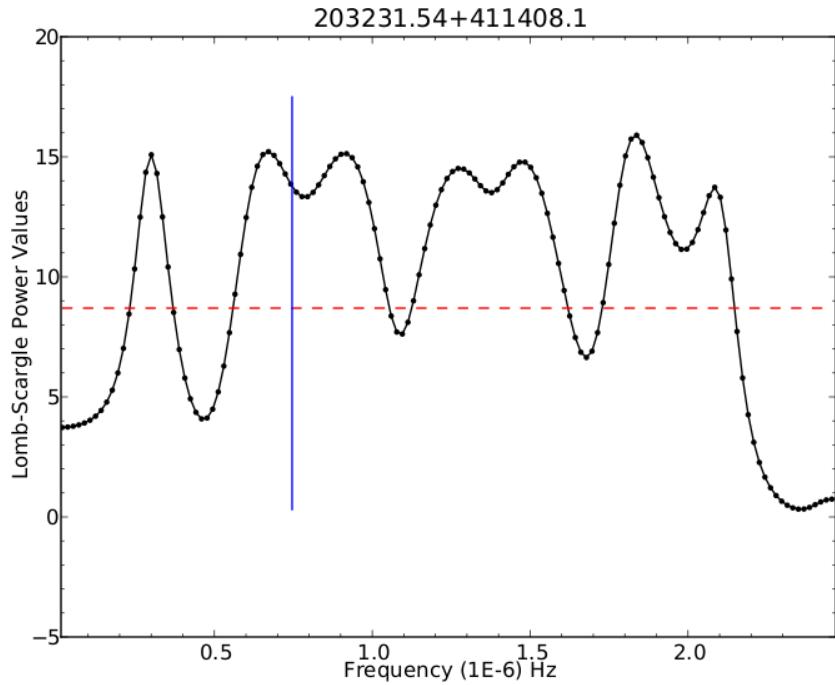


Figure 6.13: Lomb-Scargle periodogram of MT91 267. The y-axis represents the Lomb power value (z_0) and the x-axis is the frequency in μHz . The blue line denotes the reported period of 15.511 days, and the red dashed line shows the $\text{FAP} = 0.005$. This period lies between two local maxima with the longer period (smaller frequency) being more preferred due to proximity.

The results from the MC simulations are displayed in Figure 6.14. For MT91 267, 79.91% of the simulations have a $\text{FAP} < 0.30$, indicating a positive result for the period of 15.511 days $\pm 5\%$ error.

In order to determine the most frequent period found by the MC studies, the strongest peak from each MC run is recorded in a histogram (Figure 6.15). The largest fraction of the total number of MC runs is clearly around 15.9 days with around 93% of the simulations. The weighted mean, with the FAP values for each run representing the weights gives the period to be at 15.95 days which is slightly larger than the value of 15.511 days found by Kobulnicky et al. (2012), but within the error of this investigation of $\pm 5\%$ (defined by the search window for the MC simulations).

The pile-up fraction for MT91 267 is considered to be $< 8\%$ (Rauw et al. 2014), ruling out pile-up as a possible cause for the variability. The LS periodogram and MC analysis

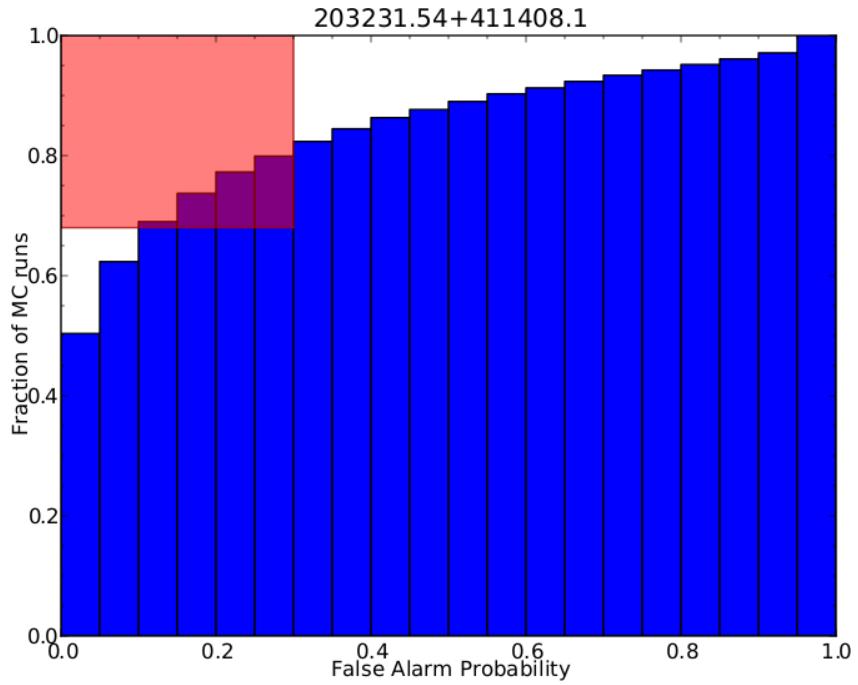


Figure 6.14: Cumulative histogram of the MC simulations for MT91 267. The y-axis shows the fraction of the total (10,000) MC simulations and the x-axis shows the FAP. The x-axis is binned to 0.05 FAP per bin. The red box in the top left corner represents the area when 68% of the MC simulations have an $FAP < 0.30$, i.e. a positive result. The cumulative histogram of the MC runs breaks the red box marker for a positive result for variability for MT91 267.

provide strong evidence that the variability arises from the orbital motion of the binary MT91 267. The light curve folded on the period found by this investigation is shown in Figure 6.16

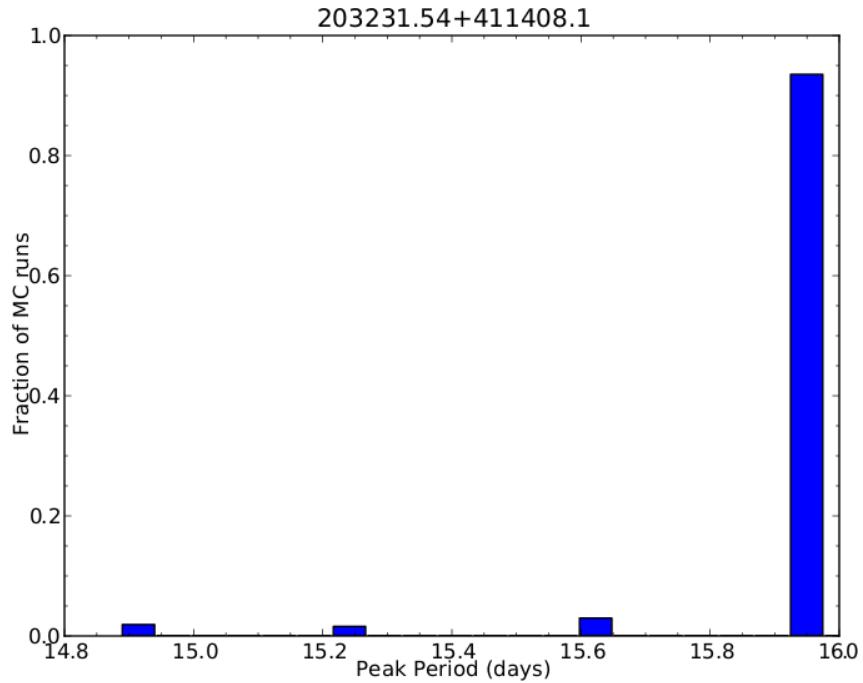


Figure 6.15: Histogram of the strongest periods found from each MC simulation for MT91 267. The y-axis shows the fraction of the total (10,000) MC simulations and the x-axis shows the strongest frequencies as periods (units of days). The largest histogram bar is at 15.9 days.

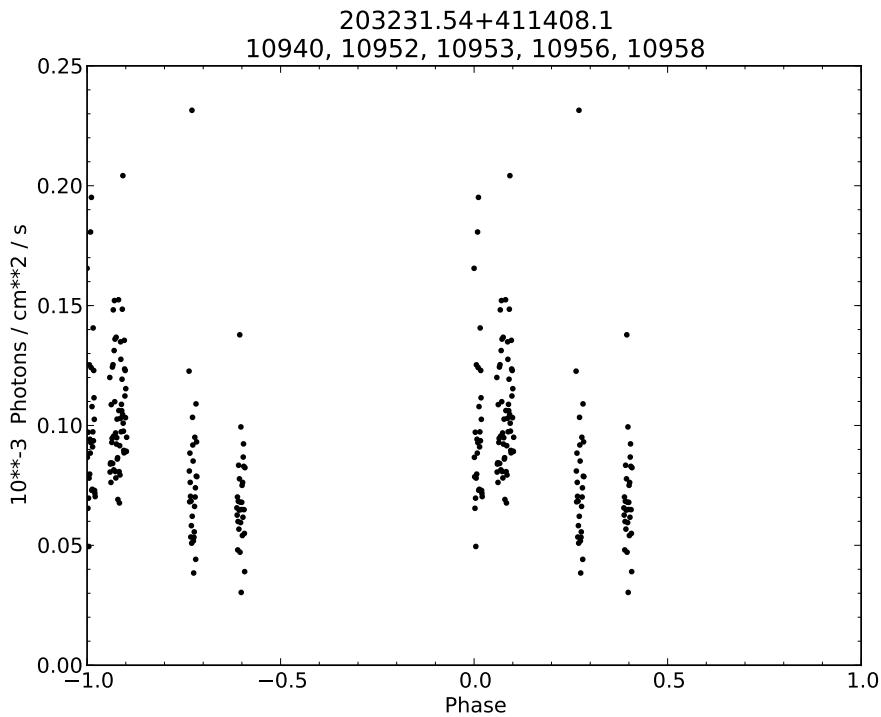


Figure 6.16: Light curve of MT91 267 folded onto the period found by this study of 15.95 days.

203308.77+411318.7 - Schulte 22

Schulte 22 is a massive binary of spectral type: O3If+O6V, with a period of 4.61 days quoted by Rauw et al. (2014), which combined the Chandra and XMM-Newton datasets. The Chandra data set contains initially nine observed pointings, however, two of these epochs (10957/58) suffer from high background counts and are excluded from further analysis. Figure 6.17 shows the binned Chandra pointings adjacent to one another. Small variations can be seen between pointings, particularly in epoch 10953 and 10955. A small 4 ks spike is noted around the 80 ks mark in epoch 4511, which has a single epoch $P_{\text{KS}} = 0.015$ (possibly variable) and will be discussed later (Section 6.6.4).

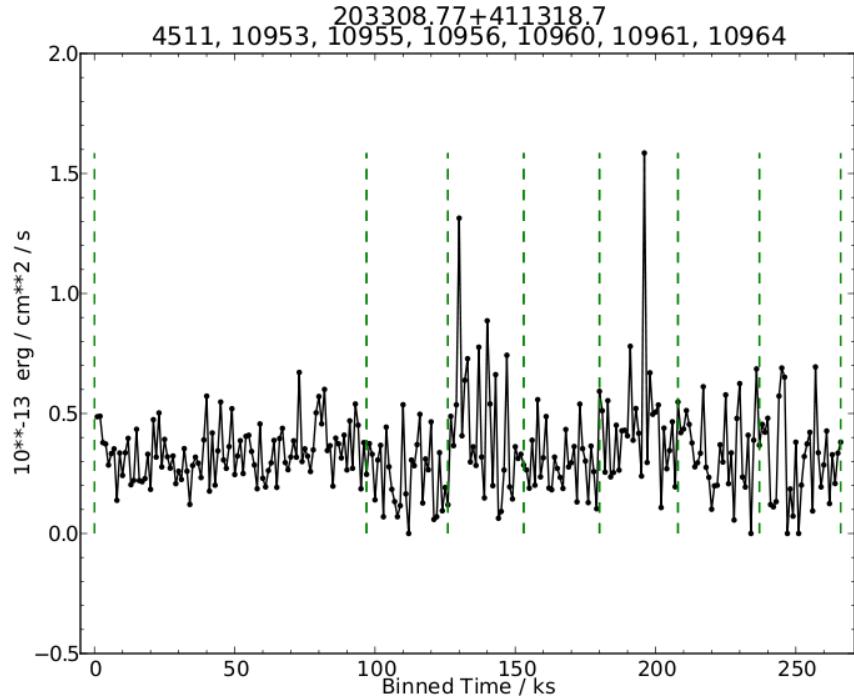


Figure 6.17: Binned light curve of Schulte 22 with the observed epochs: 4511, 10953, 10955, 10956, 10960, 10961 and 10964. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Along with a few spikes in flux, some small variations are apparent between epochs.

Excluding epochs 4511 for the high number of period cycles elasped between then and the Legacy observations, and 10957/58 for the high background counts, LS periodogram analysis on Schulte 22 is presented in Figure 6.18. The strongest peak in the periodogram coincides with the period of 4.61 days and is above the $\text{FAP} = 0.005$ significance line.

The MC analysis for Schulte 22 is shown in Figure 6.19, with 23.77% of the simulations having a $\text{FAP} < 0.30$. This does not meet the requirements imposed by this study, and

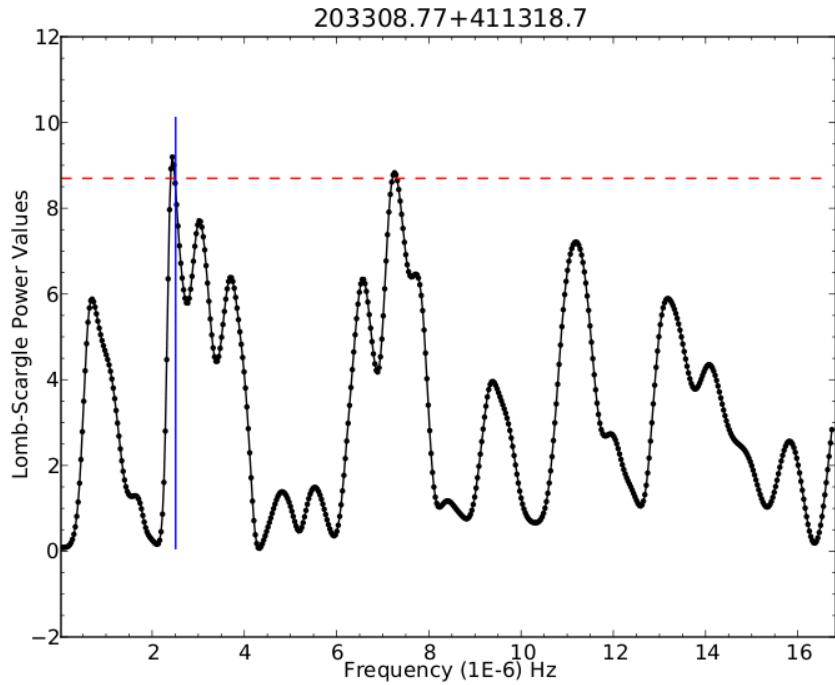


Figure 6.18: Lomb-Scargle periodogram of Schulte 22. The y-axis represents the Lomb power value (z_0) and the x-axis is the frequency in μ Hz. The blue line denotes the reported period of 4.61 days, and the red dashed line shows the FAP = 0.005. The strongest peak in the periodogram is only slightly offset from the reported period of 4.61 days.

therefore indicates a negative result. The weighted average from the MC analysis finds the period to be 4.72 days.

Despite the negative result from the MC analysis, the LS periodogram did suggest that the strongest frequency for periodicity is the reported period, with a high significance. Additional observations (with Chandra or XMM-Newton) are required to confirm whether this X-ray variability is modulated on the quoted period. Rauw et al. (2014) argue that this can not be the orbital period as it does not correlate with the visual components of the system. Instead they propose that Schulte 22 is an oblique magnetic rotator like θ^1 Ori C, but without the hard X-ray emission component (2 - 8 KeV). They calculate the rotation period to lie between 3.3 and 9.21 days.

The light curve from the Chandra data is folded onto the reported period of 4.61 days and shown in Figure 6.20. The conclusion here is that Schulte 22 shows small inter-pointing variability, with suggestions that this is fixed on the cited period with the possibility that this is the rotation period, but without statistical confirmation.

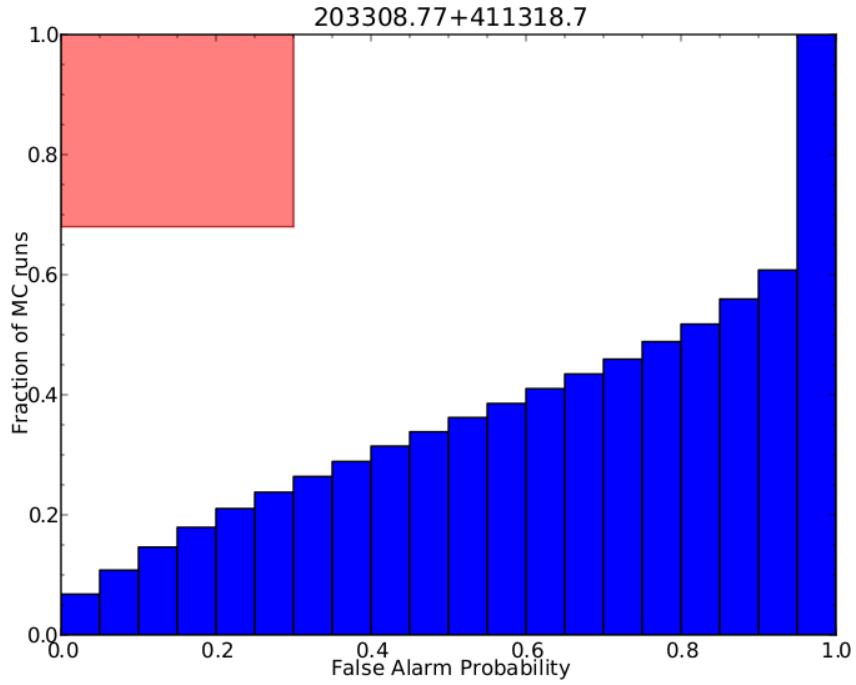


Figure 6.19: Cumulative histogram of the MC simulations for Schulte 22. The y-axis shows the fraction of the total (10,000) MC simulations and the x-axis shows the FAP. The x-axis is binned to 0.05 FAP per bin. The red box in the top left corner represents the area when 68% of the MC simulations have an $FAP < 0.30$, i.e. a positive result. The cumulative histogram of the MC runs does not break the red box required by this study.

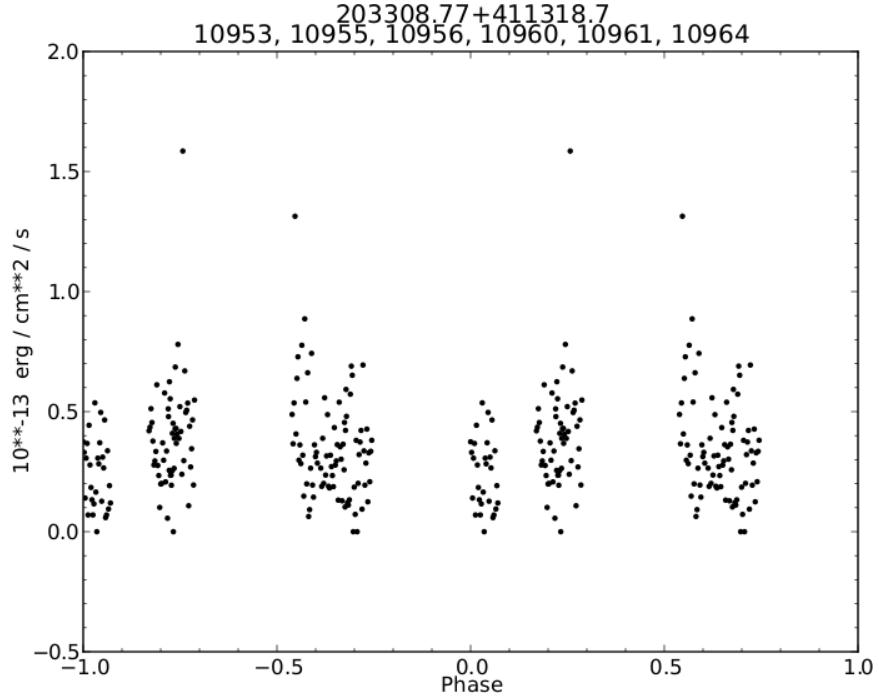


Figure 6.20: Light curve of Schulte 22 folded onto the reported period of 4.61 days.

