# Mapping the Diffuse Universe: Integral Field Spectroscopy of Galaxy Environments

Thesis by Dónal Brendan O'Sullivan

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

# Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY Pasadena, California

> 2020 Defended July 24, 2020

© 2020

Dónal Brendan O'Sullivan ORCID: 0000-0002-4959-9179

All rights reserved

## ACKNOWLEDGEMENTS

No individual can build a modern telescope, launch a satellite, or explain the origins of galaxies by themselves. This is one of the things I love most about modern science: it represents the best part of human nature, working together to accomplish something incredible which no one of us could accomplish alone. Over the past seven years, I have learned that a PhD is much the same. I could not have come this far without the guidance and support of a large number of people. I will do my best to thank as many of them as I can here, but this is by no means an exhaustive list. If anyone should find themselves missing from these acknowledgements, they have my sincere apology. It has been an intensely caffeine-fueled week.

First and foremost, of course, I would like to thank my advisor, Christopher Martin. From the moment I joined Chris' group, he has treated me as a friend and colleague. Chris has always entrusted me with the freedom to choose my own path, make mistakes, and learn from them. I believe this is one of the most important freedoms a scientist can have, but it is not one that can be achieved without a firm bedrock of support. Chris provided that for me throughout my PhD experience, so that I could explore and grow as a researcher. Thank you, Chris, it has been an incredible experience.

The research group Chris has assembled also happens to be one of the nicest groups of people I have ever had the pleasure of meeting. Among the members of this group are some people to whom I owe extra special thanks. Keri Hoadley, for her mentorship and relentless support in getting the results of the FLASHES Survey written up and published, getting my thesis together, and generally staying sane. I do not know how I would have managed the past six or seven months, in particular, without her support. To Erika Hamden, now in the distant land of Arizona, I owe similar thanks. As the (insanely efficient) project manager for FIREBall, Erika trusted me with the responsibility of building the communications system, mentored me, and supported me in writing the FLASHES pilot survey paper. Deflated balloons aside, our team's time in Fort Sumner, New Mexico remains an absolute highlight of my PhD experience.

Mateusz (Matt) Matuszewski, the 'ghost in the machine' of the cosmic web imager instruments, has been a mentor to me from the moment I joined Chris' group. Throughout every observation I have conducted for the past six years, the design and implementation of the FLASHES survey, my work on the Keck Cosmic Web Imager, and the communications side of FIREBall-2's guider software, Matt's expertise and guidance has been a critical ingredient in any success I have had.

If Matt is the ghost in the machine, Don Neill is the ghost in the code. I have never encountered a data reduction issue Don could not fix. He has been a critical element any time I have had to make sense of confusing observational data, which was a frequent occurence. He also plays a mean guitar.

Thanks also to Gillian Kyne, Behnam Darvish, and Zeren Lin, who have provided perspective, input and the occasional late-night shift at Palomar on many occasions over the years. Our team lunches have always been a highlight of any week. Cherie and Debby also deserve a huge thank you for their help (with basically everything) over the years.

Thank you to my committee members, Chuck Steidel, Gregg Hallinan, Dimitri Mawet, and Phil Hopkins, whom I have had the pleasure of knowing since my first year at Caltech. Chuck and Phil have always answered my questions about the circumgalactic medium, dark matter halos, or  $Ly\alpha$  radiative transfer with incredible insight, and given me a lot to think about each time we have talked. Although not in the same observational field, Gregg and Dimitri have always offered me great advice on pursuing instrumental projects in graduate school, and have always been available for a friendly chat. Thank you all.

No set of acknowledgements would be complete without the astronomy and astrophysics graduate students. Who could possibly remember which day Thursday is without donut hour? Thank you all for building a wonderful community.

Three years of my graduate student life were spent as a graduate resident associate, living in the Catalina apartments. During those years, I had immense support from my team in Nilza-Marie Santana-Castillo, Annelise Thompson, Emily Wyatt, Sina Booeshaghi, and Miggy Chuapoco. Thank you all for being wonderful teammates.

