

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

At High Resolution: The Crab Pulsar, Kinematics of ULIRGs, and Linear Quadratic Gaussian Control for Adaptive Optics

Permalink

<https://escholarship.org/uc/item/0q65j9b9>

Author

Rudy, Alexander Raymond

Publication Date

2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**AT HIGH RESOLUTION: THE CRAB PULSAR, KINEMATICS OF
ULIRGS, AND LINEAR QUADRATIC GAUSSIAN CONTROL FOR
ADAPTIVE OPTICS**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

ASTRONOMY & ASTROPHYSICS

by

Alexander R. Rudy

September 2017

The Dissertation of Alexander R. Rudy
is approved:

Professor Claire E. Max, Chair

Doctor Lisa A. Poyneer

Assistant Professor Andrew Skemer

Tyrus Miller
Vice Provost and Dean of Graduate Studies

Copyright © by
Alexander R. Rudy
2017

Table of Contents

List of Figures	vii
List of Tables	xi
Abstract	xiii
Dedication	xiv
Acknowledgments	xv
1 Introduction	1
I On the origin of γ-ray flares from the Crab Nebula	4
2 Characterization of the Inner Knot of the Crab: The Site of the γ-ray Flares?	5
2.1 Introduction	5
2.2 The Observations	7
2.2.1 Keck	7
2.2.2 HST	8
2.2.3 Chandra	9
2.3 Properties of the Knot	10
2.3.1 Keck	10
2.3.2 Chandra	12
2.4 Comparison amongst the knot observations	13
2.5 Comparison amongst observations of the inner knot	13
2.6 Comparison with γ -ray fluxes	14
2.7 Implications for Theoretical Models	17
2.7.1 Variability	19
2.7.2 Time averaged properties of the knot	20
2.8 Conclusions	23
2.A Appendix - Data Analysis	23
2.A.1 Keck	23

2.A.2	HST	25
2.A.3	Chandra	28
II	Searching for Outflows in the kinematics of nearby ULIRGs	53
3	Kinematics of the Nuclear Regions of ULIRG F20414-1651	54
3.1	Introduction	54
3.2	Methods	58
3.2.1	Observations	58
3.2.2	Data Reduction	59
3.2.3	Line Fitting	60
3.2.4	Modeling the Kinematics	61
3.2.5	Kinometry	61
3.2.6	Modeling the OH 119 μ m emission	61
3.3	Results	62
3.3.1	Hydrogen Recombination Lines	63
3.3.2	H ₂ Emission	69
3.3.3	OH 119 μ m Absorption Fit	72
3.3.4	Disk Parameters	72
3.3.5	Disk Mass	73
3.3.6	Non-disklike Emission	76
3.4	Discussion	77
3.4.1	Comparison of the Integrated Emission	78
3.4.2	Possible Explanations of Kinematic Structure	79
3.4.3	Implications for fast Outflows	81
3.5	Conclusions	81
4	A Survey of Nearby ULIRGs with Outflow Signatures	83
4.1	Introduction	83
4.2	Sample	85
4.3	Methods	85
4.3.1	Data Reduction	87
4.3.2	Line Fitting	87
4.4	Results	88
4.4.1	F00188-0856	88
4.4.2	F00456-2904	88
4.4.3	F03158+4227	91
4.4.4	F05189-2524	91
4.4.5	F08572+3915	100
4.4.6	F10565+2448	100
4.4.7	F14348-1447	102
4.4.8	F15206+3342	112
4.4.9	F19297-0406	112
4.5	Discussion	120

4.6	Conclusions	121
4.6.1	ULIRGs in group 1	122
4.6.2	ULIRGs in group 2	122
4.6.3	ULIRGs in group 3	123
4.6.4	Concluding thoughts	123
 III Linear Quadratic Gaussian Control for Adaptive Optics		125
5	A laboratory demonstration of an LQG controller for correcting frozen flow turbulence in adaptive optics systems	126
5.1	Introduction	126
5.2	Experimental Setup	128
5.2.1	Modifications to the Experimental Setup to Support Predictive Fourier Control	130
5.2.2	Error Budget for the Experimental Setup	131
5.3	Theoretical Framework	133
5.3.1	Fourier Wind Identification	134
5.3.2	LQG-Based Predictive Fourier Control	138
5.4	Experimental Results	139
5.4.1	Modeling Transfer Functions	139
5.4.2	Integrating the LQG Controller	140
5.4.3	Demonstrating Stable LQG Control on the MCAO Test Bench	141
5.4.4	Applying FWI on the MCAO Test Bench	142
5.4.5	Applying PFC on the MCAO Test Bench	142
5.5	Discussion and Future Directions	144
5.6	Conclusions	145
6	Linear Quadratic Gaussian Control with ShaneAO	146
6.1	Introduction	146
6.1.1	Overview of ShaneAO	146
6.1.2	Reconstruction Scheme	148
6.1.3	Software Architecture	153
6.2	Verifying the ShaneAO Reconstructor	154
6.2.1	Self-Consistent Error Transfer Model	156
6.2.2	ShaneAO Simulator	158
6.2.3	Performance on the Internal Light Source	160
6.2.4	On Sky Performance	163
6.2.5	Additional Verification	164
6.3	Restructuring the ShaneAO Reconstructor	169
6.3.1	Motivation	169
6.3.2	New Reconstructor Framework	170
6.3.3	Streaming Telemetry to Disk	177
6.3.4	New Reconstructor Verification	177

6.4	Modifying ShaneAO for PFC	187
6.4.1	Reconstructing Phase	187
6.4.2	Splitting phase onto Tweeter and Woofer	188
6.4.3	Verifying the Combined Matrix	189
6.5	Hybrid Matrix Reconstructor	189
6.5.1	Hybrid Matrix Design	190
6.5.2	Unit Test Verification	191
6.5.3	Verification with the ShaneAO Simulator	193
6.5.4	Performance with the Internal Light Source	193
6.5.5	On Sky ETFs and Performance	194
6.6	Integrator-Differentiator Reconstructor	198
6.6.1	Verification	200
6.6.2	Performance with the Internal Light Source	200
6.6.3	On Sky ETFs and Performance	200
6.7	Kalman Filter Reconstructor	206
6.7.1	Kalman Reconstructor Design	206
6.7.2	Unit Test Verification	208
6.7.3	Performance with the Internal Light Source	209
6.7.4	On Sky ETFs and Performance	210
6.7.5	Kalman Reconstructor Verification	212
6.7.6	Future Directions with the Kalman Filter	213
6.8	Conclusions	218
7	Conclusions	219
	Bibliography	222

List of Figures

2.1	Images of Pulsar and reference star from Keck NIRC2	39
2.2	SVD processed images of the Pulsar from HST ACS/WFC3	41
2.3	Keck knot properties vs. separation between knot and pulsar . . .	42
2.4	HST-SVD knot properties vs. separation between knot and pulsar	42
2.5	Chandra image of pulsar at pulse minimum	43
2.6	Fermi/LAT flux vs time with HST and Keck observations.	44
2.7	Fermi/LAT flux during flares vs time	45
2.8	Fermi/LAT flux vs pulsar-knot separations	46
2.9	Pulsar-knot separations vs time	46
2.10	Termination shock geometry	47
2.11	Knot geometric properties and coordinates on sky	47
2.12	PSF Subtraction from HST/ACS	48
2.13	Chandra image during flares	49
2.14	Further processing of HST/ACS images	50
2.15	Mapping knot onto $\theta - r$ grid	51
2.16	Variation of the HST derived properties of the knot, comparing SVD and traditional analyses.	52
3.1	IRAS F20414-1651 image from HST WFPC3 F814W	56
3.2	Kbb continuum flux from fit	60
3.3	Kinometry Ellipse Fit	62
3.4	Spatially-integrated Pa- α velocity profile	64
3.5	Integrated Pa- α profile	65
3.6	Pa- α flux from the three components.	66
3.7	Comparison of the Pa- α velocity in the disk (A) vs component B .	67
3.8	Comparison of the Pa- α dispersion in the disk vs component B .	68
3.9	Pa- α ratio of velocity to velocity dispersion in the disk (A)	68
3.10	Radial profile for the disk (A)	69
3.11	H ₂ 1 – 0 S(3) flux	70
3.12	H ₂ 1 – 0 S(3) Velocity	71
3.13	H ₂ 1 – 0 S(3) Velocity Dispersion	71
3.14	H ₂ spatially-integrated velocity profile	71
3.15	OH 119 μ m Velocity Profile	72

3.16	Posterior for OH 119 μm v_{98}	73
3.17	Kinemetry model for the primary Pa- α emission component.	74
3.18	Pa- α Rotation Curve from Kinemetry	75
3.19	Kinemetry flattening and position angle	75
3.20	Uniform density disk fit to A component	76
3.21	Cutout of velocity in Pa- α from component B	77
3.22	OH 119 μm and Pa- α integrated absorption and emission profile.	78
4.1	F00188-0856 Pa- α emission from component A	89
4.2	F00188-0856 Pa- α emission from component B	90
4.3	F00456-2904 Pa- α emission from component A	92
4.4	F00456-2904 H ₂ 1 – 0 $S(3)$ emission from component A	93
4.5	F03158+4227 Pa- α emission from component A	94
4.6	F03158+4227 Pa- α emission from component B	95
4.7	F03158+4227 H ₂ 1 – 0 $S(3)$ emission from component A	96
4.8	F05189-2524 Br- γ emission from component A	97
4.9	F05189-2524 Br- γ emission from component B	98
4.10	F05189-2524 H ₂ 1 – 0 $S(1)$ emission from component A	99
4.11	F08572+3915 Pa- α emission from component A	100
4.12	F08572+3915 Pa- α emission from component B	101
4.13	F10565+2448 Br- δ emission from component A	103
4.14	F10565+2448 Br- γ emission from component A	104
4.15	F10565+2448 H ₂ 1 – 0 $S(3)$ emission from component A	105
4.16	F10565+2448 H ₂ 1 – 0 $S(1)$ emission from component A	106
4.17	F14348-1447 Pa- α emission from component A	107
4.18	F14348-1447 Pa- α emission from component B	108
4.19	F14348-1447 Br- γ emission from component A	109
4.20	F14348-1447 H ₂ 1 – 0 $S(3)$ emission from component A	110
4.21	F14348-1447 H ₂ 1 – 0 $S(1)$ emission from component A	111
4.22	F15206+3342 Pa- α emission from component A	113
4.23	F15206+3342 Pa- α emission from component B	114
4.24	F15206+3342 Br- δ emission from component A	115
4.25	F15206+3342 Br- δ emission from component B	116
4.26	F19297-0406 Pa- α emission from component A	117
4.27	F19297-0406 H ₂ 1 – 0 $S(3)$ emission from component A	118
4.28	F19297-0406 H ₂ 1 – 0 $S(1)$ emission from component A	119
5.1	Testbed layout for wind prediction tests	129
5.2	High-level block diagram for Predictive Fourier Control	134
5.3	Temporal PSD of three Fourier modes which show evidence of frozen flow turbulence	136
5.4	Wind Identification Metric on test bench using the FWI algorithm	136
5.5	FWI frozen flow identification metric	137
5.6	An empirical and model error transfer function for the test bench	140

5.7	The RMS residual wavefront error for both the PFC Kalman filter and a standard integral controller	143
6.1	Diagram of the ShaneAO Vector Matrix Multiply (VMM)	149
6.2	Modal gain of the reconstructor matrix	151
6.3	Theoretical ETFs for the VMM reconstructor	159
6.4	Estimated phase (\hat{p}) Residual Timeseries from ShaneAO Simulator	161
6.5	\hat{p} Residual Timeseries from ShaneAO ILS	162
6.6	\hat{p} Residual Pupil Images from ShaneAO ILS	162
6.7	Histogram of \hat{p} Values from ShaneAO ILS	163
6.8	PSD of a single \hat{p} value from ShaneAO ILS	164
6.9	ETF of a single \hat{p} value from ShaneAO ILS	165
6.10	Median \hat{p} Residuals Timeseries from ShaneAO on sky	166
6.11	PSD of a single \hat{p} value from ShaneAO on sky	166
6.12	ETF of a single \hat{p} value from ShaneAO on sky	167
6.13	Phase shifted WFS signal measuring system delay	168
6.14	ShaneAO Reconstructor Block Diagram	170
6.15	Full ShadyAO Structure	173
6.16	New Reconstructor WFS GUI	175
6.17	New Reconstructor Mirror GUI	176
6.18	New Reconstructor settings GUI	176
6.19	Histogram of inter-frame arrival times for the new VMM reconstructor	179
6.20	Histogram of computation times for the new VMM reconstructor	180
6.21	PSD of \hat{p} from the new VMM on Sauce	181
6.22	ETF of \hat{p} from the new VMM on Sauce	182
6.23	Periodogram of \hat{p} averaged for the new VMM on the ILS	183
6.24	ETF of \hat{p} from ILS with the new VMM	184
6.25	Periodogram of \hat{p} averaged for the new VMM on Sky	185
6.26	ETF of \hat{p} on Sky with the new VMM	186
6.27	Modal gain of the reconstructor matrix	190
6.28	Diagram of the ShaneAO Hybrid Matrix Multiply (HMM)	191
6.29	Histogram of computation times for HMM	192
6.30	Histogram of inter-frame arrival times for HMM	193
6.31	PSD of \hat{p} averaged for the HMM on the ILS.	194
6.32	ETF of \hat{p} for the HMM on the ILS.	195
6.33	Periodogram of \hat{p} averaged for the HMM on Sky	196
6.34	ETF of \hat{p} on Sky with the HMM	197
6.35	Structure of the Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) reconstructor	198
6.36	Pairs of integrators and differentiators in IDMM	199
6.37	Internal Transfer Function for the IDMM	201
6.38	PSD of \hat{p} averaged for the IDMM on the ILS.	202
6.39	ETF of \hat{p} for the IDMM on the ILS.	203
6.40	Periodogram of \hat{p} averaged for the IDMM on Sky	204

6.41	ETF of \hat{p} on Sky with the IDMM	205
6.42	Structure of the LQG reconstructor	206
6.43	Theoretical Error Transfer Function for the LQG Filter	208
6.44	Histogram of computation times for the Kalman Filter	209
6.45	Histogram of inter-frame arrival times for the Kalman Filter	210
6.46	Internal ETF for the Integrator Kalman Filter Matrix Multiply Re- constructor (KFMM)	211
6.47	Internal ETF for the DC-only KFMM	211
6.48	PSD of $F(\hat{p})$ averaged for the Passthrough KFMM on the TelSim.	212
6.49	ETF of $F(\hat{p})$ for the Passthrough KFMM on the ILS.	214
6.50	PSD of $F(\hat{p})$ averaged for the Passthrough KFMM on Sky	215
6.51	ETF of Fourier transform of the estimated phase ($F(\hat{p})$) on Sky with the Passthrough KFMM	216
6.52	Internal ETFs for the KFMM	217

List of Tables

2.1	Summary of the Keck observations	29
2.2	Summary of the HST observations	30
2.3	Summary of the Chandra observations	30
2.4	Knot properties based on the analysis of the Keck data	31
2.5	Results of regression analyses for the 12 Keck observations.	32
2.6	Results of regression analyses for 11 Keck observations removing K07.	33
2.7	Results of HST-SVD analysis of the knot	34
2.8	Results of regression analysis HST-SVD analysis of the knot	35
2.9	Chandra analysis results	36
2.10	Fermi/LAT fluxes and knot separations	37
2.11	Results of HST-Traditional analysis of the knot	38
2.12	Results of regression analyses for HST-Traditional	40
3.1	Observations of IRAS F20414-1651	59
3.2	Emission lines in IRAS F20414-1651	63
4.1	Sample of Nearby ULIRGs	85
4.2	List of Observations	86
4.3	List of emission lines detected for each target	88
5.1	Error Budget for the test bench.	133
6.1	Facility Values for the ShaneAO Reconstructor	156
6.2	Model Parameters for the ETF from the ILS	165
6.3	Model Parameters for the ETF on Sky	167
6.4	Model Parameters for the ETF of \hat{p} the new VMM on Sauce	182
6.5	Model Parameters for the ETF from the ILS	184
6.6	Model Parameters for the ETF from the ILS	186
6.7	Model parameters for the ETF of \hat{p} for the HMM on the ILS.	195
6.8	Model Parameters for the ETF of the HMM on Sky	197
6.9	Model parameters for the ETF of \hat{p} for the IDMM on the ILS.	203
6.10	Model Parameters for the ETF of the IDMM on Sky	205
6.11	Model for the ETF of $F(\hat{p})$ for the Passthrough KFMM on the ILS	214

6.12 Model Parameters for the ETF of the Passthrough KFMM on Sky 216

Abstract

At High Resolution: The Crab Pulsar, Kinematics of ULIRGs, and Linear Quadratic Gaussian Control for Adaptive Optics

by

Alexander R. Rudy

Adaptive Optics is a technique to correct for the aberrations introduced by the Earth's atmosphere in telescope images and spectroscopy. This work covers three projects which are facilitated by adaptive optics. Adaptive optics (AO) enables high resolution imaging from the ground, which we used to explore the origins of γ -ray flares from the Crab Nebula. These flares represent some of the brightest transient events ever observed. Using AO, we were able to identify features of the Crab Nebula which appear to correlate with the large γ -ray flares. We searched for the presence and mechanisms behind high velocity molecular outflows from 10 nearby gas-rich galaxies. These galaxies were a subset of galaxies found to have evidence of high velocity outflows in spatially un-resolved spectra of the OH 119 μ m absorption line observed by the Herschel Space Telescope. We were able to identify spatially resolved outflows in only 3 of the 10 galaxies observed, and found that 3 additional galaxies had such extreme velocity rotation curves that no outflow was required to explain the spatially un-resolved observations that motivated the original study. Finally, this work develops and implements a novel algorithm for control of adaptive optics systems. Using the fact that the dynamism of atmospheric turbulence is driven by the flow of wind across the telescope pupil, we developed a predictive control algorithm using linear-quadratic gaussian control. We were able to demonstrate stable linear quadratic gaussian control on sky with the ShaneAO adaptive optics system at Lick Observatory.

To my wife, Esther, and Sokka, the rabbit.

Acknowledgments

I owe a great deal to the people in my life who have propelled me to this point. Thank you to my wife, Esther, thank you for your unwavering support. I owe a debt of gratitude to my parents, who have believed in me, and pushed me to pursue my dreams.

Pomona College provided an incredible grounding to launch my scientific career. My undergraduate advisor, Professor Phil Choi, convinced me that Astronomy is better than Physics, and let me climb all over a telescope, which pushed me towards building instruments. The Physics and Astronomy department taught me how to approach problems, how to make discoveries on my own, and showed me a standard of teaching that I have yet to see elsewhere.

My advisor, Claire Max, has been invaluable in my development as a scientist. She has provided an incredible mixture of advice, opportunity, and helpful prodding. Lisa Poyneer has helped me to develop into a competent engineer. I am still in awe of her abilities, and grateful for the chance to learn from such an amazing engineer. Don Gavel has been invaluable in our many conversations about adaptive optics and his willingness to let me take ShaneAO apart. Connie Rockosi has provided a level head, and advice from outside of my usual advisors. Reni Kupke and Daren Dillion have provided incredible instrumentation support, and the wisdom of years of hands on experience that is invaluable.

My scientific development has been shaped and guided by the graduate students who came before me. Special thanks to Anne Medling, whose invaluable expertise facilitated all of my work with OSIRIS. Rosalie McGurk provided the guidance and motivation I needed in my first few years.

My graduate student colleagues, especially Emily Cunningham, Zach Jennings, Katie Hamren, Claire Dorman, and Caroline Morley, were wonderfully supportive, great friends and helpful colleagues throughout this work. All of the graduate students in the UCSC Astronomy & Astrophysics department make it a truly special place to be.

Paper Acknowledgments

The text of this dissertation includes reprints of the following published and in-preparation material by Rudy: Rudy et al. (2015), Rudy et al. (2014).

Chapter 2 contains a reprint of Rudy et al. (2015). I was responsible for obtaining the Keck observations, conducting the analysis of the Keck observations, and drafting the paper. I would like to thank my co-authors for their work and dedication to the paper, especially Claire Max, Martin Weisskopf, Jeff Kolodziejczak, Rolf Buehler and Jon Arons, whose varied contributions were invaluable in bringing the paper to fruition.

Chapters 3 and 4 represent in-preparation material for which I am the primary author, and for which I was responsible for obtaining the Keck observations, conducting the data analysis, and drafting the papers. This work has been extensively supported by Claire Max who has provided immeasurable advice and guidance, Anne Medling, Vivian U, and Rosalie McGurk, who taught me how to use OSIRIS effectively and who laid the analysis groundwork for the material presented here.

Chapter 5 contains a preprint of a paper for which I am the first author, but for which significant work was contributed by Srikar Srinath, Lisa Poyneer, Don Gavel, Marc Reinig, Luke Johnson and S. Mark Ammons. I am grateful to all of them for their incredible support on this project.

Chapter 6 builds on work that was completed with Lisa Poyneer and Don Gavel, whose help has been invaluable in the process of implementing linear quadratic gaussian control on ShaneAO. That chapter would not have been possible without the hard work of the many people who built ShaneAO. I owe extreme gratitude for their tireless efforts on ShaneAO to Don Gavel, Daren Dillon, Reni Kupke, Constance Rockosi, Elinor Gates and Rosalie McGurk. Chapter 6 would not have been possible without the support of Kyle Lanclos, Will Deich, and Steve Allen of the UCO Software and Programming Group. Working with the incredible staff on Mt. Hamilton has been a pleasure, especially the two amazing support astronomers, Elinor Gates and Paul Lynam. As well, I owe thanks to the telescope technicians, Erik Kovacs, Shawn Stone, Matt Brooks and Donnie Redel, who have put up with my batty observing strategies.

Science Acknowledgments

Data presented herein were obtained at Lick Observatory, operated by the University of California Observatories, and supported by the University of California. I wish to extend thanks to the Muwekma Ohlone tribe, on whose ancestral land Lick Observatory is built.

Data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory and the Keck II Laser Guide Star AO system were both made possible by the generous financial support of the W. M. Keck Foundation. This research has also made use of the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExSci), under contract with the National Aeronautics and Space Administration.

I wish to extend special thanks to those of Hawaiian ancestry, on whose sacred mountain we are privileged to be guests. Without their generous hospitality, the observations would not have been possible. The observatories of Mauna Kea have contributed significantly to our understanding and wonder of the universe.

Data in this dissertation were obtained with Herschel. Herschel is an ESA space observatory with science instruments provided by European-led Principal Investigator consortia and with important participation from NASA.

Some of the data presented in this dissertation were obtained from the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Support for MAST for non-HST data is provided by the NASA Office of Space Science via grant NNX09AF08G and by other grants and contracts

This research has made use of the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

This research made extensive use of *Astropy*, a community-developed core Python package for Astronomy (Astropy Collaboration et al., 2013).

Financial Support

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1339067.

Some of the work presented here was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was funded by the UC Lab Fees Research Program grant 12-LF-236852.

The research presented in Part III was funded in part by the National Science Foundation, Major Research Instrumentation grant No. 0923585 and by the Gordon and Betty Moore Foundation.

The *Fermi* LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l’Energie Atomique and the Centre National de la Recherche Scientifique / Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden.

Chapter 1

Introduction

AO corrects for the aberrations introduced by the Earth's atmosphere in telescope images and spectroscopy. Using AO, ground based astronomy can provide diffraction-limited imaging through Earth's atmosphere, taking full advantage of the large aperture telescopes which represent the state of the art in ground based astronomy. This thesis presents three projects which rely on the unique capabilities and advantages of AO.

In Part I, we try to identify the origin of γ -ray flares from the Crab Nebula. In 2007, the Fermi γ -ray telescope observed large γ -ray flares from the Crab Nebula. These flares were some of the brightest transient events ever observed. However, the detection at γ -ray wavelengths provided only limited position and spectral information. The flares have not been directly detected in wavelengths other than the γ -ray. We search for secondary effects of the γ -ray flares. Simple causality arguments suggest that the γ -ray flares must originate from an area subtending an angle $\sim 0.3''$ at the distance to the Crab Nebula (2 kpc). This makes the origin of the γ -ray flares from the Crab Nebula a good target for AO imaging in the near-infrared (NIR), where we can resolve objects which are $\sim 0.3''$ in size.

Part II uses AO at the W.M. Keck observatory to study the kinematics of the inner regions of nearby ultra-luminous infrared galaxies (ULIRGs). Much effort has been expended to understand the way that massive, gas-rich galaxies quench star formation. This quenching is an important piece of many cosmological models, and helps to explain why we observe a peak in star for-

mation around $z \sim 2$, followed by the decrease in star formation in massive galaxies. Spoon et al. (2013) and Veilleux et al. (2013b) presented evidence of numerous high velocity molecular outflows, observed in OH $119\mu\text{m}$ absorption with Herschel. Although OH $119\mu\text{m}$ absorption provides an interesting clue as to the presence of outflowing gas, the results from Spoon et al. (2013) and Veilleux et al. (2013b), which showed strikingly high velocities (many greater than 1000 km/s) cannot confirm the presence, orientation, or origins of the outflows seen in nearby ULIRGs. We conducted a survey of 10 of the ULIRGs observed by Spoon et al. (2013) and Veilleux et al. (2013b) to determine whether a significant number contained kinematically and spatially distinct outflows, and to understand the nature and origins of those outflows.

Parts I and II present scientific inquiry facilitated by high resolution AO observations. In Part III, we describe work to advance AO for the next generation of scientific instruments. The most recent generation of extreme adaptive optics systems have started to reach the limits of performance possible with standard control techniques for AO. Many of these systems are limited by temporal delay error, the time that it takes an AO system to respond to an input and make a corresponding correction. As an alternative to building a faster AO system, Part III describes work to eliminate the impact of the temporal delay error, using a predictive controller.

The observations in Part I proved to be more accurate than the equivalent observations from Hubble Space Telescope (HST). The survey of 10 nearby ULIRGs in Part II was only possible due to the unique capabilities of the Oh-Suppressing Infrared Integral Field Spectrograph (OSIRIS) at the W.M. Keck Observatory, an instrument which takes unique advantage of the high resolution capabilities of AO on a 10-m telescope. Finally, Part III aims to push the state-of-the-art in AO techniques forward, hopefully enabling scientific projects akin to those from Parts I and II in the future.

This is a thesis in three parts, held together by a deep understanding of the ways that adaptive optics can facilitate scientific discovery. The work presented here represents a range of scientific interest and avenues, all of which I have pursued with the facilities of the UC Observatories. The Keck telescopes, and their reliable adaptive optics systems made Chapters 2, 3 and 4 possible.

The Lab for Adaptive Optics provided the test bench for Chapter 5 and laid the groundwork for Chapter 6. Lick observatory provided the facility to develop my improvements for ShaneAO and to demonstrate linear quadratic gaussian (LQG) control on sky.

Part I

On the origin of γ -ray flares from the Crab Nebula

Chapter 2

Characterization of the Inner Knot of the Crab: The Site of the γ -ray Flares?

2.1 Introduction

The Crab pulsar and its nebula is one of the most studied targets in the sky at all wavelengths. It serves as a test bed for pulsar theories as well as, more generally, for astrophysical non-thermal processes. A very nice review of the Crab system, including high spatial resolution observations with HST and Chandra, showing the complex, dynamic interaction of the pulsar wind with the surrounding medium, has been compiled by Hester (2008). A more recent review by Bühler & Blandford (2014) also includes the saga of the γ -ray flares.

A crucial question implied by the observations of the γ -ray flares is “how are the particles being accelerated to sustain the flaring behavior?” Radiating particles must have PeV energies to emit synchrotron radiation at the observed γ -ray wavelengths. A mechanism is needed that can accelerate particles to such high energies in less than one gyration time in the ambient magnetic field. For example, the large flare of April 2011 produced a radiant energy equivalent to the energy contained within a region of order 2×10^{14} m subtending an angle $\sim 0.3''$ at 2 kpc, consistent with the emission-region size estimated

from light-travel-time arguments (Weisskopf et al., 2013).

An intense theoretical effort is currently being devoted to explain the variable Crab behavior (e.g., Uzdensky et al., 2011; Yuan et al., 2011; Arons, 2012; Sturrock & Aschwanden, 2012; Lyubarsky, 2012; Lyutikov et al., 2012; Bykov et al., 2012; Clausen-Brown & Lyutikov, 2012; Cerutti et al., 2012, 2013; Baty et al., 2013; Teraki & Takahara, 2013). However, the mechanism driving the flares, their impulsive nature, the ~ 12 -month recurrence time, and the location, remain unknown. Possible explanations include: plasma instabilities in the nebula, magnetic reconnection, discontinuity in the pulsar wind acceleration, emission from the anvil, emission from the inner knot, or portions of the termination shock. See, for example, the review by Hester (2008) for the association of specific features with this nomenclature.

There are a number of reasons to suspect that the feature known as the inner knot should be considered as a possible site for the origin of the flares. First, the luminosity of the flares after accounting for a relatively high fraction of the pulsar spin-down luminosity (assuming a radiation efficiency of 3-5%) suggesting an origin in a region close to the pulsar (Tavani et al., 2011). Second, the very short variability time scales combined with causality arguments, strongly constrain the size of the flare’s emitting region (Buehler et al., 2012; Striani et al., 2011).

A multi-wavelength campaign to study the Crab, described in Weisskopf et al. (2013), is being performed. Another γ -ray flare—the second largest observed to date—was detected with Fermi/LAT in March 2013 (Mayer et al., 2013). This triggered further observations with Keck, with HST, and with Chandra. Here we concentrate on what was learned about the inner knot as a consequence of the campaign. We note that very-high-energy ($E \geq 100$ GeV) contemporaneous data from the HESS (H. E. S. S. Collaboration et al., 2014) and VERITAS (Aliu et al., 2014) Cherenkov-telescope arrays found no significant correlation of TeV flux with the flare observed with Fermi/LAT.

The observations are presented in §2.2. The measured properties of the knot are given in §2.3. Comparison and examination of possible correlations of the knot’s properties with γ -rays detected using the Fermi/LAT are discussed in §2.6. §2.7 discusses theoretical implications. §2.8 summarizes our conclusions.

2.2 The Observations

The feature we refer to as the “inner knot” (aka “knot 1”, or “synchrotron knot”) was discovered by Hester et al. (1995) in a detailed study of the Crab using high resolution images taken with HST. The feature, although resolved in those data, is small ($\sim 0.1''$) and close to the pulsar ($\sim 0.6''$). The feature was first suggested to be associated with a shock or instability in the jet ~ 1500 AU from the pulsar, due to its good alignment with the jet and elongation perpendicular to the jet. Later Komissarov & Lyubarsky (2003) proposed, based on two dimensional MHD simulation of the pulsar wind nebula, that the knot can be explained by Doppler-boosted emission from an oblique termination shock. Synthetic synchrotron images from these 2D simulations (Komissarov & Lyutikov, 2011), and more recently 3D simulations (Porth et al., 2014), indeed show a knot-like feature when the system is viewed from the proper (and reasonable) orientation. Moreover, the observed high degree of polarization and a position angle aligned with the symmetry axis (Moran et al., 2013) seems to lend further support to the oblique shock scenario.

The salient features of the inner knot deduced prior to our observations are that it has a power-law spectrum (Sollerman, 2003) and, as most recently measured by Moran et al. (2013), the flux varies, the position appears to change, the degree of polarization is very large but does not appear to change with the flux (but see the discussion at the end of §2.6), and the magnetic field (determined from the position angle of the optical polarization) is orthogonal to the axis of symmetry (presumably the spin axis) and consistent with the direction of the field for the rest of the inner nebula.

2.2.1 Keck

We obtained K' and H band near infrared images of the inner $40''$ of the Crab nebula using the Keck Near Infrared Camera, NIRC2 (PI: K. Matthews), with Laser Guide Star Adaptive Optics (Wizinowich et al., 2006). From February 2012 to January 2014 we obtained 11 observations of the nebula during γ -ray quiescent periods, and 1 target-of-opportunity observation (K8) triggered by γ -ray flares. Table 2.1 summarizes the Keck observations.

Data reduction was carried out using the reduction pipeline developed by Ghez et al. (2008). Individual frames were sky-subtracted, flat-fielded, bad-pixel-corrected, and corrected for NIRC2’s camera distortion. We used the solution developed by Yelda et al. (2010), which matched the positions of stars in the globular cluster M92 measured with HST to those measured in the NIRC2. Individual frames were checked to ensure that they showed no irregularities in the point-spread function (PSF), and that the Strehl ratio was at least ~ 0.10 . Individual frames were then combined using the Drizzle algorithm (Fruchter & Hook, 2002) to provide sub-pixel accuracy. Further analysis was then done on the combined frames to provide the highest signal to noise.

The observations were conducted with two different pixel scales, using the NIRC2 “wide” camera, with $0.04''$ pixels and a $40'' \times 40''$ field of view, and the NIRC2 “narrow” camera, with $0.01''$ pixels and a $10'' \times 10''$ field of view. The total effective exposure times ranged from 420 to ~ 2600 seconds. The shorter exposures result in optimal use of 1-hour Target-of-Opportunity interrupts, when acquisition, calibration, and sky measurements are included.

The spatial resolution of Keck’s NIRC2 with Laser Guidestar Adaptive Optics proved to be very effective in resolving the detailed structure and position of the knot. A typical narrow-camera NIRC2 image is shown in Figure 2.1. In the narrow camera, the center of the pulsar and the center of the knot are separated by ~ 60 pixels. Typical values for the signal-to-noise ratio of the knot are 5. The PSF of NIRC2 with Laser Guidestar Adaptive Optics (AO) varies slightly in time due to the variable performance of the AO system. As the pulsar and knot are well resolved from each other in the Keck images, we chose not to deconvolve the PSF from the Keck images.

2.2.2 HST

We obtained 17 HST/ACS Wide-Field Channel (WFC) observations since 2012 January. Each observation comprises 4 exposures totaling 2000 seconds and uses a standard 4-point BOX dithering pattern to fill the inter-chip gap as well as to allow efficient cosmic-ray cleaning. We employed the F550M filter, which is well suited to sampling continuum emission from the Crab with

almost no contamination from line emission.

Table 2.2 summarizes the observations, listing the date, and the mean observing time. The appendix discusses our analyses of the 17 HST/ACS images. Figure 2.2 displays the processed images, which have removed effects of the HST/ACS WFC point spread function using singular-value decomposition (SVD) and Richardson-Lucy deconvolution.

2.2.3 Chandra

Our approach for searching for X-ray emission from the inner knot with Chandra involved a special use of both the high time resolution and high spatial resolution of the High-Resolution Camera Spectroscopy Array (HRC-S) in timing mode. So doing combines the best spatial resolution and the best time resolution afforded by the Chandra instrumentation. To achieve the time resolution it is vital that the counting rate be below the rate where telemetry is saturated (see, e.g., Tennant et al., 2001) and so the Low-Energy Transmission Grating (LETG) was inserted into the optical path to reduce the flux. In addition, the trigger threshold of the HRC-S was raised in order to avoid saturation as inserting the grating is insufficient to avoid telemetry saturation. This approach was first successfully used in a trial experiment (ObsID 11245) on 2010 November 16.

The next set of observations in this mode was initiated by us on 2013 March 5 as part of a sequence triggered by a gamma-ray flare (Ojha et al., 2013). Unfortunately, the gain of the HRC had dropped since 2010. Consequently, the trigger threshold was too high during our first observation on March 5-6 and the counting rate from the pulsar was about a quarter of that expected, dramatically decreasing the number of detected counts. The next HRC observation was to take place on March 10 and the observation had already been fully planned with the commands to execute on board the spacecraft. To replan at this stage would require an extremely short turnaround, yet it was accomplished as described in the Acknowledgments. Finally, beginning on 2013 October 22, we triggered another sequence of HRC observations in response to the announcement of another γ -ray flare (Buson et al., 2013).

We also consider ObsID 9765 taken on 2008 January 22. For this observation the HRC was operated in a more standard configuration with the threshold set at the nominal value. However, to reduce the count rate, a blade was inserted into the optical path (something that is no longer allowed). This blade also reduces the diffracted flux in one direction. By chance, this reduction is greatest near the nominal location of the knot. This reduction, together with a long integration time for this observation, results in ObsID 09765 providing the best data for study of the knot. Table 2.3 summarizes the relevant Chandra observations.

2.3 Properties of the Knot

2.3.1 Keck

Table 2.4 lists the best-fit values of the model parameters for the inner knot for each of the 12 Keck observations based on the analysis and modeling described in the Appendix. The parameters are: θ_0 , the direction from the pulsar to the center of the knot measured positive, E of N; r_0 , the separation between the pulsar and the knot; FWHM_r , the full width at half maximum of the knot measured along the line of sight from the pulsar to the knot; FWHM_t , the FWHM of the knot in the tangential direction; and the ratio of the knot flux to the pulsar flux, F_k/F_p . The variable nebular background in the infrared, coupled with the narrow field of view of the adaptive optics system, limits our ability to measure the total flux of the Pulsar. Due to the difficulty of determining absolute fluxes (and thus peak surface brightness) only the ratio of the knot flux could be measured. In Figure 2.3 we plot these properties as a function of the knot separation.

We seek possible correlations of the these properties with the separation r_0 from the pulsar, of the following form: $p(r_0) = p(r_1)[r_0/r_1]^q$. We choose a reference distance $r_1 = 0.6692''$, obtained from the logarithmic average of the 12 values of r_0 . The position angle of the knot, θ_0 , was fit to a linear regression $\theta_0 = b + m[r_0 - r_1]$. Table 2.5 tabulates the results of a power-law regression analysis of the inner knot's intrinsic properties with its (projected) separation

from the pulsar.

Based upon the F-test probability, the power-law regression analysis finds marginally statistically significant correlations of the parameters with separation. The strongest relationship is a 3% correlation of the flux ratio with separation, in somewhat contradiction to the previously known and opposite effect in the optical ((Moran et al., 2013), and see also §2.3.1.1 and §2.4).

Figure 2.3 shows that F_k/F_p as a function of the pulsar-knot separation is most extreme for measurement K07, and this single point has an influence on the $F_k/F_p - r_0$ correlation discussed above. Table 2.6 tabulates the impact of removing this point from the correlation analysis. Now the correlations with the $fwhm$ have improved slightly, whereas the the surface brightness and flux ratio correlations have become even less significant.

2.3.1.1 HST

Table 2.7 lists the best-fit values of the model parameters for the inner knot for each of the 17 HST observations based on the modeling and analysis described in the Appendix (§2.A). These properties are the position angle θ_0 of the knot with respect to the pulsar, the (projected) distance r_0 from the pulsar, radial width $FWHM_r$ and tangential width $FWHM_t$ of the knot, the peak surface brightness S_k , the knot flux F_k , the pulsar flux, F_p , and flux ratio F_k/F_p .

In the Appendix we discuss two ways of analyzing these data. Table 2.7 represents the results of one of these methods for the reasons discussed in the Appendix. However, as discussed in the Appendix, none of our general conclusions concerning the knot’s optical properties depends strongly upon the method of analysis. In Figure 2.4 we plot properties as a function of the knot separation.

As with the Keck data, we calculate the results of both a power-law and a linear regression analysis of the inner knot’s intrinsic properties with its (projected) separation from the pulsar (Table 2.8). We choose a reference distance $r_1 = 0.6559''$, obtained from the logarithmic average of the 17 measured values of r_0 . The position angle of the knot, θ_0 , was fit to a linear regression $\theta_0 = b + m[r_0 - r_1]$.

Based upon the F-test probability, the power-law regression analysis finds statistically significant correlations of the tangential and radial widths with separation from the pulsar, as well as the previously known (Moran et al., 2013) anti-correlations of flux and peak surface brightness with separation. The correlations between radial width FWHM_r and θ_0 with separation are the least strong but both are still statistically significant at better than 99% confidence. Note however, there is a small systematic bias in the power-law index of the FWHM_r correlation, in that the radial width of the knot at times approaches the resolution limit of the instrument and analysis procedures. Applying our analysis procedures (see Appendix) to simulated knots of known true radial width, we estimate that the smallest measured values of FWHM_r are oversized by about 7% and the largest by about 3%, which would increase the power-law index for the $\text{FWHM}_r - r_0$ correlation from $q = 0.72$ to about 0.8. Finally, we emphasize that we find similarly strong correlations using the more traditional HST analysis methods described in the Appendix.

2.3.2 Chandra

In contradistinction with the optical properties, the leading question with the X-ray emission is whether or not emission from the knot can be detected, especially in the vicinity of the bright, X-ray-emitting pulsar. The answer, unfortunately, is no and here we set an upper limit to the X-ray flux of the inner knot relative to that of the pulsar.

Table 2.9 shows the number of counts in each of three extraction regions: one centered on the approximate location of the optical knot, one centered on the pulsar, and one used to estimate the background. Details may be found in the Appendix and are illustrated in Figure 2.5. The analysis was performed both for the full phase-averaged data and for the data at pulse minimum when the pulsar is the faintest. Consider, for example, the phase-averaged data for ObsID 11245. Comparing the inner knot counts to the background (Table 2.9) we find an excess of 167 counts. If the excess was really due to the knot, then we would also expect an excess of 50 ($= 167 \times 0.3$) counts at phase minimum. Since this was not the case, we attribute the excess to the wings of the PSF from

the pulsar and posit that it is due to the spring/fall asymmetry discussed in the Appendix. Since the impact of that asymmetry is below the statistical noise for the data taken at phase minimum, we then use the phase minimum data to derive upper limits. We convert our minimum phase upper limit into a phase average upper limit by accounting for the 0.3 phase-duration of the minimum and then divide by the phase average pulsar flux. Results are listed in the last column in Table 2.9.

2.4 Comparison amongst the knot observations

Our Chandra observations were unable to detect X-ray flux from the knot, setting an upper limit to the ratio of knot flux to pulsar flux that is 3 – 4 times lower in the X-ray band than the flux ratios measured in the infrared and in the optical. Both our infrared and optical observations detected time variability in the pulsar–knot separation, with roughly consistent separations when measured contemporaneously (Figure 2.9). Correlations of knot properties with pulsar–knot separation are very strong when measured in the optical but much weaker in the infrared. We attribute this difference to a combination of factors—including poorer statistics for the much shorter IR measurements and a number of systematic effects, resulting in part from the non-stationary PSF and narrow field of view of an AO system. As Figure 2.9 and Table 2.4 suggest, IR-measured separations at essentially the same epoch exhibit a scatter substantially larger than the statistical errors in the measurement.