203310.73+411508.2 - Schulte 9

Schulte 9 is an O5I+O3.5III binary with an orbital period of 2.355 years (Nazé et al. 2010). However, a previous study by Pigulski and Kołaczkowski (1998) found variability associated with two different periods of 1.22 and 5.6 days, the origins of which are unknown. Figure 6.21 shows the binned light curve for five epochs. The latter four epochs all have a smaller flux than the 4511 epoch which has an average flux of $\sim 5 \times 10^{-13}$ erg cm $^{-2}$ s $^{-1}$ across the epoch. Rauw et al. (2014) report a pile-up fraction for Schulte 9 of $\sim 18\%$, which is large and can account for much of the variability seen in Figure 6.21. Although as previously stated these fractions are for on-axis sources, whereas most of the massive stars are off-axis.

The number of observed epochs are too few to provide enough phase coverage to conduct searches for the orbital period of 2.355 years, but these suggested variable periods of 1.22 and 5.6 days can be investigated. The early epoch 4511 is excluded as to not to contaminate the light curves due to the many phase cycles between the two groups of pointings.

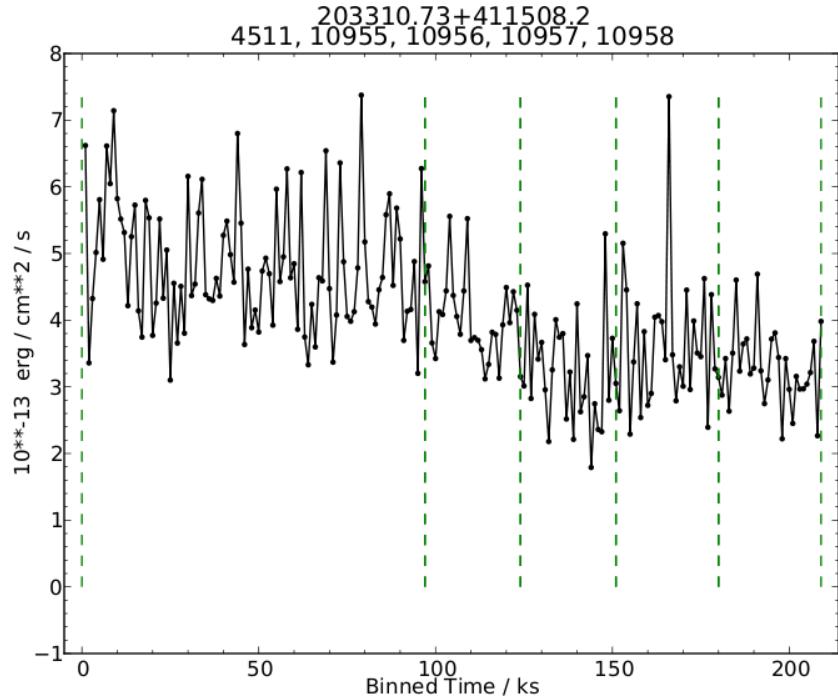


Figure 6.21: Binned light curve of Schulte 9 with the observed epochs: 4511, 10955, 10956, 10957, 10958. The y-axis is the flux in 10^{-13} erg cm $^{-2}$ s $^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Some variability is seen between epochs, but is most likely due to pile-up.

LS periodograms are performed on both of the suspected periods and the MC analysis for both periods demonstrates the non-detection of variability on these periods, with 0.01% of the MC runs having a FAP < 0.30 for 1.22 days and 0.02% of the MC runs having a FAP < 0.30 for 5.6 days.

Given the high pile-up fraction and the very negative results from the LS periodogram and MC analysis, the conclusion is that Schulte 9 does not demonstrate X-ray variability on the timescale of days.

203315.07+411850.5 - Schulte 8A

Schulte 8A is an O6I+O5.5III binary with a period of 21.908 days (De Becker et al. 2004). This is a very bright X-ray source where all of the epochs average over 9.0×10^{-13} erg cm $^{-2}$ s $^{-1}$, reaching an average of 14.0×10^{-13} erg cm $^{-2}$ s $^{-1}$ in epoch 10958. Because Schulte 8A is one of the brightest X-ray sources in this survey, the pile-up fraction is very high with Rauw et al. (2014) citing 60% pile-up. This amount of pile-up can easily account for all of the variability in the binned light curve (Figure 6.22).

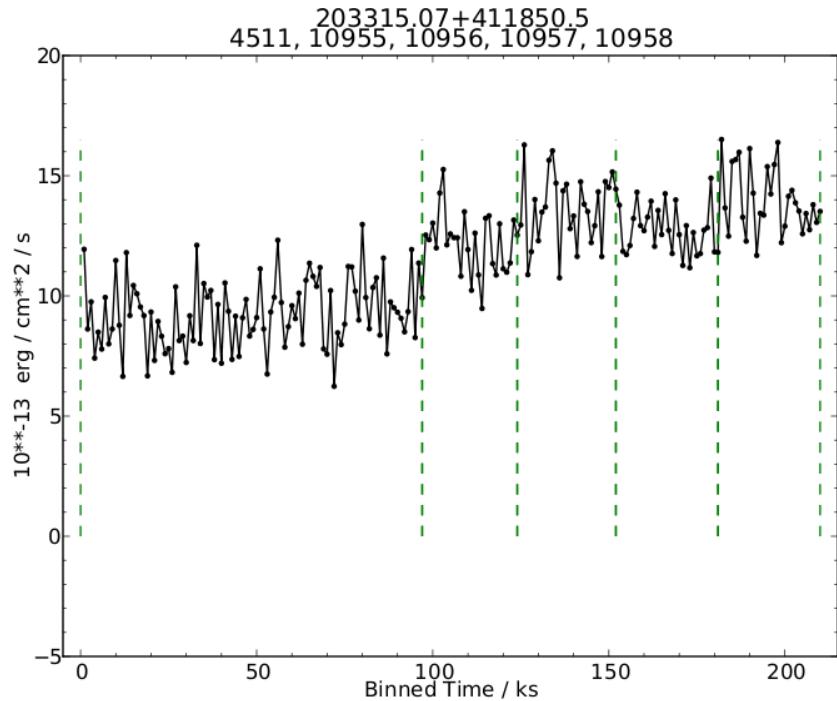


Figure 6.22: Binned light curve of Schulte 8A with the observed epochs: 4511, 10955, 10956, 10957, 10958. The y-axis is the flux in 10^{-13} erg cm $^{-2}$ s $^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Some variability is seen between epochs, but is most likely due to pile-up.

Omitting the early epoch 4511, LS periodograms are performed on Schulte 8A in an

attempt to see through the pile-up and detect the period. However, given the long period of 21.908 days (De Becker et al. 2004), the periodogram does not find any periods in the region of 21 days.

The MC analysis also shows no detection of the long period, with 0.00% of the MC runs having a FAP < 0.30. This leads to the conclusion that Schulte 8A shows no evidence of variability on the period of 21.908 days, with the variability seen in Figure 6.22 probably a result of pile-up.

203323.48+410912.6 - MT91 516

MT516 has the spectral classification of a single star (O5.5V), but has new evidence that it may possibly be part of a binary system with a period of a few weeks (Rauw et al. 2014). However, a radial velocity study by Kiminki et al. (2007) states that there is only an 11.5% chance of MT91 516 being a binary. The binned light curve of MT91 516 is given in Figure 6.23 with small amounts of variability present. With ‘only a moderate pile-up’ (Rauw et al. 2014), this variability could be real fluctuations.

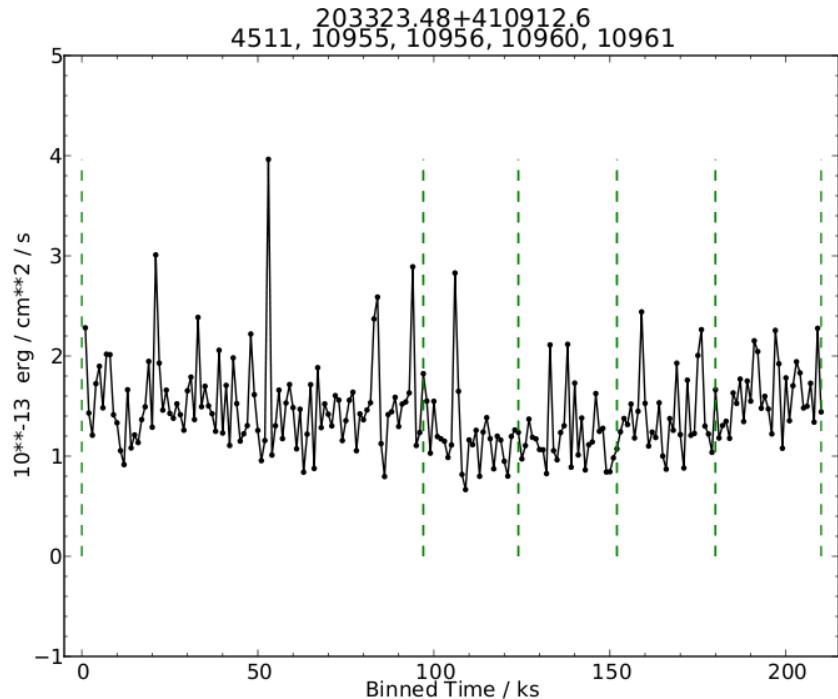


Figure 6.23: Binned light curve of MT91 516 with the observed epochs: 4511, 10955, 10956, 10960, 10961. The y-axis is the flux in 10^{-13} erg cm^{-2} s^{-1} and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. Small amounts of variability is seen between epochs.

As the proposed period is on the order of a few weeks with a large possible range,

no further periodic studies are performed. This is because the many aliases produced by the LS periodogram, make searching for an unknown period difficult. The conclusion for MT91 516 is that small variability is seen, which is unusual for a main sequence O-type star.

203408.52+413659.3 - Schulte 11

Schulte 11 is an O5If+B0V binary with a period of 72.4 days (Kobulnicky et al. 2012). Figure 6.24 shows all the epochs stacked against one another, with a moderate amount of variability present. The first epoch, 10947, shows a large amount of scatter as does epoch 10966. Epoch 10949 shows a 3 ks and 4 ks profile at the beginning of the epoch which will be investigated later (Section 6.6.4).

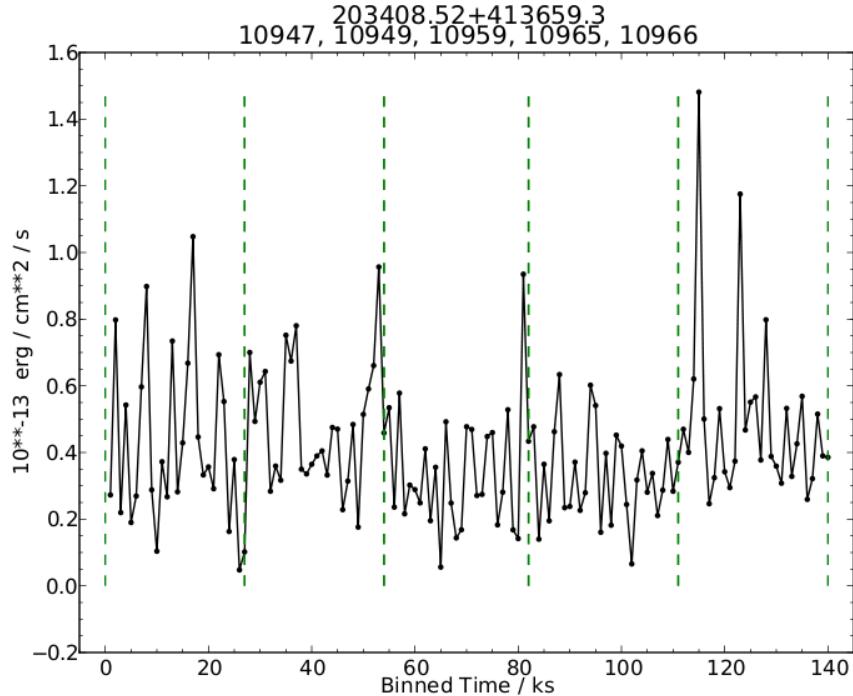


Figure 6.24: Binned light curve of Schulte 11 with the observed epochs: 4511, 10955, 10956, 10960, 10961. The y-axis is the flux in 10^{-13} erg cm^{-2} s^{-1} and the x-axis is the binned time in ks. The green dashed lines mark the different pointings.

As almost half of the phase is covered by the Chandra observations, LS periodograms are deployed to determine if the orbital period can be found. The only noticeable peak does not correspond to the quoted period and is well below the FAP = 0.005 line.

The non-detection of the period of 72.4 days by the LS periodogram may be because of the long period and the poor phase coverage for almost half the period. The MC

simulations also show no detection, with only 0.12% of MC runs having a FAP < 0.30 around the given period.

Therefore, the conclusion for Schulte 11 is that variability is present in the Chandra dataset. However, this study can not confirm that the variability is phase-locked on the orbital period, but suggests that Schulte 11 is a good candidate for further studies.

203409.51+413413.9 - Schulte 75

Schulte 75 is classified in the Chandra catalogue as a single O9V star, but qualifies for investigation due to the significant KS p-value of 4.946E-17. However, with epochs 10949, 10959 and 10966 suffering from high background counts, only epoch 10965 has a source flux greater than the background level. Despite the high background level in three out of four epochs, epoch 10965 is noticeably more active than these epochs and is what the KS statistic indicates. Figure 6.25 shows the binned light curve only for epoch 10965. There appears to be a 2 ks peak around 15 ks.

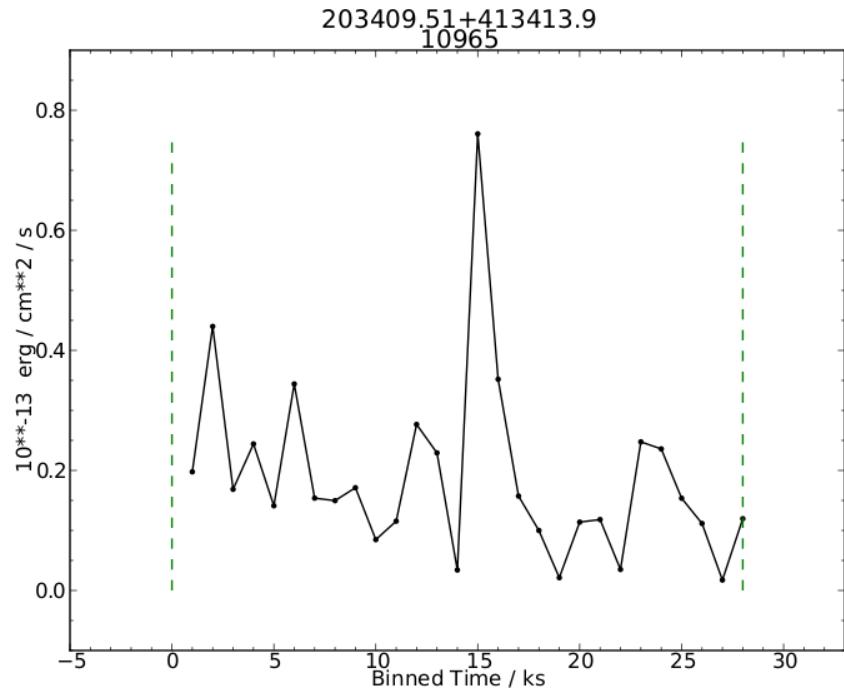


Figure 6.25: Binned light curve of Schulte 75 with the observed epoch: 10965. The y-axis is the flux in 10^{-13} erg cm^{-2} s^{-1} and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. A small 2 ks long peak is seen around 15 ks.

As only one epoch is available, it is difficult to suggest any hint of inter-pointing variability within this Chandra dataset for Schulte 75.

203421.95+411701.5 - Schulte 73

Schulte 73 is classified in the Chandra catalogue as an O8III + O8III binary, but with no associated period and a KS p-value of 9.041E-5. Epochs 10955 and 10960 suffer from high background counts, while the epochs 10954 and 10963 show small detectability. Figure 6.26 shows the binned light curve with a large amount of non-detections in the early 10954 epoch. The later epoch (10963) is seemingly slightly stronger with a peak around 36 ks. The narrow profile of this peak, causes difficulty in determining whether this is a real feature or an erroneous point.

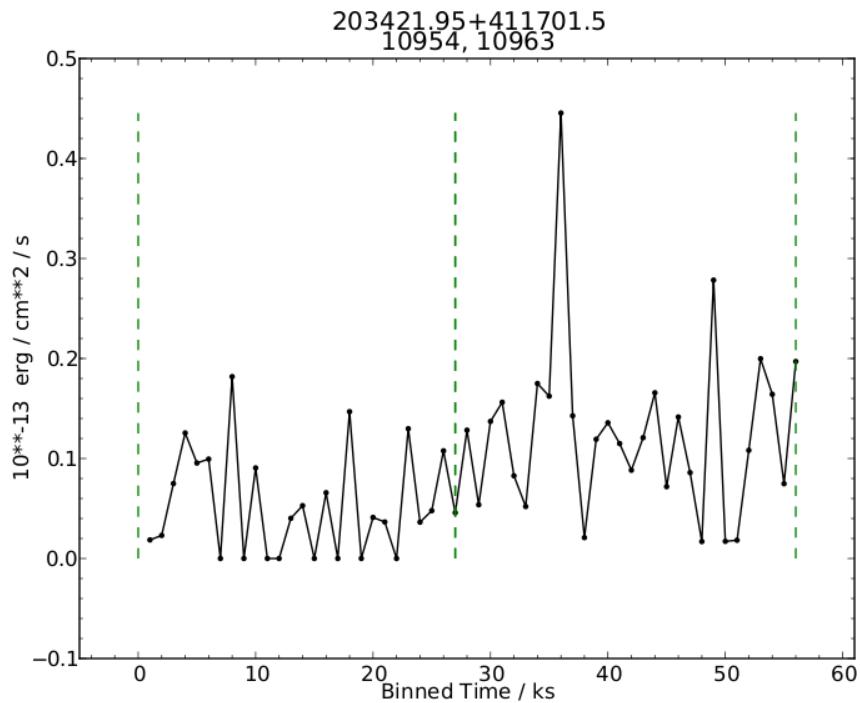


Figure 6.26: Binned light curve of Schulte 73 with the observed epochs: 10954, 10963. The y-axis is the flux in 10^{-13} erg $\text{cm}^{-2} \text{s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The fluxes appear to be stronger in the later epoch, with the early epoch containing non-detections.

The X-ray flux levels for Schulte 73 are very low for both detectable epochs, with the later seemingly stronger. There does not appear to be much, if any, evidence of variability from the Chandra dataset.

203429.60+413145.3 - MT91 771

MT771 is an O7V+O9V binary with a period of 2.8635 days (Kiminki et al. 2012b). The binned light curve with all epochs shown adjacent to one another is presented in Figure

6.27. Despite some scattering within each epoch, the average flux values are fairly constant across all epochs. Pile-up is not expected to be a problem for MT91 771 with any pile-up fraction being $< 3\%$.