As an RA, and in various other roles, I had the pleasure of getting to know some of the wonderful people who work hard to make the Caltech community a better place to live, work, and study. To Doug Rees, Kate McAnulty, Felicia Hunt, Erin-Kate Escobar and Taso Dimitriadis from CCD, Daniel Yoder and Laura Flower Kim from ISP, thank you for supporting me and other graduate students every day. Particular thanks to Grace Ho from Occupational Therapy, whose advice and support helped me write this thesis. Although I have not seen many of them in quite a while (thanks, COVID-19), my friends - both here in Pasadena, back in Ireland, and beyond - have helped keep me sane, or some version of it. Thank you Noah Olsman, Ajay Bawa, Scott Barenfeld, Ryan Monroe, Heather Curtis, Eoin Miller, Graham Toomey, Mark Murray, Fran O'Connor, and many more.

Almost last, but nowhere near least, thank you to my family. My parents, Deirdre and Brendan, and my sisters Roisín, Niamh, and Siobhán, all of whom I miss every day. Thanks to the newest generation of our family, my nephews and nieces, Cillian, Cara, and Tommy. They did not help write a single word, to be completely honest, but the pictures and videos of them have provided some much needed distraction and comic relief from time to time.

Finally, thank you to my best friend and incredible girlfriend, Sam. Of all the people to be quarantined with or stuck writing a thesis with, I cannot believe how lucky I am to have you. To Bagheera and Zuko, thank you for your continued attempts to contribute to this thesis, even if one of those contributions caused my computer to freeze.

## ABSTRACT

The population of galaxies we see today is the result of billions of years of gas inflows, outflows, mergers, and feedback. To develop any holistic picture of the origin and evolution of galaxies, we thus need to understand their environments. The circumgalactic and intergalactic media (CGM and IGM) - the gas around and between galaxies, respectively - represent a large part of this environment. However, this gas is extremely faint and thus difficult to observe, and only recently have we been able to image it directly. This thesis presents instrumental and observational work focused on revealing galaxy environments in the early universe.

Chapter 1 presents a brief history of our understanding of galaxies and an overview of our current picture of galaxy formation, including the role played by galaxy environments. In particular, it focuses on presenting the evolution of baryonic structures within a cosmological density field dominated by dark matter.

Chapter 2 presents instrumental work on the Keck Cosmic Web Imager (KCWI -Morrissey et al., 2018), a new integral field spectrograph (IFS) for the Keck-2 10m telescope designed to study faint, extended emission. As an introduction, I discuss the advantages and disadvantages of integral field spectroscopy for the application of studying galaxy environments, as well as an overview of the prototype instrument - the Palomar Cosmic Web Imager (PCWI - Matuszewski et al., 2010). This chapter focuses primarily on engineering work during the development and testing of KCWI, though I conclude with a brief comparison of PCWI and KCWI performance in measuring the CGM around a high-redshift QSO.

Chapter 3 presents the development of a software package designed to extract and analyze faint, extended emission in PCWI and KCWI data: CWITools. Although software is often an afterthought in astronomical and observational work, it is likely to become a primary barrier to conducting large IFS surveys of the CGM and IGM. This semi-automated analysis pipeline is presented and released publicly to empower future PCWI and KCWI studies.

Chapter 4 presents the FLASHES (Fluorescent Lyman- $\alpha$  Structures in High-z Environments) pilot survey, published as O'Sullivan et al., 2020. The FLASHES pilot survey is an IFS study of extended HI Lyman- $\alpha$  emission in the environments of 48 z = 2.3 - 3.0 QSOs. The FLASHES Survey is the core project of this thesis, enabled by the instrumentation in Chapter 2 and the analysis pipeline developed in

Chapter 3. The pilot survey represents the first statistically significant ( $N \gtrsim 30$ ) sample of direct CGM observations in its redshift range. As such, it provides the first direct constraints on the 2D morphology, surface brightness profiles, and spatially resolved kinematics of the CGM during this period.

Chapter 5 presents the first FLASHES follow-up study; deep IFS observations targeting extended Ly $\alpha$  1216Å, NV 1240Å, CIV 1549Å, and HeII 1640Å emission from a subset of FLASHES pilot targets (O'Sullivan et al., in prep). Emission from metals in the CGM is expected to be an order of magnitude or more fainter than its Ly $\alpha$  counterpart, yet is a crucial ingredient in understanding the composition of the gas. Detecting this emission still requires multiple hours on 10m class telescopes. As such, large surveys of the multi-phase CGM remain extremely difficult to conduct. In this chapter, I present detections and upper limits of CGM metal emission around 8 FLASHES targets.

