


Optical Pulse-Phased Observations of Faint Pulsars with a Phase-Binning CCD Camera

Thesis by
Brian Kern

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2002
(Submitted May 29, 2002)

Abstract

We have constructed a phase-binning CCD camera optimized for optical observations of faint pulsars. The phase-binning CCD camera combines the high quantum efficiency of a CCD with a pulse-phased time resolution capable of observing pulsars as fast as 10 ms, with no read noise penalty. The phase-binning CCD can also operate as a two-channel imaging polarimeter, obtaining pulse-phased linear photopolarimetric observations.

We have used this phase-binning CCD to make the first measurements of optical pulsations from an anomalous X-ray pulsar. We measured the optical pulse profile of 4U 0142+61, finding a pulsed fraction of 27%, many times larger than the pulsed fraction in X-rays. From this observation, we concluded that 4U 0142+61 must be a magnetar, an ultramagnetized neutron star ($B > 10^{14}$ G). The optical pulse is double-peaked, similar to the soft X-ray pulse profile.

We also used the phase-binning CCD to obtain the photometric and polarimetric pulse profiles of PSR B0656+14, a middle-aged isolated rotation-powered pulsar. The optical pulse profile we measured significantly disagrees with the low signal-to-noise profile previously published for this pulsar. Our results show that the optical flux is entirely pulsed, with optical peaks at phases 0.2 and 0.8 with respect to the radio peak, and a bridge of emission between the peaks. The significance of the detection of pulsed polarized flux is low, but the position angles match the extrapolation of the radio polarization profile. The optical data, both photometric and polarimetric, are consistent with the polar cap model of pulsar magnetospheric emission. The fit of the optical data with the competing emission model, the outer gap model, has not yet been determined.

We have developed a number of statistical tools, both to estimate the errors in our measurements and to identify systematic errors present in the pulse profiles. The statistical tools, when applied to the data presented here, show that the systematic errors are negligible, bolstering the claims of significance of these results.

Acknowledgements

An instrumentation project, by necessity, involves a number of people who contribute labor, insight, and support, in significant ways that often go unthanked and unrewarded. The person most deserving of thanks and rewards is Steve Kaye. As an engineer, Steve was not only responsible for a significant effort in designing the electronics that make this instrument get up and go, but also for providing an extra hand when aligning optics, an extra eye when analyzing noise, whether it be on an oscilloscope or, just as likely, generated by the data analysis software itself, an extra ear when trying to sort through complex problems, and an indiscriminating palate when trying to feed ourselves at 4 a.m. the night before packing up to drive to Palomar. For that matter, he also provided his own hallucinations during long, cold winter observing runs. Steve has always been willing to fill whatever roles are in need of filling, and has always stepped up to the task when asked to do so. Perhaps even more importantly, he has been as agreeable and entertaining a friend as I could hope for, whose value cannot be underestimated when trying to keep up the battle after years upon years of failure to get data. Steve deserves everything I can thank him for and more.

My advisor, Chris Martin, has provided the idea, the motivation, the support, and the encouragement to make this project happen, from before the time I even arrived on the scene. He has never lost faith in my abilities and my good sense, even long after I have lost that same faith. With a small number of words, he has been able to maintain my self-confidence and my desire to see this project through. From its inception (as a three-month side project) through the excitement of finally getting real data, Chris has provided the calm dedication and conviction that this would eventually bear fruit. But most importantly, when I am frustrated that things aren't coming along as I had hoped, Chris shares his understanding that these things are hard; in fact, if they weren't hard, somebody else would have already done it.

Many of the people that provide the support needed to keep this project moving never receive any feedback at all. The machinists in the Physics Shop deserve an enormous amount of credit and thanks for putting up with me when I want something by, "hmm . . . How about tomorrow?" Just as important, the friendly and supportive environment provided when I'm working in the shop along with them is another way this project kept moving along. Rick, Terry, Richard, Armando, Ramón, Isaac, Stan, José, Phil, and Nathaniel deserve far more credit than they generally receive for Caltech continuing to plod along in its pursuit of new goodies. Likewise, without the Palomar Observatory staff making that big telescope swing when we ask it to, without dropping our instrument, none of this would mean anything.

Through my embarrassingly long stay as a grad student at Caltech, a number of people have donated pieces of themselves to my experience here. My contemporaries, Brad Behr, Ben Oppenheimer, Roy Gal, Marc Kushner, and Andrew Baker, were good people to share offices and years with as we assimilated into Caltech life. Dave Schiminovich, Ryan McLean, John Klemic, Mark Mechtlen, and Rev. Jim have been daily components of my more recent years' experience, and have all added something extra to my understanding or enjoyment, or both. Dave deserves enormous thanks for donating a week of sleepless days and nights in our largely unsuccessful Chile campaign. I have Bob Rutledge to thank for teaching me everything there is to know about X-ray astronomy (or, at least, everything *I* know about X-ray astronomy).

Professional associations aside, my deep personal thanks go to to my father, who armed me with a solid education and instilled in me the self-confidence and the drive to find my own way in this world. His unswerving confidence in me was the extra push I needed at several points in my wanderings. More importantly, he taught me that it doesn't mean much unless you enjoy life along the way. I wish there were a way to share my feelings with my mother, and for her to know I've followed in her footsteps.

I also single out my brother for credit, for the odd reason that I always want to tell him interesting things about what I'm working on. Is that brotherly or just competitive?

Despite the magnitude of thanks I owe to everyone else, the one person to whom I owe the biggest debt is Genna, my fiancée. Every single night of every year, whether I was frustrated by mechanical or software problems, returning from observing through clouds or 5 arcsec seeing, Genna welcomed me home and warmed my heart. It is not possible for me to imagine making it through every day without the uncomplicated love she shows me.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Optical Pulsar Observations	1
1.2 Scientific Framework	3
1.3 History of This Effort	5
1.4 Organization of the Thesis	7
2 Instrumentation	8
2.1 Principles of Operation	8
2.1.1 Phase Binning	8
2.1.2 Timing	9
2.1.3 Systematic errors	11
2.1.4 Polarization capabilities	13
2.1.5 Nonperiodic short-exposure imaging	14
2.2 Optical/Mechanical	16
2.2.1 Overview	16
2.2.2 Wide-field optical layout	17
2.2.3 Dual-channel polarimeter optical layout	18
2.2.4 Thermal	20
2.3 Electronics	20
2.3.1 Overview	20
2.3.2 GPS System	20
2.3.3 DAQ System	22
2.3.4 CCD Logic System	22
2.3.5 Camera Head	23
2.3.6 CCD	25
2.4 Software	25
2.5 Performance	28

3	Data Analysis	31
3.1	Bias and Flat-Field Correction	31
3.2	Error Estimation and Statistics of Correlated Image Wander	38
3.3	Statistics of Polarization	48
3.4	Statistics of Pulsed Fraction	53
4	4U 0142+61	57
4.1	Nature paper on 4U 0142+61	57
5	PSR B0656+14	64
5.1	Astrophysical Journal submission on PSR B0656+14	64
5.1.1	Introduction	65
5.1.2	Observations	66
5.1.3	Data analysis	67
5.1.4	Discussion	70
5.2	Emission Models	76
6	Low-Mass X-ray Binaries and Cataclysmic Variables	80
7	Conclusions	82
	Appendix A	84
	Bibliography	101

List of Figures

2.1	CCD illumination schematic	9
2.2	Wide-field phase offset and dither schematic	13
2.3	Polarization illumination schematic	14
2.4	Polarimetric position angle, phase offset and dither schematic	15
2.5	Multiple frame transfer mode illumination schematic	16
2.6	Wide-field optical layout	17
2.7	Slit with keyhole.	18
2.8	Polarization optical layout	19
2.9	Electronics systems	21
2.10	CCD logic flow chart.	24
2.11	Control software screen shot	27
2.12	Sensitivity curves	29
2.13	Integration duty cycle	30
3.1	Iterative blank-sky calibration algorithm	35
3.2	Image wander	40
3.3	Probability distribution function of image wander	43
3.4	Scaling correction from pulse residual cross term	44
3.5	Histogram of cumulative probability of errors, assuming χ^2	45
3.6	Histogram of cumulative probability of errors, incorporating image wander	45
3.7	Map of cumulative probability of pulse residuals	46
3.8	Ratio of measured errors to predicted errors	47
3.9	Bias in polarization estimators	50
3.10	Confidence bands for polarized flux	51
3.11	Confidence bands and estimators for pulsed fraction	55
3.12	Bias in pulsed fraction estimators	56
4.1	Time-averaged optical image of field around 4U 0142+61	62
4.2	Map of pulsation amplitudes	62
4.3	Optical and X-ray pulse profiles of 4U 0142+61	63
5.1	Average intensity image, eliminating all phase and polarization information	73
5.2	Phase-binned total intensity images of PSR B0656+14	73

5.3	Total-intensity pulse profile of PSR B0656+14	74
5.4	Linearly polarized flux	74
5.5	Linear polarization position angles	75
5.6	Rotating Vector Model constraints	77
7.1	Skymap of outward-going emission for $\alpha=10^\circ$	85
7.2	Skymap of outward-going emission for $\alpha=20^\circ$	86
7.3	Skymap of outward-going emission for $\alpha=30^\circ$	87
7.4	Skymap of outward-going emission for $\alpha=40^\circ$	88
7.5	Skymap of outward-going emission for $\alpha=50^\circ$	89
7.6	Skymap of outward-going emission for $\alpha=60^\circ$	90
7.7	Skymap of outward-going emission for $\alpha=70^\circ$	91
7.8	Skymap of outward-going emission for $\alpha=80^\circ$	92
7.9	Skymap of inward-going emission for $\alpha=10^\circ$	93
7.10	Skymap of inward-going emission for $\alpha=20^\circ$	94
7.11	Skymap of inward-going emission for $\alpha=30^\circ$	95
7.12	Skymap of inward-going emission for $\alpha=40^\circ$	96
7.13	Skymap of inward-going emission for $\alpha=50^\circ$	97
7.14	Skymap of inward-going emission for $\alpha=60^\circ$	98
7.15	Skymap of inward-going emission for $\alpha=70^\circ$	99
7.16	Skymap of inward-going emission for $\alpha=80^\circ$	100

Chapter 1

Introduction

1.1 Optical Pulsar Observations

The history of optical pulsar observations began very shortly after the discovery of pulsars in late 1967. Searches for optical pulsations began as early as May 1968, just three months after the announcement of the discovery of the first pulsar, PSR B1919+21. The Crab Pulsar, discovered by radio astronomers in November 1968 (Staelin and Reifenstein, 1968), was seen to pulsate in optical light in January 1969 (Cocke *et al.*, 1969), by three groups within a matter of days.

The subsequent development of the field of optical pulsar observations progressed much more slowly. It would be another eight years until optical pulsations were detected from the Vela Pulsar (Wallace, 1977), and yet another eight years until pulsations were detected in the LMC Pulsar, PSR B0540-69 (Middleditch and Pennypacker, 1985).

In the 17 years since the detection of pulsations from PSR B0540-69, there have been reports of optical pulsations in three other isolated pulsars, SN 1987A (Kristian *et al.*, 1989), Geminga (Shearer *et al.*, 1998) and PSR B0656+14 (Shearer *et al.*, 1997). The reports of pulsations from Geminga and PSR B0656+14 are at low levels of significance, and as we show in data presented in this work, our results for PSR B0656+14, at a high signal-to-noise ratio, contradict the Shearer *et al.* (1997) results. This disagreement must cast some doubt on the Geminga results, by the same group, on the same observing run, and with the same instrument as the PSR B0656+14 results. There has been a concerted effort over the last 15 years to monitor the remnant of SN 1987A, in the hope that a pulsar may appear at optical wavelengths. There have been reports (Kristian *et al.*, 1989) and retractions (Kristian, 1991) and further reports (Middleditch *et al.*, 2000) of pulsations from SN 1987A, none of which can be claimed to have been confirmed (not even the retraction, in a manner of speaking). The difficulties associated with these observations highlight the need for careful error analysis and treatment of systematic errors.

The reason optical observations of pulsars are so difficult is by no means unique to this branch of astronomy—it is a lack of photons. In Table 1.1, we list the optical luminosities of all pulsars that have been detected at optical wavelengths. The difficulty in making the transition from a time-averaged optical detection to detecting optical pulsations lies both in the necessity to make multiple observations (a separate observation for each time resolution element) and in the detector technology required. With the occasional exception of the Crab Pulsar, all of the previous

optical pulsation detections were made using photomultiplier tubes (PMTs) or (for Geminga and PSR B0656+14) MAMAs (an imaging analogue of PMTs), all of which are limited by their photocathode sensitivity. The time-averaged optical detections, without time resolution, have been made using back-illuminated charge-coupled devices (CCDs), which provide near-perfect sensitivity, but operate with no temporal resolution on the timescale of pulsar periods.

Table 1.1: Optical observations of pulsars

Pulsar	P [s]	$\log T$ [yr]	m	
Crab	0.033	3	$V = 16.5$	Opt pulse
PSR B0540-69	0.050	3.2	$V = 22.4$	Opt pulse
Vela	0.089	4.1	$V = 23.7$	Opt pulse
PSR B0656+14	0.385	5.0	$V = 25.0$	This work
Geminga	0.237	5.4	$V = 25.4$? pulse
PSR B1055-52	0.197	5.7	$U = 24.9$	
PSR B1929+10	0.227	6.5	$U = 25.7$	
PSR B0950+08	0.253	7.2	$m_{130} = 27$	
4U 0142+61	8.7	4	$R = 25.0$	AXP, this work
1E 2259+586	7.0	4	$K \approx 21.7?$	AXP
RX J0720.4-3125	8.4		$R = 26.9$	INS
PSR B1509-58	0.151		$V \geq 22.0??$	wrong ID?

The heart of this thesis is an instrument which combines the sensitivity and imaging capabilities of CCDs with time resolution on the order of milliseconds, extending the capacity to observe pulsations in isolated pulsars to much fainter limits than previously attainable. Aware of the hidden perils of systematic issues when dealing with pulsations in faint objects, we have tailored the operation of this instrument to minimize the systematic errors introduced by the hardware, and have designed data reduction and statistical analysis routines to ensure that the resulting data does not contain artifacts that could compromise any of our findings. It has been extremely gratifying to see that the level of errors in the final stages of our data analysis is entirely consistent with fundamental error limits.

Our pulsar camera’s major achievements have been the measurement of optical pulsations from the Anomalous X-ray Pulsar (AXP) 4U 0142+61, and from the radio pulsar PSR B0656+14, both using the Palomar 200-inch telescope. Our findings on 4U 0142+61 are particularly exciting, because AXPs are relatively recent discoveries, and our findings have a significant impact on the fundamental understanding of this class of pulsar. Our measurements of PSR B0656+14 are significant in that they are the first concrete optical pulsations seen from a middle-aged ($> 10^5$ years) pulsar. Our observations contradict the previous measurement of pulsations from PSR B0656+14, but what is most exciting about these observations is that we have obtained pulse-phased linear polarization measurements on PSR B0656+14. While the polarization results are of debatable significance, this

previews the insight that will be available using high-quality optical polarization data, which should be readily observable with 10 m class telescopes (and, needless to say, our instrument). By opening the door to pulse-phased polarization observations, we should be able to explore completely new territory in the field of high-energy pulsar emission mechanisms.

1.2 Scientific Framework

Since their discovery 35 years ago, pulsars have been objects of intense interest, often because their unusual nature and environment allow theorists to explore exotic phenomena that are unavailable anywhere else in the universe. The solid-state properties of neutron star interiors and crusts, the properties of their magnetospheres, and radiation mechanisms in these strong magnetic fields have no other laboratories in the universe. The timing accuracy allowed by the stable rotation of pulsars has provided the opportunity to test general relativity in unique ways, including the first concrete evidence for gravitational radiation (worth a Nobel prize), and has provided the detection of the first extrasolar planet (now the lonely outcast from the bustling crowd of extrasolar planets around ordinary stars).

Studies of emission mechanisms in isolated, rotation-powered pulsars have generally been split into radio and high-energy (infrared through gamma-ray) regimes, because (with the exception of the Crab) it appears that the radio emission is not closely connected to the high-energy emission, except through the gross energetics of the pulsar spindown. Due to the superior relative quality of the radio observations, the radio emission mechanism has been explored more thoroughly than the high-energy emission mechanisms. While there are still a number of debates concerning the nature of the radio emission mechanism, a general consensus holds that the location of the emission region is relatively close to the surface of the neutron star, some 10–100 NS radii high (Kijak and Gil, 1997), as contrasted with the light cylinder, which is another factor of ~ 100 in radius. This proximity to the NS surface simplifies matters somewhat, permitting theories to neglect a number of effects (relativistic aberration, light travel time across emission regions, *etc.*). One consequence of this has been the success of a remarkably simple model, the Rotating Vector Model (Radhakrishnan and Cooke, 1969; Everett and Weisberg, 2001), in estimating the geometrical configuration of the rotation axis, magnetic axis, and line-of-sight to the observer from the change of polarization position angle across the radio pulse.

The high-energy emission mechanisms of isolated, rotation-powered pulsars have not been examined as comprehensively as the radio mechanisms. The soft X-rays have a distinct mechanism from the rest of the infrared through gamma-ray spectrum, in that the surface of the neutron star is sometimes visible in soft X-rays, which are also modulated by the rotation of the NS. The general belief is that the rest of the high-energy spectrum is generated in a cascade process, which begins

when charged particles accelerating through a huge potential emit curvature gamma rays as they follow the magnetic field lines, and the gamma rays then spawn a number of other processes (pair production, inverse compton scattering), eventually resulting in the observed high-energy emission. There remains a fundamental question regarding the location of this accelerating region and the torrent of energetic processes that result. The competing theories split into two broad categories, the polar cap models, with the acceleration near the NS surface (at the magnetic poles) (Baring, 2001), and the outer gap models (Romani and Yadigaroglu, 1995), with acceleration farther out in the magnetosphere (at significant fractions of the light cylinder radius). There are also sprinklings of other models in existence, including those invoking plasma instabilities near the light cylinder as the emission mechanism, possibly producing both the radio and high-energy emission (Lyutikov *et al.*, 1999; Machabeli *et al.*, 2000).

Optical observations of isolated, rotation-powered pulsars provide a means of testing the well-developed models of high-energy emission mechanisms. Simple questions, such as the number of peaks in the optical light curve, separation between peaks, separation between optical and radio or optical and gamma-ray peaks, flux ratio in peaks, *etc.* can go a long way in determining the validity of the underlying models. As a very simple example, in the polar cap model, where the optical emission region is near the NS pole and therefore near the radio emission region, one would expect the rotational phases of optical emission features to be symmetric about the radio emission, as there are no relativistic or light travel time factors that would introduce a phase offset. In fact, the two peaks in the optical light curve from the Vela Pulsar are not symmetric about the radio pulse, casting doubt on the polar cap model for the optical emission from the Vela Pulsar.

Observations of optical pulsations can act as a proxy for the high-energy emission in pulsars where the emission is unavailable at other frequencies. This is particularly true in PSR B0656+14, for which we present here high-quality measurements of optical pulsations, while the gamma-ray measurements of PSR B0656+14 (Ramanamurthy *et al.*, 1996) are of very low significance. An extremely exciting corollary to this is the opportunity for high-quality optical polarization measurements. Radio polarization data, as applied to the Rotating Vector Model, have been the only tool available to determine geometric parameters in most radio pulsars. There is little hope for X-ray or gamma-ray polarimetry to be applied to pulsars in the immediate future, and so the only avenue for high-energy polarization measurements lies with optical polarimetry. The agreement or disagreement between radio and high-energy polarization interpretations could offer very precise metrics of the structure of pulsar magnetospheres.

The current state of AXP emission models is far less advanced than that of radio pulsars. The question of the fundamental nature of the objects must be addressed before continuing to develop theories of the emission from these objects. Following the work of Hulleman *et al.* (2000a), three models remained to explain the AXP 4U 0142+61: a magnetar (neutron star with surface magnetic

field $> 10^{14}$ G) (Duncan and Thompson, 1992), a hot (5×10^5 K) magnetized (10^8 G) white dwarf (the result of a merger of two white dwarfs) (Paczynski, 1990), and an accreting isolated neutron star (with an ordinary magnetic field) (Marsden *et al.*, 2001). Our data, using the detection of optical pulsations from 4U 0142+61, was inconsistent with the accretion and white dwarf models, leaving (by the process of elimination) only the magnetar model as a viable explanation of what is known about 4U 0142+61. Investigations into the nature of the emission mechanisms of AXPs can now progress, with a better understanding of the physical nature underlying these objects.

1.3 History of This Effort

This project began in the winter of 1995, with the development of a prototype logic system capable of performing the frame transfers that define our technique for time resolution. This first prototype had one observing run at the Palomar 200-inch telescope in January 1996, where we were able to observe the Crab Pulsar, but light clouds prevented any sensitive measurements of Geminga (our primary target).

With the experience gained by taking the instrument to the telescope, and a \$30,000 grant from the NSF, a complete redesign was undertaken to create a robust instrument. This new pulsar camera went back to Palomar in December 1998, but again clouds prevented any observation of Geminga. At that time, the pulsar camera had no polarimetric capabilities, and was a narrow-field instrument only. During the next year, two-channel imaging polarization capabilities were introduced (with the addition of another \$30,000 grant from the NSF). The instrument was also modified to allow an interface with the CTIO Blanco 4 m telescope, and we took it for an observing run there in January 2000. Interface problems with the telescope and electronic noise pickup prevented a high-sensitivity measurement of Vela, but we were able to see pulsations, although at a lesser significance than those published some 20 years prior. We had hoped to obtain polarization data on Vela, but the data was not sufficient for any significant measurements. This experience was followed by another Palomar run in February 2000, where we had a break from our experiences with clouds, only to find 5 arcsec seeing conditions.

It was not until December 2000 that we obtained high-quality data on anything other than the Crab. We observed PSR B0656+14 for two nights under good conditions (1.3 arcsec seeing, nonphotometric conditions), using the camera in its polarimetric mode. The results of this observing run are part of a paper that is included as Chapter 5.

The pulsar camera was reconfigured in the first half of 2001 (with a third, and final, NSF grant of \$15,000) to allow for an alternate mode of operation, the multiple frame transfer mode, which allows bursts of imaging with very short exposures (\sim milliseconds), alternating with read cycles (while not exposing). This mode trades the polarization capabilities for a wider field of view, necessary

to include comparison stars with which to perform differential photometry. The camera was used in this mode for a Palomar observing run in July 2001, with Guillaume Dubus as the principal investigator, observing low-mass X-ray binaries. The data from this observing run has been reduced but not examined. It is given a very brief overview in Chapter 6.

The pulsar camera returned once again to Palomar in November 2001, observing 4U 0142+61 for two nights, in wide-field mode (obtaining no polarization information). The observing conditions were similar to those of our December 2000 run, allowing a solid measurement of pulsations from 4U 0142+61. The results of this observing run are presented in Chapter 4.

As is true of any instrumentation thesis, the bulk of my labor seems to have gone into mechanical drawings, machining, computer software programming, electronic logic design, optical design and alignment, and above all else, debugging. Because of the particular systematic errors associated with our periodic frame transfers (charge traps), we wanted to use as small a CCD as possible (for better charge transfer efficiency), which runs counter to the interests of nearly every other astronomer I have met. As such, we built a system literally from scratch, starting, with the prototype, from a surplus electronics classroom CCD (front-illuminated), a surplus thermoelectric cooler, and discrete logic components wired into a breadboard.

With the exception of the logic electronics and a first iteration of the autoguider optics, I was solely responsible for every component of the construction and operation of this camera. The bulk of my early labor went into creating software in a Windows environment that would operate with real time accuracy, to control the timing of our observations to microsecond precision. As much as I expected to be able to turn all of the machining over to the Physics Shop after making of all of the mechanical drawings, I ended up doing the majority of the machining myself, generally because of time constraints and delays in finishing the mechanical drawings. It is always surprising to look back and recount the weeks of work that go into each and every component of a successful system.

I have also been solely responsible for developing the data analysis routines described in Chapter 3. The broad features of the calibration routine were fleshed out early in the project, as the first cloudy data began to roll in. It was not until the December 2000 observing run, where we obtained reliable polarimetric data, that the polarization statistics were examined, and it was not until after the November 2001 observing run that the full image wander and pulsed fraction statistics were examined.

This instrument is relatively specialized, and is optimized for use on sources that are periodic on timescales of milliseconds to tens of seconds, *i.e.*, pulsars. It does offer capabilities for general high-time-resolution studies of nonperiodic sources, and has been used on one occasion for just this. The true capabilities of this instrument will be realized only by access to the Keck telescopes, or other 8–10 m class telescopes with excellent seeing. With two Palomar observing runs of pulsars living entirely up to their expectations, the time is ripe to graduate to a larger telescope. However,

it is not likely to appeal to a wide audience of observers, and as such, there has not been a genuine attempt to train other observers in its use, or develop a document for the exclusive purpose of instructing future observers on its use.

1.4 Organization of the Thesis

This thesis is divided by the logical distinctions between the directions of my efforts as a graduate student. The operation of the instrument itself will be discussed first, followed by a description of the data analysis and statistical evaluation routines. The data from our two successful pulsar observing runs are presented separately, as they differ in their scientific context, and have been written as two papers, one to *Nature* (to be published 2002 May 30), the other in preparation for submission to the *Astrophysical Journal*. There is a brief mention of the use of this instrument in its multiple frame transfer mode (for nonperiodic sources) in Chapter 6. Finally, a few words will conclude the thesis. An appendix gives computational details that accompany a discussion of the PSR B0656+14 paper.

The highlights of the thesis, if one wanted an impression of this thesis in the shortest possible time, begin with the images of PSR B0656+14 in Fig. 5.2, the clearest example of what we are interested in obtaining. The other big scientific result comes in Fig. 4.3, where we measure optical pulsations from an anomalous X-ray pulsar for the first time. From a data analysis standpoint, the real achievements shine through in Fig. 3.6–3.8, where we show that the measured errors are perfectly well explained by the fundamental error limits (Poisson background noise) and that the pulsations we measure are clearly not instrumental in origin. These figures would compose the five-minute tour of my work as a graduate student.

Chapter 2

Instrumentation

This chapter describes the operation of the phase-binning CCD camera. Apart from the principles of operation (describing the three modes of operation of this camera) and the performance sections, much of this information is very detailed and is of interest only for a complete understanding of the physical developments involved in making this camera do what it does.

2.1 Principles of Operation

2.1.1 Phase Binning

The phase-binning CCD is the most sensitive instrument to date for observing periodic sources, such as a pulsar, with periods between 10 ms and 30 s. The phase-binning CCD operates as a periodic frame transfer device, as shown in Fig. 2.1. The field of view of the CCD is masked off with a slit, so that most of the CCD is shielded from light. The inactive regions of the CCD are used as storage for multiple images of the slit, taken at different times. Data is commonly taken in ten phase bins ($n = 10$), resulting in ten images residing on the chip. A single long exposure then consists of alternating periods of integration, when one image is being exposed and nine images are in storage, and frame transfer, when the most recently exposed image is transferred into storage and the stored images are shifted so that one is brought into the region illuminated by the slit. During frame transfer, no mechanical components move, only the accumulated charge in the CCD moves. Fig. 2.1 shows the CCD illumination schematic, in opposing extremes of the transfer cycles.

The transfer of images is repeated in a cyclic pattern, with five transfers shifting the images upward on the CCD and five shifting downward, so that after beginning by integrating Image 0, then repeating the integration and transfer steps ten times, Image 0 is being reilluminated. The shutter which defines the entire exposure is opened at the beginning of the exposure, and left open during all of the integrations and transfers. No pixels are read out of the CCD at any time during the exposure. After an arbitrary number of cycles of integration and transfer, perhaps illuminating each of the ten images 100–1000 times, the shutter is closed, and the entire CCD (containing the ten images) is read out. An individual exposure may last 120 s, which means a full observing run may comprise several hundred exposures.

The output of a single exposure is an array of n phase-binned images (where n is generally 10). With no moving parts, objects appearing in the images are recorded in the same physical pixels

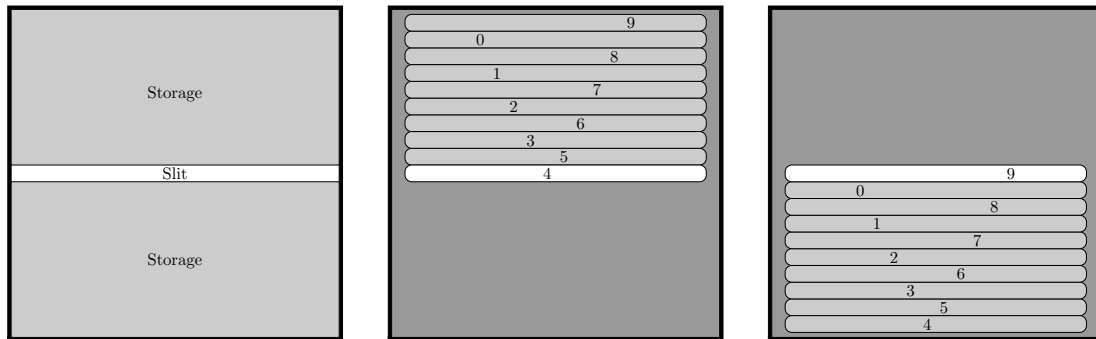


Figure 2.1: CCD illumination schematic. *Left*: Division of CCD into the region illuminated by the slit and the storage regions. *Middle*: Location of images during exposure of Image 4. The images are numbered temporally from 0 to 9, with labels horizontally offset purely for illustrative purposes; there is no horizontal offset between the images. Note that the images 0–4 are interleaved with the images 5–9, so that temporally consecutive images are either one or two distant. *Right*: Location of images during exposure of Image 9.

(in the region of the CCD illuminated by the slit) for all images, so that no variations in optical throughput or Quantum Efficiency (QE) will be imprinted on the differential light curve of objects in the images. The ability to reilluminate images hundreds or thousands of times allows the signal (and background) to grow before being read out, so that the read noise ($5\text{--}8\text{ e}^-/\text{pixel rms}$) is small compared to the Poisson noise of the signal (for bright sources) or the sky background. It is this combination of imaging, the sensitivity of the CCD, and the elimination of read noise (in a relative sense) that underlies the power of this technique.

The CCD operates in one of three different modes, two of which are phase-binning modes (as described here) and one of which is a nonperiodic frame transfer mode. The two phase-binning modes are the wide-field phase-binning and two-channel polarimetric modes, which differ in the width of the images (wider for wide-field) and the measurement of polarization information. The third mode, which we call a multiple frame transfer mode, is a wide-field mode that does not re-illuminate any exposures, resulting in a series of short-exposure images. If not explicitly stated, the assumption throughout this thesis is that one of the phase-binning modes are being used. These are the modes that are unique to this instrument, while there are a handful of frame transfer instruments in general astronomical use already.

2.1.2 Timing

The timing of the integration and transfers is controlled to match the period of the observed pulsar, when phase binning, or to provide an absolute time reference for multiple frame transfers. In a simple implementation, the timing of the phase binning (with $n = 10$) would ensure that Image 0 is always illuminated during pulse phase 0.0–0.1, Image 1 during phase 0.1–0.2, and so on, with

Image 9 illuminated from phase 0.9–1.0 (see Section 2.1.3). Because the phase binning is performed in real time, the period of the pulsar must be known *a priori*, from radio or X-ray (or gamma-ray) observations. The absolute timing of the sequence is controlled by a GPS receiver and time code generator, accurate to 1 μs . Barycenter-geocenter and geocenter-topocentric (*i.e.*, observatory) corrections must be computed beforehand so that they can be applied to the start and stop times of an individual exposure. General relativistic effects (both Shapiro delay and varying redshift due to the earth’s elliptical orbit) are small and vary slowly, and so are not important here. A simple description of these corrections is given in Lyne and Graham-Smith (1998).

There are two levels of timing accuracy required for the phase-binning CCD, the intraexposure accuracy and the interexposure accuracy. The intraexposure accuracy, the timing of integration and transfer cycles during an exposure, must be high enough to ensure that smearing of the pulse phases is small. With a pulsar period P_{obs} , as observed on the surface of the earth, the accumulated timing error during a single exposure should be a small fraction of P_{obs}/n (the duration of one phase bin). Quoted in terms of the fractional error,

$$\delta P_{\text{obs}}/P_{\text{obs}} \ll (nN_{\text{pulse}})^{-1}, \quad (2.1)$$

where δP_{obs} is the error in the observed pulsar period and N_{pulse} is the number of full pulses in an exposure, *i.e.*, $N_{\text{pulse}} = T_{\text{exp}}/P_{\text{obs}}$, where T_{exp} is the single-exposure time. Taking PSR B0656+14 as an example, with $P_{\text{obs}} \sim 385$ ms, $n = 10$, and $T_{\text{exp}} \sim 120$ s, $\delta P_{\text{obs}}/P_{\text{obs}} \ll 3 \times 10^{-4}$. For the Crab pulsar, with $P = 33$ ms, the limiting fractional error becomes 3×10^{-5} for a 120 s exposure. If the errors are larger than this, then the time series recorded in a single exposure is convolved with the error distribution, resulting in a loss of temporal resolution.