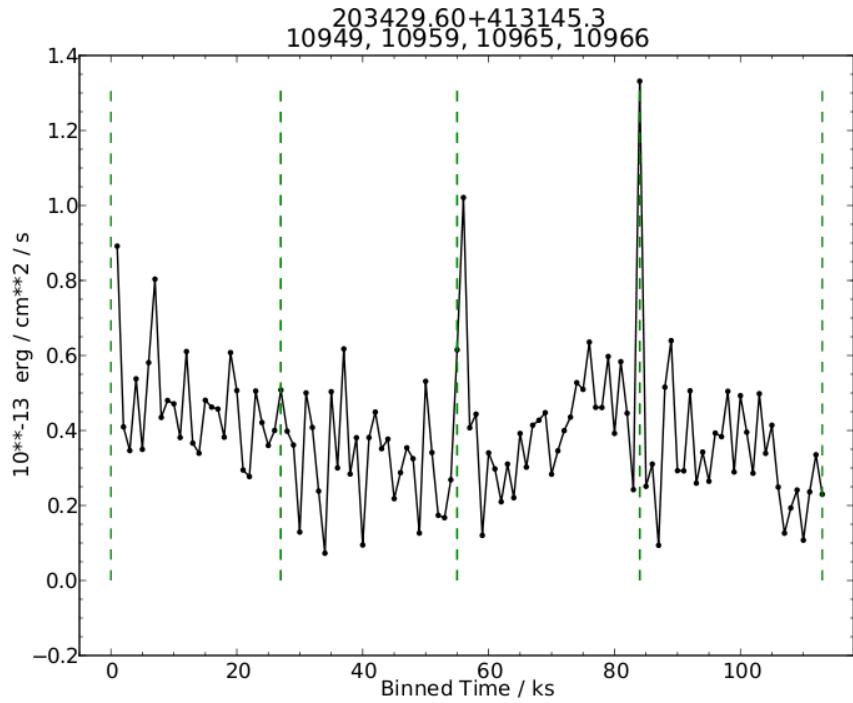


Figure 6.27: Binned light curve of MT91 771 with the observed epochs: 10949, 10959, 10965, 10966. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The fluxes across the epochs appear fairly constant.

The LS periodogram result showed a peak slightly offset from the cited period but below the $\text{FAP} = 0.005$ confidence line. MC simulations around this peak only found 9.45% of MC runs have a $\text{FAP} < 0.30$, which indicates that the period was not determined by this study.

Thus there is no evidence of X-ray variability from this Chandra dataset, on the associated period or otherwise for MT91 771.

203547.08+412244.7 - WR 146

WR 146 is a Wolf-Rayet star with an O8III companion (classification: WC6 + O8III), with a singular period estimation from a multi-frequency study by Dougherty et al. (1996), stating a period of ~ 300 years. The KS test produced a p-value of 2.733E-2, which signifies a ‘possibly’ variable object, but with a period on the order of 100’s of years, any variability will not arise from the orbital motion of the binary. Figure 6.28 displays the binned light

curve of WR 146 with the three epochs of 7426, 10967, 10968. Other than a few spikes, the general flux level is constant across all three epochs which represents a constant baseline flux over a number of years. The Chandra datasets do not demonstrate any periodic behaviour in the X-rays from WR 146.

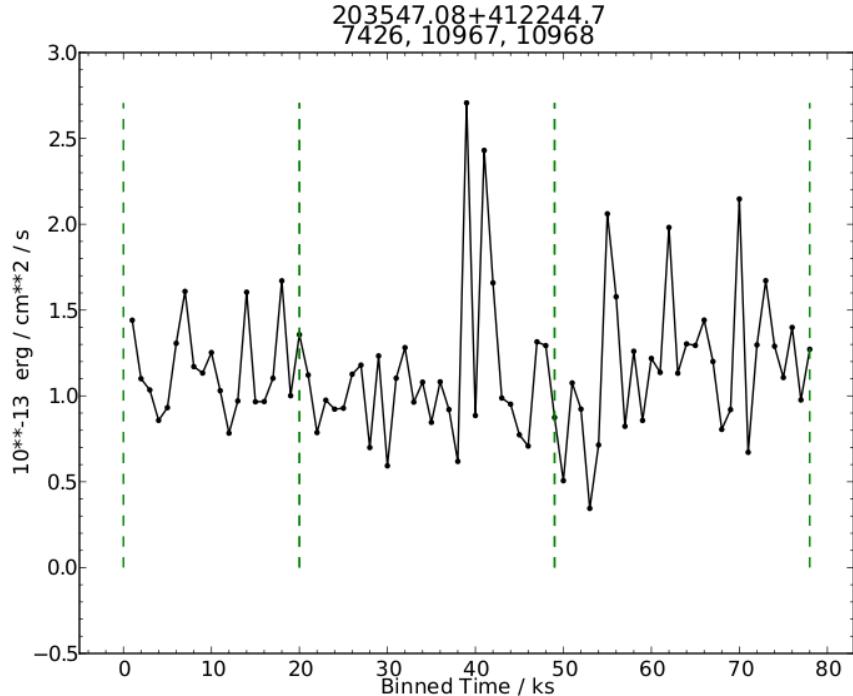


Figure 6.28: Binned light curve of WR 146 with the observed epochs: 7426, 10967, 10968. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The general flux trend across the epochs appears fairly constant, with a few spikes in the latter two epochs (10967, 10968).

6.6.4 Intra-pointing X-ray Variability

To qualify for further investigation into intra-pointing variability, a star has to demonstrate either a significant KS p-value < 0.005 for a certain epoch, or have a significant feature profile (see below for description) after visual inspection of the light curves in Section 6.6.3, lasting at least 3 ks. This amount of prolonged, heightened flux suggests the increase is physically connected with each data point and not just a spurious point. Such spurious points are defined as extreme flux outliers which last for 1 or 2 ks i.e. 1 or 2 seemingly extreme data points. Events containing heightened flux profiles lasting for less than 5 ks are deemed ‘mini-flares’, and events lasting longer than 5 ks are deemed as flares.

To further constrain the criteria for a flare or mini-flare, the peak flux must be greater than 3 standard deviations of the median (used instead of the mean to nullify any outliers)

of the tested sample. This is equivalent to the p-value of 0.005 which has also been referred to as ‘definitely variable’.

203137.50+411321.1 - Schulte 3

The binned light curve (Figure 6.5) shows two possible ‘mini-flares’ in epochs 10940 and 10953. The flux of the feature in epoch 10940 is not significantly greater than the average flux level, but the profile is broad. Considering the other flux points in the same epoch, they are of similar strength and therefore this possible mini-flare is probably just stochastic variations.

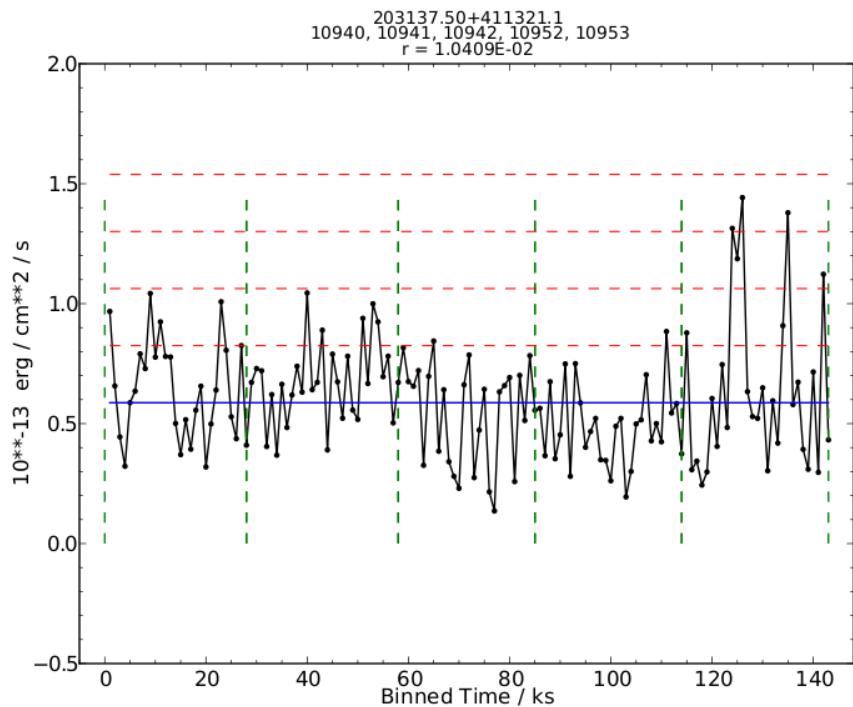


Figure 6.29: Binned light curve of Schulte 3 with the observed epochs: 10940, 10941, 10942, 10952, 10953. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The blue line represents the median flux of all epochs, and thred dashed lines represent increasing standard deviations from the median. A mini-flare is seen at around 125 ks.

However, the mini-flare in epoch 10953 lasting for 3 ks with the peak flux reaching 3.6 sigma may resemble a real feature. Figure 6.29 shows a light curve with the median and standard deviations plotted in blue and red lines respectively. Epoch 10962 is omitted because of the spurious points which reside in that epoch. Another peak > 3 sigma is also present in the same epoch, but only lasts for 1 ks at a significant flux level and is therefore not considered as a mini-flare event.

203231.54+411408.1 - MT91 267

The most prominent possibility for intra-pointing variability for MT91 267 comes from the 4-5 ks mini-flare in epoch 4511. Figure 6.30 shows the epoch 4511 binned light curve with the mean and standard deviations plotted in blue and red lines respectively. The mini-flare occurs around 56 ks in Figure 6.30, with the peak flux reaching the 3.3 sigma level.

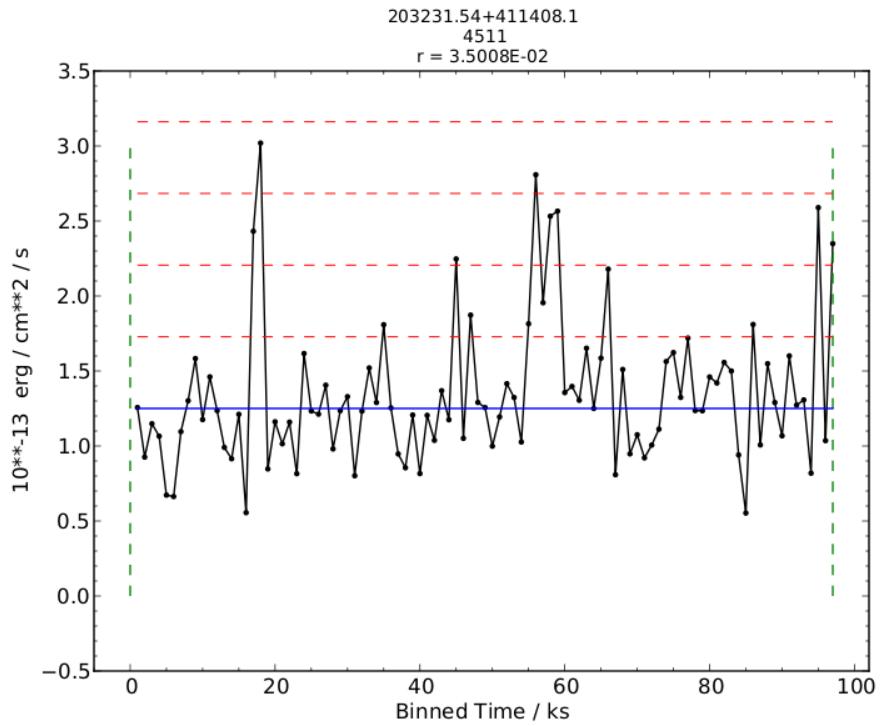


Figure 6.30: Binned light curve of MT91 267 with the observed epoch: 4511. The y-axis is the flux in 10^{-13} erg cm^{-2} s^{-1} and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The blue line represents the median flux of all epochs, and thred dashed lines represent increasing standard deviations from the median. A mini-flare is seen at around 56 ks.

Another mini-flare candidate lies around 19 ks, and reaches the 3.7 sigma level, the highest flux in the epoch. However, this mini-flare candidate only lasts for 2 ks, and therefore does not qualify as a mini-flare.

203308.77+411318.7 - Schulte 22

The binned light curve for Schulte 22 (Figure 6.17) suggests the most promising mini-flare candidates lie in epoch 4511 and 10964. However, the peak strengths of these mini-flare candidates do not reach the required 3 sigma significance level but only 2.1 and 2.7 sigmas

for events in epochs 4511 and 10964 respectively. Therefore there is no evidence for intra-pointing variability from Schulte 22.

203408.52+413659.3 - Schulte 11

The binned light curve for Schulte 11 (Figure 6.24) demonstrates large amounts of scatter within each epoch. The most promising mini-flare candidate lies in epoch 10949. However, the peak strength of this mini-flare candidate does not reach the required 3 sigma significance level but only 1.9 sigma. Therefore there is no evidence for intra-pointing variability from Schulte 11.

203206.26+404829.6 - WR 145

WR 145 is a binary with the Chandra survey spectral classification of WN7o/CE+O7V((f)) and a known period of 22.54977 days (Muntean et al. 2009). The intra-pointing KS test revealed a p-value of 5.754E-7 significance, warranting an investigation for variable activity from WR 145. Figure 6.31 shows the binned light curve for the sole epoch 10969. A general increase is visible from the beginning and end of the epoch.

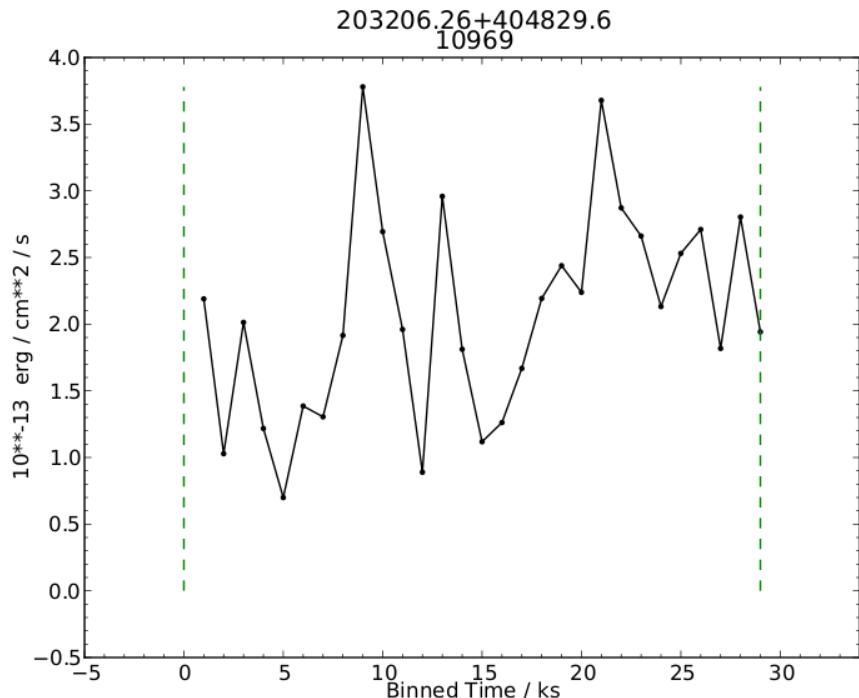


Figure 6.31: Binned light curve of WR 145 with the observed epoch: 10969. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. A general increase in flux is noticeable across the epoch.

Regressional analysis finds this positive correlation to have an r^2 value of 0.17. As such, when determining the significance of any peaks, the standard deviation should be taken from the regression fit rather than the median of the epoch. Figure 6.32 shows that with this fit, the largest peak reaches 2.5 sigma.

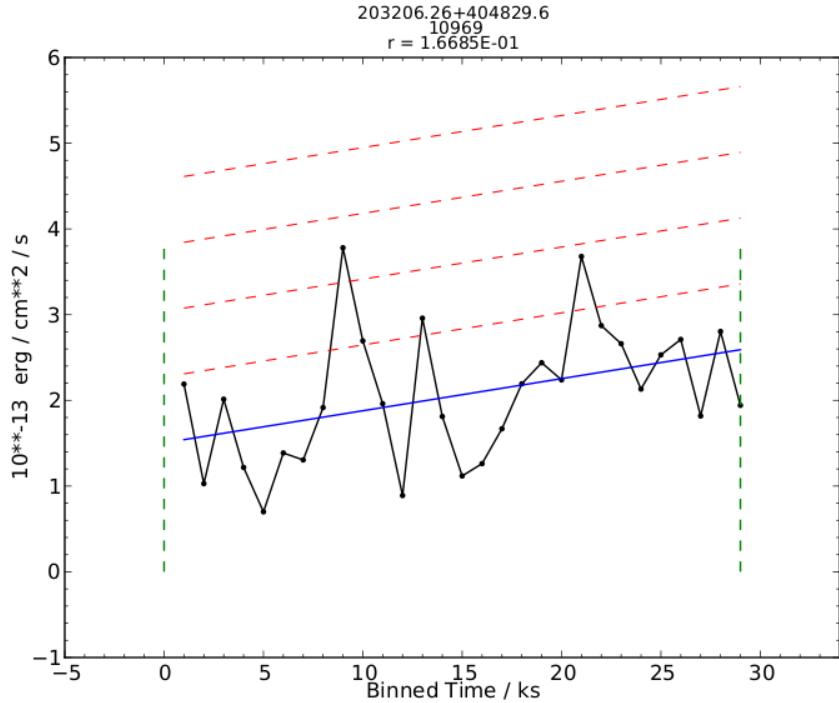


Figure 6.32: Binned light curve of WR 145 with the observed epoch: 10969. The y-axis is the flux in 10^{-13} erg $\text{cm}^{-2} \text{s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The blue line is the regressional fit applied to the data, and the red dashed lines are multiples of the standard deviations from the fit.

This does not fulfil the criteria of a mini-flare and therefore, the conclusion is that there is no evidence for any flaring activity, but the regressional analysis and KS test p-value does suggest that there is variability in the X-ray flux from WR 145. The nature of this variability is gradual, and because of the lack of observations, it is impossible to correlate this to the orbital period of 22.5 days.

203326.74+411059.4 - MT91 534

MT91 534 is a main sequence O-type star with the spectral classification of O7.5V. The KS test produced a p-value of 1.803E-65, which is a phenomenally significant value. Figure 6.33 shows the binned light curve of MT91 534, with the flare clearly visible in epoch 10960. The flare occurs in the middle of the observation and continues to the end of the epoch, making it unclear as to when the flaring activity stops. A minimum flare time of

16 ks can be deduced from the figure. The following epoch 10961 is observed only a few hours after epoch 10960 and has returned to the pre-flare flux level.

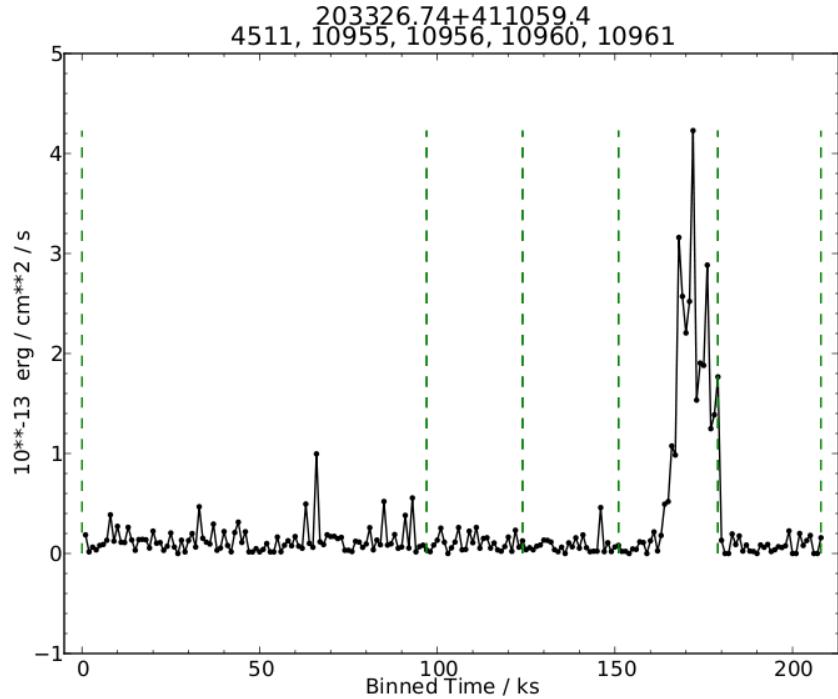


Figure 6.33: Binned light curve of MT91 534 with the observed epoch: 10969. The y-axis is the flux in 10^{-13} erg $\text{cm}^{-2} \text{s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. A flare can be seen at around 170 ks lasting at least 16 ks.