Chapter 6 presents engineering work on FIREBall-2 (the Faint Intergalactic Redshifted Emission Balloon, second generation), a high-altitude UV telescope and IFS targeting CGM emission in the low-redshift universe ( $z \approx 0.7$ ). FIREBall-2 is an ambitious project deploying a novel, electron-multiplying CCD designed to achieve  $\geq 50\%$  quantum efficiency in the UV. This technology represents an order of magnitude increase in sensitivity from the microchannel plates used in the GALEX (Martin et al., 2005) space telescope. FIREBall-2 serves as both an observational project in its own right, studying the low-z CGM, and a pathfinder mission for future UV space missions.

Finally, in Chapter 7, I summarize the contributions from this thesis and present a brief outlook on a few topics related to observations of galaxy environments.

## PUBLISHED CONTENT AND CONTRIBUTIONS

- O'Sullivan, D. and Y. Chen (2020). "CWITools: A Python3 Package for the Analysis of Integral Field Spectroscopy Data." In: *in prep*.D.O.S. led the design and implementation of CWITools and wrote the manuscript.
- O'Sullivan, D., K. Hoadley, et al. (2020). "The FLASHES Survey II: The Multiphase CGM around 8 QSOs at  $z \sim 2.7$ ." In: *in prep*. D.O.S. designed and led the FLASHES Survey II, conducted observations, performed data analysis, and wrote the manuscript.
- O'Sullivan, D., D. C. Martin, et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 z ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.
  D.O.S. designed and led the FLASHES Survey I, conducted observations, performed data analysis, and wrote the manuscript.
- Hamden, E. et al. (May 2019). "FIREBall-2: Advancing TRL While Doing Proof-of-concept Astrophysics on a Suborbital Platform." In: *Proc. SPIE*. Vol. 10982. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 1098220. DOI: 10.1117/12.2518711. arXiv: 1905.00433.
  D.O.S. designed and built the telemetry / telecommand communications system for the 2017 and 2018 flights.
- Martin, D. C. et al. (July 2019). "Multi-filament Gas Inflows Fuelling Young Starforming Galaxies." In: *Nat. Astron.* 3, pp. 822–831. DOI: 10.1038/s41550-019-0791-2. arXiv: 1904.11465.

D.O.S. led the observations, performed data reduction and analysis, and assisted in writing the manuscript.

Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597. D.O.S. designed and built the arc-lamp unit for spectral calibration of the instrument, assisted in verification of the grating exchanger, nod-and-shuffle mask, and articulation stage mechanisms, as well as characterizing detector performance and spectral focus.

## TABLE OF CONTENTS

Acknowledgements	1
Abstract	l
Published Content and Contributions	
Table of Contents	
List of Illustrations	
List of Tables	/
Chapter I: Introduction	
1.1 A Human History of Galaxies	,
1.2 Galaxy Formation in the ACDM Paradigm	,
1.3 The Circumgalactic and Intergalactic Medium	,
1.4 The Co-Evolution of Galaxy and Environment	•
1.5 Observational Signatures of the CGM and IGM	)
1.6 Integral Field Spectroscopy as an Exploratory Tool	
1.7 Outline of Thesis	, )
Chapter II: The Keck Cosmic Web Imager	
2.1 Advantages and Limitations of Integral Field Spectroscopy 31	
2.2 The Architecture of the Cosmic Web Imagers	
2.3 From Palomar to Keck: KCWI Upgrades	
2.4 Spectral Calibration Unit Design	,
2.5 Instrument Verification	
2.6 On Sky Performance: PCWI and KCWI Comparison	
Chapter III: CWITools	,
3.1 Introduction 53	,
3.2. Package Architecture 54	_
3.3 Cosmic Web Imager Data Format	
3.4 Module: Reduction 58	
3.5 Module: Extraction 72	,
3.6 Module: Synthesis 79	,
3.7 Module: Modeling 81	
3.8 Module: Measurement 83	
3.9 Summary and Discussion 88	{
Chapter IV: The FLASHES Survey I: Pilot Study 92	,
$41  \text{Overview} \qquad \qquad$	,
4.1 Overview	,
4.2 Introduction	,
4.5 Survey Methodology $\dots$	
4.4 Obscivations and Anchiary Data	,
4.5 Data Reduction and Analysis	
4.0 Results	)
4./ Discussion	)