2.5 Comparison amongst observations of the inner knot

Our Chandra observations were unable to detect X-ray flux from the knot, setting an upper limit to the ratio of knot flux to pulsar flux that is 3 – 4 times lower in the X-ray band than the flux ratios measured in the infrared and in the optical. Both our infrared and optical observations detected time variability in the pulsar–knot separation, with roughly consistent separations when measured contemporaneously (Figure 2.9). Correlations of knot properties with pulsar–knot separation are very strong when measured in the optical

but weaker in the infrared. We attribute this difference to a combination of factors—including poorer statistics for the much shorter IR measurements and a number of systematic effects, resulting in part from the varying PSF, across the field and in time. As Figure 2.9 and Table 2.4 suggest, IR-measured separations at essentially the same epoch exhibit a scatter substantially larger than the statistical errors in the measurement.

2.6 Comparison with γ -ray fluxes

It is clear from the previous sections that the properties of the optical knot are well correlated with the projected separation from the pulsar and also, albeit with weaker evidence, the IR knot. Thus, in searching for any relationship between the knot and the γ -ray flux we concentrate on the separation versus γ -ray flux. The analysis performed to reduce the Fermi/LAT data is the same as described in Mayer et al. (2013).

Table 2.10 lists the dates and separations measured in the Keck and HST observations, followed by the corresponding 12-hour-average Fermi/LAT ≥ 100 -MeV fluxes. Figure 2.6 shows the Fermi/LAT flux as a function of time covering the interval when our Keck and HST observations were made. Red (Keck) and Green (HST) vertical lines are placed at the times when observations with these two observatories were performed. Figure 2.7 expands Figure 2.6 at the times of the largest γ -ray flares.

Figure 2.8 plots the γ -ray flux versus the measured radial separation of the knot from the pulsar. The results captured in this figure clearly do not represent a significant correlation of the knot properties with the γ -ray flux with a formal correlation coefficient of only 0.2.

Inspection of the Fermi/LAT γ -ray light curve (Figure 2.6) and the time series of pulsar-knot separations (Figure 2.9) indicates that the most energetic γ -ray flares occurred at a time which is not precisely coincident with the time of greatest pulsar-knot separation. It is reasonable to ask how probable a coincidence would be as a chance occurrence. A simple answer is that the probability of a chance coincidence with the observed time scales is just the difference in the central times of the two events divided by the total duration of this observing

program.

We next estimate the time and duration of a γ -ray flare and the excursion in r_0 (Figure 2.9), using a simple Gaussian model in both cases and including a linear baseline trend for the γ -ray flare. This is not to suggest that a Gaussian properly describes the underlying physical processes, but simply to enable comparison of the two time series, one for the flares, one for the separations. Fitting separation data in the MJD range 56180-56400 and γ -ray fluxes in the range 56345-56370 places the midpoint of the inner knot excursion at 56331.8, followed by the γ -ray flare at 56357.3, 25.5 days later. The Gaussian width is a measure of the event duration with FWHMs of 70.7 and 7.6 days for the separation excursion and γ -ray flare data respectively.

The time between the our first and last pulsar-knot separation measurements is 826.1 days. However, on 4 occasions the gap between separation measurements was longer than the 70.7 day duration of the observed separation excursion above. We assume that a separation excursion would be observable if an observation fell within the FWHM window and not observable if it fell outside the window. To obtain the total observing program under this assumption, we remove the time of 4 gaps (140.1 days) that exceeded 70.7 days making the total duration of the observing program 686 days. Thus an estimate of the chance probability is $25.5/686 = 3.7\%$, which is approximately at the 2σ level.

It is worth noting that the largest pulsar-knot separation incursion (toward the pulsar) occurs near the second most energetic γ -ray flare. However, the knot separation data are sparser at this time. Therefore, we use the previous FWHM separation value, 70.7 days to establish the time of the incursion. With this prior, and fitting separation data in the MJD range 56390-56670 and γ -ray fluxes in the range 56575-56590 places the midpoint of the inner knot excursion at 56566.0, followed by the γ -ray flare at 56582.6, 16.6 days later.

In this second case, the comparable probability estimate that the excursions occurred by chance is $16.6/686 = 2.4\%$ ($\sim 2.5\sigma$). Since each variation of the knot-pulsar separation is of the opposite sign we must multiply the probabilities by a factor of 2 when combining them. So our estimate of the chance probability that the 2 largest knot excursions/incursions would occur within the observed intervals of the 2 largest γ -ray flares is $0.037 \times 0.024 \times 4 = 0.4\%$,

which is approximately the 2.9σ level. There are additional γ -ray flares; however, the 3rd most energetic flare occurs only 12 days later than the one on MJD 56582.6 and therefore is well within the assumed 70 day range of the associated knot incursion. The other flares are significantly weaker and less well-sampled by the optical/infrared observations.

In conclusion, and although our coverage is sparse, there appears to be no compelling evidence in the pulsar-knot separations for large ($> \sim .01''$) excursions from the mean separation at time-scales comparable to the flare durations. In fact knot incursion/excursion time scales appear to be a factor of $> 10\times$ flare durations, so direct correlation of the γ -ray light curve with the separation time-evolution produces correlation coefficients which are rather difficult to interpret without associated simulations. Such simulations would require assumptions regarding an underlying physical model relating the 2 curves. Since we have no such model, any correlation analysis would be purely ad hoc.

We do not place much significance on the 0.4% value for the chance probability, which is based on a rather simplistic analysis involving only 2 flares. However, in continuing to monitor the Crab at multiple wavelengths with the goal of understanding the origin of its γ -ray flares, the inner knot will continue to be a feature of interest. Based upon the r_0 time scale of 71 days (a 30 day Gaussian 1σ), a monitoring program which includes monthly optical observations is necessary to detect significant inner knot excursions from its mean pulsar separation. Monitoring every ~ 3 weeks should be sufficient to fully characterize these variations and verify the small statistical sample from which these parameter estimates were derived and is something we plan to do.

We note that inner knot incursion with a ~ 2 -month FWHM timescale and $0.075'' \pm 0.025''$ pulsar-knot separation inward amplitude was also reported by Moran et al. (2013) on page 8. This event appears to be quite similar in duration and amplitude to our observations. The reference also reports an $\sim 2\sigma$ apparent change in the knot polarization parameters, which occurred 22 ± 5 days after the maximum observed incursion time. The degree of linear polarization steps from $59.9 \pm 6.9\%$ on 2005 Nov 25 to $42.8 \pm 6.2\%$ 2005 Dec 5 and the position angle changes from $125.4^\circ \pm 3.3^\circ$ to $121.6^\circ \pm 4.2^\circ$. Amusingly, the time frame of this polarization shift occurred in a window closely matching

both the 25.5 and 16.6 day gamma-ray flare delays after peak inner knot excursion/incursion times which we quote above. While the claimed 2σ is not very significant, we observe that prior to the step near the start of December 2005, all 10 of the reported measurements of the degree of linear polarization fall into a population with mean 60.74% and small standard deviation 1.0%, whereas the 2 subsequent measurements had a mean of 43.4% with standard deviation of 0.8%. If we suppose that the quoted errors in the reference apply to the absolute, but not relative corollary measured values, then the size of the step change would be $> 10\sigma$ instead of $\sim 2\sigma$. This may be a reasonable assumption, since the point-to-point measurement variations are smaller by a factor of > 6 than the quoted errors. Similar arguments would apply to the polarization position angle. From this we merely conclude that polarization should also be of interest in future attempts to relate inner knot behavior to gamma-ray flaring.

2.7 Implications for Theoretical Models

In the basic model of the Crab Nebula (e.g., Rees & Gunn, 1974), the pulsar wind passes through a shock at a radius $\sim 3 \times 10^{17}$ cm where the wind momentum flux balances the nebular pressure. However, the wind should be anisotropic. For example, the momentum flux in a split-monopole $\propto \sin^2 \theta$ (Bogovalov, 1999) and recent simulations suggest $\propto \sin^4 \theta$ (Tchekhovskoy et al., 2013) where θ is the polar angle. As a result, the shock is likely to be quite oblate. The sections of the shock near the poles would be oblique and much closer to the pulsar than the equatorial part. The observed radiation presumably comes from the relativistic particles accelerated behind the shock. Since the outflow from an oblique shock can remain relativistic, we would be able to see a compact emitting feature (the inner knot) if the relativistic outflow happens to be aligned with our line of sight so that its emission is beamed. The inner knot should appear to us as having some offset from the pulsar due to the deflection of the outflow from the radial direction (Figure 2.10). If the shock is approximately axisymmetric as one would expect, the projected emitting site would fall on the symmetry axis, leading to an alignment with the jet. The shape of the inner knot should also be more or less symmetric about the axis ; whether it's

elongated parallel or perpendicular to the axis depends on the geometry of the oblique shock. Furthermore the synchrotron emission should be linearly polarized (Komissarov & Lyutikov, 2011, Yuan & Blandford 2015 in preparation).

The shock model predicts that the scaling between the observed properties of the knot-projected knot-pulsar separation r_0 , tangential angular width FWHM_t , radial angular width FWHM_r and surface brightness S_k —should be decided by the upstream magnetization $\sigma \equiv B_1^2/\mu_0 n_1 \gamma_1^2 m c^2$ (where m is electron mass and B_1 , n_1 , γ_1 are the upstream magnetic field, the proper density, the bulk Lorentz factor, respectively) plus three more parameters that characterize the shape of the shock near the emitting site: the unprojected knot-pulsar separation r_k , the incident angle δ_1 , defined to be the angle between upstream velocity and the shock surface as shown in Figure 2.10, and the shock meridional radius of curvature R_c .

First, noticing that the shock outflow is deflected from radial direction by an angle Δ , where, Δ is the angle between upstream velocity and downstream velocity, we have $r_0 = r_k \Delta$, and the shock tangential size is determined by Doppler beaming: $\text{FWHM}_t \sim r_k/\gamma$, where γ is the outflow Lorentz factor. In the simplest shock model with an isotropic plasma, both Δ and γ are simply functions of σ and δ_1 : $\Delta = \delta_1 - \arctan(\chi \tan \delta_1)$ and $\gamma = 1/(\sin \delta_1 \sqrt{1 - \chi^2})$, where $\chi \equiv v_{2\perp}/v_{1\perp} = B_1/B_2 = (1 + 2\sigma + \sqrt{16\sigma^2 + 16\sigma + 1})/6(1 + \sigma)$ is the compression ratio at the shock, $v_{1,2\perp}$ is the component of upstream/downstream velocity which is perpendicular to the shock (Komissarov & Lyutikov, 2011). The radial size FWHM_r should also be proportional to r_k/γ but will have an additional factor involving R_c .

As to the intensity, we assume that the emitting particles have a power law distribution $\propto n' \gamma'^{-p}$ in the fluid rest frame, where n' , γ' are the particle density and the Lorentz factor in that frame, and $p \sim 2.2 - 2.6$ for the nebula in the IR to optical; it is important (but challenging) for what follows to determine this observationally.

The emissivity in the fluid rest frame is $j'_{\nu'}(\nu') \propto n' B'^{(p+1)/2} \nu'^{-(p-1)/2}$, where B' and ν' are the magnetic field and the radiation frequency measured in this frame. After transforming to the nebula frame, the emissivity has a dependence on the flow Lorentz factor γ , particle (improper) number density

n and magnetic field B as follows: $j_\nu \propto \mathcal{D}^{2+(p-1)/2}(n/\gamma)(B/\gamma)^{(p+1)/2}$, where $\mathcal{D} = \nu/\nu' = \sqrt{1 - \beta^2}/(1 - \vec{\beta} \cdot \vec{n})$ is the Doppler factor, in which $\vec{\beta}$ is the flow velocity and \vec{n} is a unit vector along the direction of line of sight. The surface brightness corresponds to the emissivity integrated along the line of sight in the nebular frame: $S_\nu = \int j_\nu dl$. To estimate the peak surface brightness, one takes the maximum Doppler factor $\mathcal{D} \sim 2\gamma$, and we get $S_{\nu,peak} \propto nB^{(p+1)/2}\ell$ where ℓ is an estimation of the length of the emitting region along the line of sight. Since, roughly, $B \propto r_k^{-1}$ and $n \propto r_k^{-2}$, we have $S_{\nu,peak} \propto r_k^{-2-(p+1)/2}\ell$.

These relations allow us to set constraints using both the steady and variable properties of the knot.

2.7.1 Variability

In the shock model, variability can arise either upstream or downstream of the shock. Firstly, stress tensor variation in the nebula can cause the shock radius and shape to vary over time. Three dimensional MHD simulations (Porth et al., 2014) show that even if the upstream condition is fixed, the post shock flow can be quite variable: there is vortex-shedding from the termination shock, and the shock constantly interacts with waves and vortices in the nebula. Those variations are expected to be mostly magnetic in origin, and the interactions between the shock and the downstream flow are quite nonlinear. The picture is similar to earlier 2D simulations (e.g.(Camus et al., 2009)) except that kink instability sets in so that the hoop compression is less prominent as in the 2D case; also short term variability is less pronounced. The 3D simulations give a typical variation time scale of less than one year.

Now suppose that a change of downstream pressure causes the shock radius r_k to change. From the scaling relations above we can see that if Δ , ℓ , γ , p and stays more or less the same, then $S_{\nu,peak} \propto r_0^{-2-(p+1)/2}$, $\text{FWHM}_t \propto r_0$, and knot flux $F_\nu \propto S_{\nu,peak} \text{FWHM}_t \text{FWHM}_r \propto r_0^{-(p+1)/2}$, assuming FWHM_r to be also roughly proportional to r_k . In reality these properties depend on more parameters besides the shock radius, for example, upstream incident angle δ_1 , so the correlation won't be clean, but the general trend that we observe is consistent with the expectation of the shock model.

Secondly, variations initiated by the pulsar can also cause the knot properties to change. In such a scenario, the shortest possible variation time scale is $t_v \sim r_k(1 - \cos(1/\gamma))/c \sim r_k/(2\gamma^2c) \sim 1 \text{ day}(r_k/10^{17} \text{ cm})(\gamma/5)^{-2}$ up to 10% inaccuracy assuming $\gamma \gtrsim 3$. However, there is no evidence in the pulsar timing for this actually happening.

2.7.2 Time averaged properties of the knot

The properties of the knot may be used to set constraints on the flow composition and how much dissipation happens at the shock.

One of the constraints comes from geometrical relations among the three measured angles r_0 , FWHM_t and FWHM_r . Since $r_0 = r_k\Delta$ and $\text{FWHM}_t \sim r_k/\gamma$, where Δ and γ only depend on σ and δ_1 in the simplest shock model, we find that in order to get a knot-pulsar separation $r_0 \sim 0.65''$ with a shock radius $r_k \leq 10''$, σ has to be less than 4. Meanwhile $\text{FWHM}_t/r_0 = 1/(\gamma\Delta) \geq 1.41$, with the minimum happening at $\sigma = 0$ and $\delta_1 = 0.39$. This seems to be somehow inconsistent with the observed fact that $\text{FWHM}_t \sim 0.5r_0$. However, some other factors may be important in deciding the size and deflection of the knot, for example, a downstream distribution that is more concentrated near 90° pitch angle will give a smaller tangential size. More detailed investigations will be described elsewhere. Furthermore, the observed optical aspect ratio, $\text{FWHM}_t/\text{FWHM}_r \sim 2$, indicates that the radius of curvature of the shock on the meridional cross-section should be sufficiently small: $R_c \lesssim r_k(1 - \chi)/[(2 - \chi) \sin \delta_1]$.

The shape of the knot also has some interesting implications. In both the Hubble and Keck images, the knot is not just a simple ellipse that only involves quadratic terms of coordinates; it has curvature and appears to be convex away from the pulsar (“smile”), so at least a third moment of the coordinates is needed to describe the shape. A simple intuitive picture is based on the geometry that the shock surface is similar to a donut centered on the pulsar and we look through the hole from the bottom; if the emission from the shock outflow roughly follows the local toroidal magnetic field, one would imagine the knot to be convex toward the pulsar (“frown”)—contradictory to what we see.

In fact, we should not overlook the possibility that the outflow may have significant bending downstream. If the outflow bends toward the equator by an angle $\sim 1/\gamma$ within the emission length, we would see the knot more extended (contracted) on the far (near) side with respect to the pulsar, since additional flow lines enter (originally aligned flow lines leave) the emission cone on the corresponding side due to the bending.

Another constraint comes from the observed flux of the knot. In the K' band, the typical knot-pulsar flux ratio is ~ 0.1 (2.4). If we take the photometry of the pulsar from Sandberg & Sollerman (2009), the pulsar K_s band magnitude is 13.80 ± 0.01 , corresponding to a flux of $1.87 \times 10^{-26} \text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$. So the inner knot flux in K' band should be roughly $F_\nu^{\text{knot}} = 2 \times 10^{-27} \text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$. If we define $\epsilon \equiv \nu L_{\nu, \Omega} / (\dot{E} / 4\pi)$, where $\dot{E} = 5 \times 10^{38} \text{erg s}^{-1}$ is the pulsar spin down power, $L_{\nu, \Omega} = D^2 F_\nu^{\text{knot}}$ is the knot luminosity per steradian and D the distance between earth and the Crab, ϵ can be used as a measure of the radiative efficiency of the shock along the direction of line of sight and in the particular waveband with frequency ν . From the data in K' band, we have $\epsilon = 3 \times 10^{-7}$. Similarly, optical data gives a flux of $F_\nu^{\text{knot}} = 2.25 \times 10^{-28} \text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$ (Table 2.8), indicating a radiative efficiency of $\epsilon = 1.2 \times 10^{-7}$.

We can get an estimation of the enthalpy fraction of IR/optical emitting particles in the downstream flow

$$\eta(\gamma') = \frac{8}{3} \frac{\gamma'^{t'_{\text{cool}}}}{t_{\text{flow}}} \frac{\nu L_{\nu, \Omega}}{\dot{E} / (4\pi)} = \frac{8}{3} \frac{\gamma'^{t'_{\text{cool}}}}{t_{\text{flow}}} \epsilon \quad (2.1)$$

where t_{flow} is the flow time scale and

$$t'_{\text{cool}} = 6\pi\epsilon_0 m^3 c^3 / (e^4 B'^2 \gamma') \sim 10^2 \gamma'^2 B_{-3}^{-3/2} \nu_{14}^{-1/2} \text{ years} \quad (2.2)$$

is the synchrotron cooling time of the particles in the fluid rest frame (we have adopted $B = B_{-3} \text{mG}$ and $\nu = \nu_{14} 10^{14} \text{Hz}$ for the numerical value). Thus the particle injection rate per unit steradian in corresponding energy band is

$$\dot{N}(\gamma') = \frac{3\eta(\gamma')}{4\gamma' m c^2} \frac{\dot{E}}{4\pi\gamma} \quad (2.3)$$

Adopting typical values for the magnetic field strength and the flow time scale, $B = B_{-3}$ mG, $\gamma = 5\gamma_5$, $t_{\text{flow}} = r_{17}10^{17}$ cm/ c , we find that for IR emitting particles, $\eta_{\text{IR}} = 0.08\gamma_5^3 B_{-3}^{-3/2} r_{17}^{-1}$, and $\dot{N}_{\text{IR}} = 3 \times 10^{36} \gamma_5^2 B_{-3}^{-1} r_{17}^{-1} \text{ s}^{-1} \text{ sr}^{-1}$; for optical emitting particles, $\eta_{\text{O}} = 0.02\gamma_5^3 B_{-3}^{-3/2} r_{17}^{-1}$, and $\dot{N}_{\text{O}} = 3 \times 10^{35} \gamma_5^2 B_{-3}^{-1} r_{17}^{-1} \text{ s}^{-1} \text{ sr}^{-1}$. We do not yet have a very good measurement of spectral index of the knot emission, but if the spectral index in the IR band is close to that of the nebula ($\sim 0.6 - 0.8$), the IR emitting particles should comprise the majority of particle pressure downstream. From the above estimations, it seems that the particle injection rate at the shock is marginally consistent with the scenario that most of the IR/optical emitting particles are provided by the shock.

Regarding particle acceleration mechanisms, we notice that IR emitting particles usually go through $N = t_{\text{flow}}\nu_g \sim 6 \times 10^4 B_{-1}^{3/2} r_{17} \nu_{14}^{-1/2}$ orbits within the flow time scale, where ν_g is the gyro frequency, so there's enough time for acceleration mechanisms that need sufficient development of instabilities to operate. However, for particles emitting γ rays of energy 300 MeV, N is $\lesssim 1$, thus requiring special acceleration mechanisms, if they are indeed produced at the shock.

One more constraint comes from the polarization of the knot. Most recent polarimetry performed by Moran et al. (2013) gives a high polarization degree $\sim 60\%$, with position angle aligned with the symmetry axis, indicating strongly a toroidal magnetic field. Here we need to be careful about the depolarization effect due to relativistic kinematics (Lyutikov et al., 2003). What happens then is that when we receive the emission of a relativistically moving plasma element by Lorentz transformation from the comoving frame to the lab frame, the polarization vector rotates. For a curved emitting surface, neighboring fluid elements have slightly different velocities, thus their polarization vectors experience different amounts of rotation. As one sums the contribution from the visible surface (this is essentially what we observe), the result is some degree of depolarization. We find that for the oblique shock in the Crab Nebula, under an ultra-relativistic approximation, the upper limit of polarization degree is similar to that in Lyutikov et al. (2003): 56.25% for particle spectral index $p = 3$ and 43.4% for $p = 2$. Thus, the observed high degree of polarization and its possible variation is highly constraining on theoretical models.

2.8 Conclusions

We have: (1) Introduced a new approach to analyzing time series of images; (2) Discovered that key properties of the knot (radial width, tangential width, flux) are correlated to the time-variable separation between the knot and the pulsar especially in the visible band; (3) shown that with the available data it is not possible to determine a strong correlation between the knot separation and the occurrence of γ -ray flares; (4) set an upper limit to the low-energy X-ray flux from the inner knot; (5) discussed the implications of our observations to set constraints on particular elements of the standard shock modeling of the relativistic outflow from the pulsar. These include inferences as to wind magnetization σ , shock shape parameters such as incident angle δ_1 and meridional radius of curvature R_c , as well as IR/optical emitting particle enthalpy fraction. We found that the standard shock model tends to favor $\sigma \lesssim$ a few; while it gives good agreements with the observation in many aspects, there remain two puzzles: (a) the small angular size of the knot as compared with the knot-pulsar separation seems hard to achieve; (b) the observed high degree of polarization is challenging for a highly relativistic outflow and the possible significant variation of polarization is still a mystery. We also found that the IR/optical flux of the inner knot is marginally consistent with the scenario that the shock provides most of the optical emitting particles in the nebula, while the acceleration of radio and γ -ray emitting particles is not so obvious.

2.A Appendix - Data Analysis

2.A.1 Keck

To study the knot using the data from the Keck AO system and the NIRC2 NIR camera, we characterize the region around the pulsar using an empirically measured point-spread function and an analytical model of the knot geometry.

2.A.1.1 PSF Subtraction

We subtracted the Point Spread Function of the pulsar using a nearby PSF star. During our observations, the laser guide star, and so the center of the anisoplanatic patch, was aimed at the point half-way between the PSF star and the pulsar. We fit the PSF star, a background plus a power-law based seeing-disk to the Pulsar. For the fit, we excluded the region in the pulsar subimage that surrounds the knot, and used the other 270° . The background was fit with a sloping plane model, to allow for the variable nebular background in the pulsar and PSF star subimages. The fitting was performed with a Levenburg-Marquadt fitter.

2.A.1.2 Deconvolution

We chose not to deconvolve the knot with the empirically measured point spread function. Tests on several images showed that deconvolution of the shape of the knot made insignificant differences to the fit values and uncertainty.

2.A.1.3 Knot Model

To fit the geometric parameters of the knot, we masked out the central diffraction-limited core of the pulsar, and used a Levenburg-Marquadt fitting technique to fit a two dimensional Gaussian to the shape of the Knot.

We used a model of the form:

$$Knot(r, \theta) = B + S_k \times \exp\left[\frac{(r - r_0)^2}{\sigma_r^2} + \frac{(\theta - \theta_0)^2}{\sigma_\theta^2}\right] \quad (2.4)$$

We also allowed the center of the coordinate system (x_0, y_0) to vary slightly ($0.05''$). To determine the fit uncertainties, we used the $1-\sigma$ errors from the self-covariance of the fit parameters. Results are were tabulated in Table 2.4

2.A.2 HST

2.A.2.1 HST - Traditional Analysis

In order to measure the knot properties (position, flux, tangential and radial width) we used the SExtractor package (Bertin & Arnouts, 1996). SExtractor has been extensively used for the analysis of HST data, in particular for the ACS camera (e.g., for the Hubble Ultra Deep Field project, Beckwith et al., 2006). For source detection we require a minimum of five contiguous pixels with a detection threshold 5σ above the root-mean-square (rms) background, with a total of 32 deblending subthresholds, and a contrast parameter of 0.005, setting the background mesh size to 16×16 ACS pixels.

The position of a source is evaluated by SExtractor as the barycenter of the source brightness distribution. Flux is computed within an elliptical aperture, using an implementation of the method by Kron (1980). The parameters of this ellipse (semiaxes and orientation), evaluated using the second moments of the object’s brightness distribution also yield a measure of the object morphology. Indeed, for the case of the knot, the direction of the minor axis of the ellipse turned out to be consistent with the pulsar-knot direction in all images. Thus, the minor axis and the major axis of the ellipse are a measure of the radial and tangential width, respectively.

In order to assess systematic errors, we performed simulations with the ESO / MIDAS software.¹ We added to the ACS images a “synthetic knot”. To generate such an artificial source we assumed a two-dimensional Gaussian brightness distribution, with the minor axis aligned with the pulsar-true knot direction. The synthetic knot was positioned to the NW of the pulsar, opposite but along the true pulsar-true direction, and at an angular distance comparable to the one of the true knot. We repeated the exercise by varying the flux, position and morphology of the synthetic knot and we estimated the uncertainties in the parameters recovered using SExtractor.²

Table 2.11 lists the best-fit values of the model parameters for the inner knot for each of the 17 observations based on the this traditional analysis and

¹<https://www.eso.org/sci/software/esomididas/>

²<http://www.astromatic.net/software/sextractor>

modeling. Table 2.12 shows the results of a regression analysis on the listed parameters.

2.A.2.2 HST Singular Value Decomposition (SVD)

As the inner knot is within 1 arcsec of the Crab pulsar and varies in position and size, we also developed special procedures for characterizing it using the central 121×121 -pixel ($6.05'' \times 6.05''$) images of the 17 HST/ACS observations. The major steps are (1) to remove the pulsar from each image, (2) to use the residual image to generate an image of the inner knot, and (3) to characterize the properties of the inner knot.

Remove pulsar and generate an image.

In order to remove the pulsar from each of the 17 central 121×121 -pixel ($6.05'' \times 6.05''$) HST images, we first determine the point response function (PSF) of the HST/ACS, using 121×121 -pixel images around each of 19 isolated stars in the field. In doing this, we account for the fact that some observations were at a roll angle $\sim 180^\circ$ opposite that of the others. After subtracting a fitted linear gradient from the image of each (Figure 2.12) of the 19 stars and for each of the 17 observations, we register the resulting $304 = 16 \times 19$ star images and use singular-value decomposition (SVD) to generate a linear basis (lowest term shown in Figure 2.12) describing the PSF. The model pulsar image Figure 2.12 uses the first 72 components of the SVD basis. After subtracting a linear gradient from each of the 17 central (pulsar) images (Figure 2.14), we fit each using the PSF basis to generate 17 residual (pulsar-removed) images (Figure 2.14), each now dominated by the inner knot.

Characterize inner-knot properties.

The measured extrinsic properties of the inner knot are the radial separation r_p and polar angle θ_p of the peak surface rightness with respect to the pulsar. To facilitate characterization of the intrinsic properties of the inner knot, we map each inner-knot sub-image (Figure 2.14) onto a $\theta - r$ grid (Figure 2.15). For an initial model of the surface brightness (intensity) distribution of the inner

knot, we use a simple bivariate normal distribution

$$S(r, \theta) = S_k e^{-1/2 \times [(r-r_p)/\sigma_r]^2 + [(\theta-\theta_p)/\sigma_d]^2}$$

The major intrinsic properties of the inner knot are its radial width σ_r , azimuthal width σ_θ (or tangential width $\sigma_t = r_p \sigma_d$), and peak intensity (surface brightness) S_k . Integrating the intensity over solid angle gives the flux $F_k = 2\pi\sigma_r\sigma_t S_k$. In fitting this model to the $\sigma - r$ image of the inner knot, we include a constant-surface-brightness background S_b as a model parameter.

While this simple model provides very adequate estimates to characterize the primary properties of the knot (Table 2.7), as a check we also introduced a somewhat more complicated model to deal with minor asymmetries in the $\sigma - r$ plane. In particular, we took the radial location and radial width to be weak functions of the azimuthal angle, which we expand as a Taylor series to second order: $r_p \sim r_p(\theta) \sim r_0 + r'_0(\theta - \theta_0) + (1/2)r''_0(\theta - \theta_0)^2$ and $\sigma_r \sim \sigma_r(\theta) \sim \sigma_{r_0} + \sigma'_{r_0}(\theta - \theta_0) + (1/2)\sigma''_{r_0}(\theta - \theta_0)^2$. However, the derivatives proved not to be statistically significant.

Comparison between methods

The comparison between methods and the reasons for somewhat emphasizing the SVD-based results considers the various measured parameters individually. Figure 2.16 compares measurements of similar variables between the two methods and illustrates the fact that the two methods, apart from biases, typically reach the same conclusions and lead to the demonstration (and new result) that the radial separation determines all of the other properties of the knot. From the figure we see that the biases for r_0 and FWHM_r are small whereas the knot flux and transverse size estimates differ by factors of two and three respectively. That they differ is not surprising as measures of these latter two properties are more sensitive to the method of background subtraction than are the measures of the separation and the radial width. We would also expect a difference in the transverse size estimate as the two analysis methods are fitting the data to a different shapes. Perhaps the most significant difference between the two methods are the smaller uncertainties in measured parameters afforded by the SVD approach. This too is to be expected as the approach removes those

SVD components that represent noise, thus enhancing signal to noise which is why we want to bring it to everyone's attention.

2.A.3 Chandra

Figure 2.5 shows the summed image at pulse minimum from the six “flaring” observations (ObsIDs 14684-16247 in Table 2.3. Prior to the binning used to produce the figure, we extract the counts in the 3 regions shown. It is difficult, if not impossible, to make use of the absolute number of counts as the different instrument configurations are all uncalibrated. However, we can make use of the relative numbers of counts from the different regions. Thus, the central region serves as an indicator of the number of pulsar counts which can then be compared with the other observations to derive physical fluxes (assuming of course that the phase averaged pulsar flux has not varied). The region to the southeast provides the upper limit for the flux in the knot and we use the region to the northwest to estimate the background, comprised mainly of that due to the wings of the PSF.

The alert reader will notice a slight excess of counts roughly $0.6''$ to the southwest of the pulsar. This feature is seen both in the phase-averaged pulse and in the data from ObsID 11245. (The asymmetric effect of the blade inserted for Obsid 9765 does not allow these data to be used in making this comparison.) Therefore, we conclude that this feature is part of the PSF. A more subtle effect is also present as there is a slight asymmetry in the PSF that causes the region we associate with the inner knot to be somewhat brighter than the background during observations taken in the fall. In the spring, when the spacecraft has (naturally) been set at a roll angle that differs by $\sim 180^\circ$ from that in the fall, the background region is brighter than that we associate with the location of the inner knot. This effect is mainly seen in the phase averaged data, and by comparing data from the ObsID 9765 (spring) and ObsID 11245 (fall) observations. Both effects described in this paragraph are relatively minor, are mainly due to the fact we are working slightly below the spatial resolution of Chandra, and have thus been ignored in setting the upper limits in Table 2.9

Table 2.1: Summary of the Keck observations

#	Date	MJD ^a	Instrument	Scale ("/>	
K1	2012-02-08	55965.285	NIRC2-wide	0.04	420
K2	2012-03-05	55991.320	NIRC2-wide	0.01	420
K3	2012-12-22	56283.406	NIRC2-wide	0.04	1250
K4	2012-12-23	56284.410	NIRC2-narrow	0.01	870
K5	2012-12-24	56285.387	NIRC2-narrow	0.04	1350
K6	2012-12-25	56286.390	NIRC2-wide	0.01	1125
K7	2013-02-06	56329.310	NIRC2-narrow	0.04	2020
K8	2013-10-22	56587.450	NIRC2-narrow	0.01	640
K9	2014-01-09	56666.273	NIRC2-wide	0.04	1560
K10	2014-01-09	56666.375	NIRC2-narrow	0.01	2580
K11	2014-01-17	56674.330	NIRC2-wide	0.01	2270
K12	2014-01-17	56674.390	NIRC2-narrow	0.04	1260

^a At the midpoint of the observation

Table 2.2: Summary of the HST observations

#	Date	MJD ^a
H1	2012-01-08	55934.7642
H2	2012-02-10	55967.1355
H3	2012-03-12	55998.3969
H4	2012-04-22	56039.7955
H5	2012-08-16	56155.2996
H6	2012-09-10	56180.9468
H7	2013-01-10	56302.9440
H8	2013-02-24	56347.0202
H9	2013-03-06	56356.9842
H10	2013-04-01	56383.3085
H11	2013-04-14	56396.0818
H12	2013-08-13	56517.6763
H13	2013-10-20	56585.7784
H14	2013-10-29	56594.6843
H15	2013-12-01	56627.2337
H16	2014-01-20	56677.5138
H17	2014-04-13	56760.8712

^a At the beginning of the observation

Table 2.3: Summary of the Chandra observations

#	ObsID	Date	MJD ^a	Integration Time(s)
1	09765	2008-01-22	54487.6700	95082
2	11245	2010-11-16	55516.2943	22039
3	14684	2013-03-05	56356.9718	19485
4	14686	2013-03-10	56361.3384	20052
5	14687	2013-03-17	56368.3239	16718
6	16244	2013-10-19	56584.6581	20048
7	16246	2013-10-22	56587.5153	20013
8	16247	2013-10-24	56589.0067	19770

^a At the beginning of the observation

Table 2.4: Knot properties based on the analysis of the Keck data

#	θ_0 °	τ_0 "	FWHM _r "	FWHM _t "	F_{ratio}
K1	118.504 ± 0.157	0.67106 ± 0.00147	0.3931 ± 0.0040	0.504 ± 0.005	0.069 ± 0.035
K2	120.999 ± 0.131	0.65768 ± 0.00146	0.4121 ± 0.0035	0.449 ± 0.004	0.070 ± 0.056
K3	122.639 ± 0.112	0.68927 ± 0.00108	0.3215 ± 0.0028	0.412 ± 0.004	0.093 ± 0.051
K4	120.961 ± 0.181	0.71181 ± 0.00212	0.3401 ± 0.0051	0.426 ± 0.006	0.094 ± 0.015
K5	121.199 ± 0.071	0.69811 ± 0.00075	0.2983 ± 0.0019	0.391 ± 0.002	0.066 ± 0.011
K6	120.788 ± 0.116	0.71641 ± 0.00101	0.3375 ± 0.0029	0.444 ± 0.004	0.070 ± 0.059
K7	121.926 ± 0.117	0.77564 ± 0.00132	0.3355 ± 0.0034	0.419 ± 0.004	0.155 ± 0.023
K8	120.333 ± 0.062	0.57225 ± 0.00040	0.2216 ± 0.0011	0.330 ± 0.002	0.069 ± 0.008
K9	121.464 ± 0.077	0.65427 ± 0.00063	0.2958 ± 0.0018	0.424 ± 0.003	0.080 ± 0.022
K10	115.773 ± 0.033	0.62575 ± 0.00025	0.2689 ± 0.0007	0.385 ± 0.001	0.082 ± 0.012
K11	121.487 ± 0.098	0.63707 ± 0.00081	0.3119 ± 0.0022	0.397 ± 0.003	0.069 ± 0.050
K12	118.991 ± 0.092	0.64381 ± 0.00089	0.3248 ± 0.0023	0.394 ± 0.003	0.066 ± 0.003

Table 2.5: Results of regression analyses for the 12 Keck observations.

Power-law for $r_1 = 0.6692''$				
Property	Unit	$p(r_1)$	q	Probability
FWHM _r	arcsec	0.3180 ± 0.0132	1.19 ± 0.54	$5.4E - 02$
FWHM _t	arcsec	0.4126 ± 0.0108	0.73 ± 0.35	$6.1E - 02$
S_k	F_p/as^2	0.5316 ± 0.0580	0.07 ± 1.39	$9.6E - 01$
F_k	F_p	0.0794 ± 0.0047	1.99 ± 0.77	$2.7E - 02$
Linear for $< r_0 > = 0.6711''$				
Property		b	m	
θ	<i>degrees</i>	120.4 ± 0.5	15.2 ± 10.2	$1.7E - 01$

Table 2.6: Results of regression analyses for 11 Keck observations removing K07.

Power-law for $r_1 = 0.6603''$				
Property	Unit	$p(r_1)$	q	Probability
FWHM _r	arcsec	0.3164 ± 0.0135	1.65 ± 0.67	$3.6E - 02$
FWHM _t	arcsec	0.4120 ± 0.0108	1.08 ± 0.41	$2.8E - 02$
S_k	F_p/as^2	0.5034 ± 0.0441	-2.14 ± 1.35	$1.5E - 01$
F_k	F_p	0.0747 ± 0.0030	0.59 ± 0.64	$3.8E - 01$
Linear for $< r_0 > = 0.6711''$				
Property		b	m	
θ	<i>degrees</i>	120.3 ± 0.6	15.7 ± 13.9	$2.8E - 01$

Table 2.7: Knot properties based on the HST SVD analysis (See the Appendix for details)

#	θ_0 °	r_0 "	FWHM _r "	FWHM _t "	S_k $e/s/as^2$	F_k e/s	F_p $e/s/$	F_k/F_p
H1	120.5 ± 0.4	0.6819 ± 0.0023	0.153 ± 0.006	0.322 ± 0.011	1064 ± 16	59.8 ± 3.6	1425 ± 82	0.042 ± 0.004
H2	119.5 ± 0.4	0.6535 ± 0.0024	0.154 ± 0.006	0.307 ± 0.011	1008 ± 16	53.8 ± 3.4	1433 ± 80	0.038 ± 0.003
H3	120.4 ± 0.3	0.6708 ± 0.0022	0.175 ± 0.005	0.312 ± 0.009	902 ± 11	55.0 ± 2.7	1471 ± 64	0.037 ± 0.002
H4	121.7 ± 0.4	0.6356 ± 0.0024	0.146 ± 0.006	0.294 ± 0.011	1053 ± 17	52.3 ± 3.2	1276 ± 80	0.041 ± 0.004
H5	117.9 ± 0.4	0.6166 ± 0.0020	0.141 ± 0.005	0.304 ± 0.010	1256 ± 18	61.1 ± 3.4	1231 ± 84	0.050 ± 0.004
H6	117.6 ± 0.4	0.6577 ± 0.0024	0.146 ± 0.006	0.303 ± 0.011	1102 ± 18	56.7 ± 3.5	1184 ± 86	0.048 ± 0.005
H7	121.2 ± 0.6	0.7240 ± 0.0035	0.170 ± 0.008	0.354 ± 0.018	675 ± 14	44.7 ± 3.8	1422 ± 77	0.031 ± 0.003
H8	119.9 ± 0.5	0.7516 ± 0.0031	0.175 ± 0.008	0.353 ± 0.016	667 ± 12	44.5 ± 3.4	1421 ± 65	0.031 ± 0.003
H9	120.9 ± 0.5	0.7491 ± 0.0033	0.174 ± 0.008	0.336 ± 0.017	728 ± 15	45.8 ± 3.8	1359 ± 73	0.034 ± 0.003
H10	120.6 ± 0.3	0.6614 ± 0.0019	0.160 ± 0.005	0.314 ± 0.009	1098 ± 14	63.3 ± 3.0	1295 ± 72	0.049 ± 0.004
H11	119.4 ± 0.4	0.6557 ± 0.0021	0.164 ± 0.005	0.322 ± 0.010	1090 ± 14	64.9 ± 3.3	1509 ± 77	0.043 ± 0.003
H12	119.3 ± 0.4	0.6324 ± 0.0023	0.169 ± 0.006	0.330 ± 0.011	1028 ± 14	63.4 ± 3.6	1283 ± 83	0.049 ± 0.004
H13	117.8 ± 0.4	0.5687 ± 0.0019	0.143 ± 0.005	0.274 ± 0.008	1668 ± 22	75.6 ± 3.8	1222 ± 106	0.062 ± 0.006
H14	118.8 ± 0.4	0.5849 ± 0.0022	0.144 ± 0.005	0.289 ± 0.010	1583 ± 24	75.6 ± 4.4	1285 ± 118	0.059 ± 0.006
H15	118.6 ± 0.5	0.6624 ± 0.0026	0.148 ± 0.006	0.339 ± 0.014	1105 ± 20	61.8 ± 4.4	1353 ± 100	0.046 ± 0.005
H16	118.8 ± 0.4	0.6299 ± 0.0026	0.164 ± 0.006	0.296 ± 0.011	1084 ± 17	62.1 ± 3.7	1419 ± 98	0.044 ± 0.004
H17	120.9 ± 0.4	0.6448 ± 0.0024	0.161 ± 0.006	0.321 ± 0.011	1164 ± 18	69.6 ± 4.0	1177 ± 99	0.059 ± 0.006

Table 2.8: Results of regression analyses for the 17 HST observations based on the data in Table 2.7.