The strength of the flare is roughly 20 times greater in flux than the continuum level with the peak of the flare reaching 7.4 sigma of the median flux for all epochs. The sigma levels are plotted in Figure 6.34.

The nature of this flare is similar to a number of flaring events seen in a previous Chandra study (Albacete Colombo et al. 2007a; Albacete Colombo et al. 2007b) from low mass members of Cyg OB2, where the authors suggest that the X-ray flares are a result of magnetic reconnection events (Favata and Micela 2003). Figure 6.35 is a plot taken from Albacete Colombo et al. (2007a) of 20 sources, 14 of which show flaring activity.

Only two OB stars in the 97 ks Chandra study (both B-type stars, one a binary system), demonstrated gradual variability, which Albacete Colombo et al. (2007a) attributed to magnetically confined wind shocks, possibly like the oblique magnetic rotator θ^1 Orionis C (Gagné et al. 2005).

None of the OB stars in the Albacete Colombo et al. (2007a) study revealed flare-like activity, and none are expected to show the features of the flare seen in Figure 6.33.

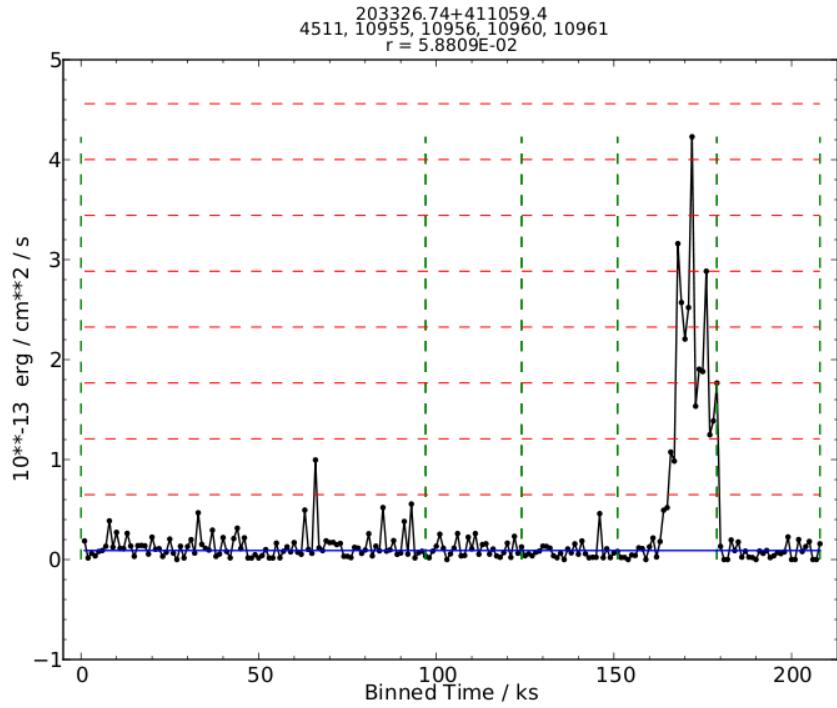


Figure 6.34: Binned light curve of MT91 534 with the observed epoch: 10969. The y-axis is the flux in $10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ and the x-axis is the binned time in ks. The green dashed lines mark the different pointings. The blue line is the median flux of all the epochs and the red dashed lines are the subsequent sigma levels. A flare can be seen reaching 7.4 sigmas at around 170 ks lasting at least 16 ks.

Moreover, the flaring activity of MT91 534 portrays the same morphology as the flares from many low mass stars in the Albacete Colombo et al. (2007a) study (Figure 6.35). Therefore the conclusion for MT91 534 is that the X-ray flare of this nature could not arise from the O7.5V star, and that a hidden low mass companion is present exhibiting a coronal magnetic reconnection event.

6.6.5 Summary of X-ray Variability

A summary of the X-ray variability of massive stars found in this study is presented here for the inter-pointing and intra-pointing variability. Table 6.3 shows all the sources investigated with notes on whether the variability was in the form of a positive KS statistic result, greater than 3σ flaring activity (intra-pointing) or gradual flux increase, or periodic variability determined by the LS periodograms and MC simulations (inter-pointing).

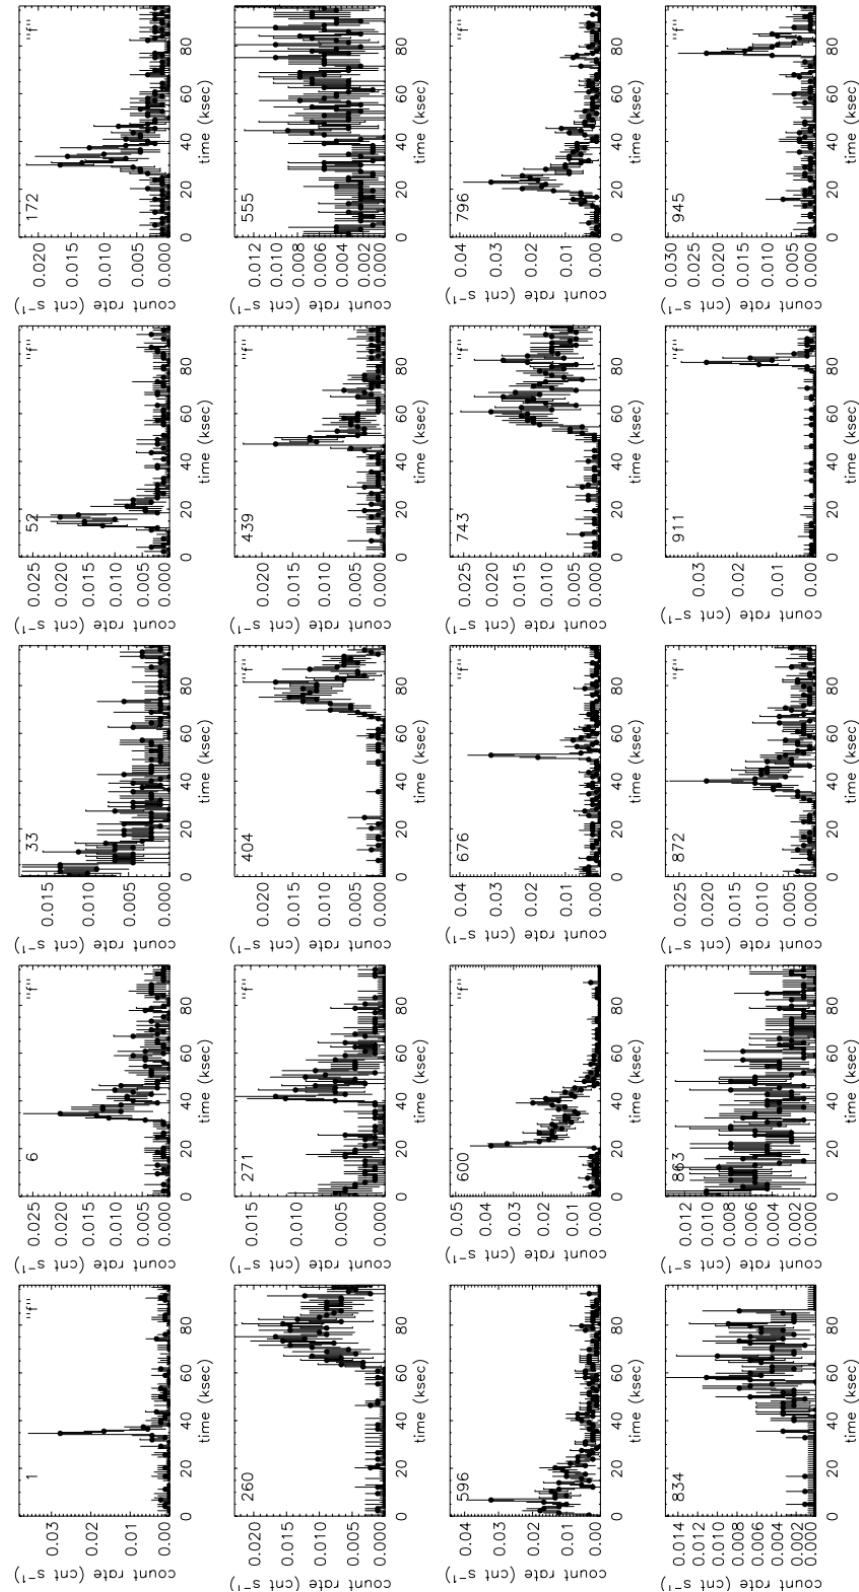


Figure 6.35: 600s binned light curves for 20 sources from the Albacete Colombo et al. (2007a) 97 ks Chandra survey. Source classified as flares are denoted with an “f” in the top right corner of each plot. This figure is taken from Albacete Colombo et al. (2007a) (Figure 6).

Table 6.3: Summary of the X-ray Variability in this Study

Chandra Name	Common Catalogue Designation	KS Statistic $P_{KS} < 0.005$	Intra-pointing Variability	Inter-pointing Variability
203137.50+411321.1	Schulte 3	No	Yes	No
203213.84+412711.4	Schulte 4	Yes	No	No
203222.42+411819.0	Schulte 5	Yes	No	Yes
203231.54+411408.1	MT91 267	Yes	Yes	Yes
203234.87+405617.0	A38	Yes	No	No
203308.77+411318.7	Schulte 22	No	No	No
203310.73+411508.2	Schulte 9 (MT91 431)	Yes	No	No
203315.07+411850.5	Schulte 8A (MT91 465)	N/A	No	No
203323.48+410912.6	MT91 516	Yes	No	No
203359.56+411735.5	Schulte 27	No	No	No
203408.52+413659.3	Schulte 11 (MT91 734)	Yes	No	No
203409.51+413413.9	Schulte 75	Yes	No	No
203421.95+411701.5	Schulte 73	Yes	No	No
203429.60+413145.3	MT91 771	No	No	No
203547.08+412244.7	WR 146	Yes	No	No
203206.26+404829.6	WR 145	Yes ^a	Yes	No
203326.74+411059.4	MT91 534	Yes ^a	Yes	No

All KS statistics are ‘MERGED’ KS statistics (inter-pointing) unless otherwise stated. ^a KS statistics from the single epoch (intra-pointing).

6.7 Conclusions and Discussions

This variability study of O-type and WR stars from the Chandra Cyg OB2 Legacy dataset has found 5 out of the final 15 candidates to be variable or possibly variable on the inter-pointing timescale. Two stars; Schulte 5 and MT91 267, reveal periodic variability on known periods corresponding to the orbital motions of the binaries which were statistically verified. The Schulte 22 analysis offered suggestive evidence for periodic behaviour based on 23% of the MC simulations fulfilling the $FAP < 0.3$ criteria imposed by Hoffman et al. (2012). MT91 516 showed small variations in the Chandra dataset, but this remains speculative.

For the intra-pointing variability, 4 out of the 15 candidates demonstrated a variety of variable behaviour, including mini-flares and flares (see Section 6.6.4 for a description), and gradual flux increases. Schulte 3 and MT91 267 revealed possible ‘mini-flares’ lasting for 3 and 4–5 ks respectively at the 3σ significance level. MT91 534 displayed a strong flare lasting for > 16 ks at the 7σ significance level. The single epoch of WR 145 showed a gentle increase in flux values over the observation, the origin of which could not be identified.

All of the stars investigated, demonstrated flux scattering within the epochs of observations. For some of the sources, pile-up may affect the variability profile but this effect does not make a constant source appear variable. However, many of the stars have low or

negligible pile-up fractions. This leads to the discussion on where this scatter originates from. The LS periodograms for many stars did reveal many strong peaks at a number of frequencies which do not correspond to known periods. These are most likely due to aliases, and therefore searching for high frequency (short time periodicities) would be difficult because of data sampling effects in the periodogram. A robust study of these scattering effects and whether they are a result of stochastic or short timescale periodic variability would be difficult, not only for these reasons, but also because the expected variable population from OB star wind instabilities on timescales such as these (hours) is very small, less than 1% (Nazé 2009; Massa et al. 2014; in press).

Similar short timescale X-ray studies on the O star ζ Pup were conducted by Nazé et al. (2013), who did not find any flare-like activity, or short timescale (< 1 day) variations whose optical counterparts are attributed to stellar pulsations (Reid and Howarth 1996). Nazé et al. (2013) do not rule out the existence of such short timescale X-ray variability, however, state that their observations are the most extensive and sensitive data on a specific target to that date. They further emphasise that more sensitive observations may detect variable features as described above.

From the Chandra catalogue and this variability study, there are 49 O stars and 3 WR stars in total with 106 known OB stars in Cyg OB2 (Albacete Colombo et al. (2007a)). Of the 49 O-type stars, 12 are known, or thought to be in binaries, making up 24% of the X-ray population of Cyg OB2. Eleven of these (22%) are massive binaries (made up of O and/ or B stars) and one O star - low mass star binary to account for the flaring activity seen in MT91 534. Two of the 11 massive binaries have been recently confirmed; MT91 267 (Kobulnicky et al. 2012), MT91 516 (Rauw et al. 2014), implying that this binary fraction of Cyg OB2 is far from complete with new binary confirmations being made every year. Moreover, the O-type star population of Cyg OB2 is not complete, with more discoveries and confirmations expected with the completion of the COBRaS Legacy survey.

The binary fractions of massive stars in young open clusters and OB associations are uncertain. Sana et al. (2013) (and references therein); summarises a large number of studies on a number of clusters, with the binary fraction ranging between 30% and 60%. These differences are accountable to random fluctuations expected from the sample sizes. Knowing the binary fraction is very important in order to understand the formation and evolution of massive stars. Sana et al. (2013) estimate that 53% of all stars born as O-type

belong to a binary system with a period < 1500 days. They go on to say; 18% of the O stars will merge with a companion, 27% will be stripped from their envelope and 8% are expected to be spun up. This clearly indicates the significance of binarity on evolution which makes up for over half of the population of O stars (assuming longer period binaries > 1500 days, do not have the same level of interaction due to the greater orbital distance).

The age of the cluster (and therefore of the member stars) must also be considered, because 5% of the O star population is running away either a result of binary interaction or supernova kicks, leading to the conclusion that a fraction of the O star population has undergone dynamical or evolutionary interaction (Sana et al. 2013). This suggests that clusters or associations with different ages, may have a slightly different binary fraction due to the evolutionary history of the population.

In light of the evidence that MT91 534, an O-type main sequence star, may have a low mass companion discovered from a ‘chance’ observation of an X-ray flare, more seemingly single O-type stars could host hidden lower mass companions. This could have significant repercussions for the perceived evolution of massive stars, by hiding the true nature and reason for an O star’s particular evolutionary path and biasing the frequency of binaries.

More extensive surveys may find more flaring events from massive stars which could be attributed to an unseen lower mass star. However, because of the serendipitous nature of the detecting event (the flare) it may be unlikely to realistically discover all of the lower mass companions via this method, but it may also be the only way of detecting late-type companions around O-type stars.

A number of methods are needed to discover all the massive binaries in clusters and associations, due to the orbital period/ distance, inclination of the systems, and the mass ratios all result in different physical scenarios. Spectroscopic techniques ($H\alpha$ studies) are better suited to determining short period binaries (< 1 year), but only if the system is not face-on (i.e. low inclination, $\sim 0^\circ$). Whereas intermediate period binaries (1 - 100 years) are more likely to be detected by non-thermal radio emission from wind-wind collision regions or by X-rays from the collision shocks. Longer periods (> 100 years) will only be visible to astrometry studies.

Therefore, COBRaS non-thermal radio and Chandra X-ray observations cover the large range of intermediate period binaries and are essential studies to robustly detect all of the binaries possible by the physical dependencies imposed by the binary systems (such as orientation during the observations). This further states the importance of multi-

wavelength surveys on Cyg OB2 and other clusters and associations to reveal the entire picture and nature of the population within.

While the X-ray variability fixed on the orbital period is a physical consequence of large shocks from wind-wind interactions (Stevens et al. 1992; Rauw et al. 2014), other origins for X-ray variability arise from Co-rotating Interaction Regions (CIRs, Mullan 1986). As a star rotates emitting its wind in a non-circularly symmetric with respect to the rotation axis, fast wind emitted along a certain direction will catch up to a slower wind emitted in the same direction at an earlier interval. If the disparity in velocities is great enough (on order of the local sound speed) a shock forms between the two wind regions (Mullan 1986). The interaction region rotates with the rotation period of the star, leading to the name CIR.

The timescale of the X-ray variability from CIRs should be comparable to the stellar rotation period and should also correlate with Ultra-Violet (UV) wind line variability (Massa et al. 2014; in press). Studies on the O star ζ Oph has shown rotation modulation of X-rays in conjunction with UV discrete absorption components (Osokinova et al. 2001). Other possible evidence for CIRs comes from ζ Pup, although additional data is required to sample the full rotation period to refine the origin of the variability (Nazé et al. 2013).

From this study, Schulte 22 showed tentative evidence that variability was present between the calculated timescale limits of the star's rotation period (3.3 and 9.2 days, Rauw et al. 2014). This could indicate that if the variability is significant, a possible origin of this variability could be from the CIR around Schulte 22, although further X-ray and additional UV observations are necessary to confirm these assertions.

Another method of detecting periodicities in X-ray datasets is to apply the Kuiper test (Kuiper 1960). Paltani (2004) searched the ROSAT archive with the Kuiper's test over a range of frequencies which the data is then folded on. The authors comment on problems with extrinsic contamination and note that firm confirmation depends on detection in independent datasets.

The Kuiper test was attempted in this investigation, with cumulative plots and relevant statistics determined. Despite visual cumulative plots appearing to show periodic behaviour, this was not statistically confirmed. The amplitude of the cumulative data to a constant model for the number of data points present, was not profound enough to trigger a Kuiper p-value of the necessary significance. Arbitrarily increasing the number of events used in calculating the significance but using the same dataset, yielded a drastic change

in the p-values. It appears that when there are a few number of events in the cumulative histogram, the amplitudes need to be quite large to reach the required significance.

The LS periodogram utilised in this study, is equivalent to a non-linear least squares model fit to the data (Scargle 1982). Such fits were made to the datasets for stars with known periods using the Levenberg-Marquardt method of minimisation (Levenberg 1944; Marquardt 1963), which uses the χ^2 statistic as the minimisation function. This method locates the nearest local minimum which may not be the global minimum and is not therefore robust. This is evident when the fit function is trigonometric where identical values occur for every $2n\pi$. Therefore, in some cases an initial estimate of the input variables for the function used in the fit may be needed to help the algorithm.

An improvement to this fitting process could be to use the Cash statistic (Cash 1979) as the minimisation function, as this works well when the number of counts per bin is low. This may result in a better fit to the dataset, although the uneven sampling may still affect the fit. However, the main issue with using the fit is finding a goodness-of-fit statistic which is reliable. The Cash statistic does not offer this, strengthening the argument for other methods which give statistical verification to any fit, such as those deployed in this thesis (LS periodograms and MC simulations).

6.8 Future Work

Future X-ray variability work can be conducted on the larger available B star sample from the Chandra Cyg OB2 Legacy survey. This can consist of similar studies presented here, with the LS periodograms and MC methods offering strong arguments for statistical verification (discussed above) for known or candidate variable periods. Alternative avenues of investigation into time series analysis with the Cash statistic are also possible, although the limitations of these methods have also been discussed above.

The larger sample may reveal more candidates and evidence for short timescale (\sim hours) variability. Any well observed stars including fundamental stellar parameters such as the stellar rotation period, may offer more evidence for CIRs and modulated variability on the rotation period.