4.8 Conclusions
4.9 Acknowledgements
Chapter V: FLASHES II: The Multiphase CGM at $z = 2.3 - 3.0$
5.1 Introduction
5.2 Observations
5.3 Data Reduction and Analysis
5.4 Preliminary Results
5.5 Summary and Ongoing Work
Chapter VI: FIREBall-2 - The Faint Intergalactic Redshifted Emission Balloon 165
6.1 Constraining the Low Redshift CGM
6.2 The History and Architecture of FIREBall
6.3 Telemetry and Telecommand System Overview
6.4 CSBF Radio System and Interface
6.5 Ground Control Stations and Flight Devices
6.6 Multiplexing Layer
6.7 TC Process Layer
6.8 Integration and Testing
6.9 2018 Flight
6.10 Future Work
Chapter VII: Summary and Outlook
7.1 Summary
7.2 Outlook
7.3 Resolving Galaxy and Environment Simultaneously
Appendix A: Extended Photometry Data for FLASHES Targets

Х

# LIST OF ILLUSTRATIONS

Number		Page
1.1	The Milky Way, seen here in a long-exposure image, has been visible	
	to the naked eye throughout history	. 2
1.2	A wood engraving showing Thomas Wright's model of the Milky	
	Way. An observer near the star at the center, labelled A, would see	
	the plane of stars the same way we see the Milky Way - as a band	
	across the sky.	. 4
1.3	Left: the initial sketch of 'spiral nebula' M51 by William Parsons in	
	1845 based on his observations through his 72-inch telescope. Right:	
	a modern image of M51 taken with the Hubble Space Telescope	. 5
1.4	Snapshot from the Illustris TNG simulation (Nelson et al., 2019),	
	showing the large scale structure of dark matter in the background.	
	The panels on the left show zoom-ins of interacting galaxies. The	
	panel in the top right shows the simulated visible baryonic matter	
	while the two panels in the lower-right show zoom-ins on the visible	
	material in two massive halos.	. 9
1.5	A two-dimensional slice of the universe from the 2dF Survey (Colless,	
	1999). Each dot is a galaxy, shown in a plane of redshift (distance	
	from the center) and right-ascension (angle). The 2dF Survey pro-	
	vided one of the first direct observations of the large-scale structure	
	of the universe. The filamentary structure of the cosmic web can be	
	seen, yet mapping the intergalactic and circumgalactic gas on this	
	scale remains a distant goal.	. 10
1.6	Cartoon from Tumlinson, Peeples, and Werk, 2017 showing the differ-	
	ent components of the CGM, including accreting gas from the IGM.	
	The mix of pristine inflows from the IGM, outflows, recycled/re-	
	accreting gas from outflows, and diffuse gas in the halo make the	
	CGM a complex environment.	. 14

1.7	From Thoul and Weinberg, 1995: the cooling rate of primordial	
	gas (76% hydrogen, 24% helium, by mass). The long dashed lines	
	show the cooling due to recombination. The short dashed lines show	
	the cooling due to <i>collisional</i> excitation, and the dotted line repre-	
	sents collisional ionization. These mechanisms become important	
	at T $\sim 10^4$ K. The thin solid line, representing radiation from free-	
	free electron interactions (i.e. Bremsstrahlung) begins to dominate	
	at $T \gtrsim 10^{5.5}$ K	16
1.8	From Bertone and Schaye, 2012, 'How the diffuse universe cools.'	
	Dashed lines show contributions from emission lines and solid lines	
	show cooling from continuum emission. The colors indicate dif-	
	ferent portions of the electromagnetic spectrum. At large tempera-	
	tures (T $\gtrsim 10^6$ K), x-ray emission continuum from Bremsstrahlung	
	radiation dominates the cooling process. At lower temperatures	
	$10^4 \le T \le 10^{5.5}$ K, UV emission from both lines and continuum	
	dominate the cooling process. IR line emission becomes significant	
	at T ~ $10^3$ K	17
1.9	From Bertone, Aguirre, and Schaye, 2013, average emission per unit	
	volume from different atomic species in the diffuse universe across	
	redshifts $z = 0 - 8$	18
1.10	From Dijkstra, 2017. The top panel shows the probability of recom-	
	bination events producing a Ly $\alpha$ photon as a function of temperature	
	for both Case A and Case B recombination. The bottom panel shows	
	the recombination rate for each Case as a function of temperature. $\ . \ .$	20
1.11	From Bertone and Schaye, 2012, simulated emission levels of dif-	
	ferent UV emission lines in the $z = 2 - 5$ CGM and IGM. The	
	second column shows hydrogen $Ly\alpha$ under the assumption that the	
	gas becomes self-shielding.	23
2.1	The Keck Cosmic Web Imager (KCWI) being installed on the Nas-	
	myth platform in summer 2017 (Morrissey et al., 2018)	31