Power-law for $r_1 = 0.6559''$				
Property	Unit	$p(r_1)$	Power-law index q	Probability
FWHM_r	arcsec	0.1575 ± 0.0022	0.72 ± 0.19	$1.7E - 03$
FWHM_t	arcsec	0.3152 ± 0.0029	0.80 ± 0.13	$1.4E - 05$
S_k	$F_p/a.s^2$	1.11 ± 0.05	-3.68 ± 0.61	$2.2E - 05$
F_k	F_p	0.0629 ± 0.0024	-2.15 ± 0.51	$7.2E - 04$
Linear for $< r_0 > = 0.6577''$				
Property		b	m	
θ	degrees	119.64 ± 0.27	13.5 ± 5.52	$2.72E - 02$

Table 2.9: Results of the Chandra analysis

ObsID	Pulse Average			Pulse Minimum			3 σ Limit $\times 10^3$
	Pulsar	Knot	Background	Pulsar	Knot	Background	
09765	51844	2963	3704	845	57	72	2.2
11245	5630	1018	851	88	16	12	9.4
last 6	13549	2622	2503	203	39	44	6.7

Table 2.10: Fermi/LAT 12hr-average fluxes and pulsar-knot separations as measured with Keck and HST.

#	Date	t_{avg} MJD	r_0 "	Fermi/LAT flux $\times 10^6$ $phs^{-1}cm^{-2}$ ($E \geq 100$ MeV)
H01	2012-01-08	55934.797	0.6819 ± 0.0023	2.45 (+0.47, -0.43)
K01	2012-02-08	55965.280	0.6711 ± 0.0014	2.58 (+0.60, -0.54)
H02	2012-02-10	55967.171	0.6535 ± 0.0024	2.79 (+0.65, -0.58)
K02	2012-03-05	55991.311	0.6577 ± 0.0015	2.78 (+0.46, -0.43)
H03	2012-03-12	55998.414	0.6708 ± 0.0022	3.55 (+0.52, -0.49)
H04	2012-04-22	56039.850	0.6356 ± 0.0024	2.18 (+0.56, -0.50)
H05	2012-08-16	56155.336	0.6166 ± 0.0020	3.25 (+0.56, -0.52)
H06	2012-09-10	56180.960	0.6577 ± 0.0024	3.86 (+0.50, -0.47)
K03	2012-12-22	56283.381	0.6893 ± 0.0011	3.13 (+0.68, -0.63)
K04	2012-12-23	56284.403	0.7118 ± 0.0021	3.51 (+0.67, -0.62)
K05	2012-12-24	56285.379	0.6981 ± 0.0008	3.29 (+0.67, -0.62)
K06	2012-12-25	56286.370	0.7164 ± 0.0010	2.74 (+0.62, -0.56)
H07	2013-01-10	56302.992	0.7240 ± 0.0035	3.22 (+0.66, -0.66)
K07	2013-02-06	56329.248	0.7756 ± 0.0013	3.93 (+0.75, -0.70)
H08	2013-02-24	56347.055	0.7516 ± 0.0031	5.44 (+0.72, -0.68)
H09	2013-03-06	56357.023	0.7491 ± 0.0033	10.1 (+0.52, -0.51)
H10	2013-04-01	56383.320	0.6614 ± 0.0019	3.50 (+0.68, -0.63)
H11	2013-04-14	56396.113	0.6557 ± 0.0021	2.89 (+0.59, -0.54)
H12	2013-08-13	56517.710	0.6324 ± 0.0023	2.65 (+0.46, -0.43)
H13	2013-10-20	56585.793	0.5687 ± 0.0019	3.71 (+0.26, -0.25)
K08	2013-10-22	56587.443	0.5723 ± 0.0004	1.25 (+1.73, -1.73)
H14	2013-10-29	56594.700	0.5849 ± 0.0022	7.46 (+1.16, -1.07)
H15	2013-12-01	56627.273	0.6624 ± 0.0026	2.15 (+0.53, -0.47)
K09	2014-01-09	56666.316	0.6543 ± 0.0006	3.38 (+0.53, -0.50)
K10	2014-01-09	56666.438	0.6258 ± 0.0003	3.54 (+0.51, -0.48)
K11	2014-01-17	56674.400	0.6371 ± 0.0008	3.47 (+0.44, -0.42)
K12	2014-01-17	56674.356	0.6438 ± 0.0009	3.84 (+0.53, -0.50)
H16	2014-01-20	56677.530	0.6299 ± 0.0026	2.85 (+0.43, -0.41)
H17	2014-04-13	56760.900	0.6448 ± 0.0024	3.96 (+0.78, -0.71)

Table 2.11: Results of the Traditional HST Analysis

#	r_0 arcsec	$width_r/2$ arcsec	$width_t/2$ arcsec	F_k $e s^{-1}$
H1	0.660 ± 0.011	0.143 ± 0.009	0.276 ± 0.018	116.40 ± 2.90
H2	0.634 ± 0.012	0.141 ± 0.009	0.269 ± 0.018	110.80 ± 2.70
H3	0.648 ± 0.013	0.166 ± 0.011	0.297 ± 0.019	131.60 ± 3.20
H4	0.615 ± 0.012	0.127 ± 0.009	0.244 ± 0.018	90.30 ± 2.20
H5	0.588 ± 0.012	0.139 ± 0.009	0.280 ± 0.019	112.80 ± 2.80
H6	0.635 ± 0.012	0.144 ± 0.009	0.276 ± 0.018	111.50 ± 2.70
H7	0.706 ± 0.010	0.144 ± 0.011	0.294 ± 0.021	97.50 ± 2.40
H8	0.734 ± 0.017	0.168 ± 0.011	0.329 ± 0.023	109.20 ± 2.70
H9	0.732 ± 0.013	0.160 ± 0.011	0.325 ± 0.022	108.00 ± 2.70
H10	0.641 ± 0.008	0.149 ± 0.010	0.290 ± 0.019	117.50 ± 2.90
H11	0.636 ± 0.010	0.157 ± 0.010	0.281 ± 0.018	129.10 ± 3.20
H12	0.613 ± 0.012	0.129 ± 0.009	0.289 ± 0.019	121.20 ± 3.00
H13	0.546 ± 0.008	0.112 ± 0.007	0.245 ± 0.016	119.10 ± 2.90
H14	0.552 ± 0.009	0.117 ± 0.007	0.255 ± 0.017	124.80 ± 3.10
H15	0.631 ± 0.012	0.134 ± 0.009	0.294 ± 0.019	115.10 ± 2.80
H16	0.613 ± 0.011	0.133 ± 0.013	0.260 ± 0.025	116.80 ± 2.90
H17	0.623 ± 0.011	0.130 ± 0.011	0.275 ± 0.025	127.70 ± 3.10




Figure 2.1: The Pulsar (left) and a nearby star (right) from Keck data taken in 2005 November. (a) An H filter image of the pulsar and knot (in the lower-left) as seen by the NIRC2 narrow camera ($0.01''$ pixels and a $10'' \times 10''$ field of view). The structure of the knot is resolved and separated from the pulsar by a statistically significant valley. (b) A nearby ($5''$) comparison star, used to establish the point-response-function. Both images are shown with a logarithmic color stretch.

Table 2.12: Results of regression analyses for the 17 HST observations based on the traditional data analysis.

Power-law for $r_1 = 0.6342''$				
Property	Unit	$p(r_1)$	Power-law index q	Probability
$width_r$	arcsec	0.2801 ± 0.0046	1.14 ± 0.21	$6.0E - 05$
$width_t$	arcsec	0.5605 ± 0.0067	0.85 ± 0.15	$4.9E - 05$
S_k	$e/s/as^2$	465.6 ± 9.1	-2.39 ± 0.24	$6.9E - 08$
F_k	e/s	114.8 ± 2.6	-0.39 ± 0.29	$1.9E - 01$




Figure 2.2: HST-SVD-Processed images of the Crab pulsar and inner knot after removing effects of the HST/ACS WFC point spread function.

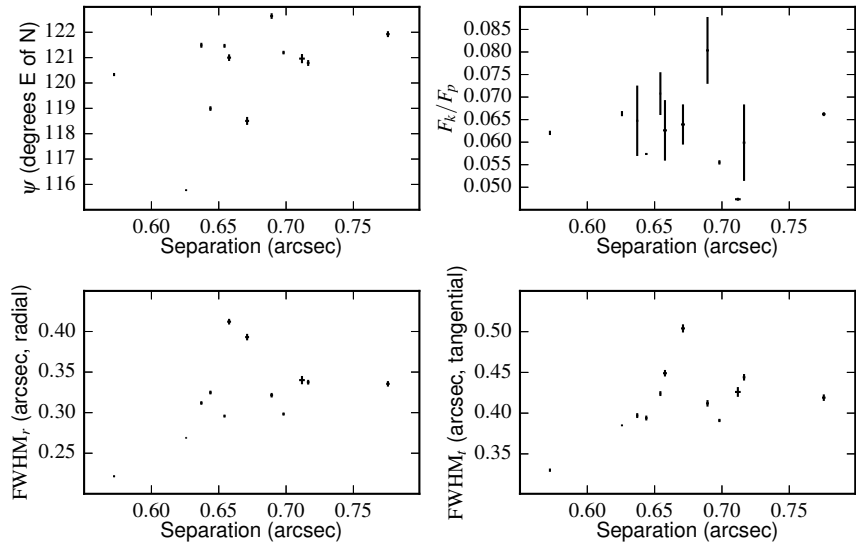


Figure 2.3: Variation of the Keck-measured knot properties versus the separation between the knot and the pulsar.

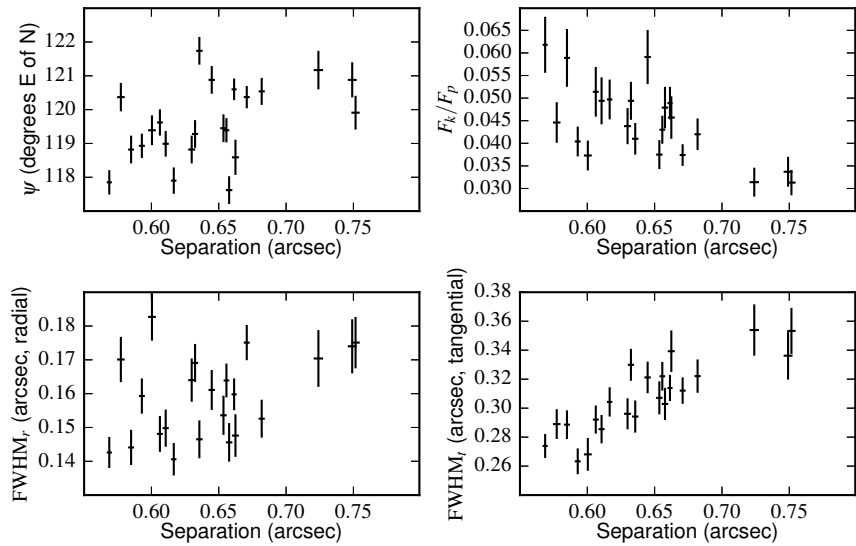


Figure 2.4: Variation of the HST-SVD-measured properties of the inner knot.




Figure 2.5: Chandra image at pulse minimum from sum of the last six “flaring” data sets in Table 2.3. The three small black circles show $0.3''$ -radius, extraction regions, the central of which is used to estimate the pulsar flux. The circle, at a position angle 120° east of north and $0.65''$ from the pulsar, is roughly centered on the average optical location of the knot. Data in the opposite small circle (the one to the NW where N is up) is used to estimate background. Large circles show the best fitting Gaussian to the phase average pulsar data at intensity levels of 10, 20, 40, 80, 160, 320 and 640 cts/pixel illustrating the level of impact of the PSF at the site of the knot i.e. slight but non-negligible. The intensity level is shown by the grayscale bar on the right.

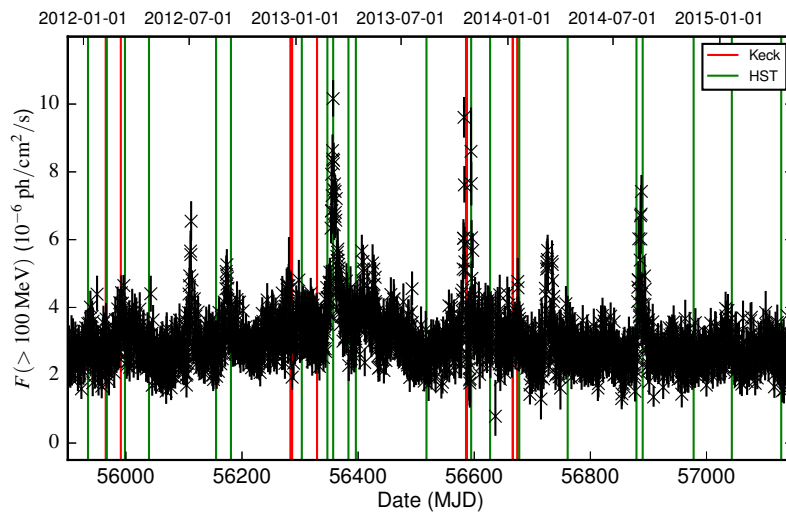


Figure 2.6: Fermi/LAT 12-hr average flux versus time with HST (green) and Keck (red) observations indicated.




Figure 2.7: Fermi/LAT 12-hr average fluxes during the Keck and HST observations at the time of the largest γ -ray flares in 2013 Mar (top) and Oct (bottom). Green indicates the time of a HST observation, red a Keck observation.




Figure 2.8: Fermi/LAT 12hr-average fluxes centered on the times of the Keck and HST observations versus pulsar-knot separations as measured with HST (green) and Keck (red).

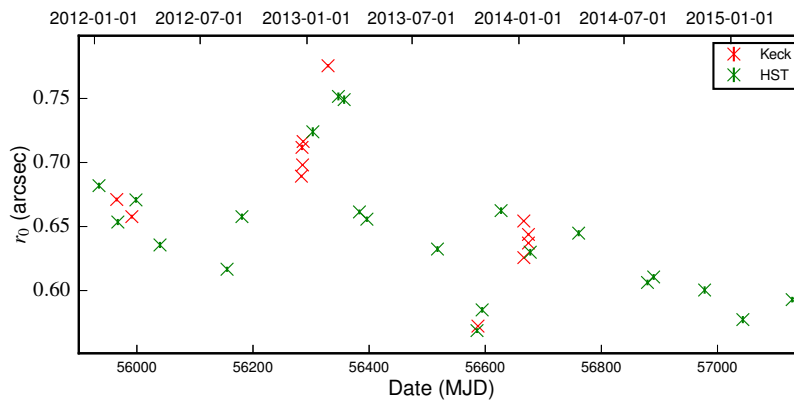


Figure 2.9: Pulsar-knot separation versus time with HST (green) and Keck (red) observations indicated.




Figure 2.10: Geometry of the termination shock on the meridional plane decided by the pulsar rotation axis and our line of sight toward the pulsar. The pulsar is located at point P. The angle between line of sight and pulsar rotation axis is $\theta_{\text{ob}} = 2\pi/3$. The knot is projected on the plane of the sky at an angle $r_0 \sim 0.65''$ southeast of the pulsar (represented by point K here). The angle between the upstream velocity and the shock surface is δ_1 . The outflow is deflected from radial direction by an angle Δ .




Figure 2.11: Geometric properties of the inner knot: r_0 , the projected separation of the knot from the pulsar; ψ_0 , the position angle (east of north) on the sky of the pulsar-knot separation; FWHM_r , the full width at half maximum of the knot in the radial direction; and FWHM_t , the FWHM of the knot in the tangential direction. The underlying image is from the 2013 October 20 HST observation (H13), which exhibits the minimum separation and maximum optical flux observed during this campaign.




Figure 2.12: Images illustrating the steps in determining the HST/ACS PSF for this analysis: (a) Extract 323121×121 -pixel ($6.05'' \times 6.05''$) images and subtract a linear-gradient background for 19 isolated stars in these 17 HST/ACS observations of the Crab; (b) register images of stars and use singular value decomposition (SVD) to determine a basis describing the PSF (first SVD component shown); and (c) model the pulsar image using the first 72 SVD components.




Figure 2.13: Chandra image at pulse minimum from sum of the last six “flaring” data sets in Table 2.3. The three small black circles show $0.3''$ -radius, extraction regions, the central of which is used to estimate the pulsar flux. The circle, at a position angle 120° east of north and $0.65''$ from the pulsar, is roughly centered on the average optical location of the knot. Data in the opposite small circle (the one to the NW where N is up) is used to estimate background. Large circles show the best fitting Gaussian to the phase average pulsar data at intensity levels of 10, 20, 40, 80, 160, 320 and 640 cts/pixel illustrating the level of impact of the PSF at the site of the knot i.e. slight but non-negligible. The intensity level is shown by the grayscale bar on the right.




Figure 2.14: Images further illustrating the steps in processing HST/ACS images of the Crab pulsar and inner knot: (a) Subtraction of a linear-gradient background in central 121×121 pixel ($6.05'' \times 6.05''$) image of pulsar and inner knot; (b) remove pulsar using SVD model of PSF leaving inner knot; (c) remove inner knot using its SVD model leaving residual background; (d) add SVD model of inner knot to residual background; (e) apply Richardson-Lucy algorithm to generate deconvolved image of inner knot (and residual background); and (f) synthesize 25×25 pixel ($1.25'' \times 1.25''$) sub-image of registered (unresolved) pulsar and SVD model of inner knot.




Figure 2.15: Mapping of processed images of the inner knot onto a $\theta(\text{radian}) - r(\text{pixel})$ grid. The pixel size is $0.05''$.

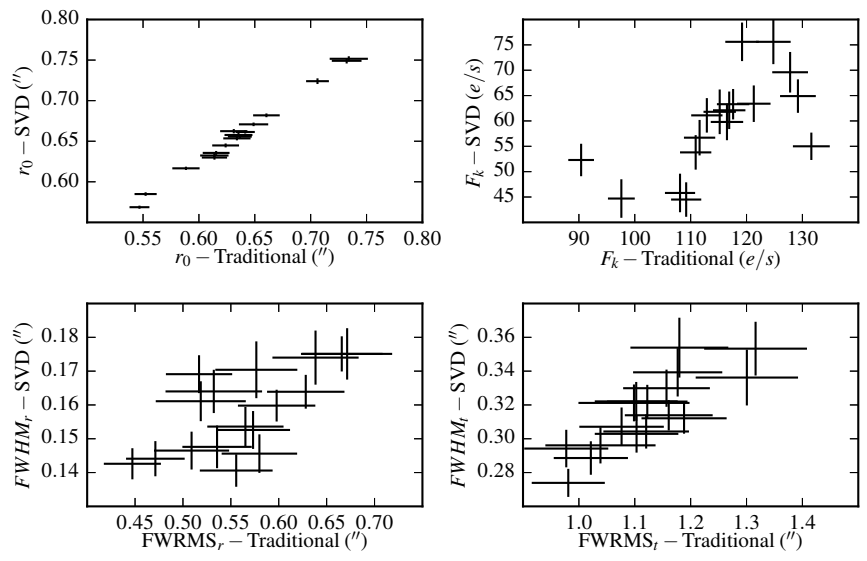


Figure 2.16: Variation of the HST derived properties of the knot, comparing SVD and traditional analyses.

Part II

Searching for Outflows in the kinematics of nearby ULIRGs

Chapter 3

Kinematics of the Nuclear Regions of ULIRG F20414-1651

3.1 Introduction

Quenching star formation in massive, gas-rich galaxies in the early universe is a key component of many cosmological models. Outflowing gas from massive galaxies provides a regulating feedback mechanism for the active galactic nuclei (AGN), limits the potential for star formation in these systems, and helps to give rise to the observed relationships between black hole mass and the properties of the host galaxy's bulge (e.g. Silk & Rees, 1998; Di Matteo et al., 2005a,b; Hopkins & Hernquist, 2009). Without these outflows, it is difficult to explain the old stellar populations and low gas content of local massive galaxies, or their steeply declining number at high masses.

Large-scale outflows have been identified in several nearby ULIRGs. Signatures of high velocity outflows have been observed in ionized (e.g. Rupke & Veilleux, 2011; Spoon et al., 2013; Spoon & Holt, 2010), neutral (e.g. Martin, 2005; Soto & Martin, 2012; Rupke & Veilleux, 2011; Soto et al., 2012), warm (e.g. U et al., 2013; Medling et al., 2015b) and cold molecular gas tracers (e.g. Feruglio et al., 2010, 2013; Sturm et al., 2011; Cicone et al., 2014). Sometimes, in the presence of a strong AGN, these outflows appear to be accelerated by that central AGN (U et al., 2013; Gupta et al., 2013; Gofford et al., 2013). These outflows

serve as a proxy for the mechanisms which drive the transition from gas rich to gas poor galaxies around $z \sim 2$, where it is difficult to study the physical mechanisms driving outflows in detail.

Several recent observational advances have detected numerous outflows in cold molecular gas. In addition to the detection of many outflows in CO (Feruglio et al., 2010, 2013; Sturm et al., 2011; Cicone et al., 2014), the Herschel observatory has provided a large sample of potential outflows detected in OH $119\mu\text{m}$ (Sturm et al., 2011; Spoon et al., 2013; Veilleux et al., 2013b; Contursi et al., 2013; González-Alfonso et al., 2013; Rangwala et al., 2011). These outflows are detected in spatially unresolved spectra of the OH $119\mu\text{m}$ doublet, where high velocity blue wings are apparent, as well as in some examples of P-Cygni profiles (Spoon et al., 2013; Veilleux et al., 2013b). The outflows detected in Spoon et al. (2013) and Veilleux et al. (2013b) are detected in a single Herschel spatial pixel (spaxel). Although the OH $119\mu\text{m}$ profiles are attributable to outflows, we need ancillary, spatially resolved data to confirm the presence of an outflow and to understand the mechanism driving them.

Evidence of some of the outflows observed by Herschel has been corroborated by optical emission lines observed via long-slit spectroscopy (Soto & Martin, 2012; Martin & Soto, 2016). Long-slit spectroscopy confirms that for some systems, there is gas at large distances from the center of the system, which is traveling at high velocities. However, even with long slit optical spectroscopy, we cannot constrain the dynamics and origins of individual outflows. In fact, without spatially resolved studies, it is possible for kinematically unusual arrangements to appear like an outflow, such as a superimposed tidal tail or merger remnant. This is unlikely in the highest velocity cases (~ 1500 km/s), but many of the observed outflows in Spoon et al. (2013) and Veilleux et al. (2013b) are only at speeds between 250 km/s and 600 km/s.

With spatially-resolved spectroscopy, either in the optical (Davis et al., 2012) or the NIR (Rupke & Veilleux, 2013), we can constrain the location, kinematics and excitation mechanisms for outflows from nearby ULIRGs (U et al., 2013; Medling et al., 2015b) where integral-field spectroscopy previously identified outflows. Seeing the distinct spatial and kinematic structure of an outflow (Rupke & Veilleux, 2013) confirms the signatures of outflows previously

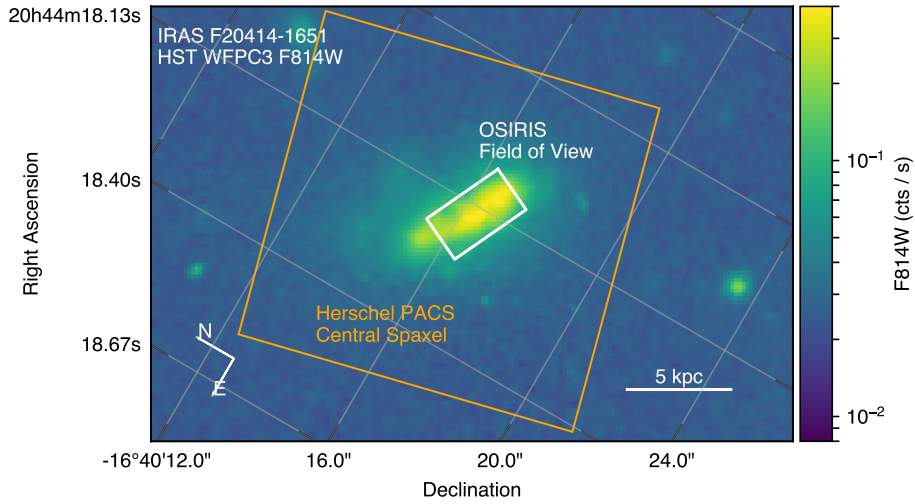


Figure 3.1: IRAS F20414-1651 image from HST WFPC3 F814W (Bushouse et al., 2002). The image is displayed on a logarithmic color scale, in counts per second for the WFPC3 camera’s F814W filter. The ULIRG shows evidence of disruption, and clear signs of being mid-merger. The OSIRIS field of view is plotted over the object as a white box. The Herschel PACS central spaxel, used for the measurements in Spoon et al. (2013), is shown in orange. The scale bar shows the physical scale for the redshift of IRAS F20414-1651 ($z = 0.0876$)

observed only spectroscopically. In the case of IRAS F08572+3915, there is an observed OH $119\mu\text{m}$ blue wing at $v_{84} = -832 \text{ km s}^{-1}$ (Veilleux et al., 2013a) and spatially resolved studies in the near-infrared (NIR) clearly show this outflow (Rupke & Veilleux, 2013). However, some galaxies which show signatures of outflows in unresolved studies don’t show strong evidence for outflows in their spatially resolved spectra. One such example, IRAS F23365+3604, has an observed OH $119\mu\text{m}$ blue wing at $v_{84} = -604 \text{ km s}^{-1}$ (Veilleux et al., 2013a), but doesn’t directly show signs of a high velocity outflow in spatially resolved NIR spectroscopy (Martin & Soto, 2016).

Here we present spatially resolved observations of the ULIRG IRAS F20414-1651 which shows evidence of a -250 km/s outflow in OH $119\mu\text{m}$ from Herschel PACS (Spoon et al., 2013). IRAS F20414-1651 has been imaged with HST before (see Figure 3.1 Bushouse et al., 2002), which shows it to be a late-stage merger with evidence of tidal disruption. It has also been observed with

the Keck ESI spectrograph (Soto et al., 2012), which shows evidence of higher-velocity gas at -300 km/s, but no signs of a distinct outflow.

This chapter presents NIR OSIRIS observations of IRAS F20414-1651. Section 3.2 describes the data acquisition, reduction, and analysis. Section 3.3 discusses the results of the data analysis. Section 3.4 interprets the results in the broader context of OH $119\mu\text{m}$ outflows observed with Herschel. Throughout the chapter, we adopt $H_0 = 70\text{km s}^{-1} \text{Mpc}^{-1}$, $\Omega_m = 0.28$ and $\Omega_\Lambda = 0.72$ (Hinshaw et al., 2009).

3.2 Methods

This chapter explores the inner regions of IRAS F20414-1651. Our NIR integral field spectroscopy was acquired with OSIRIS as part of our larger survey studying outflows from the nuclear regions of nearby ULIRGs. We targeted IRAS F20414-1651 based on its moderate outflow velocity detected from Herschel (Spoon et al., 2013). We also reference ancillary long-slit ESI data collected by Soto et al. (2012) and used wide-field HST WFPC3 imaging from Bushouse et al. (2002).

3.2.1 Observations

We observed IRAS F20414-1651 with OSIRIS, the near-infrared integral field unit spectrograph at the W.M. Keck Observatory (Larkin et al., 2006) using laser guide star adaptive optics (LGS AO). Our observations were conducted on 13 October 2014 and 14 October 2014 (summarized in Table 3.1). OSIRIS uses a lenslet array to subdivide the field of view into ~ 3000 individual spectra at $R \sim 3000$. Our observations were conducted with the $0.035''/\text{px}$ scale lenslet in the Kbb filter. We conducted observations at 3 pointings which overlapped, creating a mosaic across the center of IRAS F20414-1651. The data were collected in an object-sky-object pattern to capture changes in sky brightness. Each exposure was 600 s long.

OSIRIS produces better spatial-resolution than an equivalent seeing-limited instrument through the use of the Keck adaptive optics system (Wizinowich et al., 2000, 2006; van Dam et al., 2004, 2006). We used the sodium laser guide star to provide an artificial guide reference on axis with the spectrograph. Along with the sodium laser guide star, we used a nearby star to correct for overall image motion. Our tip/tilt star had an R-band magnitude of 13.3 and was $30.2''$ from the center of our pointings at IRAS F20414-1651.

At the start and end of our observations, we observed the UKIRT faint standard star FS-34 (EG141, $M_{K'} = 13.000$), which we used to flux-calibrate our data. We also observed telluric stars once every two hours, at the average air-mass of our objects to provide a reference for atmospheric absorption.

The full field of view of our OSIRIS observations is shown as a box on

Table 3.1: Observations of IRAS F20414-1651

Date	Filter	Scale	Exposure Time	Sky Time
2014/10/13	Kbb	0.035 "/px	6 × 600 s	3 × 600 s
2014/10/14	Kbb	0.035 "/px	4 × 600 s	2 × 600 s

a HST WFPC3 image (Bushouse et al., 2002) in Figure 3.1. The three pointings we used overlapped 4 lenslets along the edge of the field of view, to provide a continuous measurement across the area of interest.

3.2.2 Data Reduction

Our observations were reduced using the OSIRIS data reduction pipeline, v4.0b¹ and the wavelength solution from February 2013. The pipeline contains modules to subtract sky frames, adjust channel levels, remove inter-pixel crosstalk, identify glitches, clean cosmic rays, extract a deconvolved spectrum for each spatial pixel (spaxel), assemble a data cube of spaxels, correct for atmospheric dispersion, subtract scaled residual sky lines (from a technique described by Davies (2007)), and perform a telluric correction. Our telluric correction was done with the telluric star observed closest to each individual observation. Flux calibration was done based on the photometry measured from the UKIRT faint standard star FS-34 in the reduced OSIRIS data cubes to set a zero-point for the instrument.

Before each observation, we imaged the tip-tilt star, which we use as an estimate of the PSF during observations of each target galaxy. Medling et al. (2014) found that the tip-tilt star PSF did not vary significantly from the start to the end of each observation. However, the PSF will vary based on the distance between the tip-tilt guidestar and the target, and based on the isokinetic angle at Mauna Kea. We model the observed PSF as a convolution of the measured PSF, broadened by a gaussian consistent with a degradation in Strehl of $1/e$ for a separation of $75''$ (from van Dam et al., 2006) between the tip/tilt star and the center of the pointing on the galaxy. This correction results in a 15% degradation in the FWHM of the reference PSF.

¹The beta version was unreleased when we did this work, but it is now released on GitHub at <https://github.com/Keck-DataReductionPipelines/OSIRISDRP>.

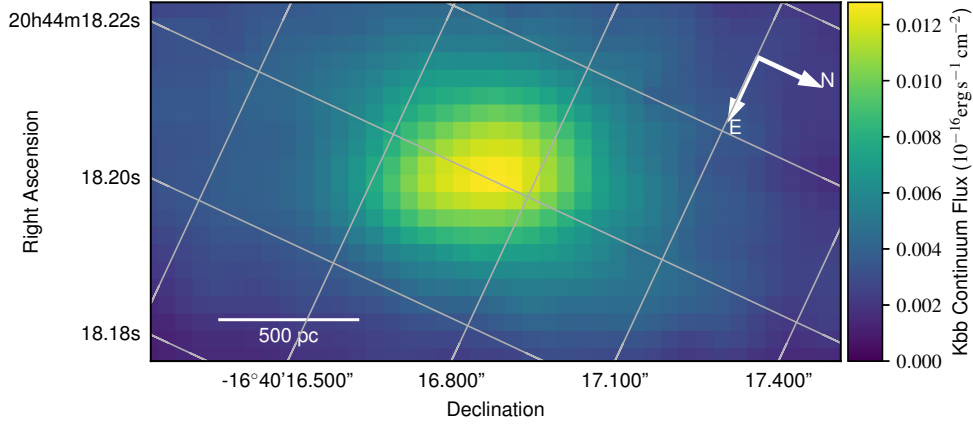


Figure 3.2: We fit a power-law model of the continuum to the full Kbb data cube. The power-law amplitude is shown on a linear stretch in this figure. The continuum shows relatively little structure, and looks significantly different from the F814W image of IRAS F20414-1651 from HST.

3.2.3 Line Fitting

To fit the datacube, we fit a simultaneous continuum and emission model to each spaxel. The continuum is fit as a simple power-law across the entire wavelength range. We fit emission in 2 HII lines and 3 H₂ lines. The emission model assumes that species of lines (e.g. Pa- α and Br- γ) are related in velocity and velocity dispersion, but we allow the emission amplitude to vary among the individual realizations of a species. The entire cube is fit simultaneously, and we require that spaxels produce fits which are related to their neighbors, by convolving the model with the PSF described in Section 3.2.2. This has the visual effect of producing smooth fits (as in Figure 3.2) even when there is significant noise in the data.

To reduce the complexity of fitting the entire data cube simultaneously, we only consider wavelength ranges within 3000 km/s of the rest-frame wavelength of each line of interest. We allow for the continuum to float between these disconnected windows, although this is a better measure of the quality of the telluric correction rather than a statement about the way the continuum deviates from a power-law.

The fitting is done with a modified Levanburg-Marquadt fitter (Mark-

wardt, 2009; Astropy Collaboration et al., 2013) which takes into account the error in each spaxel, and which iteratively applies fits to successively less-smooth spectra to overcome the sensitivity to initial conditions.

3.2.4 Modeling the Kinematics

For the emission components which showed significant signs of rotation, we fit a model of a thin gas disk with a central point mass, using the model in U et al. (2013). Our disk model is applied to the fitted velocity and velocity dispersion from 3.2.3. We use an inclined thin disk model, with a constrained floating central point. Our fit uses the emission line flux to fit the center of the disk, and the emission line velocity to fit for a uniform density disk.


3.2.5 Kinemetry

In addition to our thin-disk model, we fit a non-parametric disk model which does not depend on the physical properties observed, using Kinemetry (Krajnović et al., 2006). Kinemetry works by fitting a series of concentric ellipses to the observed velocity and flux distribution. Each ellipse is fit by minimizing the deviation from the even (velocity) or odd (flux) moments of the Fourier transform of the values along the ellipse. Figure 3.3 shows what this process looks like for a single ellipse at $a = 0.60$ kpc. Kinemetry provides an estimate of the disk eccentricity, as well as a velocity profile for the disk (obtained by measuring the peak velocity for each concentric ellipse).

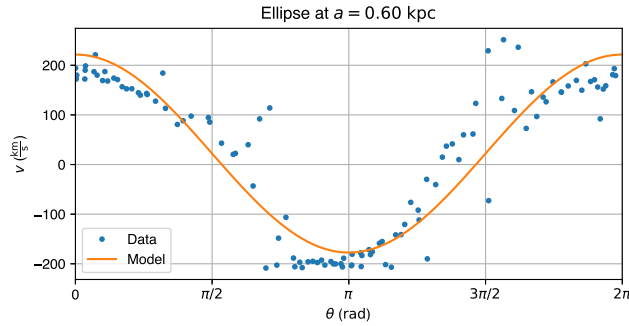
3.2.6 Modeling the OH $119\mu\text{m}$ emission

Along with our analysis of the OSIRIS observations, we chose to re-fit the data from OH $119\mu\text{m}$ described in Spoon et al. (2013) so that we could get an accurate measure of the uncertainty in the v_{98} quantity (the 98th percentile blue velocity) quoted by Spoon et al., and determine whether these values were consistent with the observed gas emission in the NIR.

To fit the OH $119\mu\text{m}$ absorption, we tried to match the technique described by Spoon et al. (2013). We fit the OH $119\mu\text{m}$ doublet as a pair of lines with identical absorption and emission, and allowed for the presence of two



(a) Velocity in Pa- α of spaxels included in the kinemetry fit.



(b) Best fit to a radially symmetric ellipse at $a = 0.6$ kpc.

Figure 3.3: Kinemetry fit for an ellipse with a semi-major axis $a = 0.6$ kpc. The best fit has a velocity amplitude of $k_1 = 199.5$ km/s, a PA of 85° and a flattening of $q = 0.677$. The top panel (a) shows the velocity spaxels selected for the best fit. The bottom panel (b) shows the best fit ellipse as velocity vs. angle around the ellipse.

absorbers/emitters with independent velocities and velocity dispersions. We added a polynomial-valued continuum. To understand the uncertainties involved in this fit, we sampled our model with a Markov-chain Monte Carlo method.

3.3 Results

We detect HII, H₂, and HeI emission lines across IRAS F20414-1651. The kinematic structure of Pa- α shows evidence of two components (Figures 3.4 and 3.5). The Pa- α and H₂ lines in IRAS F20414-1651 show a compact gas disk, with a strong rotation curve visible in Pa- α and H₂ emission. We marginally detect HeI emission consistent with the presence of ionized Helium in the gas disk.

Table 3.2: List of emission lines detected in our observations of IRAS F20414-1651.

Line	Species	Wavelength
Pa- α	HII	1.8764 μm
H ₂ 1 – 0 $S(3)$	H ₂	1.9575 μm
H ₂ 1 – 0 $S(2)$	H ₂	2.0338 μm
HeI	HeI	2.0586 μm
H ₂ 1 – 0 $S(1)$	H ₂	2.1218 μm
Br- γ	HII	2.1661 μm

We find that the Pa- α emission has three components across the field of view (see Figure 3.4 and Figure 3.6). There are two primary components detected in Pa- α (hereafter **A** (disk) and **B**). Additionally, we find a weakly detected third component in Pa- α emission to the north of the primary nucleus, which we label **C**, and which is characterized by high velocities ($v \geq 1000$ km/s). The total velocity distribution we detect in Pa- α and H₂ emission is broader than the velocity distribution detected in OH 119 μm by Spoon et al. (2013).

We identify emission from three families of lines, which are summarized in Table 3.2. Two HII lines are detected, Pa- α and Br- γ (§ 3.3.1). Three rotational H₂ transitions are detected in the integrated data cube, however we do not detect one (H₂ 1 – 0 $S(2)$) with enough signal to noise to fit its parameters independently. We also detect HeI emission, but without sufficient signal-to-noise to detect an independent source.

3.3.1 Hydrogen Recombination Lines

We detect a kinematically complex Pa- α emission, which we fit with three gaussian emission components. The velocity profile of Pa- α , spatially integrated across the field of view, is shown in Figure 3.4. This profile shows clear indicators of multiple distinct sources for Pa- α . The emission line fitting shows two primary peaks, and a broader secondary component underneath. Although the integrated emission shows two peaks in the same velocity profile, these peaks move spatially across the object. Figure 3.5 shows the velocity profile for the north and south halves of IRAS F20414-1651, showing that the two larger peaks separate, and could belong to the same emission source since they




Figure 3.4: Spatially-integrated velocity profile of Pa- α for IRAS F20414-1651. The black line shows the raw, spatially-integrated spectrum. The green line shows the sum of all of the fits to the individual spaxels across IRAS F20414-1651. The spatially-integrated velocity profile spans 800 km/s. The spatially-integrated Pa- α fit fails to capture the red wings of the emission around 400 km/s. This dip is a known instrumental feature of bright emission lines created from the current OSIRIS data reduction pipeline. The two components are shown as green dotted lines, emphasizing that the bulk of the integrated flux is attributable to the bright disk component, and only a small portion of the remaining spatially-integrated flux can be attributed to the **B** component.

appear on opposite sides of IRAS F20414-1651. The secondary emission component (labeled Pa- α **B** in Figures 3.4 and 3.5) is broader and fainter. It primarily appears in the south half of IRAS F20414-1651, as can be seen in Figure 3.5.

The spatial distribution of Pa- α emission is shown in Figure 3.6. The primary component (Figure 3.6b) is associated with the disk of IRAS F20414-1651, and a secondary extended component (**B**, Figure 3.6c) is shown to the south of the center of the disk. The contours of the primary Pa- α component, shown on top of the Kbb continuum in Figure 3.6a, shows that the disk is offset slightly north from the center of the Kbb continuum. This primary component shows a large degree of radial symmetry. The flux from the secondary component (**B**, see Figure 3.6c) is south of the primary Pa- α component, but consistent with the extended Kbb continuum in that direction. The second component contains 16% of the Pa- α flux in spaxels where it is detected.

The velocity of the first component shows Keplerian rotation out to the

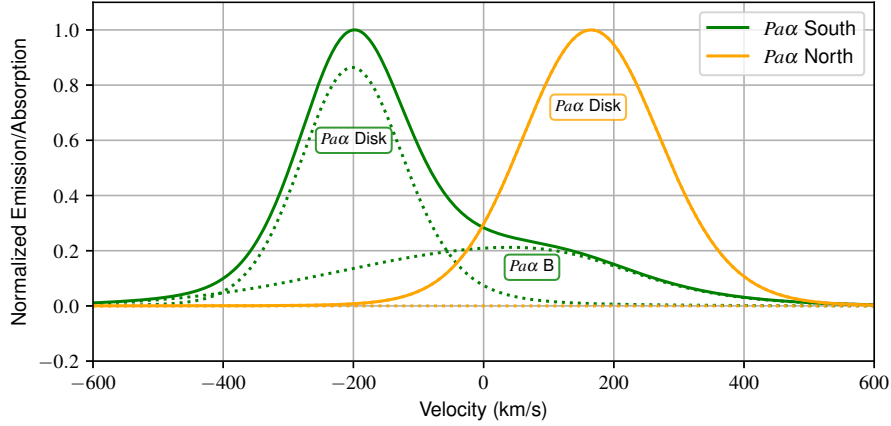




Figure 3.5: Integrated velocity profile of Pa- α for IRAS F20414-1651 split between the north and south regions. The green line shows the average velocity profile of Pa- α for a region 10×5 spaxels south of the line-of-nodes for the primary nucleus. The orange line show the average velocity profile of Pa- α for a region 10×5 spaxels north of the line-of-nodes for the primary nucleus. The secondary emission component slightly above $v = 0$ km/s is only visible in the south portion (component **B**, in green), while the north region is well fit by a single emission peak. We interpret this to mean that the lower, higher dispersion component (**B**) is only present in the southern half of the disk, and that the two peaks in the spatially-integrated spectrum

edge of the field of view (Figure 3.7a). A rotation curve is fit using the Kinemetry technique described in § 3.2.5. The second component shows a velocity 6.9 km/s extending 1 kpc to the south of the primary nucleus (Figure 3.7b). Although the second component is at an average of 6.9 km/s, it could still be rotating, just not moving along the line-of-sight relative to the center of the IRAS F20414-1651 nucleus.


The two Pa- α components show differing velocity dispersion signatures (see Figure 3.8). The disk component has a velocity dispersion consistent with ordered rotation, showing low velocity dispersion at high radii from the center of the disk, with a higher velocity dispersion along the line of nodes. This is consistent with the higher velocity dispersion arising from viewing the emitting gas through a longer line of sight in the center of the disk, if the disk is at a low inclination. The high values of v/σ (Figure 3.9) suggest highly ordered rotation in the disk. The second component has a mean velocity dispersion of



(a) Kbb Continuum with flux contours from Pa- α disk (A) component.