Chapter 7

Summary, Conclusions and Future Work

Knowledge speaks, but wisdom listens.

Jimi Hendrix

This chapter contains a summary of the technical and science work completed in this thesis and thus providing the basis for the ultimate goals of the COBRaS project. Future work with COBRaS Legacy datasets is reviewed. The final section discusses the wider context of software development, engineering and applications to future major projects such as the SKA.

7.1 Summary of Thesis

Chapter 1 introduces the theory of radio interferometry including the different aspects of the interferometric heterodyne system and aperture synthesis algorithms. The chapter also contains a description of e-MERLIN and the technical capabilities as a fully operational telescope. Lastly, an introduction to COBRaS includes the technical and science goals of the Legacy project.

Chapter 2 presents the SERPent algorithm to reduce and flag radio interferometric datasets from e-MERLIN. SERPent contains three reduction passages: the Lovell station-

ary scan removal, zero-level amplitude removal and RFI-mitigation. SERPent interacts with AIPS to read in visibilities for sources and appends a flag (FG) table to the data in AIPS. SERPent has been tested on a large number of datasets and is currently used by a number of international institutions and is included in the e-MERLIN calibration pipeline.

Chapter 3 describes the COBRaS calibration pipeline, including delay and phase calibration, flux and amplitude calibration, spectral fitting and bandpass calibration, and self-calibration of the phase calibrator J2007+4029. The automated algorithm includes passages which conduct otherwise manual tasks such as setting the delay window for delay calibration on a suitable section of data, bootstrapping the fluxes from 3C286 onto OQ208 (if that option in the pipeline is chosen), automated spectral fitting unbaised by obvious outliers, and selecting a CLEAN box around the phase calibrator for self-calibration. The pipeline has been tested on a L-band Legacy dataset and a commissioning C-band dataset. Parts of the COBRaS pipeline have been adopted in the general e-MERLIN calibration pipeline (Megan Argo, University of Manchester, UK).

Chapter 4 includes a number of programs to detect, extract and classify sources. Firstly, a catalogue amalgamation script cross-correlates a number of previous studies on the Cyg OB2 region into one definitive catalogue. Subsequent specific catalogues are compiled from this one catalogue to create an OB star catalogue and candidate catalogue. The next part of the chapter presents a source detection and flux extraction program which determines the fluxes of point sources and resolved sources in a pixel-by-pixel (PP) manner. This method does not assume any particular source structure (such as Gaussian) and performs better or equivalent flux extraction compared to existing methods in AIPS (JMFIT) on simulated data. The last part of this chapter contains a cross referencing algorithm to match sources in a map with those in a catalogue. Using Bayes' Theorem, a source significance boosting module is also presented which can statistically increase the detection significance of faint sources.

Chapter 5 introduces mass loss and the winds of massive stars, the radio emission from massive stars and the effects of clumping on the radio fluxes. The chapter presents predicted smooth mass loss rates and radio fluxes based on the predicted fundamental stellar properties of massive stars from their spectral type. Building on these smooth wind predictions, a study on how a varying clumping factor affects the radio fluxes is included, comparing these clumped radio fluxes with previously published radio observations of massive stars in Cyg OB2. Finally, the first radio maps from COBRaS are presented with

source and flux lists and some initial analysis.

Chapter 6 contains an introduction to X-rays from massive stars, the Chandra Cyg OB2 Legacy survey and the Chandra space observatory. A background description of the statistical methods used in characterising variable sources in X-ray astronomy is also given. A study on X-ray variability in O-type stars and WR stars is presented. Long term variability of inter-pointing observations is investigated with the Lomb-Scargle periodograms and Monte Carlo simulations, and short term intra-pointing variability in the form of flaring events and continuum increases.

7.2 Thesis in the Context of COBRaS

In Section 1.3.4, a number of technical requirements were stated as part of the COBRaS project. These include: data processing, calibration, imaging, analysis and data archiving. This thesis addresses the data processing with SERPent (Chapter 2), calibration with the COBRaS calibration pipeline (Chapter 3), source analysis with the source and flux extraction program (Chapter 4) and the foundations of the data archiving with the catalogue scripts (Chapter 4). Therefore, this thesis provides a substantial amount of the technical ground work for COBRaS, with imaging and mosaicing the only major technical aspects still to be addressed.

In Section 1.3.3, a number of astrophysical themes and goals are described. These include: mass loss rates from massive stars, massive binary fractions and the incidence of non-thermal radiation, and on-going and triggered star formation. This thesis contains the predicted smooth radio fluxes and clumped radio fluxes for O stars and early B supergiants and a first analysis of the radio maps is given with potential sources (Chapter 5). The topic of mass loss rates and how clumping affects the radio fluxes and mass loss rates determined from radio observations is a key area of COBRaS, and is led by UCL.

Another application of COBRaS is the synergy with other Cyg OB2 datasets, observed at other wavelengths. Surveys such as the Chandra Cyg OB2 Legacy survey, can provide additional insight into the science goals of COBRaS by studying different phenomena associated with the population of Cyg OB2. The variable X-ray emission from massive stars traces behaviour such as binarity, CIRs and flares. This thesis contains studies on the variability of O stars and WR stars (Chapter 6) whilst demonstrating the importance of multi-wavelength studies of the same region.

The scientific output from COBRaS is only just beginning, with the first radio maps now being created (April 2014). In the coming year, the rest of the L-band Legacy dataset will be observed allowing for investigations into low frequency, non-thermal phenomena. This includes the search for known massive binaries with cross correlations with the radial velocity studies (Kiminki et al. 2007), and also unknown or candidate binaries with comparisons with the catalogues created in Chapter 4. Other non-thermal emitters will consist of weak T-Tauri stars, tracing star formation in the Cyg OB2 region, and the active galactic nuclei (AGN) of any background galaxies. Whilst the AGNs are out of the context of COBRaS which focuses on Galactic stellar objects, these are still interesting objects for the extra-galactic astronomy community.

The second phase of COBRaS (2015) will consist of the significantly larger C-band dataset of 42 pointings with a total integration time of 252 hours. Despite the larger amount of data, the C-band should have considerably less RFI than the L-band observations, resulting in a smoother passage through the calibration pipeline. The data can therefore be reduced and calibrated by SERPent and the calibration pipeline with no manual intervention or checks with a higher confidence of good results. This is vital because of the large data volume (tens of TBs).

The C-band dataset will exploit the full 2 GHz bandwidth producing sensitive mosaic maps of thermal sources with spectral indices (in conjunction with the L-band maps), allowing the detection of a large number of massive stars. This will produce a statistical sample of massive star radio fluxes, from which constraints on the clumping factor and mass loss rates can be determined. COBRaS will conclusively demonstrate whether clumping exists in the outer winds of massive stars and will play a major role in future studies of clumping at other wavelength ranges. The technical and scientific ground work conducted in this thesis will be pivotal in helping COBRaS achieve this and the other science goals.

To add to the investigation of clumping in the winds of massive stars, a recent ALMA cycle 2 proposal was accepted (P.I. Dr Danielle Fenech) to observe Westerlund 1, the most massive stellar cluster in the Galaxy, located in the Southern hemisphere. This will obtain a number of millimetre fluxes for massive stars and will refine the amount of clumping in the middle part of the wind, where it is predicted from hydrodynamic simulations that the peak amount of clumping occurs (Runacres and Owocki 2002). This ALMA dataset along with the e-MERLIN COBRaS dataset will provide the most extensive investigation of clumping in the middle and outer winds of massive stars.

7.3 Thesis in the Wider Context of Software Engineering

This thesis includes three major pieces of software: SERPent, the COBRaS calibration pipeline, and the PP source and flux extraction program. All three were created from scratch, each with a specific technical goal to accomplish. There are many conclusions gleaned from developing these programs, which are now discussed.

All of these programs are automated, requiring minimal or no manual interaction, but each has its own breakdown limit. SERPent incorporates robust statistical variance estimators, but will breakdown if the amount of RFI for any given sample (total integration time on-source versus channels in an IF) approaches $\sim 50\%$. SERPent can remove a number of errors in the dataset, but there are no guarantees it will successfully remove any new future problems. The calibration pipeline is particularly vulnerable to any bad data remaining in the datasets, because of the high number of variables in the pipeline and therefore potential errors which can propagate through the system. Including fail-safes for every conceivable eventuality is not possible. The PP extraction performance is in many cases limited by the noise distribution, not just for low SNR sources but also from errors which may arise during deconvolution. Therefore automated programs require a certain level of data quality in order to be effective or even successful in their operations.

The testing and development of software can be broken down into four general stages:

1. Creation of the algorithm which performs a specific task.
2. Making the algorithm robust by testing it on multiple datasets, editing the software if required so that it completes the tasks on all of the datasets to a sufficient level every time. This is known as test-driven development.
3. Optimisation of the algorithm performance both in output results and perhaps execution time.
4. Maintaining the algorithm, including updates of any new modules or ideas in the field.

SERPent is currently the most developed piece of software presented in this thesis, and is at the stage of requiring optimisation in its flagging and computational performances (see discussions in Section 2.5). The calibration pipeline now requires the rest of the L-band Legacy observations to tune each calibration passage and become more robust and

ready for the C-band observations. After running the pipeline on L-band and C-band data, it appears the calibration performance is limited to the amount of RFI still remaining in the data. The exact break down point (i.e. how much RFI still exists in the dataset, and the morphology of that RFI) is not known, but can be tested with the rest of the L-band Legacy data. The PP flux extraction algorithm has been robustly tested on simulated data, and now requires more real maps from COBRaS to push further development.

This last point also raises the consideration of simulated data in testing the performance of software. Simulated data is incredibly useful in creating and testing the performance of certain software as is seen in Chapter 4, as the inputs are known. However, simulated data does not include many unknowns which can occur in real data. For software which acts upon real outputs from real instruments, performance testing on real data has obvious merits, as the real data is itself the final and ultimate performance test.

These considerations are important when planning to develop the next generation of software for future projects such as the SKA. Such endeavours will rely upon previous knowledge and experience from precursor projects such as e-MERLIN and MeerKat, to allow for the construction and implementation of hardware and software to be successful. The SKA will benefit from these pathfinders as well as the vast amount of infrastructure surrounding the project to aid its development, which is sometimes unavailable to smaller projects. With technology becoming more sophisticated and acquiring massive amounts of data, there is an increased importance on software engineering, particularly automated programs to process this data. Software development for the SKA will play an essential part for the project and will become one of the most important aspects of processing and analysing future astronomical surveys.

Appendix A

AIPS Nomenclature

This thesis contains many abbreviated tasks and commands from the Astronomical Image Processing System (AIPS), where a priori knowledge is assumed in the main body of the thesis. A detailed list of all of the tasks and commands (also known as ‘verbs’ and ‘adverbs’) referred to in this thesis is given here for clarity.

AIPS - A reduction, calibration, imaging and analysis package for interferometric data. Originally created for the VLA, it has been widely used since the 1970s as the processing program for other arrays such as MERLIN and e-MERLIN. It contains a large range of algorithms, tasks and commands to manually manipulate datasets and images.

BPASS - Bandpass calibration task. BPASS determines the antenna based complex gain variations as a function of frequency, flattening the bandpass response across the bandwidth. An option to include spectral fitting is available with the SPECindx and SPECurve abverbs.

CALIB - Complex gain calibration task. CALIB uses the Fourier transform and a least squares algorithm to find the best solutions for the electronic gains. It therefore requires an input model to help the least squares algorithm converge onto a solution. No input models are currently available for the calibration sources from e-MERLIN, therefore a point source model is assumed.

CLCAL - Task which applies the solutions determined from calibration tasks (such as FRING and CALIB) and applies them to the individual sources (within a multi-source file) to create calibration tables (CL). These CL tables can then be applied to the sources to calibrate the data “on the fly”.

COMB - A task to combine two overlapping images.

DBCON - A task which concatenates two visibility datasets together, and copies the extension tables over from each dataset.

FLATN - A task which interpolates a series of images created with IMAGR into a single image.

FRING - Delay, rates and phase calibration task. FRING, as with CALIB, uses a least squares algorithm and Fourier Transform to determine the antenna based components of the fringe rate and delays. Therefore a point source calibrator is required for the delays and rates, and a strong point source phase calibrator in the vicinity of the target field is required for the full fringe fitting algorithm. FRING outputs a solution (SN) table which can be converted into a calibration (CL) table to be applied to the visibility data.

GETJY - Task which interpolates flux densities from a primary flux calibrator onto the other sources. GETJY is used in conjunction with SETJY which determines the fluxes for the primary calibrator. GETJY modifies the source (SU) table fluxes for each IF and I, U, V and Q polarisations.

IMAGR - The wide-field imaging task for AIPS. IMAGR contains the Cotton-Schwab (Schwab 1984) version of the CLEAN algorithm which is described in Section 1.1.5. IMAGR is a visibility based CLEAN task which contains a large range of imaging options, including gridding and visibility weightings.

IMAGRPRM - An adverb within the task IMAGR. This adverb performs additional corrections and enhancements to the CLEAN algorithm, notably IMAGRPRM(1) which makes a frequency dependent primary beam correction, where the input is the antenna diameter.

IMEAN - A task to perform statistics on an image. IMEAN determines the mean and RMS brightness of the image or a specified portion of the image. It also creates

a histogram of the image and calculates the noise part of the signal from the peak and RMS of the histogram.

JMFIT - A task to fit Gaussian models to sources in an image via a least squares algorithm. Can fit up to four Gaussians per portion of an image, with or without constraints on the model. It requires an initial guess on the peak position and search area for the least squares algorithm.

KNTR - A task to create contour plots and greyscale plots. KNTR uses a contour tracing algorithm with inputs on the base contour level and sequential levels.

PBCOR - A task to apply the primary beam correction. PBCOR is to be used after the final image has been made to apply the corrections.

POSSM - A plotting task to show complex gains (phases and amplitudes) as a function of frequency. POSSM uses scalar or vector averaging of data and is useful for showing the effects of calibration (CL) tables or bandpass calibration (BP) tables on the visibility dataset.

QUACK - A task to flag beginning or end parts of a scan. The user has to input the length to be removed and the whereabouts in a scan.

REFLG - A task to condense the number of rows in the FG extension table and to flag visibility data. The task is controlled by a number of parameters.

SETJY - A task to set the flux densities of one or more sources. SETJY modifies the flux densities in the source (SU) table for individual IFs and polarisations.

SNEDT - A task which allows manual editing of solution (SN) and calibration (CL) tables. The task plots the solutions or calibration for phases, delays, rates and amplitudes as a function of time for each baseline, IF and polarisation.

SOLINT - The solution interval. An adverb for a number of tasks such as FRING, CALIB etc. and is used to determine how often a solution is made, or how frequently to average the data.

SOUSP - A task to determine the spectral index and/or curvature of a source. Fits a polynominal (up to third order) to a spectrum and outputs the orders to the adverbs SPECindx and SPECURVE.

SPECINDEX - The zeroth order coefficient of the spectral fit.

SPECURVE - The first to third order coefficient of the spectral fit.

SPLAT - A task to split out individual sources from a multi-source file. Calibration (CL), bandpass calibration (BP) and flags (FG) tables may be applied during the splitting process. SPLAT differs from SPLIT by also providing a source (SU) table to the individual source and copies any unapplied tables to the new source file.

SPLIT - A task to split out individual sources from a multi-source file. Calibration (CL), bandpass calibration (BP) and flags (FG) tables may be applied during the splitting process.

UVCON - A task to generate visibility datasets. UVCON generates a simulated visibility dataset for a given array model with antenna positions, sizes etc. Simple models of point sources and resolved sources can be simulated, and system noise can be added in an attempt to replicate a real interferometric system.

This page was intentionally left almost blank

Appendix B

Derivations

This appendix includes the derivations of some of the equations which were not explicitly given in the chapters.

B.1 Derivation of Equation 1.8

Equation 1.4 defines the correlator output for the even part of the sky and equation 1.7 for the odd part of the sky. These are reproduced here

$$r_c(\tau_g) = v_i v_j \cos(2\pi\nu\tau_g)$$

$$r_s(\tau_g) = v_i v_j \sin(2\pi\nu\tau_g).$$

These equations are modified by equation 1.3 for the time delay τ_g and combined to produce the complex correlator output:

$$\begin{aligned} r(\mathbf{b}) &= \int_{\Omega} B(\mathbf{s}) \cos\left(\frac{2\pi\nu \mathbf{b} \cdot \mathbf{s}}{c}\right) d\Omega \\ &\quad - \int_{\Omega} B(\mathbf{s}) \sin\left(\frac{2\pi\nu \mathbf{b} \cdot \mathbf{s}}{c}\right) d\Omega. \end{aligned}$$

Defining the visibilities as a complex function with the form $V = r_c - i r_s$ and using Euler's Formula $e^{i\phi} = \cos\phi + i \sin\phi$ where $\phi = 2\pi\nu(\mathbf{b} \cdot \mathbf{s}/c)$, updates the correlator

function to

$$\begin{aligned}
r(\mathbf{b}) &= \int_{\Omega} B(\mathbf{s}) \left(\frac{e^{i\phi} + e^{-i\phi}}{2} - \frac{e^{i\phi} - e^{-i\phi}}{2} \right) d\Omega \\
&= \int_{\Omega} B(\mathbf{s}) \frac{1}{2} 2e^{-i\phi} d\Omega \\
&= \int_{\Omega} B(\mathbf{s}) e^{-i\phi} d\Omega \\
&= \int_{\Omega} B(\mathbf{s}) e^{-2\pi i \left(\frac{\mathbf{b} \cdot \mathbf{s}}{c} \right)} d\Omega.
\end{aligned}$$

The coordinate system for the baseline vector \mathbf{b} is defined by (u, v, w) , and the corresponding sky distribution vector \mathbf{s} coordinate system is defined by (l, m, n) . Therefore

$$\frac{\nu \mathbf{b} \cdot \mathbf{s}}{c} = ul + vm + wn,$$

and the solid angle of the sky surface $d\Omega$ is given by

$$d\Omega = \frac{dl dm}{n}.$$

Substituting these into the correlator function gives the visibility function or measurement equation:

$$V(u, v, w) = \int_l \int_m \int_n B(l, m, n) e^{-2\pi i (ul + vm + wn)} \frac{dl dm}{n}, \quad (\text{B.1})$$

quod erat demonstrandum.

B.2 Error Propagation Laws

The error propagation laws are given for reference.

Multiplication and Division: $Z = A \cdot B$ or $Z = A/B$:

$$\left(\frac{\Delta Z}{Z} \right)^2 = \left[\left(\frac{\Delta A}{A} \right)^2 + \left(\frac{\Delta B}{B} \right)^2 + \dots \right]. \quad (\text{B.2})$$

Products of powers: $Z = A^X \cdot B^Y$:

$$\left(\frac{\Delta Z}{Z}\right)^2 = \left[(X)^2 \left(\frac{\Delta A}{A}\right)^2 + (Y)^2 \left(\frac{\Delta B}{B}\right)^2 + \dots \right]. \quad (\text{B.3})$$

Logarithmic: $Z = \log_{10} A$:

$$\left(\frac{\Delta Z}{Z}\right)^2 = \left[0.434 \left(\frac{\Delta A}{A}\right)^2 + \dots \right]. \quad (\text{B.4})$$

B.3 Derivation of Mass Loss Rate Errors ($\Delta\dot{M}$)

The individual errors on the O star parameters (M_{spec} , T_{eff} , R and L_*) are given by Martins et al. (2005) to be: $\Delta M = 35$ to 50% , $\Delta T_{\text{eff}} \pm 1000$ to 2000 K, $\Delta R = 10$ to 20% and $L_* = 17$, 10 and 15% for I, III, and V respectively.