The required interexposure accuracy, the accuracy with which successive exposures must be timed, is significantly stricter than the intraexposure accuracy, but can be largely corrected *a posteriori*. Similar to the constraint for the intraexposure accuracy,

$$\delta P_{\text{obs}}/P_{\text{obs}} \ll (nN_{\text{tot}})^{-1}, \quad (2.2)$$

where N_{tot} is the total number of pulses elapsed during the entire observing run. For a three-night observation, for instance, N_{tot} may be 5×10^5 for PSR B0656+14, or 5×10^6 for the Crab, resulting in a fractional error requirement of $\delta P_{\text{obs}}/P_{\text{obs}} \ll 2 \times 10^{-7}$ or 2×10^{-8} . Pulsar periods, as measured at other wavelengths, are generally accurate to levels well beyond this requirement, as are the barycenter-geocenter and geocenter-topocentric corrections. In the event that the actual accuracy lies between the intraexposure and interexposure accuracy limits, the individual exposures properly preserve the timing information, which can be retained by properly coadding successive exposures (*i.e.*, interexposure corrections can be applied in software).

2.1.3 Systematic errors

Several steps are taken to reduce the effects of potential systematic errors in the time series of objects imaged by the phase-binning CCD. The nature of the illumination scheme, with no moving parts, ensures that QE and throughput variations will not be imprinted on the measured (differential) pulse profile, unless the throughput variations occur synchronously with the pulse period. In addition, the imaging capabilities of the instrument allow a number of “control” objects to be recorded in each image. By examining the pulse profile of every star in the field of view, one can ensure that there are no instrumental effects that mimic pulsations. This imaging capability provides the surest confirmation of the significance of measured pulsations, by showing that the pulsations happen nowhere else.

While each image is observed by illuminating the same pixels on the CCD, eliminating pixel-to-pixel QE effects, the same cannot be said for the frame transfer paths of each image. For example, in Fig. 2.1, Image 4 is shuttled between the center of the CCD (illuminated by the slit) and the bottom of the CCD, while Image 9 moves between the center and the top of the CCD. Charge Transfer Efficiency (CTE) differences across the CCD will mimic a time-varying signal. Because the number of parallel shifts in an individual exposure is large (on the order of 10^5 – 10^6), CTE that differs from unity by more than 10^{-7} will have a large effect on the resulting pulse profiles. In addition, localized traps ($\leq 1 e^-$) will redistribute large amounts of charge between neighboring pixels across the CCD (Janesick, 2001). While these two effects have different effects and causes, they will generally be discussed as sources of differential charge redistribution.

An additional effect of the different frame transfer paths for the images is that the time spent in integration and transfer differs slightly for different images. As shown in Fig. 2.1, the on-chip locations of the images are interleaved, so that temporally consecutive images are either next to each other or separated by only one other image. The time to transfer from illuminating Image 9 to illuminating Image 0, or from Image 4 to Image 5, is half that of the transfer between any other temporally consecutive images.

In general, the integration time for a given image is defined by the pulsar period, the number of phase bins, and the transfer time,

$$t_{\text{trans},i} + t_{\text{int},i} = P_{\text{obs}}/n \quad (2.3)$$

for each $i = 0, 1, \dots, n-1$, where $t_{\text{trans},i}$ is the transfer time and $t_{\text{int},i}$ is the integration time for Image i . The parallel shift time is $42 \mu\text{s}/\text{line}$, and a typical image height is 26 pixels, so the transfer time is either 1.1 ms (for Image 9–0 or Image 4–5) or 2.2 ms. For PSR B0656+14 this is small compared to $P/10 = 38.5$ ms, giving a duty cycle of 97% or 94%. For the Crab, however, $P/10 = 3.3$ ms, giving a duty cycle of 67% or 33%, creating a large differential signal in the pulse profile, with Images 0 and

5 receiving twice the integration time as the other images. Because the shutter is open throughout the exposure, the time spent in transfer, in addition to reducing the integration time, results in some vertical smearing of the images.

The opening and closing of the shutter introduces an extra nonuniformity in the total integration time of images in an exposure. The shutter open and close time is approximately 10 ms, but because the shutter is an iris style shutter, the effects are different at different places in the field of view. The shutter opening time results in a reduced total integration time for Image 0 (and possibly other images if P is small), and the shutter closing time results in an increased total integration time for Image 9 (with a small amount of smearing as readout begins).

The most reliable way to mitigate the effects of charge redistribution, transfer time and shutter effects is to ensure that, in the final analysis, these effects are imprinted equally on each phase bin. While the paths traveled by each phase-binned image are different in a single exposure, by altering the start time of successive exposures with respect to the pulsar’s pulse phase, the mapping between images and pulse phase can be altered in such a way that in a series of ten exposures, each pulse-phase interval appears in every image exactly once. To be specific, we define a phase offset, between 0.0 and 1.0 that is applied to the start time of an exposure. The actual start of the exposure is delayed by the phase offset (multiplied by the period, P_{obs}). For a given exposure, Image 0 begins at a pulse phase equal to the phase offset of the start of the exposure, with this offset persisting in all of the successive images. In other words,

$$\varphi_{\text{lo}} = i/n + \varphi_{\text{off}}, \quad (2.4)$$

where φ_{lo} is the lower limit to the pulse phase in a given phase bin (*i.e.*, $\varphi_{\text{lo}} = 0$ corresponds to pulse phase 0.0–0.1), i is the image index (0–9), n is the number of phase bins (10), and φ_{off} is the phase offset. Throughout this discussion, we will express pulse phase in units of cycles, *i.e.*, as fractions of 2π radians. By incrementing φ_{off} by 0.1 in successive exposures, after ten exposures, pulse phase 0.0–0.1 will appear in each of Images 0–9 exactly once, as does pulse phase 0.1–0.2, *etc.* This ensures that any effects specific to individual images appear identically in each phase bin. To the extent that the data analysis yields a differential measurement, charge redistribution, transfer time, and shutter effects have no effect on the measured pulse profiles.

Charge redistribution effects pose difficult flat-fielding problems, relevant when examining variations in the sky dimensions (rather than the pulse phase dimension). The effects of the localized traps are slightly dependent on the charge contained in a pixel, so for a given point in an image, following the same transfer path through an exposure, the trap effects depend on the incident flux. This means that proper “dark” and flat-field calibration frames can only be obtained under identical illumination conditions as the observations they are intended to calibrate. The only practical solu-




Figure 2.2: Wide-field phase offset and dither schematic. This figure represents the grouping of the sequence of successive exposures. The (x, y) plane represents the CCD itself, which contains n images ($n = 10$ in this case), as represented in Fig. 2.1. Successive images are taken with different phase offsets, which are represented as the depth dimension (into the page) in this figure. After a complete set of n exposures, started with each of all possible values of the phase offset, the telescope dither is changed, and the process is started again. At the end of an observing run, the observer has collected some n_D different dither positions, each composed of n images at different phase offsets.

tion to this is to use a self-calibration technique, dithering the telescope pointing between exposures, in the hope that every point in an image will be alternately exposed to blank sky and to sky plus objects in the field of view. In this fashion, a composite flat-field calibration can be constructed using blank sky wherever it may land on the CCD, to patch together a complete calibration. This is discussed in detail in Section 3.1. To preserve the differential immunity to charge redistribution, transfer time, and shutter effects, the dithering is performed only every n exposures, *i.e.*, after a full cycle of phase offsets. In other words, for a set of n exposures, the phase offset is increased between exposures, then the telescope pointing is changed, and the process is repeated. In a night of observing a single object, 200 exposures may be taken, in 20 different dither positions with 10 phase offsets at each dither position. This is shown in Fig. 2.2.

2.1.4 Polarization capabilities

The second mode of operation of the phase-binning CCD adds a polarimetric capability to the phase-binning imaging performance. With the addition of polarizing beam splitters and an achromatic half-wave plate, as described more fully in Section 2.2.3, the incoming light is split into orthogonal linear polarizations, which land in separate columns on the CCD. This is shown schematically in Fig. 2.3.

The orientation of the fast axis of the half-wave plate determines the linearly polarized position angle (as projected on the sky), so by rotating the half-wave plate between exposures, the position




Figure 2.3: Polarization illumination schematic. The CCD is split into two columns, each measuring orthogonal linear polarizations. The hatch patterns represent polarization angles of -45° and 45° . The arrangement of images in each column is the same as in Fig. 2.1. The two columns observe the same field of view.

angles measured by each of the two columns rotate as well. In a series of four consecutive exposures, with the position angles rotated 45° between each exposure, each column independently measures the Stokes parameters I , Q and U . Because this is a two-channel polarimeter, each column independently measures the Stokes parameters, but because orthogonal polarizations are measured simultaneously, any throughput variations affect the Q and U measurements oppositely, and so cancel when the two separate measurements are combined. For observations of faint pulsars, where the sky background dominates the measured flux, the timescale for making successive polarization measurements should be as fast as possible, because the sky polarization changes with altitude and azimuth. As such, the observing scheme is to rotate the position angle between each successive exposure (keeping the same phase offset), change the phase offset after each set of four position angles, and dither the telescope pointing after ten sets of four exposures. In other words, the ordering of the variations, from most frequent to least, is the polarization position angle, the phase offset, and then the dither position. This is illustrated in Fig. 2.4.

2.1.5 Nonperiodic short-exposure imaging

The third mode of operation of this camera is a wide-field multiple frame transfer mode, similar to the phase-binning modes described above, but without re-illuminating any images. In this mode, images are accumulated by exposing for short periods of time (5 ms–2 s) and transferred away from the region of the CCD illuminated by the slit, until the CCD is full of images, at which point the shutter is closed and the entire CCD is read out. The transfers are performed in only one direction in this mode, resulting in a series of temporally consecutive images. A schematic diagram is shown in Fig. 2.5

As shown in Fig. 2.5, the slit illuminates a strip of the CCD near the top, rather than across the center, as is done with the periodic frame transfer mode. This maximizes the usable storage area

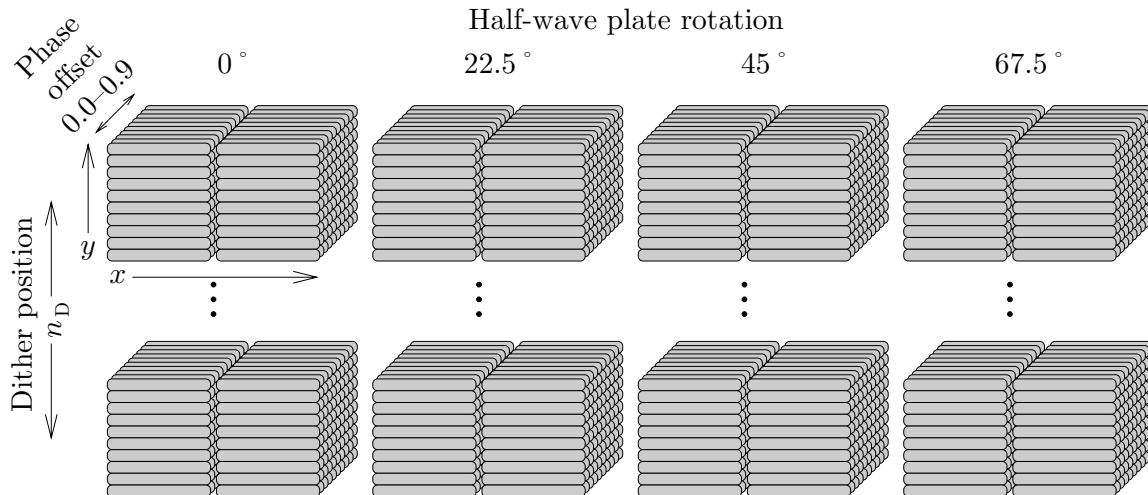


Figure 2.4: Polarimetric position angle, phase offset and dither schematic. The (x, y) plane is the CCD, containing two columns of n images, which compose a single exposure. The next exposure is at a different half-wave plate position, *i.e.*, to the right in this figure. After all four position angles have been measured, the phase offset is changed, *i.e.*, the next exposure into the page is taken. After all position angles and phase offsets have been visited once ($4n$ exposures), the dither position is changed, *i.e.*, downward in this figure. The order of exposures is left-to-right, front-to-back, top-to-bottom.

of the CCD, while minimizing the transfer time. Temporally sequential images are adjacent in this mode.

The transfer time is identical for all images in this mode, eliminating one source of systematic error. There is no robust solution to eliminate the effects of the shutter open and close time, so in practice, the first (or first few, if the integration time is short) and last images are discarded. To maximize efficiency, the first image (or images) can be discarded simply by accumulating more images than will fit on the CCD at one time. As the last images are integrated and transferred, the first images will be shifted off the bottom of the CCD. Since the image would be discarded even if it were read out, this practice ensures that no time is spent reading it out.

The efficiency of this technique is dependent on the read time of the CCD, during which time the shutter is closed. To obtain a measurable signal in a single integration of 5 ms–2 s, the objects must be bright, $V \sim 10$ –16. Few objects in the field of view will be comparably bright—in most cases, there is only one comparison star in the slit. Rather than read out the entire width of the slit, filled with empty space, the CCD logic defines serial windows which are read out, with the rest of the pixels discarded without digitizing. The serial clocking time ($2.4 \mu\text{s}$) is much shorter than the digitization time ($20 \mu\text{s}$, using the “old” data acquisition electronics), so the total read time varies almost linearly with the number of serial pixels read in each line. Standard operation in multiple frame transfer mode involves defining two serial windows (one for the science object and




Figure 2.5: Multiple frame transfer mode illumination schematic. In this mode, the slit illuminates a strip across the top of the CCD rather than the center. Images are numbered temporally, with indices offset horizontally for purely illustrative purposes; there is no horizontal offset of the images. This figure shows integration of the last image before closing the shutter and reading out the CCD, as the storage region is full.

one for a comparison star), reducing the number of digitizations per line from 512 to 80. This serial windowing greatly improves the overall efficiency of the instrument in this mode.

It deserves to be mentioned that the multiple frame transfer mode is not unique to this instrument. The technique has been used on a number of different instruments in different configurations, including the use of LRIS on Keck I in a continuous-read mode, effectively losing one spatial dimension to obtain a time dimension. Other instruments (Dhillon *et al.*, 2001; O’Brien *et al.*, 2001), optimal for this mode of observation, show superior efficiency to our camera. However, our camera is perfectly capable of operating in this mode, and in some isolated situations, can show superior performance to any other instrument.

2.2 Optical/Mechanical

2.2.1 Overview

The phase-binning CCD has been designed for use for the F/16 and F/9 Cassegrain foci at the Palomar 200-inch telescope, and the F/8 Cassegrain focus of the CTIO Blanco 4 m telescope. For all of these arrangements, the optical system is similar, with a shutter, entrance slit, field lens, collimating lens, filter (optional), and a camera lens. The plate scale at the entrance slit, in all cases, is reimaged to 0.4–0.5 arcsec/pixel at the CCD. Some of the basic imaging parameters are listed in Table 2.1. For this table, we assume that the minimum usable image height is 6 pixels. Figures for multiple frame transfers, as opposed to the phase-binning modes, are given in parentheses.

The instrument is constructed on a 24 × 36 inch breadboard, with the optical axis located 5 inches above the plane of the breadboard. This modular configuration has proven extremely convenient, as the optical layout has been reconfigured at several stages of the instrument design.

Table 2.1: Imaging parameters of instrument modes

	Wide-field	Dual-channel Polarimeter
Max. number of images	25 (50)	25 (50)
Image width, W	250 arcsec	50 arcsec
Image height, H	3–13 arcsec	3–13 arcsec
Bandwidth	380–1000 nm	400–700 nm
Throughput	45%	25%




Figure 2.6: Wide-field optical layout.

2.2.2 Wide-field optical layout

The wide-field operating mode, in both phase-binning and multiple frame transfer modes, has been used only at the F/9 Cassegrain focus of the Palomar 200-inch telescope. The optical train consists of the entrance slit, field lens, collimating lens, shutter, and camera lens. A 1-inch diameter filter can be placed at the entrance of the shutter if desired. This is shown in Fig. 2.6.

The entrance slit is a machined slot in a sheet of aluminum, with a black anodized finish. The dimensions of the slit can only be changed by exchanging pre-machined slits, which necessitates a simple re-alignment of the optics. In practice, this has not been done during a night of observing, although if an observing run required observations with different slit dimensions (*e.g.*, 10 phase bins and 20 phase bins with different image heights), this could be done.

The camera lens used is a commercial F/1.3 25 mm compound lens with manual focus and iris. As the actual components in the lens (7 elements) are unknown, we do not attempt to model the performance of the camera lens in our optical design software (Code V), but rather assume it is a perfect imaging system.




Figure 2.7: Slit with keyhole. *Left*: Outline of 230×5 arcsec slit, with 20×70 arcsec keyhole. *Right*: 30 s image of Her X-1, with no frame transfers, stretched to show sky level. The entire CCD is shown in this figure. Compare to Fig. 2.5, which illustrates multiple frame transfers.

Simulations of the optical system (shown in Fig. 2.6) using Code V optical analysis software, including an accurate prescription for the telescope mirror (but no atmosphere), give an rms spot diameter of 8–11 μm (0.17–0.23 arcsec) across the full width of the slit, compared to a pixel size of 24 μm . In practice, the best images recorded at Palomar have FWHM of 1.0 arcsec (≥ 2 pixels), from which we conclude that the imaging properties of the instrument alone do not limit the final resolution.

For the multiple frame transfer mode, we move the CCD so that the slit illuminates a region at the top of the CCD, rather than in the center (see Fig. 2.5). In this mode, all of the optical elements remain in exactly the same place as for the phase-binning mode, but the dewar and CCD (not the camera lens) are moved ~ 250 pixels farther from the breadboard, which, because the CCD is actually inverted in this case, moves the slit image upward. In this way, the optics do not operate off-axis, maintaining the best imaging performance, and the modification to the mechanical layout of the phase-binning mode is simply the addition of spacers to the mounts holding the dewar.

To maximize the exposure duty cycle (relative to the readout time) in multiple frame transfer mode, we have used a relatively small slit, with a 5 arcsec height (10 pixels). This allows us to record 50 images on the CCD before reading out. Using this small of a slit poses some difficulty in acquiring the correct field, so we have used a modified slit with a “keyhole” for field acquisition (in a frame taken with only one image, to avoid overlap), shown in Fig. 2.7. The confusion resulting from the overlap of the keyhole in successive images renders the keyhole width unusable, but it is straightforward to arrange the objects of interest along the slit in such a way that this is not a problem.

2.2.3 Dual-channel polarimeter optical layout

The dual-channel polarimeter mode has been used at both the F/16 Cassegrain focus of the Palomar 200-inch telescope and the F/8 Cassegrain focus of the CTIO Blanco 4 m telescope. The optical train consists of a shutter, entrance slit, field lens, achromatic half-wave plate, collimating lens,




Figure 2.8: Polarization optical layout.

bandpass filter, polarizing cube beam splitters, folding mirrors, and a camera lens. This is shown diagrammatically in Fig. 2.8.

The achromatic half-wave plate is a custom 2-inch optic (from Meadowlark Optics), mounted in a custom rotating mount. The rotation is driven by a small stepper motor, connected by a worm gear to the half-wave plate mount. The mount has an optical sensor and a pinhole, to locate an absolute zero-point of rotation when powered up. The half-wave plate is designed to give a retardation of $180 \pm 2^\circ$ for wavelengths between 410 and 750 nm. Because the performance becomes more unreliable outside this wavelength range, a bandpass filter is required to maintain good performance. In practice, a simple Schott colored glass BG38 filter is used in this mode, giving a bandpass of 400–600 nm.

The polarizing beam splitters are cube beam splitters with a broadband dielectric polarizing coating on the hypotenuses. This coating provides a transmission of 90–95% for *p*-polarization and a reflection of >99% for *s*-polarization, with <0.5% absorption, for wavelengths from 400–700 nm. The discrimination between *s*- and *p*-polarizations is limited by the 5–10% of the *p*-polarization that is reflected by the beam splitters, but because there are two beam splitters in the optical path, this impurity is reduced to the 1% level. These numbers are only accurate for on-axis light rays, however, and near the edge of the field of view, the mixture of polarization states is more pronounced. While the absorption in the beam splitters is very low, some *p*-polarized light leaves through the “inactive” side of the second beam splitter. As this amount is angle-dependent, this introduces some throughput variations across the field of view. These variations are static, and can be calibrated using the normal calibration routine (see Section 3.1). The use of dielectric cube beam splitters for

polarization separation limits the accuracy of this technique to relatively large polarizations, but there is little hope of reaching $\sim 1\text{--}3\%$ polarization accuracy on faint pulsars, even with the 10 m telescopes and a perfectly clean polarization measurements.

The key to creating the two columns of images shown in Fig. 2.3 is the alignment of the folding mirrors. Because the two orthogonal polarization states travel through separate arms of the polarizing assembly, the propagation direction of each polarization state is independently controllable by positioning the fold mirrors. Each of the fold mirrors is 3-axis adjustable, so that not only tip-tilt is controlled, but the pupil image through each arm can be positioned on the camera lens.

2.2.4 Thermal

The CCD is enclosed in an Infrared Labs 5-liter dewar. It maintains the CCD at a temperature of 195 K for 15–20 hours between LN₂ fills, requiring no refilling during a night of observing. The dewar has an offset fill tube, giving it a preferred orientation when it is not perfectly vertical. It can be operated either upright or inverted, with the insertion of a “stinger” in the fill tube for inverted operation. Typically, the dewar is operated in inverted configuration when in wide-field mode and in upright configuration in dual-channel polarimeter mode. The temperature is maintained with a closed-loop heater control, tied to a semiconductor temperature sensor located near the CCD.

2.3 Electronics

2.3.1 Overview

There are multiple electronic systems at work in this instrument, all coordinated by the data acquisition computer. A block diagram of these systems appears in Fig. 2.9. The systems can be divided into three groups, those in the data acquisition computer, those that are part of a NIMBIN crate, and those in the camera head. The data acquisition computer houses the GPS system, which provides an absolute time and frequency reference, and the DAQ (data acquisition) system, which provides the digitization and analog computer interface to the rest of the system. The NIMBIN crate houses the CCD logic system, which combines the CCD control and timing logic, and three simple, self-contained subsystems: the shutter control, heater control, and half-wave plate control.

2.3.2 GPS System

The GPS system consists of three components: a GPS receiver, an IRIG-B time code generator, and a frequency reference and match pulse generator. Palomar Observatory, CTIO, and our laboratory have GPS receivers and IRIG-B time code generators, which are standardized and essentially interchangeable. The IRIG-B time code, an amplitude-modulated 1 kHz signal containing absolute

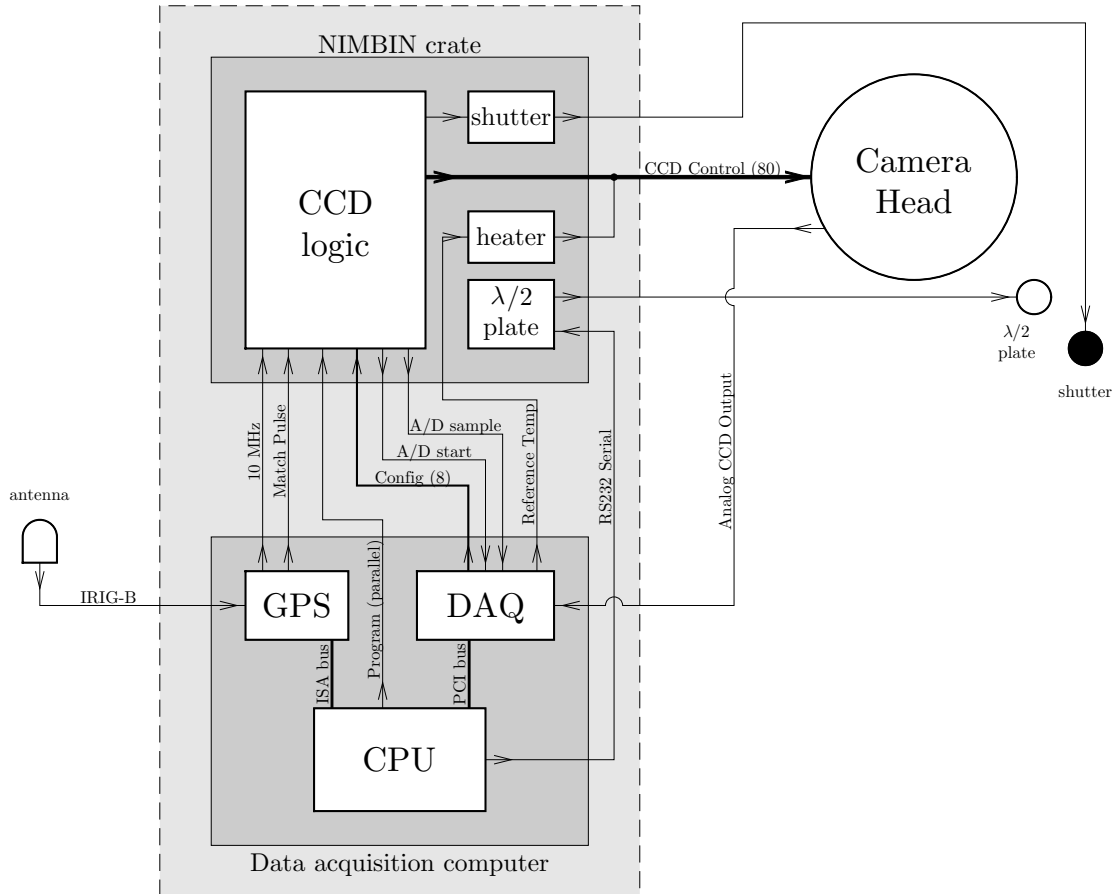


Figure 2.9: Electronics systems. The data acquisition computer and NIMBIN crate are rack-mounted in a single electronics cabinet (denoted by the dashed line). The camera head, half-wave plate, and shutter are separately located on the optical breadboard. An external GPS antenna and receiver deliver an IRIG-B time code to the GPS system. For each electrical connection shown (except the computer buses), the direction of information flow is represented by arrows, and when multiple lines are present, the number of lines are shown in parentheses.

timing information, is input into an ISA-bus computer-based frequency reference and match pulse generator. The frequency reference provides a 10 MHz square-wave pulse, disciplined to the IRIG-B time code. Based on commands from the data acquisition computer, this ISA-bus card will generate a TTL signal at a specified time with absolute accuracy of 1 μ s. The match pulse is the signal to start and stop individual exposures, and the frequency reference is used to control the timing of integrate and transfer states during an exposure.

2.3.3 DAQ System

The DAQ (data acquisition) system is a National Instruments PCI-6052E multipurpose data acquisition board. It provides eight channels of 16-bit differential-input analog-to-digital converters (ADCs), two 16-bit digital-to-analog converters (DACs), eight digital I/O (DIO) lines, and two internal 24-bit counter/timers. The board currently in use recently replaced an older board that was damaged; the ADCs on the new board acquire samples at a rate of 333 kSamples/s, while the old board (AT-MIO16XE10) acquired samples at 100 kSamples/s. The older board was used for all of the observing runs to date. All other inputs/outputs to/from the DAQ system are quasi-static.

All eight DIO lines are used as outputs to the CCD logic system, to set configuration bits. They select from pre-defined phase binning configurations and signal conditioning filters, and enable or disable the shutter and pixel-binning options. One DAC is used to supply a reference voltage to be used as a temperature reference for the heater system.

2.3.4 CCD Logic System

The CCD logic system is based on three Altera field-programmable gate arrays (FPGAs). These FPGAs are high-speed (10 MHz) logic devices with a large number of peripheral input-output lines. The FPGAs operate according to programs written in a high-level programming language, AHDL, with an interactive graphical programming environment. The devices are programmed via a parallel-port interface with the data acquisition computer; reprogramming requires only a few seconds. The flexibility to alter the operating logic is a great benefit both in the design/debugging stages of development and in the general operation of the instrument.

The CCD logic system is responsible for providing both the static and dynamic signals to the camera head electronics that control the CCD. These include voltage supplies, the three parallel and three serial phase controls, the transfer gate and reset gate controls, and a configuration word containing all of the settings for the camera head DACs. The CCD logic system also sends signals to the ADCs in the DAQ system, controlling the timing of the ADC samples of the analog CCD output, as well as a shutter open/close signal to the shutter control system.

The sequence of events that take place in a single phase-binning exposure is shown graphically

in Fig. 2.10. A negative-going pulse from the GPS system begins an exposure. The timing counters in the CCD logic are reset, so that all timing during the exposure is referenced to the falling edge of the match pulse. The CCD logic performs a fast frame erase of the CCD, removing all of the charge on the CCD in approximately 42 ms. The CCD logic enters a wait state until an integral number of phase bin intervals have passed since the match pulse went low. At the end of the wait state, the shutter is opened and the exposure begins. The exposure consists of alternating integrate and transfer states, until the match pulse goes high again. After every phase bin has been exposed the same number of times, the exposure ends. In other words, the exposure time (time between the shutter opening and closing) is always an integral multiple of the pulsar period, P .

Operation in the multiple frame transfer mode is identical to the operation described above, with the number of phase bins, n , set to unity. In multiple frame transfer mode, the frame transfers all shift charge in the same direction (down, in Fig. 2.5), unlike in phase-binning mode where there are an equal number of frame transfers in each direction.

In addition to controlling the frame transfer options, the CCD logic system determines the amount (if any) of pixel binning during readout. The pixel binning can be programmed to sum any number of pixels along the serial and parallel directions (independently). This reduces both the read noise in a given region of the CCD and the readout time, in exchange for the loss of spatial resolution. In practice, this option is only used in the laboratory, when aligning optics or establishing a coarse focus.

The timing of the ADC sampling is also controlled by the CCD logic system. Correlated double sampling (CDS) is performed on the analog CCD output, in order to eliminate the reset noise in each pixel (Janesick, 2001). The CDS is performed digitally, by taking an ADC sample after resetting each pixel and a second sample after clocking that pixel's accumulated charge onto the sense node of the readout amplifier. These two samples, the reset sample and the Charge sample, are then subtracted in software to produce a Difference signal, which is free of reset noise. The storage of this data is described in Section 2.4. This variety of CDS, as opposed to real time CDS methods, allows for more diagnostics of the noise properties of the CCD, by having access to the low-frequency signals in the reset and Charge samples without differencing. The downside to performing a digital CDS is that the read time is twice as long and the data storage requirements are twice as great. Neither of these considerations limits the performance of this instrument.

2.3.5 Camera Head

The camera head contains three circuit boards, a power board, driver board, and signal board. The power board has a number of DACs, allowing all of the operating voltages to be set programmably. The driver board converts the dynamic logic signals from the CCD logic system into analog voltages, as delivered by the power board. The signal board amplifies the analog signal from the CCD, and

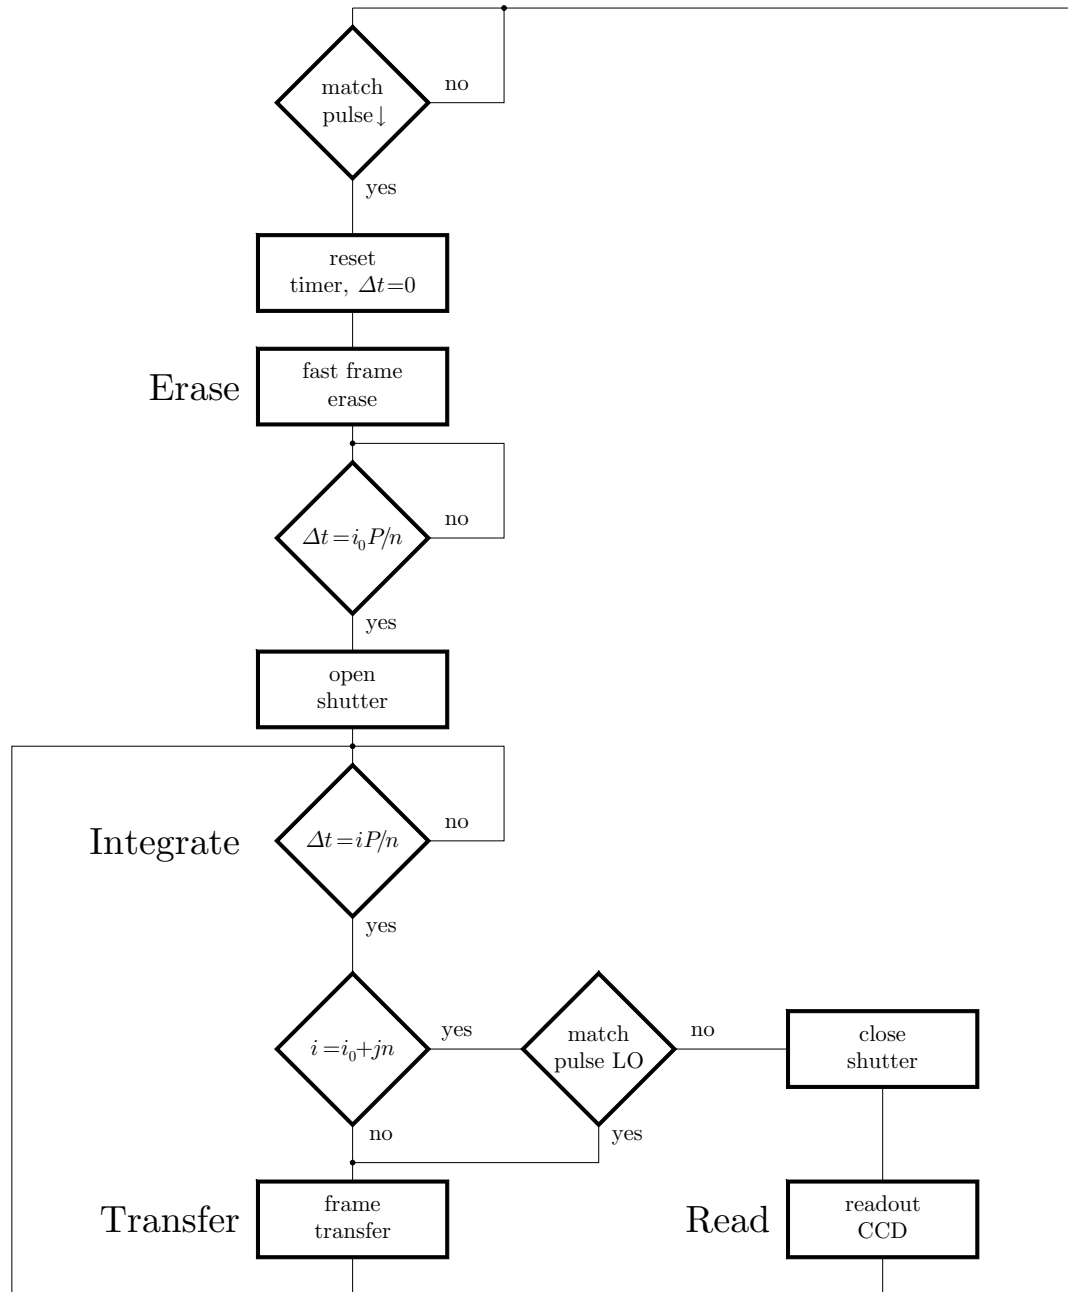


Figure 2.10: CCD logic flow chart. The indices i_0 , i , and j are integers. P is the pulsar period, n is the number of phase binned images.

supplies sufficient current to drive a transmission line to the ADCs in the DAQ system.