(b) Pa- α flux from the disk (A) component.




(c) Pa- α flux from the B component.

Figure 3.6: The top panel (a) shows the Kbb continuum from Figure 3.2 overlaid with contours showing the flux in Pa- α for the disk (A) emission component. The center panel (b) shows the flux from the Pa- α disk (A) emission component. The lower panel (c) shows the flux of the B Pa- α emission component. We fit two emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel. The Pa- α emission is concentrated around the center of the object, suggesting a relatively round central object.



(a) Pa- α velocity in the disk component **(A)**.




(b) Pa- α velocity in component **B**.

Figure 3.7: Pa- α velocity comparison between the disk component and the “B” component in IRAS F20414-1651. The color scale shows the emission velocity, centered around $z = 0.087084$, the nominal redshift from Strauss et al. (1992). Black solid contours represent the Pa- α flux in this component (see Figure 3.6c). Some spaxels with flux were removed from the line fit as they didn’t produce satisfactory fits. This component has a relatively constant velocity of ~ 0 km/s, which is red-ward of the disk rotation velocity at this location (-200 km/s)


140 km/s, which is higher than the high-radius velocity dispersion observed in the disk (101 km/s).

The velocity dispersion in the disk (Figure 3.8a) helps to explain the two-peaked structure of the spatially integrated disk shown in Figure 3.5. The high dispersion in the center of the disk, where the line-of-sight through the disk is largest, reduces the amplitude of Pa- α observed in these spaxels, but still keeps the flux (see Figure 3.6b) symmetric and peaked at the center of the disk.

We fit a radial profile to the Pa- α emission from the disk (Figure 3.10) and to the continuum, both in an azimuthally averaged bins. Although we al-



(a) Pa- α velocity dispersion in the disk component (A)



(b) Pa- α velocity dispersion in the B component

Figure 3.8: Pa- α velocity dispersion comparison between the disk component (a) and the B component (b) in IRAS F20414-1651. The color scale shows the emission velocity dispersion, σ , on a linear stretch, in km/s. Black solid contours represent the Pa- α flux in this component (see Figure 3.6c). Some spaxels with flux were removed from the line fit as they didn't produce satisfactory fits.




Figure 3.9: The ratio of velocity to velocity dispersion in the disk component of the Pa- α emission. Black solid contours represent the Pa- α flux in this component (see Figure 3.6c). The high values ≥ 3 at large radii from the center of the disk suggest that the disk is undergoing ordered, “cool” rotation.




Figure 3.10: Radial profile of the flux from Pa- α emission in the disk (**B**, orange) and continuum (blue). The points are azimuthally averaged flux values with error bars. The line shows a fit to a Sersic profile. For Pa- α we found a Sersic index of $n = 0.94$, consistent with an exponential profile. The effective radius is 505 pc. For the continuum, the best fit has a Sersic index of $n = 0.81$, which is shallower than would be expected for an exponential profile. The effective radius for the Kbb continuum 617 pc. The continuum image suggests that the continuum emission is not solely due to the disk, and in fact is flattened in the southern half due to emission from component **B**. The Sersic fit to both the continuum and the Pa- α disk flux is consistent with the disk having an exponential flux profile.

lowed the Sersic index (n) to vary, for Pa- α , we find the best fit to the flux is an exponential profile ($n = 1$). The continuum shows a flatter profile ($n = 0.81$), although this is expected from the extended, non-symmetric continuum emission shown in Figure 3.6a

We detect Br- γ in all spaxels where we detect Pa- α . We fit for the ratio between Br- γ and Pa- α emission in a given spaxel. For the disk component, we find that the amplitude of Br- γ emission is $0.089 \times A_{\text{Pa-}\alpha}$ in the disk, and $0.16 \times A_{\text{Pa-}\alpha}$ in component **B**. This corresponds to a total flux in Br- γ of $0.231 \times 10^{-16} \text{ erg s}^{-1} \text{ cm}^{-2}$ from the disk and $0.0394 \times 10^{-16} \text{ erg s}^{-1} \text{ cm}^{-2}$ from the second component (**B**).

3.3.2 H₂ Emission

We fit three H₂ transitions to IRAS F20414-1651, and observed emission consistent with H₂ co-located with the disk component observed in Pa- α . The flux observed in H₂ 1-0 S(3) is shown in Figure 3.11, and the corresponding velocity is shown in Figure 3.12. In all spaxels where we detect a significant




Figure 3.11: Flux from H_2 $1-0$ $S(3)$, shown in a logarithmic color scale identical to the one in Figure 3.6b. The distribution of flux is consistent with that of the primary $\text{Pa-}\alpha$ component in Figure 3.6b (presented here as the white contours) suggesting that the primary H_2 $1-0$ $S(3)$ emission comes from the observed gas disk.

amount of H_2 $1-0$ $S(3)$ we also detect a similar amplitude emission from H_2 $1-0$ $S(1)$. The ratio of H_2 $1-0$ $S(3)$ to H_2 $1-0$ $S(1)$ corresponds to an excitation temperature of 1200 K. We also detect H_2 $1-0$ $S(2)$ in many spaxels where we can successfully detect H_2 $1-0$ $S(3)$. The ratio of H_2 $1-0$ $S(2)$ to H_2 $1-0$ $S(3)$ corresponds to an excitation temperature of 1000 K, consistent with the excitation temperature suggested by H_2 $1-0$ $S(3)$ / H_2 $1-0$ $S(1)$.

We detect velocities in H_2 (Figure 3.12) which are consistent with the velocities observed in the disk in $\text{Pa-}\alpha$ (Figure 3.7a). As we don't measure reliable velocities at very low fluxes, the velocity map in H_2 is more limited than the same map in $\text{Pa-}\alpha$. Nevertheless, the velocities that we do detect are consistent with those found in $\text{Pa-}\alpha$. The flux and rotation curve in H_2 are slightly offset from the flux in $\text{Pa-}\alpha$ (see Figure 3.11). We don't consider this significant, as the zero-velocity line in both $\text{Pa-}\alpha$ and H_2 are spatially consistent.

The lack of two identified peaks in H_2 is consistent with the relatively uniform velocity dispersion in H_2 $1-0$ $S(3)$ (Figure 3.13). The velocity dispersion for H_2 $1-0$ $S(3)$ has a flux-weighted average value of 154 km/s.

Figure 3.14 shows the spatially-integrated velocity profile in H_2 and $\text{Pa-}\alpha$. The velocity profiles cover similar widths, but the H_2 does not show the double-peaked structure in the disk that is apparent in $\text{Pa-}\alpha$.




Figure 3.12: Velocity from $H_2 1 - 0 S(3)$ in km/s. The black lines show flux contours in $H_2 1 - 0 S(3)$. There is evidence of rotational velocities consistent with those observed in the disk in Pa- α (Figure 3.7a).




Figure 3.13: Velocity Dispersion from $H_2 1 - 0 S(3)$ in km/s. The black lines show flux contours in $H_2 1 - 0 S(3)$. The velocity dispersion doesn't show much structure, and is at an average value of 150 km/s.




Figure 3.14: Velocity profile for Pa- α and H_2 fits in IRAS F20414-1651. The green line shows a normalized, spatially-integrated velocity profile for Pa- α . The red line shows a normalized, spatially-integrated velocity profile for H_2 .




Figure 3.15: The velocity profile of OH $119\mu\text{m}$ absorption in IRAS F20414-1651 from Spoon et al. (2013). The blue dotted line shows the median parameter fit. The blue shaded lines show 100,000 samples of the posterior distributions of the fit parameters, converted to a velocity profile. There is not a significant variance the realized parameters in the posterior distribution, as the shaded region does not extend far from the median values. However, even at this level, it is evident that there can be significant variance in the v_{98} outflow velocity parameter.

3.3.3 OH $119\mu\text{m}$ Absorption Fit

As described in § 3.2.6, we fit the OH $119\mu\text{m}$ absorption observed in Spoon et al. (2013) with a two-component absorption profile. The MCMC fits yield a posterior distribution which is well concentrated around the median values of the fit (see Figure 3.15), suggesting that there is not significant uncertainty in the fit.

We used the MCMC sampling to determine the posterior distribution of v_{98} for the combined velocity profile. The posterior distribution of the v_{98} outflow velocity parameter (Figure 3.16) suggests that this measure of outflow velocity is uncertain at the ~ 40 km/s level.

3.3.4 Disk Parameters

We used kinemetry (Krajinović et al., 2006) to model the shape and rotation curve of the disk observed in Pa- α . Kinemetry provides disk parameters without any physical interpretation, and so is not dependent on any assumptions about the density of gas in the disk or the presence of a supermassive black hole. As well, kinemetry will always fit a well-ordered, symmetric disk, so the residuals are indicative of areas in the disk which are not undergoing well or-




Figure 3.16: The posterior distribution of the 98th percentile velocity of OH $119\mu\text{m}$ absorption in IRAS F20414-1651 from Spoon et al. (2013). The black line shows the median value, and the blue histogram shows values drawn from 100,000 samples of the posterior. The histogram has been normalized so that the posterior integrates to 1. The posterior distribution has an inter-quartile range of 40 km/s

dered rotation.


The velocity field fit with kinemetry is shown in Figure 3.17. The disk is well fit in the high velocity wings. There are higher residuals (due perhaps to the higher dispersion) along the minor axis of the disk (see Figure 3.17c).

The velocity curve derived from kinemetry is shown in Figure 3.18. It flattens after 250 pc, suggesting the presence of dark matter. We note that the flattening of the velocity curve occurs at a lower radius than the scale radius of the disk as measured by the continuum flux (606 pc, see Figure 3.10).


Kinemetry also allows us to derive the position angle and flattening of the disk. Figure 3.19 shows both of these parameters as a function of the disk radius.

3.3.5 Disk Mass


We fit a simple, uniform density inclined disk model to the Pa- α rotation curve, as described in § 3.2.4. We find that the disk is well fit by a disk with $2.55 \times 10^{10} M_{\odot}$ and an inclination of 9.5° . The fit velocity profile, and residuals from the fit are shown in Figure 3.20



(a) Pa- α velocity in the disk component (A).



(b) Kinemetry Model of Pa- α Velocity



(c) Kinemetry Model Residuals compared to Pa- α Velocity

Figure 3.17: Velocity map from the kinemetry model (Krajnović et al., 2006) for the primary Pa- α emission component. The top panel (a) shows the same velocity field described in Figure 3.7a. The middle panel (b) shows the velocity derived from the best-fit kinemetry model. The black lines show the best-fit ellipsoids for each disk radius in this model. The lower panel (c) shows the residual between the disk model and the observed velocity. The fit captures the velocity profile of the system well, but fails to account for the velocity and dispersion along the line-of-nodes of the disk, where it generally underestimates the rotation curve.






Figure 3.18: Rotation curve in Pa- α for a symmetric kinematics fit to the disk. The rotation curve shows a steep slope out to 300 pc, after which it is flat.




(a) Kinematics ellipse flattening q as a function of radius




(b) Kinematics ellipse position angle as a function of radius

Figure 3.19: Kinematics flattening (a) and position angle (b) as a function of radius for IRAS F20414-1651. The high q values at most radii suggests that the disk is inclined relative to the line of sight, consistent with the high inclination value found in the disk fit (see Section 3.3.5). The position angle of the component ellipses stays within a 10° range for most radii, although the large deviation at low radii might be indicative of a warp. The presence or absence of a warp in the disk can't be confirmed without higher resolution kinematics, but is seen in other similar ULIRGs (Martin & Soto, 2016).



(a) Pa- α velocity from disk fit (A)



(b) Residuals from disk fit

Figure 3.20: Thin disk model of the Pa- α emission from component **A** for IRAS F20414-1651. The top panel (a) shows the modeled velocity for the disk. The disk fit finds a total mass in the observed portion of the disk of $2.55 \times 10^{10} M_{\odot}$, and an inclination of 9.5° . The bottom panel (b) shows the velocity residuals, with higher residuals around the line of nodes.

3.3.6 Non-disklike Emission

In Pa- α , (§ 3.3.1) we observe two components, one which clearly belongs to a disk, and a second (**B**) which is not disk-like. To consider whether this second component could be the remnant of a merger, we examined the second component for evidence of rotation. Figure 3.21 shows just this second component. There is not strong evidence for a rotation gradient. However, the emission and velocity is extended away from the brightest point to the south, suggesting that the source (**B**) is extended in that direction and may have been tidally stripped, creating the observed tail.




Figure 3.21: Cutout of velocity in Pa- α from component **B**. Although there is a slight gradient across the brighter emission (along the black arrow), this gradient is not significant. The flux is extended to the south east (see Figure 3.6c)

3.4 Discussion

IRAS F20414-1651 provides an important example of the way complicated kinematics revealed by spatially resolved observations can inform the interpretation of integrated velocities discovered through other instruments. Both Spoon et al. (2013) and Veilleux et al. (2013a) found numerous examples of similar galaxies with blue-shifted velocity wings ranging from -100 km/s to -2000 km/s. From Herschel PACS observations of the OH 119 μm absorption doublet (Spoon et al., 2013), there was compelling evidence of a blue-shifted velocity component from IRAS F20414-1651 at ~ -500 km/s. This kind of blue absorption wing has been widely cited as evidence for high velocity outflows from ULIRGs (Sturm et al., 2011; Spoon et al., 2013; Veilleux et al., 2013b; Conrursi et al., 2013; González-Alfonso et al., 2013; Rangwala et al., 2011). Unfortunately, integrated kinematics are difficult to interpret, and don't provide concrete evidence of a particular physical process.

For IRAS F20414-1651, we find a kinematic structure consistent with the integrated kinematics observed in Spoon et al. (2013), but which shows that the high velocity OH 119 μm absorption cannot be directly interpreted as a galactic-scale outflow. Figure 3.22 clearly shows emission in Pa- α at -500 km/s in integrated light. However, the resolved kinematics and emission features tell a richer story.




Figure 3.22: The green lines show the integrated emission profile from Pa- α across the OSIRIS field of view for IRAS F20414-1651. The solid line is the total emission, and the dotted lines show the two emission components described in § 3.3.1. The blue lines show integrated velocity profile for OH 119 μ m from Spoon et al. (2013). The OH 119 μ m absorption was observed over a much wider field of view. The emission in Pa- α is consistent with the absorption seen in OH 119 μ m, although the double peaked nature of the Pa- α emission is not visible in the OH 119 μ m profile.

The integrated kinematics of Pa- α for IRAS F20414-1651 in Figure 3.22 clearly show two components, which are well separated and described in Section 3.3. Neither of these components is consistent with any substantial velocity of outflowing gas. The primary emission component (see Figure 3.6b) shows clear signs of ordered rotation, and is consistent with the central nuclear disk found in many nearby ULIRGs (Medling et al., 2014). The secondary emission component does not belong to the disk and must be alternatively explained.

3.4.1 Comparison of the Integrated Emission

Figure 3.22 shows both the spatially-integrated Pa- α emission and the the observed OH 119 μ m absorption as measured in Spoon et al. (2013). The spatially-integrated Pa- α emission covers a region $1.4'' \times 0.73''$, which is much smaller than the $9.4'' \times 9.4''$ spaxel used for the Herschel PACS observations in Spoon et al. (2013) (see Figure 3.1 for both footprints). We find that the Pa- α emission can easily explain the velocities observed in OH 119 μ m absorption,

despite the more limited spatial coverage. In fact, we observe Pa- α at higher velocities than OH 119 μm , with a v_{98} (the 98th-percentile velocity as defined in Spoon et al. 2013 and Veilleux et al. 2013a) in Pa- α (-388 km/s) almost a factor of two larger than the v_{98} measured from OH 119 μm absorption (-256 km/s). Our H₂ emission shows a similar velocity profile to Pa- α , and also easily explains the velocity of observed OH 119 μm .

Although we observe Pa- α emission red of the equivalent OH 119 μm absorption, we do not consider this to be significant, as many objects similar to IRAS F20414-1651 show P-cygni profiles in OH 119 μm emission (Spoon et al., 2013; Veilleux et al., 2013a). If a P-cygni profile were present in IRAS F20414-1651² it would explain the lack of red absorption in OH 119 μm but would not suppress the emission visible in both Pa- α and H₂.

At least one other object, IRAS F08572+3915:NW (Rupke & Veilleux, 2013) shows H₂ and Pa- α emission consistent with the observed OH 119 μm absorption. In the case of IRAS F08572+3915:NW, the H₂ emission shows a blue ~ 1000 km/s outflow, which is consistent with a blue absorption component observed in OH 119 μm . However, there is no evidence from IRAS F08572+3915:NW that the v_{98} blue-wing velocity corresponds to an individual outflowing component, rather a second, independent absorption component is observed at -1000 km/s in OH 119 μm which corresponds to a spatially-constrained H₂ outflow.

3.4.2 Possible Explanations of Kinematic Structure

The kinematic structure of Pa- α and H₂ emission is not consistent with outflowing gas as suggested in the Herschel observations (Spoon et al., 2013). Below we present two possible explanations of the secondary emission in IRAS F20414-1651:

- § 3.4.2.1 describes the case for an active star forming region, driving the secondary component of Pa- α emission, possibly a nuclear star forming region belonging to the host galaxy.

²We find that when we fit two components to the OH 119 μm profile in IRAS F20414-1651, we cannot rule out the presence of a P-cygni profile.

- § 3.4.2.2 explores the case for this emission coming from the remnants of the minor component of a merger in the history of IRAS F20414-1651.

Each explanation is sufficient to explain the observed kinematics and emission from IRAS F20414-1651. However, we are unable to constrain which explanation best fits the observed data. Additional resolved observations of the cool molecular gas (e.g. with ALMA) would greatly help to disentangle the kinematics in this system.

3.4.2.1 Super star cluster

The emission from the second component shows a H_2 to $\text{Br-}\gamma$ ratio of 1.03, which is consistent with emission coming from star formation ($\text{H}_2/\text{Br-}\gamma < 2$). The relative paucity of H_2 emission belonging to either the primary or secondary emission component means that IRAS F20414-1651 is more likely dominated by star forming regions. From the $\text{Pa-}\alpha$ and $\text{Br-}\gamma$ emission, we estimated a star formation rate of in § 3.3.1. These values are consistent with an actively star-forming, gas-rich region with a kinematic signature separate from the central galaxy.

Such super star clusters are observed in the inner regions of ULIRGs with turbulent merger histories (Medling et al., 2014; Davies et al., 2014, 2016). However, usually super star clusters are more compact, and only partially resolved at this redshift in ULIRGs. The secondary component in IRAS F20414-1651 would be one of the most extended super-star cluster, as it spans more than 500 pc. In Davies et al. (2016), all identified star clusters had a FWHM of less than 400 pc, which is smaller than the object observed in IRAS F20414-1651, though on the same order of magnitude.

3.4.2.2 Minor merger remnant

Evidence that this object is more than just a simple star forming region is primarily due to the slight velocity gradient observable across the second component (Figure 3.7b). The velocity gradient goes from -100 km/s at the southern edge of the object to -100 km/s at the northern edge of the object. Rotation is an indicator that this object is not a nuclear star forming region, but

rather has angular momentum which it acquired from some additional source. If this is the case, a likely candidate for that angular momentum would be an original host galaxy.

The wide field images of IRAS F20414-1651 (Figure 3.1) suggest that it has undergone a merger and is in a disturbed state. This is indicative of a recent merger history. The luminosity from the second emission component in Pa- α is consistent with a mass ratio of 1 : 10 relative to the primary nucleus.

3.4.3 Implications for fast Outflows

The high resolution observations presented here paint a complex picture of the structure of IRAS F20414-1651. We observed integrated emission profiles in Pa- α and H₂ consistent with both OH119 μ m from Spoon et al. (2013) and optical emission and absorption from Soto & Martin (2012). Although unresolved, this emission is consistent with the blue-shifted 500 km/s outflow observed by Spoon et al. (2013), we find that the resolved kinematics do not support such an explanation for the integrated emission and absorption. Rather, we proposed three alternative explanations which could readily explain kinematic signatures seen in many ULIRGs at \sim 500 km/s. IRAS F20414-1651 provides an important counter-example to the unresolved studies of outflows from massive galaxies, suggesting that there is more than meets the (unresolved) eye for many ULIRGs.

3.5 Conclusions

We present spatially-resolved NIR observations of the central few kpc of IRAS F20414-1651. We found significant Pa- α Br- γ and H₂ emission, indicative of active star formation. We find that IRAS F20414-1651 shows evidence of a well-ordered ($v/\sigma \sim 4$) disk, and an additional co-moving component. Both components appear to be actively star forming, based on their Br- γ luminosities. Although we cannot conclusively identify the nature of the co-moving component, we suggest that it is either a very large nuclear star forming re-

gion, or the stripped nucleus from a recent merger. Although this object shows $v = 250$ km/s blue absorption in OH $119\mu\text{m}$ from Herschel, we find that we can explain the entire velocity profile through the kinematics of the system, without invoking outflows. This suggests that for systems observed to have blue absorption wings in OH $119\mu\text{m}$ at moderate velocities, there is not necessarily a corresponding outflow.

Chapter 4

A Survey of Nearby ULIRGs with Outflow Signatures

4.1 Introduction

In this chapter, we present a survey of ULIRGs which were observed by the Herschel Space Telescope to have high velocity outflows in OH $119\mu\text{m}$ (Sturm et al., 2011; Spoon et al., 2013; Veilleux et al., 2013b; Contursi et al., 2013; González-Alfonso et al., 2013; Rangwala et al., 2011). These outflows are detected in un-spatially-resolved spectra of the OH $119\mu\text{m}$ doublet, where high velocity blue wings are apparent, as well as some examples of P-Cygni profiles (Spoon et al., 2013; Veilleux et al., 2013b). The outflows detected in Spoon et al. (2013) and Veilleux et al. (2013b) are detected in a single Herschel Spaxel. Although the OH $119\mu\text{m}$ profiles are attributable to outflows, we need ancillary, spatially resolved data to confirm the presence of an outflow and to understand the mechanism driving them.

There is additional evidence which helps to confirm the presence of the outflows observed in OH $119\mu\text{m}$. For some objects, the OH $119\mu\text{m}$ outflows are corroborated by optical emission lines observed via long-slit spectroscopy (Soto & Martin, 2012; Martin & Soto, 2016). Long-slit spectroscopy confirms that for some systems, there is gas at large distances from the center of the system, which is traveling at high velocities. However, even with long slit optical spectroscopy,

we cannot constrain the the dynamics and origins of individual outflows. In fact, without spatially resolved studies, it is possible for kinematically unusual arrangements to appear like an outflow, such as a superimposed tidal tail or merger remnant. This is unlikely in the highest velocity cases (~ 1500 km/s), but many of the observed outflows in Spoon et al. (2013) and Veilleux et al. (2013b) are only at speeds between 250 km/s and 600 km/s.

With spatially-resolved spectroscopy, either in the optical (Davis et al., 2012) or the NIR (Rupke & Veilleux, 2013), we can constrain the location, kinematics and excitation mechanisms for outflows from nearby ULIRGs (U et al., 2013; Medling et al., 2015b) where spectroscopy previously identified outflows. Observing the distinct spatial and kinematic structure of an outflow (Rupke & Veilleux, 2013) confirms the signatures of outflows observed only spectroscopically. In the case of IRAS F08572+3915, there is an observed OH $119\mu\text{m}$ blue wing, where the 84th% of the blue absorption (v_{84}) has a velocity of $v_{84} = -832$ km s⁻¹ (Veilleux et al., 2013a) and spatially resolved studies in the NIR clearly show this outflow (Rupke & Veilleux, 2013). However, some galaxies which show signatures of outflows in unresolved studies don't show strong evidence for outflows when observed with NIR spatially resolved spectroscopy. One such example, IRAS F23365+3604, has an observed OH $119\mu\text{m}$ blue wing at $v_{84} = -604$ km s⁻¹ (Veilleux et al., 2013a), but doesn't directly show signs of a high velocity outflow in spatially resolved NIR spectroscopy (Martin & Soto, 2016).

Several studies have looked at collections of ULIRGs with spatially resolved spectroscopy. Medling et al. (2014) studied the central supermassive black hole and nuclear disks in 17 nearby ULIRGs. Hicks et al. (2013) surveyed 10 nearby Seyfert galaxies in matched AGN and quiescent pairs, searching for evidence of the mechanisms which fuel the central AGN. Müller-Sánchez et al. (2017) conducted the KONA survey of 40 nearby galaxies with central AGN, some of which are luminous enough in the infrared to be classified as ULIRGs.

In order to understand the rate at which OH $119\mu\text{m}$ identifies a spatially resolved, independent outflow, we conducted a survey of 10 ULIRGs. In section 4.2 we present the sample. Section 4.3 describes our data reduction and analysis methods. Section 4.4 presents the emission line kinematics obtained

Table 4.1: The sample of nearby ULIRGs observed for this paper with OSIRIS

Name	Redshift	OH 119 μ m v_{98} km/s	$\log L_{\text{bol}}$ L_{\odot}	α_{AGN}
IRAS F00188-0856	0.128	-1805	12.39	0.51
IRAS F00456-2904	0.110	–	12.32	0.83
IRAS F03158+4227	0.134	-2044	12.63	0.47
IRAS F05189-2524	0.042	-850	12.22	0.71
IRAS F08572+3915	0.058	-1100	12.20	0.70
IRAS F10565+2448	0.042	-950	12.10	0.47
IRAS F14348-1447	0.083	-900	12.40	0.17
IRAS F15206+3342	0.125	–	12.30	0.64
IRAS F19297-0406	0.086	-1000	12.44	0.23
IRAS F20414-1651	0.087	-240	12.22	0.00

in this survey. In Section 4.5 we explore the implications of this survey for the findings from Herschel. Section 4.6 contains our conclusions. Throughout the chapter, we adopt $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.28$ and $\Omega_{\Lambda} = 0.72$ (Hinshaw et al., 2009).

4.2 Sample

This chapter presents spatially resolved near-infrared observations of 10 ULIRGs (Table 4.1). These ULIRGs were chosen to span a range of observed OH 119 μ m velocities, but with a particular focus on ULIRGs where the 98th-percentile velocity of the absorption or emission feature (v_{98}) measurement was ≥ 800 km/s. The sample spans a range of AGN dominance (α_{AGN} , the fraction of the bolometric luminosity attributable to AGN activity). The sample was constrained to objects which have a suitable tip-tilt star for use with laser guide star adaptive optics (LGS AO) within 60'' and $m_R \geq 18$.

4.3 Methods

We observed 10 ULIRGs with OSIRIS, the near-infrared integral field unit spectrograph at the W.M. Keck Observatory (Larkin et al., 2006) using LGS AO. Observations for this survey were conducted on various nights from October 2014 through June 2016. Table 4.2 lists the observations and scales

Table 4.2: OSIRIS observations for this chapter.

Name	RA	Dec	Obs Date	Filter	Scale	Time min
F00188-0856	00:21:26.5	-08:39:26	2015-08-07	Kbb	0.035''	40
F00456-2904	00:48:06.8	-28:48:19	2014-10-13	Kbb	0.035''	80
F03158+4227	03:19:12.4	+42:38:28	2015-09-22	Kbb	0.035''	60
F05189-2524	05:21:01.4	-25:21:45	2015-10-25	Kbb	0.100''	140
F08572+3915	09:00:25.4	+39:03:54	2015-05-11	Kbb	0.035''	80
F10565+2448	10:59:18.1	+24:32:34	2016-04-20	Kbb	0.100''	60
F14348-1447	14:37:38.4	-15:00:20	2016-04-20	Kbb	0.035''	70
F15206+3342	15:22:38.0	+33:31:36	2016-06-25	Kbb	0.100''	50
F19297-0406	19:32:21.2	-03:59:56	2015-09-22	Kbb	0.100''	80
F20414-1651	20:44:18.2	-16:40:16	2014-10-13	Kbb	0.035''	150

for each target. OSIRIS uses a lenslet array to subdivide the field of view into ~ 3000 individual spectra at $R \sim 3000$. Our observations were conducted with the $0.035''/\text{px}$ scale lenslet and the $0.100''/\text{px}$ scale lenslet in the Kbb filter. We chose the lenslet scaling to maximize the coverage of parts of the ULIRG in which were able to achieve a good (~ 30) signal-to-noise ratio for the observed emission lines. These observations spanned the upgrade of the OSIRIS detector in January 2016. Data from before January 2016 used the old OSIRIS detector, which had worse noise properties. Data after January 2016 uses the new OSIRIS spectrograph detector.

Most of our observations were conducted at 2 pointings which overlapped, creating a mosaic across the center of each target. The data were collected in an object-sky-object pattern to capture changes in sky brightness. Each exposure was 600 s long.

We observed a UKIRT faint standard star at the start of each observing night. When practical, we also observed a UKIRT faint standard at the end of the night. These were used to flux calibrate our data. We also observed telluric stars once every two hours, at the average airmass of our objects to provide a reference for atmospheric absorption.

4.3.1 Data Reduction

Our observations were reduced using the Keck OSIRIS data reduction pipeline, v4.0b1¹ and the wavelength solution from early 2015² for data from before the OSIRIS detector upgrade, or with the wavelength solutions from early 2016 for data taken after the OSIRIS detector upgrade. The pipeline contains modules to subtract sky frames, adjust channel levels, remove inter-pixel crosstalk, identify glitches, clean cosmic rays, extract a deconvolved spectrum for each spatial pixel (spaxel), assemble a data cube of spaxels, correct for atmospheric dispersion, subtract scaled residual sky lines (from a technique described by Davies (2007)), and perform a telluric correction. Our telluric correction was done with the telluric star observed closest to each individual observation.

Before each observation, we imaged our target’s tip-tilt star, which we use as an estimate of the PSF during observations of each target galaxy. Medling et al. (2014) found that the tip-tilt star PSF did not vary significantly from the start to the end of each observation. However, the PSF will vary based on the distance between the tip-tilt guidestar and the target, and based on the isokinetic angle at Mauna Kea. We model the observed PSF as a convolution of the measured PSF, broadened by a gaussian consistent with a degradation in Strehl of $1/e$ for a separation of $75''$ (from van Dam et al., 2006) between the tip/tilt star and the center of the pointing on the galaxy. This correction results in a 15% degradation in the reference PSF for the average tip-tilt star separation. We convolve this PSF into our raw data, which also serves to reduce the noise in each spaxel by cross-correlating neighboring spaxels.

4.3.2 Line Fitting

To fit emission lines, we used the techniques described in Section 3.2.3, where we fit a simultaneous continuum and emission model to each spaxel. The continuum is fit as a simple power-law across the entire wavelength range. The

¹The beta version was unreleased when we did this work, but it is now released on GitHub at <https://github.com/Keck-DataReductionPipelines/OSIRISDRP>.

²The so-called “median combined scans” version, which eliminated bad pixels and cosmic rays in the rectification matrix

Table 4.3: Emission lines detected for each target in the survey.

Name	Pa- α	Br- δ	Br- γ	H ₂ 1 – 0 S(3)	H ₂ 1 – 0 S(1)
IRAS F00188-0856	X				
IRAS F00456-2904	X			X	
IRAS F03158+4227	X			X	
IRAS F05189-2524			X		X
IRAS F08572+3915	X				
IRAS F10565+2448		X	X	X	X
IRAS F14348-1447	X		X	X	X
IRAS F15206+3342	X	X			
IRAS F19297-0406	X			X	X
IRAS F20414-1651	X		X	X	X

emission lines fit varied for each object based on it’s redshift. The particular emission lines fit are shown in table 4.3

4.4 Results


We successfully measure kinematics for emission lines in every target. In one target, we are not able to measure H₂ emission (F00188-0856) because the H₂ emission falls directly on top of a telluric feature.

4.4.1 F00188-0856


The Pa- α emission for F00188-0856 is shown in Figure 4.1 and Figure 4.2. The primary component (**A**) shows evidence of a rotating disk. The secondary component appears to be a higher-dispersion wing in the rotating disk. There is a small anomalous velocity feature in the north east corner of component **A** that could be consistent with a ~ 300 km/s inflow, but there is no evidence of the -1850 km/s in the OSIRIS field of view.

4.4.2 F00456-2904


The Pa- α emission for F00456-2904 is shown in Figure 4.3. The primary component (**A**) shows strong evidence for a rotating disk, with elevated velocity



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.1: Pa- α emission from F00188-0856 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. This is the primary Pa- α component. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel. The velocity of this component is consistent with a rotating disk.



(a) Pa- α flux from component **B**



(b) Pa- α velocity from component **B**



(c) Pa- α velocity dispersion from component **B**

Figure 4.2: Pa- α emission from F00188-0856 The top panel (a) shows the flux in Pa- α for component **B**. The center panel (b) shows the velocity of Pa- α for component **B**. The lower panel (c) shows the velocity dispersion of the Pa- α for component **B**. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel. The velocity of this component is consistent with a rotating disk.

dispersion along the line of nodes of the disk. This is expected when we under-resolve a large rotation curve (see Section 3.3.1).

The H_2 $1 - 0$ $S(3)$ emission for F00456-2904 is shown in Figure 4.4. The H_2 $1 - 0$ $S(3)$ emission shows weaker signatures of rotation than the Pa- α emission in Figure 4.3, and has a much lower velocity dispersion.

4.4.3 F03158+4227


The Pa- α emission for F03158+4227 is shown in Figure 4.5 and Figure 4.6. Neither detected component shows signs of rotation. However, **B** does have -400 km/s blue-shifted emission, which could be indicative of the blue-shifted emission observed in OH 119 μm by Herschel, though the Herschel emission is at -2044 km/s (Spoon et al., 2013).

The H_2 $1 - 0$ $S(3)$ emission for F03158+4227 is shown in Figure 4.7. H_2 does show evidence of rotation, as well as a high-dispersion, blue shifted component in the north-west corner of the OSIRIS field of view. This could be indicative of the base of an outflow which is accelerated to the full -2044 km/s (Spoon et al., 2013) beyond the field of view of OSIRIS.


4.4.4 F05189-2524

The Br- γ emission for F05189-2524 is shown in Figure 4.8 and Figure 4.9. The brighter Br- γ component in Figure 4.8 does not show any evidence of rotation, and the fainter component in Figure 4.9 shows only slight evidence for rotation. However, both components have high ~ 500 km/s velocity dispersions in some locations, indicating that F05189-2524 does not have a cool, well-ordered disk.


The H_2 $1 - 0$ $S(1)$ emission for F05189-2524 is shown in Figure 4.10. The H_2 $1 - 0$ $S(1)$ does not show a significant velocity gradient, or any evidence of a kinematically separate component. The dispersion rises towards the center of the H_2 $1 - 0$ $S(1)$ emission in F05189-2524, consistent with the rising velocity dispersion observed in Br- γ .



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.3: Pa- α emission from F00456-2904 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. This component shows strong evidence for a rotating disk, with elevated velocity dispersion along the line of nodes of the disk. This is expected when we under-resolve a large rotation curve (see Section 3.3.1)



(a) $H_2 1 - 0 S(3)$ flux from component A




(b) $H_2 1 - 0 S(3)$ velocity from component A




(c) $H_2 1 - 0 S(3)$ velocity dispersion from component A


Figure 4.4: $H_2 1 - 0 S(3)$ emission from F00456-2904 The top panel (a) shows the flux in $H_2 1 - 0 S(3)$ for component A. The center panel (b) shows the velocity of $H_2 1 - 0 S(3)$ for component A. The lower panel (c) shows the velocity dispersion of the $H_2 1 - 0 S(3)$ for component A. We fit 2 emission components to $H_2 1 - 0 S(3)$ in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel. This component shows weaker signatures of rotation than the Pa- α emission in Figure 4.3, and has a much lower velocity dispersion.



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.5: Pa- α emission from F03158+4227 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Pa- α flux from component **B**




(b) Pa- α velocity from component **B**




(c) Pa- α velocity dispersion from component **B**


Figure 4.6: Pa- α emission from F03158+4227 The top panel (a) shows the flux in Pa- α for component **B**. The center panel (b) shows the velocity of Pa- α for component **B**. The lower panel (c) shows the velocity dispersion of the Pa- α for component **B**. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel. The -400 km/s blue-shifted emission SE of the central nucleus could be indicative of the blue-shifted emission observed in OH $119\mu\text{m}$ by Herschel, though the OH $119\mu\text{m}$ emission is at -2044 km/s (Spoon et al., 2013).



(a) $H_2 1 - 0 S(3)$ flux from component A




(b) $H_2 1 - 0 S(3)$ velocity from component A




(c) $H_2 1 - 0 S(3)$ velocity dispersion from component A


Figure 4.7: $H_2 1 - 0 S(3)$ emission from F03158+4227 The top panel (a) shows the flux in $H_2 1 - 0 S(3)$ for component A. The center panel (b) shows the velocity of $H_2 1 - 0 S(3)$ for component A. The lower panel (c) shows the velocity dispersion of the $H_2 1 - 0 S(3)$ for component A. H_2 does show evidence of rotation, as well as a high-dispersion, blue shifted component in the north-west corner of the OSIRIS field of view. This could be indicative of the base of an outflow which is accelerated to the full -2044 km/s (Spoon et al., 2013) beyond the field of view of OSIRIS.



(a) Br- γ flux from component A




(b) Br- γ velocity from component A




(c) Br- γ velocity dispersion from component A

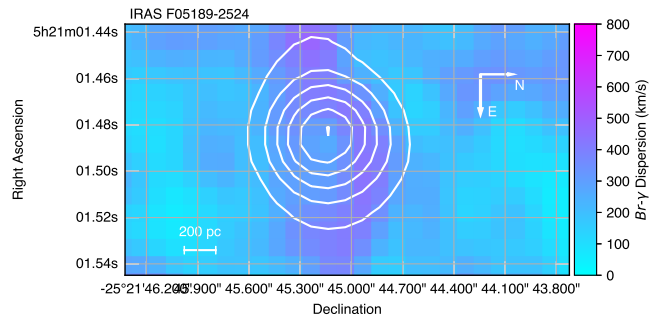
Figure 4.8: Br- γ emission from F05189-2524 The top panel (a) shows the flux in Br- γ for component A. The center panel (b) shows the velocity of Br- γ for component A. The lower panel (c) shows the velocity dispersion of the Br- γ for component A. We fit 2 emission components to Br- γ in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Br- γ flux from component **B**




(b) Br- γ velocity from component **B**




(c) Br- γ velocity dispersion from component **B**


Figure 4.9: Br- γ emission from F05189-2524 The top panel (a) shows the flux in Br- γ for component **B**. The center panel (b) shows the velocity of Br- γ for component **B**. The lower panel (c) shows the velocity dispersion of the Br- γ for component **B**. We fit 2 emission components to Br- γ in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) H_2 1 – 0 $S(1)$ flux from component A




(b) H_2 1 – 0 $S(1)$ velocity from component A




(c) H_2 1 – 0 $S(1)$ velocity dispersion from component A

Figure 4.10: H_2 1 – 0 $S(1)$ emission from F05189-2524 The top panel (a) shows the flux in H_2 1 – 0 $S(1)$ for component A. The center panel (b) shows the velocity of H_2 1 – 0 $S(1)$ for component A. The lower panel (c) shows the velocity dispersion of the H_2 1 – 0 $S(1)$ for component A.



(a) Pa- α flux from component A



(b) Pa- α velocity from component A


Figure 4.11: Pa- α emission from F08572+3915 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.

4.4.5 F08572+3915


The Pa- α emission for F08572+3915 is shown in Figure 4.11 and Figure 4.12. The H₂ 1 – 0 $S(3)$ and H₂ 1 – 0 $S(1)$ emission for this object are presented in Rupke & Veilleux (2013). The Pa- α shows a rotation curve consistent with the rotation observed in H₂ 1 – 0 $S(1)$ in Rupke & Veilleux (2013).

4.4.6 F10565+2448


We observed Br- δ and Br- γ in F10565+2448, both of which show evidence for rotation. The Br- δ emission for F10565+2448 is shown in Figure 4.13. The high dispersion on the eastern side of the disk suggests that it has been re-



(a) Pa- α flux from component **B**



(b) Pa- α velocity from component **B**



(c) Pa- α velocity dispersion from component **B**

Figure 4.12: Pa- α emission from F08572+3915 The top panel (a) shows the flux in Pa- α for component **B**. The center panel (b) shows the velocity of Pa- α for component **B**. The lower panel (c) shows the velocity dispersion of the Pa- α for component **B**. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.

cently disrupted. The Br- γ emission is shown in Figure 4.14. There are several spots where Br- γ is observed with significant red velocities and dispersions, which could be indicative of inflows in the line of sight, or outflows on the far side of the observed disk.


The H₂ 1 – 0 $S(3)$ emission for F10565+2448 is shown in Figure 4.15. The H₂ 1 – 0 $S(1)$ emission is shown in Figure 4.16. Both H₂ transitions show clear evidence of rotation, and high velocity dispersion on the western side of the disk. Taken together with the velocity dispersion in Br- δ , this might suggest a disruption history for F10565+2448. If the eastern edge was disrupted first, the disruption may have triggered a burst of star formation, giving rise to the observed excess Br- δ emission and higher dispersions. The higher dispersions on the western edge of the disk in H₂ then represent more recently disrupted gas, where the warm H₂ has been shocked.

4.4.7 F14348-1447


The Pa- α emission for F14348-1447 is shown in Figure 4.17 and Figure 4.18. The primary Pa- α emission shows less dispersion, and is blue shifted from the rest-frame velocity of the system.

The Br- γ emission for F14348-1447 is shown in Figure 4.19. There is only very limited evidence for ordered rotation in Br- γ . There is evidence for a secondary source of emission in the North-East corner of the OSIRIS field of view, which has both a significant dispersion and a higher velocity than the surrounding Br- γ emission.


The H₂ 1 – 0 $S(3)$ emission for F14348-1447 is shown in Figure 4.20. The H₂ 1 – 0 $S(1)$ emission is shown in Figure 4.21. Both H₂ components do not show any evidence for rotation in the central regions. Each has pockets of higher-dispersion gas, blue-shifted and traveling at ~ 500 km/s. Combined with the blue-shifted Pa- α emission, this is consistent with the OH 119 μ m absorption observed in Veilleux et al. (2013a).