The individual errors on the B Supergiant parameters (M_{spec} , T_{eff} , R and L_*) are given by Searle et al. (2008) to be: $\Delta M \pm 14.52$, $\Delta T_{\text{eff}} = 2000$ K, $\Delta R = 10\%$ and $L_* = 57\%$.

The errors on v_∞ is taken as the average range of v_∞ for each luminosity class from Prinja et al. (1990); OI $\Delta v_\infty = 750$ km s $^{-1}$, OIII $\Delta v_\infty = 680$ km s $^{-1}$, OV $\Delta v_\infty = 750$ km s $^{-1}$, BI $\Delta v_\infty = 725$ km s $^{-1}$. These values are roughly equal and are therefore generalised into one error for all stars to be $\Delta v_\infty \pm 360$ km s $^{-1}$.

The final parameter v_{esc} is a function of M , R , and Γ_e which is also a function of L and M .

The errors on Γ_e follow Equation B.2:

$$\left(\frac{\Delta\Gamma_e}{\Gamma_e}\right)^2 = \left[\left(\frac{\Delta L}{L}\right)^2 + \left(\frac{\Delta M}{M}\right)^2 \right]. \quad (\text{B.5})$$

This is then substituted into the error equation for v_{esc} (also following Equation B.2 and B.4):

$$\left(\frac{\Delta v_{\text{esc}}}{v_{\text{esc}}}\right)^2 = \left(\frac{1}{2}\right)^2 \left[\left(\frac{\Delta M}{M}\right)^2 + \left(\frac{\Delta R}{R}\right)^2 + \left(\frac{\Delta\Gamma_e}{\Gamma_e}\right)^2 \right]. \quad (\text{B.6})$$

All of these are propagated to define the error on the mass loss rate $\Delta\dot{M}$ for stars

earlier than and including B0:

$$\begin{aligned}
\left(\frac{\Delta\dot{M}}{\dot{M}}\right)^2 = & 0.434 \left(\left[\frac{0.061}{6.697} \right] + \left[\left\{ \frac{0.021}{2.194} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta L}{L} \right) \right\}^2 \right] \right. \\
& + \left[\left\{ \frac{0.046}{1.313} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta M}{M} \right) \right\}^2 \right] \\
& + \left[\left\{ \frac{0.037}{1.226} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta v_\infty/v_{\text{esc}}}{v_\infty/v_{\text{esc}}} \right) \right\}^2 \right] \\
& + \left[\left\{ \frac{0.064}{0.933} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta T_{\text{eff}}}{T_{\text{eff}}} \right) \right\}^2 \right] \\
& + \left[\left\{ \frac{0.90}{10.92} \right\}^2 + (2)^2 \left\{ 0.434 \left(\frac{\Delta T_{\text{eff}}}{T_{\text{eff}}} \right) \right\}^2 \right] \\
& \left. + \left[\left\{ \frac{0.10}{0.85} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta Z}{Z} \right) \right\} \right] \right), \tag{B.7}
\end{aligned}$$

and for stars later than B0:

$$\begin{aligned}
\left(\frac{\Delta\dot{M}}{\dot{M}}\right)^2 = & 0.434 \left(\left[\frac{0.061}{6.697} \right] + \left[\left\{ \frac{0.021}{2.194} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta L}{L} \right) \right\}^2 \right] \right. \\
& + \left[\left\{ \frac{0.046}{1.313} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta M}{M} \right) \right\}^2 \right] \\
& + \left[\left\{ \frac{0.037}{1.226} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta v_\infty/v_{\text{esc}}}{v_\infty/v_{\text{esc}}} \right) \right\}^2 \right] \\
& + \left[\left\{ \frac{0.064}{0.933} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta T_{\text{eff}}}{T_{\text{eff}}} \right) \right\}^2 \right] \\
& \left. + \left[\left\{ \frac{0.10}{0.85} \right\}^2 + \left\{ 0.434 \left(\frac{\Delta Z}{Z} \right) \right\} \right] \right). \tag{B.8}
\end{aligned}$$

B.4 Derivation of Flux Errors

Equation 5.21 contains four parameters with significant uncertainty, \dot{M} , v_∞ , d , and g_ν . $\Delta\dot{M}$ and Δv_∞ are given above, the error on the distance Δd is ± 0.08 kpc (Rygl et al. 2012) and the error on the Gaunt factor Δg_ν is defined as the difference in the cited electron temperature coefficient 0.85 (Scuderi et al. 1998) defined in Section 5.5 from other values in the literature ranging from 0.5 to 1. Therefore $\Delta g_\nu / g_\nu \sim 0.35$.

The error on the predicted flux ΔS_ν is given by:

$$\left(\frac{\Delta S_\nu}{S_\nu} \right)^2 = \left(\frac{4}{3} \right)^2 \left(\frac{\Delta \dot{M}}{\dot{M}} \right)^2 + \left(\frac{4}{3} \right)^2 \left(\frac{\Delta v_\infty}{v_\infty} \right)^2 + \left(\frac{2}{3} \right)^2 \left(\frac{\Delta g_\nu}{g_\nu} \right)^2 + (2)^2 \left(\frac{\Delta d}{d} \right)^2. \quad (\text{B.9})$$

This page was intentionally left almost blank

Bibliography

Abbott, D. C., Bieging, J. H., Churchwell, E., May 1984. The detection of variable, non-thermal radio emission from two O type stars. *Astrophysical Journal* 280, 671–678.

Abbott, D. C., Bieging, J. H., Churchwell, E., Cassinelli, J. P., May 1980. VLA radio continuum measurements of mass loss from early-type stars. *Astrophysical Journal* 238, 196–202.

Albacete Colombo, J. F., Caramazza, M., Flaccomio, E., Micela, G., Sciortino, S., Nov. 2007a. X-ray flaring from the young stars in Cygnus OB2. *Astronomy & Astrophysics* 474, 495–504.

Albacete Colombo, J. F., Flaccomio, E., Micela, G., Sciortino, S., Damiani, F., Mar. 2007b. Unveiling the Cygnus OB2 stellar population with Chandra. *Astronomy & Astrophysics* 464, 211–227.

Athreya, R., May 2009. A New Approach to Mitigation of Radio Frequency Interference in Interferometric Data. *Astrophysical Journal* 696, 885–890.

Austin, M. J., Jan. 2011. The nature and structure of the winds of galactic O stars. Ph.D. thesis, University College London.

Baan, W. A., Fridman, P. A., Millenaar, R. P., Aug. 2004. Radio Frequency Interference Mitigation at the Westerbork Synthesis Radio Telescope: Algorithms, Test Observations, and System Implementation. *Astrophysical Journal* 128, 933–949.

Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., Witzel, A., Oct. 1977. The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators. *Astronomy & Astrophysics* 61, 99–106.

Babel, J., Montmerle, T., Jul. 1997. X-ray emission from Ap-Bp stars: a magnetically confined wind-shock model for IQ Aur. *Astronomy & Astrophysics* 323, 121–138.

Bell, A. R., Jan. 1978. The acceleration of cosmic rays in shock fronts. I. *Monthly Notices of the Royal Astronomical Society* 182, 147–156.

Berghoefer, T. W., Schmitt, J. H. M. M., Danner, R., Cassinelli, J. P., Jun. 1997. X-ray properties of bright OB-type stars detected in the ROSAT all-sky survey. *Astronomy & Astrophysics* 322, 167–174.

Bertin, E., Arnouts, S., Jun. 1996. SExtractor: Software for source extraction. *Astronomy & Astrophysics Supplement Series* 117, 393–404.

Bieging, J. H., Abbott, D. C., Churchwell, E. B., May 1989. A survey of radio emission from Galactic OB stars. *Astrophysical Journal* 340, 518–536.

Blomme, R., Mar. 2007. Structure in the Winds of OB Stars: Radio and Millimeter Observations. In: Okazaki, A. T., Owocki, S. P., Stefl, S. (Eds.), *Active OB-Stars: Laboratories for Stellare and Circumstellar Physics*. Vol. 361 of *Astronomical Society of the Pacific Conference Series*. p. 186.

Blomme, R., Jan. 2011. Radio observations of massive stars. *Bulletin de la Societe Royale des Sciences de Liege* 80, 67–80.

Blomme, R., De Becker, M., Volpi, D., Rauw, G., Sep. 2010. Non-thermal radio emission from O-type stars. IV. Cygnus OB2 No. 8A. *Astronomy & Astrophysics* 519, A111.

Blomme, R., Prinja, R. K., Runacres, M. C., Colley, S., Feb. 2002. Radio and submillimetre observations of epsilon Ori. *Astronomy & Astrophysics* 382, 921–934.

Blomme, R., van de Steene, G. C., Prinja, R. K., Runacres, M. C., Clark, J. S., Sep. 2003. Radio and submillimetre observations of wind structure in zeta Puppis. *Astronomy & Astrophysics* 408, 715–727.

Bonnarel, F., Fernique, P., Bienaymé, O., Egret, D., Genova, F., Louys, M., Ochsenbein, F., Wenger, M., Bartlett, J. G., Apr. 2000. The ALADIN interactive sky atlas. A reference tool for identification of astronomical sources. *Astronomy & Astrophysics Supplement Series* 143, 33–40.

Bridle, A. H., Schwab, F. R., 1999. Bandwidth and Time-Average Smearing. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 371.

Briggs, D. S., Dec. 1995. High Fidelity Interferometric Imaging: Robust Weighting and NNLS Deconvolution. In: *American Astronomical Society Meeting Abstracts*. Vol. 27 of *Bulletin of the American Astronomical Society*. p. 112.02.

Briggs, D. S., Schwab, F. R., Sramek, R. A., 1999. Imaging. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 127.

Briggs, F. H., Bell, J. F., Kesteven, M. J., Dec. 2000. Removing Radio Interference from Contaminated Astronomical Spectra Using an Independent Reference Signal and Closure Relations. *Astrophysical Journal* 120, 3351–3361.

Brisken, W. F., 2004. Cross Correlators. In: *9th Synthesis Imaging Workshop*, Socorro, New Mexico.

Broos, P. S., Townsley, L. K., Feigelson, E. D., Getman, K. V., Bauer, F. E., Garmire, G. P., May 2010. Innovations in the Analysis of Chandra-ACIS Observations. *Astrophysical Journal* 714, 1582–1605.

Bryc, W., 2002. A uniform approximation to the right normal tail integral. *Applied Mathematics and Computing* 127, 365–374.

Carlson, B., 2012. Cross-correlators for Radio Astronomy. In: *13th Synthesis Imaging Workshop*, Socorro, New Mexico.

Carvalho, P., Rocha, G., Hobson, M. P., Mar. 2009. A fast Bayesian approach to discrete object detection in astronomical data sets - PowellSnakes I. *Monthly Notices of the Royal Astronomical Society* 393, 681–702.

Cash, W., Mar. 1979. Parameter estimation in astronomy through application of the likelihood ratio. *Astrophysical Journal* 228, 939–947.

Castor, J. I., Abbott, D. C., Klein, R. I., Jan. 1975. Radiation-driven winds in Of stars. *Astrophysical Journal* 195, 157–174.

Cazorla, C., Nazé, Y., Rauw, G., Jan. 2014. Wind collisions in three massive stars of Cygnus OB2. *Astronomy & Astrophysics* 561, A92.

Chen, W., White, R. L., Nov. 1994. Nonthermal radio emission from hot star winds: Its origin and physical implications. *Astrophysics and Space Science* 221, 259–272.

Clark, B. G., Sep. 1980. An efficient implementation of the algorithm 'CLEAN'. *Astronomy & Astrophysics* 89, 377.

Clark, G. W., Jan. 1965. Balloon Observation of the X-Ray Spectrum of the Crab Nebula Above 15 keV. *Physical Review Letters* 14, 91–94.

Comerón, F., Pasquali, A., Jul. 2012. New members of the massive stellar population in Cygnus. *Astronomy & Astrophysics* 543, A101.

Comerón, F., Pasquali, A., Rodighiero, G., Stanishev, V., De Filippis, E., López Martí, B., Gálvez Ortiz, M. C., Stankov, A., Gredel, R., Jul. 2002. On the massive star contents of Cygnus OB2. *Astronomy & Astrophysics* 389, 874–888.

Comerón, F., Torra, J., Aug. 2001. Near-infrared imaging of compact HII regions in Cygnus X. *Astronomy & Astrophysics* 375, 539–552.

Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J., May 1998. The NRAO VLA Sky Survey. *Astronomical Journal* 115, 1693–1716.

Conway, J. E., Cornwell, T. J., Wilkinson, P. N., Oct. 1990. Multi-Frequency Synthesis - a New Technique in Radio Interferometric Imaging. *Monthly Notices of the Royal Astronomical Society* 246, 490.

Cooper, B. F. C., Aug. 1970. Correlators with two-bit quantization. *Australian Journal of Physics* 23, 521.

Cornwell, T., Braun, R., Briggs, D. S., 1999. Deconvolution. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 151.

Cornwell, T. J., Evans, K. F., Feb. 1985. A simple maximum entropy deconvolution algorithm. *Astronomy & Astrophysics* 143, 77–83.

Cornwell, T. J., Golap, K., Bhatnagar, S., Nov. 2008. The Noncoplanar Baselines Effect in Radio Interferometry: The W-Projection Algorithm. *IEEE Journal of Selected Topics in Signal Processing* 2, 647–657.

Cox, A. N., 2000. *Allen's Astrophysical Quantities*. New York: AIP Press; Springer.

Crowther, P. A., Sep. 2007. Physical Properties of Wolf-Rayet Stars. *Annual Review of Astronomy & Astrophysics* 45, 177–219.

Crowther, P. A., Lennon, D. J., Walborn, N. R., Jan. 2006. Physical parameters and wind properties of galactic early B supergiants. *Astronomy & Astrophysics* 446, 279–293.

Crowther, P. A., Schnurr, O., Hirschi, R., Yusof, N., Parker, R. J., Goodwin, S. P., Kassim, H. A., Oct. 2010. The R136 star cluster hosts several stars whose individual masses greatly exceed the accepted $150M_{\odot}$ stellar mass limit. *Monthly Notices of the Royal Astronomical Society* 408, 731–751.

De Becker, M., Nov. 2007. Non-thermal emission processes in massive binaries. *Astronomy & Astrophysics Review* 14, 171–216.

De Becker, M., Rauw, G., Nov. 2005. Evidence for phase-locked X-ray variations from the colliding wind massive binary Cyg OB2 #8A. In: Rauw, G., Nazé, Y., Blomme, R., Gosset, E. (Eds.), *Massive Stars and High-Energy Emission in OB Associations*. pp. 73–76.

De Becker, M., Rauw, G., Manfroid, J., Sep. 2004. A Spectroscopic study of the non-thermal radio emitter Cyg OB2 #8A: Discovery of a new binary system. *Astronomy & Astrophysics* 424, L39–L42.

De Becker, M., Rauw, G., Manfroid, J., Eenens, P., Sep. 2006. Early-type stars in the young open cluster IC 1805. II. The probably single stars HD 15570 and HD 15629, and the massive binary/triple system HD 15558. *Astronomy & Astrophysics* 456, 1121–1130.

Dessart, L., Owocki, S. P., Jul. 2005. 2D simulations of the line-driven instability in hot-star winds. II. Approximations for the 2D radiation force. *Astronomy & Astrophysics* 437, 657–666.

Dionne, D., Robert, C., Apr. 2006. Evolutionary Synthesis Models of Young Star-forming Regions: The Influence of Binary Stars. *Astrophysical Journal* 641, 252–267.

Dixon, R. S., Jul. 1970. A Master List of Radio Sources. *Astrophysical Journal Supplement* 20, 1–503.

Dougherty, S. M., Beasley, A. J., Claussen, M. J., Zauderer, B. A., Bolingbroke, N. J., Apr. 2005. High-Resolution Radio Observations of the Colliding-Wind Binary WR 140. *Astrophysical Journal* 623, 447–459.

Dougherty, S. M., Williams, P. M., van der Hucht, K. A., Bode, M. F., Davis, R. J., Jun. 1996. Multifrequency observations of the Wolf-Rayet star WR 146: another colliding-wind binary? *Monthly Notices of the Royal Astronomical Society* 280, 963–970.

Drake, J., Sep. 2009. The Chandra Cygnus OB2 Survey. In: Chandra Proposal. p. 2753.

Drake, J. E., Wright, N. J., Guarcello, M. G., Flaccomio, E., Kashyap, V., in prep. The Chandra Cygnus OB2 Legacy Survey.

Drake, S. A., Linsky, J. L., Nov. 1989. Radio-continuum emission from the ionized stellar winds of warm supergiants. *Astronomical Journal* 98, 1831–1841.

Drew, J. E., Greimel, R., Irwin, M. J., Aungwerojwit, A., Barlow, M. J., Corradi, R. L. M., Drake, J. J., Gänsicke, B. T., Groot, P., Hales, A., Hopewell, E. C., Irwin, J., Knigge, C., Leisy, P., Lennon, D. J., Mampaso, A., Masheder, M. R. W., Matsuura, M., Morales-Rueda, L., Morris, R. A. H., Parker, Q. A., Phillipps, S., Rodriguez-Gil, P., Roelofs, G., Skillen, I., Sokoloski, J. L., Steeghs, D., Unruh, Y. C., Viironen, K., Vink, J. S., Walton, N. A., Witham, A., Wright, N., Zijlstra, A. A., Zurita, A., Sep. 2005. The INT Photometric H α Survey of the Northern Galactic Plane (IPHAS). *Monthly Notices of the Royal Astronomical Society* 362, 753–776.

Eichler, D., Usov, V., Jan. 1993. Particle acceleration and nonthermal radio emission in binaries of early-type stars. *Astrophysical Journal* 402, 271–279.

Eldridge, J. J., Izzard, R. G., Tout, C. A., Mar. 2008. The effect of massive binaries on stellar populations and supernova progenitors. *Monthly Notices of the Royal Astronomical Society* 384, 1109–1118.

Elsmore, B., Kenderdine, S., Ryle, Sir, M., 1966. The operation of the Cambridge one-mile telescope. *Monthly Notices of the Royal Astronomical Society* 134, 87.

Favata, F., Micela, G., Oct. 2003. Stellar Coronal Astronomy. *Space Science Reviews* 108, 577–708.

Fazio, G. G., Hora, J. L., Allen, L. E., Ashby, M. L. N., Barmby, P., Deutsch, L. K., Huang, J.-S., Kleiner, S., Marengo, M., Megeath, S. T., Melnick, G. J., Pahre, M. A., Patten, B. M., Polizotti, J., Smith, H. A., Taylor, R. S., Wang, Z., Willner, S. P., Hoffmann, W. F., Pipher, J. L., Forrest, W. J., McMurtry, C. W., McCreight, C. R., McKelvey, M. E., McMurray, R. E., Koch, D. G., Moseley, S. H., Arendt, R. G., Mentzell, J. E., Marx, C. T., Losch, P., Mayman, P., Eichhorn, W., Krebs, D., Jhabvala, M., Gezari, D. Y., Fixsen, D. J., Flores, J., Shakoorzadeh, K., Jungo, R., Hakun, C., Workman, L., Karpati, G., Kichak, R., Whitley, R., Mann, S., Tollestrup, E. V., Eisenhardt, P., Stern, D., Gorjian, V., Bhattacharya, B., Carey, S., Nelson, B. O., Glaccum, W. J., Lacy, M., Lowrance, P. J., Laine, S., Reach, W. T., Stauffer, J. A., Surace, J. A., Wilson, G., Wright, E. L., Hoffman, A., Domingo, G., Cohen, M., Sep. 2004. The Infrared Array Camera (IRAC) for the Spitzer Space Telescope. *Astrophysical Journal Supplement* 154, 10–17.

Feldmeier, A., Puls, J., Pauldrach, A. W. A., Jun. 1997. A possible origin for X-rays from O stars. *Astronomy & Astrophysics* 322, 878–895.

Fermi, E., Apr. 1949. On the Origin of the Cosmic Radiation. *Physical Review* 75, 1169–1174.