A number of operating parameters in the camera head can be set via a configuration string, delivered by the CCD logic system. The operating voltages that are delivered to the CCD are all outputs from a set of 12-bit DACs on the power board. In addition, the waveforms sent to the CCD are bandpass filtered by a number of resistors and capacitors that can be switched in or out of the driver circuits. The combinations of resistors and capacitors in the circuit are also set by the configuration string. This programmability allows for significant flexibility in optimizing the operating parameters of the CCD.

2.3.6 CCD

The CCD used is a SITe SI-502A 512×512 -pixel back-illuminated, antireflection coated multipinned phase CCD. The quantum efficiency (QE) of the CCD is $> 80\%$ from 400–800 nm, and above 60% to 900 nm. The multipinned phase (MPP) operation of the CCD involves keeping all of the parallel phases at negative voltages when integrating, ensuring that a population of charge holes is maintained at the surface interface, absorbing any surface dark current generated.

The small size of the CCD (12×12 mm, with $24 \mu\text{m}$ pixels) gives excellent charge transfer efficiency (CTE) performance. In laboratory tests, an LED was focused on the CCD for a brief period of time, then the charge was transferred repeatedly for a large number (10^6) of parallel shifts, and the spreading of the CCD spot was measured. Uniformly poor CTE is analogous to a loss of vertical spatial resolution, *i.e.*, poorer seeing. The CTE measured depends on the timing of the parallel phases causing the shifting. At a parallel shift rate of $42 \mu\text{s}/\text{line shift}$, the CTE was in excess of $1-10^{-8}$. For slightly faster shift rates, the CTE performance begins to degrade.

Because the CTE is so high on this chip, the charge traps are a more serious problem. To a first approximation, small spurious charge traps will be filled during every integrate state and emptied during every transfer state. Because the number of transfer states in an ordinary exposure is large (10^3-10^4), the charge redistributed by small traps delivers a large signal. More importantly, this signal grows linearly with the exposure time, as does the optically integrated signal. The treatment of these traps is fully discussed in Section 3.1.

2.4 Software

The control software for this instrument is written in C++ in a Windows environment, making extensive use of the Microsoft Foundation Class Library. The software provides a graphical interface for the observer to control and monitor sequences of exposures, selecting the operating mode (binning patterns and periods, exposure time) as well as several hardware parameters (ADC ranges, operating temperature). Once acquired, exposures are stored in 16-bit FITS format. The control software has

an interface with DS9, an image display program with a number of built-in image analysis features. When an exposure is acquired, DS9 is commanded to display the current image. A sample screen shot of the main window is shown in Fig. 2.11.

In practice, it is often important to be able to switch rapidly between exposures with no frame transfers at all, and images with either phase binning or multiple frame transfers. The software offers a simple check box to disable frame transfers. There is also a check box to enable and disable pixel binning, as described in Section 2.3.4.

Great care was taken to ensure that the start time of each exposure, set by the timing of the match pulse (see Fig. 2.3.4), is accurate. This can be difficult in a Windows environment, which has few provisions for real time accuracy in its operation. The main complication comes from an operating quirk in the GPS system. The GPS system operates the match pulse by matching the current time to a 16-bit match pulse register. This 16-bit register contains a binary-coded decimal (BCD) representation of the 10^{-3} – 10^{-6} s digits of the time at which the match pulse will toggle. Because of the limited resolution of the match pulse register, if the register is filled with the desired start time 10 ms or more before that time passes, the match pulse register will be satisfied too early. For instance, when the match pulse register is filled with any valid entry (an entry is valid if every 4-bit digit is between 0 and 9, excluding a–f), the match pulse will toggle exactly once every 10 ms (forming a 50 Hz square wave). Likewise, if the match pulse register is not cleared within 10 ms after the match pulse, the match pulse will toggle again. Therefore, the match pulse must be filled less than 10 ms before the desired match time, and cleared less than 10 ms after the desired match time. Therefore, the software must ensure that the Windows operating system executes the proper code during those time intervals. Note that the match pulse is the only function of the GPS system that suffers from this limited resolution; all other timing functions are accurate to a resolution of 10 years.

The execution thread responsible for setting the match pulse registers takes on the highest priority possible under Windows some time before the desired match time, and maintains that priority until after the match pulse register is cleared. The software repeatedly polls the GPS system for the current time, until the current time is less than 10 ms before the match time, sets the match time register, continues polling for the current time until after the match pulse, and clears the match pulse register. The sequence of current-time stamps is stored during this process, and is analyzed after the match pulse register is cleared, to determine if there is any possibility that the match pulse toggled at the wrong time (10 ms early, 10 ms late, or any multiple of these), and if so, terminates the current exposure to start again later. In addition, the match pulse from the GPS system is fed back into a time tag input to the GPS system, which the computer retrieves after the match pulse register is cleared, to verify yet again that the match pulse toggled at the proper time. If any of these timing diagnostics fail, the exposure is terminated. In practice, this generally happens in <

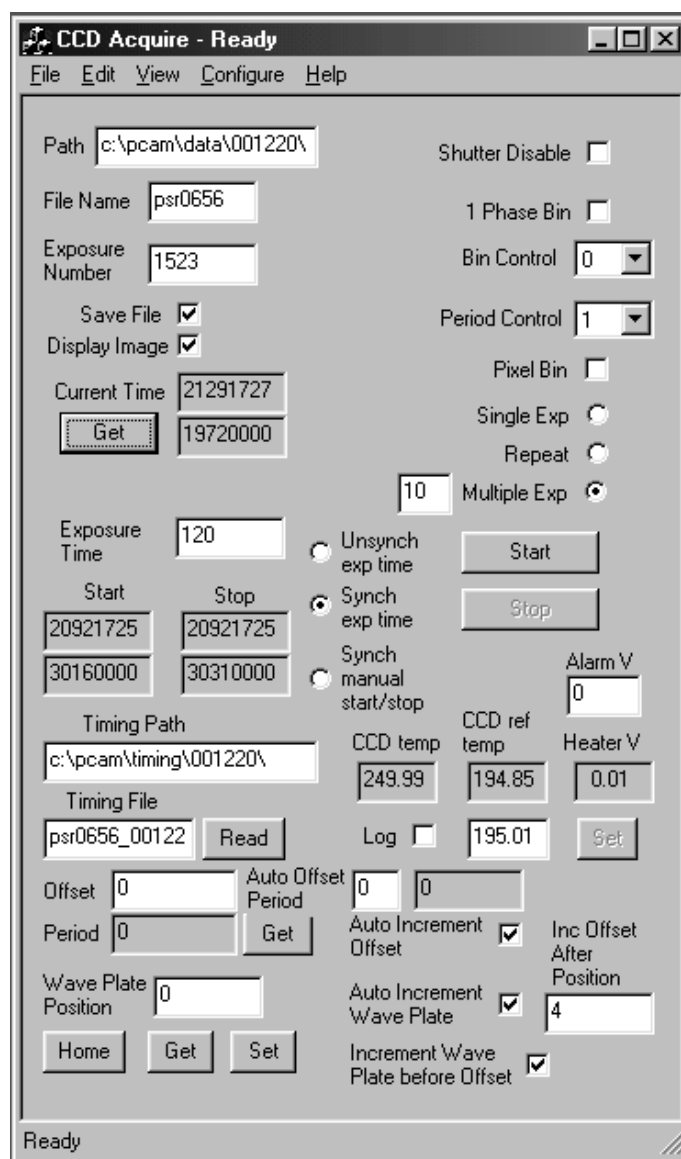


Figure 2.11: Control software screen shot.

1% of exposures.

Once the CCD logic system enters its Read mode (see Fig. 2.10), it remains in Read mode, continually reading out pixels, until the match pulse signal begins the next exposure. The control software sets the DAQ system to acquire only as many pixels as are desired, which causes the ADC to ignore any further signals from the CCD logic system. When phase binning, every exposure ends with just over half of the CCD containing accumulated charge (see Fig. 2.1). As such, the ADCs generally only acquire only about half of the CCD.

The acquired data is reorganized to make a FITS file for storage. Two samples, the reset and charge samples, are taken for every pixel (see Section 2.3.4). The Difference sample is free of reset noise, and is a 17-bit number. The Difference samples are stored in the FITS file as 16-bit numbers and 1-bit overflows. The FITS file that is output contains the lowest 16 bits of the Difference samples as the primary image, and the 16-bit reset samples and 1-bit Difference overflows as image extensions. The charge samples can be reconstructed from these images. All of the operating parameters are stored in the FITS header.

2.5 Performance

The first step in evaluating the performance of this instrument is to set forth the priorities under which this instrument operates. The two most important criteria for this instrument are its sensitivity and time resolution. These are the features that separate this instrument from the photon counting devices used by other groups.

The AR-coated, back-illuminated CCD we use has a quantum efficiency (QE) greater than 80% over the visible band (compare to $\sim 10\text{--}15\%$ efficiency for photocathodes of photon counting devices). The transmissive optics are all AR coated as well. Unfortunately, the system has never been used under photometric conditions, and since we are predominately interested in differential photometry (differential in the rotational phase domain), we have not been overly concerned with absolute calibration. The pulsar camera has been tested using a monochromator and calibrated photodiode assembly in Robinson Lab (equipment used by the Palomar Observatory staff), with measured system throughput in wide-field unfiltered mode of over 70%. Naturally, this number does not include the effects of atmospheric extinction and telescope transmission. A rough estimate of the end-to-end on-telescope throughput is 45%, from photons outside the atmosphere to detected photoelectrons.

The time resolution can, at some level, be traded for field of view in the frame transfer direction on the CCD. The time resolution is determined by the time spent performing transfers, which takes $42\ \mu\text{s}$ per line transfer. In phase binning mode, the bins are interleaved, so the number of lines in each transfer is twice the number of lines in an individual image. Slit edge effects would overwhelm the performance of the system if the images were less than 6 pixels high, so the minimum transfer




Figure 2.12: Sensitivity curves. The total integration time required to reach a signal-to-noise ratio of 10 in at least one bin, for flat pulses (solid lines) and peaked pulses (dashed lines), where all of the flux arrives in a single phase bin. The Palomar curves assume a sky background of $R = 20.5$ mag/arcsec² with 1.5 arcsec seeing, the Keck curves have $R = 21.5$ mag/arcsec² and 1.0 arcsec seeing. Compare this to Table 1.1 for known pulsars.

time is approximately 500 μ s. If one were to image in 10 phase bins, and insist on a 50% integration duty cycle (equal time integrating and transferring), the minimum pulsar period is 10 ms. This default configuration is then not useful for observations of millisecond pulsars.

We summarize the capabilities of this instrument in two plots, the sensitivity plot and the integration duty cycle plot. For the sensitivity plot (Fig. 2.12), we assume that the camera is being operated in wide-field phase-binning mode, with 10 phase bins, and that detection of pulsations requires a signal-to-noise ratio of 10 in at least one bin. We extrapolate these estimates from the results obtained on our observing runs. We plot the time required for detection for two extremes of the possible pulse profiles, one in which the pulse profile is essentially flat, and one where all of the intensity arrives in a single phase bin. We also plot this for two observatories, for Palomar with 1.4 arcsec seeing and $R=20.5$ mag/arcsec² sky, and for Keck with 1.0 arcsec seeing and $R=21.5$ mag/arcsec² sky. Both of the pulsars observed for this thesis (at Palomar) have magnitudes ~ 25 .

It is interesting to note the powerful effect that highly peaked pulse profiles have on the detectability of pulsations. For faint pulsars with highly peaked profiles, pulsations are easier to detect than the integrated light, because the sky background is divided evenly across the phase bins even though the pulsar intensity is not. With the Keck telescopes, it is entirely possible that currently undetected pulsars will be detectable because their profiles concentrate the flux into a small number of phase bins, without requiring the “normal” sequence of optical counterpart identification in pulse-averaged light followed by a search for pulsations.




Figure 2.13: Integration duty cycle. The solid line represents a 26-pixel high image, the dotted line a 6-pixel high image. The periods of three observed pulsars are labeled.

The integration duty cycle, shown in Fig. 2.13, shows the fraction of the total observation time spent integrating, versus the pulsar period. Two curves are shown, for different image heights, 6 pixels ($\sim 2.5\text{--}3$ arcsec) and 26 pixels ($\sim 10\text{--}13$ arcsec). The observation time required for a given signal-to-noise ratio is the integration time required divided by the integration duty cycle. For very low duty cycles, smearing becomes a problem, where the light that falls on a given pixel during transfer is comparable to the light that falls on that pixel during integration. Unless the field is very crowded, this is a small effect until the duty cycle becomes very low, as the transfer time distributes charge across a broad path, *e.g.*, across 52 vertical pixels. The effect of the smeared light, as it falls on an individual object, is related to the ratio of image height to seeing height, so for 26-pixel-high images, this smearing is generally lower by a factor of 10. So even for the 26-pixel high images viewing the Crab, with a duty cycle of $1/3$, the effect of the smearing is small (but visible), giving approximately $\frac{1/10 \cdot 2/3}{1/3} = 20\%$ as much light as the integrated image.

Chapter 3

Data Analysis

The objective of the data analysis routines is to manipulate a series of exposures to deliver a background-subtracted pulse profile for every point in the field of view. It is critical that the statistical routines be able to estimate the error in the measurements, and assess the overall confidence in the absence of systematic errors. In particular, the imaging properties allow one to make a map of pulsation amplitudes (at the pulsar period), which can alert the observer to the presence of unexplained errors, if many points (other than the pulsars) show pulsations.

The calibration and statistical analysis routines have proven to be extremely successful, delivering Poisson-limited errors with no systematic imprints. Fig. 3.6–3.8 provide a summary of the performance issues.

The calibration routines and statistical analysis presented here are unusual, and address the particular requirements posed by this technique. Charge traps, which dominate the systematic errors before calibration, cannot be calibrated out by simple dark and flat-field subtraction. The calibration routine is an iterative self-calibration routine which, while complicated, produces good results. We are not aware of any other treatment of image wander in a statistical analysis, which, if improperly performed, would produce large spurious systematic signals, destroying our confidence in the error analysis.

The treatment presented here is quite complete, but tedious. An attempt has been made to produce the important points graphically (or diagrammatically, in the case of the blank-sky detection algorithm).

3.1 Bias and Flat-Field Correction

The calibration routines are based on a specific model with a set of assumptions, which generally concern the separability of the measured flux into spatially and temporally varying components. Defining the counts (in units of electrons) measured at the CCD as $C(\mathbf{x}, t)$, where $\mathbf{x} = (x, y)$ is the CCD-pixel location and t is an exposure index (*i.e.*, a time variable, but not relative to the pulse phase), then the model specifies

$$\begin{aligned}
 C(\mathbf{x}, t)R(\mathbf{x}) + Z(\mathbf{x}) &= S(\mathbf{x} - \mathbf{x}_t)\eta_S(t) + B(t)\eta_B(t) + \varepsilon(\mathbf{x}, t), \\
 F(\mathbf{x}, t) &= C(\mathbf{x}, t)R(\mathbf{x}) + Z(\mathbf{x}), \\
 F(\mathbf{x}, t) &= S(\mathbf{x} - \mathbf{x}_t)\eta_S(t) + B(t)\eta_B(t) + \varepsilon(\mathbf{x}, t).
 \end{aligned}
 \tag{3.1}$$

Here, $S(\mathbf{x} - \mathbf{x}_t)$ is the source term, *i.e.*, the flux of the objects in the field of view. The variable $\mathbf{x}_t = (x_t, y_t)$ represents the time-dependent pointing of the telescope, and enters only in the source term. The implicit assumption here is that only the source term changes as the telescope pointing changes, and that the seeing does not change through the observations. $\eta_S(t)$ is the time-dependent source throughput factor, which includes loss terms due to extinction. $B(t)$ is the time-dependent background term, which includes the sky brightness, average dark current, and average spurious charge generation (as distinct from the charge redistributed by spurious traps). $\eta_B(t)$ is the time-dependent background throughput factor, which, in general, will vary differently than $\eta_S(t)$. In practice, it is not possible to independently determine $B(t)$ and $\eta_B(t)$. $R(\mathbf{x})$ is the spatially dependent response term, which includes static optical throughput variations, pixel-to-pixel QE variations, and the effect of spurious charge traps. $Z(\mathbf{x})$ is the spatially dependent zero-offset term, which includes the effect of spurious charge traps and the detector bias (even though the detector bias is not spatially dependent). In practice, there is a degeneracy between $\langle B(t)\eta_B(t) \rangle_t$, where $\langle \dots \rangle_t$ denotes a temporal average, and $\langle Z(\mathbf{x}) \rangle_{\mathbf{x}}$, where $\langle \dots \rangle_{\mathbf{x}}$ denotes a spatial average. $\varepsilon(\mathbf{x}, t)$ is an error term, which includes Poisson errors, Gaussian read noise errors, and the charge due to cosmic rays. Because the charge due to cosmic rays is positive-definite, $\varepsilon(\mathbf{x}, t)$ is not zero-mean. $F(\mathbf{x}, t)$ is the measured flux, which is a useful term when regrouping the other variables.

The goal of the calibration routines is to determine $R(\mathbf{x})$, $Z(\mathbf{x})$, and $B(t)\eta_B(t)$. This allows a direct estimation of $S(\mathbf{x} - \mathbf{x}_t)\eta_S(t)$, which is sufficient for our analysis, but precludes absolute photometry of the field of view (because of the lack of knowledge of η_S). The calibration is performed by finding regions where $S(\mathbf{x} - \mathbf{x}_t) \approx 0$, with no cosmic rays, which we call blank sky. We then separate the spatial and temporal components of the remaining terms. The determination of regions where $S(\mathbf{x} - \mathbf{x}_t) \approx 0$ depends on knowledge of all the other parameters, and separating the temporal and spatial components of those regions is easiest when the complementary components are already known, so these determinations are made iteratively. The process involves construction of several different pixel masks, which define the regions of blank sky suitable for use in calibration.

The first step in the calibration process is to take spatial and temporal averages of the measured flux. We define the spatially averaged flux, F_{sp} , as

$$\begin{aligned}
 F_{\text{sp}}(t) &= \langle F(\mathbf{x}, t) \rangle_{\mathbf{x}}, \\
 &= \langle S\eta_S \rangle_{\mathbf{x}} + B\eta_B, \\
 &= \langle CR \rangle_{\mathbf{x}} + \langle Z \rangle_{\mathbf{x}},
 \end{aligned}
 \tag{3.2}$$

where we have dropped the arguments of each of these functions for simplicity. By dropping the $\langle \varepsilon \rangle_{\mathbf{x}}$ term, we have assumed that pixels containing cosmic rays have not been included in the average, and replaced the mean value with the expected value (zero). When referring to the temporally averaged

flux, F_{te} , we construct it as a zero-mean average,

$$\begin{aligned}
F_{\text{te}}(\mathbf{x}) &= \langle F(\mathbf{x}, t) \rangle_t - \langle F(\mathbf{x}, t) \rangle_{\mathbf{x}, t}, \\
&= \langle S\eta_S \rangle_t - \langle S\eta_S \rangle_{\mathbf{x}, t}, \\
&= \langle C \rangle_t R + Z - \langle \langle C \rangle_t R + Z \rangle_{\mathbf{x}}.
\end{aligned} \tag{3.3}$$

The next calibration step is to identify pixels where $S(\mathbf{x} - \mathbf{x}_t) \approx 0$, and pixels with cosmic ray strikes. This can be done by comparing the observed flux to the sum of the averaged fluxes, $F_{\text{sp}} + F_{\text{te}}$. In general, we have estimates to the ‘‘true’’ parameters R and Z , which we index R_j and Z_j , that lead to an estimate F_j of the true flux F . This flux estimate, relative to the true flux, is best described by

$$F_j = \frac{R_j}{R} [F - Z] + Z_j. \tag{3.4}$$

The difference between the estimated flux and the average estimated fluxes is

$$\begin{aligned}
F_j - (F_{\text{sp},j} + F_{\text{te},j}) &= \frac{R_j}{R} [S\eta_S - \langle S\eta_S \rangle_t] - \left\langle \frac{R_j}{R} [S\eta_S - \langle S\eta_S \rangle_t] \right\rangle_{\mathbf{x}} + \\
&\quad \left[\frac{R_j}{R} - \left\langle \frac{R_j}{R} \right\rangle_{\mathbf{x}} \right] [B\eta_B - \langle B\eta_B \rangle_t] + \varepsilon \\
&= \frac{R_j}{R} [F - \langle F \rangle_t] - \left\langle \frac{R_j}{R} [F - \langle F \rangle_t] \right\rangle_{\mathbf{x}},
\end{aligned} \tag{3.5}$$

This equation need only be evaluated to an order-of-magnitude accuracy. In practice, $R \approx 1$, with $[\langle R^2 \rangle_{\mathbf{x}}]^{1/2} / R \sim 0.02$. This makes the $\frac{R_j}{R} - \left\langle \frac{R_j}{R} \right\rangle_{\mathbf{x}}$ term small, of the order of 10^{-2} . If the field is not particularly crowded, $\langle S\eta_S \rangle_{\mathbf{x}} / S\eta_S$ will be of the order of the source ‘‘fill factor,’’ typically $< 10^{-2}$. If the telescope tracking is good and the seeing conditions are relatively stable, the only temporal variation in $S(\mathbf{x} - \mathbf{x}_t)$ comes when the dither position changes, which is once every n exposures for wide-field imaging mode, or once every $4n$ exposures in two-channel polarimetric mode (with 4 exposures at different half-wave plate angles for every phase offset). We define the number of dither positions, n_{D} ,

$$n_{\text{D}} = \begin{cases} N_{\text{tot}}/n & \text{in wide-field mode,} \\ N_{\text{tot}}/(4n) & \text{in polarimetric mode,} \end{cases} \tag{3.6}$$

where N_{tot} is the total number of exposures in an observing run. Assuming a typical two-channel polarimetric observing run with $N_{\text{tot}} = 320$ and $n = 10$, $n_{\text{D}} = 8$. This implies that $\langle S\eta_S \rangle_t \sim (S\eta_S)/8$. With these numbers, in regions where $S\eta_S \gtrsim B\eta_B/100$,

$$F_j - (F_{\text{sp},j} + F_{\text{te},j}) \approx 0.9S\eta_S + \varepsilon. \tag{3.7}$$

Using Eq. 3.7, we can determine those regions where cosmic rays are present ($\varepsilon \gg 0$) or where $S > 0$. The sensitivity of this technique to small values of ε and S depends on the noise. The easiest way to reduce the noise is to sum over many samples, which in this case, can mean summing in time over $N_{\text{tot}}/n_{\text{D}}$ exposures, *i.e.*, summing over all of the exposures taken at a given dither position, as well as summing over all n phase-binned images (which all see the same field of view). All of the estimates of the magnitudes of the terms in Eq. 3.5 remain the same when averaging over exposures at a single dither position and over n phase-binned images, and so Eq. 3.7 is also valid after either form of averaging.

The ultimate goal of this calibration algorithm is not only to find blank sky regions, but also to determine the R and Z parameters of the model as described in Eq. 3.1. Both the iterative improvements in blank sky determination and in R and Z estimation rely on eliminating regions with $S > 0$ or $\varepsilon \gg 0$. This is done by excluding those regions when taking the spatial and temporal averages defined in Eqs. 3.2 and 3.3. We construct a mask, M^X , which is 1 in blank sky regions and 0 in regions with cosmic rays or sources. We denote the use of this mask in an average as F_{sp}^X or F_{te}^X . The iterative algorithm used to make masks and estimate R and Z is shown in Fig. 3.1.

The algorithm in Fig. 3.1 begins by creating masks that define Images 0–9 (see Figs. 2.1 and 2.3), for a single polarimetric column if the data is taken in polarimetric mode. (All of the calibration analysis is done separately on the two polarimetric columns, and only after the calibration is complete are the two columns considered together.) These image masks have superscripts I, i where i is the image number. The single-exposure masks (superscript 1, i) and source mask (superscript S , the same for all images) begin with no information. The $R_0(\mathbf{x})$ and $Z_0(\mathbf{x})$ estimates are initialized to unity and to the CCD bias, as estimated by the overscan regions of the CCD. For each image, the calibration algorithm goes through four steps.

The first step is to create an initial mask, $M_j^{0,i}(\mathbf{x}, t)$, which is a product of the appropriate image mask, the accumulated single-exposure mask (which generally only flags cosmic rays and very bright sources) and the accumulated source mask. In the first iteration ($j = 0$), this mask is identical to the current image mask.

The second step is to use the initial mask to create the single-exposure residual, $\Delta_j^{1,i}(\mathbf{x}, t)$, and a new single-exposure mask, $M_j^{1,i}(\mathbf{x}, t)$. The criterion for rejection by the single-exposure mask is that the residual must be more than five times the expected single-exposure error, $\sigma_1(t)$. The single-exposure error is defined as the quadrature sum of the Poisson noise from the background counts (as estimated by $F_{\text{sp}}^{1,i}(t)$, see Eq. 3.2) and the read noise (which is typically negligible).

The third step is to use the new single-exposure mask to create the intra-dither residual, $\Delta_j^{\text{D},i}(\mathbf{x}, t)$, and intra-dither mask, $M_j^{\text{D},i}(\mathbf{x}, t)$. These intra-dither variables are created by summing a given image over all of the exposures for which the telescope pointing, \mathbf{x}_t , is the same. For each image and each dither position, this is a sum over either n or $4n$ exposures, as defined in Eq. 3.6. As such, the

Initialize:

$$M^{I,i} = \begin{cases} 1, & x \text{ inside Image } i \\ 0, & x \text{ outside Image } i \end{cases}$$

$$M_{-1}^{1,i} = 1$$

$$M_{-1}^S = 1$$

$$R_0 = 1$$

$$Z_0 = \text{bias from overscan}$$

$$j = 0$$

Repeat:

$$F_j = CR_j + Z_j$$

Image 0	...	Image $n-1$
$i = 0$		$i = n-1$
$M_j^{0,i} = M^{I,i} M_{j-1}^{1,i} M_{j-1}^S$		$M_j^{0,i} = M^{I,i} M_{j-1}^{1,i} M_{j-1}^S$
$\Delta_j^{1,i} = F_j - (F_{\text{sp},j}^{0,i} + F_{\text{te},j}^{0,i})$		$\Delta_j^{1,i} = F_j - (F_{\text{sp},j}^{0,i} + F_{\text{te},j}^{0,i})$
$M_j^{1,i} = M_j^{0,i} \begin{cases} 1, & \Delta_j^{1,i} \leq 5\sigma_1 \\ 0, & \Delta_j^{1,i} > 5\sigma_1 \end{cases}$		$M_j^{1,i} = M_j^{0,i} \begin{cases} 1, & \Delta_j^{1,i} \leq 5\sigma_1 \\ 0, & \Delta_j^{1,i} > 5\sigma_1 \end{cases}$
$\Delta_j^{\text{D},i} = \sum_{N_{\text{tot}}/n_{\text{D}}} F_j - (F_{\text{sp},j}^{1,i} + F_{\text{te},j}^{1,i})$		$\Delta_j^{\text{D},i} = \sum_{N_{\text{tot}}/n_{\text{D}}} F_j - (F_{\text{sp},j}^{1,i} + F_{\text{te},j}^{1,i})$
$M_j^{\text{D},i} = M_j^{1,i} \begin{cases} 1, & \Delta_j^{\text{D},i} \leq 5\sigma_{\text{D}} \\ 0, & \Delta_j^{\text{D},i} > 5\sigma_{\text{D}} \end{cases}$		$M_j^{\text{D},i} = M_j^{1,i} \begin{cases} 1, & \Delta_j^{\text{D},i} \leq 5\sigma_{\text{D}} \\ 0, & \Delta_j^{\text{D},i} > 5\sigma_{\text{D}} \end{cases}$
Fit F_j to $aF_{\text{sp},j}^{\text{D},i} + b$		Fit F_j to $aF_{\text{sp},j}^{\text{D},i} + b$
$R_{j+1} = R_j/a$		$R_{j+1} = R_j/a$
$Z_{j+1} = Z_j - b/a$		$Z_{j+1} = Z_j - b/a$
$\Delta_j^S = \sum_{i = \text{Image } 0}^{\text{Image } n-1} F_j - F_{\text{sp},j}^{\text{D},i}$ $M_j^S = \begin{cases} 1, & \Delta_j^S \leq 5\sigma_S \\ 0, & \Delta_j^S > 5\sigma_S \end{cases}$ $j = j + 1$		

Figure 3.1: Iterative blank-sky calibration algorithm.

relevant error, $\sigma_D(t)$, grows as the square root of the number of exposures summed,

$$\sigma_D(t) = (N_{\text{tot}}/n_D)^{1/2}\sigma_1(t). \quad (3.8)$$

It should be noted that the intra-dither variables change only every n or $4n$ exposures, *i.e.*, $\Delta_j^{\text{D},i}(\mathbf{x}, t) = \Delta_j^{\text{D},i}(\mathbf{x}, t+1)$, if t and $t+1$ do not cross a dither boundary. It is also important to recognize that a given image samples all of the different pulsar phases in successive exposures, as well as all different polarizations if operating in polarization mode. By combining all of the intra-dither exposures of a given image, the assumption is that any actual pulse-phase or polarized flux variations between these exposures is negligible at the level of σ_D . To the extent that this assumption is violated, this algorithm is more likely to reject pixels with the extra variation. For sources that show a pulse-phase variation, this is good, because the goal is to eliminate regions where $S > 0$. If the background is highly polarized and the position angle changes significantly over the field of view, this will result in the mask rejecting all of the background except that with an arbitrary position angle. A highly polarized background with a constant position angle poses no problem, as flux changes between exposures (when the half-wave plate rotates) are perfectly represented by $F_{\text{sp}}(t)$. In principle, if a highly differentially polarized background posed such a problem, this algorithm could be updated to separate exposures at different position angles and sum only over exposures with similar position angles.

The fourth step is to use the intra-dither mask to estimate the blank-sky background flux alone, and improve the R_j and Z_j estimates. The intra-dither mask, $M_j^{\text{D},i}(\mathbf{x}, t)$, represents the most accurate estimate of the blank-sky region specific to an individual image (with time resolution equal to the number of dither positions, n_D). The blank-sky background flux is estimated over the entire image, *i.e.*, $F_{\text{sp},j}^{\text{D},i}(t)$ depends only on t . Assuming that the intrinsic background flux (*i.e.*, top-of-the-atmosphere flux) is uniform, any variations in $F_j(\mathbf{x}, t)$ in the blank-sky regions are due to errors in the estimates of $R_j(\mathbf{x})$ and $Z_j(\mathbf{x})$. As the background flux varies through a night of observing (as the airmass changes, the moon rises, *etc.*), a large number of samples of $F_{\text{sp},j}^{\text{D},i}(t)$ are accumulated, and the linear relationship with $F_j(\mathbf{x}, t)$ can be established at every \mathbf{x} .

After these four steps have been completed for all n images, the image information can be combined to form a composite mask from all of the images. This is done by mapping \mathbf{x} into $\boldsymbol{\vartheta} = (\vartheta_x, \vartheta_y)$, where $\boldsymbol{\vartheta}$ is angle on the sky. This simplest form of the transformation is $\vartheta_x = x$, $\vartheta_y = y - y_0 \bmod \Delta y$, where y_0 is the bottom of the lowest image (Image 4, see Fig. 2.1), and Δy is the distance between images, measured in pixels on the CCD. The source residuals and mask, $\Delta_j^{\text{S}}(\boldsymbol{\vartheta}, t)$ and $M_j^{\text{S}}(\boldsymbol{\vartheta}, t)$, are formed in the same way as the other masks, summing the images with

their appropriate intra-dither masks. The source error term is

$$\sigma_S(t) = n^{1/2}\sigma_D(t). \quad (3.9)$$

It would be possible to extend this technique to sum over all exposures, shifting the images by \mathbf{x}_t to account for the telescope dither, and produce an estimate with a lower noise ($\sigma_S/n_D^{1/2}$). In practice, not only is this not necessary, but the lower noise levels result in rejecting large regions of the images due to small systematic errors rather than actual source terms.