(a) Br- δ flux from component A




(b) Br- δ velocity from component A




(c) Br- δ velocity dispersion from component A


Figure 4.13: Br- δ emission from F10565+2448 The top panel (a) shows the flux in Br- δ for component A. The center panel (b) shows the velocity of Br- δ for component A. The lower panel (c) shows the velocity dispersion of the Br- δ for component A.



(a) Br- γ flux from component A




(b) Br- γ velocity from component A




(c) Br- γ velocity dispersion from component A


Figure 4.14: Br- γ emission from F10565+2448 The top panel (a) shows the flux in Br- γ for component A. The center panel (b) shows the velocity of Br- γ for component A. The lower panel (c) shows the velocity dispersion of the Br- γ for component A.



(a) H_2 1 – 0 $S(3)$ flux from component A




(b) H_2 1 – 0 $S(3)$ velocity from component A




(c) H_2 1 – 0 $S(3)$ velocity dispersion from component A


Figure 4.15: H_2 1 – 0 $S(3)$ emission from F10565+2448 The top panel (a) shows the flux in H_2 1 – 0 $S(3)$ for component A. The center panel (b) shows the velocity of H_2 1 – 0 $S(3)$ for component A. The lower panel (c) shows the velocity dispersion of the H_2 1 – 0 $S(3)$ for component A.



(a) H_2 1 – 0 $S(1)$ flux from component A




(b) H_2 1 – 0 $S(1)$ velocity from component A




(c) H_2 1 – 0 $S(1)$ velocity dispersion from component A


Figure 4.16: H_2 1 – 0 $S(1)$ emission from F10565+2448 The top panel (a) shows the flux in H_2 1 – 0 $S(1)$ for component A. The center panel (b) shows the velocity of H_2 1 – 0 $S(1)$ for component A. The lower panel (c) shows the velocity dispersion of the H_2 1 – 0 $S(1)$ for component A.



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.17: Pa- α emission from F14348-1447 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Pa- α flux from component **B**




(b) Pa- α velocity from component **B**




(c) Pa- α velocity dispersion from component **B**


Figure 4.18: Pa- α emission from F14348-1447 The top panel (a) shows the flux in Pa- α for component **B**. The center panel (b) shows the velocity of Pa- α for component **B**. The lower panel (c) shows the velocity dispersion of the Pa- α for component **B**. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Br- γ flux from component A




(b) Br- γ velocity from component A




(c) Br- γ velocity dispersion from component A


Figure 4.19: Br- γ emission from F14348-1447 The top panel (a) shows the flux in Br- γ for component A. The center panel (b) shows the velocity of Br- γ for component A. The lower panel (c) shows the velocity dispersion of the Br- γ for component A. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) H₂ 1 – 0 S(3) flux from component A




(b) H₂ 1 – 0 S(3) velocity from component A




(c) H₂ 1 – 0 S(3) velocity dispersion from component A


Figure 4.20: H₂ 1 – 0 S(3) emission from F14348-1447 The top panel (a) shows the flux in H₂ 1 – 0 S(3) for component A. The center panel (b) shows the velocity of H₂ 1 – 0 S(3) for component A. The lower panel (c) shows the velocity dispersion of the H₂ 1 – 0 S(3) for component A.



(a) H_2 1 – 0 $S(1)$ flux from component A



(b) H_2 1 – 0 $S(1)$ velocity from component A



(c) H_2 1 – 0 $S(1)$ velocity dispersion from component A

Figure 4.21: H_2 1 – 0 $S(1)$ emission from F14348-1447 The top panel (a) shows the flux in H_2 1 – 0 $S(1)$ for component A. The center panel (b) shows the velocity of H_2 1 – 0 $S(1)$ for component A. The lower panel (c) shows the velocity dispersion of the H_2 1 – 0 $S(1)$ for component A.

4.4.8 F15206+3342


The Pa- α emission for F15206+3342 is shown in Figure 4.22 and Figure 4.23. The primary Pa- α component shows a compact core, with evidence for rotation, and high dispersion on the eastern side of the disk. The secondary component shows almost no significant velocity structure, and a very low velocity dispersion, perhaps indicative of a well ordered nuclear disk viewed face-on. The Br- δ emission for F15206+3342 is shown in Figure 4.24 and Figure 4.25. The primary component shows high dispersion, and an extended region to the west of the core of the ULIRG.

4.4.9 F19297-0406


The Pa- α emission for F19297-0406 is shown in Figure 4.26. There is a clear rotation signature, centered around the central emission point for Pa- α in this object.

The H₂ 1 – 0 $S(3)$ emission for F19297-0406 is shown in Figure 4.27. This H₂ transition shows rotation consistent with that observed in Pa- α .


The H₂ 1 – 0 $S(1)$ emission for F19297-0406 is shown in Figure 4.28. There is an extended tidal object to the south of the primary Nucleus which contains significant H₂ 1 – 0 $S(1)$ emission. This object shows evidence of rotation rather than an outflow, since the blue velocities are not centered on the flux from the tidal object, and the higher dispersion region also does not correlate directly with the blue velocities. However, it is possible that this is a partially obscured outflow where we are seeing only the blue velocities from H₂ where it has broken through another cloud of gas.



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.22: Pa- α emission from F15206+3342 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Pa- α flux from component **B**




(b) Pa- α velocity from component **B**




(c) Pa- α velocity dispersion from component **B**


Figure 4.23: Pa- α emission from F15206+3342 The top panel (a) shows the flux in Pa- α for component **B**. The center panel (b) shows the velocity of Pa- α for component **B**. The lower panel (c) shows the velocity dispersion of the Pa- α for component **B**. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Br- δ flux from component A




(b) Br- δ velocity from component A




(c) Br- δ velocity dispersion from component A


Figure 4.24: Br- δ emission from F15206+3342 The top panel (a) shows the flux in Br- δ for component A. The center panel (b) shows the velocity of Br- δ for component A. The lower panel (c) shows the velocity dispersion of the Br- δ for component A. We fit 2 emission components to Br- δ in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Br- δ flux from component **B**




(b) Br- δ velocity from component **B**




(c) Br- δ velocity dispersion from component **B**


Figure 4.25: Br- δ emission from F15206+3342 The top panel (a) shows the flux in Br- δ for component **B**. The center panel (b) shows the velocity of Br- δ for component **B**. The lower panel (c) shows the velocity dispersion of the Br- δ for component **B**. We fit 2 emission components to Br- δ in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) Pa- α flux from component A




(b) Pa- α velocity from component A




(c) Pa- α velocity dispersion from component A


Figure 4.26: Pa- α emission from F19297-0406 The top panel (a) shows the flux in Pa- α for component A. The center panel (b) shows the velocity of Pa- α for component A. The lower panel (c) shows the velocity dispersion of the Pa- α for component A. We fit 2 emission components to Pa- α in each spaxel. The primary component for each spaxel was the component with a larger flux in that spaxel.



(a) $H_2 1 - 0 S(3)$ flux from component A




(b) $H_2 1 - 0 S(3)$ velocity from component A




(c) $H_2 1 - 0 S(3)$ velocity dispersion from component A


Figure 4.27: $H_2 1 - 0 S(3)$ emission from F19297-0406 The top panel (a) shows the flux in $H_2 1 - 0 S(3)$ for component A. The center panel (b) shows the velocity of $H_2 1 - 0 S(3)$ for component A. The lower panel (c) shows the velocity dispersion of the $H_2 1 - 0 S(3)$ for component A.



(a) H_2 1 – 0 $S(1)$ flux from component A



(b) H_2 1 – 0 $S(1)$ velocity from component A



(c) H_2 1 – 0 $S(1)$ velocity dispersion from component A

Figure 4.28: H_2 1 – 0 $S(1)$ emission from F19297-0406 The top panel (a) shows the flux in H_2 1 – 0 $S(1)$ for component A. The center panel (b) shows the velocity of H_2 1 – 0 $S(1)$ for component A. The lower panel (c) shows the velocity dispersion of the H_2 1 – 0 $S(1)$ for component A.

4.5 Discussion

We obtained spatially resolved NIR emission line kinematics for 10 galaxies. Broadly, they can be classified into three groups:

1. Galaxies with clear evidence of outflowing gas from the nuclear regions observed with OSIRIS.
2. Galaxies with an NIR spatially resolved velocity profile that matches the velocity profile observed by Herschel, but with no evidence of spatially resolved molecular outflows.
3. Galaxies with an NIR spatially resolved velocity profile that cannot fully explain the velocity profile observed by Herschel.

We find evidence for spatially resolved high velocity molecular outflows in F08572+3915, and our findings are consistent with the outflows presented in Rupke & Veilleux (2013). In F08572+3915 there is a spatially-resolved, high velocity component observed in H₂ at -1000 km/s, which matches the -1000 km/s blue wing observed in OH 119 μ m in Veilleux et al. (2013b) (see Figures 4.11 and 4.12). F08572+3915 is a clear example of a system in group 1. We also observe evidence of high velocity molecular outflows in the ULIRG F14348-1447 (see Figure 4.19) and F19297-0406 (see Figure 4.28) which we classify in group 1. These three galaxies are the only 3 in our sample of 10 galaxies which have clear evidence of a spatially resolved outflow.

Some of the ULIRGs observed show an NIR spatially resolved velocity profile that matches the velocity profile observed by Herschel, but with no evidence of spatially resolved molecular outflows. These systems show NIR rotation curves which span the velocities observed by Herschel. The spatially resolved NIR emission line kinematics of F20414-1651 (see chapter 3 and specifically Figure 3.22) show a velocity profile consistent with the velocities observed in OH 119 μ m (Spoon et al., 2013) in a single large Herschel spaxel, however, there is no evidence of a molecular outflow in the spatially resolved NIR observations which observed only the nuclear regions. The ULIRGs F00188-0856 (Section 4.4.1) and F00456-2904 (Section 4.4.2) also show high velocity rotation

curve kinematic profiles but do not show evidence of high velocity molecular outflows. These galaxies are classified in group 2.

That we find 3 ULIRGs in group 2 suggests that spectroscopy which is spatially integrated over the $9.4'' \times 9.4''$ Herschel spaxel is not enough to uniquely identify outflowing molecular gas from nearby ULIRGs. The OH $119\mu\text{m}$ spectroscopy of these three ULIRGs presented in Spoon et al. (2013) and Veilleux et al. (2013b) shows a blue wing, however, we can explain all of the observed velocity in the blue wing by examining the spatially resolved NIR kinematics. These ULIRGs have high velocities present in their rotation curves and represent cases where OH $119\mu\text{m}$ spectroscopy is not sufficient to identify a molecular outflow. Spatially resolved NIR spectroscopy does not support the conclusion that these systems have high velocity molecular outflows.

The remaining four ULIRGs (F03158+4227, F05189-2524, F10565+2448, and F15206+3342) do not have evidence of outflows or rotation curves with velocity profiles that match the velocities observed in OH $119\mu\text{m}$ by Herschel. These ULIRGs all show velocity rotation curves with amplitudes ranging from 100 km/s to 250 km/s. The velocities observed in the spatially resolved NIR observations for these systems are less than the velocities observed in OH $119\mu\text{m}$. We classify these ULIRGs group 3.

ULIRGs in group 3 are systems where the spatially resolved NIR spectroscopy presented here does not eliminate the possibility that there are high velocity molecular outflows elsewhere the system outside of the OSIRIS field of view ($1.4'' \times 0.73''$). Since the spatially resolved NIR spectroscopy does not cover an area as large as the OH $119\mu\text{m}$ spectroscopy presented in Spoon et al. (2013) and Veilleux et al. (2013b), it is possible that the high velocity outflows are present outside the field of view of OSIRIS.

4.6 Conclusions

We presented the kinematics of 10 ULIRGs observed with OSIRIS. Our conclusions are split according to the grouping in Section 4.5. We found evidence for high velocity outflows consistent with those observed in Spoon et al. (2013) and Veilleux et al. (2013b) in 3 systems (group 1). In three other sys-

tems (group 2), we found that the spatially resolved NIR kinematics observed with OSIRIS showed rotation curves with velocities when spatially integrated, were enough to explain the velocity profile observed in OH $119\mu\text{m}$, and the spatially resolved NIR observations did not provide direct evidence of outflows. In the remaining four systems (group 3), the spatially resolved NIR observations were unable to account for the spatially integrated velocity profile observed in OH $119\mu\text{m}$, and no evidence for outflows was seen in the NIR.

4.6.1 ULIRGs with clear evidence of outflows in NIR spectroscopy and the spatially integrated OH $119\mu\text{m}$ velocity profile

Our observations of F08572+3915 confirmed the presence of an outflow which was previously presented in Rupke & Veilleux (2013). For the two other systems with molecular outflows observed with spatially-resolved NIR spectroscopy (F14348-1447 and F19297-0406) we do not see kinematically distinct components centered at the observed OH $119\mu\text{m}$ velocity from Herschel. Rather the three observed outflows suggest that the 98th percentile measurement of the blue emission in OH $119\mu\text{m}$ from the Herschel space telescope is indeed the extreme blue velocity wing of the outflowing gas when that outflowing gas is present, and that the spatially resolved outflows show slower velocities (-220 km/s and -310 km/s) than were originally suggested by the OH $119\mu\text{m}$ observed in Spoon et al. (2013) and Veilleux et al. (2013b).

4.6.2 ULIRGs with rotation curves which explain the spatially integrated OH $119\mu\text{m}$ velocity profile

In three systems, we observed rotation curves in the NIR which can explain the spatially integrated velocity profile observed in OH $119\mu\text{m}$. These ULIRGs represent cases where blue-shifted velocities in OH $119\mu\text{m}$ can be confused for outflows, where they appear in spatially-resolved spectroscopy to arise due to the rotation curve kinematics present in many ULIRGs. This kind of ambiguity in interpreting the results of Spoon et al. (2013) and Veilleux et al. (2013b) limits our ability to interpret spatially-unresolved, extreme blue-shifted velocities in ULIRGs as a complete statistical sample of ULIRGs with outflow-

ing gas. The molecular outflows could be present outside of the rather narrow field of view of OSIRIS ($1.4'' \times 0.73''$), but their presence is not required to explain the origins of the OH $119\mu\text{m}$ observations. Rather, further study is needed to determine the level at which large samples of spatially un-resolved outflows could be explained by ULIRGs with rotation curves which match the OH $119\mu\text{m}$ velocity profiles.

4.6.3 ULIRGs with no outflow and where the rotation curve is insufficient to explain the integrated OH $119\mu\text{m}$ velocity profile

In the remaining four ULIRGs, we were unable to detect the presence of a high velocity component (either an outflow, or a high velocity rotation curve) in the spatially resolved NIR spectroscopy of the nuclear regions, despite the presence of a blue wing in the OH $119\mu\text{m}$ observations. A high velocity outflow may be present in these systems outside the field of view of OSIRIS. Although we cannot draw a direct conclusion about the presence or absence of a molecular outflow in these systems, we can conclude that the outflow observed in OH $119\mu\text{m}$ does not currently extend to the nucleus of these systems.

4.6.4 Concluding thoughts

When we launched this survey, we were hoping to use the nature of the outflow in the nucleus of the system to constrain the driving mechanisms behind the spatially-unresolved outflows observed in OH $119\mu\text{m}$ by Herschel. That so many of our targets had no direct evidence of a nuclear molecular outflow suggests that we must observe a significantly larger sample of ULIRGs in order to capture statistically significant examples of molecular outflows in nuclear regions of ULIRGs, and study the spatially-unresolved outflows' driving mechanisms.

Even in our small sample, some of the velocity profiles observed by Herschel in OH $119\mu\text{m}$ can be explained by rapid rotation in the nuclear region of the ULIRG, without requiring the presence of a molecular outflow. A larger sample should address what fraction of velocity profiles observed in OH $119\mu\text{m}$ with Herschel can be explained by rapid rotation in the nucleus of the ULIRG,

and what impact this ambiguity has on the broader statistical conclusions about the prevalence of OH $119\mu\text{m}$ outflows (e.g. Spoon et al., 2013; Veilleux et al., 2013b; Contursi et al., 2013; González-Alfonso et al., 2013).

Part III

Linear Quadratic Gaussian Control for Adaptive Optics

Chapter 5

A laboratory demonstration of an LQG controller for correcting frozen flow turbulence in adaptive optics systems

5.1 Introduction

We present the implementation of a predictive Fourier control (PFC) algorithm (Poyneer et al., 2007; Poyneer & Véran, 2008) in the lab to correct for frozen flow turbulence in adaptive optics systems. Recently, linear-quadratic Gaussian (LQG) controllers (Gavel & Wiberg, 2003; Le Roux et al., 2004) have been used for the stable control of adaptive optics systems (Petit et al., 2008), and for the suppression of specific frequency vibrations in operating adaptive optics systems (Poyneer et al., 2014; Beuzit et al., 2008). Predictive Fourier Control is an LQG method which explicitly corrects for frozen-flow turbulence crossing the telescope aperture. It diagonalizes that problem by working in Fourier space, making the method computationally feasible for high-order AO systems.

Several groups have formulated LQG controllers for the optimal control of adaptive optics systems. Gavel & Wiberg (2003) initially developed a Kalman filter for optimal AO control. Le Roux et al. (2004) formulated an LQG

controller for adaptive optics. Petit et al. (2008) demonstrated the use of LQG in the laboratory to suppress vibrations. Sivo et al. (2014) have applied this on sky with the CANARY demonstrator. LQG control has progressed from the laboratory to the latest generation of astronomical instruments. The recently commissioned Gemini Planet Imager (Macintosh et al., 2014) and SPHERE (Beuzit et al., 2008) instruments have both demonstrated the use of optimal LQG control (Poyneer et al., 2014; Petit et al., 2014). GPI’s AO system uses an LQG controller to suppress vibrations in both tip-tilt and focus (Poyneer et al., 2014). Similarly, SPHERE uses an LQG controller to suppress vibrations in only the tip-tilt modes (Petit et al., 2014). These controllers are implemented with the same LQG framework that we will apply to correct frozen flow turbulence across spatial frequencies beyond the tip-tilt regime.

Our experimental setup provides a tightly integrated way to test the use of an LQG controller both in the lab and on-sky. The ShaneAO system with the ShARCs Camera on the Shane 3-m telescope at Lick Observatory (Gavel et al., 2014; Kupke et al., 2012) provides an on-sky instrument. At the Lab for Adaptive Optics we have configured an adaptive optics test bench (Laag et al., 2008) to mimic the hardware and physical conditions for ShaneAO. These closely coupled systems allow us to test and develop techniques in a controlled lab before proceeding to the instrument on-sky for functional tests.

This chapter describes a laboratory implementation of the full Predictive Fourier Control (PFC) algorithm. In Section 5.2 we provide a brief overview of the test bench, a full, physical simulator designed to test adaptive optics hardware, software and techniques. In Section 5.3 we describe the PFC algorithm, and the way this algorithm interacts with an operating adaptive optics system. In Section 5.4 we describe the implementation of the PFC algorithm in the lab, and the measured performance improvement. Section 5.5 discusses the significance of this laboratory demonstration.

Our framework, described in Section 5.3 and diagramed in Figure 5.2, consists of three components:

1. Phase reconstruction as performed by a standard adaptive optics system.
2. The Predictive Fourier Controller, which uses a Kalman Filter to provide

a best estimate of the correction in the future.

3. Fourier Wind Identification, which provides the state space model to the Predictive Fourier Controller.

Fourier Wind Identification is described theoretically in Section 5.3.1, and demonstrated in the lab in Section 5.4.4. We discuss the theoretical framework of the Predictive Fourier Controller in Section 5.3.2, and demonstrate its operation in Sections 5.4.2 and 5.4.5.

5.2 Experimental Setup

Our experiments are performed on a modified and optimized version of the UCSC Laboratory for Adaptive Optics (LAO) MCAO/MOAO testbed (Laag et al., 2008). In the past, its utility and versatility has been shown by simulating MCAO and open-loop MOAO on a 30-m equivalent aperture telescope (Ammons et al., 2006), demonstrating the positive effect of linearity calibrations to Shack-Hartmann wavefront sensors during open-loop MOAO operation on a 10-m telescope (Ammons et al., 2007) and for developing wavefront reconstruction and control algorithms for MCAO on an 8-m class telescope (Laag et al., 2008). Currently, we have reconfigured the testbed to match as closely as possible the new adaptive optics system on the Shane 3-m telescope at Lick Observatory (ShaneAO) (Gavel et al., 2014) so that the transition from lab to a real system is as seamless as possible. The optical layout of the reconfigured system is shown in Figure 5.1.

For our purposes, the testbed is operated in single conjugate (SCAO) mode with two deformable mirrors (DMs): a low (spatial and temporal) frequency, high-stroke Alpao DM52 “woofer” and a high-frequency, lower-stroke, Hamamatsu SLM-X8267 spatial light modulator (SLM) “tweeter”. As the test bench does not operate in real time, we treat both mirrors as having instantaneous temporal properties and ignore the slower temporal performance of the “woofer” mirror.

Both DMs are conjugate to the equivalent of the atmosphere’s ground layer represented by a phase plate mounted on a linear stepper motor. The




Figure 5.1: Testbed layout for wind prediction tests. The light source is a superluminescent diode laser. The spatial light modulator is a high-frequency deformable mirror. In single conjugate mode only one of the Shack-Hartmann wavefront sensors is used.

phase plate is 160 subapertures across with a minimum wind speed of 0.01 subapertures per timestep. The linear stepper can be reliably repositioned so that repeated runs can be performed across the same section of a phase plate and a higher signal-to-noise ratio is achieved by co-adding runs. The size of the phase plate ensures that each test run contains 4.4 pupil crossings.

The woofer’s actuator layout is 4×4 across the pupil. The edges of the woofer are pinned, so outer actuators have reduced stroke and are not used. This is nearly identical to the configuration of the woofer used in ShaneAO. The SLM has a resolution of 768×768 , however pixels are binned to emulate a 36×36 actuator configuration similar to that of the MEMS tweeter on ShaneAO. The wavefront sensor has a 36×36 lenslet array layout, and is aligned with the centers of the lenslets matched to the centers of the mirror actuators.

Prior to every set of extended runs (typically once a day), the woofer and tweeter are re-registered to the wavefront sensor and a new woofer influence matrix is generated. Woofer edge effects are minimized by performing a fit to individually measured influence functions. A singular value decomposition (SVD) reconstructor zeroes out potentially problematic (high frequency) woofer modes. Each DM has a separately set integrator leak (up to 0.01 for the woofer and 0.005 for the tweeter), similar to ShaneAO, to allow each mirror to handle

its own actuator clipping and windup. Detailed telemetry (residuals, DM commands) and configuration is saved during each run to ensure reproducibility.

The test bench control law mimics the integrators and reconstructors used in ShaneAO, however the system delay can be arbitrarily shortened or lengthened to match the operating control delay measured with ShaneAO or to test the effects of longer or shorter time delays.

5.2.1 Modifications to the Experimental Setup to Support Predictive Fourier Control

In order to use the test bench as a scale representation of the ShaneAO system, and to integrate the LQG controller into the test bench software architecture, several changes to the test bench were required.

We first re-mapped the pupil of the test bench to the WFS lenslet arrays to simulate subapertures which are scaled equivalents to the 10 cm subapertures found in the ShaneAO system. The remapping resulted in 36 active subapertures across the pupil, slightly more than are found on the ShaneAO system, which will have 30 subapertures across the pupil in its highest performance configuration.

The test bench operates at approximately 1 Hz in the lab, in discrete steps. As such, it has essentially no servo-lag (both mirrors respond in a manner of a few ms), and we introduce an artificial delay, which applies an older correction to the mirrors each timestep. Consistent with both ShaneAO and the GPI AO system, we apply a 2-frame delay between wavefront sensor measurements and applying corrections. Throughout this chapter, we use test bench units, timesteps and subapertures, as our test bench does not operate in real time. All of these numbers can be easily scaled to match the 3 m ShaneAO system, an 8 m or a 30 m telescope. For the ShaneAO system, we convert from the test bench to an operating frequency of 1 kHz and a subaperture size of 10 cm. On the test bench, we use windspeeds between 0.01 and 0.38 subapertures/timestep, which correspond to physical windspeeds on ShaneAO between 1 m/s and 25 m/s.

As well, we implemented an SVD reconstructor for the system's high

stroke woofer. This allows us to eliminate problematic high frequency woofer modes, but to use the natural actuator basis set for the woofer directly. This also allows us to easily re-measure the woofer influence functions for each actuator, recalibrating the system for minor changes in alignment and optimizing the system's long term performance.

Although the SVD reconstructor for the woofer improved its performance and registration, the SLM registration algorithm had to be improved to fix actuator mis-registration errors that would cause very high spatial frequency content to build up on the mirror. We solved this with an improved actuator registration method which uses actuators across the full SLM mirror working surface, and automatically fits the pixel locations instead of relying on user input to identify actuator pokes.

The original fiber laser was replaced with a 589 nm super luminescent diode (SLD) to fix non-uniform pupil intensity effects that manifested as an 8 Hz power peak due to laser power fluctuations. Over long runs (10,000 iterations or greater), we observe a correlation between RMS error and measured Strehl indicating low non-common path errors. The far-field image is sharpened using an algorithm that iteratively maximizes Strehl by successively placing Zernike modes on the woofer and tweeter.

We also improved the control loop architecture of the test bench to facilitate the use of the Kalman reconstructor. The new architecture allows us to transition quickly between reconstructors, to perform calibration with one reconstructor (usually a simple Fourier-transform reconstructor) and then transition to a different reconstructor (the full LQG controller) for operation. These modifications also permit the system to enable, disable and reset integrator states during normal operation.

5.2.2 Error Budget for the Experimental Setup

We have developed a consistent error budget for the test bench in its current configuration, based on the error budget presented in Ammons et al. (2010). Similar to Ammons et al., we derive many of our error budget terms from the total atmospheric turbulence as measured by the wavefront sensor's

RMS total wavefront error measured with the phase plate held still ($\sigma_{\text{atmosphere}} = 180.0$ nm). From this measurement, we can derive the fitting error ($\sigma_{\text{fit}} = 0.0441\sigma_{\text{atmosphere}}$, see Ammons et al. (2010) section 4.3.2) and the wavefront sensor aliasing error ($\sigma_{\text{alias}} = 0.4\sigma_{\text{fit}}$, see Ammons et al. (2010) section 4.3.2 and references therein). The measurement error from the wavefront sensor was measured directly on the test bench, with no atmosphere present, and found to be 10.0 nm.

We explicitly separate bandwidth and delay temporal errors. As the test bench operates in discretized timesteps, waiting for each mirror movement to complete, there is no bandwidth or servo-lag error. The software does introduce an artificial timestep delay, which holds the computed correction for 2 timesteps to simulate the increased computation time. The error due to the artificial delay in the system was computed from equation 9.56 in Hardy (1998),

$$\sigma_{\text{delay}}^2 = 28.4 (f_g \tau_s)^{5/3} \quad (5.1)$$

where f_g is the Greenwood frequency, which can be computed from the wind velocity, and τ_s is the delay between measurement and correction, which for the experiments described in this chapter was set to $\tau_s = 2$ timesteps.

To measure the non-common path errors in our Strehl measurement, we applied the extended Maréchal approximation (see Hardy 1998 equation 4.40, and Maréchal 1947) to the far-field Strehl after image sharpening but with no aberrations in the system. After image sharpening, the far-field Strehl was 0.80, contributing $\sigma_{NCP} = 49.5$ nm of error to the system. Examining the error budget during experiments suggests that our far-field correction did not remain stable, and that an additional 44.0 nm of calibration error exist in the system. All of these error terms are presented in Table 5.1. The RMS far-field wavefront error in Table 5.1 includes terms which are not seen by the wavefront sensor. When considering the RMS residual wavefront error as measured by the wavefront sensor, we ignore the Mirror fitting, spatial aliasing, static uncorrectable and calibration error terms. When measuring the far-field Strehl and the RMS Far-Field Wavefront Error (computed via the Strehl and the extended Maréchal approximation), these terms are included.

Table 5.1: Error Budget for the test bench.

Error Term	Value
Wavefront Sensing	10.0 nm
Mirror Fitting	8.8 nm
Spatial Aliasing	3.5 nm
Time Delay	40.5 nm
Static Uncorrectable	49.5 nm
Calibration	44.0 nm
RMS Far-Field Wavefront Error	78.8 nm
Predicted Strehl Ratio (589; nm)	0.471

For the remainder of this chapter, when we use RMS residual wavefront error, we are referring to only those terms which are seen by the wavefront sensor, or to measurements taken directly from the wavefront sensor. In contrast, the RMS total wavefront error is the wavefront error measured through our simulated atmosphere but with no correction applied. When considering the wavefront error in the far field, we will always convert to the Strehl ratio using the extended Maréchal approximation.

5.3 Theoretical Framework

We build on a predictive Fourier control (PFC) framework for Adaptive Optics Systems (Poyneer et al., 2007; Poyneer & Véran, 2008) that aims to minimize temporal errors, including servo-lag and computation time, which manifest as a misestimation of the current atmospheric phase. This predictive controller uses a slow (~ 0.1 Hz) loop to identify frozen flow turbulence seen by the AO system, and an LQG controller in the Fourier domain to suppress the specific errors which arise from frozen flow turbulence across the telescope aperture.

The PFC algorithm uses telemetry from the adaptive optics system to identify frozen flow. Telemetry is recorded over the course of 1-10 seconds and then post-processed. The post-processing uses the Fourier Wind Identification scheme (Section 5.3.1, and Poyneer et al. (2009)), as well as a vibration and aliasing analysis to identify specific temporal frequencies which should be masked

and eliminated. This information is used to construct a state space model that captures the dynamics of the atmosphere. The relationship between the Predictive Fourier Controller and the Fourier Wind Identification is shown in Figure 5.2.

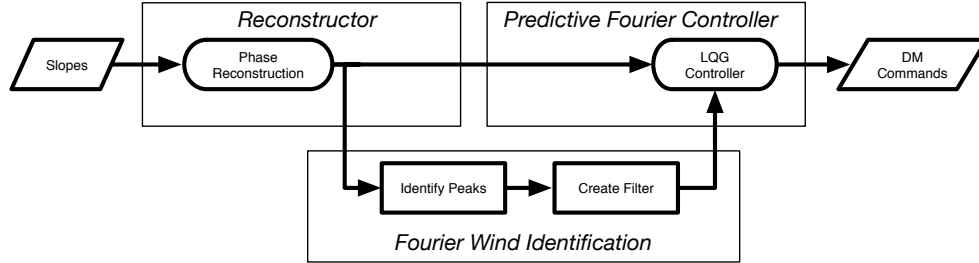


Figure 5.2: A block diagram showing the dual-loop arrangement used with Predictive Fourier Control. The usual AO loop runs at kHz frequencies, while a separate analysis loop measures the frozen flow velocities and generates a new Kalman filter for the system. Frozen flow turbulence is generally stable for more than 10s of seconds (Poyneer et al., 2009), so updating the filter on these timescales should not affect the quality of the predictive control. However, even if the wind velocity does change before the filter is updated, we have shown (Section 5.4.3) that the correction does not degrade significantly.

This section describes the framework for Predictive Fourier Control, first describing Fourier Wind Identification (Section 5.3.1), then describing the Predictive Fourier Controller (Section 5.3.2). The application of this framework to the test bench is discussed in Section 5.4.

5.3.1 Fourier Wind Identification

Using pseudo-open loop phase telemetry from the system, the Fourier wind identification (FWI) technique (Poyneer et al., 2007) identifies atmospheric frozen flow for correction. The Fourier basis set provides a convenient way to examine translating frozen-flow, as the individual Fourier modes each oscillate at a specific temporal frequency.

The temporal power spectral density (PSD) of each Fourier mode is estimated from the data. Individual Fourier modes are split into segments of length S , which are windowed to emphasize the middle of the segment, and

which overlap with neighboring segments. The overall frame rate, f_s sets the maximum estimated temporal frequency at $f_s/2$. The length of the segments sets the frequency sampling spacing at f_s/S . The half overlapped segments increase the signal-to-noise in the resulting PSD (for a more detailed discussion of this method, see Poyneer et al. (2007)). By increasing or decreasing S , we can control the signal-to-noise and frequency resolution of the PSDs.

Once we have created a PSD for each Fourier mode, we can look for temporal peaks which are indicative of frozen flow. In Fourier space, temporal peaks will appear with frequencies given by

$$f_t = f_x v_x + f_y v_y. \quad (5.2)$$

We model each peak as an oscillation at the temporal frequency f_t , with an added white noise broadening term (Poyneer & Véran, 2008). Three sample PSDs from the test bench, which show peaks that obey Equation (5.2), are shown in Figure 5.3. Each peak corresponds to a potential match to Equation. The peaks that appear close to 0 timesteps⁻¹ are eliminated as they correspond to the slowly varying steady-state errors found in every system. Peaks are fit for each spatial Fourier mode separately.

Using the identified peaks from all the Fourier modes, the FWI algorithm works backwards through Equation (5.2). The frequencies in Equation (5.2), when shown on an f_x, f_y grid, appear as a plane in frequency space, with the $f_x = 0, f_y = 0$ (piston) term always at 0 timesteps⁻¹. Figure 5.4 shows the process of matching found peaks in a PSD to a theoretical plane in Fourier space. FWI then produces a metric in velocity space that shows the percentage of matched peaks at each velocity. Areas with high metric scores are velocities at which frozen flow has been detected. The full metric is shown in Figure 5.5.

Once peaks have been identified, they can be formulated into a state-space model. Solving the Discrete Algebraic Riccati Equation (DARE) generates a controller that suppresses those specific temporal frequencies in its error transfer functions. The gains, leaks and high-pass coefficients of the LQG controller are then updated once the frozen flow layers and vibrations have been identified.




Figure 5.3: A temporal power spectral density (PSD) for three separate Fourier modes ($k = 3, 6, 9, l = 0$) which show clear evidence of frozen flow turbulence. The system ran using an LQG controller which did not model any atmospheric frozen flow. Peaks at $\sim -0.021, -0.042,$ and -0.062 timesteps $^{-1}$ are due to the effect of frozen flow turbulence caused by a phase screen mounted on a stepper motor passing in front of the entrance pupil to the test bench. The non-zero peak positions are set by Equation (5.2). The peaks at very high frequencies are spatial aliases of the frozen-flow wind.




Figure 5.4: Identification of frozen flow turbulence on the test bench using the Fourier Wind Identification algorithm. This metric shows a frozen flow layer that is moving in the x -direction ($v_x = 0.38$ and $v_y = 0.0$ subapertures/timestep). The color bar shows the percentage of theoretically expected peaks identified for a given wind vector. Higher percentage matches correspond to wind vectors which appear in more Fourier modes in the raw data.




Figure 5.5: Identification of frozen flow turbulence on the test bench using the Fourier Wind Identification algorithm. This match shows a frozen flow layer that is moving in the x -direction ($v_x = 0.38$ and $v_y = 0.0$ subapertures/timestep). The left panel shows the detected peaks in telemetry, as described in section 5.3.1. The color bar corresponds to the temporal frequency of the detected peaks, which satisfy Equation (5.2). Modes where no matching peak was found are shown as white. The right panel shows which peaks match the detected frozen flow layer. It is red for modes which were not detected in the telemetry, and white for modes where the detected peak would be too close to 0 timesteps⁻¹ to be detected. For this set of telemetry data, 97% of modes where peaks could be detected showed evidence of frozen flow.

5.3.2 LQG-Based Predictive Fourier Control

The wind parameters derived in the Fourier Wind Identification step are then used to calculate an optimal LQG controller for the AO system. The controller uses a Kalman state-space framework with an independent state for each frozen flow layer. A derivation of the appropriate Kalman filter for predictive control in a variable delay adaptive optics system is done analytically and numerically in Poyneer et al. (2007) and Poyneer & Véran (2008), and only a brief overview is provided below.

Frozen flow in the Fourier domain results in each Fourier mode shifting with a frequency given by Equation (5.2). To apply this shift over a fixed delay between the phase measurement on the WFS and the application of the phase to the DM, we use a complex valued gain (g_α in the derivations from (Poyneer et al., 2007)). The complex gain serves to translate each Fourier mode to account for the delay between WFS sensing and command application, usually somewhere between 1 and 2 timesteps. A purely real α corresponds to no frozen flow and a complex α determines the shift applied to each Fourier mode, with the complex phase¹ of α given by

$$\text{Phase}(\alpha) = -2\pi(kv_x + lv_y)\frac{T}{Nd} \quad (5.3)$$

where k and l are the respective Fourier mode numbers, $[v_x, v_y]$ (m/s) is the wind vector for a single frozen flow layer, T is the delay between the WFS measurement and the DM command application (s), N is the phase grid size ($N = 36$ for the test bench) and d is the subaperture diameter (m). The magnitude of α , $|\alpha|$, sets the “memory” for an individual layer, and is less than 1. Along with α , the noise covariance matrix of the frozen flow layers (including the static term) is used to solve the DARE, and yield an optimal controller with g_α tuned for each Fourier mode.

The LQG controller can be applied directly with the Kalman Filter ma-

¹This should not be confused with the “phase” of the wavefront.

trix equations. For our model, it can be simplified into the controller $C(z)$

$$C(z) = \left(Q^{-1} \sum_{k=0}^L \frac{p_{L+1,k} \alpha_k}{1 - \alpha_k z^{-1}} \right) / \left(1 - z^{-1} Q^{-1} \sum_{k=0}^L p_{L+1,k} \right), \quad (5.4)$$

which provides insight into how the prediction works. In Equation (5.4) $Q = p + \sigma_v^2$ is the total layer power, and σ_v^2 is the noise variance of the particular Fourier mode.

This filter produces transfer functions which show a “notch” at the particular layer frequency, as well as good rejection around 0 Hz, the DC term. The theoretical stability of these filters is examined in detail in Poyneer et al. (2007) and Poyneer & Véran (2008).

5.4 Experimental Results

5.4.1 Modeling Transfer Functions

With the improvements described in Section 5.2.1, we first demonstrate that the test bench can reproduce theoretical transfer functions across a wide range of spatial frequencies for our system. To measure the error transfer functions (ETFs), we recorded the wavefront sensor residual error measurements for 4096 iterations with the control loop active and closed, and with the control loop uncontrolled and open. For each telemetry data set, we converted to the Fourier basis, and then created periodograms using the technique described in Section 5.3.1. To compute the error transfer function, we took the ratio of the active, closed loop PSD for a given mode with the uncontrolled, loop open PSD. An ETF for a single Fourier mode ($k = 5, l = 5$) is shown in Figure 5.6 along with a model for that ETF.

Figure 5.6 shows the transfer function measured and computed for an LQG controller which does not model any frozen flow turbulence. The LQG controller was implemented in a “DC-only” mode (simulating 0 layers of atmospheric turbulence) which has similar performance characteristics to an integral controller used in many AO systems.

The model ETF was computed using Equation (5.4), but for a single,




Figure 5.6: An empirical (*green*) and model (*red*) error transfer function for the test bench. The empirical measurements were taken with an LQG controller which did not model any frozen-flow layers. The model was computed using Equation (5.4) for a 0-layer atmosphere (i.e. correcting for static errors only). The transfer function shows the correct roll-off slope at low frequency. The model also correctly predicts the controller bandwidth on the real system (the lowest frequency where the ETF first crosses 1), though the model underestimates the overshoot (the degree to which the system makes frequencies higher than the bandwidth worse).

static layer, and so the model behaves like a standard integral controller. The improvements to the test bench allowed us to match the roll-off (the slope of the ETF at low frequencies), a feature which is universal to all first-order integral controllers, as well as the gain (which sets the left-right position of the model in Figure 5.6) and bandwidth (the lowest frequency at which the ETF crosses 1.0) between the model and the measured transfer function.

The well matched ETFs show that the system is performing nearly as expected at all spatial frequencies, and that more advanced controllers should be stable, with a predictable effect on the ETF of the system.

5.4.2 Integrating the LQG Controller

In order to support the LQG controller on our test bench, we made a few modifications to the control flow of the test bench. The LQG controller

relies on an internal state vector and produces an estimate of the entire, open-loop phase at the next timestep. However, both the test bench, and the ShaneAO system use the phase error internally and apply the integrators only to the commands sent to the deformable mirrors.

In ShaneAO (and by default on the test bench) the temporal integration is done on the mirror commands directly so that each integrator can consider and provide feedback for actuator clipping internally, without affecting the reconstructor or the state of the whole system. This is especially important when using multiple deformable mirrors to allow the system to take advantage of the high-stroke of the woofer and the many actuators on the tweeter.

To accommodate this architecture with LQG, we inserted a temporal differentiator into the phase reconstructor. The differentiator is applied after the LQG controller is used to predict the phase at the next step. Since the commands to each deformable mirror are independently integrated, the differentiation at the end of the LQG controller and the integration on each of the mirror commands are designed to cancel each other out. This method (differentiating the open-loop correction, and applying to the woofer and tweeter separately) is used in GPI (Poyneer et al., 2014).

However, as the integrators for the woofer and tweeter DM have differing leak settings, and respond differently when actuators clip, the temporally differencing the output of the LQG controller is not guaranteed to produce a stable controller. When operating in this mode, the effective gain of the system is controlled by the coefficients of the LQG controller, while the individual command integrators are set to have a gain of 1, but with their usual leak terms to prevent actuator saturation and windup.

5.4.3 Demonstrating Stable LQG Control on the MCAO Test Bench

Once the LQG controller was integrated into the test bench, we were able to demonstrate simple stable LQG control. The LQG controller achieved comparable RMS residual wavefront error (103.0 nm RMS residual wavefront error from 180.0 nm RMS total wavefront error from the atmosphere) to the original integral controller over the course of 4,000 iterations. The far field Strehl

ratio remained comparably stable during this period. The controller is stable for a variable frame delay of 1, 2, or 3 frames (higher delays were not tested).

Stability was easily maintained even when a wind suppression model was applied with no wind translation in the system. The decrease in performance due to an incorrect frozen flow assumption was minimal, causing an RMS residual wavefront error of 50 nm.

5.4.4 Applying Fourier Wind Identification on the MCAO Test Bench

We measured the Error Transfer Functions of the test bench with a fast artificial wind speed (0.38 subapertures/timestep). The system was run for 2048 iterations in both closed loop, and loop open (no correction) mode. Using the closed loop telemetry data, we were able to see evidence of frozen flow turbulence at every measured spatial frequency. Figure 5.3 shows the temporal PSD for three Fourier modes ($k = 3, 6, 9$ $l = 0$); each shows clear evidence for frozen flow turbulence.