Fisher, R. A., 1966. *The Design of Experiments*, 8th Edition. Hafner Publishing Company, New York.

Fomalont, E. B., Perley, R. A., 1999. Calibration and Editing. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 79.

Fridman, P. A., May 2008. Statistically Stable Estimates of Variance in Radio-Astronomy Observations as Tools for Radio-Frequency Interference Mitigation. *Astrophysical Journal* 135, 1810–1824.

Fullerton, A. W., Massa, D. L., Prinja, R. K., Feb. 2006. The Discordance of Mass-Loss Estimates for Galactic O-Type Stars. *Astrophysical Journal* 637, 1025–1039.

Gagné, M., Oksala, M. E., Cohen, D. H., Tonnesen, S. K., ud-Doula, A., Owocki, S. P., Townsend, R. H. D., MacFarlane, J. J., Nov. 2005. Erratum: “Chandra HETGS Multi-phase Spectroscopy of the Young Magnetic O Star θ^1 Orionis C”. *Astrophysical Journal* 634, 712–713.

Garrington, S. T., Anderson, B., Baines, C., Battilana, J. A., Bentley, M. N., Brown, D., Burgess, P., Diamond, P. J., Kitching, G. J., McCool, R., Muxlow, T. W., Noble, R. G., Roddis, N., Spencer, R. E., Thomasson, P., Oct. 2004. e-MERLIN. In: Oschmann, Jr., J. M. (Ed.), Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Vol. 5489 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. pp. 332–343.

Getman, K. V., Feigelson, E. D., Broos, P. S., Townsley, L. K., Garmire, G. P., Jan. 2010. Methods for Estimating Fluxes and Absorptions of Faint X-ray Sources. *Astrophysical Journal* 708, 1760–1771.

Güdel, M., 2002. Stellar Radio Astronomy: Probing Stellar Atmospheres from Protostars to Giants. *Annual Review of Astronomy & Astrophysics* 40, 217–261.

Hales, C. A., Murphy, T., Curran, J. R., Middelberg, E., Gaensler, B. M., Norris, R. P., Sep. 2012. BLOBCAT: software to catalogue flood-filled blobs in radio images of total intensity and linear polarization. *Monthly Notices of the Royal Astronomical Society* 425, 979–996.

Hamann, W.-R., Oskinova, L. M., Feldmeier, A., Apr. 2008. Spectrum formation in clumpy stellar winds. In: Hamann, W.-R., Feldmeier, A., Oskinova, L. M. (Eds.), *Clumping in Hot-Star Winds*. p. 75.

Hancock, P. J., Murphy, T., Gaensler, B. M., Hopkins, A., Curran, J. R., May 2012. Compact continuum source finding for next generation radio surveys. *Monthly Notices of the Royal Astronomical Society* 422, 1812–1824.

Hanson, M. M., Nov. 2003. A Study of Cygnus OB2: Pointing the Way toward Finding Our Galaxy’s Super-Star Clusters. *Astrophysical Journal* 597, 957–969.

Hillier, D. J., Davidson, K., Ishibashi, K., Gull, T., Jun. 2001. On the Nature of the Central Source in η Carinae. *Astrophysical Journal* 553, 837–860.

Hillier, D. J., Lanz, T., Heap, S. R., Hubeny, I., Smith, L. J., Evans, C. J., Lennon, D. J., Bouret, J. C., May 2003. A Tale of Two Stars: The Extreme O7 Iaf+ Supergiant AV 83 and the OC7.5 III((f)) star AV 69. *Astrophysical Journal* 588, 1039–1063.

Hillier, D. J., Miller, D. L., Mar. 1998. The Treatment of Non-LTE Line Blanketing in Spherically Expanding Outflows. *Astrophysical Journal* 496, 407.

Hobson, M. P., McLachlan, C., Jan. 2003. A Bayesian approach to discrete object detection in astronomical data sets. *Monthly Notices of the Royal Astronomical Society* 338, 765–784.

Hoffman, J., Günther, H. M., Wright, N. J., Nov. 2012. Constraints on the Ubiquity of Coronal X-Ray Cycles. *Astrophysical Journal* 759, 145.

Högbom, J. A., Jun. 1974. Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. *Astronomy & Astrophysics Supplement Series* 15, 417.

Howarth, I. D., Brown, A., 1991. Radio Observations of Massive OB Stars. In: van der Hucht, K. A., Hidayat, B. (Eds.), *Wolf-Rayet Stars and Interrelations with Other Massive Stars in Galaxies*. Vol. 143 of *IAU Symposium*. p. 315.

Jackson, J. C., Jannetta, A. L., Nov. 2006. Legacy data and cosmological constraints from the angular-size/redshift relation for ultracompact radio sources. *Journal of Cosmology and Astroparticle Physics* 11, 2.

Jannetta, A., Jackson, J., Birch, I. P., Kotre, C. J., Robson, K. J., Padgett, R., November 2004. Mammographic image restoration using maximum entropy deconvolution. *Physics in Medicine and Biology* 49 (21), 4997–5010.
URL <http://nrl.northumbria.ac.uk/2580/>

Jansky, K. G., Jul. 1933. Radio Waves from Outside the Solar System. *Nature* 132, 66.

Jennison, R. C., 1958. A phase sensitive interferometer technique for the measurement of the Fourier transforms of spatial brightness distributions of small angular extent. *Monthly Notices of the Royal Astronomical Society* 118, 276.

Johnson, H. L., Morgan, W. W., Mar. 1954. A Heavily Obscured O-Association in Cygnus. *Astrophysical Journal* 119, 344.

Keller, C. U., 1995. X-rays from the sun. *Experientia* 51 (7), 710–720.
URL <http://dx.doi.org/10.1007/BF01941268>

Kennedy, M., Dougherty, S. M., Fink, A., Williams, P. M., Feb. 2010. Modeling the Radio Emission from Cyg OB2 No. 5: A Quadruple System? *Astrophysical Journal* 709, 632–643.

Kharb, P., Lister, M. L., Cooper, N. J., Feb. 2010. Extended Radio Emission in MOJAVE Blazars: Challenges to Unification. *Astrophysical Journal* 710, 764–782.

Kiminki, D. C., Kobulnicky, H. A., May 2012. An Updated Look at Binary Characteristics of Massive Stars in the Cygnus OB2 Association. *Astrophysical Journal* 751, 4.

Kiminki, D. C., Kobulnicky, H. A., Ewing, I., Bagley Kiminki, M. M., Lundquist, M., Alexander, M., Vargas-Alvarez, C., Choi, H., Henderson, C. B., Mar. 2012a. Additional Massive Binaries in the Cygnus OB2 Association. *Astrophysical Journal* 747, 41.

Kiminki, D. C., Kobulnicky, H. A., Ewing, I., Bagley Kiminki, M. M., Lundquist, M., Alexander, M., Vargas-Alvarez, C., Choi, H., Henderson, C. B., Mar. 2012b. Additional Massive Binaries in the Cygnus OB2 Association. *Astrophysical Journal* 747, 41.

Kiminki, D. C., Kobulnicky, H. A., Gilbert, I., Bird, S., Chunev, G., Jun. 2009. Five More Massive Binaries in the Cygnus OB2 Association. *Astronomical Journal* 137, 4608–4620.

Kiminki, D. C., Kobulnicky, H. A., Kinemuchi, K., Irwin, J. S., Fryer, C. L., Berrington, R. C., Uzpen, B., Monson, A. J., Pierce, M. J., Woosley, S. E., Aug. 2007. A Radial Velocity Survey of the Cyg OB2 Association. *Astrophysical Journal* 664, 1102–1120.

Kiminki, D. C., McSwain, M. V., Kobulnicky, H. A., Jun. 2008. New Massive Binaries in the Cygnus OB2 Association. *Astrophysical Journal* 679, 1478–1489.

Klochkova, V. G., Chentsov, E. L., Dec. 2004. The Optical Spectrum of an LBV Candidate in the Cyg OB2 Association. *Astronomy Reports* 48, 1005–1018.

Knödlseder, J., Aug. 2000. Cygnus OB2 - a young globular cluster in the Milky Way. *Astronomy & Astrophysics* 360, 539–548.

Kobulnicky, H. A., Gilbert, I. J., Kiminki, D. C., Feb. 2010. OB Stars and Stellar Bow shocks in Cygnus-X: A Novel Laboratory Estimating Stellar Mass Loss Rates. *Astrophysical Journal* 710, 549–566.

Kobulnicky, H. A., Smullen, R. A., Kiminki, D. C., Runnoe, J. C., Wood, E. S., Long, G., Alexander, M. J., Lundquist, M. J., Vargas-Alvarez, C., Sep. 2012. A Fresh Catch of Massive Binaries in the Cygnus OB2 Association. *Astrophysical Journal* 756, 50.

Kudritzki, R.-P., Puls, J., 2000. Winds from Hot Stars. *Annual Review of Astronomy & Astrophysics* 38, 613–666.

Kuehr, H., Nauber, U., Pauliny-Toth, I. I. K., 1979. A Catalogue of radio sources. Bonn: Max-Planck-Institut (MPI) für Radioastronomie, 1979.

Kuiper, N. H., 1960. Tests concerning random points on a circle. *Proceedings of the Koninklijke Nederlandse Akademie Van Wetenschappen* 63, 38–47.

Kundu, M. R., Dec. 1961. Bursts of Centimeter-Wave Emission and the Region of Origin of X Rays from Solar Flares. *Journal of Geophysical Research* 66, 4308–4312.

Lamers, H. J. G. L. M., Cassinelli, J. P., 1999. Book Review: *Introduction to stellar winds* / Cambridge U Press, 1999. *Irish Astronomical Journal* 26, 171.

Lamers, H. J. G. L. M., Leitherer, C., Aug. 1993. What are the mass-loss rates of O stars? *Astrophysical Journal* 412, 771–791.

Lamers, H. J. G. L. M., Snow, T. P., Lindholm, D. M., Dec. 1995. Terminal Velocities and the Bistability of Stellar Winds. *Astrophysical Journal* 455, 269.

Lawrence, L. C., Reddish, V. C., 1965. The Cygnus II association. I. Intercomparison of photometries with the Edinburgh and Hamburg Schmidt telescopes. *Publications of the Royal Observatory of Edinburgh* 3, 275–308.

Lee, H.-T., Chen, W. P., Mar. 2007. Triggered Star Formation by Massive Stars. *Astrophysical Journal* 657, 884–896.

Lépine, S., Moffat, A. F. J., Aug. 2008. Direct Spectroscopic Observations of Clumping in O-Star Winds. *Astronomical Journal* 136, 548–553.

Levenberg, K., 1944. A Method for the Solution of Certain Non-Linear Problems in Least Squares. *Quarterly of Applied Mathematics* 2, 164–168.

Lister, M. L., Aller, H. D., Aller, M. F., Cohen, M. H., Homan, D. C., Kadler, M., Kellermann, K. I., Kovalev, Y. Y., Ros, E., Savolainen, T., Zensus, J. A., Vermeulen,

R. C., Mar. 2009. MOJAVE: Monitoring of Jets in Active Galactic Nuclei with VLBA Experiments. V. Multi-Epoch VLBA Images. *Astronomical Journal* 137, 3718–3729.

Lomb, N. R., Feb. 1976. Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science* 39, 447–462.

Lucy, L. B., Solomon, P. M., Mar. 1970. Mass Loss by Hot Stars. *Astrophysical Journal* 159, 879.

Luhman, K. L., D'Alessio, P., Calvet, N., Allen, L. E., Hartmann, L., Megeath, S. T., Myers, P. C., Fazio, G. G., Feb. 2005. Spitzer Identification of the Least Massive Known Brown Dwarf with a Circumstellar Disk. *Astrophysical Journal Letters* 620, L51–L54.

Maeder, A., Meynet, G., Sep. 2000. Stellar evolution with rotation. VI. The Eddington and Omega -limits, the rotational mass loss for OB and LBV stars. *Astronomy & Astrophysics* 361, 159–166.

Maíz-Apellániz, J., Walborn, N. R., Galué, H. Á., Wei, L. H., Mar. 2004. A Galactic O Star Catalog. *Astrophysical Journal Supplement* 151, 103–148.

Markova, N., Puls, J., Repolust, T., Markov, H., Jan. 2004. Bright OB stars in the Galaxy. I. Mass-loss and wind-momentum rates of O-type stars: A pure H α analysis accounting for line-blanketing. *Astronomy & Astrophysics* 413, 693–709.

Markova, N., Puls, J., Scuderi, S., Markov, H., Sep. 2005. Bright OB stars in the Galaxy. II. Wind variability in O supergiants as traced by H α . *Astronomy & Astrophysics* 440, 1133–1151.

Marquardt, D., 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM Journal on Applied Mathematics* 11, 431–441.

Martí, J., Paredes, J. M., Ishwara Chandra, C. H., Bosch-Ramon, V., Sep. 2007. Deep radio images of the HEGRA and Whipple TeV sources in the Cygnus OB2 region. *Astronomy & Astrophysics* 472, 557–564.

Martins, F., Plez, B., Oct. 2006. UVBJHK synthetic photometry of Galactic O stars. *Astronomy & Astrophysics* 457, 637–644.

Martins, F., Schaerer, D., Hillier, D. J., Jun. 2005. A new calibration of stellar parameters of Galactic O stars. *Astronomy & Astrophysics* 436, 1049–1065.

Masias, M., Freixenet, J., Lladó, X., Peracaula, M., May 2012. A review of source detection approaches in astronomical images. *Monthly Notices of the Royal Astronomical Society* 422, 1674–1689.

Massa, D., Fullerton, A. W., Sonneborn, G., Hutchings, J. B., Apr. 2003. Constraints on the Ionization Balance of Hot-Star Winds from FUSE Observations of O Stars in the Large Magellanic Cloud. *Astrophysical Journal* 586, 996–1018.

Massa, D., Oschinova, L. M., Fullerton, A. W., Prinja, R. K., Bohlender, D. A., Morrison, N. D., Blake, M., Pych, W., 2014; in press. CIR Modulation of the X-ray Flux from the O7.5III(n)(f) Star Xi Persei. *Monthly Notices of the Royal Astronomical Society*.

Massey, P., Thompson, A. B., Apr. 1991. Massive stars in CYG OB2. *Astronomical Journal* 101, 1408–1428.

McCready, L. L., Pawsey, J. L., Payne-Scott, R., Aug. 1947. Solar Radiation at Radio Frequencies and Its Relation to Sunspots. *Royal Society of London Proceedings Series A* 190, 357–375.

McEwen, J. D., Wiaux, Y., May 2011. Compressed sensing for wide-field radio interferometric imaging. *Monthly Notices of the Royal Astronomical Society* 413, 1318–1332.

Milne, E. A., May 1926. On the possibility of the emission of high-speed atoms from the sun and stars. *Monthly Notices of the Royal Astronomical Society* 86, 459–473.

Morgan, W. W., Johnson, H. L., Roman, N. G., Apr. 1954. A Very Red Star of Early Type in Cygnus. *Publications of the Astronomical Society of the Pacific* 66, 85.

Morgan, W. W., Keenan, P. C., Kellman, E., 1943. An atlas of stellar spectra, with an outline of spectral classification. Chicago, Ill., The University of Chicago press [1943].

Morton, D. C., Mar. 1967. The Far-Ultraviolet Spectra of Six Stars in Orion. *Astrophysical Journal* 147, 1017.

Muijres, L. E., Vink, J. S., de Koter, A., Müller, P. E., Langer, N., Jan. 2012. Predictions for mass-loss rates and terminal wind velocities of massive O-type stars. *Astronomy & Astrophysics* 537, A37.

Mullan, D. J., Sep. 1986. Displaced narrow absorption components in the spectra of mass-losing OB stars - Indications of corotating interaction regions? *Astronomy & Astrophysics* 165, 157–162.

Münch, L., Morgan, W. W., Jul. 1953. Notes: A Probable Clustering of Blue Giants in Cygnus. *Astrophysical Journal* 118, 161–162.

Muntean, V., Moffat, A. F. J., Chené, A. N., de La Chevrotière, A., Nov. 2009. The Galactic hybrid Wolf-Rayet WN7o/CE + O7V((f)) binary system WR145. *Monthly Notices of the Royal Astronomical Society* 399, 1977–1987.

Muxlow, T. W., 2007. Wide-field Imaging. In: European Radio Interferometry School 2007, Bonn, Germany.

Napier, P. J., 1999. The Primary Antenna Elements. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 37.

Napier, P. J., Thompson, A. R., Ekers, R. D., Nov. 1983. The Very Large Array - Design and performance of a modern synthesis radio telescope. *IEEE Proceedings* 71, 1295–1320.

Narayan, R., Nityananda, R., 1986. Maximum entropy image restoration in astronomy. *Annual Review of Astronomy & Astrophysics* 24, 127–170.

Nazé, Y., Nov. 2009. Hot stars observed by XMM-Newton. I. The catalog and the properties of OB stars. *Astronomy & Astrophysics* 506, 1055–1064.

Nazé, Y., Broos, P. S., Osokinova, L., Townsley, L. K., Cohen, D., Corcoran, M. F., Evans, N. R., Gagné, M., Moffat, A. F. J., Pittard, J. M., Rauw, G., ud-Doula, A., Walborn, N. R., May 2011. Global X-ray Properties of the O and B Stars in Carina. *Astrophysical Journal Supplement* 194, 7.

Nazé, Y., Damerdji, Y., Rauw, G., Kiminki, D. C., Mahy, L., Kobulnicky, H. A., Morel, T., De Becker, M., Eenens, P., Barbieri, C., Aug. 2010. First Orbital Solution for the Non-thermal Emitter Cyg OB2 No. 9. *Astrophysical Journal* 719, 634–641.

Nazé, Y., De Becker, M., Rauw, G., Barbieri, C., May 2008. A binary signature in the non-thermal radio-emitter Cyg OB2 #9. *Astronomy & Astrophysics* 483, 543–546.

Nazé, Y., Oskinova, L. M., Gosset, E., Feb. 2013. A Detailed X-Ray Investigation of ζ Puppis. II. The Variability on Short and Long Timescales. *Astrophysical Journal* 763, 143.

Negueruela, I., Marco, A., Herrero, A., Clark, J. S., Aug. 2008. New very massive stars in Cygnus OB2. *Astronomy & Astrophysics* 487, 575–581.

Niamsuwan, N., Johnson, J. T., Ellingson, S. W., Jun. 2005. Examination of a simple pulse-blanking technique for radio frequency interference mitigation. *Radio Science* 40, 5.

Noordam, J. E., Smirnov, O. M., Dec. 2010. The MeqTrees software system and its use for third-generation calibration of radio interferometers. *Astronomy & Astrophysics* 524, A61.

Offringa, A. R., Jun. 2012. Algorithms for radio interference detection and removal. Ph.D. thesis, University of Groningen.

Offringa, A. R., de Bruyn, A. G., Biehl, M., Zaroubi, S., Bernardi, G., Pandey, V. N., Jun. 2010a. Post-correlation radio frequency interference classification methods. *Monthly Notices of the Royal Astronomical Society* 405, 155–167.

Offringa, A. R., de Bruyn, A. G., Zaroubi, S., Biehl, M., Jul. 2010b. A LOFAR RFI detection pipeline and its first results. in proceedings of “RFI mitigation workshop” PoS(RFI2010)036.

O’Neal, D., Feigelson, E. D., Mathieu, R. D., Myers, P. C., Nov. 1990. A radio survey of weak T Tauri stars in Taurus-Auriga. *Astronomical Journal* 100, 1610–1617.

Ortiz-León, G. N., Rodríguez, L. F., Tapia, M., Apr. 2012. Extended Non-Thermal Emission Possibly Associated with Cyg OB2 #5. *Revista Mexicana de Astronomía y Astrofísica* 48, 85–93.