After construction of the source mask, M_j^S , the iterative process is repeated, until the improvement in R_j and Z_j is small. At each successive iteration, the initial mask, $M_j^{0,i}$ incorporates the source mask from the previous iteration, with the transformation from \mathbf{x} to $\boldsymbol{\vartheta}$ inverted for each image. Once the process has converged, the last (*i.e.*, most comprehensive) intra-dither mask, $M_j^{D,i}$, is used to determine $F_{\text{sp},j}^{D,i}(t) \approx B(t)\eta_B(t)$ (see Eq. 3.2).

With good estimates of $R(\mathbf{x})$, $Z(\mathbf{x})$, and $B(t)\eta_B(t)$, along with knowledge of the location of cosmic rays, Eq. 3.1 allows a direct estimation of $S(\mathbf{x}, t)\eta_S(t)$. We have not yet attempted to perform absolute photometry, which would require knowledge of $\eta_S(t)$ (or verification that it is constant). We only assume that $\eta_S(t)$ does not change appreciably on the time scale of the pulse period, or more importantly, that changes on the time scale of the pulse period are not themselves periodic, so that by averaging over 10^5 – 10^6 cycles ensures that the differential photometry (between phase-binned images) is not affected by nonphotometric observing conditions.

The last step in the image reduction is to transform the arguments of the source fluxes to variables appropriate for analysis. This entails converting \mathbf{x} to $\boldsymbol{\vartheta}$, accounting for the telescope dither \mathbf{x}_t , and i (the image index). In addition, the image indices, i , are transformed into phase bin indices, p , by applying the phase offset, φ_{off} (see Eq. 2.4),

$$p = (i - \varphi_{\text{off}}) \bmod n. \quad (3.10)$$

If the images are taken in polarimetric mode, then the exposures for a given polarimetric column are arranged with respect to the half-wave plate rotation angle, giving an index l , and the exposure number is reordered and indexed by k ,

$$k = \begin{cases} t, & \text{wide-field mode} \\ t/4, & \text{polarimetric mode} \end{cases} \quad (3.11)$$

$$l = \begin{cases} 0, & \text{wide-field mode} \\ t \bmod 4, & \text{polarimetric mode} \end{cases}. \quad (3.12)$$

For convenience, we define N_k to be the number of reordered exposures, $N_k = nm_D$ (see Eqs. 3.6 and 3.11). We also denote the two polarimetric columns, which were independently calibrated, by index m (and assign $m = 0$ to all wide-field data). This leads to a background-subtracted source intensity I , where we have included the throughput $\eta_S(t)$,

$$I(\boldsymbol{\vartheta}, p, k, l, m) = S\{\mathbf{x}(\boldsymbol{\vartheta}, p, l, m), t(k, l)\}\eta_S\{t(k, l)\}, \quad (3.13)$$

and a background intensity I_0 ,

$$I_0(k, l, m) = B\{t(k, l)\}\eta_B\{t(k, l)\}, \quad (3.14)$$

where $\boldsymbol{\vartheta}$, p , k , l , and m are the sky position, phase bin, exposure, position angle, and polarization column variables, respectively.

3.2 Error Estimation and Statistics of Correlated Image Wander

The dominant noise term in the pulse profiles of any point in the field of view is Poisson noise, due to the sky background, or due to bright individual sources. Because the number of detected photoelectrons per pixel is large ($> 10^3$), even in blank sky regions, the Poisson distribution is well represented by the normal distribution. The read noise (5 e⁻ rms) is generally small compared to the Poisson noise, and is also normally distributed. The simplest theory is that the errors are normally distributed with variance

$$\sigma_{\text{th}}^2 = [I + I_0] + \sigma_{\text{RN}}^2, \quad (3.15)$$

where σ_{RN}^2 is the read noise variance. For most locations, this is well approximated by $\sigma_{\text{th}}^2(\boldsymbol{\vartheta}, p, k, l, m) \approx I_0(k, l, m)$, *i.e.*, the noise is background-limited.

An important step in verifying the performance of the imaging system is to examine the pulse profiles of every point in the field of view, comparing the null hypothesis (constant flux) to the observed variability. A significant violation of the null hypothesis, at locations away from the pulsar, would indicate either the presence of a systematic error mimicking pulsations, or an underestimate of the random errors.

A more realistic error estimate can be obtained by using a technique very similar to that used in Sec. 3.1, where the measured flux was separated into spatial and temporal averages. To estimate errors, the calibrated, background-subtracted intensities, $I(\boldsymbol{\vartheta}, p, k, l, m)$ are separated into phase-averaged intensities, I_{ph} , exposure-averaged intensities, I_{ex} , and phase- and exposure-averaged

intensities, $I_{\text{ph,ex}}$,

$$I_{\text{ph}}(\boldsymbol{\vartheta}, k, l, m) = \langle I(\boldsymbol{\vartheta}, p, k, l, m) \rangle_p, \quad (3.16)$$

$$I_{\text{ex}}(\boldsymbol{\vartheta}, p, l, m) = \langle I(\boldsymbol{\vartheta}, p, k, l, m) \rangle_k, \quad (3.17)$$

$$I_{\text{ph,ex}}(\boldsymbol{\vartheta}, l, m) = \langle I(\boldsymbol{\vartheta}, p, k, l, m) \rangle_{p,k}. \quad (3.18)$$

These averages are then recombined to form intensity residuals, Δ_I ,

$$\Delta_I = I - [(I_{\text{ph}} - I_{\text{ph,ex}}) + (I_{\text{ex}} - I_{\text{ph,ex}})]. \quad (3.19)$$

The unbiased estimate of the single-pixel variance, $\sigma_{I,\text{pix}}^2$, is then

$$\sigma_{I,\text{pix}}^2(\boldsymbol{\vartheta}, l, m) = \sum_{p,k} \Delta_I^2(\boldsymbol{\vartheta}, p, k, l, m) / [nN_k - (n-1) - (N_k - 1)]. \quad (3.20)$$

Implicit in Eq. 3.19 is the assumption that, for a given $\boldsymbol{\vartheta}$, l , and m , the intensity $I(\boldsymbol{\vartheta}, p, k, l, m)$ is separable into variations in phase p and exposure k , but that there is no change in the pulse profile between exposures. There is a second-order violation of this assumption, through the variation of $\eta_S(t)$. In cases where the true pulse profile is not uniform, the measured variations from the mean flux level will depend linearly on $\eta_S(t)$, and so the variance measured by this technique will be higher than the true “error” in the exposure-averaged pulse profile. To first order, the fact that the estimated variance is properly mean-subtracted ensures that the estimate is correct to first order. Because we expect that only the pulse profile of the pulsar will be variable, then this caveat means there is a second-order overestimate of the variance at the location of the pulsar.

The variances that are of interest in the final analysis are those affecting the pulse profiles, *i.e.*, the exposure-averaged intensity, $\langle I \rangle_k$. The unbiased exposure-averaged variance is

$$\sigma_{\langle I \rangle_k, \text{pix}}^2(\boldsymbol{\vartheta}, l, m) = \sigma_{I,\text{pix}}^2(\boldsymbol{\vartheta}, l, m) / (nN_k). \quad (3.21)$$

In some regions, the variances, $\sigma_{I,\text{pix}}^2$ and $\sigma_{\langle I \rangle_k, \text{pix}}^2$, contain a component due to image wander. Image wander, due to the tip/tilt component of the seeing, will cause intensity variations wherever the intensity has a nonzero spatial gradient. Specifically, if the image motion is $\Delta\boldsymbol{\vartheta}(\boldsymbol{\vartheta}, p, k, l)$, then the intensity variation is

$$\Delta I = \Delta\boldsymbol{\vartheta} \cdot \nabla I, \quad (3.22)$$

where $\nabla = (\partial/\partial\vartheta_x, \partial/\partial\vartheta_y)$. While the contributions of image wander to the intensity variations will be properly included in the intensity residuals, Δ_I , these variations are likely to be correlated

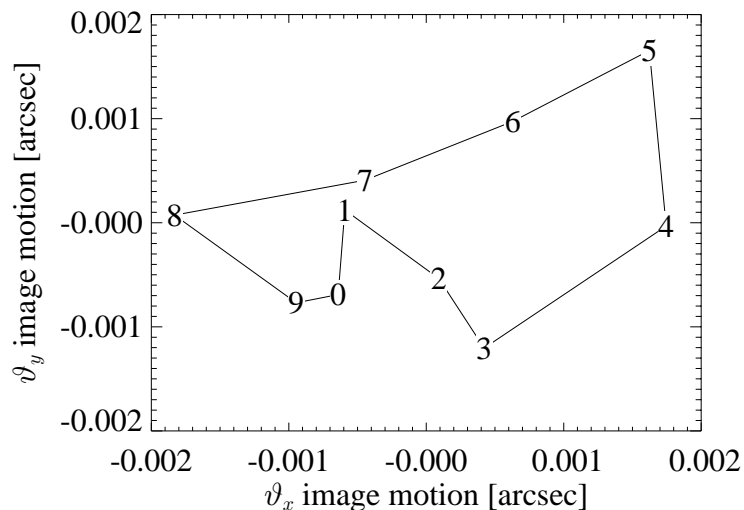


Figure 3.2: Image wander of Star 3 in field of 4U 0142+61. Image motion is measured in arcsec, with a plate scale of 0.4 arcsec/pixel. Numbers are phase bin indices.

between phase bins. As such, the mean squared residuals do not follow a χ^2 distribution.

We idealize the intensity variations from image wander by assuming that the image motion is sinusoidal between phase bins. The specific two-dimensional distribution of the image motion is not important, as only the component along the intensity gradient affects the measured intensity (see Eq. 3.22). The sinusoidal aspect of the image motion is the lowest-order approximation given that the image motion must be periodic. A plot of the measured image motion in the field of 4U 0142+61 is shown in Fig. 3.2. We also assume that the image wander is the same everywhere in the field of view, so that the relative contribution of the image motion to the measured variance depends only on the local intensity gradient.

To further simplify the analysis of the effects of image wander, we assume that the amplitude, a_0 , of the image motion is the same in all exposures, but that the phase of the sinusoid that defines the actual $\Delta\boldsymbol{\vartheta}$ observed in an exposure is random (*i.e.*, uncorrelated between successive exposures). When the image wander is averaged over all exposures (N_k exposures), the average image wander will still be sinusoidal, but with a reduced amplitude, a . If the number of exposures over which the image motion is averaged is large, the probability density function of the squared amplitude, $p_{a^2}(a^2)$ is exponential,

$$p_{a^2}(a^2) = N_k/a_0^2 e^{-a^2 N_k/a_0^2}. \quad (3.23)$$

This is most easily seen by performing the average in Fourier phase-bin space, where the the image motion in an exposure is always at the same frequency (1 cycle per n phase bins), with modulus

a_0 , and a random phase (and its Hermitian duplicate). Averaging a large number, N_k , of these components is an example of the classic random walk problem, resulting in an average squared modulus of a_0^2/N_k . The important issue is that these errors are non-normal, *i.e.*, the squared modulus does not follow a χ^2 distribution.

We separate the errors into errors with a Gaussian distribution, g , and sinusoidal errors, s . The definition of the measured residuals in Eq. 3.5 ensures that only errors that are uncorrelated between bins and the sinusoidal (zero-mean) errors need to be considered. The errors in the exposure-averaged intensity, $I_{\text{ex}}(\boldsymbol{\vartheta}, p, l, m)$, at a given $(\boldsymbol{\vartheta}, l, m)$, are then

$$\varepsilon(p) = g(p) + s(p), \quad (3.24)$$

$$s(p) = a \sin(2\pi p/n - \varphi_0), \quad (3.25)$$

$$p_g(g) = N(0, \sigma_g^2), \quad (3.26)$$

$$p_{a^2}(a^2) = 1/a_m^2 e^{-a^2/a_m^2}, \quad (3.27)$$

where a is the amplitude of the intensity variations in Eq. 3.22 (not simply the amplitude of the image motion), φ_0 is a random phase offset (which we take to be zero), $p_g(g)$ is the probability density function of g , $N(0, \sigma_g^2)$ is the normal distribution with variance σ_g^2 , $p_{a^2}(a^2)$ is the probability density function of a^2 , and a_m^2 is the mean a^2 . The pulse residuals will then be

$$\begin{aligned} \varepsilon_{\text{pulse}}^2 &= \sum_p [\varepsilon(p) - \bar{\varepsilon}]^2, \\ &= \sum_p [g(p) - \bar{g}]^2 + 2 \sum_p [g(p) - \bar{g}] p(g) - \sum_p [p(g)]^2, \\ &= \sum_p [g(p) - \bar{g}]^2 + 2 \sum_p [g(p) - \bar{g}] p(g) - a^2/2. \end{aligned} \quad (3.28)$$

The first term, divided by σ_g^2 , is $\sum [g - \bar{g}]^2 / \sigma_g^2$, which is distributed as χ^2 with 9 degrees of freedom. The second term, $2 \sum [g - \bar{g}] p$, is the product of two uncorrelated, zero-mean variates. This term is small, and will be considered as a correction after accounting for the first and third terms. The third term, $a^2/2$, is distributed exponentially.

The distribution of a sum of a χ^2 -distributed variable and an exponentially distributed variable can be analytically determined. Neglecting the second term in Eq. 3.28, we define the pulse residuals normalized to σ_g^2 , ε_g^2 , and a normalization parameter, b , as

$$\varepsilon_g^2 = \varepsilon_{\text{pulse},1,3}^2 / \sigma_g^2, \quad (3.29)$$

$$b = 1/2 - \sigma_g^2 / (a_m^2 / 2). \quad (3.30)$$

The probability density function of ε_g^2 is then,

$$p_{\varepsilon_g^2}(\varepsilon_g^2) = (1/2 - b) e^{-\varepsilon^2/2} / 16 b^4 \left\{ \begin{array}{l} e^{b\varepsilon^2} \operatorname{erf}[(b\varepsilon^2)^{1/2}] / (2b)^{1/2} - \\ [105 + 70b\varepsilon^2 + 28(b\varepsilon^2)^2 + 8(b\varepsilon^2)^3] (2\varepsilon^2)^{1/2} / 105\pi^{1/2} \end{array} \right\}, \quad (3.31)$$

where $\operatorname{erf}[\dots]$ is the error function, and the subscript on ε_g^2 was dropped on the right-hand side of the equation. In the limit of $a_m^2 \rightarrow 0$ (*i.e.*, $b \rightarrow -\infty$), $p_{\varepsilon^2}(\varepsilon^2)$ reduces to χ^2 with 9 degrees of freedom. Eq. 3.31 is valid for all values of b (including when $b \rightarrow 0$), understanding that $\operatorname{erf}[(b\varepsilon^2)^{1/2}] / (2b)^{1/2}$ becomes $\operatorname{erfi}[(-b\varepsilon^2)^{1/2}] / (-2b)^{1/2}$ when $b < 0$, where erfi is the imaginary error function.

The normalization of ε_g is convenient from the standpoint of writing Eq. 3.31, but does not directly correspond to a measured quantity. We have a measurement of $\sigma_{(I)_k, \text{pix}}^2$ (Eq. 3.21), which encompasses both of the idealized error sources (Eq. 3.24). To normalize ε to the measured variance, we introduce a variance ration r , and a scaling factor f , to arrive at the probability density function of the pulse residuals normalized to the measured variance $\varepsilon_{\text{meas}}^2$,

$$r = \sigma_g^2 / (a_m^2 / 2), \quad (3.32)$$

$$f = 1 + (a_m^2 / 2) / (9\sigma_g^2), \quad (3.33)$$

$$\varepsilon_{\text{meas}}^2 = f\varepsilon_g^2, \quad (3.34)$$

$$p_{\varepsilon_{\text{meas}}^2}(\varepsilon_{\text{meas}}^2) = fp_{\varepsilon_g^2}(\varepsilon_{\text{meas}}^2). \quad (3.35)$$

Plots of $p_{\varepsilon_{\text{meas}}^2}(\varepsilon_{\text{meas}}^2)$, for several values of f (and therefore b), appear in Fig. 3.3. The most notable feature of Fig. 3.3 is the increase in the high-end tail of the distribution for small values of r . This implies that a few of the regions of the image with large intensity gradients will show very large residuals.

All of the analytic expressions describing the distribution of pulse residuals have ignored the cross term in Eq. 3.28, which we claimed was small. To evaluate the effect of this cross term, we make a simple assumption, that the resulting probability distribution is a member of the family of distributions described by Eq. 3.31, with a different value of b . We define $r_2(r)$, the effective ratio r (which also defines b) that describes the full probability distribution incorporating the cross term. This effective ratio is determined by constructing a frequency distribution of pulse residuals for each true ratio r using a Monte-Carlo routine, comparing the frequency distributions to the distributions defined by Eq. 3.31, and defining $r_2(r)$ as the ratio input into Eq. 3.31 that best matches the Monte-Carlo frequency distribution. In the course of comparing Monte-Carlo frequency distributions to Eq. 3.31, we also define r_{min} and r_{max} , the minimum and maximum r for which the frequency distributions differ significantly from the limiting distributions. For $r < r_{\text{min}}$, the




Figure 3.3: Probability distribution function of image wander. Identifying the curves by their peak positions, from left to right, are values of $r = 0, 0.01, 0.03, 0.1,$ and ∞ . $r = 0$ is an exponential distribution, and $r = \infty$ is a χ^2 distribution with 9 degrees of freedom. The important point is that all of the distributions have larger tails in the high error ranges than the χ^2 distribution, *i.e.*, there will be more outliers under these distributions.

exponential distribution ($r = 0$) is applicable, and for $r > r_{\max}$, the χ^2 distribution ($r = \infty$) is applicable. A plot of $r_2(r)$ is shown in Fig. 3.4. As predicted, the effect of the cross term is small, varying r_2 by a factor of 3 over 3 orders of magnitude in r .

To apply this probability distribution to the measured pulse residuals, the correct value of r for a given point $(\boldsymbol{\vartheta}, l, m)$ must be determined, with which a value of $r_2(r)$ is calculated and fed into Eq. 3.31. Knowing $I_0(\boldsymbol{\vartheta}, p, k, l, m)$ from Eq. 3.14, we estimate the Gaussian variance $\sigma_g^2(\boldsymbol{\vartheta}, l, m)$ by averaging over k and p and assuming that the errors are Poissonian. We also calculate the intensity gradient (Eq. 3.22) at each point. We then fit a global amplitude conversion factor α to the equation

$$\sigma_{\langle I \rangle_k, \text{pix}}^2(\boldsymbol{\vartheta}, l, m) = \sigma_g^2(\boldsymbol{\vartheta}, l, m) + \alpha |\nabla \langle I \rangle_k(\boldsymbol{\vartheta}, l, m)|. \quad (3.36)$$

Using Eq. 3.28, we now set

$$a_m^2(\boldsymbol{\vartheta}, l, m) = 2\alpha |\nabla \langle I \rangle_k(\boldsymbol{\vartheta}, l, m)|. \quad (3.37)$$

We now have all the components necessary to compare the measured pulse residuals with the distribution we expect in the absence of pulsations. The best visual metric for determining whether the errors have been realistically calculated is to plot a histogram of the cumulative probability




Figure 3.4: Scaling correction from pulse residual cross term. r is the variance ratio defined by Eq. 3.32, and $r_2(r)$ is the parameter that, when used in Eq. 3.31, best matches a Monte-Carlo frequency distribution calculated using r . The line marked r_{\min} denotes the r for which the best-fitting probability distribution is exponential. The line marked r_{\max} denotes the r for which the best-fitting probability distribution is χ^2 .

distribution $P_{\varepsilon_{\text{meas}}^2}(\varepsilon_{\text{meas}}^2)$,

$$P_{\varepsilon_{\text{meas}}^2}(\varepsilon_{\text{meas}}^2) = \int_0^{\varepsilon_{\text{meas}}^2} p_{\varepsilon_{\text{meas}}^2}(\varepsilon^2) d\varepsilon^2. \quad (3.38)$$

If the errors have been properly calculated, a histogram of the cumulative probability distribution will be uniform. We compare the histograms of the cumulative probability distributions of the measured errors assuming the population distribution is a χ^2 with 9 degrees of freedom, $P_{\chi^2}(\varepsilon_{\text{meas}}^2)$, in Fig. 3.5, to that assuming the probability distribution is as calculated in Eq. 3.38, $P_{\varepsilon_{\text{meas}}^2}(\varepsilon_{\text{meas}}^2)$, in Fig. 3.6. As can be seen in Fig. 3.5, the simple χ^2 interpretation of the errors leads to an overidentification of regions of large variability. Looking at the 0.999–1.000 bin of Fig. 3.5, we would identify roughly 20 pulsars, expecting 2.7 of those to be spurious identifications. Looking at Fig. 3.6, on the other hand, we identify 2 pulsars, expecting 2.7 of these to be spurious. The difference between Fig. 3.5 and Fig. 3.6 is a graphical representation of how the differences in the tails of the distributions in Fig. 3.3 result in large differences in the interpretation of the residuals.

For identification of the regions of largest variation (*i.e.*, the real pulsar), the histogram ceases to be the most efficient means of identifying the variations. A variability image of the field of 4U 0142+61 is shown in Fig. 3.7, where the two pixels that form the 0.999–1.000 bin of Fig. 3.6 are readily identifiable. The cumulative probability at the location of 4U 0142+61 is $1 - 6 \times 10^{-5}$, and the cumulative probability at the peak 2 arcsec E of star 3 is $1 - 1.2 \times 10^{-4}$. The probability that,




Figure 3.5: Histogram of cumulative probability distribution of errors, assuming a χ^2 distribution with 9 degrees of freedom. Data is taken from field of 4U 0142+61. Left and right panels differ only in the x -axis range, and in the bin size. The left panel has bin sizes of 0.01, the right panel 0.001. The dotted line denotes the level expected from a uniform histogram. There is a large overpopulation at high values of $\varepsilon_{\text{meas}}^2$.




Figure 3.6: Histogram of cumulative probability distribution of errors, assuming a $P_{\varepsilon_{\text{meas}}^2}$ distribution, as defined in Eq. 3.38. Data is taken from field of 4U 0142+61. Left and right panels differ only in the x -axis range, and in the bin size. The left panel has bin sizes of 0.01, the right panel 0.001. The dotted line denotes the level expected from a uniform histogram. The histogram is very uniform at all scales. The range of N in both panels is the same as in Fig. 3.5.




Figure 3.7: Map of cumulative probability of pulse residuals. The top and middle panels are direct images of the field of 4U 0142+61, with logarithmic (top) and linear (middle) stretch. The numbering of stars comes from Reid *et al.* (1980) Coe and Pightling (1998), and Hulleman *et al.* (2000a). 4U 0142+61 is star A. The bottom panel is $P_{\epsilon_{\text{meas}}^2}(\epsilon_{\text{meas}}^2)$ for each point, displayed logarithmically, where dark is a high cumulative probability. The cumulative probability at the location of 4U 0142+61 is $1 - 6 \times 10^{-5}$. The signal-to-noise ratio at Star 3 (the brightest region in this field) is over 150 per phase bin.

given 2720 samples (the number of pixels in Fig. 3.7), at least one pixel will have a value of P greater than $1 - 1.2 \times 10^{-4}$, the value of P at the spurious peak, is 28%. This is in excellent agreement with the hypothesis that at every point in the field except at 4U 0142+61, the residuals are distributed according to Eq. 3.31, with no spurious systematic measured variability. The signal-to-noise ratio at the brightest star in the field is over 150 per phase bin, with no measurable systematic fluctuations between phase bins.

As a last analysis of the errors present in the data reduction, we examine the relationship between the measured errors and the errors predicted from first principles. Because the variance due to image wander was globally fit to compensate for the level of measured errors, we do not consider the variance due to image wander when analyzing the error from first principles. Assuming that the background intensity, I_0 , has only Poissonian errors, we compare ϵ_{meas}^2 directly to I_0 in Fig. 3.8, assuming that $\sigma_g^2 = \langle I_0 \rangle_k$ (this neglects read noise, for instance). The blank sky regions of the image will have no image wander variance (because the intensity gradients are zero), and so we see a floor to the plot of $\epsilon_{\text{meas}}^2 / \sigma_g^2$ at a ratio of unity, *i.e.*, the errors are Poissonian. The large values of this ratio are due to regions where the image wander adds to the variance. For comparison, Fig. 3.8 also shows the ratio




Figure 3.8: Ratio of measured errors to predicted errors. The top panel shows the ratio of measured errors to Poissonian errors, where $\sigma_g^2 = \langle I_0 \rangle_k$, the Poisson errors in the background. The bottom panel shows the ratio of measured errors to errors modeled in Eq. 3.28, including the image wander variance. In each plot, the ratios are plotted for cuts along ϑ_x , at each value of ϑ_y in the field of view.

of measured errors to modeled errors, *i.e.*, $\epsilon_{\text{meas}}^2 / (\sigma_g^2 + a_m^2/2)$ (see Eq. 3.28), where we ignore the cross term. The inclusion of the image wander variance makes the ratio of measured to predicted errors very nearly uniform.

It has not been demonstrated that this model of image wander is the correct explanation for the measured variance. The appeal of this explanation is that it is a single-parameter model, which fits only the global amplitude conversion factor α (see Eq. 3.36) to the entire dataset, giving a resulting distribution of expected variances that matches the measured variances quite well. The image wander, as described in this section, must be a factor at some level, although the assumption that the wander is well-described by a single sinusoidal term (rather than a sum of Fourier components) will be violated for long-period pulsars (in which the image wander becomes less correlated between bins). In addition, the measured variance shows a pattern which closely matches the intensity gradient, showing a “donut” with a minimum at the center of each bright object (where the intensity gradient passes through zero). However, the temporal components of the image wander can be exploited in a simple experiment, where exposures are taken with different pulsar periods, in which the image wander variance (as defined by the global amplitude conversion factor α) should reflect the truncation of the atmospheric tip-tilt power spectrum at different frequencies. This experiment has not yet been done, and should provide clearer evidence for or against the interpretation of the variance presented in this section.

3.3 Statistics of Polarization

When polarization is measured with low signal-to-noise ratios, systematic biases can greatly affect the measurements. The linearly polarized flux, L , is computed from the Stokes parameters Q and U by

$$L = (Q^2 + U^2)^{1/2}. \quad (3.39)$$

Because L is positive-definite, when Q and U are both near zero, noise in Q and U results in a systematic overestimate of L .

To conform with the traditional labeling of Stokes parameters, we use I in this Section in reference to the Stokes I parameter, which can be understood as $I_{\text{Stokes}} = 2 \langle I(\boldsymbol{\vartheta}, p, k, l, m) \rangle_{k,l,m}$, where we will not use the subscript on I_{Stokes} further. To facilitate the labeling in equations in this section, we define S_0 , S_{45} , S_{90} , and S_{135} with a different (although related) definition than the S defined in Eq. 3.1. We define these S variables as

$$S_0(\boldsymbol{\vartheta}, p) = [\langle I(\boldsymbol{\vartheta}, p, k, 0, 0) \rangle_k + \langle I(\boldsymbol{\vartheta}, p, k, 2, 1) \rangle] / 2, \quad (3.40)$$

$$S_{45}(\boldsymbol{\vartheta}, p) = [\langle I(\boldsymbol{\vartheta}, p, k, 1, 0) \rangle_k + \langle I(\boldsymbol{\vartheta}, p, k, 3, 1) \rangle] / 2, \quad (3.41)$$

$$S_{90}(\boldsymbol{\vartheta}, p) = [\langle I(\boldsymbol{\vartheta}, p, k, 2, 0) \rangle_k + \langle I(\boldsymbol{\vartheta}, p, k, 0, 1) \rangle] / 2, \quad (3.42)$$

$$S_{135}(\boldsymbol{\vartheta}, p) = [\langle I(\boldsymbol{\vartheta}, p, k, 3, 0) \rangle_k + \langle I(\boldsymbol{\vartheta}, p, k, 1, 1) \rangle] / 2. \quad (3.43)$$

In all of the polarization analysis, we will deal only with averages over exposures k . See Eq. 3.12 for a definition of the l variable. The variation in the l variable denotes observations taken at different half-wave plate angles, which defines different polarization position angles of the observations. Because two orthogonal polarizations are sampled simultaneously (the two-channel operation, see Fig. 2.4), the $m = 0$ images at a half-wave plate angle of 0° (part of S_0) sample the same polarization as the $m = 1$ images at a half-wave plate angle of 45° (also part of S_0). Recall that the polarization position angle is rotated by twice the half-wave plate angle.

The Stokes parameters I , Q , and U are determined from the S variables, and L is determined from Q and U by Eq. 3.39

$$\widehat{Q}(\boldsymbol{\vartheta}, p) = S_0(\boldsymbol{\vartheta}, p) - S_{90}(\boldsymbol{\vartheta}, p), \quad (3.44)$$

$$\widehat{U}(\boldsymbol{\vartheta}, p) = S_{45}(\boldsymbol{\vartheta}, p) - S_{135}(\boldsymbol{\vartheta}, p), \quad (3.45)$$

$$\begin{bmatrix} Q \\ U \end{bmatrix} = \begin{bmatrix} \cos 2\phi & \sin 2\phi \\ -\sin 2\phi & \cos 2\phi \end{bmatrix} \begin{bmatrix} \widehat{Q} \\ \widehat{U} \end{bmatrix}, \quad (3.46)$$

$$I(\boldsymbol{\vartheta}, p) = [S_0(\boldsymbol{\vartheta}, p) + S_{90}(\boldsymbol{\vartheta}, p) + S_{45}(\boldsymbol{\vartheta}, p) + S_{135}(\boldsymbol{\vartheta}, p)] / 2, \quad (3.47)$$

where ϕ is the angle between the 0° instrument polarization axis and N on the sky, which is generally constant throughout an observing run. \widehat{Q} and \widehat{U} are Stokes parameters defined only along the instrument polarization axes, while Q and U are defined on the sky. The arguments of $Q(\boldsymbol{\vartheta}, p)$ and $U(\boldsymbol{\vartheta}, p)$ were dropped in Eq. 3.46. Note that all of the Stokes parameters used here (including L) are not normalized to the total intensity I , but rather all of them are in units of intensity.

We will assume that the errors in Q and U are uncorrelated and Gaussian, with equal variance, σ_Q^2 ,

$$\sigma_Q^2(\boldsymbol{\vartheta}) = \left\langle \sigma_{(I)_{k,\text{pix}}}^2(\boldsymbol{\vartheta}, l, m) \right\rangle_{l,m} \quad (3.48)$$

with $\sigma_{(i)_{k,\text{pix}}}^2$ as defined in Eq. 3.21. The statistical distribution of measured L values follows a Rice distribution, which has important deviations from the simple (Gaussian) approximation when $L \sim \sigma_Q$.

An excellent overview of the statistics of L is given in Simmons and Stewart (1985), of which we summarize what is important to this analysis. In discussing the probability distributions, we normalize L to σ_Q , to give

$$L_n(\boldsymbol{\vartheta}, p) = L(\boldsymbol{\vartheta}, p)/\sigma_Q(\boldsymbol{\vartheta}), \quad (3.49)$$

and define the true, normalized linearly polarized flux L_0 (and the associated unnormalized true linearly polarized flux $L_{\text{true}} = L_0\sigma_Q$). The Rice distribution gives a probability density function of

$$p_{L_n}(L_n|L_0) = L_n e^{-(L_n^2 L_0^2)/2} I_0(L_n L_0), \quad (3.50)$$

where I_0 is a modified Bessel function of zero order (of the first kind), not to be confused with the background intensity (also I_0) defined earlier. Given an L_0 , the estimated L_n is

$$E(L_n) = \int_0^\infty L' p_{L_n}(L'|L_0) dL'. \quad (3.51)$$

The bias, $E(L_n) - L_0$, is shown in Fig. 3.9. This bias is large for small values of L_0 (*i.e.*, low signal-to-noise measurements).

We use a Wardle-Kronberg estimator (Wardle and Kronberg, 1974), which is also described in Simmons and Stewart (1985), to reduce the bias. The Wardle-Kronberg estimator, L_W , is defined as the L_0 which satisfies the equation

$$\partial p_{L_n}(L'|L_W)/\partial L'|_{L'=L_n} = 0, \quad (3.52)$$

i.e., L_W is the value of L_0 for which the observed L_n maximizes $p_{L_n}(L_n|L_0)$. The bias in the




Figure 3.9: Bias in polarization estimators. The X in L_X denotes the estimator used. The zero-order estimator is labeled L_n (the upper curve), the Wardle-Kronberg estimator is labeled L_W .

Wardle-Kronberg estimator is also shown in Fig. 3.9.

The Wardle-Kronberg estimator is not without bias, but it does reduce the bias significantly. It is easy to show, in fact, that it is impossible to create an estimator (which is continuous with respect to L_n) that is without bias near $L_0 = 0$. This can be seen by taking the derivative of $E(L_X) - L_0$ with respect to L_0 . If the bias is zero at $L_0 = 0$ and remains zero for $L_0 = \delta > 0$, then

$$\partial [E(L_X) - L_0] / \partial L_0 = 0, \quad (3.53)$$

for some estimator L_X . However, at $L_0 = 0$,

$$\begin{aligned} \partial E(L_X) / \partial L_0 |_{L_0=0} &= \partial / \partial L_0 \left[\int_0^\infty L_X L' e^{-(L'^2 + L_0^2)/2} I_0(L' L_0) dL' \right] \Big|_{L_0=0} \\ &= \int_0^\infty L_X L' e^{-(L'^2 + L_0^2)/2} [-L_0 I_0(L' L_0) + L' I_1(L' L_0)] dL' \Big|_{L_0=0} . \quad (3.54) \\ &= \int_0^\infty L_X L' e^{-(L'^2 + L_0^2)/2} [0 I_0(L' L_0) + L' 0] dL' \\ &= 0 \end{aligned}$$

To interchange the partial derivative and integration operators, it is sufficient that $L_X(L_n)$ be continuous. Therefore, $\partial [E(L_X) - L_0] / \partial L_0 = 1$, and so no estimator can be unbiased at more than one point near $L_0 = 0$. The choice of the Wardle-Kronberg estimator over other estimators, such as a maximum likelihood (ML) estimator, is specific to our desire to reduce the bias for moderately low signal-to-noise ratios. As discussed in Simmons and Stewart (1985), the ML estimator performs best at extremely low signal-to-noise ratios, $L_n < 1$, but we are less interested in those values, as they




Figure 3.10: Confidence bands for polarized flux. The bands, from darkest to lightest, are the 1-, 2-, and 3- σ confidence bands. The line in the middle of the 1- σ band is the Wardle-Kronberg estimator; at every L_0 , this is the L_n which maximizes $p_{L_n}(L_n|L_0)$, as defined by Eq. 3.52.

are of the least significance. We are interested in the values with marginal significance, $1 < L_n < 4$, where the Wardle-Kronberg estimator suffers less bias than the ML estimator.