xii

- 2.2 An example PCWI image, cropped to 1/3 of the vertical/wavelength axis. The 2D spectra of the slices are arranged from left to right. Each 2D spectrum thus represents one 'pixel' in the spatial y-axis of the field of view. Within each 2D spectrum, the spatial x-axis (i.e. distance along a slice) runs horizontally, while the wavelength axis runs vertically. The quasar HS1700+6416 can be seen as a bright continuum source on slices 9-13, while emission lines from the night sky can be seen as horizontal bands in each slice.
- 2.3 The integral field unit of the Palomar Cosmic Web Imager, taken from Matuszewski et al., 2010. The left panel shows a zoomed-in view of the image slicer; a 24mm × 16mm block with 1mm reflective slices. It can be seen mounted in the bottom-right corner of the right panel. Each slice is at a slightly different angle, such that the reflected image is directed to one of the 24 pupil-array mirrors seen in the right panel. These mirrors reflect the slice images to a curved mirror which collimates the beam before it is passed through the grating. . . . . . 37

- the CAD model in SolidWorks. The only component that cannot be clearly seen in this view is the shutter, mounted on the hidden side of the front panel. Right: a photograph of the assembled unit in the lab. 41

xiii

34

Summarized figure from the verification report for Red Arm compat-2.7 ibility weight testing. Panel (a) shows an example of how the weight was applied to the instrument. Panel (b) shows a 'loaded vs. unloaded' subtraction of an arc-bars image, which is an arc lamp shone through a series of vertical slits, resulting in an image of 2D point sources (i.e. discrete spatially and spectrally). Panel (c) shows a similar subtraction, except with an arc-flat image instead of an arc-bars image. An offset can clearly be seen in both between the original 43 2.8 Repeatability tests for the KCWI Blue Exchanger (BEX) Mechanism. The top panel shows the change in the x and y centroids over the course of the test. Each point represents the average value measured from 5 images, each containing 2500 sources. The bottom panel shows the 44 same thing for the 1D full-width-at-half maximum of each source. . . 2.9 Stability tests for KCWI's Blue Exchanger Mechanism. Both panels show the standard deviations of the centroid position and fullwidth-at-half-maximum (FWHM) of sources in the images over a five minute period. The position appears to have a typical standard deviation of  $\sigma(X) \simeq \sigma(Y) \sim 0.03$  and the FWHM appears to have a typical standard deviation of 0.1px in both axes. 45 2.10 Stability of the flux within each source over five-minute periods. The black histogram shows the standard deviation in the flux values for all measured sources as a percentage of the initial flux. The red histogram shows the expected distribution from shot noise and a read noise of 5 electons/pixel, while the blue distribution shows the same for twice the read noise. 46 2.11 KCWI Articulation Stage verification test results. The top left panel shows the time taken, in seconds, as a function of the move size, in degrees. The time taken at the required speed (1 deg/s) is shown as a dashed red line. The top right panel shows the net speed (deg/s) for a move versus the move size, with the minimum speed shown as a blue dashed line. The bottom panel shows the RMS deviation of the stage position, in arcseconds, as a function of position, in degrees. 47