Oskinova, L. M., Clarke, D., Pollock, A. M. T., Oct. 2001. Rotationally modulated X-ray emission from the single O star ζ Ophiuchi. *Astronomy & Astrophysics* 378, L21–L24.

Oskinova, L. M., Hamann, W.-R., Feldmeier, A., Dec. 2007. Neglecting the porosity of hot-star winds can lead to underestimating mass-loss rates. *Astronomy & Astrophysics* 476, 1331–1340.

Owocki, S. P., Gayley, K. G., Shaviv, N. J., Nov. 2004. A Porosity-Length Formalism for Photon-Tiring-limited Mass Loss from Stars above the Eddington Limit. *Astrophysical Journal* 616, 525–541.

Owocki, S. P., Rybicki, G. B., Sep. 1984. Instabilities in line-driven stellar winds. I - Dependence on perturbation wavelength. *Astrophysical Journal* 284, 337–350.

Paltani, S., Jun. 2004. Searching for periods in X-ray observations using Kuiper's test. Application to the ROSAT PSPC archive. *Astronomy & Astrophysics* 420, 789–797.

Panagia, N., Felli, M., Feb. 1975. The spectrum of the free-free radiation from extended envelopes. *Astronomy & Astrophysics* 39, 1–5.

Pavlov, G. G., Zavlin, V. E., Aschenbach, B., Trümper, J., Sanwal, D., Mar. 2000. The Compact Central Object in Cassiopeia A: A Neutron Star with Hot Polar Caps or a Black Hole? *Astrophysical Journal Letters* 531, L53–L56.

Pearson, T. J., 1999. Non-Imaging Data Analysis. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 335.

Peck, L. W., Fenech, D. M., Aug. 2013. SERPent: Automated reduction and RFI-mitigation software for e-MERLIN. *Astronomy and Computing* 2, 54–66.

Perley, R. A., Butler, B. J., Feb. 2013. An Accurate Flux Density Scale from 1 to 50 GHz. *Astrophysical Journal Supplement* 204, 19.

Persi, P., Ferrari-Toniolo, M., Tapia, M., Roth, M., Rodriguez, L. F., Jan. 1985. Time-variable, excess radio emission from CYG OB2 No. 5. *Astronomy & Astrophysics* 142, 263–267.

Pigulski, A., Kołaczkowski, Z., Aug. 1998. Mapping the β Cephei instability strip: photometric variability of stars in the central part of the Cygnus OB2 association. *Monthly Notices of the Royal Astronomical Society* 298, 753–762.

Pittard, J. M., May 2007. A Clumping-independent Diagnostic of Stellar Mass-Loss Rates: Rapid Clump Destruction in Adiabatic Colliding Winds. *Astrophysical Journal Letters* 660, L141–L144.

Pittard, J. M., Dougherty, S. M., Oct. 2006. Radio, X-ray, and γ -ray emission models of the colliding-wind binary WR140. *Monthly Notices of the Royal Astronomical Society* 372, 801–826.

Prinja, R. K., Barlow, M. J., Howarth, I. D., Oct. 1990. Terminal velocities for a large sample of O stars, B supergiants, and Wolf-Rayet stars. *Astrophysical Journal* 361, 607–620.

Prinja, R. K., Massa, D., Searle, S. C., Jan. 2005. The ionization structure of early-B supergiant winds. *Astronomy & Astrophysics* 430, L41–L44.

Prinja, R. K., Massa, D. L., Oct. 2010. Signature of wide-spread clumping in B supergiant winds. *Astronomy & Astrophysics* 521, L55.

Prinja, R. K., Massa, D. L., Nov. 2013. Ultraviolet diagnostic of porosity-free mass-loss estimates in B stars. *Astronomy & Astrophysics* 559, A15.

Puls, J., Markova, N., Scuderi, S., Stanghellini, C., Taranova, O. G., Burnley, A. W., Howarth, I. D., Aug. 2006. Bright OB stars in the Galaxy. III. Constraints on the radial stratification of the clumping factor in hot star winds from a combined $H\alpha$, IR and radio analysis. *Astronomy & Astrophysics* 454, 625–651.

Puls, J., Repolust, T., Hoffmann, T. L., Jokuthy, A., Venero, R. O. J., 2003. Advances in radiatively driven wind models. In: van der Hucht, K., Herrero, A., Esteban, C. (Eds.), *A Massive Star Odyssey: From Main Sequence to Supernova*. Vol. 212 of IAU Symposium. p. 61.

Puls, J., Vink, J. S., Najarro, F., Dec. 2008. Mass loss from hot massive stars. *Astronomy & Astrophysics Review* 16, 209–325.

Rau, U., Cornwell, T. J., Aug. 2011. A multi-scale multi-frequency deconvolution algorithm for synthesis imaging in radio interferometry. *Astronomy & Astrophysics* 532, A71.

Rauw, G., Naze, Y., Wright, N. J., Drake, J. J., Guarcello, M. G., Prinja, R. K., Peck, L. W., Albacete Colombo, J. F., Herrero, A., Kobulnicky, H. A., Sciortino, S., Vink, J. S., Jan. 2014. X-ray emission from massive stars in Cyg OB2. *Astrophysical Journal Supplement Special Issue devoted to the Chandra Cygnus OB2 Legacy Survey*.

Rauw, G., Vreux, J.-M., Bohannan, B., May 1999. The Interacting Early-Type Binary BD +40 deg4220 (V729 Cyg): Modeling the Colliding Winds Region. *Astrophysical Journal* 517, 416–430.

Reddish, V. C., Lawrence, L. C., Pratt, N. M., 1966. The Cygnus II association 2. The Distribution of Stars and Interstellar Matter. *Publications of the Royal Observatory of Edinburgh* 5, 111–180.

Reddish, V. C., Lawrence, L. C., Pratt, N. N., 1967. The Cygnus II association-II. The distribution of stars and interstellar matter. *Monthly Notices of the Royal Astronomical Society* 136, 428–429.

Reid, A. H. N., Howarth, I. D., Jul. 1996. Optical time-series spectroscopy of the O4 supergiant ζ Puppis. *Astronomy & Astrophysics* 311, 616–630.

Repolust, T., Puls, J., Herrero, A., Feb. 2004. Stellar and wind parameters of Galactic O-stars. The influence of line-blocking/blanketing. *Astronomy & Astrophysics* 415, 349–376.

Rieke, G. H., Young, E. T., Engelbracht, C. W., Kelly, D. M., Low, F. J., Haller, E. E., Beeman, J. W., Gordon, K. D., Stansberry, J. A., Misselt, K. A., Cadien, J., Morrison, J. E., Rivlis, G., Latter, W. B., Noriega-Crespo, A., Padgett, D. L., Stapelfeldt, K. R., Hines, D. C., Egami, E., Muzerolle, J., Alonso-Herrero, A., Blaylock, M., Dole, H., Hinz, J. L., Le Floc'h, E., Papovich, C., Pérez-González, P. G., Smith, P. S., Su, K. Y. L., Bennett, L., Frayer, D. T., Henderson, D., Lu, N., Masci, F., Pesenson, M., Rebull, L., Rho, J., Keene, J., Stolovy, S., Wachter, S., Wheaton, W., Werner, M. W., Richards, P. L., Sep. 2004. The Multiband Imaging Photometer for Spitzer (MIPS). *Astrophysical Journal Supplement* 154, 25–29.

Roshi, D. A., Perley, R. A., 2003. A New Technique to Improve RFI Suppression in Radio Interferometers. In: Y. C. Minh (Ed.), *Astronomical Society of the Pacific Conference Series*. Vol. 306 of *Astronomical Society of the Pacific Conference Series*. p. 109.

Rousseeuw, P. J., Croux, C., 1993. Alternatives to the Median Absolute Deviation. *Journal of the American Statistical Association* 88 (424), 1273–1283.

URL <http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1993.10476408>

Runacres, M. C., Owocki, S. P., Jan. 2002. The outer evolution of instability-generated structure in radiatively driven stellar winds. *Astronomy & Astrophysics* 381, 1015–1025.

Rygl, K. L. J., Brunthaler, A., Sanna, A., Menten, K. M., Reid, M. J., van Langevelde, H. J., Honma, M., Torstensson, K. J. E., Fujisawa, K., Mar. 2012. Parallaxes and proper motions of interstellar masers toward the Cygnus X star-forming complex. I. Membership of the Cygnus X region. *Astronomy & Astrophysics* 539, A79.

Ryle, M., Mar. 1952. A New Radio Interferometer and Its Application to the Observation of Weak Radio Stars. *Royal Society of London Proceedings Series A* 211, 351–375.

Sana, H., de Koter, A., de Mink, S. E., Dunstall, P. R., Evans, C. J., Hénault-Brunet, V., Maíz Apellániz, J., Ramírez-Agudelo, O. H., Taylor, W. D., Walborn, N. R., Clark, J. S., Crowther, P. A., Herrero, A., Gieles, M., Langer, N., Lennon, D. J., Vink, J. S., Feb. 2013. The VLT-FLAMES Tarantula Survey. VIII. Multiplicity properties of the O-type star population. *Astronomy & Astrophysics* 550, A107.

Sana, H., Gosset, E., Nazé, Y., Rauw, G., Linder, N., May 2008. The massive star binary fraction in young open clusters - I. NGC 6231 revisited. *Monthly Notices of the Royal Astronomical Society* 386, 447–460.

Sault, R. J., Wieringa, M. H., Dec. 1994. Multi-frequency synthesis techniques in radio interferometric imaging. *Astronomy & Astrophysics Supplement Series* 108, 585–594.

Savage, R. S., Oliver, S., Jun. 2007. Bayesian Methods of Astronomical Source Extraction. *Astrophysical Journal* 661, 1339–1346.

Scargle, J. D., Dec. 1982. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *Astrophysical Journal* 263, 835–853.

Schulte, D. H., Nov. 1956a. New Members of the Association VI Cygni. *Astrophysical Journal* 124, 530.

Schulte, D. H., Mar. 1956b. Some Recent Results of Low-Dispersion Spectral Classification. *Astrophysical Journal* 123, 250.

Schwab, F. R., 1984. Optimal Gridding of Visibility Data in Radio Interferometry. In:

Roberts, J. A. (Ed.), *Indirect Imaging. Measurement and Processing for Indirect Imaging*. p. 333.

Schwab, F. R., Cotton, W. D., May 1983. Global fringe search techniques for VLBI. *Astronomical Journal* 88, 688–694.

Scuderi, S., Panagia, N., Stanghellini, C., Trigilio, C., Umana, G., Apr. 1998. Radio observations of stellar winds from early type stars. *Astronomy & Astrophysics* 332, 251–267.

Searle, S. C., Prinja, R. K., Massa, D., Ryans, R., Apr. 2008. Quantitative studies of the optical and UV spectra of Galactic early B supergiants. I. Fundamental parameters. *Astronomy & Astrophysics* 481, 777–797.

Setia Gunawan, D. Y. A., de Bruyn, A. G., van der Hucht, K. A., Williams, P. M., Nov. 2003. A Westerbork Synthesis Radio Telescope 1400 and 350 MHz Continuum Survey of the Cygnus OB2 Association, in Search of Hot Massive Stars. *Astrophysical Journal Supplement* 149, 123–156.

Shepherd, M. C., 1997. Difmap: an Interactive Program for Synthesis Imaging. In: Hunt, G., Payne, H. (Eds.), *Astronomical Data Analysis Software and Systems VI*. Vol. 125 of *Astronomical Society of the Pacific Conference Series*. p. 77.

Simpson, D. E., Fleming, J. S., Aldous, A. J., Daniell, G. J., 1995. Deconvolution of planar scintigrams by maximum entropy. *Physics in Medicine and Biology* 40, 153–162.

Singh, N., Browne, L.-M., Bulter, R., Aug. 2013. Parallel astronomical data processing with Python: Recipes for multicore machines. *Astronomy and Computing* 2, 1–10.

Sivia, D. S., Skilling, J., 2006. *Data Analysis: A Bayesian Tutorial*, 2nd Edition. Oxford University Press, Great Clarendon Street, Oxford OX2 6DP.

Skiff, B. A., Nov. 2010. Catalogue of Stellar Spectral Classifications (Skiff, 2009-2012). VizieR Online Data Catalog 1, 2023.

Skinner, S. L., Brown, A., Stewart, R. T., Jul. 1993. A high-sensitivity survey of radio continuum emission from Herbig Ae/Be stars. *Astrophysical Journal Supplement* 87, 217–265.

Skrutskie, M. F., Cutri, R. M., Stiening, R., Weinberg, M. D., Schneider, S., Carpenter, J. M., Beichman, C., Capps, R., Chester, T., Elias, J., Huchra, J., Liebert, J., Lonsdale, C., Monet, D. G., Price, S., Seitzer, P., Jarrett, T., Kirkpatrick, J. D., Gizis, J. E., Howard, E., Evans, T., Fowler, J., Fullmer, L., Hurt, R., Light, R., Kopan, E. L., Marsh, K. A., McCallon, H. L., Tam, R., Van Dyk, S., Wheelock, S., Feb. 2006. The Two Micron All Sky Survey (2MASS). *Astronomical Journal* 131, 1163–1183.

Stevens, I. R., Blondin, J. M., Pollock, A. M. T., Feb. 1992. Colliding winds from early-type stars in binary systems. *Astrophysical Journal* 386, 265–287.

Taylor, A. R., Mar. 2013. The Square Kilometre Array. In: van Leeuwen, J. (Ed.), IAU Symposium. Vol. 291 of IAU Symposium. pp. 337–341.

Taylor, M. B., Dec. 2005. TOPCAT & STIL: Starlink Table/VOTable Processing Software. In: Shopbell, P., Britton, M., Ebert, R. (Eds.), *Astronomical Data Analysis Software and Systems XIV*. Vol. 347 of *Astronomical Society of the Pacific Conference Series*. p. 29.

Taylor, M. B., Jul. 2006. STILTS - A Package for Command-Line Processing of Tabular Data. In: Gabriel, C., Arviset, C., Ponz, D., Enrique, S. (Eds.), *Astronomical Data Analysis Software and Systems XV*. Vol. 351 of *Astronomical Society of the Pacific Conference Series*. p. 666.

Thompson, A. R., 1999. Fundamentals of Radio Interferometry. In: Taylor, G. B., Carilli, C. L., Perley, R. A. (Eds.), *Synthesis Imaging in Radio Astronomy II*. Vol. 180 of *Astronomical Society of the Pacific Conference Series*. p. 11.

Thompson, A. R., Moran, J. M., Swenson, G. W., 1986. *Interferometry and synthesis in radio astronomy*. New York, Wiley-Interscience, 1986, 554 p.

ud-Doula, A., Owocki, S. P., Sep. 2002. Dynamical Simulations of Magnetically Channeled Line-driven Stellar Winds. I. Isothermal, Nonrotating, Radially Driven Flow. *Astrophysical Journal* 576, 413–428.

Vacca, W. D., Garmany, C. D., Shull, J. M., Apr. 1996. The Lyman-Continuum Fluxes and Stellar Parameters of O and Early B-Type Stars. *Astrophysical Journal* 460, 914.

Van Loo, S., 2005. Non-thermal radio emission from single. Ph.D. thesis, Royal Observatory of Belgium Ringlaan 3 B-1180 Brussels.

van Loo, S., Runacres, M. C., Blomme, R., Jun. 2006. Can single O stars produce non-thermal radio emission? *Astronomy & Astrophysics* 452, 1011–1019.

Vink, J. S., de Koter, A., Lamers, H. J. G. L. M., Oct. 1999. On the nature of the bi-stability jump in the winds of early-type supergiants. *Astronomy & Astrophysics* 350, 181–196.

Vink, J. S., de Koter, A., Lamers, H. J. G. L. M., Oct. 2000. New theoretical mass-loss rates of O and B stars. *Astronomy & Astrophysics* 362, 295–309.

Vink, J. S., de Koter, A., Lamers, H. J. G. L. M., Apr. 2001. Mass-loss predictions for O and B stars as a function of metallicity. *Astronomy & Astrophysics* 369, 574–588.

Vink, J. S., Drew, J. E., Steeghs, D., Wright, N. J., Martin, E. L., Gänsicke, B. T., Greimel, R., Drake, J., Jun. 2008. IPHAS discoveries of young stars towards Cyg OB2 and its southern periphery. *Monthly Notices of the Royal Astronomical Society* 387, 308–318.

Watson, S. K., Davis, R. J., Williams, P. M., Bode, M. F., Aug. 2002. Thermal and non-thermal radio variations in WR 147. *Monthly Notices of the Royal Astronomical Society* 334, 631–636.

Weber, R., Faye, C., Biraud, F., Dansou, J., Nov. 1997. Spectral detector for interference time blanking using quantized correlator. *Astronomy & Astrophysics Supplement Series* 126, 161–167.

Weisskopf, M. C., Tananbaum, H. D., Van Speybroeck, L. P., O'Dell, S. L., Jul. 2000. Chandra X-ray Observatory (CXO): overview. In: Truemper, J. E., Aschenbach, B. (Eds.), *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 4012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. pp. 2–16.

Werner, M. W., Roellig, T. L., Low, F. J., Rieke, G. H., Rieke, M., Hoffmann, W. F., Young, E., Houck, J. R., Brandl, B., Fazio, G. G., Hora, J. L., Gehrz, R. D., Helou, G., Soifer, B. T., Stauffer, J., Keene, J., Eisenhardt, P., Gallagher, D., Gautier, T. N., Irace, W., Lawrence, C. R., Simmons, L., Van Cleve, J. E., Jura, M., Wright, E. L., Cruikshank, D. P., Sep. 2004. The Spitzer Space Telescope Mission. *Astrophysical Journal Supplement* 154, 1–9.

Wessolowski, U., Feb. 1996. X-ray emission from (putatively) single WN-type W-R stars. In: Zimmermann, H. U., Trümper, J., Yorke, H. (Eds.), *Roentgenstrahlung from the Universe*. pp. 75–76.

White, R. L., Feb. 1985. Synchrotron emission from chaotic stellar winds. *Astrophysical Journal* 289, 698–708.

White, R. L., Becker, R. H., Helfand, D. J., Aug. 2005. New Catalogs of Compact Radio Sources in the Galactic Plane. *Astronomical Journal* 130, 586–596.

Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., Vandergheynst, P., May 2009. Compressed sensing imaging techniques for radio interferometry. *Monthly Notices of the Royal Astronomical Society* 395, 1733–1742.

Wolff, S. C., Strom, S. E., Dror, D., Venn, K., Mar. 2007. Rotational Velocities for B0-B3 Stars in Seven Young Clusters: Further Study of the Relationship between Rotation Speed and Density in Star-Forming Regions. *Astronomical Journal* 133, 1092–1103.

Wright, A. E., Barlow, M. J., Jan. 1975. The radio and infrared spectrum of early-type stars undergoing mass loss. *Monthly Notices of the Royal Astronomical Society* 170, 41–51.

Wright, N. J., May 2011. Chandra Observations Of The Massive Young Association Cygnus OB2. In: *American Astronomical Society Meeting Abstracts #218*. p. 133.08.

Wright, N. J., Drake, J. J., Sep. 2009. The Massive Star-Forming Region Cygnus OB2. I. Chandra Catalog of Association Members. *Astrophysical Journal Supplement* 184, 84–99.

Wright, N. J., Drake, J. J., Guarcello, M. G., Aldcroft, T. L., Kashyap, V. L., Damiani, F., De Pasquale, J., Fruscione, A., in prep. The Chandra Cygnus OB2 Legacy Survey: Design and X-ray Point Source Catalog.

“The world is a thing of utter inordinate complexity and richness and strangeness that is absolutely awesome. I mean the idea that such complexity can arise not only out of such simplicity, but probably out of nothing, is the most fabulous extraordinary idea. And once you get some kind of inkling of how that might have happened, it’s just wonderful. And . . . the opportunity to spend 70 or 80 years of your life in such a universe is time well spent as far as I am concerned.”

Douglas Adams