Confidence intervals for L_0 can be generated from the probability density functions in Eq. 3.50. Given a confidence level, α , for every value of L_0 , we define an interval of L_n between the confidence limits L_{α_0} and L_{α_1} , such that

$$\alpha = \int_{L_{\alpha_0}}^{L_{\alpha_1}} p_{L_n}(L'|L_0)dL'. \quad (3.55)$$

For a given value of α , the confidence limits L_{α_0} and L_{α_1} are not uniquely specified. We arbitrarily choose to define L_{α_0} and L_{α_1} as the limits that give a minimum-length interval (*i.e.*, minimize $L_{\alpha_0} - L_{\alpha_1}$), satisfying Eq. 3.55. This is equivalent to insisting that $p_{L_n}(L_{\alpha_0}|L_0) = p_{L_n}(L_{\alpha_1}|L_0)$. The set of these intervals, for all values of L_0 , define a confidence band, shown in Fig. 3.10. Given a measured value of L_n , a confidence interval for L_0 can be found by the intersection of the boundary of the appropriate confidence band with a line at the measured L_n . The value of L_0 at which $L_{\alpha_1}(L_0) = L_n$ becomes $L_{0,lo}$, and the value of L_0 at which $L_{\alpha_0}(L_0) = L_n$ becomes $L_{0,hi}$. The confidence interval of L_0 is then $[L_{0,lo}, L_{0,hi}]$. In addition, Fig. 3.10 gives a graphical representation of the Wardle-Kronberg estimator, which gives the estimated L_0 for a measured L_n . Fig. 3.10 is constructed vertically, by solving for properties of L_n at a fixed L_0 , but is read horizontally, by specifying L_n and finding the values of L_0 that correspond to the estimator and confidence limits.

It should be noted that the confidence intervals as defined here are not unbiased intervals. For

a given $L_{0,\text{true}}$, the true value of L_0 , there exists another L_0 which satisfies

$$P[L_0 \in C(L_n)|L_{0,\text{true}}] > \alpha, \quad (3.56)$$

where $C(L_n) = [L_{0,\text{lo}}(L_n), L_{0,\text{hi}}(L_n)]$ is the confidence interval for a sampled L_n , and $P[\dots]$ is the probability, as determined over a random sampling of L_n . The fact that these intervals are biased implies that they are not Uniformly Most Accurate (UMA) intervals. The bias in the intervals, however, is not as serious as the bias in the estimators of L_0 . We do not attempt to reduce the bias in the confidence intervals.

All of the issues discussed here, concerning the bias in the estimators and the accurate analysis of confidence intervals when the errors are non-Gaussian, are of importance only when the signal-to-noise ratio is low. In the presence of low signal-to-noise ratios, a more important question is the significance of the measurements themselves, relative to a null hypothesis. At a given point in the field of view, we have 10 (we assume $n = 10$) phase-binned sets of Stokes parameters. The significance of individual nonzero values in these bins can be determined using a χ^2 test, $\chi^2 = \sum_p L_n^2(\boldsymbol{\vartheta}, p)$, distributed as χ^2 with 20 degrees of freedom. The 20 degrees of freedom from 10 samples are due to Eq. 3.39, where the null hypothesis is that both Q and U are zero, with Gaussian errors of variance σ_Q^2 .

The χ^2 test, as applied here, is designed to test independent, identically distributed Gaussian variates against a null hypothesis. However, the measures of $L(\boldsymbol{\vartheta}, p)$ are not independent, as they form a periodic time series. Because the ordering of the samples has an intrinsic meaning, the χ^2 test is not sufficient to test for the randomness of L . The most robust technique to test this variety of time series for a random ordering is a non-parametric serial correlation test, as described by Wald and Wolfowitz (1943). This technique defines a serial correlation statistic $R = \sum_{p=0}^9 L(\boldsymbol{\vartheta}, p)L(\boldsymbol{\vartheta}, p+1)$, where $L(\boldsymbol{\vartheta}, 10)$ is replaced by $L(\boldsymbol{\vartheta}, 0)$, because L is periodic. The test involves calculating R for every permutation of the phase bin indices. We define an index permutation function $p_i(j)$, where for each i , $p_i(j)$ takes on values of 0–9 for values of the index j between 0 and 9, with a different permutation for every i . For notational simplicity, we define $p_0(j) = j$, so that when $i = 0$, we recover the true ordering. In all, there are $10!$ values of i , representing all of the permutations of the 10 phase bins. A frequency distribution is then formed from the set of values of R_i , defined by

$$R_i = \sum_{j=0}^9 L(\boldsymbol{\vartheta}, p_i(j))L(\boldsymbol{\vartheta}, p_i(j+1)), \quad (3.57)$$

where $p_i(10)$ is replaced by $p_i(0)$. We form a cumulative probability distribution, $P_R(R)$, by integrating the frequency distribution, and base our test on the value of $P_R(R_0)$. We are able to perform this test exactly, without approximations to the permutations, for n (number of phase bins)

up to 10, which involves 10! permutations, or 4×10^6 values of R_i . This is within the comfortable limits of current computational power, but increasing n to 20, for instance, would preclude the direct computation of each permutations.

The Wald-Wolfowitz serial correlation test is a non-parametric test. This means that the test is not based on any assumptions about the distributions underlying the observed numbers. With no assumptions of normality or measurements of variance, the test is free of many potential systematic errors. The serial correlation, as defined by Eq. 3.57, is a measure of the tendency for values of L to be correlated between neighboring phase bins. The relationship between values of R_i calculated under different permutations shows how randomized samples (where only the ordering in time is random) would behave, and so the size of R_0 with respect to the set of R_i is a measure of the randomness of the observed ordering.

The combination of χ^2 and the Wald-Wolfowitz serial correlation tests assess both the randomness of the size of the samples and the randomness of the ordering of the samples in time. The results of these two tests give a true measure of the significance of a given sample set.

3.4 Statistics of Pulsed Fraction

Pulsation amplitudes in time series are commonly measured in one of two ways, using order statistics, *i.e.*, finding maxima and minima, or by fitting some function (usually sinusoidal) to the data. Each of these measurement techniques has associated systematic problems. Fitting the time series to a sinusoid tests only pulsations at a particular harmonic frequency. When the duty cycle of a pulse profile is low (for instance, if in 10 phase bins the intensity is constant for 9 bins and doubles in the 10th bin), the Fourier amplitudes are distributed over many harmonics, and the measurement of any one harmonic frequency will not adequately represent the pulsation amplitude. The proper treatment of order statistics, on the other hand, also involves assumptions about the true underlying pulse profile (*i.e.*, the profile in the absence of measurement noise). We choose to define pulsation amplitudes by the maximum and minimum of the pulse profile, defining a pulsed fraction h by

$$h = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \quad (3.58)$$

where the maximum and minimum are taken over all p for a given ϑ .

Measurement of h , as defined in Eq. 3.58, is tied intimately with the full treatment given in Section 3.2, where the intent was to determine the significance of pulsations. Here, on the other hand, we are interested in developing an estimator of and confidence intervals for h , rather than the test of significance of the pulsations. To evaluate the distribution of errors in h , we need to make some distinction between the errors and the estimate of the underlying pulse profile. In testing

the null hypothesis, the assumption is that the pulse profile is uniform, and that all deviations from uniformity are due to sampling error. In this case, the probability distributions involve no assumptions, because the intensities across all phase bins are identically distributed, and so the ordered intensity values are distributed according to the Beta distribution of the first kind. When evaluating h , however, when $h \neq 0$, the assumption is that the mean of the probability distribution varies across the phase bins.

The measured intensity in each phase bin is the best estimate of the mean of the intensity distribution in that phase bin (at least in normal and Poissonian distributions, which are the only distributions we consider for intensity distributions). However, using the measured intensities as the estimators of the true intensities introduces a bias. Any noise in the measurements will tend to increase the extrema of the pulse profile, resulting in an overestimate of h . A straightforward example of the bias is seen if the true pulse profile is in a low (uniform) state for 5 phase bins, and in a high (also uniform) state for the remaining 5 phase bins, *i.e.*, the pulse profile is a step function with duty cycle 1/2. In this case, it is intuitive, and simple to show, that sampling noise leads to an increase in $E(h)$, the expectation of the measured pulsed fraction.

The approach we take is to assume that the measured intensities across phase bins represent the best estimates of the relative deviations from a uniform pulse profile, but that the amplitude of the deviations, as measured by h , is unknown. This is equivalent to stating that the shape of the pulse profile is defined by the measured intensities, but not assuming that h is estimated directly from the measured extrema in Eq. 3.58. Taking the true pulsed fraction, h_0 , to be variable, for each h_0 , given the assumed true pulse profile shape, it is straightforward to compute the probability density function of measured pulsed fractions, h_{meas} , assuming that the measured intensities have either normal or Poissonian distributions. From these probability density functions, Wardle-Kronberg estimators, maximum likelihood (ML) estimators, and confidence bands can be generated by the same techniques as in Section 3.3. Fig. 3.11 shows the confidence bands and estimators for the pulsed fraction of the optical light from 4U 0142+61, as described in Chapter 4. In Chapter 4, we use the maximum likelihood estimator, $h_{\text{ML}}(h_{\text{meas}})$, defined by

$$\partial p_h(h_{\text{meas}}|h_0)/\partial h_0|_{h_0=h_{\text{ML}}} = 0. \quad (3.59)$$

This definition is similar to the Wardle-Kronberg estimator in Eq. 3.52, except the maximization is performed over values of h_0 , rather than over the measured quantities. In practice, the Wardle-Kronberg and ML estimators are indistinguishable, giving results that differ by less than 0.01 when applied to the data in Fig. 3.11. We therefore consider them to be interchangeable for these purposes, and we use the ML estimator in the paper presented in Chapter 4 simply because it requires less explanation.




Figure 3.11: Confidence bands and estimators for pulsed fraction. The bands, from darkest to lightest, are the 1-, 2-, and 3- σ confidence bands. The line in the middle of the 1- σ band is the Wardle-Kronberg estimator; at every h_0 , this is the h which maximizes $p_h(h|h_0)$, similar to the definition in Eq. 3.52. The maximum likelihood (ML) estimator is indistinguishable from the Wardle-Kronberg estimator in this case. The noise in the lines reflects the precision of the Monte-Carlo routines used to generate the probability density functions.

The estimators and confidence bands shown in Fig. 3.11 are specific to the dataset used, *i.e.*, the measured pulse profile of 4U 0142+61. For every measured pulse profile and error estimates on the individual intensity measurements, there is a new assumption regarding the true pulse shape. This is in contrast to the data in Fig. 3.10, which applies to any polarization measurements (properly normalized, with normally distributed errors).

For completeness, we also show the bias of the ML estimator in Fig. 3.12. The bias of the naïve estimator, $h = h_{\text{meas}}$, is also shown. For the dataset used in this analysis, $h_{\text{meas}} = 29.3\%$, giving $h_{\text{ML}} = 26.9\%$, with 1- σ confidence limits of 21.4% and 35.3%. Assuming that the true pulsed fraction $h_0 = 27\%$, the remaining bias due to the ML estimator is 0.2%, while the bias due to the zero-order estimator would have been 2.8%. The residual bias, 0.2%, is insignificant compared to the size of the confidence interval, and so we feel no need to improve on the estimator.

The technique we have chosen to use to estimate the pulsed fraction and its confidence intervals, while relatively straightforward computationally, relies on a very specific assumption regarding the underlying pulse profiles. We feel that these assumptions are reasonable, however, they must be carefully sized up relative to other techniques of measurement for similar properties. The specifics of this technique become important, as with other statistical techniques, when the signal-to-noise ratio is low. This is the case for the 0.5–1.0 keV X-ray pulse profile of 4U 0142+61, as described in Chapter 4.




Figure 3.12: Bias in pulsed estimators. The X in h_X denotes the estimator used. The zero-order estimator is labeled h_{meas} (the upper curve), the Maximum Likelihood (ML) estimator is labeled h_{ML} .

Chapter 4

4U 0142+61

This paper was accepted by Nature, for publication on 2002 May 30. The paper is included here in its original form, with the exception of the bibliographic references, numbering, and figure font sizes. Because I lack both the prior knowledge and the willpower to overcome the formidable L^AT_EX issues associated with maintaining separate referencing styles in different sections of this thesis, the citations in this chapter are in (Author, year) format, and the bibliography for this paper is merged with the overall thesis bibliography. The numbering of equations and figures conforms to the numbering scheme of the thesis. The font in the figures appears large, because the figures were generated with the understanding that the figures would be printed in two-column format, but here they appear in single-column format.

4.1 Nature paper on 4U 0142+61

Optical pulsations from the anomalous X-ray pulsar

4U0142+61

B. Kern & C. Martin

*Division of Physics, Mathematics, and Astronomy, California Institute of Technology, MS 405-47,
Pasadena, California 91125, USA*

Anomalous X-ray pulsars (Mereghetti, 2001) (AXPs) differ from ordinary radio pulsars in that their X-ray luminosity is orders of magnitude greater than their rate of rotational energy loss, and so they require an additional energy source. One possibility is that AXPs are highly magnetized neutron stars (Duncan and Thompson, 1992)—or ‘magnetars’—having surface magnetic fields greater than 10^{14} G. This would make them

similar to the soft γ -ray repeaters (SGRs) (Kouveliotou *et al.*, 1998), but alternative models that do not require extreme magnetic fields also exist. An optical counterpart to the AXP 4U 0142+61 was recently discovered (Hulleman *et al.*, 2000a), consistent with emission from a magnetar, but also from a magnetized hot white dwarf (Paczyński, 1990), or an accreting isolated (Wilson *et al.*, 1999) neutron star (Marsden *et al.*, 2001). Here we report the detection of optical pulsations from 4U 0142+61. The pulsed fraction of optical light (27 percent) is five to ten times greater than that of soft X-rays, from which we conclude that 4U 0142+61 is a magnetar. Although this establishes a direct relationship between AXPs and the soft γ -ray repeaters, the evolutionary connection between AXPs, SGRs and radio pulsars remains controversial.

We observed the field around the faint ($R = 25$) optical counterpart proposed by Hulleman *et al.* (2000a) for the 8.7-second Anomalous X-ray Pulsar (AXP) 4U 0142+61 (Figure 4.1) on 15–16 November 2001 at the Palomar Observatory Hale 5-m telescope, using a novel phase-binning CCD. The average seeing was 1.4 arcsec, in non-photometric conditions. The light curves of every point in the field-of-view (200×10.4 arcsec) are binned on-chip into 10 phase bins. The imaging properties of this technique allow us to ensure that the observed time variability of 4U 0142+61 is not instrumental in origin. In Figure 4.2 we show that the variability (at the X-ray period of 4U 0142+61) everywhere but the optical counterpart is consistent with the noise. The light curve of 4U 0142+61 is shown in Figure 4.3.

The remarkable aspect of the optical light curve is that its modulation amplitude is very large compared to that of the X-ray light curves. We define pulsed fraction h as

$$h = \frac{F_{\max} - F_{\min}}{F_{\max} + F_{\min}}, \quad (4.1)$$

where F_{\max} and F_{\min} are the maximum and minimum flux in any phase bin. We use a maximum likelihood estimator for h to reduce the systematic bias caused by random errors. We find $h_{\text{opt}} = 27_{-6}^{+8}\%$ ($1\text{-}\sigma$ errors), with a probability of spurious detection of 6×10^{-5} (one-tailed). Due to the finite temporal resolution of the optical data, this estimate is a lower limit to the true pulsed fraction.

The X-ray spectrum of 4U 0142+61 is generally interpreted as an absorbed sum of power-law and blackbody components (White *et al.*, 1996), with a power-law index ~ 3.7 and temperature 0.4 keV. With this ultrasoft power-law index, the observed luminosity, corrected for absorption, is dominated by soft X-rays. Between 0.5–10.0 keV, the 0.5–1.0, 1.0–2.0, and 2.0–10.0 keV X-rays provide 65%, 25%, and 10% of the unabsorbed flux. To compare the optical and soft X-ray pulse profiles, we use archival *Chandra* ACIS-S data (Figure 4.3), which provide the most sensitivity to 0.5–1.0 keV X-rays of any observations to date. Following Equation 4.1 to define $h_{0.5}$ and $h_{1.0}$ as the pulsed fractions for 0.5–1.0 and 1.0–2.0 keV X-rays, respectively, we find $h_{0.5} = 2.9_{-2.9}^{+1.9}\%$ and

$h_{1.0} = 5.0 \pm 0.5\%$. The quoted errors are $1\text{-}\sigma$ errors, which show that $h_{0.5}$ is consistent with zero. For these estimates, the X-ray data were binned into 10 phase bins to match the time resolution of the optical data. Binning the 1.0–2.0 keV X-rays into 20 phase bins instead of 10 increases the measured pulsed fraction by 2%.

We first consider the optical pulsations in the framework of the magnetar model. There is compelling evidence that the soft gamma-ray repeaters (SGRs) are magnetars (Kouveliotou *et al.*, 1998), but the SGRs are heavily absorbed at visible wavelengths, and there are no detailed models predicting the optical emission from magnetospheres of magnetars. We turn to the closest relatives of magnetars, the young ($10^3\text{--}10^4$ yr) isolated radio pulsars with measured optical pulsations, *i.e.*, Crab (Romani *et al.*, 2001; Tennant *et al.*, 2001), Vela (Helfand *et al.*, 2001; Manchester *et al.*, 1980) and PSR B0540-69 (Gotthelf and Wang, 2000; Boyd *et al.*, 1995). The optical light curves for these pulsars are similar in morphology to their X-ray light curves, and they have $L_X/L_{\text{opt}} \sim 10^3\text{--}10^4$, similar to 4U 0142+61. At present, the data is consistent with the optical emission that might be expected from magnetars.

We next consider the possibility that 4U 0142+61 is a neutron star (NS) fed by accretion, without a stellar companion (Wilson *et al.*, 1999). We assume that beamed X-rays are emitted at the NS surface, and that rotation of the NS leads to the observed X-ray pulsations. Optical pulsations at the X-ray pulsation frequency can arise only from thermal reprocessing of X-rays illuminating a disk. In binary systems, where the X-rays illuminate a companion star, the pulsed fraction of the reprocessed light cannot exceed the pulsed fraction of the incident X-rays (Chester, 1979). The reprocessing mechanism is similar in accretion disks (Pedersen and *et al.*, 1982), giving

$$h_{\text{opt}} \leq h_X, \quad (4.2)$$

where h_X is the pulsed fraction of the X-rays illuminating the disk. Assuming that the softest X-rays (0.5–1.0 keV) dominate the energetics of reprocessing and that the X-rays illuminating the disk are the same as the observed X-rays (*i.e.*, $h_X = h_{0.5}$), this is violated, because $h_{\text{opt}} \sim 10h_{0.5}$.

To reconcile Equation 4.2 with the observations, two conditions must be satisfied; $h_X > h_{0.5}$, and the disk axis and the NS rotation axis must be misaligned, so that the X-ray illumination of the disk is not azimuthally symmetric (or else there would be no optical pulsations). The first condition, $h_X > h_{0.5}$, could be satisfied if the X-rays were more strongly beamed toward the disk than toward the observer. This is unlikely to be significant, because general relativistic effects spread out beamed X-rays from the NS surface (Pechenick *et al.*, 1983). For a $1.4 M_\odot$ NS with a 10-km radius, the flux ratio between the center of the X-ray beam and a line-of-sight 90° away will be < 4 (Perna and Hernquist, 2000). An alternate reason h_X may be greater than $h_{0.5}$ is that the pulsed fraction of X-rays with energies < 0.5 keV (unobservable because of absorption) may be higher than $h_{0.5}$.

While there is no observational evidence for or against this possibility, we note that the trend of pulsed fraction with energy is to lower pulsed fractions with lower energies (White *et al.*, 1996).

All of the preceding NS discussions relate the X-rays to the reprocessed optical flux. The *VRI* optical fluxes measured by Hulleman *et al.* (2000a) limit the possible size of an accretion disk to $R_{\text{out}} < 0.05R_{\odot}$. For such a small disk, the disk heating is expected to be dominated by viscous dissipation rather than X-ray irradiation (Hulleman *et al.*, 2000b). The optical flux is then dominated by a steady component due to the viscous dissipation. For these reasons, we conclude that our measurement of optical pulsations is inconsistent with the NS accretion model.

The third model we consider is that 4U 0142+61 is a $B = 5 \times 10^8$ G, $T = 4 \times 10^5$ K white dwarf, possibly the result of a double-degenerate merger (Hulleman *et al.*, 2000a). The white dwarf with the most similar physical properties is RE J0317-853, with $B = 5 \times 10^8$ G, $T = 5 \times 10^4$ K, $P = 725$ s (Barstow *et al.*, 1995; Ferrario *et al.*, 1997). RE J0317-853 shows optical pulsations with $h_{\text{opt}} \sim 10\%$ (Barstow *et al.*, 1995) and 100-Å EUV pulsations with $h_{\text{EUV}} \sim 20\%$ (Ferrario *et al.*, 1997). Two other highly magnetized, rotating white dwarfs have shown photometric optical variability—PG 1031+234 with $h \sim 20\%$, 15% and 10% for U , B and V (Piirola and Reiz, 1992), and Feige 7 with $h \sim 5\%$ in V (Achilleos *et al.*, 1992). Surface variations in magnetic field strength in these objects change the continuum opacity (magnetic dichroism), and may cause variations in surface temperature. In addition, the surface abundances may be nonuniform. Models have successfully explained the photometric variability of these white dwarfs by adjusting their geometric and chemical parameters.

We know of no model atmospheres that accommodate both the high temperatures and large magnetic fields proposed for 4U 0142+61. Model atmospheres with temperatures of 10^5 – 10^6 K (but no magnetic fields) have X-ray spectra that are very sensitive to both temperature and chemical abundance. Assuming that the observed X-rays are related to the EUV / soft X-rays emitted near the peak of the blackbody spectrum, the X-rays scale as T^4 , while the optical luminosity scales as T^1 , giving $h_X > h_{\text{opt}}$. Even in a pure hydrogen atmosphere, there is a temperature-dependent photoionization opacity that will cause more variability in the X-ray fluxes (Heise *et al.*, 1994). Ionization edges of highly ionized He, C, N, O, and Ne are very sensitive to changes in temperature, and line blanketing by these elements results in EUV / soft X-ray spectra that vary by orders of magnitude for temperature differences of a factor of 2 (Rauch, 1996). We conclude that the mechanisms that have explained optical photometric variability in other white dwarfs are likely to produce $h_{\text{opt}} < h_X$ (as in RE J0317-853, with $h_{\text{opt}} = h_{\text{EUV}}/2$).

The X-ray spectrum of 4U 0142+61 is difficult to explain in the context of a white dwarf. If the blackbody component is to be believed, then the emission region is confined to a radius of 12 km (White *et al.*, 1996) (assuming a distance of 5 kpc). Assuming $M \sim 1.3M_{\odot}$, $R \sim 0.01R_{\odot}$, the surface gravity is $\log g \sim 9$, which means that, assuming the blackbody radiation is emitted at the

surface, the flux in the emitting region is $\sim 16\times$ the Eddington flux.

The evidence that 4U 0142+61 is a magnetar is strengthened by the disagreement between the accretion and white dwarf models and the optical observations. Optical observations of magnetars open up new avenues for exploring exotic effects, such as QED in strong magnetic fields resulting in magnetic photon splitting (Baring and Harding, 2001) and polarization of the vacuum index of refraction (Shaviv *et al.*, 1999). High-quality pulse-phased photopolarimetry is possible with 8–10 m telescopes, opening a new field of physical exploration. Future multi-color optical and infrared phase, morphology, and polarization observations, when combined with X-ray data, will provide crucial diagnostics of magnetar physics in much the same way that multi-wavelength pulse-phased observations have proved invaluable for expanding our understanding of radio pulsars. In addition to the physical properties of magnetars themselves, AXPs may share an evolutionary connection between SGRs and radio pulsars, and could represent a significant fraction of all stellar remnants. Establishing 4U 0142+61 as a magnetar is the first step in this investigation.

Acknowledgements

We thank S. Kaye, the Palomar Observatory staff, D. Kaplan, S. Kulkarni, and R. Rutledge for invaluable assistance. This work was supported by grants from the National Science Foundation and the California Institute of Technology.




Figure 4.1: Time-averaged optical image of field around 4U 0142+61. This image is 170×16 -pixels (68×6.4 -arcsec), displayed with linear (top panel) and logarithmic (bottom panel) scaling. The phase-binning CCD has 512×512 pixels, masked by a 500×26 -pixel (200×10.4 -arcsec) slit. Intensity is binned on-chip into 10 phase intervals using a periodic frame transfer of the CCD. Every $1/10$ of the period of 4U 0142+61 (known *a priori*), the accumulated image is transferred away from the slit onto a region of the CCD used for storage, but no pixels are read out of the CCD. At the end of 10 transfers (*i.e.*, 1 period of 4U 0142+61), which alternate between shifting the accumulated charge both up and down on the CCD, the image accumulated during the first phase interval is shifted back into the CCD region illuminated by the slit, and the image is re-exposed. At any time, there are 10 accumulated images residing on the CCD, one of which is being exposed. This exposure pattern is repeated for 95 s (11 cycles of 4U 0142+61), at which point the shutter is closed and the entire array is read out. Each readout delivers 10 phase-binned images, each of which has been exposed for 9.5 seconds, long enough for the read noise of the CCD to be small compared to the Poisson noise of the sky background. The total exposure time in two nights was 60,000 s. No filter was used, giving a 400–1000 nm spectral response. The images shown here are sums of all 10 phase-binned images. The numbering of stars follows the convention of Reid *et al.* (1980) and Coe and Pightling (1998), with 4U 0142+61 labeled star A, as in Hulleman *et al.* (2000a).




Figure 4.2: Map of pulsation amplitudes. At each point of the field in Figure 4.1, the square of the intensity fluctuations divided by the variance is summed over all 10 phase bins. Each point in this map shows the probability of chance occurrence of the residuals (plotted logarithmically). In a field this large (68×6.4 arcsec) we expect residuals as large as the spurious peak (2 arcsec E of star 3) 24% of the time. The residuals in 4U 0142+61 have a chance probability of 6×10^{-5} .




Figure 4.3: Optical and X-ray pulse profiles of 4U 0142+61. Each panel is normalized to the mean flux at that energy, with $1\text{-}\sigma$ error bars. The optical flux is measured in a 4×4 -pixel (1.6×1.6 -arcsec) aperture centered on 4U 0142+61. The timing information used for these exposures was taken from the “second span” entry in Table 2 of Gavriil and Kaspi (2002), with an arbitrary choice of absolute zero-phase. These parameters are: $\nu = 0.1150969336$ Hz, $\dot{\nu} = -2.687 \times 10^{-14}$ Hz/s, on MJD 51704.0 TDB, with zero-phase barycentric time-of-arrival MJD 51704.00006816 TDB. The timing of the optical observations is controlled by a GPS receiver, that delivers start and stop pulses accurate to $1\ \mu\text{s}$. The Crab Pulsar was observed each night to verify the accuracy of the timing. The middle and lower panels are pulse profiles of 1.0–2.0 keV and 0.5–1.0 keV X-rays, respectively, with the intensity axis zoomed by a factor of 4 relative to the optical pulse profile panel. The X-ray pulse profiles are 5 ks of *Chandra* ACIS-S continuously-clocked data taken on 21 May 2000, folded onto the same timing solution as the optical data. The background level is $< 10^{-3}$ the total flux level in both plots. We do not correct the continuously-clocked data for dither, which results in periodic timing errors of 50 ms ($0.006P$). The absolute accuracy of *Chandra* continuously-clocked timing has been estimated at < 20 ms. The Gavriil and Kaspi (2002) timing solution is from *RXTE* data up to June 2001, 5 months before the Palomar observations. Because the X-ray and optical data are not contemporaneous, and the optical observations were taken outside the span of the timing solution used, the relative phases are unconstrained.

Chapter 5

PSR B0656+14

This paper is, at the time of this writing, in a near-final (presumably) stage of preparation, for submission to the *Astrophysical Journal*. It is included here, in its present state, as a stand-alone body of work. As with Chapter 4, the references are merged with the thesis bibliography.

Apart from the paper, an effort has begun to calculate detailed outer gap models, which may provide very specific predictions with which we can compare our observations. A synopsis of these additional calculations is presented after the paper, in Section 5.2.

5.1 *Astrophysical Journal* submission on PSR B0656+14

Optical Pulse-Phased Photopolarimetry of PSR B0656+14

B. Kern, C. Martin, B. Mazin

California Institute of Technology, Pasadena, CA 91125

J. Halpern

Columbia University, New York, NY 10027

Abstract

We have observed the optical pulse profile of PSR B0656+14 in 10 phase bins at a high signal-to-noise ratio, and have measured the linear polarization profile over 30% of the pulsar period with some significance. The pulse profile is double-peaked, with a bridge of emission between the two peaks, similar to some gamma-ray profiles observed in other pulsars. There is no detectable unpulsed flux, to a $1\text{-}\sigma$ limit of 16% of the pulse-averaged flux. The emission in the bridge is highly polarized,

with a position angle sweep that agrees well with the prediction of the Rotating Vector Model as determined from radio polarization observations. We have not computed outer-gap emission models to compare with the optical observations, but polar cap emission models are broadly consistent with the data.

5.1.1 Introduction

Radio pulsars, as their name implies, are well studied at radio wavelengths. High time-resolution pulse profiles and polarization profiles have allowed determinations of the relative orientations of the rotation axis, magnetic axis, and observer line-of-sight (LOS), based on variations of the Rotating Vector Model (RVM), originally proposed for the Vela pulsar by Radhakrishnan and Cooke (1969). It is commonly believed that the radio emission is generated near the surface of the neutron star (Kijak and Gil, 1997), which simplifies the interpretation of the radio data. Studies by Lyne and Manchester (1988), Rankin (1990), and Everett and Weisberg (2001) use different methods and assumptions to determine these angles, with a rough agreement established between the methods.

High-energy (infrared to gamma-ray) magnetospheric emission from isolated radio pulsars presents a more complicated theoretical picture. Two categories of theories have been presented to explain the location of the emission region for high-energy photons. Polar cap models (Daugherty and Harding, 1994, 1996) claim that the emission region is close to the neutron star surface, near the emission region for the observed radio waves. Outer gap models (Romani and Yadigaroglu, 1995; Romani, 1996) place the high-energy emission regions farther out in the magnetosphere, up to $\sim 30\%$ of the light-cylinder radius. Part of the problem in distinguishing between these two models lies in the small number of pulsars observed at high energies. Pulsed gamma rays have been observed in seven radio pulsars (counting Geminga as a radio pulsar for this argument), with possible detections in another three (Thompson, 2001). Many more pulsars have been observed in X-rays, with somewhat inconclusive impact on the question of emission model (Becker and Trümper, 1997). Optical pulsations have been observed in only three radio pulsars, with weak detections in two more.

PSR B0656+14 is a middle-aged pulsar, with characteristic age 1.1×10^5 years, considered a “cooling neutron star” because its soft X-ray emission is believed to come from the surface of the neutron star (Becker and Trümper, 1997). There is a tentative detection of a gamma-ray pulse (Ramanamurthy *et al.*, 1996), as well as a claim of optical pulsations (Shearer *et al.*, 1997). Optically, PSR B0656+14 is the second-brightest radio pulsar in the Northern sky ($V = 25$), after the Crab pulsar. Optical polarization measurements hold the best promise to constrain the models for high-energy emission regions, until such time that X-ray and gamma-ray polarimetry becomes feasible.

5.1.2 Observations

We observed PSR B0656+14 on 2000 December 20–21 at the Palomar 5 m telescope, with a Phase-Binning CCD Camera. Over the two nights, we obtained 38,000 s of integration time, with an average FWHM of 1.3 arcsec under non-photometric conditions. The phase-binning CCD used is a 512×512 pixel back-illuminated CCD, masked by a 100×26 pixel (50×13 arcsec) slit. Intensity is binned on-chip into 10 phase intervals using a periodic frame transfer of the CCD. Between the slit and the CCD, an achromatic half-wave plate and broadband polarizing beam splitters form an imaging two-channel polarimeter, measuring two orthogonal linear polarizations simultaneously. The two polarized images are arranged side-by-side on the CCD in two columns. Every 1/10 of the period of PSR B0656+14 (known *a priori*), the accumulated images (both polarizations) are transferred away from the slit onto a region of the CCD used for storage, but no pixels are read out of the CCD. At the end of 10 transfers (*i.e.*, 1 period of PSR B0656+14), which alternate between shifting the accumulated charge both up and down on the CCD, the images accumulated during the first phase interval are shifted back into the CCD region illuminated by the slit, and the images are re-illuminated. At any time, there are 20 accumulated images residing on the CCD (2 columns of linearly polarized images, 10 rows of phase-binned images), two of which are being illuminated.