xiv

2.12	The three types of test images taken for the NAS mask mechanism verification. In the top row, panel (a) shows the 'open' position	
	where the NAS mask is fully removed Panel (b) shows the 'test'	
	position with the mask half inserted. Panel (c) shows the 'dark'	
	position, with the mask almost fully inserted, such that the left edge	
	falls on the first slice.	48
2.13	Results of the KCWI NAS mask mechanism test. The top panel shows	
	the mask edge position across the iterations of removing/inserting the	
	mask into the 'test' position. The second panel shows the slope of	
	the mask edge for the same images. The bottom two panels are the	
	same as the top two, except applied to the 'dark' position.	49
2.14	Photon transfer curves for the four KCWI amplifiers (UL = upper-	
	left, UR = upper-right, etc.) in $1 \times 1$ binning, slow readout mode	
	at high gain. The horizontal dashed line represents the read noise	
	component in each panel, while the sloped dashed line represents	
	shot noise. Inset are the values for the read noise $(\sigma)$ and gain $(k)$ .	50
2.15	Comparison of a 2hr PCWI observation (1hr source and 1hr sky in	
	NAS mode) with a 1hr KCWI observation (standard observing mode).	51
3.1	The architecture of CWITools, showing the broad pipeline from stan-	
	dard DRP results to scientific results, as well as the nature of the	
	input and output of each module. Arrows represent different modules	
	within the package and boxes represent the different types of data.	
	For example, at the top, the reduction module (green arrow) takes	
	standard DRP data cubes as input, and outputs corrected/coadded	
	data cubes. On the right hand side are specific examples of the kinds	
	of data products at each stage	54
3.2	Automatic cropping parameters obtained by CWITools. This view is	
	presented to the user if automatic cropping is requested. Each panel	
	presents a summed one-dimensional profile for a different axis $(x, y, y)$	
	z from top to bottom). Data within the cropped range, delineated by	
	vertical red lines, is highlighted in blue, while data outside the range	
	is black. This is most useful as a first step, from which the user can	
	determine the best cropping parameters to suit their needs	59

XV

Automatic WCS correction using source fitting. As one option for 3.3 spatial WCS correction, CWITools assumes that the initial WCS is approximately correct, then identifies and fits the nearest source to that location. The above view is the view presented to the user during this step if requested, so that the user can inspect the fit visually. The top left panel shows the full field of view, with a white cross indicating the expected location of the primary source. The white box indicates the search area, which can be adjusted by the user. The red cross indicates the fitted location of the source. The top right panel shows a zoom in on the white box. The bottom two panels show the one-dimensional PSF of the source along each axis and a 3.4 Cropping and WCS correction applied to an individual cube. Left (a): a spatial 2D snapshot of the data cube prior to cropping and WCS correction. The red circle has the correct coordinates for the source, SDSS0958+4703. Right (b): the cropped and WCS-corrected cube, Spatial 2D slices of coadded frames using different drizzle factors. 3.5 In panel (a) on the left, the coadd was performed normally, with no drizzling (i.e.  $f_{drz} = 1.0$ ), while in panel (b), a very low Drizzle factor of  $f_{drz} = 0.4$  was used. Normally, one would use a factor in the range  $f_{drz} \sim 0.6 - 0.8$ , but we use a low factor here to make the visual difference clear. The image in panel (b) is clearly sharper with higher frequency noise, while that in panel (a) looks smooth in comparison. 65 Illustration of the exposure and pixel coverage thresholds in CWITools' 3.6 coadd function. Panel (a) shows an illustration of two overlapping fields of view with equal exposure time. The numbers in each area represent the local stacked exposure time relative to the maximum stacked exposure time. These are the values considered when applying the exposure threshold. Panel (b) illustrates the pixel coverage threshold. The white grid represents the coadd pixel grid, while the red rectangle represents the footprint of an input pixel. . . . . . . . . 68