Given the indication of frozen flow turbulence in the PSDs from the test bench, we were able to construct a wind likelihood metric from test bench data showing that the method described in Section 5.3.1 can successfully retrieve the input frozen flow wind velocity from the system telemetry. Figure 5.4 shows the identification of the 0.38 subapertures/timestep wind in test bench telemetry.

The identified wind vector (or a-priori known wind vector(s)) can be provided to the filter generation algorithm described in Section 5.3.2 and derived in (Poyneer et al., 2007) and (Poyneer & Véran, 2008) which will place a notch in the error transfer function for each Fourier mode to remove the temporal peak due to the frozen flow layer. The ETF for $k = 6$, $l = 0$ will be notched near -0.042 timesteps⁻¹ and will reduce the large power peak visible in Figure 5.3 for that mode.

5.4.5 Applying Predictive Fourier Control on the MCAO Test Bench

To test Predictive Fourier Control, we generated a state-space model for a known windspeed, assuming that the wind impacted all Fourier modes.

The filter was generated for test bench wind speeds of 0.05, 0.125, and 0.38 subapertures/timestep, and assumed a two frame delay to match our system design (see Section 5.2). To generate the filter, we assumed perfect frozen flow wind was detected in all Fourier modes. This assumes that modes where the Fourier Wind Identification scheme (outlined in Section 5.3.1) did not detect peaks due to frozen flow turbulence are nevertheless worth including in the filter. For the test bench, this proved to be the correct assumption, but it should be tested on-sky as well.

We compared the Predictive Fourier Control Kalman filter to a standard integral controller in the Fourier Domain, with a gain of 0.38 and a leak of 0.99. The RMS residual wavefront error for both the Kalman Filter and the standard Integral Controller is shown in Figure 5.7. The RMS total wavefront error (with no correction) is 180.0 nm. The integral controller has a 103.0 nm RMS residual wavefront error. Predictive Fourier Control improves the RMS residual wavefront error by a factor of 3 to 50.2 nm.




Figure 5.7: The RMS residual wavefront error for both the Predictive Fourier Controller Kalman filter (*blue*) and a standard integral controller (*green*) over 4096 iterations. For the standard integral, the RMS residual wavefront error was 103.0 nm, and for the Predictive Fourier Controller, it was 50.2 nm. The large spikes which occur every 640 iterations occur when the phase plate reaches the end of its movement range and must be reset to its initial position. With no correction, the RMS total wavefront error is 180.0 nm.

Unfortunately, when we went to measure the Error Transfer Function of the test bench using Predictive Fourier Control, we found that the system had a bug that recorded telemetry in an internal state, rather than as the final result. This means that we were unable to identify the notches which should be created by the LQG filter.

We were not able to re run the experiment, as the hardware which runs the test bench has critically failed. A power outage to the laboratory destroyed the computer which manages the bench hardware, both wavefront sensor cameras, and one of the SLM drivers. Although we tried to replace some hardware, we've found that many of the bench interfaces need to be re-built from the ground up in order to return to working order, and that time would be better spent assembling a new bench, which is out of the scope of this work.

5.5 Discussion and Future Directions

This chapter demonstrates the feasibility of "Predictive Fourier Control", but leaves a true characterization of system performance improvements and stability to future work.

The LQG filter was able to correct for frozen flow turbulence using our Fourier mode state-space model. Correction was applied in some modes beyond the controller bandwidth, and improved performance at very high temporal frequencies.

Although we have performed cursory checks of the LQG filter stability using the test bench, several questions remain.

We plan to explore the stability of the LQG filter during real-world operation, where wind vectors can vary during the application of the filter. Although the filter appears stable for moderate changes in the wind vector, we will require an on-sky application to test the performance gains and losses that occur due to a naturally varying wind vector.

As the LQG filter maintains a state-space model of the system, it has some "burn-in" time during which it must build up a "memory" of the wavefront error. We will characterize, on-sky, the effect of this "burn-in" time during filter transitions, and use this to optimize the trade-off between the speed of the

Fourier Wind Identification loop (the bottom loop in Figure 5.2) and the penalty incurred by changing the LQG filter in real time.

The PFC algorithm can also be extended to eliminate wind aliasing features, which appear as similar planes in Fourier space to those described by frozen flow. In the same way that we can reduce delay error in the presence of frozen flow turbulence, we can use the LQG controller with the PFC algorithm to correct for the effects of wind aliasing when it is detected in system telemetry (Poyneer & Véran, 2010).

Finally, once we have demonstrated the PFC algorithm on-sky with ShaneAO, we intend to adapt the LQG controller to account for actuator clipping and the woofer-tweeter architecture. A full controller would penalize actuators which are close to clipping within the Kalman filtering framework, and so would provide an end-to-end optimal control that would be less reliant on the absolute calibration of the individual woofer and tweeter mirrors.

5.6 Conclusions

This chapter presents experimental verification of a computationally efficient method for removing the effects of frozen-flow turbulence from the residual errors in an adaptive optics system. Predictive Fourier Control has the potential to reduce the effect of time delay errors, and to correct for specifically modeled and known errors beyond the bandwidth of the adaptive optics system.

We demonstrate that in the presence of wind-blown frozen flow turbulence Predictive Fourier Control is a feasible method for gaining 10% improvement in an AO system in the far-field performance and a factor of 3 in the RMS residual wavefront error. The laboratory experiments suggest that we will see significant gains in AO system performance on-sky. This level of performance improvement is promising for an algorithm which only requires a software adjustment, and which will not add significantly to the computational complexity of an existing adaptive optics system.

Chapter 6

Linear Quadratic Gaussian Control with ShaneAO

6.1 Introduction

ShaneAO is a high-order AO system commissioned in 2014 for the Shane 3-m telescope at Lick observatory (Gavel et al., 2014). It is a facility instrument designed to perform across a wide wavelength range ($0.8\mu\text{m}$ r-band through $2.2\mu\text{m}$ K'). The AO system feeds a Near-IR imager (McGurk et al., 2014) which includes standard and narrow band NIR filters as well as two low-resolution ($R \sim 200$) grisms. ShaneAO is commissioned to operate both in natural guide star mode (NGS) and with a sodium laser beacon (LGS).

ShaneAO was designed both as a facility instrument, in regular use at Lick Observatory, and as a pathfinder for a new generation of AO systems. The major features and design of the system is described in § 6.1.1, and we start by verifying the design and performance of ShaneAO (§ 6.2) For this work, we decided to branch significantly from the facility architecture for ShaneAO, in order to enable alternative reconstructors (§ 6.3).

6.1.1 Overview of ShaneAO

ShaneAO is mounted on the Cassegrain Focus of the 3-meter Shane Telescope at Lick Observatory. It provides AO correction over a $20''$ field of

view. In natural guide star mode, ShaneAO is designed to operate down to a limiting magnitude of $m_V = 13$, and can reliably acquire natural guide stars within $10''$ of the science target.

ShaneAO corrects for atmospheric turbulence using a woofer-tweeter deformable mirror pair. The tweeter is a Boston Micromachines MEMS device with 32×32 actuators, a $1.5\mu\text{m}$ stroke, and a 30 kHz bandwidth. The woofer is a 52 actuator AlpAO continuous phase sheet mirror with a $\pm 50\mu\text{m}$ stroke and a 200 Hz bandwidth. There is no separate tip-tilt mirror in ShaneAO, as the large stroke of the woofer can be used to control for the residual tip, tilt and focus in the system. Each of these modes is offloaded to the telescope control structure every 2 s.

The wavefront sensing is done with a Shack-Hartmann wavefront sensor (WFS). The WFS has a selectable lenslet array which can provide either 8, 16 or 32 subapertures across the pupil of the telescope, corresponding approximately to $2r_0$, r_0 and $\frac{1}{2}r_0$ sampling at $1.5\mu\text{m}$. For this experiment, unless otherwise stated, the 16-across lenslet array was used. The WFS is backed by a SciMeasure BigJoe Camera controlling a Lincoln Labs CCID66 180×180 pixel array. The WFS is programed to sample at a range of operating rates from 50 Hz to 1.5 kHz, but for this experiment, only 250 Hz was used.¹ The WFS is fed by light split from the science path. Additional light can be split from the WFS to feed the Tip/Tilt sensor. In LGS mode, only the 589 nm light from the sodium beacon is fed to the WFS subsystem. However, in NGS mode, a beam splitter can be inserted to send some of the NGS starlight to the Tip/Tilt sensor. We left this beam splitter in place during our experiments so that the tip-tilt sensor could serve as a high-speed diagnostic camera.²

The Tip-Tilt sensor is a Marconi CCD39 80×80 pixel array driven by a SciMeasure LittleJoe camera controller. It is used primarily in LGS mode. Along with the Tip-Tilt Sensor, there is a laser uplink controller to correct for laser tip

¹Experimentation and characterization described in this chapter lead us to believe that the 500 Hz and 1 kHz facility modes for ShaneAO do not properly perform AO correction due to either mis-calibration or incorrectly set default parameters. As we could not achieve the desired error transfer function (ETF)s which correspond to the the expected performance of the system at these operating rates, we limited experiments to 250 Hz.

²The science camera is limited to 0.1 s exposures and takes 2 – 3 seconds to read an image to disk and display it, making it suboptimal for real time performance feedback.

and tilt during LGS operation.

The ShaneAO real time controller is a linux PC running CentOS 6. The reconstruction software is written in C, and relies heavily on the OpenBLAS³ framework. There is no realtime linux or kernel level code in ShaneAO.⁴

ShaneAO is integrated into the Lick Observatory Environment, and so uses the Keck Telescope Library (KTL) system of keywords and values to pass information about system state between the controlling computer and the observatory environment. Graphical user interfaces for operating the system in facility mode are provided, and run on a host separate from the ShaneAO real time control computer.

6.1.2 Reconstruction Scheme

ShaneAO uses a modal reconstruction scheme using the natural modes of the tweeter and woofer mirrors, implemented around a single vector-matrix multiply. The natural mirror modes were measured using an interferometer to measure the poke positions of each actuator, and decomposing the poke matrix for each mirror into modes in approximate order of spatial frequency. The reconstruction matrix is generated from a poke matrix of mirror modes, pseudo-inverted to suppress local waffle, piston and additional invisible modes, and combined to output commands to both the tweeter and woofer with a temporal filter between the two.

The vector matrix multiply is fed by a flexible centroider which uses a 5×5 pixel box for each subaperture. One row and column is used as guard space between subapertures. The remaining pixels can be binned down to a standard quadcell, or used with a center-of-gravity or centered quadcell, by changing the pixel weighting used to compute slopes (Gavel et al., 2016).

The matrix is a compact representation of the tweeter, woofer and temporal filter matrices (Figure 6.1). Combining the three reconstruction matrices, and applying the temporal filter within the matrix, means that the bulk of the computation time in the reconstructor is contained to a single linear-algebra call.

³<https://openblas.net>

⁴See § 6.3.4.2 for an analysis of the frequency and impact of OS interruptions on the real time performance of the system.

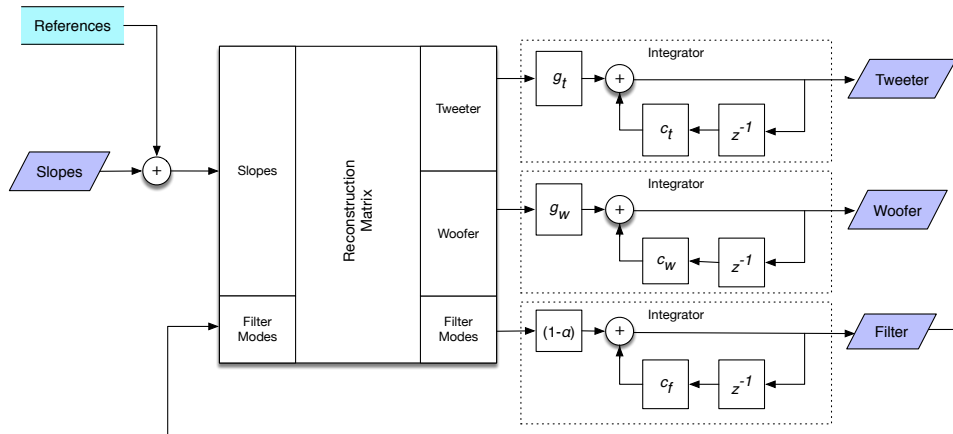


Figure 6.1: Diagram of the ShaneAO Vector Matrix Multiply reconstruction scheme, in NGS mode. Purple boxes are input and output data from the core reconstructor. The white boxes are the reconstruction steps. The reconstructor accepts inputs from the centroider, and subtracts reference values from the measured slopes. The matrix multiply uses an inverted modal matrix for the tweeter and woofer to independently reconstruct actuator positions for each mirror. It produces filter coefficients, which are just the 14 controlled woofer modal coefficients that are also applied to the woofer actuators. Finally, the matrix multiply subtracts the temporally filtered woofer modes from the tweeter position. Each of the output channels (Tweeter, Woofer and Filter) are integrated and clipped separately. The integrated tweeter and woofer actuator positions are then sent as output to the mirrors. The integrated filter coefficients are returned as inputs to the matrix multiply on the next timestep.

Each output from the vector matrix multiply is run through an individual integrator, with a separate gain, bleed, and clipping limit applied to all outputs. ShaneAO in facility mode is commonly operated with different bleeds on each the Tweeter, Woofer, and Filter.

6.1.2.1 Modal Calibration

To produce the reconstruction matrix for ShaneAO, we use a modal poke matrix for each mirror. The mirror modes were measured using an interferometer to measure the poke positions of each actuator, and decomposing the poke matrix for each mirror into modes in approximate order of spatial frequency. Due to the rectangular grid of actuators on the tweeter, the tweeter

modes (matrix \mathbf{A}_t) are roughly linear combinations of Fourier modes, with some evidence of the influence function of each actuator at higher spatial frequencies. The Woofer modes \mathbf{A}_w are an orthonormal basis set which resembles Zernike modes, but which cannot be directly mapped to Zernikes due to the irregular position of the actuators on the back of the Woofer’s continuous phase sheet.

The modal poke matrix is generated by producing a moderate amount of each mirror mode, one at a time, and measuring the WFS response. This matrix (which we label \mathbf{H}_t for the tweeter, \mathbf{H}_w for the woofer) is used as the input to the inversion problem for reconstruction.

Using this measured modal calibration has the advantage that the reconstruction matrix is a product only of modes which each mirror can produce and which the WFS can measure. Proper weighting and penalization during inversion removes local waffle effects (Gavel, 2003), as well as eliminating modes which are blind to the WFS (e.g. high spatial frequency modes when the system is using the 8 or 16 across lenslet arrays).

For the tweeter reconstruction, the penalized pseudoinverse takes the form

$$\mathbf{H}_t^\dagger = (\mathbf{H}_t \mathbf{P} \mathbf{H}_t^T + \mathbf{Q})^{-1} \mathbf{H}_t^T \mathbf{P} \quad (6.1)$$

where \mathbf{Q} is the a-priori covariance of the modal coefficients, and \mathbf{P} is the covariance of the measurement noise. \mathbf{P} is usually constructed without any specific subaperture weighting, but we have the ability to include a weighted least-squares measure to account for partially illuminated subapertures. This has some impact in 16-across mode, but it is more relevant in 32-across mode, an operational mode not explored in this work. For ShaneAO, we use a regularization penalty matrix to penalize waffle and piston modes. The penalty matrix is implemented in terms of k penalized (“bad”) actuator patterns a_{bad}

$$\mathbf{Q} = \mathbf{A}_t \left(\sum_k a_{bad,k} w_k a_{bad,k} \right) \mathbf{A}_t^T \quad (6.2)$$

To check the modal gain of this matrix, we can examine the diagonal terms of the matrix $\mathbf{H}_t^\dagger \mathbf{H}_t$. We expect that this will not produce a modal gain of




Figure 6.2: The modal gain of the penalized inverse modal matrix for ShaneAO. The gain is lower than 1 for all modes due to the penalization of waffle and piston modes, the penalization of actuators outside of the illuminated apertures, as well as the down-weighting of under-illuminated subapertures. The ShaneAO tweeter modes are arranged in roughly spatial frequency order.

1 due to the penalization of specific modes which are not independent tweeter modes. The modal response is shown in Figure 6.2.

For the woofer reconstruction, we limit the \mathbf{H}_w matrix to the lowest 14 modes. For these modes, we use a regularized pseudoinverse to construct the woofer reconstruction matrix \mathbf{H}_w^\dagger .

6.1.2.2 Woofer-Tweeter Control Split

Along with the woofer and tweeter reconstruction, there is the temporal filter which subtracts the woofer position from the tweeter. The temporal filter is reconstructed with the forward matrix for the woofer, \mathbf{H}_w , so that only the 14 controlled woofer modes are filtered from the tweeter. To subtract these modes from the tweeter, we construct a modal crosstalk matrix which converts

from Woofer modes to Tweeter modes

$$\mathbf{C}_{wt} = \mathbf{H}_t^\dagger \mathbf{H}_w \quad (6.3)$$

This crosstalk is applied to the filtered woofer modes, which have been through an integrator with a gain parameter α and a bleed c_f (see Figure 6.1). The discrete time implementation of the woofer filter is

$$\begin{aligned} \mathbf{v}_k &= c_f \mathbf{v}_{k-1} - (1 - \alpha) \mathbf{H}_w^\dagger \mathbf{s}_k - (1 - \alpha) \mathbf{v}_{k-1} \\ &= \mathbf{v}_{k-1} (c_f + \alpha - 1) - (1 - \alpha) \mathbf{H}_w^\dagger \mathbf{s}_k \end{aligned} \quad (6.4)$$

These filtered modes are then subtracted from the tweeter. This is done by re-projecting the woofer modes to tweeter modes with \mathbf{C}_{wt} from Equation 6.3, and then converting those tweeter modes to actuator commands with the modal projection matrix \mathbf{A}_t .

6.1.2.3 Matrix Layout

In order to combine the two inputs (slopes and filter coefficients) and three outputs (tweeter, woofer and filter coefficients), the reconstruction matrix is stacked and provided as a single linear algebra operation. The update equation is

$$\begin{pmatrix} \Delta \mathbf{a}_t \\ \Delta \mathbf{a}_w \\ \Delta \mathbf{v} \end{pmatrix}_k = \begin{pmatrix} -\mathbf{A}_t \mathbf{H}_t^\dagger & -\mathbf{A}_t \mathbf{H}_t^\dagger \mathbf{H}_w \\ -\mathbf{A}_w \mathbf{H}_w^\dagger & \mathbf{0} \\ -\mathbf{H}_w^\dagger & -\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{s}_k \\ \mathbf{v}_{k-1} \end{pmatrix} \quad (6.5)$$

We label the control matrix part of this \mathbf{R}_c . The c subscript denotes matrices which include this 2-column, 3-row structure to provide the woofer-tweeter filter and split.

This is assembled into the full reconstructor which obeys the following state equation for providing the actuator positions at time k using the state at

time $k - 1$

$$\begin{pmatrix} \mathbf{a}_t \\ \mathbf{a}_w \\ \mathbf{v} \end{pmatrix}_k = \begin{pmatrix} c_t \\ c_w \\ c_f \end{pmatrix} \begin{pmatrix} \mathbf{a}_t \\ \mathbf{a}_w \\ \mathbf{v} \end{pmatrix}_{k-1} + \begin{pmatrix} g_t \\ g_w \\ 1 - \alpha \end{pmatrix} \begin{pmatrix} \Delta \mathbf{a}_t \\ \Delta \mathbf{a}_w \\ \Delta \mathbf{v} \end{pmatrix}_k \quad (6.6)$$

This equation is implemented in two steps in the reconstructor. The first step contains the matrix multiply described by Equation 6.5. The second step is the integrator described by Equation 6.6. The block matrix in Equation 6.5 is stored as a single file on disk, and re-computed for each operating mode for ShaneAO. The gains for the second term in Equation 6.5 (g_t , g_w and α) are implemented by multiplying the appropriate rows of the block matrix.

6.1.3 Software Architecture

The original ShaneAO software (called “unreal”) uses a hybrid infrastructure written in python and c, built to be robust as a facility instrument. The core hardware interfaces and reconstructor are written in a c-extension to python. It uses OpenBLAS to perform the vector matrix multiply in parallel. The python code provides the native interface to the system as a python prompt. The cameras run in independent threads in the main process, and provide an interrupt-driven interface to the reconstructor. This means that reconstruction only happens when the camera notifies the system that an image is ready for processing. The tweeter actuator positions are updated when the reconstructor runs. The woofer is controlled in a separate execution thread, and limited to the response rate of the woofer hardware (200 Hz). The reconstructor is a simple c function, which operates on pre-allocated global arrays of data for data processing. This is robust when working with a single reconstructor, but it is difficult to switch between reconstructors, as the reconstructor state is held at a global level. The system interface is designed primarily as a python prompt, with individual python commands manipulating the state of the system.

Configuration occurs in a python module, which stores the locations of various configuration files (e.g. the control matrix) for the different operational modes of ShaneAO.

6.2 Verifying the ShaneAO Reconstructor

Implementing predictive Fourier control (PFC) requires a detailed understanding of the performance and characteristics of any adaptive optics (AO) system. ShaneAO was built without verification tests, and without a detailed documentation of performance and timing of the completed system. I undertook work to verify that the performance and characteristics of the AO system matched those intended by the system design.

ShaneAO has three operating configurations which were used for verification:

1. A full stack software simulator which tests the software against simulated mirrors, wavefront sensor (WFS), and atmosphere.
2. An internal light source which can be inserted into the instrument beam after the first fold mirror (called M1 in ShaneAO documentation) and used to operate the system with the real hardware, but with no atmosphere simulation.
3. A telescope simulator, which provides a simulated atmosphere, but which requires the physically removing M1 to mount the simulator, and so is only used when the system is not attached to the telescope.
4. On-sky, with the system used as it is when performing science observations.

To the extent possible, I performed verification tests on each configuration, in the order presented above, moving from software verification, to the internal light source, and finally, to stars on sky, progressing from the least complex instrument set up to the most complex and realistic instrument set up.

There are several factors about each setup which influence the realism and applicability of tests done at that stage, and a few considerations which ensured that each stage would be applicable to the next one.

When working with the simulator (1 above), the system runs at about 1 Hz, and behaves slightly differently than it would on-sky due to these unusual

timing characteristics. As well, the simulator has different alignment properties, and so requires a control matrix derived from the modal response of the mirrors *in the simulator* rather than on sky. As such, there is not necessarily a direct correlation between a given on-sky configuration and that configuration in the simulator. In particular, it is difficult to probe the temporal split between the woofer and tweeter in the simulator (again, because of the 1 Hz issue), but the simulator is still useful as a test of the observing procedures and general stability of the system. As well, it is possible to match the controller gain and leak to values that would be used during normal operation. The simulator also contains a realistic atmosphere model which can be tuned to match the atmospheric power at Mt. Hamilton.

The internal light source operates at the expected frequency (consistently chosen to be 250 Hz for this experiment) as the on sky system, and communicates with hardware which is used on sky. As well, the system is in its final alignment configuration, and so we can use control matrices derived from the modal response of the system mirrors, identical to those we would use for a natural guide star. Unfortunately, there is no realistic atmosphere simulation with the internal light source, so we rely on the photon shot noise from the light source itself to provide the necessary noise inputs. This means that we can easily be limited by the read noise of the system. As well, reference centroids taken with the internal light source don't have aberrations, and so the WFS spots are especially tight (Gavel et al., 2016). On sky, these spots are broadened, which reduces the response of the centroiding algorithms. This discrepancy means that reference centroids must be converted between the internal light source values and on-Sky values.

The telescope simulator can be installed in the system when it is not mounted to the telescope.⁵ The telescope simulator provides a correctly resampled input beam, with a rotating phase plate. The rotating phase plate has an adjustable speed, and a turbulent spectrum applied with acrylic paint (Rampy et al., 2010). We use the telescope simulator when we require a simulated atmo-

⁵The telescope simulator requires that we remove the first fold mirror, and mount a small optical bench with two OAPs. The optical bench extends off the top of the instrument, and would physically collide with the telescope when it is mounted.

Table 6.1: Parameter definitions and default values for the ShaneAO reconstructor. The column *Default* is the parameter value used in experiments described in this chapter. The column *Facility* is the value used by facility instrument operators at Lick Observatory.

Parameter	Symbol	Default	Facility	Notes
Tweeter Gain	g_t	0.3	0.3 – 0.7	See § 6.2.5.1.
Tweeter Memory	c_t	0.99	0.9	
Woofers Gain	g_w	0.3	0.3 – 0.7	Often controlled with g_t .
Woofers Memory	c_w	0.99	0.99	
Woofers Filter	α	0.2	0.2	
Filter Memory	c_f	0.99	0.99	

sphere to verify the performance of reconstructors.

When testing on-Sky, we select magnitude 8 or brighter natural guide stars to ensure that we are not photon-noise limited.⁶ Natural guide stars are chosen to be near zenith, to eliminate effects due to gravitational flexure in the telescope.

Nominal operating parameters for the original ShaneAO reconstructor are shown in Table 6.1. The *Facility* column shows the values for these parameters that are used when ShaneAO is operated as a facility instrument. For this work, we chose a consistent set of operating parameters, in the *Default* column, for our tests. These parameters were shown to produce good performance for ShaneAO before any modifications were done.

6.2.1 Self-Consistent Error Transfer Model

In order to understand the temporal performance of ShaneAO, we developed a transfer function model which incorporates the response of both the tweeter and woofer mirrors in the system. We assume that both mirrors have a hold-and-stare characteristic, i.e. they move to their target position and then are held fixed. For a simple integrating controller, the transfer function (Madec, 2004) is

⁶This results in at least 4000 counts per subaperture on the WFS for an $m_V = 8$ guide star, above the background of ~ 1500 counts. Most data was taken with $m_V \geq 7$, resulting a comparably brighter set of spots on the WFS, with between 6000 and 10000 counts per subaperture.

$$\frac{E(z)}{\Phi(z)} = \frac{1}{1 + dC(z)} \quad (6.7)$$

where d is the delay, $C(z)$ is

$$C(z) = \frac{g}{1 - cz^{-1}} \quad (6.8)$$

We define the delay in terms of two stare-and-hold times (one for the WFS image, and one for the mirror position) and a computational delay

$$d = \left(\frac{1 - z^{-1}}{\omega} \right)^2 \exp(-\tau s) \quad (6.9)$$

Although this control law holds well for a single-mirror AO system, ShaneAO requires additional treatment to handle the split of commands between the woofer and tweeter, since the woofer has a slower temporal response than the tweeter. The woofer's temporal bandwidth is at about 200 Hz, which is not a standard operating rate for ShaneAO. To ensure stability, there is a temporal filter applied in the control loop which serves to subtract the response of the woofer from the commands sent to the tweeter (Gavel et al., 2014).

To model this dual response, we use two mirrors, operating at different rates, but combined into a single error transfer function (ETF):

$$\frac{E(z)}{\Phi(z)} = \frac{1}{1 + d_w r_w C_w(z) + d_t r_t C_t(z) f_t} \quad (6.10)$$

The delays, d_w and d_t , for the woofer and tweeter respectively, apply like the delay term in Equation 6.9, but for the woofer the stare-and-hold delay is matched to the update rate of the woofer (200 Hz) instead of the update rate for the WFS.⁷ τ_c is the computational delay time, which applies to both the tweeter and the woofer. We introduce an additional delay parameter, τ_w , which represents the amount of time it takes the commands to reach the woofer, since

⁷We do not model the arbitrary sampling and aliasing of the woofer commands explicitly. When operating at 250 Hz, the woofer will receive four out of every five commands, causing a slight aliasing effect, which we have elected not to model.

the woofer is controlled on a separate machine and requires network traffic to function. We assume that $\tau_t = 0$, as the tweeter commands are immediately output to the DPIO2 board in the control loop and the DPIO2 card is designed to run at tens-of-kHz speeds.

The mirror responses are modeled as $r_t = 1$ and r_w is a single-pole low pass filter, with a cutoff frequency $\omega_{c,w}$ which represents the hardware response time of the ALPAO deformable mirror. The low pass filter is implemented as

$$r_w = \frac{1}{1 - s/\omega_{c,w}} \quad (6.11)$$

and $\omega_{c,w}$ is assumed to be 20 Hz, which is consistent with the performance of the system when using only the woofer.

The control law for each mirror is implemented separately according to Equation 6.8, with appropriate changes to reflect the differing hold time for the woofer (1/200 ms irrespective of the WFS rate.) The last term in the denominator of Equation 6.10 is the single-step temporal filter which is applied to the woofer modes, reprojected on the tweeter, in the vector matrix multiply. In z-transform space, this is modeled as

$$f_t = \frac{1 - c_f z^{-1}}{1 - (c_f + \alpha - 1)z^{-1}} \quad (6.12)$$

where c_f is the bleed coefficient for the filter terms from equation 6.4.

The full control law model for the default parameters of ShaneAO is shown in Figure 6.3. The full control law can be fit to measured ETFs, however, several parameters can quickly become degenerate (e.g. the computation delay, τ and the woofer delay τ_w when fitting the woofer only). As such, we hold values constant for which we have verified the performance inside the reconstructor.

6.2.2 ShaneAO Simulator

The ShaneAO simulator, called *sauce* (Gavel et al., 2014; McGurk et al., 2014; Gavel et al., 2016), can reproduce most of the behavior of the full AO system. It intercepts commands sent to both the woofer and tweeter mirrors, and




Figure 6.3: Theoretical ETFs for the VMM reconstructor, showing the transfer function for the woofer and the transfer function for the tweeter, which combine to form a single transfer function.

provides images to the WFS via an alternative driver for the WFS camera.

The simulation itself models the influence functions of each mirror, along with a simulated atmosphere, on a *fine*-sampled grid of 512 pixels on a side. The combined residual phase from the tweeter, woofer, and atmosphere is fed to a simulation of the WFS, which uses a full Fourier optics propagation to simulate the optical propagation through the lenslet array and on to the CCD. The simulator can add read noise, background, and photon noise. As such, the simulator serves as a suitable verification that the basic aspects of the system are correctly implemented. Unfortunately, the simulator operates at about 1 Hz, and so it is impractical to perform long-running verification tests. However, it is possible to use the simulator to demonstrate that the reconstructor is stable, and that the software performs as expected.

As well, the simulator runs with an interrupt every time it is done computing the next residual phase, and with no ability to model the temporal response of the woofer. As such, it essentially includes no delay for the woofer, and the woofer and tweeter temporal response functions appear identical.

The first stage of this process was to verify the reconstructor calculations. The reconstructor provides telemetry dumps which contain both the input slopes and the output actuator values. My initial verification process ensured that these input slopes, multiplied by the ShaneAO reconstruction matrix, then run through an integrator which matches the ShaneAO control law, produce identical actuator commands.

Along with ensuring that the calculation was occurring correctly in the simulator, I ensured that the simulator produced a stable, closed loop at a moderate gain. This was done by setting the loop to a moderate gain (~ 0.3) and leak (~ 0.9), and ensuring that the loop would converge on a stable solution for a given simulated atmosphere. An example of the estimated phase (\hat{p}) telemetry typical during convergence is shown in Figure 6.4.

6.2.3 Performance on the Internal Light Source

ShaneAO includes an internal, fiber-fed calibration source, which is primarily used to determine non-common path errors during image sharpening. I used this light source to verify the stability of the ShaneAO reconstructor. Although there is no phase screen in use for the internal light source, the ratio of commands applied to the noise in the AO system takes the same form as the ratio of controlled to uncontrolled pieces of the atmosphere, and so produce an equivalent error transfer function.

To verify the performance of ShaneAO, I recorded telemetry data with the loops open (effectively, gain = 0), and then with the loops closed at a moderate gain (~ 0.3 for most tests). The phase error time series for both of these sets of telemetry is shown in Figure 6.5 and Figure 6.6. The histogram in Figure 6.7 shows that the loop was closing across all points on the \hat{p} grid, as all of the closed loop values are clustered around 0.

For each telemetry data set, we converted the \hat{p} to the Fourier basis, and then created periodogram (PSD)s using the technique described in Section 6.2.1. Figure 6.8 shows the PSDs of the \hat{p} for a single grid point. To compute the ETF, we took the ratio of the active, closed loop PSD for a given mode with the uncontrolled, loop open PSD (Figure 6.9). We computed a theoretical ETF modeled

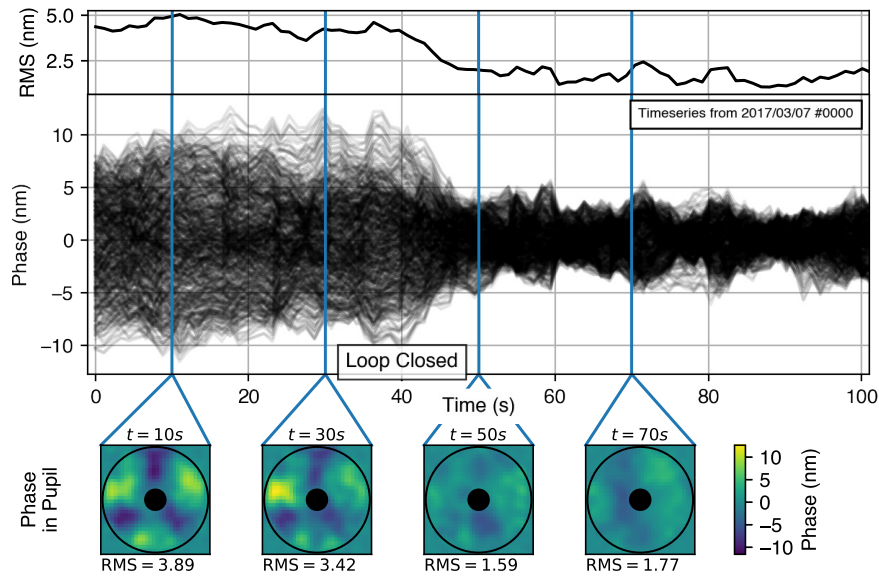


Figure 6.4: The top row shows a residual \hat{p} timeseries while the loop is closing and converging on a stable value. Each line is a separate illuminated \hat{p} point in the ShaneAO grid, where the loop was closed at around the 40 s time mark, after which the residuals converge on zero. The bottom row shows postage stamp images of the \hat{p} values at the marked blue line time points. The first two postage stamps are from before the loop has closed, and the second two are after the loop closes, where it is evident that the residual phase is flattened.

on a basic integral controller (see Section 6.2.1).

The shape of the measured ETF matches our theoretical model quite well, though it suggests that the system is slightly underperforming the nominal input values. The gap between the theory and data at low frequencies (below ~ 5 Hz) is best explained by the noise characteristics of the internal light source. Simulations demonstrate that this mismatch can occur when the data is limited by the read noise of the system. This is consistent with the fact that the internal light source on ShaneAO is quite faint to facilitate its use with the Near-InfraRed imaging camera ShARCs.

More detail on the fitting process and the match between nominal and fit parameters is discussed in Section 6.2.1.




Figure 6.5: Time Series of \hat{p} values from the ILS for phase grid position (10, 10), showing that in open loop there is significant phase residual relative to the reference point, but that the residuals center around 0 when the loop is closed.




Figure 6.6: Selected values of the \hat{p} values in the ILS, showing that in open loop there is significant phase residual relative to the reference point, but that the residuals center around 0 when the loop is closed. Both sets of residuals show significant residual error behind the secondary mirror, but this is an artifact of the reconstruction and not controlled in the system.




Figure 6.7: Histogram of median \hat{p} values during loop closed and loop open operation on ShaneAO with the ILS. The clustering of loop closed values suggests that the system correctly converged on the reference points.

6.2.4 On Sky Performance

The final performance verification for ShaneAO was done “On Sky”, using relatively bright natural guide stars (between $5 \leq V_{\text{mag}} \leq 8$, observed near zenith to minimize the atmospheric path length). The ShaneAO WFS has sufficient dynamic range when used in the 16-across mode to measure telemetry with both the loop open and the loop closed, a feature we used to compute error transfer functions for ShaneAO on sky. The residuals from both the loop open and loop closed telemetry are shown in Figure 6.10.

Again, for each telemetry data set, we converted to the Fourier basis, and then created PSDs using the technique described in Section 6.2.1. The peri-dograms are shown in Figure 6.11. To compute the ETF, we took the ratio of the active, closed loop PSD for a given mode with the uncontrolled, loop open PSD. We computed a theoretical ETF modeled on a basic integral controller (see Section 6.2.1), and fit the theoretical model to the data.

The on-sky ETFs (Figure 6.12) show similar features to the internal light source ETFs, but suggest that the system has significantly different per-




Figure 6.8: PSD of \hat{p} at position (12,12) in ShaneAO, recorded with data from the ILS. The open loop data (orange) is relatively flat, as there is no atmospheric aberration in the system. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.

formance on sky relative to using the internal light source.

6.2.5 Additional Verification

Along with on sky transfer functions, I verified additional input parameters in the ShaneAO reconstructor.

6.2.5.1 Gain and Slope Regularization

The ShaneAO centroider uses a regularization parameter to limit the effect of low-intensity WFS measurements. The centroiding algorithm, in the quad-cell case, can be reduced to the following equation in the x direction (with an analogous equation in the y direction):

$$S_x = \frac{I_{\text{left}} - I_{\text{right}}}{(I_{\text{left}} + I_{\text{right}} + R)} \quad (6.13)$$




Figure 6.9: The ETF from the ShaneAO internal light source. This ETF was computed using the \hat{p} at the 12,12 grid point in ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.19 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data.

Table 6.2: Model parameters for the model ETFs shown in Figure 6.9

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	33.0	0.2	0.99
Fit (<i>red</i>)	1.18	0.11	0.997	1.18	0.16	0.861	45.9	0.2	0.99

where I_{left} and I_{right} are the intensities in the left and right halves of the quad-cell respectively. The regularization parameter should not only prevent zero-division errors, but should also serve to damp out the slope measurements when there is a low signal-to-noise ratio on the WFS. Hardy (1998) suggests that R should be approximately 10% of the quad-cell intensity of a limiting magnitude guide star.

For ShaneAO, this parameter defaults to $R = 500$. This value is close to the usual intensity values observed on the WFS for moderate (2000 counts for $R_{\text{mag}} \sim 7$) natural guide stars, which means that the effective gain of the centroiding algorithm is artificially reduced by the regularization parameter.




Figure 6.10: Time series of the median residuals recorded with the loop open (orange) and the loop closed (blue) on sky, showing that in open loop there is significant phase residual relative to the reference point, but that the residuals center around 0 when the loop is closed.




Figure 6.11: PSD of \hat{p} at position (12,12) in ShaneAO, recorded with data from a $V_{\text{mag}} \sim 7$ guide star on sky. Both the open loop and closed loop data show a power-law roll off expected from atmospheric turbulence. The closed loop shows signs of significantly reduced power below 4 Hz




Figure 6.12: The error transfer function (ETF) from the ShaneAO on sky. This ETF was computed using the \hat{p} at the 12,12 grid point in ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.13 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data. Table 6.3 shows the model values for the plotted ETFs.

Table 6.3: Model parameters for the model ETFs shown in Figure 6.12

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	33.0	0.2	0.99
Fit (<i>red</i>)	1.47	0.04	0.009	1.57	0.06	0.977	61.5	0.2	0.99

6.2.5.2 Reconstructor Timing Analysis

To use predictive control on ShaneAO, I need a detailed understanding of the timing involved in the ShaneAO system. Predictive control predicts the open-loop phase at some time in the future, and predicting across the wrong time delay would impact the stability of the predictive controller. By injecting a known signal onto the ShaneAO mirrors, I measured the round-trip delay for the ShaneAO system. The delay was measured using the internal light source on the system, with a modified reconstructor which facilitated signal injection.

In the reconstructor, we injected a known-period sine wave of a single tweeter mirror mode (mode 10). The sine wave was injected with a period of




Figure 6.13: The input sinusoid signal (blue, applied to the tweeter) and the output WFS signal (green) which shows a phase shift due to the round trip delay of the system. The y-axis shows the modal amplitude, and the x axis shows the timestep for the recorded data. The delay was measured with the system running at 250 Hz.

50 reconstructor frames. This sine wave signal was then recorded on the WFS. Telemetry of a single actuator and a single slope is shown in Figure 6.13. These two sine waves show a phase shift due to the delay between the signal as applied to the mirror, and the signal as measured by the WFS. Using a Fourier Transform, we can measure the phase shift between the WFS and deformable mirror signals at the injected frequency.

Using three different periods for the input signal, we measured the delay in the whole system to be $8900 \mu s$, which is 2.3 frames at 250 Hz. We can account for 1.3 frames of delay. The 1 frame delay occurs while the camera is integrating, and we know from timing measurements that the reconstructor computations require about 0.3 frames (see Figure 6.20). The remaining frame delay we attribute to the readout time of the camera, as we have documentation suggesting that the WFS camera is configured to spend a single integration time reading pixels off of the image.

6.3 Restructuring the ShaneAO Reconstructor

The ShaneAO reconstructor was designed to be a compact, efficient, single vector-matrix multiply reconstruction. This leads to a robust reconstructor which is fast enough on modern commercial hardware without requiring a real-time operating system or custom mathematical hardware. In order to enable predictive control on ShaneAO, we broke the ShaneAO reconstructor into its constituent pieces, exposing an estimate of the uncorrected phase which could be used for the linear-quadratic gaussian predictor. This process was done in two steps: first, we re-architected the ShaneAO control software to be modular and testable, allowing us to insert and replace individual pieces of the ShaneAO algorithm.

6.3.1 Motivation

In order to facilitate switching algorithms in ShaneAO, we wanted to abstract away the common components so that they could be written, tested, and verified once, and then used for multiple reconstruction algorithms. Components like the centroider are common to all reconstruction algorithms which use data from the Shack-Hartmann wavefront sensor (WFS) in ShaneAO, and changing the centroider is independent of changing other aspects of the reconstruction algorithm.

We developed a system where individual components of the reconstruction algorithm can be configured and tested against known, artificial inputs and outputs to verify that they are functioning as intended. Our scheme relies on modular unit tests to test the individual algorithm components, and a framework for re-using components in a common algorithm. This ensures that we cannot introduce bugs in components like the centroider while switching the core reconstructor from a single vector-matrix multiply to a double vector-matrix multiply.