This illumination pattern is repeated for 120 s (311 cycles of PSR B0656+14), at which point the shutter is closed and the entire array is read out. Each 120 s exposure delivers 20 phase-binned images, each of which has been illuminated for 12 seconds, long enough for the read noise of the CCD ($< 10 e^-$ rms) to be small compared to the Poisson noise of the sky background. Timing of the exposures is coordinated using a GPS receiver and time-code generator, which delivers a disciplined 10 MHz oscillator signal and start and stop signals accurate to $1 \mu\text{s}$. Between 120 s exposures, the achromatic half-wave plate is rotated by 22.5° , rotating the polarization position angles the detector observes, as projected on the sky, by 45° . On each night, we observed the Crab Pulsar to calibrate our timing and check our polarization optics. We used a Schott BG38 colored-glass filter, which together with the quantum efficiency of our CCD, defines a bandpass of approximately 400–600 nm.

Because the folding of the intensity to the pulsar period is done in real time on the CCD, an accurate pulse frequency, as observed at the telescope, must be known *a priori*. We used an ephemeris for PSR B0656+14 from A. Lyne (2000, private communication), with barycentric $f = 2.5980738005954 \text{ s}^{-1}$, $\dot{f} = -3.71214 \times 10^{-13} \text{ s}^{-2}$, and $\ddot{f} = 8.33 \times 10^{-25} \text{ s}^{-3}$ at 2451687.5 JD TDB, and a radio pulse geocentric time-of-arrival 2451687.500000997 JD UTC. This was used to predict a topocentric pulse arrival time at Palomar and an observed pulse frequency for each individual exposure. The JPL DE200 solar-system ephemeris was used to compute the barycenter-geocenter and geocenter-topocenter corrections.

Care was taken to eliminate systematic effects in our measurements. The most important aspect

of this technique is that the data are a time-series of images, which contain many stars in addition to PSR B0656+14. The light curve of every object in the field can be extracted, to verify that no spurious pulsation signals come from the instrument or observing scheme. Another important aspect of this technique is that in a given exposure, the light at each point in the field falls on the same CCD pixels for every phase bin (*i.e.*, the slit does not move with respect to the CCD pixels), so that quantum efficiency and throughput fluctuations are not imprinted differentially on the light curves. Charge transfer effects, on the other hand, can introduce light curve variations, because the paths the phase bins follow as they are shifted up and down on the CCD are different. To eliminate this effect, successive exposures are started at different phases, so that by properly coadding a 10-exposure set, each phase bin will have equal contributions from each path up and down the CCD. The rotation of the half-wave plate between exposures ensures that the linear polarization measurements alternate position on the CCD. The dual-channel configuration ensures that throughput variations are not interpreted as polarization signals.

5.1.3 Data analysis

The largest systematic signal to be removed from the data is due to charge traps in the CCD. Because the number of parallel line shifts during an exposure is large ($\sim 150,000$), even small charge traps ($\ll 1 e^-$) generate large noise signals. The pattern of charge shifting caused by these traps is slightly illumination-dependent, which necessitates calibration under similar illumination as the observations. This was done by dithering the pulsar position in the field-of-view and using blank regions of the field to calibrate both additive and multiplicative nonuniformities across the field-of-view. After this self-calibrating flatfielding and bias removal, the blank regions of the sky show variances approximately 30% higher than the expected Poisson variances from the sky brightness, an acceptable level of noise.

Individual exposures are coadded to preserve the phase and polarization information, resulting in a 4×10 array of images, at 4 position angles, in 10 phase bins. The integrated image, formed by adding all of the average images to eliminate phase and polarization information, is shown in Figure 5.1. The total intensity images, formed by combining sets of 4 polarized images but keeping the phase bins separated, are shown in Figure 5.2.

To extract the flux of PSR B0656+14 in each averaged image, we use a 3×3 pixel (1.5 arcsec \times 1.5 arcsec) aperture. At any point in the field, we denote the measured flux by $S_{\chi,i}(x,y)$, where χ is the instrumental polarization position angle, i is the phase bin number (ranging from 0 to 9), and x and y are the position in the field-of-view. The instrumental position angle, χ , is measured with respect to an instrumental axis, which is rotated with respect to N on the sky. We denote the linearly polarized Stokes parameters measured with respect to the instrument axes $\hat{Q}_i(x,y)$ and $\hat{U}_i(x,y)$, and the rotation of these measurements to refer to N on the sky $Q_i(x,y)$ and $U_i(x,y)$. We

define

$$\widehat{Q}_i = S_{0,i} - S_{90,i} \quad (5.1)$$

$$\widehat{U}_i = S_{45,i} - S_{135,i} \quad (5.2)$$

$$\begin{bmatrix} Q_i \\ U_i \end{bmatrix} = \begin{bmatrix} \cos 2\phi & \sin 2\phi \\ -\sin 2\phi & \cos 2\phi \end{bmatrix} \begin{bmatrix} \widehat{Q}_i \\ \widehat{U}_i \end{bmatrix} \quad (5.3)$$

$$I_i = (S_{0,i} + S_{45,i} + S_{90,i} + S_{135,i})/2, \quad (5.4)$$

with I_i the total intensity, and ϕ the angle between the instrument and sky orientations. We do not measure V , the circularly polarized Stokes parameter. We do not normalize the Stokes parameters to the total intensity, leaving all measurements in flux units.

Background Subtraction

A nearby extended object, which is visible in Hubble Space Telescope WF/PC2 and NICMOS images of the field around PSR B0656+14 (Koptsevich *et al.*, 2001), when observed from the ground contributes some light to the photometric aperture we use for PSR B0656+14. This object is labeled o2 in Koptsevich *et al.* (2001). The *HST* WF/PC2 images were observed in the F555W filter, a similar bandpass to that of our ground-based images. After standard image reduction to remove hot/warm pixels and cosmic rays, we perform aperture photometry on the pulsar in the WF/PC2 images to determine the intensity of the pulsar alone. We then convolve the WF/PC2 images with a gaussian PSF whose FWHM equals our average ground-based seeing FWHM (1.3 arcsec), and determine how much light falls in the photometric aperture we used for the ground-based images. By referencing both the convolved HST images and the ground-based images to a “blank-sky” region of the images, we can determine the contribution of the pulsar alone, as well as that of the pulsar and extended object, to the light in our photometric aperture. Because the bandpasses are similar, we perform no color corrections between the HST WF/PC2 images and our ground-based images. We find that in our 1.5 arcsec square photometric aperture, the extended object contributes $39\% \pm 7\%$ of the pulse-averaged light in our photometric aperture.

The intensity profile I_i is shown in Figure 5.3, which has been background-subtracted as described above. The phase of our observations are referenced to the peak of the radio pulse, that arrives at phase 0.0. The possible gamma-ray pulse (Ramanamurthy *et al.*, 1996) peaks at phase 0.2. The soft X-ray pulse broadly peaks near phase 0.85, with a minimum near phase 0.3 (Marshall and Schulz, 2002).

The minimum of the background-subtracted pulse profile, I_{\min}/\bar{I} , is -0.05, normalized to the pulse-averaged flux \bar{I} . While negative fractions are unphysical, the measurement error in each bin ($\sigma_I/\bar{I} = 0.24$) makes it very likely that negative values are measured. Using all of the bins in a

joint probability measurement, we find a $1\text{-}\sigma$ upper limit (two-tailed) to the true unpulsed flux to be $I_{\text{unpulsed}}/\bar{I} < 0.16$, accounting for both measurement error in each bin and the uncertainty in the background level. The limited temporal resolution of these measurements causes I_{unpulsed} to be overestimated, so these numbers should be considered true upper limits.

The pulse profile we measure for PSR B0656+14 is very different than that measured by Shearer *et al.* (1997). Shearer *et al.* measure a pulse with a broad single peak, reaching a maximum at phase 0.2, coincident with our Peak 1, but they find a minimum near phase 0.8, coincident with our Peak 2. Basing some estimates on Figure 3 of Shearer *et al.* (1997), the quoted unpulsed flux limit must be understated, because the background flux uncertainty is $\sim 40\%$ of the average net flux, and the error on the flux in each phase bin is $\sim 80\%$ of the average net flux. Combining the flux in their four faintest phase bins, and considering the uncertainty in the background, the unpulsed flux must have an error greater than 40% , and yet the quoted $1\text{-}\sigma$ upper limit (8×10^{-31}) is only 20% of the net flux (3.9×10^{-30}). More to the point, a simple test of the null hypothesis (that the flux is unpulsed) gives only $\chi^2 = 15$ for 10 degrees of freedom (cumulative probability 0.87). The disagreement between the pulse profile reported here and that reported in Shearer *et al.* (1997) raises some doubt as to the significance of the optical pulse profile for Geminga reported by the same group (Shearer *et al.*, 1998), from data taken concurrently with their PSR B0656+14 data.

Linearly Polarized Flux

We compute the linearly polarized flux, L_i , from the Stokes parameters Q_i and U_i ,

$$L_i = \sqrt{Q_i^2 + U_i^2}, \quad (5.5)$$

$$\theta_i = \tan^{-1}(U_i/Q_i)/2. \quad (5.6)$$

Here, θ_i is the polarization position angle, defined as the direction of vibration of the electric field measured E from N on the sky. The equation for L_i uses a biased estimator, because any noise will, on average, result in a systematic overestimate of the linearly polarized flux. Following the prescriptions in Simmons and Stewart (1985), we use a Wardle and Kronberg (Wardle and Kronberg, 1974) estimator, which is equivalent to a maximum likelihood estimator. The analysis we perform on L_i uses estimates derived from the Wardle and Kronberg estimator. A plot of the linearly polarized flux is shown in Figure 5.4. From phase 0.4 – 0.7, the L_i values differ from zero at the $2\text{--}3\text{-}\sigma$ level.

We test the significance of the polarized flux measurements in two ways. If the uncertainties in Q_i and U_i are distributed as independent gaussian variables with variance σ^2 , under the null hypothesis that there is no polarized flux, $\chi^2 = \sum_{i=0}^9 P_i^2/\sigma^2$ should be distributed as χ^2 with 20 degrees of freedom. We find $\chi^2 = 50$, which has a cumulative probability of 94.5%.

Since these polarization measurements form a time series, in which the ordering of the mea-

surements is significant, we also address the question of the randomness of the distribution. We test the randomness of the time series with the Wald-Wolfowitz test of serial correlation (Wald and Wolfowitz, 1943). Because the number of data points (10) is small, we can compute the serial correlation

$$R = \sum_{i=0}^9 L_i L_{i+1}, \quad (5.7)$$

where L_{10} is replaced by L_0 , since the measured values are periodic. We then compute R for every permutation of indices i (3.6×10^6 permutations), and form a distribution function from the collection of values of R at different permutations. This forms a non-parametric test, which involves no assumptions of normality or variance, and which has no dependence at all on the time-independent distribution of the L_i being measured or on the total intensity I_i . The ordering of the polarization values measured violates the null hypothesis, that the numbers are randomly chosen, at the 97% level. This number represents the significance of the three most significant polarization flux measurements being clustered in phase.

The combination of these two tests, which are independent of one another, lend credibility to the measured polarization flux values. In addition, the position angles are correlated in the three most significant bins. The angles in phases 0.4–0.7 are $\theta_4 = 77^\circ \pm 11^\circ$, $\theta_5 = 89^\circ \pm 9^\circ$, and $\theta_6 = 137^\circ \pm 17^\circ$.

If the position angle changes by 90° on timescales comparable to the bin width, the measured linear polarization will be reduced. Without a specific model of the behavior of the position angle with phase, the measured linear polarization must be considered a lower limit. The best estimate is then that the flux is $\sim 100\%$ polarized from phase 0.4–0.7, and mostly (linearly) unpolarized at other phases.

5.1.4 Discussion

The optical light curve is double-peaked, unlike either the radio (Gould and Lyne, 1998) or the X-ray (Marshall and Schulz, 2002) curves, which are both single-peaked. We separate the optical light curve (see Fig. 5.3) into 4 phase intervals. We define Peak 1 from phase 0.2–0.3, Bridge emission from phase 0.3–0.8, Peak 2 from 0.8–0.9, and Off-pulse from 0.9–1.2.

The soft X-rays are commonly interpreted as thermal emission from the surface of the neutron star itself, with a 10% modulation. If the optical light were due to the same thermal component, it too would show a small pulsed fraction. Assuming that the thermal flux is equal to the unpulsed flux, we find a $1\text{-}\sigma$ upper limit to the thermal flux of 16% of the flux in the observed bandpass. The bulk of the observed optical emission is then of nonthermal origin, arising from the pulsar magnetosphere rather than the surface of the neutron star. The high pulsed fraction of the optical light rules out the notion that the optical emission could be due to a disk of material (i.e., fallback material from a

supernova explosion), as it would not be pulsed at the neutron star rotation frequency (Perna *et al.*, 2000).

The radio pulse peaks at phase 0.0, with FWHM 0.04–0.07 P at frequencies from 0.2–1.6 GHz (Gould and Lyne, 1998). The Rotating Vector Model (RVM) allows radio polarization profiles to determine the geometry of the magnetic poles relative to the rotation axis and the observer line-of-sight (LOS). In the case of PSR B0656+14, the deepest data (at 1.4 GHz) gives a measurement of $\alpha = 29^\circ \pm 23^\circ$, $\beta = 8.9^\circ \pm 6.1^\circ$, where α is the angle between the rotation and magnetic axes, and β is the angle of the closest approach of the observer LOS to the magnetic axis (Everett and Weisberg, 2001). The uncertainty in these two angles is highly (positively) correlated. Two earlier studies investigated PSR B0656+14, with Lyne and Manchester (1988) giving $\alpha=8.2^\circ$, $\beta=8.2^\circ$, and Rankin (1990) giving $\alpha=30^\circ$ (with no estimate of β .) These earlier studies do not estimate errors, as they are fits to empirical assumptions about the underlying geometry whose errors are not easily estimated.

The RVM is a simple model, which determines the angle of the dipole magnetic field at a constant radius in the magnetosphere, projected onto the observer’s LOS, at different rotation phases. The polarization position angle is then assumed to be parallel or perpendicular to the projected magnetic field lines. We take the position angles predicted by the fits to the radio polarization, degrade the temporal resolution to that of the optical observations (10 resolution elements), and add a constant position angle offset to best fit our polarization measurements (because the position angles of the radio data are uncalibrated). The best-fit RVM prediction is shown in Figure 5.5. The range of α and β allowed by the radio fit does not change the predicted position angle by an amount comparable to the errors in our measurement of θ_i , so our data does not help constrain the observational angles beyond the level of the radio polarization data. The fit of the optical position angles to the RVM prediction is surprisingly good; the measured χ^2 is 2.3 for 2 degrees of freedom (cumulative probability 0.69).

In the context of the polar cap model, the two peaks in the optical emission come from two edges of a hollow cone of emission, that the LOS intercepts at two phases. The separation in phase of the two optical peaks will, in principle, allows a determination of the height of the emission region. In the relevant polar cap models (Daugherty and Harding, 1996), the radio peak appears in the center of the hollow cone of optical emission, so Peak 2 occurs first, followed by the radio emission, and then Peak 1. With a peak separation of 0.4 in phase (a coarse estimate, given our limited temporal resolution), $\alpha = 29^\circ$, $\beta = 8.9^\circ$, the opening angle of the optical cone is 77° . Following the last closed field lines in a dipole field (ignoring relativistic aberration, time-of-travel, and magnetic sweepback effects), the height where the opening angle reaches 77° is $r/r_{\text{LC}} = 0.2$, where $r_{\text{LC}} = cP/2\pi$ is the radius of the light cylinder. At this height, aberration and time-of-travel effects would become large, and this calculation would be too simplistic. The emission region height decreases significantly as α

decreases. Taking $\alpha = 6^\circ, \beta = 1.9^\circ$, which is at the limit of the $1\text{-}\sigma$ confidence interval determined by the radio polarization, the opening angle becomes 17° , at a height of $0.01 r_{\text{LC}}$. This height is comparable to that assumed for the radio emission region.

In the outer gap model, the optical emission is generated farther out in the magnetosphere, where the velocities are mildly relativistic and light-travel times can be large, so the emission region observed at a given phase is not necessarily along the LOS connecting the observer and the neutron star. Therefore, the RVM should not be expected to apply as simply as in the polar cap model. For instance, the optical polarization of the Crab Pulsar shows two distinct swings that cannot be accounted for in the polar cap model but do arise naturally in outer gap models. Agreement of outer gap models with the optical polarization data in this study will require a detailed model of the emission regions observed at each phase.

The Crab Pulsar is the only other pulsar for which optical polarization measurements are available. The polarization pulse profile measured in PSR B0656+14 is different in character than that measured in the Crab. The Crab's linearly polarized flux is maximized at the peaks and minimized in the bridge, while its polarized fraction is maximized in the bridge and minimized at the peaks (Smith *et al.*, 1988). In PSR B0656+14, the linearly polarized flux and polarized fraction are both maximized in the bridge and minimized at the peaks. The low temporal resolution of our optical measurements will lead to a decrease in the measured linearly polarized flux, but our observations of the Crab at the same temporal resolution show a much greater polarized flux at the peaks (where the position angle swings rapidly) than in the bridge.

Acknowledgements

We were greatly assisted by Stephen Kaye and the Palomar engineering staff. Thanks to J. Everett for additional assistance. This work was supported by NSF Grants AST-9618880, AST-9819762, and AST-0096930. This investigation made use of observations made with the NASA/ESA Hubble Space Telescope, obtained from the data archive at the Space Telescope Science Institute. STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.




Figure 5.1: Average intensity image, eliminating all phase and polarization information. PSR B0656+14 is in the middle of the image. Note the extended object which overlaps the pulsar (above the pulsar in this plot).




Figure 5.2: Phase-binned total intensity (*i.e.*, no polarization information) images of PSR B0656+14. Each image is labeled by its phase bin index (0–9), with 0 corresponding to phase 0.0–0.1, 1 to phase 0.1–0.2, *etc.* PSR B0656+14 is in the center of each 40×10 arcsec image. The intensity is a maximum in bins 2 and 8.




Figure 5.3: Total-intensity pulse profile of PSR B0656+14. The intensity scale is normalized to the pulse-averaged total intensity. Error bars are $1\text{-}\sigma$ errors. Pulse is plotted twice for clarity. All intensities have been background-subtracted as described in Section 5.1.3. The radio pulse peaks at phase 0.0, and the possible gamma-ray peak (Ramanamurthy *et al.*, 1996) is at phase 0.2.




Figure 5.4: Linearly polarized flux. The filled circles are the maximum likelihood estimates to the linearly polarized flux, with $1\text{-}\sigma$ error bars, normalized to the pulse-averaged total intensity. The dotted line is the total-intensity pulse profile, as in Fig. 5.3.




Figure 5.5: Linear polarization position angles. Diamonds are the optical position angles with $1\text{-}\sigma$ error bars. Dashed line is predicted position angles from Rotating Vector Model fit to radio data, small dots are 1418 MHz position angles (without error bars) from Everett and Weisberg (2001), solid line is predicted position angles binned in 10 phase bins.

5.2 Emission Models

One of the exciting opportunities that arises from the data on PSR B0656+14 is the chance to test the polar cap and outer gap model predictions. The rotational phase of the optical peaks with respect to the radio peaks, and the sweep of the optical polarization position angle, provide the input to these models.

Polar cap models have fared poorly against optical data in the Crab and Vela pulsars. The Crab Pulsar, the only pulsar (other than our measurements for PSR B0656+14) with high-quality pulse-phased polarization data, shows two nearly identical 90° sweeps of position angle through the two optical peaks (Smith *et al.*, 1988). In the polar cap model, the two optical peaks would be the intersection of the observer line-of-sight with the edges of a hollow cone of emission, with the magnetic axis as the axis of the cone of emission. According to the Rotating Vector Model, the position angle sweep should be centered on the magnetic axis, with no additional features associated with the two optical peaks. This contradiction is difficult to reconcile under the polar cap model. One of the significant successes of the outer gap models, on the other hand, was the ability to account for the double-sweep of position angle in the Crab Pulsar (Romani and Yadigaroglu, 1995).

The two optical peaks of the Vela pulsar are not symmetrically located with respect to the radio pulse. Defining rotational phase such that the radio pulse is at phase 0.0, the two optical peaks occur at phases 0.15 and 0.6. These locations are not symmetric about the radio pulse (or about phase 0.5). Polar cap models have no capacity to introduce phase offsets between components, and so again, the optical measurements are difficult to reconcile with the polar cap models. Again, the outer cap models, on the other hand, can be arranged to generate the observed emission profiles (Romani and Yadigaroglu, 1995).

Applying the polar cap models to our optical observations of PSR B0656+14, as discussed in the preceding paper, poses no problems. The optical peaks are symmetric about the radio pulse, occurring at phases 0.2 and 0.8. The only other constraint on the data comes from the radio polarization data, which, through the Rotating Vector Model (RVM), constrains the geometry of the pulsar and observer. The pulsar geometry is specified by the angle between the rotation axis and magnetic axis, α , and the observer colatitude, ζ . The observer colatitude (the angle between rotation axis and observer line-of-sight) is often rearranged to give the impact parameter, β , which is the smallest angle between magnetic axis and observer line-of-sight, when the rotation phase is zero. In other words, $\zeta = \alpha + \beta$. RVM models generally constrain β much more tightly than α (and therefore ζ).

Using the 1418 MHz radio polarization data from Weisberg *et al.* (1999) and models from Everett and Weisberg (2001), we calculate the range of α and ζ allowed by fits to the RVM in Fig. 5.6. The raw radio data are shown in Fig. 5.5 in the PSR B0656+14 paper. The lack of constraint on α is




Figure 5.6: Rotating Vector Model constraints. The lines denote regions where χ^2_{ν} , the reduced χ^2 , increases by $\Delta\chi^2_{\nu} = 3$.

immediately obvious from this figure, however, β is constrained to positive (and small) values, *i.e.*, $\zeta > \alpha$. The geometrical arrangement required to describe the optical emission, assuming an exterior hollow cone emission model near the polar cap, is consistent with the radio polarization data.

Since the work of Romani and Yadigaroglu (1995), outer gap models have followed a specific prescription to form the basis of their pulse profile predictions. Beginning with the assumption of a dipolar field at the neutron star surface, the magnetic field of the rotating neutron star is calculated using the full retarded potentials (Deutsch, 1955). The surface of last closed field lines is determined by those field lines which intersect tangentially with the light cylinder. Field lines in the open field line region (*i.e.*, field lines that exit the light cylinder) are parametrized by a gap width parameter w , which measures the distance between the magnetic axis and the point where that field line intersects the NS surface, normalized to the distance between the magnetic axis and the last closed field lines at the NS surface. Radiation is assumed to be emitted by particles traveling at the speed of light along the magnetic field lines, and relativistic aberration determines the direction in which the radiation travels in the nonrotating reference frame. The direction of emission from every point along a magnetic field line can then be represented in spherical coordinates, by colatitude and longitude. The observer occupies a particular colatitude (ζ), and as the NS rotates the observer sees the longitude vary with the rotational phase. The rotational phase at which radiation from a particular point in the magnetosphere is observed is determined both by the direction of propagation and the light travel time delay. Accounting for all these factors, the emission from every magnetic field line can be mapped into the (φ, ζ) plane, where φ is the rotational phase and ζ is the observer colatitude. Examples of these skymaps are shown in Appendix A. A full description of the construction of

skymaps is given in the PhD thesis of Yadigaroglu (1997).

The skymaps can be interpreted as pulse profiles by identifying the ζ of the observer and transforming the map of emission points into observed intensities. A standard assumption is that radiation is emitted from magnetic field lines that form a cone of constant w , the simplest assumption being that the intensity is uniformly distributed across the surface area. The skymap is transformed into observed intensity through the Jacobian defining the transformation of surface area at the magnetosphere into (φ, ζ) area on the skymap. Regions where the field lines in the skymap “fold” over themselves correspond to peaks in the intensity, because the Jacobian of the area transformation becomes very large (*i.e.*, the transformation from surface area to (φ, ζ) becomes degenerate).

The generation of pulse profiles according to this prescription is a matter of representing the intensity as a function of 4 variables, φ , ζ , α , and w . Rather than attempting to distill this information into some subset of this parameter space, we display the skymaps in all four dimensions, in Appendix A. This situation is simplified somewhat by the constraint on ζ as shown in Fig. 5.6. For a given α , only values of ζ slightly larger than α (within ~ 10 degrees) are consistent with the radio data. Fitting the optical data to the polar cap model is now a question of finding pulse profiles from skymaps with small β that match the observed pulse profile.

There are some general predictions about which outer gap model parameters are appropriate. For a middle-aged pulsar such as PSR B0656+14, the energetics required to “close the gap” and end the cascade of particles are harder to achieve, and so the gap is wider than it would be in a younger pulsar. Some predictions place the gap width of PSR B0656+14 at $w=0.7$ (Cheng *et al.*, 2000), nearly filling the entire open field line region.

Because many of the particles responsible for the emission in the outer gap region are generated by pair production, there are both positively and negatively charged particles participating in the emission. The oppositely charged particles will be accelerated in opposite directions, causing emission to be both inward-going and outward-going. In the study of gamma-ray pulsars, which are commonly double-peaked, it was determined that if both inward- and outward-going emission were observable, many of these pulsars would have four peaks. The assumption is then that the inner magnetosphere is opaque to gamma-ray radiation, and the inward-going emission does not reach the observer.

We are not aware of similar opacity calculations for optical photons. In fact, outer gap models of hard X-ray emission sometimes invoke outward-going emission for gamma rays and inward-going emission for X-rays (Zhang and Cheng, 2001). As such, it seems reasonable to allow the possibility that the observed optical pulse profile may be due to either outward- or inward-going emission.

We attempt to find skymaps which reproduce the observed optical pulse profiles. Generally adhering to the restrictions on ζ given in Fig. 5.6, we attempt to find values of α and w that give peaks near phase 0.2 and 0.8, relative to the radio pulse at phase 0.0. A preliminary analysis, without calculating full pulse profiles from the plots in Appendix A, finds it hard to account for both of the

optical pulses using either inward- or outward-going emission, but that it might be possible using both.

We find general features in the emission skymaps in Appendix A. The emission is generally centered around phase 0.5, and pulse widths are almost always smaller than 0.5. The optical pulse we measure has width 0.6, which, while not significantly greater than the 0.5 just mentioned, seems hard to generate from the skymaps. We estimate that it would be possible to generate the peak at phase 0.2 from the $\alpha=50^\circ$, $w=0.05$ outward-going map at $\zeta \sim 75^\circ$, and the peak at phase 0.8 from the $\alpha=50^\circ$, $w=0.05$ inward-going map at $\zeta \sim 75^\circ$.

This analysis is very preliminary, and the reason for including the skymaps in Appendix A is to allow for a “by eye” confirmation of the ideas presented here. Calculation of the pulse profiles, which we can generate on a case-by-case basis but not by an automated routine, will be the next step in this investigation.

To summarize, we are hypothesizing that the outer gap models may fit our data, but with emission coming from a much narrower gap than predicted ($w=0.05$, compared to a prediction of 0.7), and utilizing both the inward- and outward-going emission from the gaps.

Chapter 6

Low-Mass X-ray Binaries and Cataclysmic Variables

The first observing run with the camera operating in multiple frame transfer mode was in July 2001. The primary target of this investigation was Cyg X-2, for which we had scheduled simultaneous observations with the *Rossi X-ray Timing Explorer* satellite. The other objects of this observing run were other low-mass X-ray binaries (LMXBs) and cataclysmic variables (CVs), Her X-1, RW Tri and WZ Sge, with no additional X-ray coverage.

The Cyg X-2 observations were performed alternately through an H α filter and in unfiltered white light, with integration times of 100 ms and 5 ms, respectively. The interest in the H α high-time-resolution data was to watch for signatures of coronal flares, which might have extremely short timescales and briefly dominate the H α emission. This was a speculative exploration, searching for coronal flares as a signature of magnetic field loop reconnection. This reconnection mechanism would explain the production of hard X-rays, which are unaccounted for under the standard accretion disk theories. Guillaume Dubus was the principal investigator for this observing run.

The data analysis, which will not be described in detail here, had few similarities with the data analysis procedure discussed in Section 3.1, because charge traps play an insignificant role in the error analysis when the charge is not shifted back and forth over thousands of cycles. As in the phase-binning case, the use of comparison stars observed simultaneously with the object star was essential. In fact, the field of view of the camera was specifically widened (to its current wide-field configuration) specifically for this observing run, to ensure that comparison stars of comparable brightness to the science objects were observable in each field. Because the LMXBs and CVs were bright ($V \sim 8\text{--}12$), the challenge was to find comparison stars that were of comparable brightness. In the end, the comparison stars proved absolutely essential, as the observing conditions were not photometric. Correlations with the comparison stars, which are known to be stable to < 0.01 mag (Henden and Honeycutt, 1997), allow differential photometry to obviate the need for photometric conditions.

The principal challenges of the data analysis of the multiple frame transfer data consisted of finding appropriate photometric apertures for the object and comparison stars. Read noise dominates large apertures while image wander (nonperiodic, as opposed to that seen in Section 3.2) dominates small apertures. The reduction of this data eventually resulted in images of Cyg X-2 with a signal-to-noise ratio of ~ 20 per 5 ms exposure in white light, and ~ 5 per 100 ms exposure in H α . The

5 ms exposures had an integration duty cycle of only 10%, the 100 ms exposures 70%. While these are low duty cycles, the exposures come in bursts of 50 contiguous exposures, so that the resolution exists on the shortest timescales, but with holes in the temporal coverage.

The analysis thus far has been performed by using autocorrelations of the optical and X-ray data and cross-correlations between them. No features in the correlations are obvious at a first glance, and further investigation is ongoing.

Chapter 7

Conclusions

As is true more often than not, the material presented in this thesis is a work in progress. The work of constructing the phase-binning CCD camera as an effective and reliable observing instrument is largely finished, until such time as modifications will need to be made to adapt to other telescopes (Keck?). The data analysis routines, while not automated for turnkey convenience, have proven adequate for performance near the fundamental limits of the observations. However, the interesting work, as I see it, has just begun.

The data presented in this thesis is based on only two observing runs. While the list of pulsars observable from Palomar does not stretch on very long, use of this instrument at Keck or in the Southern Hemisphere (because Vela is significantly brighter than the pulsars we observed) promises to provide a few more objects in the near future. The real promise of this technique is in pulse-phased polarimetry, which will also be possible at Keck.

On the theoretical front, a great deal of opportunities present themselves with the new data from both pulsars observed. The optical emission expected from magnetars is a topic currently under investigation by other groups (C. Thompson, private communication), which will benefit from the solid measurement presented by this work. More importantly, the confirmation that 4U 0142+61 is, in fact, a magnetar is the most significant contribution of this thesis work to the scientific community.

Our treatment of the outer gap considerations from PSR B0656+14 was not comprehensive, and we have not dealt with the polarization data (although of low significance) in the framework of the outer gap models. Our analysis of the data from PSR B0656+14 does not make or break either of the competing models for high-energy pulsar emission, but it does at least set the record straight regarding the previous measurements of pulsations from PSR B0656+14.

The gratification of seeing a project through from its initial inception on paper through the production of exciting scientific results is hard to overestimate. I sincerely expect that this project will be able to continue without the degree of mechanical and electronic coercion (and, at times, brute physical force) required to bring the camera to its presently mature state.

At present, this camera fills a niche that is largely ours to fill. However, new superconducting detectors under development promise to exceed our capabilities. Superconducting Tunnel Junctions (STJs), Transition Edge Superconductors (TESs), and kinetic inductance detectors all promise photon counting performance with near-perfect quantum efficiency, with the added bonus of energy resolution on each photon and some imaging abilities. At present, the imaging capabilities of these

technologies are extremely limited, and there are throughput issues associated with coupling each of these types of detectors to telescopes, but it is only a matter of time until these technologies dominate this field. As such, the phase-binning CCD camera is likely to have a few more years of competitive advantage, until the competing technologies have developed the imaging capabilities that will enable robust statistical analysis to be performed on the data.

Appendix A—Outer Gap Model Calculations

Collected here are the skymaps computed from the outer gap models as described by Yadigaroglu (1997). Each page of figures is calculated from a single value of α , the angle between the rotation and magnetic axes. The different plots on a page correspond to magnetic field lines originating in different regions of the polar cap of the NS, as parametrized by w . w is the ratio of the distance from the intersection of a particular field line with the NS surface to the magnetic axis, normalized by the distance from the last closed field line to the magnetic axis, measured in the sense that $w = 0$ is the location of the last closed field line, and $w = 1$ is the magnetic axis. Outer gaps form at the null charge surfaces, the locations where the Goldreich-Julian charge density vanishes. This occurs when the magnetic field is perpendicular to the rotation axis, which corresponds to $\zeta = 90^\circ$. Therefore, emission in these skymaps is only seen from $\zeta < 90^\circ$. The maps must be turned upside down to see emission for $\zeta > 90^\circ$, as that emission must come from the opposite pole.

In each skymap, the observed rotational phase is plotted as φ , and the observer colatitude is ζ . As an observer is located at only one value of ζ , the pulse profile is obtained by slicing through the skymap at a constant ζ . The basic assumption is that emitted intensity is proportional to surface area on the surface of field lines at a constant w , and so the observed intensity depends on the Jacobian of the transformation between surface area and (φ, ζ) . Regions where the field lines “fold” over themselves provide the highest intensity (*i.e.*, peaks in the pulse profile).