xvi

- Variance scaling by assertion of a standard normal distribution in 3.7 'background' regions. The grey shaded histogram shows the distribution of SNR values based on the input data and variance. Note the logarithmic scale on the y-axis. The green histogram shows the distribution of SNR values after large, contiguous 3D objects (either systematic residuals or real emission regions) have been detected and masked. The dashed black line shows a standard normal distribution, and the solid black line shows the best-fit Gaussian model used to calculate the re-scaling factor. 70
- Calibration of covariance in a single KCWI data cube. The panel 3.8 on the left shows the ratio between the observed noise and the propagated noise assuming no covariance ( $\sigma_{\rm obs}$  and  $\sigma_{\rm ideal}$ ) - after the binning the data - as a function of bin size. The red line shows the best-fit two-component model, with  $\sigma_{obs}/\sigma_{ideal} = C(1 + \alpha \log(K))$ for  $K \leq K_{thresh}$  and  $\sigma_{obs}/\sigma_{ideal} = \beta$  for  $K > K_{thresh}$ . The right panel
- PSF-subtraction of a bright source to isolate extended emission. The 3.9 top panel (a) shows a section of the 2D spectrum of a slice containing the bright source. The bottom panel shows the same 2D spectrum, with the same color map, after PSF subtraction. Bright, extended Ly $\alpha$ emission can be seen clearly after the subtraction. The small, bright spots to the left (blue) side of the extended emission are systematic residuals caused by sharp  $Ly\alpha$  forest absorption features, where the PSF shape changes rapidly as a function of wavelength. The central pixels used to scale the PSF are masked, as these cannot be used to measure signal. 74
- 3.10 Integrated spectrum of a region before (top) and after (bottom) PSF and polynomial background subtraction. Black vertical bands indicate the locations of noisy residuals from extremely sharp absorption features or bright sky line, both of which can be masked. Vertical red lines, from left to right, indicate the positions of emission lines: HI Lyα λ1216, NV λλ1239, 1243, SiIV λ1394, CIV λλ1548, 1551, HeII  $\lambda 1640$ . 75 z-Moment maps, shown in units of km  $^{-1}$ , calculated using CWITools. 3.11 From left to right: first moment, propagated error on the first moment,

xvii

3.12	Left: an illustration of the three characteristic radii provided by the	
	measurements module. From largest to smallest, they are: (i) the	
	maximum radius ( $R_{max}$ - dotted circle), (ii) the effective radius ( $R_{eff}$	
	- dashed circle), and (iii) the flux-weighted RMS radius ( $R_{rms}$ - solid	
	circle). Right: a radial profile generated by the synthesis module for	
	the same nebula. The same three radii are shown on this axis	85
4.1	IFS surveys of extended emission around high redshift galaxies. Sur-	
	veys are shown as stacked histograms representing the number of	
	targets in each.	94
4.2	The FLASHES pilot sample in redshift $(z)$ vs. absolute i-band mag-	
	nitude $(M_i)$ . Circles indicate targets for which the value of $M_i$ is	
	estimated from the given apparent magnitude, while diamonds indi-	
	cate those for which a value of $M_i$ was provided in the SDSS DR12Q.	
	The colorbar indicates the WISE infrared W2-W3 (4.6 – $11.6\mu$ m)	
	color. The blue histograms represent the distributions of $z$ and $M_i$ in	
	the SDSS DR12Q in the same redshift range	95
4.3	PSF subtraction performed for off-center (i.e. continuum) wave-	
	lengths for the first five targets. The columns, from left to right,	
	show: the scaled white-light image, the off-center pNB image, the	
	subtracted off-center pNB image, the Ly $\alpha$ pNB image (for compari-	
	son), and the subtracted Ly $\alpha$ pNB image	07
4.4	Calibration of variance measurement in FLASHES pilot data. Black	
	crosses indicate individual calibration measurements. The solid black	
	curve indicates the averaged profile, while the grey-shaded region	
	represents the $\pm 1\sigma$ uncertainty. The solid red curve indicates the	
	functional fit to $\sigma_{meas}/\sigma_{nocov} = (1 + \alpha_v Log(N_k))$ , with $\alpha_v = 0.79$ ,	
	and the horizontal dashed red line indicates the approximate asymp-	
	tote for the relationship at large $N_k$ ( $\beta_v \simeq 2.6$ )	08
4.5	Example of the optimization of pNB image parameters (wavelength	
	center and bandpass) for target HS1700+6416 (ID 1). The colormap	
	in both panels shows the integrated SNR in a circular region of radius	
	75pkpc around the QSO. The left panel shows the initial coarse	
	grid, which searches a wide velocity range of $\pm 10000$ km s <sup>-1</sup> . The	
	red lines indicate the $\pm 1000$ km s $^{-1}$ zoom-in for the high resolution	
	grid, shown on the right.	10