6.3.2 New Reconstructor Framework

The new framework breaks out each colored block in Figure 6.14 into a separate, testable component, with its own configuration, input data files, and unit tests to ensure that the implementation is correct.

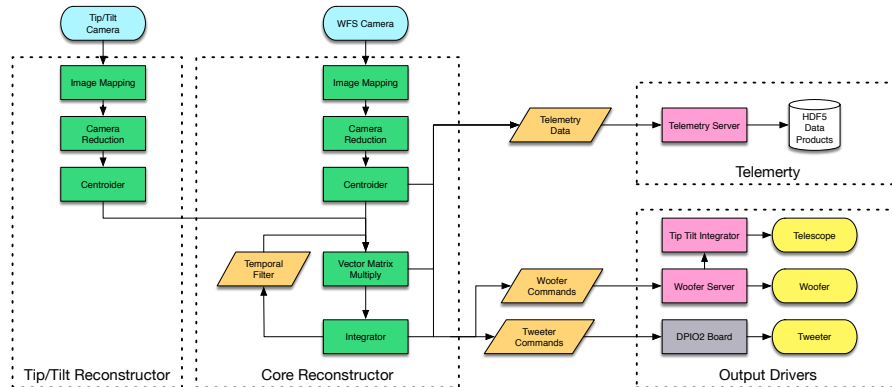


Figure 6.14: Block diagram showing the computational steps involved in the vector matrix multiply reconstructor for ShaneAO. Green blocks represent reusable computational components. Orange blocks are data storage. Pink blocks are the long-running server processes which manage interaction between the reconstructor and hardware, and which were tuned to be re-usable for differing reconstructor configurations. The grey block is the DPIO2 driver, which is a library embedded into the real time control software. The dashed boxes represent thread boundaries. In the original implementation, all of these blocks were combined into a single c function. The improved implementation splits out each step so that it can be independently tested and re-used.

Each block in Figure 6.14 becomes a paired c-structure for use in the real-time code, and a python wrapper which allows us to pass arbitrary inputs and retrieve outputs from the computational component either in the real system, or during unit tests for each component. The restructured reconstructor re-uses the output drivers (pink boxes in Figure 6.14) from the original code, but abstracts each additional component into a separate c function. Reconstruction algorithms are then implemented by connecting the input and output arrays from each of these components to each other in a camera handler object.

The WFS can be supplied with as many camera handler objects as desired. This allows the restructured ShaneAO to run two reconstructors simultaneously, so long as only one of them is connected to the actuator outputs

for the system. As well, it allows us to switch between reconstructor handlers quickly on sky, without a significant amount of overhead time for restarting or reconfiguring software.

6.3.2.1 Reusable Components

The re-usable, independently tested reconstructor components are

- An image mapping step which un-scrambles data read from the cameras. When used with a reconstructor, it only unscrambles illuminated sub-apertures, but for graphical interfaces a separate handler is used which unscrambles the full image.
- A camera reduction step which can subtract background counts and flat field the image based on known values for each camera's background and flat field.
- A centroider, which uses the full 5x5 subaperture pixels and accepts a flexible weighting function to control the effective layout of the subapertures. The centroider also provides the reference slope subtraction step.
- A generic vector matrix multiply step, which includes an external input facility for incorporating tip-tilt measurements with the full WFS slopes, and a loopback step to implement the single-step temporal filter used to subtract the woofer response from the tweeter. The vector matrix multiply also implements the loop gain using a row-multiplier on the raw matrix values.
- A generic integrator, which implements the actuator default values, actuator clipping and integrator leaks.
- A driver interface for the tweeter, used to send commands to the DPIO2 board.
- A TCP client used to send commands to the woofer processing thread.
- The telemetry writing server, for saving data about the performance of the reconstructor.

- A camera moving average, performed last, after all time critical components are done, to provide a graphical user interface with images from the cameras.

Each component is independently tested and verified against known good data, edge cases, and a re-implementation of the algorithm in python. Components that can be re-used, such as the image mapping code, camera reduction processing, and the moving average component are re-used between the WFS and tip tilt sensor.

ShaneAO is designed for the reconstructor to run in a camera-interrupt model, where the main timing loop is provided by the WFS camera, which wakes up the processing thread whenever an image is ready to perform reconstruction. To ensure that the timing is as reliable as possible, the new framework moves some processing steps, such as the camera moving average, to occur after commands have been sent to the mirrors during time spent waiting for the next image.

6.3.2.2 Application Architecture

When restructuring the reconstructor, we split the ShaneAO application into well-defined layers of responsibility:

- Graphical user interfaces and command line interfaces for controlling the adaptive optics (AO) system while it is running.
- An Application Layer to handle configuration and supervision, in python.
- A Python Middleware layer, which has components that represent individual pieces of the AO system, e.g. the tweeter, the WFS, and the reconstructor core.
- Low level c and cython⁸ code for performing real-time computations and interfacing with hardware drivers.
- Hardware drivers provided by vendors.

⁸Cython is a hybrid language which converts python-like code to compiled c code with a transparent interface layer.

This general structure is shown in Figure 6.15.

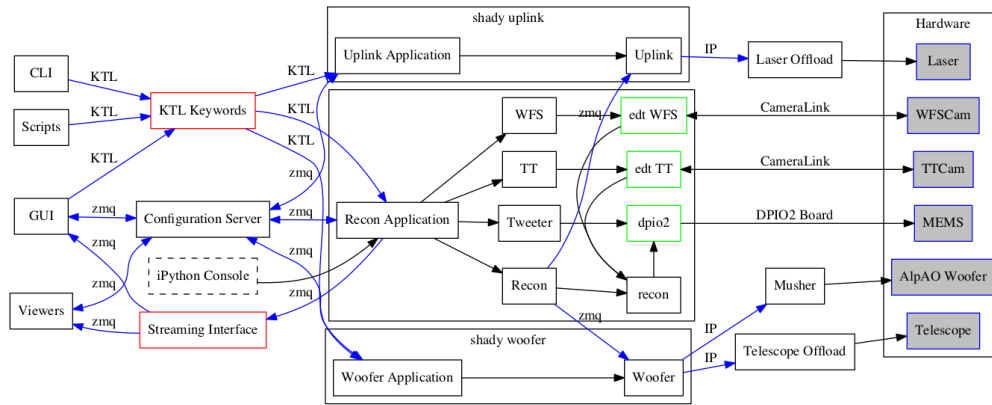


Figure 6.15: The software architecture of ShadyAO, the new ShaneAO reconstructor. On the left are the four types of user interface programs, Command Line Interfaces (CLIs), observatory scripts (e.g. dithering image pattern scripts), GUIs and Viewers, a subset of GUIs which are used to show data streamed from the AO system in near real time, such as WFS images. The next layer shows the three interfaces used to communicate between the user interface programs, KTL keywords, described in Section 6.3.2.4, a configuration server, described in Section 6.3.2.3, the engineering IPython console, discussed in this section, and the array streaming interface, described in Section 6.3.3. The central boxes in the figure represent core processes that power ShadyAO. At the top is the tip-tilt uplink server which provides the control loop for uplink tip-tilt to the laser. The box in the middle is the core reconstructor, which talks to most of the system hardware. The woofer is controlled by a separate process, shown on the bottom as shady woofer.

6.3.2.3 Repeatable Configuration

Along with the rearchitected reconstructor framework, I placed all of the relevant settings for ShaneAO into a single configuration file. This file can be read and written by the system to store most of the system state between nights, and to switch between different instrument setups. Along with the configuration file, some values are stored externally in FITS files. External values are reserved for large arrays, such as the WFS reference centroids and the mirror actuator defaults. These values, along with the configuration file, are recorded in a version control system so that the previous state of the system can be retrieved for any given night and set up.

Two tools were developed to assist in this configuration setup: one, `config.py`, which inspects a given configuration file, and can produce a list of dependent files as well as their MD5 sums⁹. The second tool downloads configuration files from each host which operates ShaneAO and stores them in a version control system, and which can then be used to restore the configuration files to all of the hosts that control ShaneAO.

To ensure that the configuration values are consistent across the different hosts which run various pieces of ShaneAO software, a central configuration server on the main computer, `real.ucolick.org`, hosts a configuration server which provides configuration values to all pieces of ShaneAO via a network request-reply pattern.

6.3.2.4 Observatory Software Interface

The new reconstructors are heavily integrated with the observatory interprocess communication system, called Keck Task Library (Lupton & Conrad, 1993) (KTL) which is used to control all aspects of the reconstructor software which are adjusted during normal operation. KTL works by exposing key-value pairs from the reconstructor, which command line and graphical clients can read, change, and listen for updates on. These key-value pairs form the primary way for additional observatory software, such as the science camera controller, and the telescope and guiding system, to interact with the AO system. As well, by exposing adaptive-optics specific details to the KTL keyword system, KTL provides a limited, well defined API over which users can interact with the reconstructor, limiting the potential for breaking changes while the system is operating.

Deeply integrating the KTL required creating a mock library which simulates the key-value interface used by KTL for use during unit and integration tests, when the AO software is running independently of the observatory environment. To do this, I wrote a relatively complete implementation of the key-value store interface in Python, called “Cauldron”¹⁰, which provides a

⁹MD5 is a hashing algorithm that produces a unique 128-bit value for a given file structure, and so is useful to differentiate between files with differing contents.

¹⁰Open source, available at <https://github.com/aLexrud/Cauldron>

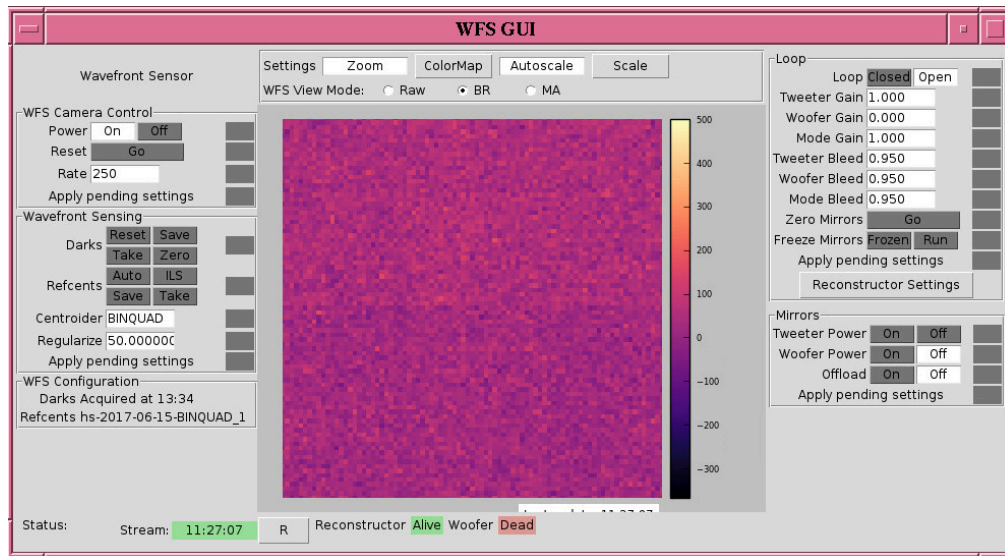


Figure 6.16: WFS GUI for the new reconstructor. The left hand panel controls the wavefront sensor camera, including basic camera data reduction. The central panel shows a live view of the wavefront sensor camera. The right panel shows the controls for managing the state of the AO loop, and the state of the mirrors. On the bottom, the status of the various background tasks is shown.

KTL-compatible interface to the AO system. It is used by the unit testing framework in an isolated mode, and by the ShaneAO simulator, “Sauce” to simulate the full observatory environment.

“Cauldron” contains three key-value backends: a local isolated backend for testing, a networked, zeromq-powered backend for use with the instrument simulator, and an observatory backend which passes values through to KTL. Cauldron has 387 unit tests which exercise the KTL application programming interface (API) and ensure that the local and networked backend provide an equivalent API to the observatory environment, facilitating an efficient testing and debugging loop without impacting observatory resources.

On top of the observatory interface, I built new GUIs which rely on the underlying KTL API. The GUIs incorporate feedback in the form of live camera and mirror views streamed from the reconstructor, and allow the user (and engineer, in my case) to control the AO system. Figure 6.16 shows the WFS GUI. Figure 6.17 shows the mirror control GUI. Figure 6.18 shows the engineering WFS reconstructor settings.

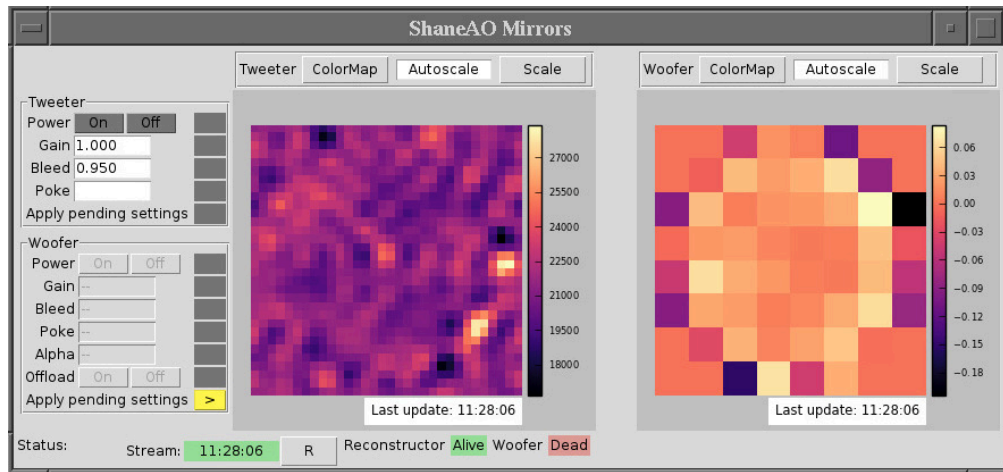


Figure 6.17: Mirror GUI for the new reconstructor. The left hand panel controls the mirrors. The middle and right panels show the actuator positions for the woofer and tweeter.

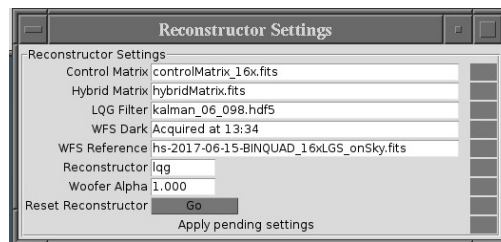


Figure 6.18: Reconstructor settings GUI for the new software system.

6.3.2.5 Sauce Simulator

ShaneAO was designed with an included simulator (called “Sauce”) which provides an optical path simulation of the deformable mirrors, atmosphere, lenslets and WFS camera, as well as the tip-tilt camera. Sauce is described in more detail in Gavel et al. (2014); McGurk et al. (2014); Gavel et al. (2016), and was largely un-modified for the new reconstructor framework.

6.3.2.6 Simulated Hardware Interfaces

One adjustment to the Sauce simulator in the new reconstructor was to provide simulated hardware drivers instead of directly accessing reconstructor products. In the original configuration, the simulator inserted data into the reconstructor’s internal memory, and extracted mirror positions from the recon-

structor memory. This had the dual disadvantage of requiring a switch inside the reconstructor which would skip sections of the reconstructor during simulation mode. By providing a simulated driver interface, we limited the scope of bugs which could be introduced in a new reconstructor which successfully runs in the simulator. The simulated drivers provide a c-api which is compatible with the real hardware drivers, but communicate with the simulator instead of physical devices.

6.3.3 Streaming Telemetry to Disk

In order to properly measure the performance of ShaneAO under multiple reconstructors, I developed a flexible telemetry recording system which streams telemetry data from the reconstructor to the disk using HDF5 files. The structured nature of HDF5 files makes it easy to add and remove telemetry arrays as they become relevant for a particular reconstructor. As well, the streaming system allows me to record an arbitrary length of telemetry data (subject to the available disk space on the recording computer). The telemetry recorder is implemented as an independently tested and verified library (“Zeeko”¹¹) which uses zeromq to stream data over the network and provides a python interface to control the high-speed Cython code used to complete high-bandwidth I/O.

The same library is used to provide the graphical user interface showing real-time mirror commands and WFS data to end users. In this case, streaming data is throttled so that it doesn’t overwhelm the GUIs or the observatory network.

6.3.4 New Reconstructor Verification

Once the new reconstructor was implemented, we verified its performance in five stages (similar to the process used in Section 6.2):

1. Unit tests of individual components (Section 6.3.4.1)
2. Functional tests of multiple components interacting (Section 6.3.4.1)
3. Verification on the ShaneAO simulator (Section 6.3.4.3)

¹¹Open source, available at <https://github.com/aLexrudy/Zeeko>

4. Performance with the internal light source (Section 6.3.4.3)
5. Performance on sky (Section 6.3.4.5)

6.3.4.1 Unit Tests

There are 580 unit tests which cover most aspects of the new reconstructor in single step mode. Each unit test creates an isolated environment in a temporary directory. Tests then construct a particular component of the reconstructor, and test that component against random inputs, known good inputs, and bad inputs, checking for the expected results in each case. The unit tests are parameterized across the various operational configurations for ShaneAO.

6.3.4.2 Performance and Timing

To ensure that the new reconstruction system was operating with acceptable performance, we incorporated timing information for the reconstructor in the telemetry stream. The reconstructor timing includes both a roundtrip reconstruction time (interval between two successive outputs to the tweeter) and the computation time (interval between camera read and tweeter output). Figure 6.19 shows that the round trip time is centered around 1.05 frames. The inter-frame times should be close to 1 frame, indicating that the reconstructor is keeping pace with the arrival of new WFS images.

ShadyAO, like ShaneAO, is compiled c code, but it is not written in real-time linux. There is also no thread affinity¹² applied in ShadyAO. Since the reconstructors heavily rely on both OpenBLAS¹³ and FFTW¹⁴, which are well developed parallel libraries, thread affinity would only serve to defeat the parallelism already present in the system. These choices leave ShadyAO exposed to operating system interrupts into the process flow for reconstruction. Fortunately, empirically, these interrupts have only a limited effect on the total computation time for the reconstructor. Figure 6.20 shows a histogram of computation times, where the effect of operating system interrupts is visible in

¹²Thread affinity allocates a particular physical processor core to a particular computational thread

¹³<https://openblas.net>

¹⁴<https://fftw.org>




Figure 6.19: Histogram of inter-frame arrival times for the ShadyAO reconstructor. The y-axis is log scaled, but shows the probability of events for each arrival time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile interval in arrival times is 1.255 frames, and the median inter-arrival time is 1.053 frames. There is a known 1 frame delay due to the readout settings on the WFS camera. The WFS exposures are handled in parallel, so although inter-arrival times are around 1 frame, the time between exposure start on the WFS and commands sent to the mirrors is 2.053 frames.

the tail of computation times. Although the median computation time is 0.0568 frames, computation times range up to an entire frame. This level of variance is an acceptable penalty for the ease of use programming in a standard linux environment.

6.3.4.3 Simulator Verification

To verify the timing and performance of the new reconstructor, we recorded telemetry using the Sauce simulator. The resulting telemetry was processed as in Section 6.2.3. Because the simulator operates at an effective rate close to 2 Hz, we only recorded two 1024-timestep segments. For consistency with the rest of this chapter, when analyzing telemetry from the Sauce simulator we assume a sampling rate of 250 Hz. Figure 6.21 shows the periodograms (PSDs) from the simulator. The slope of the PSD is driven by the atmosphere




Figure 6.20: Histogram of reconstructor computation times for the ShadyAO reconstructor. The y-axis is log scaled, but shows the probability of events for each computation time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile computational time is 0.2444 frames, and the median computation time is 0.0568 frames. Only 0.1% of computations take longer than half a frame.

simulation in Sauce.

The error transfer function shows performance that is consistent with the chosen gain (Figure 6.22). The gain $g = 0.25$ was selected, the fit suggests that the gain was $g = 0.24$. The error transfer function (ETF) does suggest that the system bleed was lower than requested ($c = 0.9$, but the fit found 0.84). This might be attributable to the atmospheric simulation in Sauce, which does not contain any net static errors.

The ETF and PSDs suggest that the new reconstructor performs as designed in the simulator environment.

6.3.4.4 Internal Light Source Verification

The new Vector Matrix Multiply (VMM) performs as designed on the internal light source. Figure 6.23 shows the PSDs from the internal light source (ILS) for the new reconstructor. A clear range of controlled frequencies is visible. The ETF associated with these PSDs is shown in Figure 6.24. The gen-




Figure 6.21: PSD of estimated phase (\hat{p}), averaged across the aperture in ShaneAO, recorded with data from the Sauce simulator. The open loop data is in orange. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.

eral characteristics of the ETF suggest that the new reconstructor is performing as designed, albeit with an additional delay of around 1.2 Frames rather than the predicted 1.1 Frames, though this additional computation time may be due to the changed reconstructor design. The ETF shows a broadened overshoot region between 5 Hz and 50 Hz which does not show the expected rounded shape. We find this extended overshoot region.

6.3.4.5 On Sky Tests

We verified the new Vector Matrix Multiply by recording images and telemetry on ShaneAO while looking at a $m_V = 7$ star near zenith. Figure 6.25 shows the PSDs, which show the expected power-law due to atmospheric turbulence, and the effect of the controller on the closed loop data. The reconstructor produces good error transfer functions on Sky, which match the predicted performance from our model ETF described in Section 6.2.1. The model ETF accurately detects the input gains and bleeds in the system. As well, the ETF re-




Figure 6.22: The error transfer function (ETF) of phase (\hat{p}) from Sauce with the new VMM. This ETF was computed by averaging the periodograms across illuminated phase points. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.25 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data. Table 6.4 shows the model values for the plotted ETFs. Because this data is recorded on the simulator, the woofer has a significantly different bandwidth from the real system. Our model detects this, and increases the woofer’s bandwidth from 10 Hz to 35 Hz.

Table 6.4: Model parameters for the model ETFs shown in Figure 6.22

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	10.0	0.2	0.99
Fit (<i>red</i>)	1.26	0.1	0.999	1.2	0.15	0.933	35.6	0.2	0.999




Figure 6.23: Periodogram estimated phase (\hat{p}), averaged over the illuminated aperture, recorded with data from the ILS. The open loop data (orange) is relatively flat, as there is no atmospheric aberration in the system. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.

produces the measured computational delay in the new VMM (see Figure 6.20). The ETF suggests that in the new reconstructor, there is an additional $\tau_w = 1.85$ frame delay communicating with the woofer, and that the woofer responds with a 33 Hz cutoff frequency.




Figure 6.24: The ETF from the ShaneAO internal light source. This ETF was computed using the pseudophase at the 12,12 grid point in ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.19 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data.

Table 6.5: Model parameters for the model ETFs shown in Figure 6.24

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	10.0	0.2	0.99
Fit (<i>red</i>)	1.4	0.04	0.999	1.3	0.07	0.974	35.1	0.2	0.999




Figure 6.25: Periodogram estimated phase (\hat{p}), averaged over the illuminated aperture, recorded with data on sky with a $m_V = 7$ star near zenith. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.




Figure 6.26: The ETF on sky using the new VMM reconstructor. This ETF was computed using the average \hat{p} for ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.07 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data.

Table 6.6: Model parameters for the model ETFs shown in Figure 6.26

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	33.0	0.2	0.99
Fit (<i>red</i>)	1.3	0.05	0.94	1.85	0.09	0.989	33.2	0.2	0.99

6.4 Modifying the ShaneAO Control Scheme for predictive Fourier control

Linear quadratic gaussian (LQG) controllers produce the “best estimate” of system state. In an AO system, this is usually set up as the best estimate of the state of the atmosphere, i.e. the open loop state of the system. This means that in order to do predictive control using an LQG controller in ShaneAO, we need to input the best estimate of the atmospheric state, and expect that wavefront correctors (in the case of ShaneAO, the woofer and tweeter) will produce the desired correction. Since ShaneAO uses a single modal control matrix (see § 6.1.2.3), and this matrix outputs split commands for the woofer and tweeter, we cannot apply LQG control directly to the existing control matrix structure. Instead of using this single control matrix, we split the matrix into two components. The first component does a “classical” reconstruction, solving the inversion problem to convert between slopes and phase. The second matrix splits the the phase between the tweeter and woofer mirrors, using the temporal filter strategy described in Section 6.1.2.3. As the two matrices are directly related, produce two pairs of control matrices, one which produces phase as the intermediate product and one which produces the complex Fourier coefficients of the phase.

6.4.1 Reconstructing Phase

To generate an estimate of the phase in the system at any given time, we used an inverse Fourier transform reconstructor (Poyneer et al., 2002). Using the Fourier transform of the natural modes defined on the tweeter mirror (see § 6.1.2.1), we applied the forward filters for the Shack Hartmann WFS, and then used the inverse Fourier transform to get the slopes back. These slopes are then compared to the slopes measured in the modal poke matrix \mathbf{H}_t (see § 6.1.2.1). From this, we determined that the modal definition matrix \mathbf{A}_t includes the influence function of individual MEMS actuators. As well, the response from each mode in the Fourier domain, compared to the expected response from the \mathbf{H}_t matrix, revealed that there is a 1% magnification between the WFS and the

tweeter.

The forward matrix that applies this 1% magnification and the influence function of individual actuators is \mathbf{D} . We also generate an inverse matrix which removes the physical effects of the MEMS is \mathbf{D}^{-1} . To reconstruct phase from measured slopes in the WFS, we combine the modal definition matrix (\mathbf{A}_t) and the actuator influence matrix (\mathbf{D}) to produce a phase reconstruction matrix

$$\mathbf{P}_t = \mathbf{D}^{-1}\mathbf{A}_t \quad (6.14)$$

We can also invert this matrix, via a regularized pseudoinverse, to produce a matrix which converts between phase and modal coefficients (\mathbf{P}^\dagger).

This pair of matrices provides the building blocks to implement the hybrid reconstructor. We apply the matrix \mathbf{P}^\dagger to the reconstruction matrix in (6.5)

$$\Delta\mathbf{p}_k = \mathbf{P}_t\mathbf{H}_t^\dagger\mathbf{s}_k \quad (6.15)$$

$$\mathbf{L} = \mathbf{P}_t\mathbf{H}_t^\dagger \quad (6.16)$$

where \mathbf{L} is the new control matrix we feed to the first vector matrix multiply in the split reconstructor.

We use this estimate of the phase, $\Delta\mathbf{p}_k$, as either an input into the LQG controller (see Section 6.7) or as input to the second part of the hybrid matrix reconstructor.

6.4.2 Splitting phase onto Tweeter and Woofer

Once we have reconstructed an estimate of the phase from WFS measurements, we must still reconstruct mirror commands for both the woofer and the tweeter. This process is a modification of the reconstruction in described in Section 6.1.2.2 and Section 6.1.2.3, where we remove the WFS to mode matrix \mathbf{H}_t^\dagger , and insert the phase-to-mode matrix \mathbf{P}_t^\dagger . For the woofer modes, we reconstruct woofer modes from phase using the forward influence matrix \mathbf{H}_t and the

phase to mode matrix \mathbf{P}_t . The full control matrix becomes

$$\begin{pmatrix} \Delta \mathbf{a}_t \\ \Delta \mathbf{a}_w \\ \Delta \mathbf{v} \end{pmatrix}_k = \begin{pmatrix} -\mathbf{A}_t \mathbf{P}_t^\dagger & -\mathbf{A}_t \mathbf{H}_t^\dagger \mathbf{H}_w \\ -\mathbf{A}_w \mathbf{H}_w^\dagger \mathbf{H}_t \mathbf{P}_t^\dagger & \mathbf{0} \\ -\mathbf{H}_w^\dagger \mathbf{H}_t \mathbf{P}_t^\dagger & -\mathbf{I} \end{pmatrix} \begin{pmatrix} \Delta \mathbf{p}_k \\ \mathbf{v}_{k-1} \end{pmatrix} \quad (6.17)$$

We label the matrix for this component \mathbf{P}_c

We apply additional transformations around the phase $\Delta \mathbf{p}$ to provide internal estimates of additional quantities. By pre-multiplying the matrix in Equation 6.15 as well as the first half of the hybrid matrix in Equation 6.15, we convert the internal quantity from phase to some other value. For example, using the matrix formulation of the discrete Fourier transform, we can provide the values between equations 6.15 and 6.15 in terms of Fourier modes.

6.4.3 Verifying the Combined Matrix

To confirm that the split matrix operates as designed, we computed a few simple tests of the hybrid matrix. First, we verified that the phase matrix \mathbf{P}_t and inverse \mathbf{P}_t^\dagger produce the identity matrix

$$\mathbf{I} = \mathbf{P}_t \mathbf{P}_t^\dagger \quad (6.18)$$

We also checked that the combined matrix $\mathbf{P}_c \mathbf{L}$ produces a modal response on the tweeter identical to that produced by the original control matrix. Figure 6.27 shows the modal response of the hybrid and standard control matrix.

6.5 Hybrid Matrix Reconstructor

The hybrid matrix reconstructor implements the matrix multiplies described in Section 6.4, to demonstrate that the woofer-tweeter split works alongside an estimate of the phase. This reconstructor is also the first test case for switching reconstructors in the ShaneAO software, as described in Section 6.3.2.

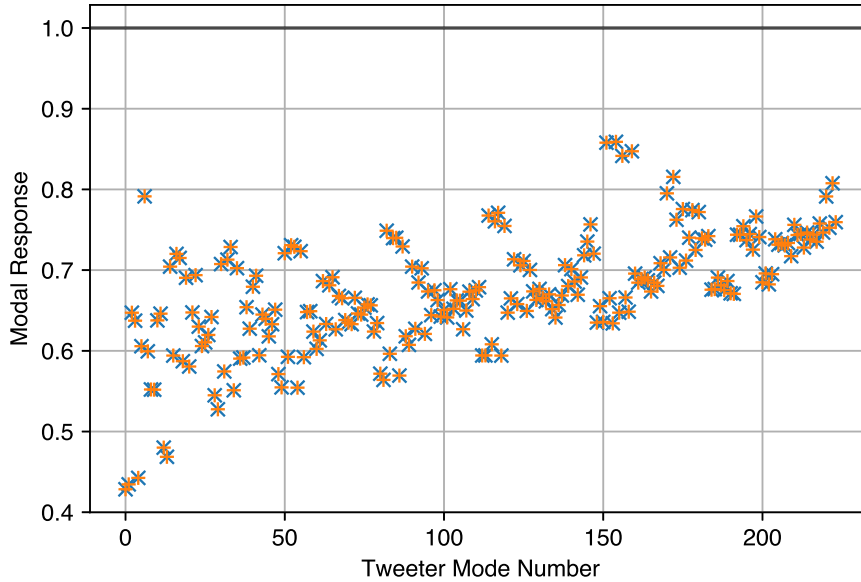


Figure 6.27: The modal gain of the penalized inverse modal matrix for ShaneAO. The blue crosses show the modal response of the standard matrix (from Figure 6.2). The orange crosses show the modal response of the hybrid matrix. The modal response of both matrices is almost identical. The gain is lower than 1 for all modes due to the penalization of waffle and piston modes, the penalization of actuators outside of the illuminated apertures, as well as the down-weighting of under-illuminated subapertures. The ShaneAO tweeter modes are arranged in roughly spatial frequency order.

6.5.1 Hybrid Matrix Design

The hybrid matrix reconstructor applies two vector-matrix multiplies in direct succession, splitting the matrix multiply described in Section 6.1.2.3 into the two parts described in Section 6.4, \mathbf{L} and \mathbf{P}_c . The hybrid reconstructor also saves the internal values between the two matrix multiplies, which gives us an estimated phase measurement from the system in real time.

Figure 6.28 shows the data flow within this reconstructor, where the two matrix multiplies are identified by the Reconstruction matrix (\mathbf{L}) and the control matrix (\mathbf{P}_c). This design allows us to insert additional reconstruction steps between the two matrix multiplies, where they will have access to an estimate of the phase error ($\Delta\hat{p}$).

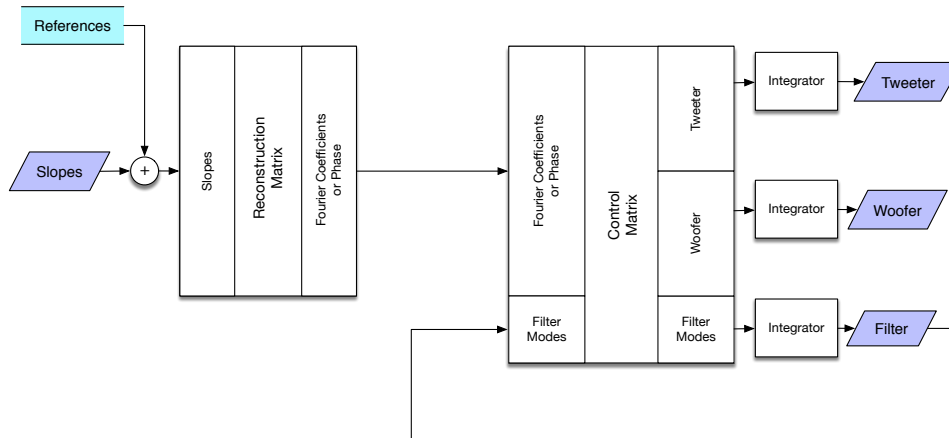


Figure 6.28: Diagram of the ShaneAO Hybrid Vector Matrix Multiply reconstruction scheme, in NGS mode. Purple boxes are input and output data from the core reconstructor. The white boxes are the reconstruction steps. The reconstructor accepts inputs from the centroider, and subtracts reference values from the measured slopes. The first matrix multiply reconstructs the phase using the L matrix (Equation 6.15). The second matrix multiply (control matrix, P_c) splits the phase into tweeter and woofer commands and applies the woofer temporal filter (Equation 6.17). The internal data between the two matrix multiplies can be either the phase values over the pupil, or the Fourier coefficients of those phase values, depending on whether a DFT has been added to L , the reconstruction matrix, and its inverse to P_c , the control matrix. Each of the output channels (Tweeter, Woofer and Filter) are integrated and clipped separately. The integrated tweeter and woofer actuator positions are then sent as output to the mirrors. The integrated filter coefficients are returned as inputs to the control matrix multiply on the next timestep.

6.5.2 Unit Test Verification

The hybrid reconstructor re-uses the core reconstructor components, replacing only the vector-matrix multiply with two vector matrix multiplies (see § 6.3.2 and Figure 6.14). We developed unit tests for the vector matrix multiplies¹⁵ The unit tests also ensure that the full reconstructor operates as designed before inserting it into the full system, by injecting a dummy wavefront sensing image and ensuring that sensible values are returned.

¹⁵There is an entire VMM framework which is re-used between the single vector matrix multiply and hybrid matrix. It only needs to restructure the dimensions of each matrix involved in the matrix multiply.




Figure 6.29: Histogram of reconstructor computation times for the ShadyAO reconstructor. The y-axis is log scaled, but shows the probability of events for each computation time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile computational time is 1.137 frames for the hybrid matrix and 0.588 frames for the VMM, and the median computation time is 0.480 frames for the hybrid and 0.175 frames for the VMM. The hybrid reconstructor shows more computation time by about a factor of two, likely due to the two calls to the OpenBLAS linear algebra library, each which incurs a penalty to parallelize the matrix multiply.

We also checked the computation time and inter-frame delay of the Hybrid Matrix Multiply (HMM) reconstructor, and compared the values to the VMM reconstructor. Figure 6.29 shows the computation time for the HMM reconstructor compared to the VMM reconstructor. The VMM is generally ~ 1 ms faster than the HMM. This is probably due to the second matrix multiply required for the HMM. Since each matrix multiply is parallelized across the 32 cores of the real time computer, each matrix multiply incurs significant overhead to set up, distribute, and then gather the computations.

Figure 6.30 shows the inter-frame times for the HMM, and VMM. The inter-frame times should be close to 1 frame, indicating that the reconstructor is keeping pace with the arrival of new WFS images. Both reconstructors perform with about the same average inter-frame time, but there is much greater variance in the inter-frame arrival times for the HMM.




Figure 6.30: Histogram of inter-frame arrival times for the ShadyAO Vector Matrix Multiply (VMM) and Hybrid Matrix Multiply (HMM) reconstructors. The y-axis is log scaled, but shows the probability of events for each arrival time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile interval in arrival times is 1.255 frames for the VMM and 1.480, and the median inter-arrival time is 1.053 frames for both reconstructors. There is a known 1 frame delay due to the readout settings on the WFS camera. The WFS exposures are handled in parallel, so although inter-arrival times are around 1 frame, the time between exposure start on the WFS and commands sent to the mirrors is 2.053 frames.

6.5.3 Verification with the ShaneAO Simulator

We ran the hybrid matrix multiply on the ShaneAO simulator (Sauce, see § 6.2.2), where we verified that the Hybrid reconstructor converges to static reference centroids, even in the presence of a simulated atmosphere.

6.5.4 Performance with the Internal Light Source

We verified the performance of the hybrid matrix using the internal light source on ShaneAO. The PSDs measured for open and closed loop performance using the ILS are shown in Figure 6.31. Figure 6.32 shows the transfer function for the phase estimate \hat{p} , derived from the individual PSDs shown in Figure 6.31. We find that the system performs similarly with the HMM as with




Figure 6.31: Periodogram (PSD) of estimated phase (\hat{p}) averaged for the Hybrid Matrix Multiply (HMM) on the internal light source (ILS). The orange line shows the temporal power of \hat{p} when the loop is not controlling the mirrors. The absence of any structured turbulence means that in the open loop case, we are essentially measuring white noise. The blue line shows the response of the system in closed loop, with good rejection at low frequencies, and a distinct overshoot at higher frequencies.

the VMM, confirming that the two matrix implementation in the HMM successfully reproduces the one-matrix design of the VMM.

6.5.5 On Sky ETFs and Performance

To confirm that the Hybrid Matrix Multiply reconstructor is working as designed, we recorded pairs of open and closed loop telemetry analogous to the data collected in Section 6.3.4.5. The PSDs are shown in Figure 6.33. The ETFs are shown in Figure 6.34. The model ETFs fit the data quite well. There is a noticeable speedup in the woofer in this dataset (note that $\tau_w = 0$ for the fit data) caused by a fix to the woofer controller scheme, which had previously tied the woofer to the outputs for telemetry recording.




Figure 6.32: Error transfer function (ETF) of estimated phase (\hat{p}) for the Hybrid Matrix Multiply (HMM) on the internal light source (ILS). The *blue* line shows the ETF for the HMM, and the *brown* line shows the ETF for the same parameters using the VMM reconstructor. Note that these were run at two separate times, so some variance due to noise is expected between the VMM and HMM transfer functions. The woofer was turned off for this experiment (as is reflected in the model parameters in Table 6.7). The fit transfer function does do a noticeably better job reproducing the performance of ShaneAO with the HMM reconstructor. Specifically, the HMM reconstructor shows more delay ($\tau = 1.81$ Frames = $7240 \mu\text{s}$) than the VMM. This is consistent with the timing shown in Figure 6.30, where the HMM shows delays which are longer than the frame delay for the Vector Matrix Multiply.

Table 6.7: Model parameters for the fiducial and fit error transfer function (ETF) of estimated phase (\hat{p}) for the Hybrid Matrix Multiply (HMM) on the internal light source (ILS).

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.0	0.99	10.0	1.0	0.99
Fit (<i>red</i>)	1.82	0.09	0.96	1.0	0.0	0.99	10.0	1.0	0.99




Figure 6.33: Periodogram estimated phase (\hat{p}), averaged over the illuminated aperture, recorded with data on sky with a $m_V = 7$ star near zenith using the HMM reconstructor. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.




Figure 6.34: The ETF on sky using the HMM reconstructor. This ETF was computed using the average \hat{p} for ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.07 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data. The ETF fit finds parameters comparable to those found with the new VMM reconstructor.

Table 6.8: Model parameters for the model ETFs shown in Figure 6.34

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{w,c}$	α_w	c_f
Model (<i>green</i>)	1.1	0.07	0.99	1.0	0.07	0.99	33.0	0.2	0.99
Fit (<i>red</i>)	1.61	0.12	0.926	2.32	0.15	0.991	102.8	0.2	0.99

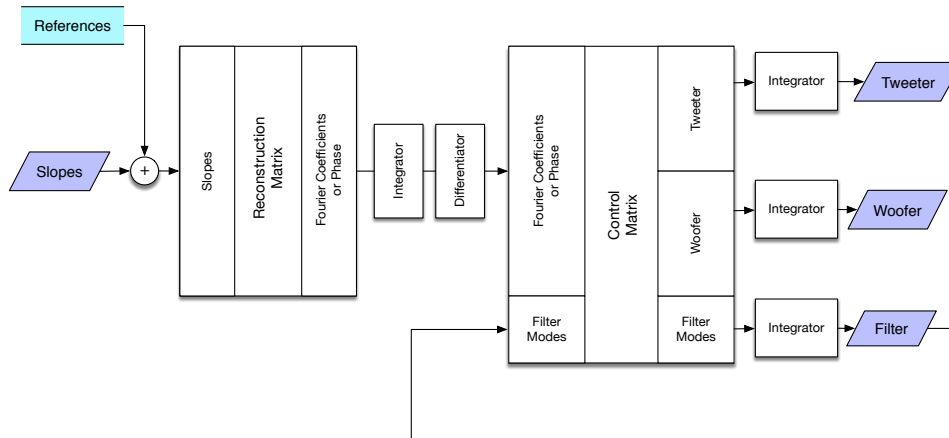


Figure 6.35: Structure of the integrator-differentiator reconstructor (SIDMM) which performs an integration and differentiation on the best-estimate of the phase error, $\Delta\hat{p}$. The integrator and differentiator are inserted between the reconstruction and control matrices. The integrator behaves like the Kalman filter, in that it accepts a measurement of the error ($\Delta\hat{p}$) and returns a measurement of the absolute phase (\hat{p}). The differentiator compensates for this effect so that we can use separate integrators applied to the actuators for each mirror, which allows us to prevent windup and actuator saturation in those mirrors.

6.6 Integrator-Differentiator Reconstructor

The Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) is the next evolution of the hybrid matrix reconstructor towards enabling an linear quadratic gaussian (LQG) controller. Since the LQG controller operates on a measured error in phase ($\Delta\hat{p}$) and produces a “best-estimate” of the open loop phase (\hat{p}), we required a differentiator to produces a “best-estimate” $\Delta\hat{p}$ which can be provided to the woofer-tweeter split and temporal filter. This architecture allows each mirror to have independent integrators and independent clipping points, preventing actuator windup and saturation.