In the outer gap models, emission may be emitted in both directions along a field line, *i.e.*, inward-going and outward-going. These directions are plotted separately here, with all of the outward-going plots appearing before the inward-going plots. Models based on outer gap emission generally only examine the outward-going emission, assuming that the inward-going emission is absorbed in the inner magnetosphere before escaping to where the observer can receive it. We present the inward-going emission skymaps because the outward-going emission alone does not easily explain the observations.

In each of the figures, the gray circles in the upper hemisphere ($\zeta < 90^\circ$) correspond to the pole emitting the observed radio emission.




Figure 7.1: Skymap of outward-going emission for $\alpha=10^\circ$.




Figure 7.2: Skymap of outward-going emission for $\alpha=20^\circ$.

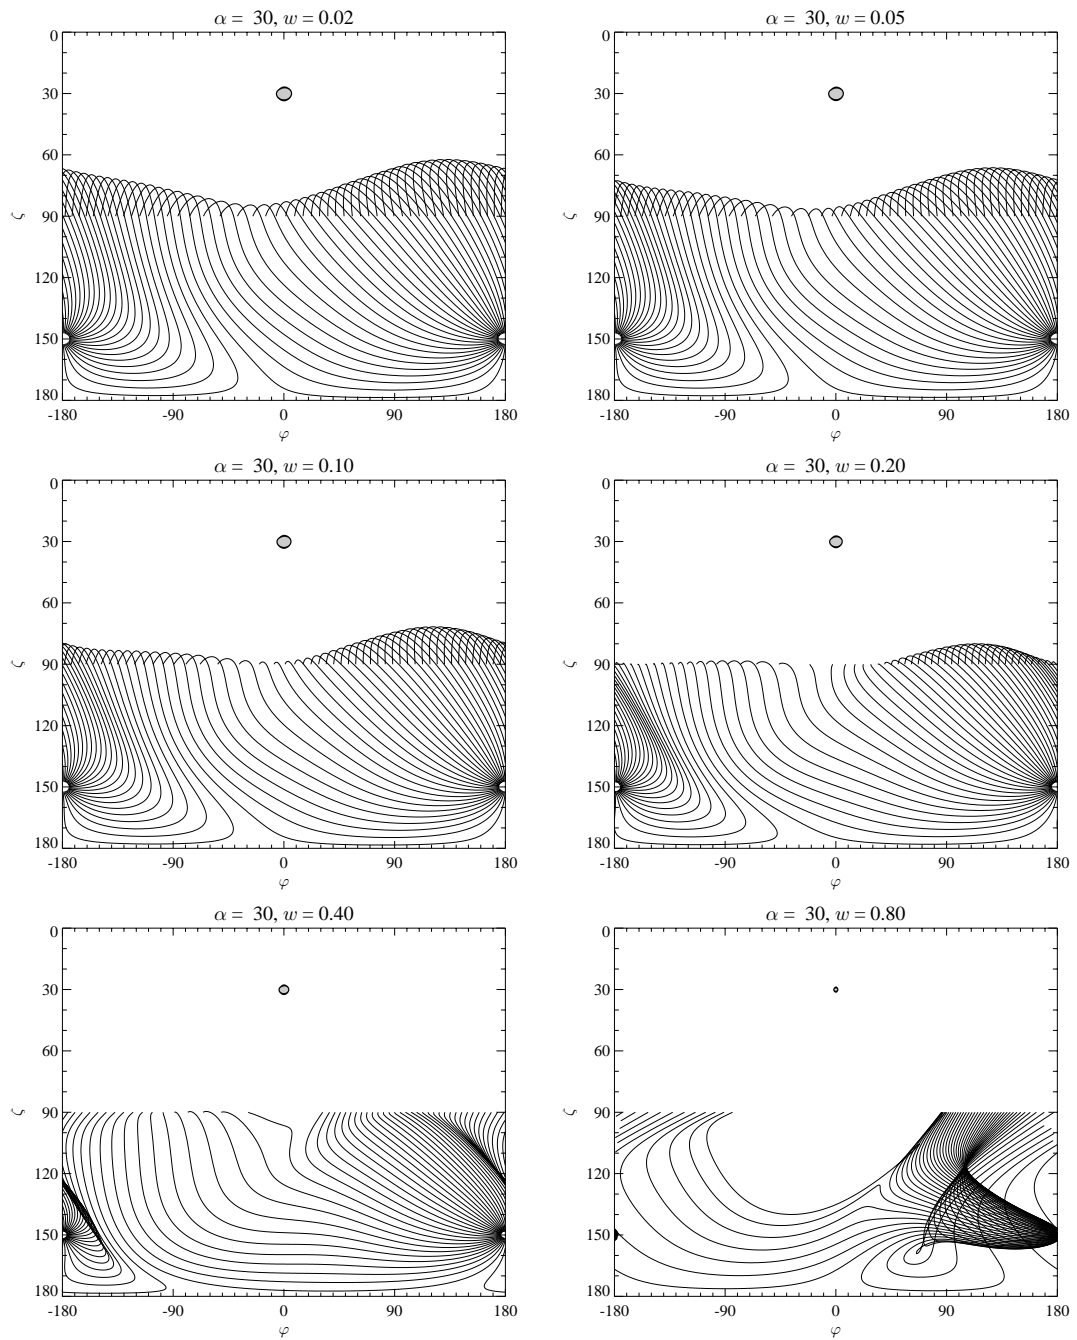


Figure 7.3: Skymap of outward-going emission for $\alpha=30^\circ$.

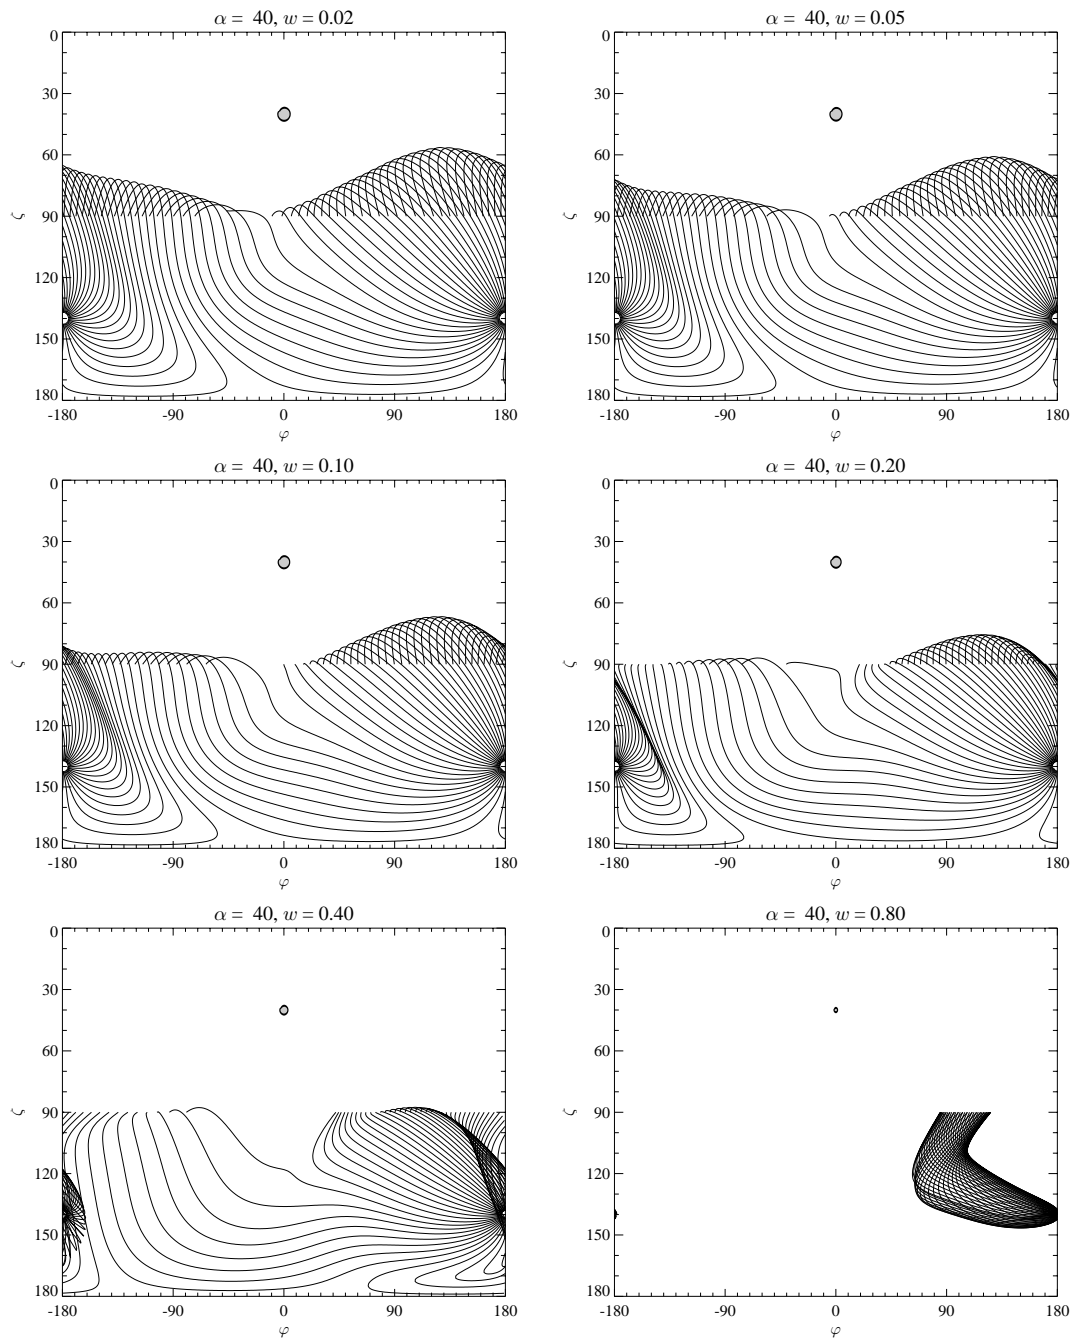


Figure 7.4: Skymap of outward-going emission for $\alpha=40^\circ$.

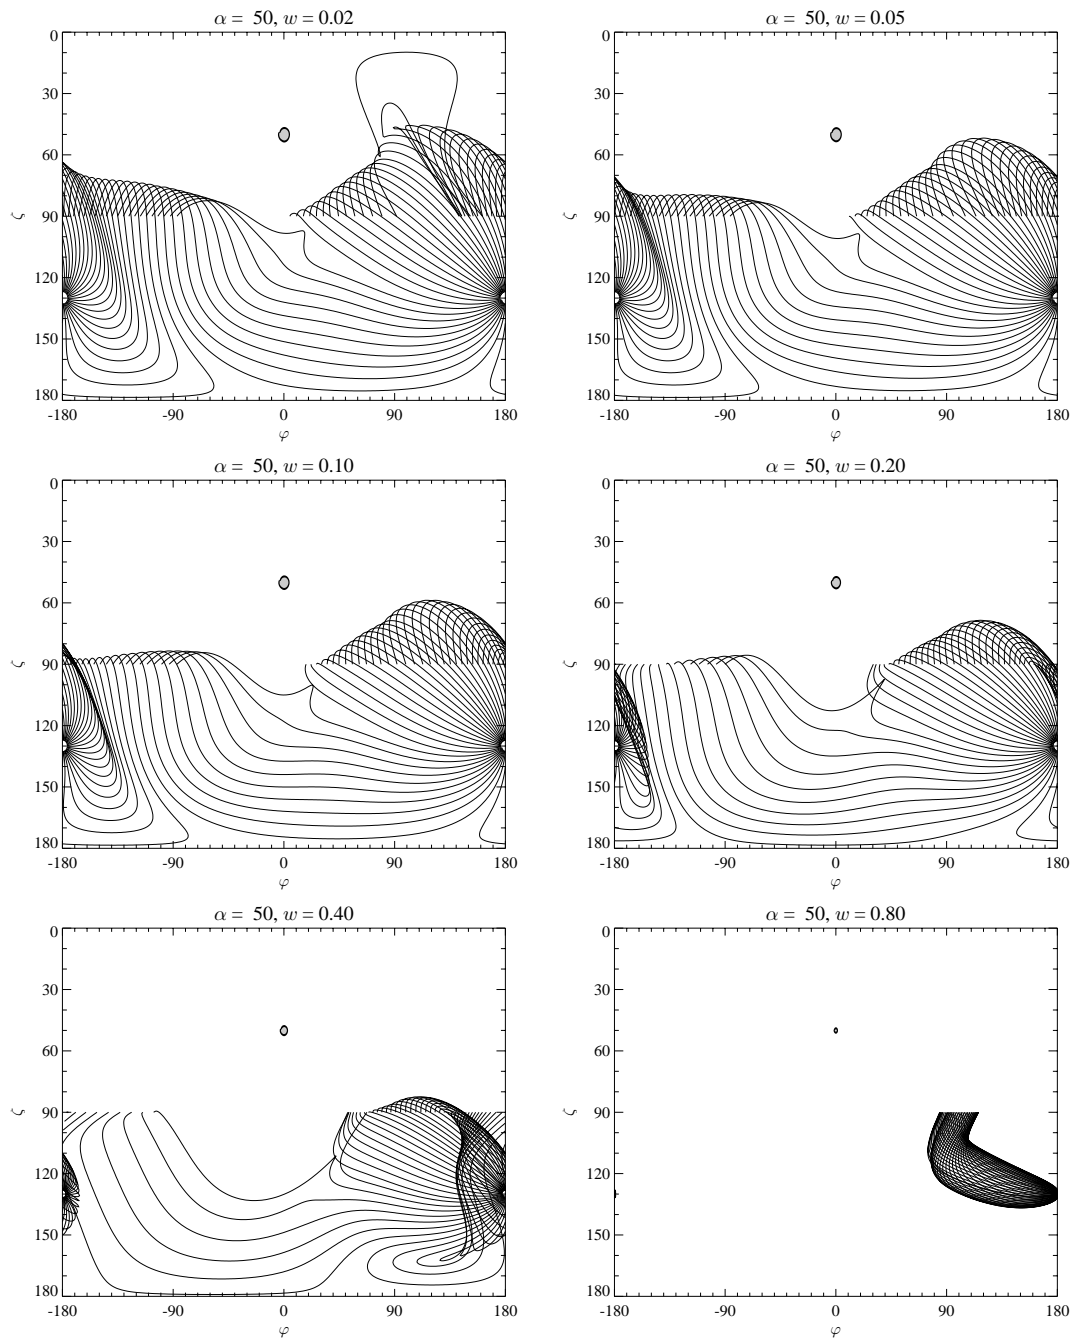


Figure 7.5: Skymap of outward-going emission for $\alpha=50^\circ$.

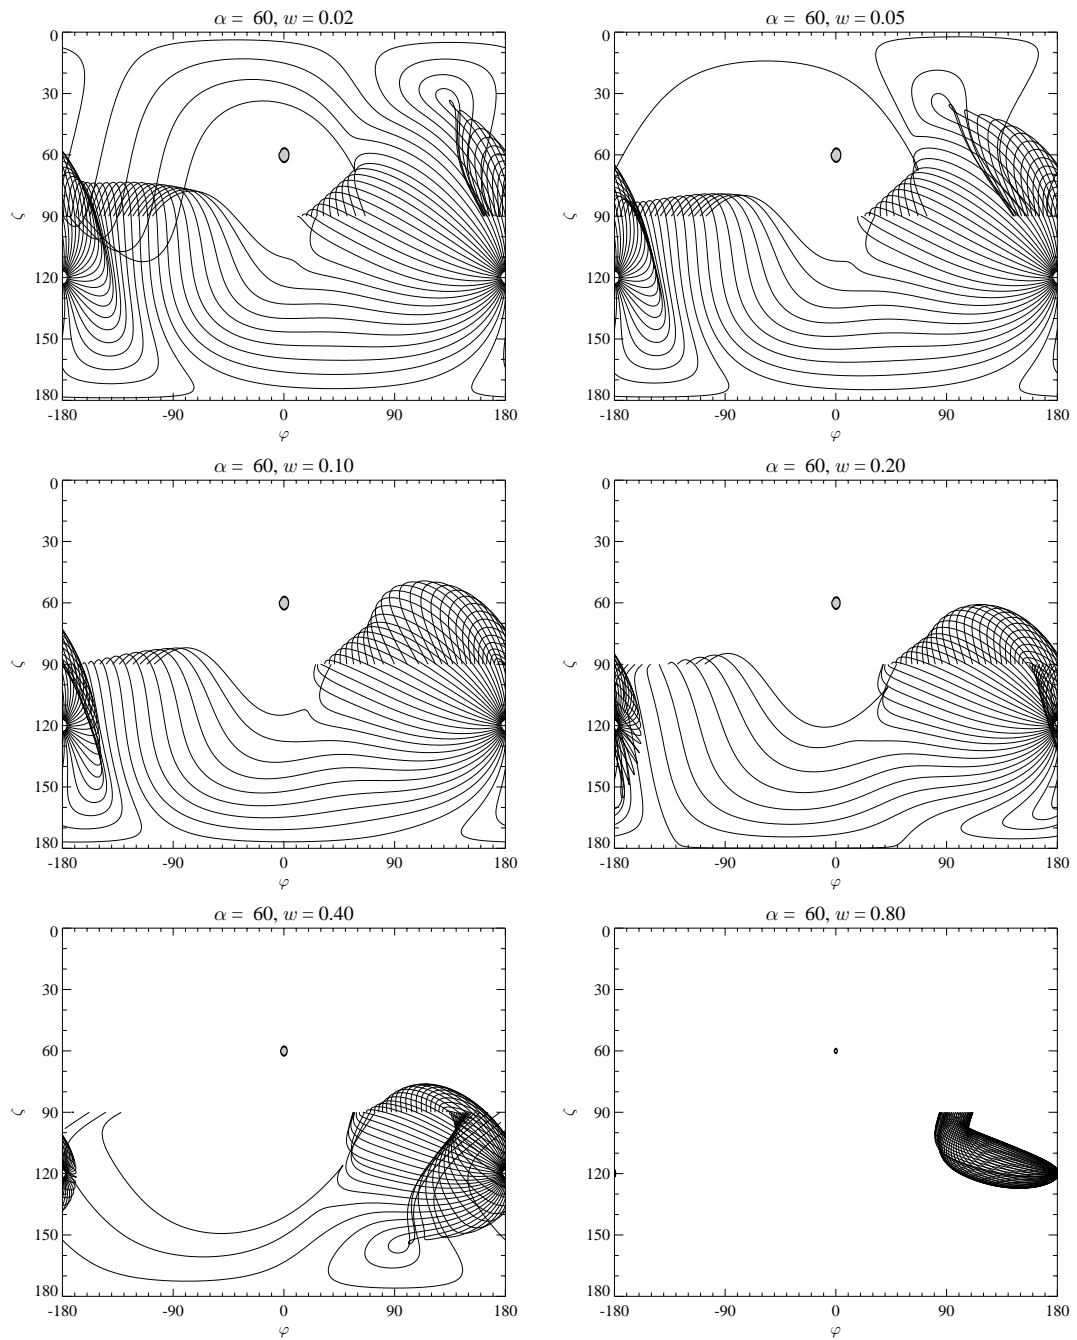


Figure 7.6: Skymap of outward-going emission for $\alpha=60^\circ$.

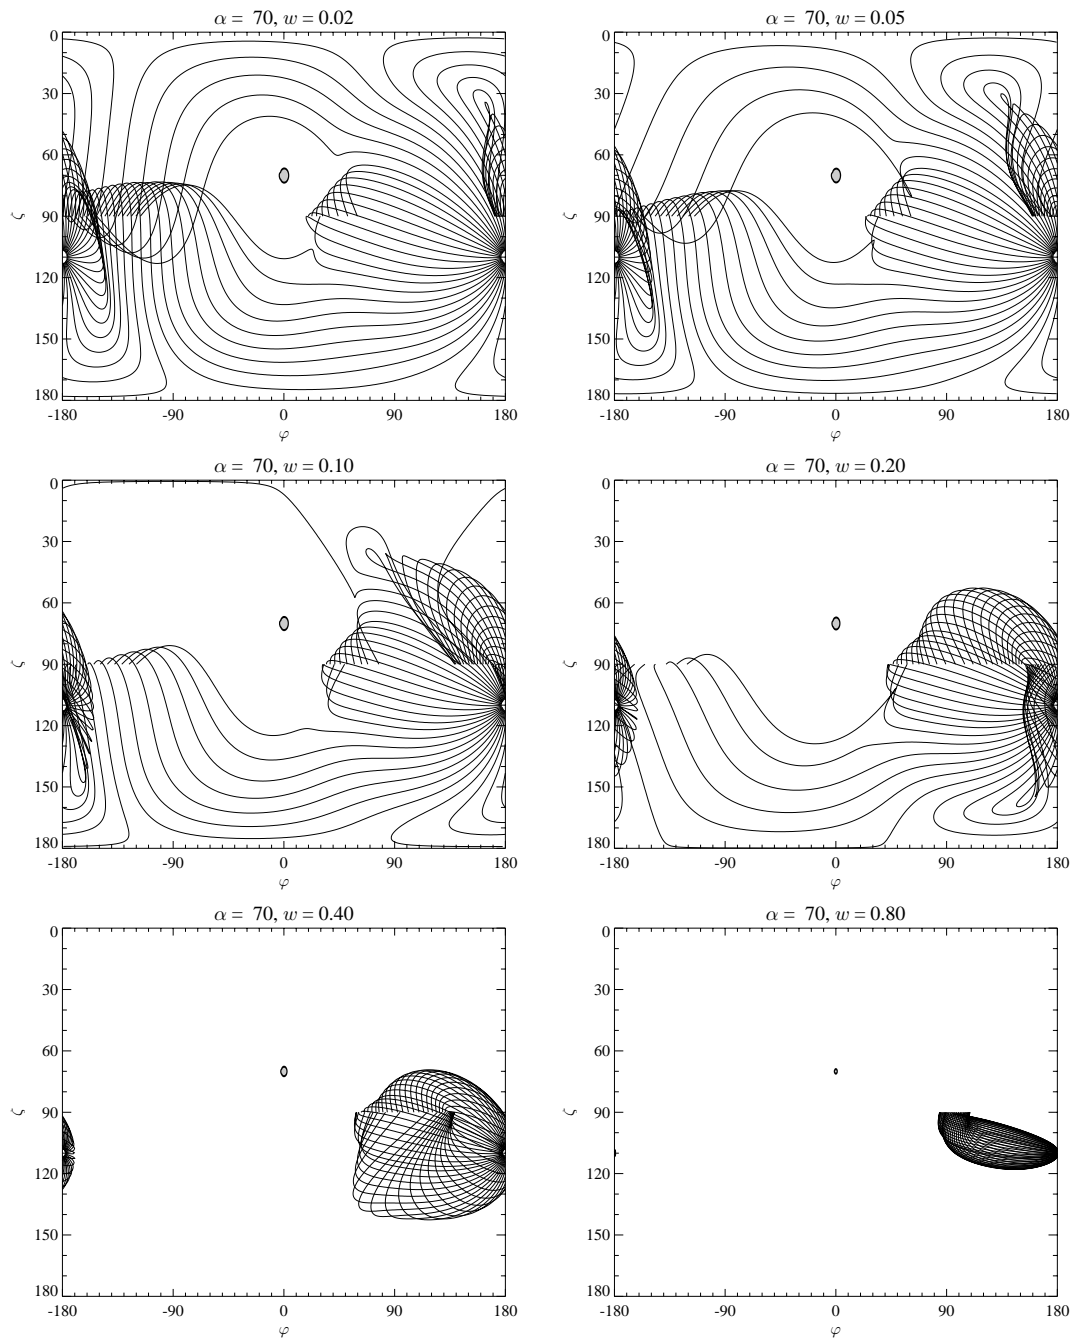


Figure 7.7: Skymap of outward-going emission for $\alpha=70^\circ$.

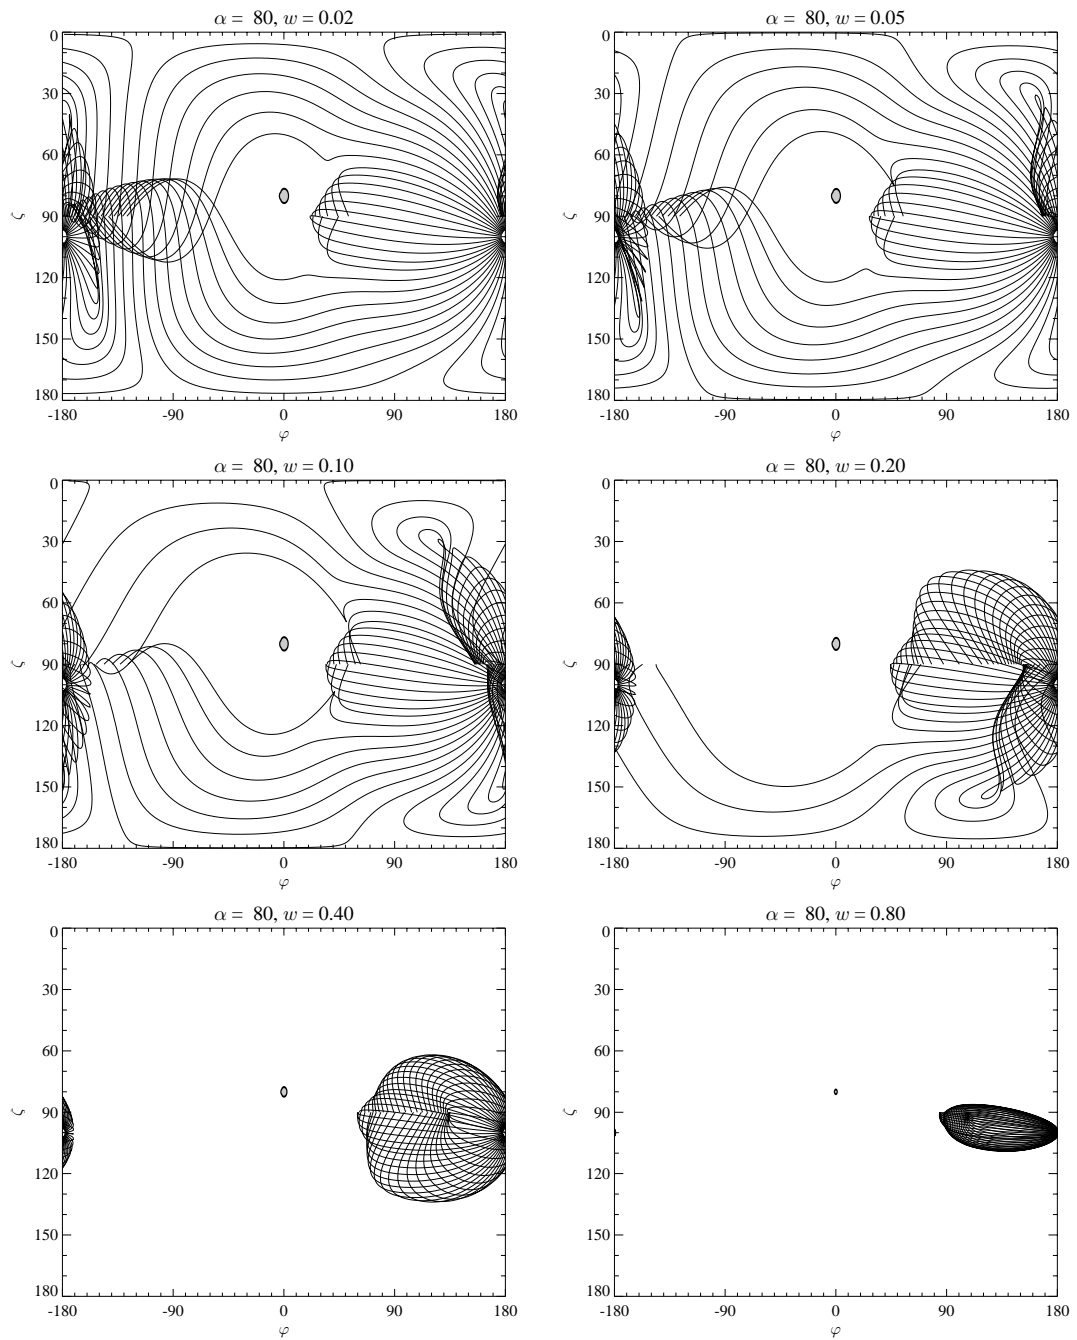


Figure 7.8: Skymap of outward-going emission for $\alpha=80^\circ$.




Figure 7.9: Skymap of inward-going emission for $\alpha=10^\circ$.




Figure 7.10: Skymap of inward-going emission for $\alpha=20^\circ$.




Figure 7.11: Sky map of inward-going emission for $\alpha=30^\circ$.

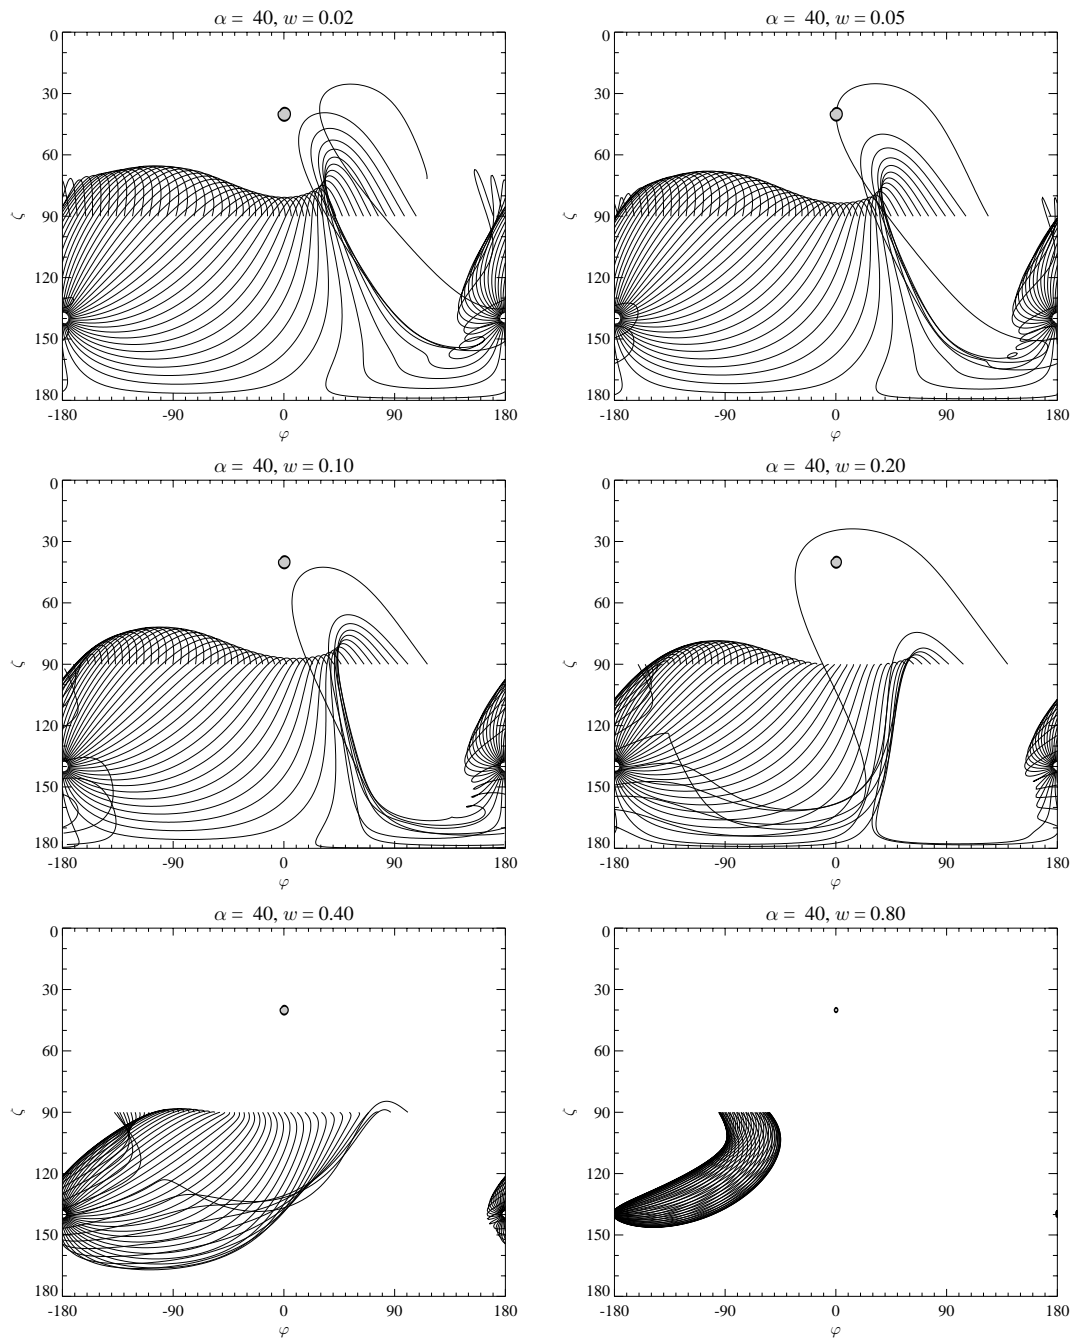


Figure 7.12: Skymap of inward-going emission for $\alpha=40^\circ$.