There are two primary modes used with the integrator-differentiator. In the first mode, the integrator and differentiator applied in coefficient space are matched to each other, with matching gains and bleeds applied to the integrator and differentiator. In this mode, the effective gain and bleed for the system is applied to the woofer-tweeter split matrix (\mathbf{P}_c). In the second mode, the effective system gain is applied to the internal integrator, and the differen-

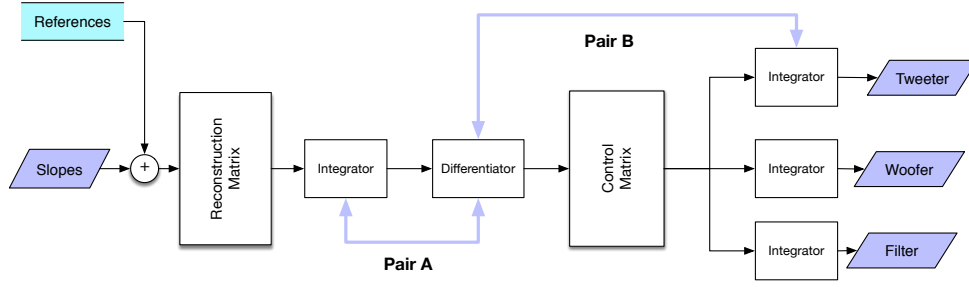


Figure 6.36: Structure of the Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) which performs an integration and differentiation on the best-estimate of the phase error, $\Delta\hat{p}$, showing the two pairs of integrators and differentiators. Pair **A** represents the trivial case, where the differentiator is matched to the integrator which operates on the phase estimate. In this case, the integrator and differentiator should exactly cancel each other. Pair **B** is more analogous to the structure of the LQG reconstructor, where the differentiator is matched to the integrators applied to the actuator commands. In case **B**, we ensure that the gain (g) and leak (c) applied to the differentiator is also applied to the mirror integrators. However, there is no feedback provided from the integrators to the differentiator to correct for large amounts of saturation or windup.

tiator is matched to the integrators which are applied to the actuator commands. These two modes are shown in the diagram in Figure 6.36.

The integrator component re-uses the integrators implemented for the basic vector matrix multiply, described by the discrete time equation:

$$\mathbf{a}_k = c(\mathbf{a}_{(k-1)} - \mathbf{a}_0) + g\Delta\mathbf{a}_k + \mathbf{a}_0 \quad (6.19)$$

The differentiator is implemented as the inversion of this process

$$\Delta\mathbf{a}_k = \mathbf{a}_k - c(\mathbf{a}_{(k-1)} - \mathbf{a}_0) - \mathbf{a}_0 \quad (6.20)$$

where in both equations, \mathbf{a}_0 is the default position of the vector \mathbf{a} . In the case of the mirrors, this is the reference voltage, which is a measured flat position in the middle of the actuator's dynamic range. When matching the differentiator to an integrator (either in pair **A** or **B**), we set the gain of the matched integrator to 1. The overall gain of the system is then controlled by the integrator that is not matched to a differentiator. The gain of the internal integrator is labeled g_c .

The bleed of for the internal integrator is $c_{c,I}$, and for the differentiator, it is $c_{c,D}$, though we guarantee in software that $c_{c,I} = c_{c,D} = c_c$.

6.6.1 Verification

To verify the IDMM reconstructor, we used the existing unit tests, as well as 50 new unit tests which exercise the integrator and differentiator separately. Along with the unit test, we measured the internal transfer function between the start of the integrator and end of the differentiator. These two reconstructor components are designed to exactly cancel each other. Figure 6.37 shows the internal transfer functions measured from the IDMM reconstructor, showing that they are consistent with the theoretical model for the combined integrator-differentiator system.

6.6.2 Performance with the Internal Light Source

The Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) performs as expected on the internal light source. The periodograms are shown in Figure 6.38. Similar to Figure 6.23 and Figure 6.31, the uncontrolled data are roughly flat, and the controlled data shows a normal controlled region, bandwidth and overshoot. There is also evidence of a small vibration in the system in both the open and closed loop data near 60 Hz. This is unsurprising for the observatory environment. However, the amplitude of the vibration is not large enough to impact the performance of the system. Figure 6.39 shows the error transfer function (ETF) of these periodograms (PSDs). The ETF of the Vector Matrix Multiply (VMM) is shown for comparison. The IDMM performs similarly to the VMM, but with a higher bandwidth, and a larger overshoot.

6.6.3 On Sky ETFs and Performance

To confirm that the Integrator-Differentiator Matrix Multiply Reconstructor reconstructor is working as designed, we recorded pairs of open and closed loop telemetry analogous to the data collected in Section 6.3.4.5. The PSDs are shown in Figure 6.40. The ETFs are shown in Figure 6.41. The model




Figure 6.37: Internal transfer function for the IDMM reconstructor. Measurements at **A** are the coefficients, taken after applying the first matrix (**L**). Measurements at **B** are recorded after the integrator has been applied, but before the differentiator is applied. Measurements at **C** are recorded after the differentiator, before the woofer-tweeter matrix split matrix (**P_c**) is applied. The blue line shows the transfer function between **A** and **C**, and so is flat at 1.0, demonstrating that the integrator and differentiator cancel each other out. The orange line shows the the transfer function due to just the integrator (between **A** and **B**). The green line shows the transfer function due to just the differentiator (between **B** and **C**). Both the orange and green dotted lines show the theoretical transfer functions for the integrator and differentiator. These are well matched to the recorded transfer functions from the IDMM reconstructor.




Figure 6.38: Periodogram (PSD) of estimated phase (\hat{p}) averaged for the Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) on the internal light source (ILS). The orange line shows the temporal power of \hat{p} when the loop is not controlling the mirrors. The absence of any structured turbulence means that in the open loop case, we are essentially measuring white noise. The blue line shows the response of the system in closed loop, with good rejection at low frequencies, and a distinct overshoot at higher frequencies.

ETFs fit the data quite well, and the model recovers the internal integrator gain used during this test.




Figure 6.39: Error transfer function (ETF) of estimated phase (\hat{p}) for the Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) on the internal light source (ILS). The *blue* line shows the ETF for the IDMM, and the *brown* line shows the ETF for the same parameters using the VMM reconstructor. Note that these were run at two separate times, so some variance due to noise is expected between the VMM and IDMM transfer functions. The woofer was turned off for this experiment (as is reflected in the model parameters in Table 6.7). The fit transfer function does do a noticeably better job reproducing the performance of ShaneAO with the IDMM reconstructor. Specifically, the IDMM reconstructor shows more delay ($\tau = 2.03$ Frames = $8120 \mu\text{s}$) than the VMM.

Table 6.9: Model parameters for the fiducial and fit error transfer function (ETF) of estimated phase (\hat{p}) for the Integrator-Differentiator Matrix Multiply Reconstructor (IDMM) on the internal light source (ILS).

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{c,w}$	g_c	$c_{c,I}$	$c_{c,D}$
Model	1.1	1.0	0.99	1.0	0.0	0.99	10.0	0.07	0.99	0.99
Fit	2.03	1.0	0.959	1.0	0.0	0.99	10.0	0.1	0.99	0.99




Figure 6.40: Periodogram estimated phase (\hat{p}), averaged over the illuminated aperture, recorded with data on sky with a $m_V = 7$ star near zenith using the IDMM reconstructor. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.




Figure 6.41: The ETF on sky using the IDMM reconstructor. This ETF was computed using the average \hat{p} for ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.07 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data. The ETF fit finds parameters comparable to those found with the new VMM reconstructor.

Table 6.10: Model parameters for the model ETFs shown in Figure 6.41

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{c,w}$	g_c	$c_{c,I}$	$c_{c,D}$
Model	1.1	1.0	0.99	1.0	1.0	0.99	33.0	0.07	0.99	0.99
Fit	1.49	1.0	0.994	0.75	1.0	0.99	200.0	0.12	0.99	0.99

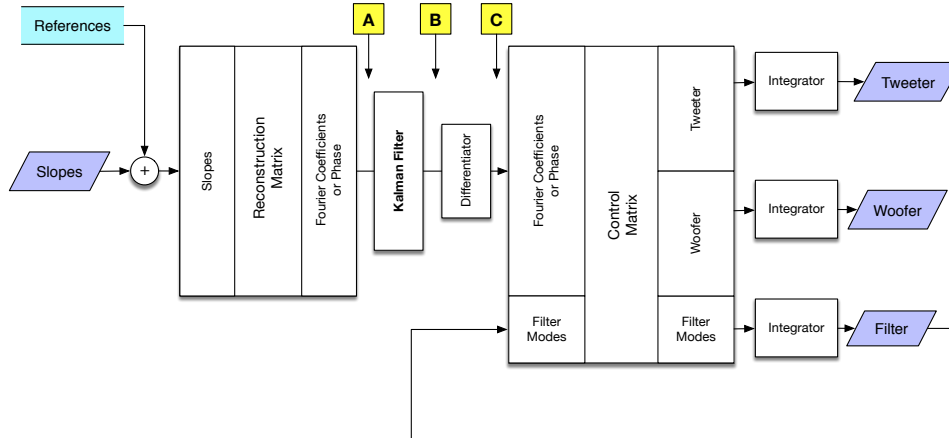


Figure 6.42: Structure of the Kalman filter reconstructor (LQG) which applies a Kalman filter and differentiation on the best-estimate of the phase error, $\Delta\hat{p}$. The Kalman filter and differentiator are inserted between the reconstruction and control matrices. The Kalman filter accepts a measurement of the error ($\Delta\hat{p}$) and returns a measurement of the absolute phase (\hat{p}). The differentiator compensates for this effect so that we can use separate integrators applied to the actuators for each mirror, which allows us to prevent windup and actuator saturation in those mirrors. During operation, the bleed of the differentiator is matched to that of the mirror integrators, and the gain on the mirror integrators is set to 1, as described in pair **B** in § 6.6. The yellow boxes labeled **A**, **B** and **C** show points where we record coefficients in the reconstructor for analysis.

6.7 Kalman Filter Reconstructor

The Kalman filter reconstructor, which implements linear quadratic gaussian control, is implemented in the hybrid reconstructor framework by replacing the integrator described in Section 6.6 with a Kalman Filter. The implementation in ShaneAO uses an external library, FTR, which has a compiled C library and companion python library for saving, loading, and managing linear quadratic gaussian (LQG) filters.¹⁶

6.7.1 Kalman Reconstructor Design

The Kalman filter reconstructor is implemented as a filter applied to the Fourier transform of the estimated phase error $F(\Delta\hat{p})$ in the middle of the

¹⁶Available at <https://github.com/alexrudy/FTR/>

hybrid reconstructor (see Figure 6.42) The reconstructor follows the discrete time equation for Kalman filtering from Poyneer & Véran (2010)

$$F(\hat{p}_k) = \sum_L [Q\alpha_L F(\Delta\hat{p}_k) + \alpha_L \hat{p}_{L,(k-1)}] - F(\hat{p}_{(k-1)})Q \quad (6.21)$$

As implemented in ShaneAO, this discrete time equation has three coefficients in use, $Q\alpha_L = g_L$, which is pre-multiplied and stored, α_L , and Q . This implementation allows us to flatten the LQG reconstructor into a standard integrating controller, by setting $Q = 0$, and then setting the gain g_L and bleed α_L appropriately as real numbers.

To model the Kalman Filter, we use the following z-transform control law

$$C(z) = \prod_L \left(\frac{\alpha_L Q}{1 - \alpha z^{-1}} \right) \left(\frac{1}{1 + Qz^{-1}} \right) \quad (6.22)$$

which can be incorporated into our transform model in equation 6.10.

Along with implementing a standard integrating controller, the LQG filter can be applied in "zero-layer" mode, where we generate a filter for only a single layer, with only real coefficients, so that no prediction is occurring. Figure 6.43 shows the theoretical transfer function for a real-valued LQG controller, compared to a standard integrating controller like that used in the original Vector Matrix Multiply (VMM).

We also checked the computation time and inter-frame delay of the LQG reconstructor, and compared the values to the VMM reconstructor. Figure 6.44 shows the computation time for the LQG reconstructor compared to the VMM reconstructor. The LQG shows a uniform distribution of computation times spanning 1.5 to 4 ms. This might be due to the larger matrices used to perform the DFT in the LQG reconstructor, which produce the full 1024 complex-valued coefficients instead of the 1024 real coefficients used by the Hybrid Matrix Multiply (HMM) reconstructor. However, since each matrix multiply is parallelized across the 32 cores of the real time computer, each matrix multiply incurs significant overhead to set up, distribute, and then gather the computations, but they should scale relatively well with the size of the matrix in the




Figure 6.43: The theoretical error transfer function of the real-valued LQG filter (blue, equation 6.22), with $Q = 0.3$ and $\alpha = 0.95$. and the theoretical transfer function for the standard integrator (orange, equation 6.8) with $g = 0.3$ and $c = 0.95$. The two controllers produce similar performance, though the LQG has a lower bandwidth.

matrix multiply. We don't have a good explanation for the uniform distribution of computation times shown in Figure 6.44

Figure 6.30 shows the inter-frame round trip times for the LQG, HMM, and VMM. The inter-frame times should be close to 1 frame, indicating that the reconstructor is keeping pace with the arrival of new wavefront sensor (WFS) images. Unlike the HMM and VMM, the longer computation times for the LQG and the inter-command times which are greater than 1 frame, suggest that the LQG is either sometimes dropping WFS frames or delaying the computation of the next WFS frame in order to complete work on the current frame.

6.7.2 Unit Test Verification

The Kalman Filter is tested as part of the unit and functional tests in FTR, which contains 407 unit tests, including benchmarks to confirm that the performance of the Kalman filter as implemented scale in the expected way. We implemented a thin wrapper around the Kalman filter in FTR to provide an




Figure 6.44: Histogram of reconstructor computation times for the ShadyAO VMM, HMM and Kalman Filter reconstructor. The y-axis is log scaled, but shows the probability of events for each computation time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile computational time is 1.358 frames for the Kalman filter compared to 0.588 frames for the VMM, and the median computation time is 0.692 frames for the Kalman filter and 0.175 frames for the VMM. The Kalman Filter reconstructor shows a large variance in the computation times, significantly greater than what would be naively expected from the computational complexity of the implementation.

interface that matches the ShaneAO reconstructor described in Section 6.3.

Along with the unit tests, we ensured that the internal transfer functions for the Kalman Filter (recorded as the ratio of the Fourier coefficients at points **A**, **B** and **C** in Figure 6.42) when behaving like an integrator (i.e. with $Q = 0$ in equations 6.21 and 6.22) in Figure 6.46.

As well as the integrator behavior, we check the performance of the Kalman filter in the 0-layer case (referred to as DC-only), shown in Figure 6.47, which is analogous to Figure 6.46.

6.7.3 Performance with the Internal Light Source

Figure 6.48 shows the periodograms (PSDs) of the DC-only Kalman Filter on the telescope simulator (TelSim) with the rotating phase plate run-




Figure 6.45: Histogram of inter-frame arrival times for the ShadyAO VMM, HMM and Kalman Filter reconstructors. The y-axis is log scaled, but shows the probability of events for each arrival time. The data was collected with the reconstructor running at 250 Hz over 4 seconds. The 99th percentile interval in arrival times is 1.255 frames for the VMM and 1.480, and the median inter-arrival time is 1.053 frames for both reconstructors. There is a known 1 frame delay due to the readout settings on the WFS camera. The WFS exposures are handled in parallel, so although inter-arrival times are around 1 frame, the time between exposure start on the WFS and commands sent to the mirrors is 2.053 frames. The Kalman filter reconstructor shows a large variance in the inter-frame arrival times, due to the large variance in the computation times (see Figure 6.44)

ning. The strong power in the open loop (orange) PSD near zero is a result of the turbulent screen on the phase plate. The closed loop PSD shows good performance, but does show evidence of an extra bump just outside the control bandwidth. We cannot currently explain the source of this bump. The resulting error transfer function (ETF) is shown in Figure 6.49

6.7.4 On Sky ETFs and Performance

To confirm that the Passthrough Kalman Filter Matrix Multiply Reconstructor reconstructor is working as designed, we recorded pairs of open and closed loop telemetry analogous to the data collected in Section 6.3.4.5. The PSDs are shown in Figure 6.50. The ETFs are shown in Figure 6.51. The model




Figure 6.46: Internal transfer functions from Fourier coefficients measured at points **A**, **B** and **C** in Figure 6.42 showing the integrator (blue) and differentiator (orange) in the middle of the KFMM, and their joint transfer function (green), which shows that both components are well matched.




Figure 6.47: Internal transfer functions from Fourier coefficients measured at points **A**, **B** and **C** in Figure 6.42 showing the DC-only Kalman Filter (blue) and differentiator (orange) in the middle of the KFMM, and their joint transfer function (green), which shows that both components are well matched.




Figure 6.48: Periodogram (PSD) of Fourier transform of the estimated phase ($F(\hat{p})$) averaged for the Passthrough Kalman Filter Matrix Multiply Reconstructor (Passthrough KFMM) on the telescope simulator (TelSim). The orange line shows the temporal power of \hat{p} when the loop is not controlling the mirrors. The absence of any structured turbulence means that in the open loop case, we are essentially measuring white noise. The blue line shows the response of the system in closed loop, with good rejection at low frequencies, and a distinct overshoot at higher frequencies.

ETFs fit the data quite well, and the model recovers the internal integrator gain used during this test.

6.7.5 Kalman Reconstructor Verification

To verify the performance of the Kalman Filter reconstructor, we examined the internal transfer functions which describe the frequency response of the Kalman Filter controller when combined with the differentiator. To do this, we recorded the Fourier coefficient telemetry from before the Kalman filter, between the Kalman filter and the differentiator, and after the differentiator (see Figure 6.42). Transfer functions showing the progression of the controller through these steps are in Figure 6.52

6.7.6 Future Directions with the Kalman Filter

Although we are successfully able to demonstrate that the Kalman Filter is performing as designed when measuring internal coefficients, we are not able to successfully verify this performance by observing the transfer functions in Figure 6.52 appear when measured in two independent sets of telemetry. When we record closed loop telemetry with the Kalman Filter, we find that it does not produce a notch in temporal space as expected, and instead causes periodic instabilities in the AO correction.




Figure 6.49: Error transfer function (ETF) of Fourier transform of the estimated phase ($F(\hat{p})$) for the Passthrough Kalman Filter Matrix Multiply Reconstructor (Passthrough KFMM) on the internal light source (ILS). The *blue* line shows the ETF for the Passthrough KFMM, and the *brown* line shows the ETF for the same parameters using the VMM reconstructor. Note that these were run at two separate times, so some variance due to noise is expected between the VMM and Passthrough KFMM transfer functions. The woofer was turned off for this experiment (as is reflected in the model parameters in Table 6.7). The fit transfer function does do a noticeably better job reproducing the performance of ShaneAO with the Passthrough KFMM reconstructor. Specifically, the Passthrough KFMM reconstructor shows more delay ($\tau = 4.3$ Frames) than the VMM.

Table 6.11: Model parameters for the fiducial and fit error transfer function (ETF) of Fourier transform of the estimated phase ($F(\hat{p})$) for the Passthrough Kalman Filter Matrix Multiply Reconstructor (Passthrough KFMM) on the internal light source (ILS).

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{c,w}$	Q	α_L	$c_{c,D}$
Model	1.1	1.0	0.99	1.0	0.0	0.99	10.0	0.07	0.99	0.99
Fit	4.39	1.0	0.993	1.0	0.0	0.99	10.0	0.08	0.993	0.99




Figure 6.50: PSD estimated phase ($F(\hat{p})$), averaged over the illuminated aperture, recorded with data on sky with a $m_V = 7$ star near zenith using the Passthrough KFMM reconstructor. The blue line shows the loop closed, controlled behavior of the system, which eliminates much of the power below 10 Hz, and shows a bandwidth of about 10 Hz with an overshoot region beyond 10 Hz.




Figure 6.51: The ETF on sky using the Passthrough KFMM reconstructor. This ETF was computed using the average \hat{p} for ShaneAO. The blue line shows the empirical data, the red line shows a simple ETF model of the system, with parameters set to the nominal values in the system (i.e. the system was running with a gain of 0.07 when the loop closed data was recorded), and the green line shows a simple model ETF with the best fit parameters from a fit to the ETF data. The ETF fit finds parameters close to the expected values, with the exception of Q , which is almost exactly a factor of 2 less than the input value. This may be related to reduced gains discussed in Section 6.2.5.1.

Table 6.12: Model parameters for the model ETFs shown in Figure 6.51

Model	τ_c	g_t	c_t	τ_w	g_w	c_w	$\omega_{c,w}$	Q	α_L	$c_{c,D}$
Model	1.1	1.0	0.99	1.0	1.0	0.99	33.0	0.2	0.995	0.99
Fit	1.78	1.0	0.996	1.43	1.0	0.985	200.0	0.09	0.988	0.99




Figure 6.52: Internal error transfer functions (ETFs) for the Kalman Filter Matrix Multiply Reconstructor (KFMM) showing the effect of the controller on the input residuals. The blue line shows the performance of the Kalman Filter. The solid blue shows the ETF of the Fourier coefficients recorded from ShaneAO, and the dashed blue line shows the z -transform theoretical ETF for the Kalman filter (see Equation 6.22). The notch frequency for this single-wind-layer filter shown as a vertical black dashed line. The orange line shows the performance of the differentiator, again where the dashed line shows the z -transform theoretical performance, and the solid line shows the ETF of the Fourier coefficients. The green line shows the combined effect of both the Kalman filter and the differentiator. As expected, the notch from the Kalman filter remains, and the differentiator has slightly different gain characteristics from the Kalman filter, resulting in the green line not following a constant value even at temporal frequencies not impacted by the filtering notch. Also, the Kalman filter is designed to have an effective gain of $g \sim 0.3$, but the differentiator in this case is matched to the actuator integrators (case **B** from Figure 6.36), and so has a gain of $g = 1$. The average value of the temporal response in the joint ETF is therefore greater than 1, reflecting that some gain has occurred in the system.

6.8 Conclusions

We presented a demonstration of linear quadratic gaussian (LQG) control on sky with ShaneAO, a split-mirror adaptive optics (AO) system. This required detailed characterization of the performance of ShaneAO, a major refactoring of the reconstructor code to improve flexibility, and a custom analysis framework to understand the properties of the split-mirror reconstruction with the LQG controller.

The characterization consisted of verifying the behavior and performance of the original AO system, understanding the delay, gain and default input parameters used for ShaneAO, and ensuring that we could re-produce measured transfer functions both on sky and with the internal light source. We developed a model of the error transfer function for the two mirror system with a temporal split between the mirrors.

We re-factored the ShaneAO control software to facilitate the use of reconstructors as modular components which could be replaced during the operation of the instrument. We also improved the acquisition of telemetry data, and the consistency of the end user interface and configuration system.

We developed a hybrid reconstruction scheme, which is capable of providing an in-the-loop estimate of either the phase, or the Fourier transform of the phase. This reconstruction scheme allowed us to split the classical reconstruction into one matrix, and the modal split between the woofer and tweeter into a second matrix.

Finally, we implemented two additional reconstructors, a matched integrator differentiator pair and a linear quadratic gaussian controller. All of our reconstructors were rigorously tested both in the laboratory and on sky against natural guide star targets.

The next step towards a demonstration of predictive Fourier control is to verify that the LQG controller can produce notched control away from 0 Hz and detect this control in the error transfer function. Then, we must connect the predictive Fourier control (PFC) controller to an accurate current estimate of the wind to produce a controller which can predict a given wind velocity.

Chapter 7

Conclusions

This dissertation contains three parts, which are held together by their interdependence on AO.

Part I presented work searching for the γ -ray flares which emanate from the Crab Nebula. Although we were not conclusive in identifying the origin of the γ -ray flares, we conducted an extensive, multi-wavelength study of the behavior of the inner knot. Our work on the inner knot has sparked a wider range of papers and debate about the nature of the knot from the theoretical community (Porth et al., 2017, 2016; Komissarov, 2013), and presents several avenues for future work exploring the influence of the pulsar on the surrounding Crab Nebula. With the new Laser Guide Star facility at Lick Observatory, the ShaneAO system will be well positioned to contribute to our understanding of the inner Knot and the Crab nebula. The latest theoretical work on the knot (Porth et al., 2017) suggests that understanding both the spectral index and the polarization of the Knot, two properties that are measurable with ShaneAO, would strongly constrain the potential origins of the inner Knot and help to identify the mechanisms which are driving the termination shock of the pulsar wind nebula.

Part II presents spatially resolved near-infrared (NIR) spectroscopic observations of 11 nearby ultra-luminous infrared galaxies (ULIRGs). The observations of IRAS F20414-1651 are presented in detail, showing that the velocities observed in the NIR with OSIRIS (Oh-Suppressing Infrared Integral Field Spectrograph) provide a consistent explanation for the full velocity profile ob-

served with Herschel in Spoon et al. (2013), despite the narrow field of view. I anticipate finishing the remaining pieces of analysis for Chapter 3 and submitting this chapter to a journal. Chapter 4 presents the remaining 10 ULIRGs in our survey. Our survey provides significant evidence that the high velocity outflows reported from Herschel in Spoon et al. (2013) and Veilleux et al. (2013b) do not necessarily correspond to spatially and kinematic distinct structure that would result from gas being ejected from the galaxy. In 4 objects we see evidence for distinct kinematic components, blue shifted with velocities consistent with those observed with Herschel. In 3 of these 4 objects, we cannot conclusively identify an outflow, but rather we identify an anomalous velocity component, inconsistent with the ordered rotation velocities observed in the host system. All together, our survey suggests that the velocities reported by Spoon et al. (2013) and Veilleux et al. (2013b) represent the extremes within individual systems, rather than evidence of high velocity, spatially and kinematically distinct components. There is a wealth of further analysis which can be done based on the data presented in Chapter 4. In the near-term, it would be good to publish an overview paper of the survey, which could be extended to include some of the GOALS galaxies previously published in Medling et al. (2015a) and Medling et al. (2014).

Part III presented a successful on sky demonstration of full LQG control for adaptive optics, and progress towards the implementation of PFC on sky and in a laboratory demonstrator. The LQG controller, as implemented in ShaneAO, compensates for the delay in the system, but does not do any of the active wind prediction. The laboratory demonstration of PFC in Chapter 5 lacks only a single data set which demonstrates the predictor in the error transfer functions (ETFs) from the controller. Unfortunately, the laboratory test bench requires significant work to be fully operational. The bench needs the integration of the new cameras to be completed, a re-test of the remaining hardware, and a new optical alignment, since the alignment changed by necessity when the old WFS cameras were replaced. If the test bench is brought online to its performance in 2015, then a demonstration of PFC should be easily completed. Otherwise, it should be possible to demonstrate PFC on ShaneAO using the telescope simulator, and achieve an analogous result.

Chapter 6 presents an in-depth characterization and re-structuring of ShaneAO. The work already completed has led to the discovery of a number of pending issues which should improve the performance of ShaneAO as a facility instrument, such as the high slope regularization value described in Section 6.2.5.1. As well, our work on PFC surfaced a number of minor issues which can be easily addressed to improve the performance of ShaneAO, such as the low effective gain and high overshoot apparent at higher wavefront sensor (WFS) camera rates. Several other aspects of the work in Chapter 6 should be contributed back to the observatory, including the software refactoring that facilitates changing reconstructors and providing real-time feedback to AO operations in the form of telemetry data. Finally, the demonstration of LQG control on sky should be extracted and published, as full-modal LQG control on sky is a novel feature for a facility AO system.

In the future, the PFC techniques described and advanced in Part III can be used to improve any existing AO system. In fact, PFC is a prime candidate algorithm to include in new reconstructor development for existing AO systems, as it can be readily retrofitted to an existing AO system. That I spent so much time and energy in Chapter 6 developing PFC without yet demonstrating it on sky speaks to the way in which PFC, and other atmospheric modeling techniques require a deeper understanding of the inner workings of the AO system.

The future for all three parts of this dissertation is bright. Each presents avenues for further discovery and analysis, and the work presented here has the potential to circle back around. Specifically, PFC can improve the performance of many AO systems, including improving the performance of ShaneAO, to help facilitate additional studies of the inner knot of the Crab nebula, and improving Keck AO, which would increase the efficiency with which ULIRGs can be studied with OSIRIS.

Bibliography

- Aliu, E., Archambault, S., Aune, T., et al. 2014, *ApJ*, 781, L11
- Ammons, S. M., Johnson, L., Laag, E. A., et al. 2010, *PASP*, 122, 573
- Ammons, S. M., Laag, E. A., Kupke, R., Gavel, D. T., & Max, C. E. 2007, in *Proc. SPIE*, Vol. 6691, *Astronomical Adaptive Optics Systems and Applications III*, 669108
- Ammons, S. M., Kupke, R., Laag, E. A., et al. 2006, in *Proc. SPIE*, Vol. 6272, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 627202
- Arons, J. 2012, *Space Sci. Rev.*, 173, 341
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
- Baty, H., Petri, J., & Zenitani, S. 2013, *MNRAS*, 436, L20
- Beckwith, S. V. W., Stiavelli, M., Koekemoer, A. M., et al. 2006, *AJ*, 132, 1729
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, in *Proc. SPIE*, Vol. 7014, *Ground-based and Airborne Instrumentation for Astronomy II*, 701418
- Bogovalov, S. V. 1999, *A&A*, 349, 1017
- Buehler, R., Scargle, J. D., Blandford, R. D., et al. 2012, *ApJ*, 749, 26
- Bühler, R., & Blandford, R. 2014, *Reports on Progress in Physics*, 77, 066901
- Bushouse, H. A., Borne, K. D., Colina, L., et al. 2002, *ApJS*, 138, 1
- Buson, S., Buehler, R., & Hays, E. 2013, *The Astronomer's Telegram*, 5485
- Bykov, A. M., Pavlov, G. G., Artemyev, A. V., & Uvarov, Y. A. 2012, *MNRAS*, 421, L67
- Camus, N. F., Komissarov, S. S., Bucciantini, N., & Hughes, P. A. 2009, *MNRAS*, 400, 1241

- Cerutti, B., Uzdensky, D. A., & Begelman, M. C. 2012, *ApJ*, 746, 148
- Cerutti, B., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2013, *ApJ*, 770, 147
- Cicone, C., Maiolino, R., Sturm, E., et al. 2014, *A&A*, 562, A21
- Clausen-Brown, E., & Lyutikov, M. 2012, *MNRAS*, 426, 1374
- Contursi, A., Poglitsch, A., Grácia Carpio, J., et al. 2013, *A&A*, 549, A118
- Davies, R. I. 2007, *MNRAS*, 375, 1099
- Davies, R. I., Maciejewski, W., Hicks, E. K. S., et al. 2014, *ApJ*, 792, 101
- Davies, R. L., Medling, A. M., U, V., et al. 2016, *MNRAS*, 458, 158
- Davis, T. A., Krajnović, D., McDermid, R. M., et al. 2012, *MNRAS*, 426, 1574
- Di Matteo, T., Springel, V., & Hernquist, L. 2005a, in *Growing Black Holes: Accretion in a Cosmological Context*, ed. A. Merloni, S. Nayakshin, & R. A. Sunyaev, 340–345
- Di Matteo, T., Springel, V., & Hernquist, L. 2005b, *Nature*, 433, 604
- Feruglio, C., Fiore, F., Piconcelli, E., et al. 2013, *A&A*, 558, A87
- Feruglio, C., Aussel, H., Le Floch, E., et al. 2010, *ApJ*, 721, 607
- Fruchter, A. S., & Hook, R. N. 2002, *PASP*, 114, 144
- Gavel, D., Kupke, R., Dillon, D., et al. 2014, in *Proc. SPIE*, Vol. 9148, *Adaptive Optics Systems IV*, 914805
- Gavel, D. T. 2003, in *Proc. SPIE*, Vol. 4839, *Adaptive Optical System Technologies II*, ed. P. L. Wizinowich & D. Bonaccini, 972–980
- Gavel, D. T., Kupke, R., Rudy, A. R., et al. 2016, in *Proc. SPIE*, Vol. 9909, *Adaptive Optics Systems V*, 99092W
- Gavel, D. T., & Wiberg, D. 2003, in *Proc. SPIE*, Vol. 4839, *Adaptive Optical System Technologies II*, ed. P. L. Wizinowich & D. Bonaccini, 890–901
- Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, *ApJ*, 689, 1044
- Gofford, J., Reeves, J. N., Tombesi, F., et al. 2013, *MNRAS*, 430, 60
- González-Alfonso, E., Fischer, J., Bruderer, S., et al. 2013, *A&A*, 550, A25
- Gupta, A., Mathur, S., Krongold, Y., & Nicastro, F. 2013, *ApJ*, 772, 66
- H. E. S. S. Collaboration, Abramowski, A., Aharonian, F., et al. 2014, *A&A*, 562, L4
- Hardy, J. W. 1998, *Adaptive Optics for Astronomical Telescopes* (Oxford, UK: Oxford University Press), 448

- Hester, J. J. 2008, *ARA&A*, 46, 127
- Hester, J. J., Scowen, P. A., Sankrit, R., et al. 1995, *ApJ*, 448, 240
- Hicks, E. K. S., Davies, R. I., Maciejewski, W., et al. 2013, *ApJ*, 768, 107
- Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, *ApJS*, 180, 225
- Hopkins, P. F., & Hernquist, L. 2009, *ApJ*, 694, 599
- Komissarov, S. S. 2013, *MNRAS*, 428, 2459
- Komissarov, S. S., & Lyubarsky, Y. E. 2003, *MNRAS*, 344, L93
- Komissarov, S. S., & Lyutikov, M. 2011, *MNRAS*, 414, 2017
- Krajnović, D., Cappellari, M., de Zeeuw, P. T., & Copin, Y. 2006, *MNRAS*, 366, 787
- Kron, R. G. 1980, *ApJS*, 43, 305
- Kupke, R., Gavel, D., Roskosi, C., et al. 2012, in *Proc. SPIE*, Vol. 8447, *Adaptive Optics Systems III*, 84473G
- Laag, E. A., Ammons, S. M., Gavel, D. T., & Kupke, R. 2008, *Journal of the Optical Society of America A*, 25, 2114
- Larkin, J., Barczys, M., Krabbe, A., et al. 2006, in *Proc. SPIE*, Vol. 6269, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 62691A
- Le Roux, B., Conan, J.-M., Kulcsár, C., et al. 2004, *Journal of the Optical Society of America A*, 21, 1261
- Lupton, W. F., & Conrad, A. R. 1993, in *Astronomical Society of the Pacific Conference Series*, Vol. 52, *Astronomical Data Analysis Software and Systems II*, ed. R. J. Hanisch, R. J. V. Brissenden, & J. Barnes, 315
- Lyubarsky, Y. E. 2012, *MNRAS*, 427, 1497
- Lyutikov, M., Balsara, D., & Matthews, C. 2012, *MNRAS*, 422, 3118
- Lyutikov, M., Pariev, V. I., & Blandford, R. D. 2003, *ApJ*, 597, 998
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proceedings of the National Academy of Science*, 111, 12661
- Madec, P. Y. 2004, in *Adaptive Optics in Astronomy*, ed. F. Roddier (Cambridge, UK: Cambridge University Press),
- Maréchal, A. 1947, *Rev. Opt*, 26, 257
- Markwardt, C. B. 2009, in *Astronomical Society of the Pacific Conference Series*, Vol. 411, *Astronomical Data Analysis Software and Systems XVIII*, ed. D. A. Bohlender, D. Durand, & P. Dowler, 251

- Martin, C. L. 2005, *ApJ*, 621, 227
- Martin, C. L., & Soto, K. T. 2016, *ApJ*, 819, 49
- Mayer, M., Buehler, R., Hays, E., et al. 2013, *ApJ*, 775, L37
- McGurk, R., Rockosi, C., Gavel, D., et al. 2014, in *Proc. SPIE*, Vol. 9148, Adaptive Optics Systems IV, 91483A
- Medling, A. M., U, V., Guedes, J., et al. 2014, *ApJ*, 784, 70
- Medling, A. M., U, V., Max, C. E., et al. 2015a, *ApJ*, 803, 61
- Medling, A. M., U, V., Rich, J. A., et al. 2015b, *MNRAS*, 448, 2301
- Moran, P., Shearer, A., Mignani, R. P., et al. 2013, *MNRAS*, 433, 2564
- Müller-Sánchez, F., Hicks, E. K. S., Malkan, M., et al. 2017, *ArXiv e-prints*, arXiv:1705.06678
- Ojha, R., Hays, E., Buehler, R., & Dutka, M. 2013, *The Astronomer's Telegram*, 4855
- Petit, C., Conan, J.-M., Kulcsár, C., Raynaud, H.-F., & Fusco, T. 2008, *Optics Express*, 16, 87
- Petit, C., Sauvage, J.-F., Fusco, T., et al. 2014, in *Proc. SPIE*, Vol. 9148, Adaptive Optics Systems IV, 91480O
- Porth, O., Buehler, R., Olmi, B., et al. 2017, *Space Sci. Rev.*, 207, 137
- Porth, O., Komissarov, S. S., & Keppens, R. 2014, *MNRAS*, 438, 278
- Porth, O., Vorster, M. J., Lyutikov, M., & Engelbrecht, N. E. 2016, *MNRAS*, 460, 4135
- Poyneer, L., van Dam, M., & Véran, J.-P. 2009, *Journal of the Optical Society of America A*, 26, 833
- Poyneer, L., & Véran, J.-P. 2008, *Journal of the Optical Society of America A*, 25, 1486
- Poyneer, L. A., Gavel, D. T., & Brase, J. M. 2002, *Journal of the Optical Society of America A*, 19, 2100
- Poyneer, L. A., Macintosh, B. A., & Véran, J.-P. 2007, *Journal of the Optical Society of America A*, 24, 2645
- Poyneer, L. A., & Véran, J.-P. 2010, *Journal of the Optical Society of America A*, 27, A223
- Poyneer, L. A., De Rosa, R. J., Macintosh, B., et al. 2014, in *Proc. SPIE*, Vol. 9148, Adaptive Optics Systems IV, 91480K

- Rampy, R., Gavel, D., Dillon, D., & Thomas, S. 2010, in Proc. SPIE, Vol. 7736, Adaptive Optics Systems II, 77362Y
- Rangwala, N., Maloney, P. R., Glenn, J., et al. 2011, ApJ, 743, 94
- Rees, M. J., & Gunn, J. E. 1974, MNRAS, 167, 1
- Rudy, A., Horns, D., DeLuca, A., et al. 2015, ApJ, 811, 24
- Rudy, A. R., Srinath, S., Poyneer, L., et al. 2014, in Proc. SPIE, Vol. 9148, Adaptive Optics Systems IV, 91481Z
- Rupke, D. S. N., & Veilleux, S. 2011, ApJ, 729, L27
- . 2013, ApJ, 775, L15
- Sandberg, A., & Sollerman, J. 2009, A&A, 504, 525
- Silk, J., & Rees, M. J. 1998, A&A, 331, L1
- Sivo, G., Kulcsár, C., Conan, J.-M., et al. 2014, Optics Express, 22, 23565
- Sollerman, J. 2003, A&A, 406, 639
- Soto, K. T., & Martin, C. L. 2012, ApJS, 203, 3
- Soto, K. T., Martin, C. L., Prescott, M. K. M., & Armus, L. 2012, ApJ, 757, 86
- Spoon, H. W. W., & Holt, J. 2010, in Astronomical Society of the Pacific Conference Series, Vol. 427, Accretion and Ejection in AGN: a Global View, ed. L. Maraschi, G. Ghisellini, R. Della Ceca, & F. Tavecchio, 80
- Spoon, H. W. W., Farrah, D., Lebouteiller, V., et al. 2013, ApJ, 775, 127
- Strauss, M. A., Huchra, J. P., Davis, M., et al. 1992, ApJS, 83, 29
- Striani, E., Tavani, M., Piano, G., et al. 2011, ApJ, 741, L5
- Sturm, E., González-Alfonso, E., Veilleux, S., et al. 2011, ApJ, 733, L16
- Sturrock, P., & Aschwanden, M. J. 2012, ApJ, 751, L32
- Tavani, M., Bulgarelli, A., Vittorini, V., et al. 2011, Science, 331, 736
- Tchekhovskoy, A., Spitkovsky, A., & Li, J. G. 2013, MNRAS, 435, L1
- Tennant, A. F., Becker, W., Juda, M., et al. 2001, ApJ, 554, L173
- Teraki, Y., & Takahara, F. 2013, ApJ, 763, 131
- U, V., Medling, A., Sanders, D., et al. 2013, ApJ, 775, 115
- Uzdensky, D. A., Cerutti, B., & Begelman, M. C. 2011, ApJ, 737, L40
- van Dam, M. A., Le Mignant, D., & Macintosh, B. A. 2004, Appl. Opt., 43, 5458
- van Dam, M. A., Bouchez, A. H., Le Mignant, D., et al. 2006, PASP, 118, 310
- Veilleux, S., Meléndez, M., Sturm, E., et al. 2013a, ApJ, 776, 27
- Veilleux, S., Trippe, M., Hamann, F., et al. 2013b, ApJ, 764, 15

Weisskopf, M. C., Tennant, A. F., Arons, J., et al. 2013, *ApJ*, 765, 56
Wizinowich, P., Acton, D. S., Shelton, C., et al. 2000, *PASP*, 112, 315
Wizinowich, P. L., Le Mignant, D., Bouchez, A. H., et al. 2006, *PASP*, 118, 297
Yelda, S., Lu, J. R., Ghez, A. M., et al. 2010, *ApJ*, 725, 331
Yuan, Q., Yin, P.-F., Wu, X.-F., et al. 2011, *ApJ*, 730, L15

**AT HIGH RESOLUTION: THE CRAB PULSAR, KINEMATICS OF
ULIRGS, AND LINEAR QUADRATIC GAUSSIAN CONTROL FOR
ADAPTIVE OPTICS**

Alexander R. Rudy

Typeset

Wednesday 30th August, 2017 at 23:11

Typeset using X_YL^AT_EX version 3.1415926-2.2-0.9997.4

Primarily using the font Palatino, as well as Monaco

228 Pages