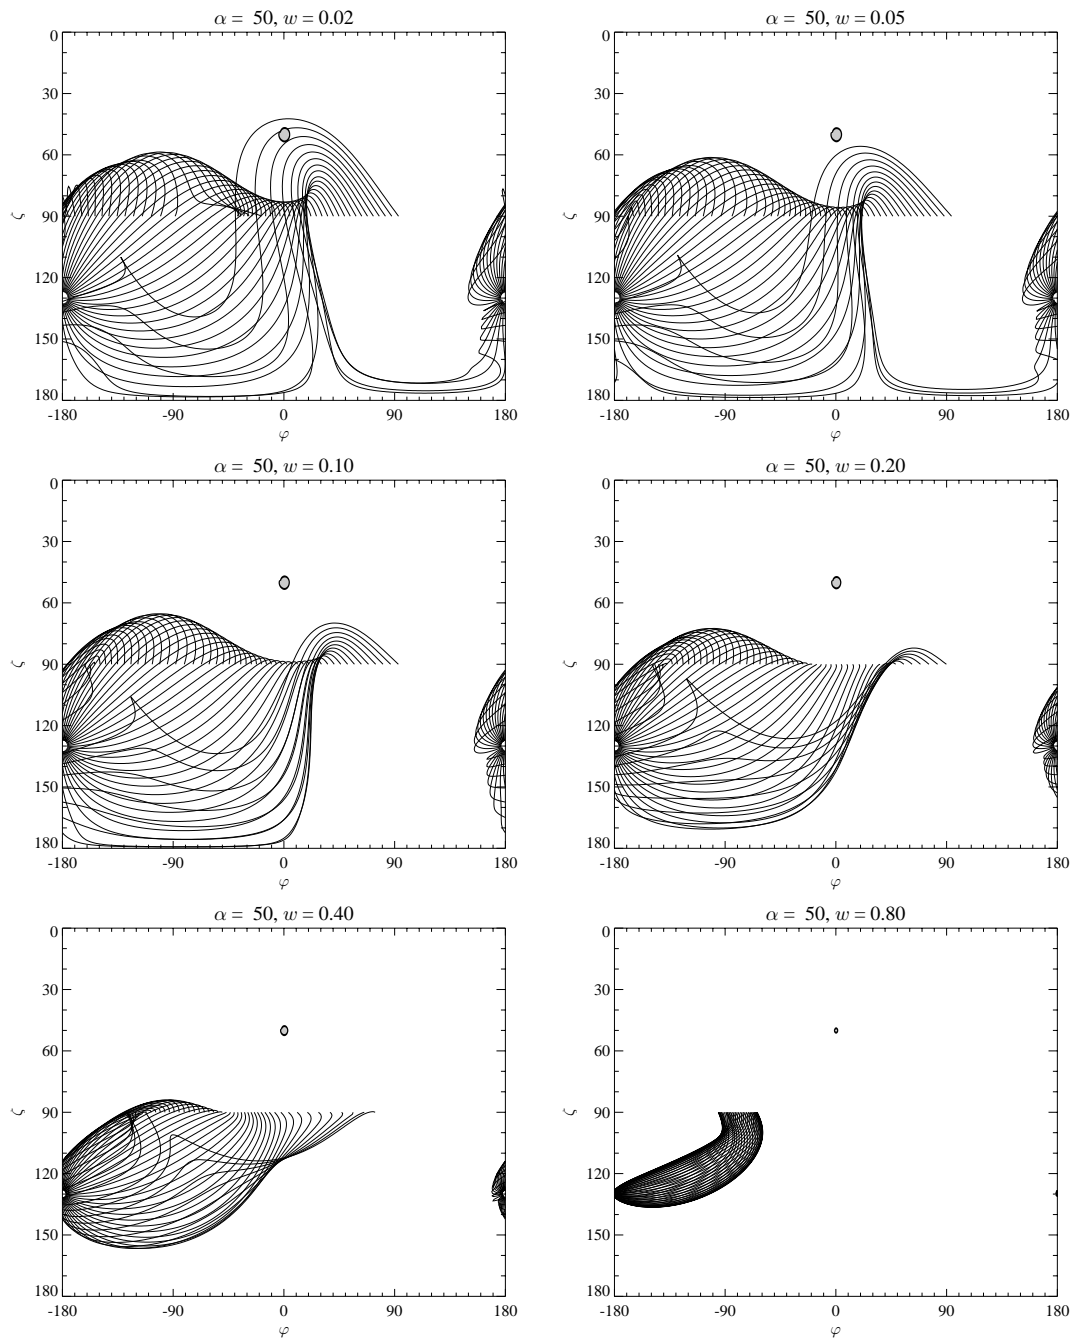


Figure 7.13: Skymap of inward-going emission for $\alpha=50^\circ$.

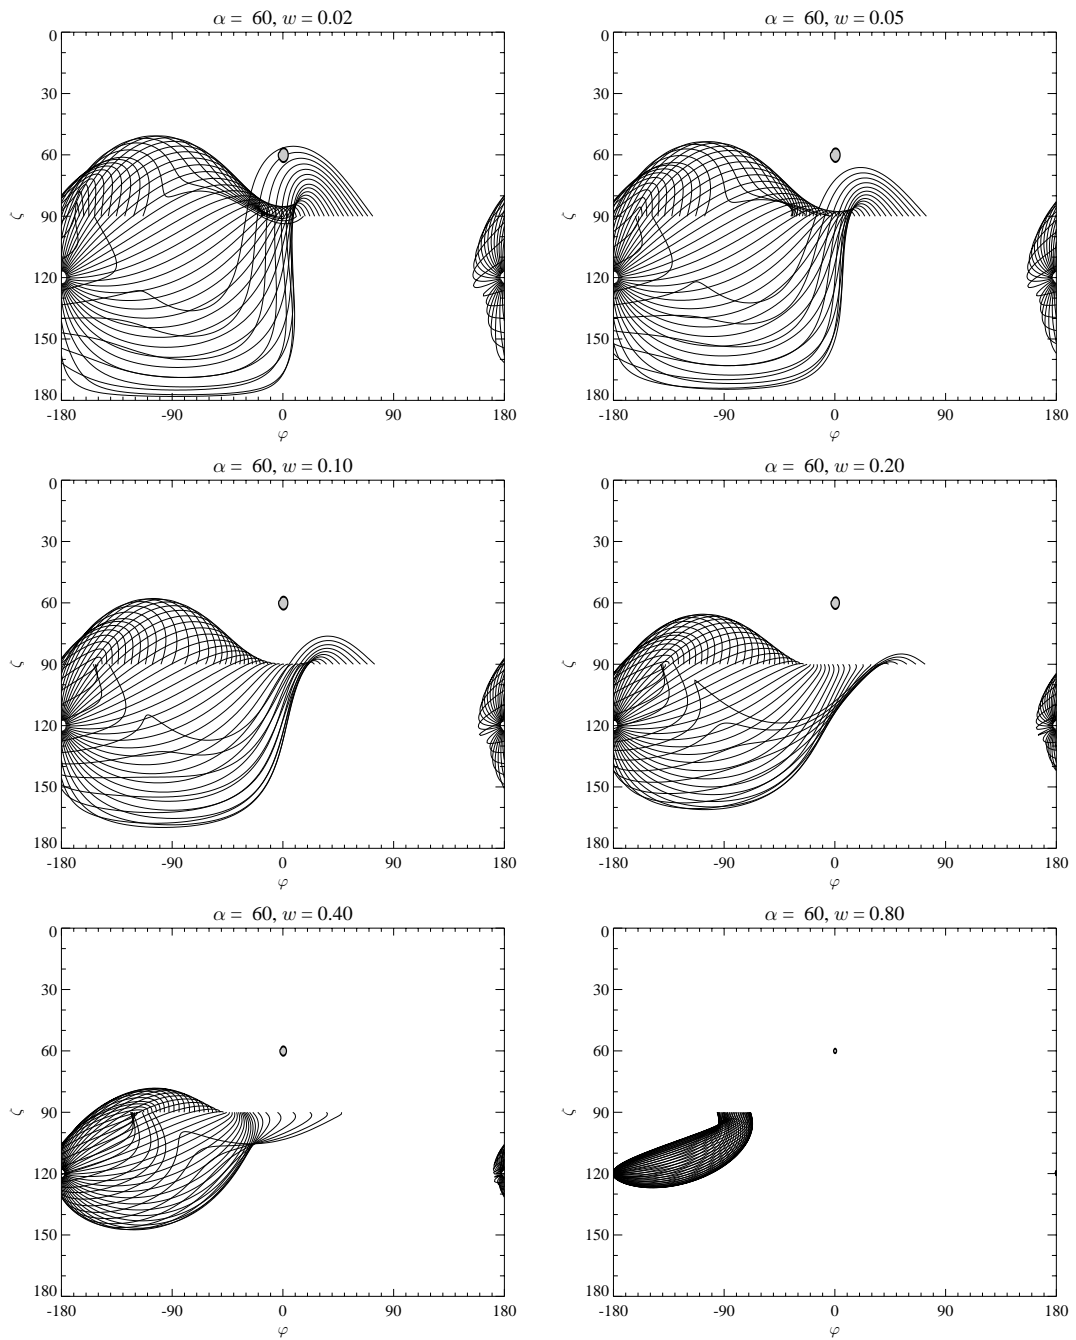
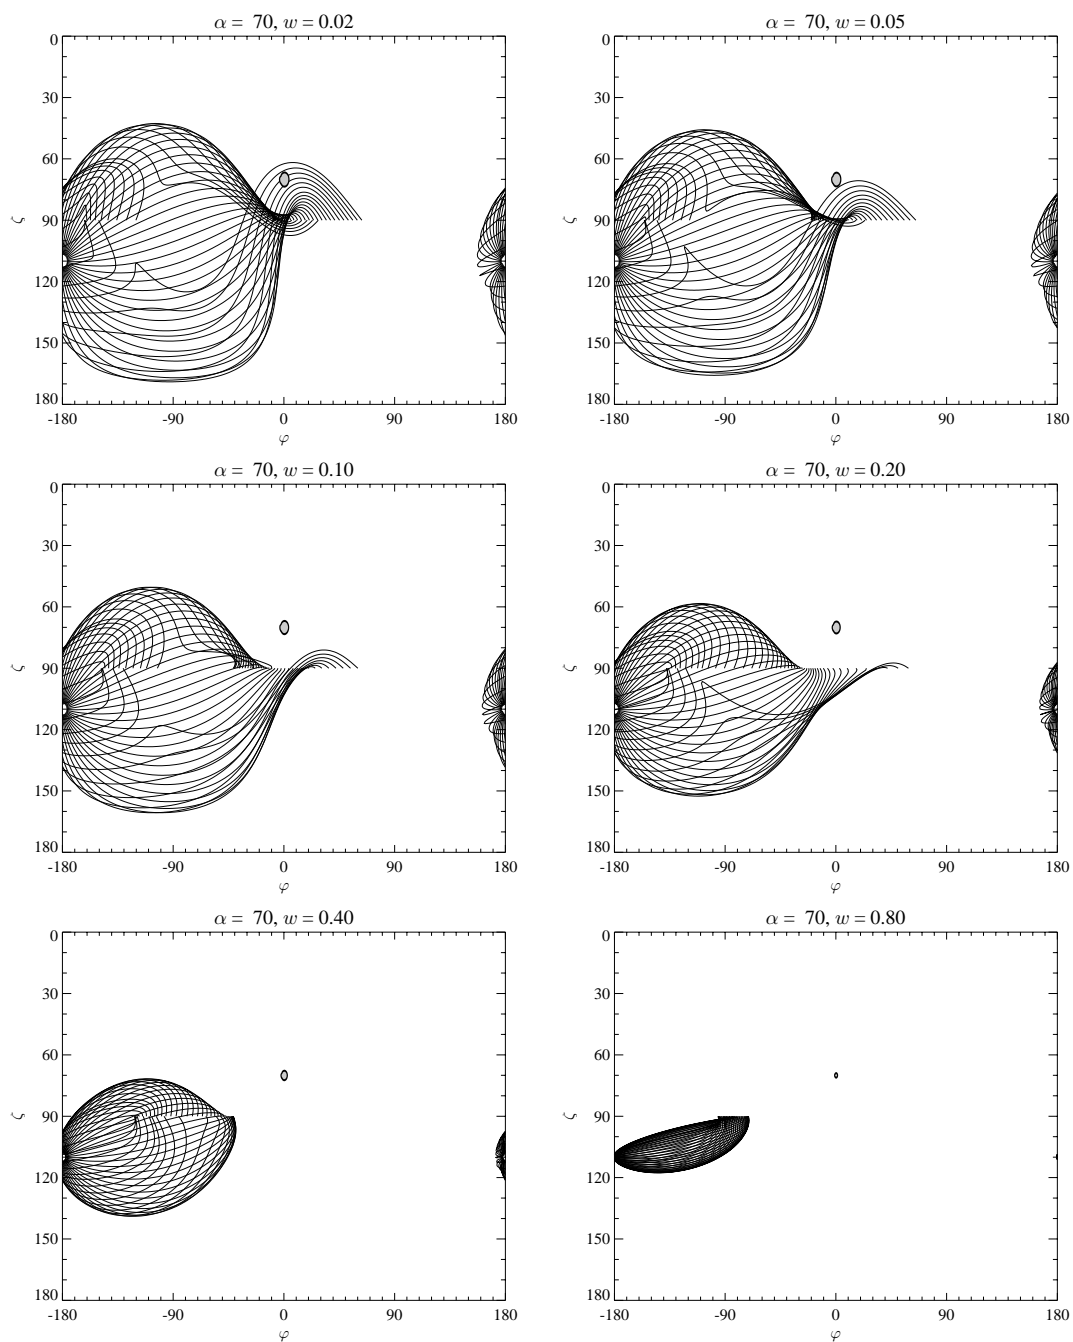
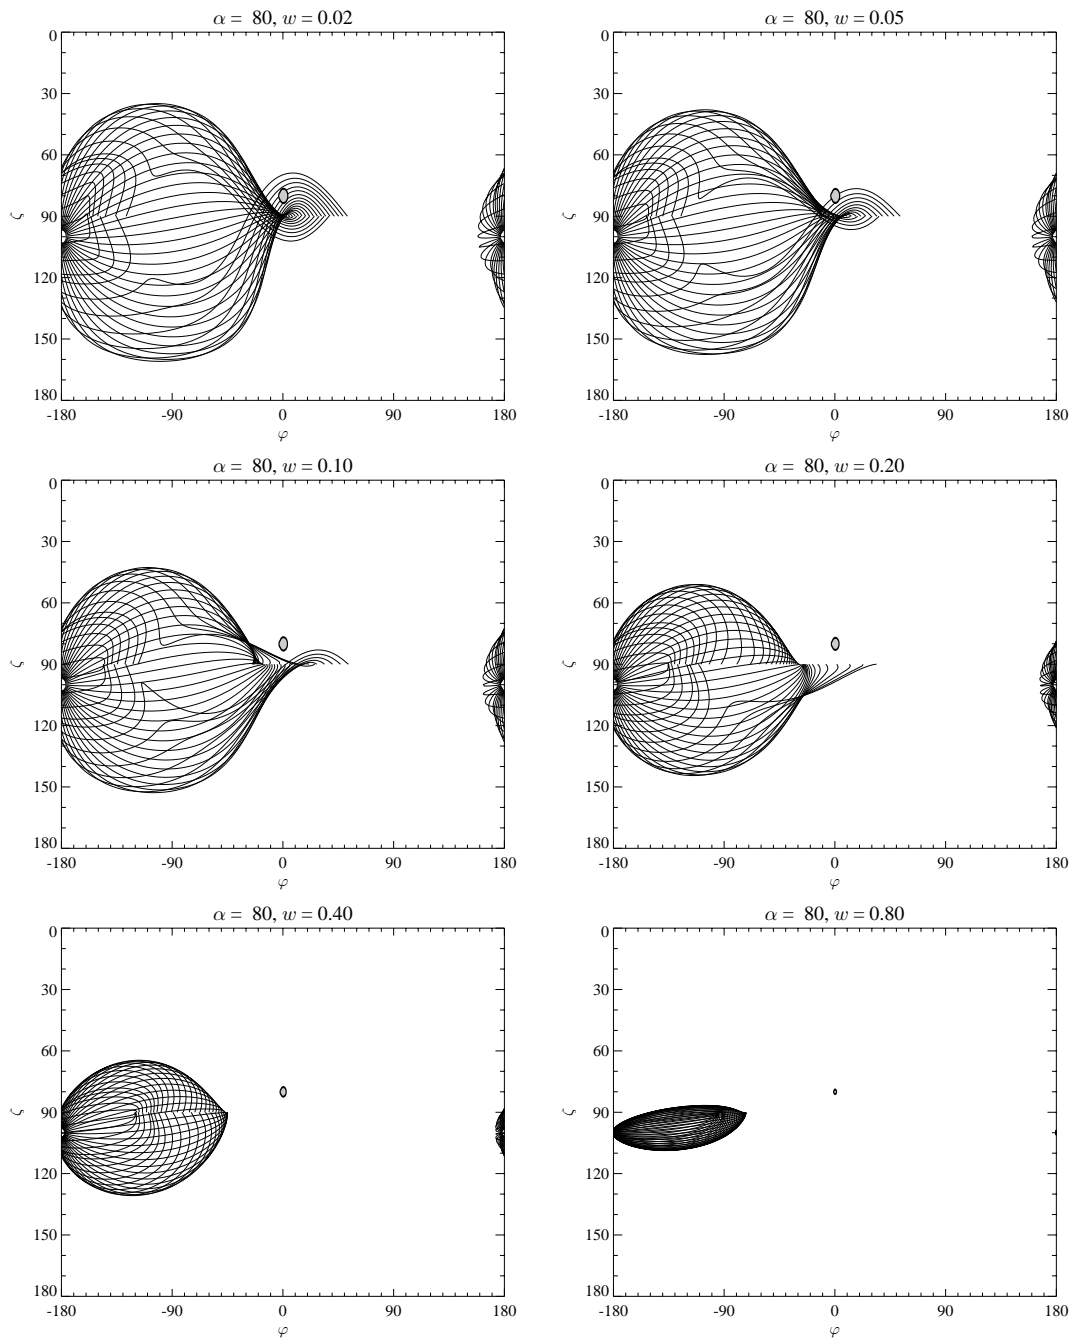


Figure 7.14: Sky map of inward-going emission for $\alpha=60^\circ$.

Figure 7.15: Skymap of inward-going emission for $\alpha=70^\circ$.

Figure 7.16: Sky map of inward-going emission for $\alpha=80^\circ$.

Bibliography

- Achilleos, N., Wickramasinghe, D. T., Liebert, J., Saffer, R. A., and Grauer, A. D. (1992). Exploring the peculiar magnetic field of Feige 7. *ApJ*, **396**, 273–288.
- Baring, M. G. (2001). Identifying the Mysterious EGRET Sources: Signatures of Polar Cap Pulsar Models. In A. Carraminana, editor, *The Nature of Unidentified Galactic High-Energy Gamma-Ray Sources*. (astro-ph/0106161).
- Baring, M. G. and Harding, A. K. (2001). Photon Splitting and Pair Creation in Highly Magnetized Pulsars. *ApJ*, **547**, 929–948.
- Barstow, M. A., Jordan, S., O’Donoghue, D., Burleigh, M. R., Napiwotzki, R., and Harrop-Allin, M. K. (1995). RE J0317-853: the hottest known highly magnetic DA white dwarf. *MNRAS*, **277**, 971–985.
- Becker, W. and Trümper, J. (1997). The X-ray luminosity of rotation-powered neutron stars. *A&A*, **326**, 682–691.
- Boyd, P. T., van Citters, G. W., Dolan, J. F., Wolinski, K. G., Percival, J. W., Bless, R. C., Elliot, J. L., Nelson, M. J., and Taylor, M. J. (1995). High-Speed Photometer Observations of the LMC Pulsar B0540-69. *ApJ*, **448**, 365–368.
- Cheng, K. S., Ruderman, M., and Zhang, L. (2000). A Three-dimensional Outer Magnetospheric Gap Model for Gamma-Ray Pulsars: Geometry, Pair Production, Emission Morphologies, and Phase-resolved Spectra. *ApJ*, **537**, 964–976.
- Chester, T. J. (1979). Continuum optical pulsation from the companions of binary X-ray pulsars. *ApJ*, **227**, 569–578.
- Cocke, W. J., Disney, M. J., and Taylor, D. J. (1969). Discovery of optical signals from pulsar NP 0532. *Nature*, **221**, 525.
- Coe, M. J. and Pightling, S. L. (1998). Deep infrared and optical studies of the fields of three proposed TŻO remnants. *MNRAS*, **299**, 223–228.
- Daugherty, J. K. and Harding, A. K. (1994). Polar CAP models of gamma-ray pulsars: Emission from single poles of nearly aligned rotators. *ApJ*, **429**, 325–330.

- Daugherty, J. K. and Harding, A. K. (1996). Gamma-Ray Pulsars: Emission from Extended Polar CAP Cascades. *ApJ*, **458**, 278+.
- Deutsch, A. J. (1955). The electromagnetic field of an idealized star in rigid rotation in vacuo. *Annales d'Astrophysique*, **18**, 1–10.
- Dhillon, V., Marsh, T., and the ULTRACAM team (2001). ULTRACAM - an ultra-fast, triple-beam CCD camera. In B. Gaensicke, K. Beuermann, and K. Reinsch, editors, *The Physics of Cataclysmic Variables and Related Objects*. (astro-ph/0110020).
- Duncan, R. C. and Thompson, C. (1992). Formation of very strongly magnetized neutron stars: Implications for gamma-ray bursts. *ApJ*, **392**, L9–L13.
- Everett, J. E. and Weisberg, J. M. (2001). Emission Beam Geometry of Selected Pulsars Derived from Average Pulse Polarization Data. *ApJ*, **553**, 341–357.
- Ferrario, L., Vennes, S., Wickramasinghe, D. T., Bailey, J. A., and Christian, D. J. (1997). EUVE J0317-855 A rapidly rotating, high-field magnetic white dwarf. *MNRAS*, **292**, 205–217.
- Gavriil, F. P. and Kaspi, V. M. (2002). Long-Term RXTE Monitoring of Anomalous X-ray Pulsars. *ApJ*, **567**, 1067–1076.
- Gotthelf, E. V. and Wang, Q. D. (2000). A Spatially Resolved Plerionic X-Ray Nebula around PSR B0540-69. *ApJ*, **532**, L117–L120.
- Gould, D. M. and Lyne, A. G. (1998). Multifrequency polarimetry of 300 radio pulsars. *MNRAS*, **301**, 235–260.
- Heise, J., van Teeseling, A., and Kahabka, P. (1994). The spectra and luminosity of super-soft X-ray sources. *A&A*, **288**, L45–L48.
- Helfand, D. J., Gotthelf, E. V., and Halpern, J. P. (2001). Vela Pulsar and Its Synchrotron Nebula. *ApJ*, **556**, 380–391.
- Henden, A. A. and Honeycutt, R. K. (1997). Secondary Photometric Standards for Northern Cataclysmic Variables and Related Objects. *PASP*, **109**, 441–460.
- Hulleman, F., van Kerkwijk, M. H., and Kulkarni, S. R. (2000a). An optical counterpart to the anomalous X-ray pulsar 4U0142+61. *Nature*, **408**, 689–692.
- Hulleman, F., van Kerkwijk, M. H., Verbunt, F. W. M., and Kulkarni, S. R. (2000b). A deep search for the optical counterpart to the anomalous X-ray pulsar 1E 2259+58.6. *A&A*, **358**, 605–611.
- Janesick, J. R. (2001). *Scientific Charge-Coupled Devices*. SPIE Press, Bellingham, Washington.

- Kijak, J. and Gil, J. (1997). Radio emission altitudes in pulsar magnetospheres. *MNRAS*, **288**, 631–637.
- Koptsevich, A. B., Pavlov, G. G., Zharikov, S. V., Sokolov, V. V., Shibano, Y. A., and Kurt, V. G. (2001). Optical photometry of the PSR B0656+14 and its neighborhood. *A&A*, **370**, 1004–1016.
- Kouveliotou, C., Dieters, S., Strohmayer, T., van Paradijs, J., Fishman, G. J., Meegan, C. A., Hurley, K., Kommers, J., Smith, I., Frail, D., and Murakami, T. (1998). An X-ray pulsar with a superstrong magnetic field in the soft γ -ray repeater SGR 1806-20. *Nature*, **393**, 235–237.
- Kristian, J. (1991). No Pulsar in Supernova 1987A. *Nature*, **349**, 747.
- Kristian, J., Pennypacker, C. R., Morris, D. E., Muller, R. A., Middleditch, J., Hamuy, M. A., Kunkel, W. E., Imamura, J. N., Lucinio, R., and Steiman-Cameron, T. Y. (1989). Submillisecond optical pulsar in supernova 1987A. *Nature*, **338**, 234–236.
- Lyne, A. G. and Graham-Smith, F. (1998). *Pulsar Astronomy*. Cambridge University Press, Cambridge.
- Lyne, A. G. and Manchester, R. N. (1988). The shape of pulsar radio beams. *MNRAS*, **234**, 477–508.
- Lyutikov, M., Blandford, R. D., and Machabeli, G. (1999). On the nature of pulsar radio emission. *MNRAS*, **305**, 338–352.
- Machabeli, G. Z., Luo, Q., Melrose, D. B., and Vladimirov, S. (2000). A new mechanism for pulsar gamma-ray emission. *MNRAS*, **312**, 51–56.
- Manchester, R. N., Wallace, P. T., Peterson, B. A., and Elliott, K. H. (1980). The optical pulse profile of the Vela pulsar. *MNRAS*, **190**, 9P–11P.
- Marsden, D., Lingenfelter, R. E., Rothschild, R. E., and Higdon, J. C. (2001). Nature versus Nurture: The Origin of Soft Gamma-Ray Repeaters and Anomalous X-Ray Pulsars. *ApJ*, **550**, 397–409.
- Marshall, H. L. and Schulz, N. S. (2002). Using the High Resolution X-ray Spectrum of PSR B0656+14 to Constrain the Chemical Composition of the Neutron Star Atmosphere. In *14 pages, 5 figures, accepted for publication in the Astrophysical Journal. Preprint no. CSR-02-12.*, pages 3463+.
- Mereghetti, S. (2001). The anomalous X-ray pulsars. In C. Kouveliotou, J. Ventura, and E. P. J. van den Heuvel, editors, *The Neutron Star - Black Hole Connection*, pages 351–368. Kluwer Academic Publishers, Dordrecht. (astro-ph/9911252).
- Middleditch, J. and Pennypacker, C. (1985). Optical pulsations in the Large Magellanic Cloud remnant 0540-69.3. *Nature*, **313**, 659–661.

- Middleditch, J., Kristian, J. A., Kunkel, W. E., Hill, K. M., Watson, R. D., Lucinio, R., Imamura, J. N., Steiman-Cameron, T. Y., Shearer, A., Butler, R., Redfern, M., and Danks, A. C. (2000). Rapid photometry of supernova 1987A: a 2.14 ms pulsar? *New Astronomy*, **5**, 243–283.
- O’Brien, K., Horne, K., Boroson, B., Still, M., Gomer, R., Oke, J. B., Boyd, P., and Vrtillek, S. D. (2001). Keck II spectroscopy of mHz quasi-periodic oscillations in Hercules X-1. *MNRAS*, **326**, 1067–1075.
- Paczyński, B. (1990). X-ray pulsar 1E 2259+586:A merged white dwarf with a 7 second rotation period? *ApJ*, **365**, L9–L12.
- Pechenick, K. R., Ftaclas, C., and Cohen, J. M. (1983). Hot spots on neutron stars: The near-field gravitational lens. *ApJ*, **274**, 846–857.
- Pedersen, H. and *et al.* (1982). Simultaneous optical and X-ray bursts from 4U/MXB 1636-53. *ApJ*, **263**, 325–339.
- Perna, R. and Hernquist, L. (2000). Disks Irradiated by Beamed Radiation from Compact Objects. *ApJ*, **544**, L57–L60.
- Perna, R., Hernquist, L., and Narayan, R. (2000). Emission Spectra of Fallback Disks around Young Neutron Stars. *ApJ*, **541**, 344–350.
- Pirola, V. and Reiz, A. (1992). The highly magnetic ($B \sim 500$ MG) white dwarf PG 1031+234: Discovery of wavelength dependent polarization and intensity variations. *A&A*, **259**, 143–148.
- Radhakrishnan, V. and Cooke, D. J. (1969). Magnetic poles and the polarization structure of pulsar radiation. *Astrophys. Lett.*, **3**, 225–229.
- Ramanamurthy, P. V., Fichtel, C. E., Kniffen, D. A., Sreekumar, P., and Thompson, D. J. (1996). Possible Evidence for Pulsed Emission of High-Energy Gamma Rays by PSR B0656+14. *ApJ*, **458**, 755+.
- Rankin, J. M. (1990). Toward an empirical theory of pulsar emission. IV - Geometry of the core emission region. *ApJ*, **352**, 247–257.
- Rauch, T. (1996). Implications of Light Metals (Li-Ca) on NLTE Model Atmospheres for Hot Stars. In J. Greiner, editor, *Supersoft X-ray Sources*, Lecture Notes in Physics 472, pages 139–142. Springer, Berlin.
- Reid, C. A., Johnston, M. D., Bradt, H. V., Doxsey, R. E., Griffiths, R. E., and Schwartz, D. A. (1980). Unidentified galactic X-ray sources positioned with the HEAO-1 scanning modulation collimator. *AJ*, **85**, 1062–1070.

- Romani, R. W. (1996). Gamma-Ray Pulsars: Radiation Processes in the Outer Magnetosphere. *ApJ*, **470**, 469+.
- Romani, R. W. and Yadigaroglu, I.-A. (1995). Gamma-ray pulsars: Emission zones and viewing geometries. *ApJ*, **438**, 314–321.
- Romani, R. W., Miller, A. J., Cabrera, B., Nam, S. W., and Martinis, J. M. (2001). Phase-resolved Crab Studies with a Cryogenic Transition-Edge Sensor Spectrophotometer. *ApJ*, **563**, 221–228.
- Shaviv, N. J., Heyl, J. S., and Lithwick, Y. (1999). Magnetic lensing near ultramagnetized neutron stars. *MNRAS*, **306**, 333–347.
- Shearer, A., Redfern, R. M., Gorman, G., Butler, R., Golden, A., O’Kane, P., Beskin, G. M., Neizvestny, S. I., Neustroev, V. V., Plokhotnichenko, V. L., and Cullum, M. (1997). Pulsed Optical Emission from PSR 0656+14. *ApJ*, **487**, L181–+.
- Shearer, A., Golden, A., Harfst, S., Butler, R., Redfern, R. M., O’Sullivan, C. M. M., Beskin, G. M., Neizvestny, S. I., Neustroev, V. V., Plokhotnichenko, V. L., Cullum, M., and Danks, A. (1998). Possible pulsed optical emission from Geminga. *A&A*, **335**, L21–L24.
- Simmons, J. F. L. and Stewart, B. G. (1985). Point and interval estimation of the true unbiased degree of linear polarization in the presence of low signal-to-noise ratios. *A&A*, **142**, 100–106.
- Smith, F. G., Jones, D. H. P., Dick, J. S. B., and Pike, C. D. (1988). The optical polarization of the Crab Pulsar. *MNRAS*, **233**, 305–319.
- Staelin, D. H. and Reifenstein, III, E. C. (1968). Pulsating radio sources near the Crab Nebula. *Science*, **162**, 1481.
- Tennant, A. F., Becker, W., Juda, M., Elsner, R. F., Kolodziejczak, J. J., Murray, S. S., O’Dell, S. L., Paerels, F., Swartz, D. A., Shibasaki, N., and Weisskopf, M. C. (2001). Discovery of X-Ray Emission from the Crab Pulsar at Pulse Minimum. *ApJ*, **554**, L173–L176.
- Thompson, D. J. (2001). Gamma-ray Pulsars: Observations. In *High Energy Gamma-Ray Astronomy*, pages 103+.
- Wald, A. and Wolfowitz, J. (1943). An exact test for randomness in the non-parametric case based on serial correlation. *Annals of Mathematical Statistics*, **14**, 378–388.
- Wallace, P. T. *et al.* (1977). Detection of optical pulses from the Vela pulsar. *Nature*, **266**, 692.
- Wardle, J. F. C. and Kronberg, P. P. (1974). The linear polarization of quasi-stellar radio sources at 3.71 and 11.1 centimeters. *ApJ*, **194**, 249–255.

- Weisberg, J. M., Cordes, J. M., Lundgren, S. C., Dawson, B. R., Despotes, J. T., Morgan, J. J., Weitz, K. A., Zink, E. C., and Backer, D. C. (1999). Arecibo 1418 MHz Polarimetry of 98 Pulsars: Full Stokes Profiles and Morphological Classifications. *ApJS*, **121**, 171–217.
- White, N. E., Angelini, L., Ebisawa, K., Tanaka, Y., and Ghosh, P. (1996). The Spectrum of the 8.7s X-Ray Pulsar 4U0142+61. *ApJ*, **463**, L83–L86.
- Wilson, C. A., Dieters, S., Finger, M. H., Scott, D. M., and van Paradijs, J. (1999). *Rossi X-Ray Timing Explorer* Observations of the Anomalous Pulsar 4U 0142+61. *ApJ*, **513**, 464–470.
- Yadigaroglu, I. A. (1997). *Theory and Modeling of Gamma-Ray Pulsars*. Ph.D. thesis, Stanford University.
- Zhang, L. and Cheng, K. S. (2001). Gamma-ray pulsars: the pulse profiles and phase-resolved spectra of Geminga. *MNRAS*, **320**, 477+.