- 4.6 FLASHES Pilot survey observations (ID 1-6). Each tile is  $250 \times$  $250 \text{ pkpc}^2$  in size, centered on the QSO. The left-most four columns show a white-light image,  $Ly\alpha$  surface brightness, velocity, and dispersion. Surface brightness is in cgs units, erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>. The black bar in the top white-light image shows 100 pkpc and the black square shows the box kernel used to smooth the WL and pNB data. Foreground sources in each field have been masked, with the masked regions shown in black. The rightmost column shows integrated nebular spectra (black) and scaled QSO spectra (grey). The spectra are summed over the object masks and shown in units of  $10^{-17}$  erg s<sup>-1</sup> cm<sup>-2</sup> Å<sup>-1</sup>. Spectra are shown rest-frame wavelength, according to the systemic QSO redshift. Blue lines indicate the fluxweighted centers of nebular emission, while black lines indicate the peak of QSO emission. (Caption continues in next Figure) . . . . . . 114 (continued) A very bright mercury sky emission line (Hg  $\lambda$ 4358.3) 4.7 is masked in some spectra and shown here as a vertical red band wherever it appears. Empty regions (shown in white) in the pNB images are outside the field of view. An ellipse representing the FWHM of the QSO's PSF is shown in each tile. Red ellipses are used for smoothed PSF (in the WL and pNB images) while black ellipses 4.8 4.9 4.10 4.11

4.14	Circularly averaged radial profiles of the detected CGM, centered on
	the QSOs. The top panel shows three examples: a bright detection
	(ID2), an intermediate detection (ID 22) and a non-detection (ID 48).
	The middle panel shows the averaged profiles in observed surface
	brightness, with a Sérsic fit and a power-law fit. The bottom panel
	shows the average of the profiles after scaling each by $(1 + z)^4$ to
	correct for cosmological surface brightness dimming, with the same
	fits. The x-axis is shown in log-scale
4.15	Covering fraction of $SNR \ge 2\sigma$ as a function of radius. The average
	for all fields in the sample is shown with black diamond markers,
	while the average of all detections is shown with blue circle markers. 129
4.16	Top panel: global dispersions of the detected nebulae as measured
	from a single-component Gaussian fit. Bottom panel: dispersions of
	individual features when fitting spectra with a composite model of
	1-4 Gaussian components
4.17	Distributions of CGM Ly $\alpha$ velocity offsets with respect to different
	redshifts. The top panel shows velocity with respect to the best-fit
	SDSS/DR12Q QSO redshift. The middle panel shows velocity with
	respect to the peak of Ly $\alpha$ emission in the QSO spectrum. The
	bottom panel shows velocity offset with respect to the HeII $\lambda 1640$
	redshift from SDSS
4.18	Stacked Ly $\alpha$ profiles of the CGM detections in the FLASHES pilot
	survey. Different colors indicate different redshifts used to convert
	from observed to rest-frame wavelengths: the redshift of the CGM
	Ly $\alpha$ emission itself, the redshift of the peak of Ly $\alpha$ emission in the
	QSO (blue), the QSO's systemic redshift from DR12Q (green), and
	the HeII $\lambda$ 1640 redshift from SDSS (red)
4.19	Comparison of $L_{Ly\alpha}/R_{max}^2$ as a function of $R_{max}$ for different surveys.
	The top panel shows the comparison for sizes in proper kiloparsecs,
	while the bottom panel shows the same comparison for comoving
	kiloparsecs. The quantity $L_{Ly\alpha}/R_{max}^2$ should depend only on the
	intrinsic radial surface brightness profile of the emission, so com-
	paring nebula of equal size under this metric provides an equitable
	comparison of the average surface brightness of <i>detected</i> regions 134

4.20	Cumulative distribution of nebular eccentricities for the FLASHES
	survey, A19, and C19. The solid black line shows the distribution as
	measured from the FLASHES $1\sigma$ contour object masks. The dashed
	black line shows those measured from $2\sigma$ contour object masks. The
	red line shows the distribution of values presented in A19. The green
	dashed line shows the distribution from C19
4.21	Change in eccentricity as a function of the increase in limiting surface
	brightness. Contours show a Gaussian kernel density estimate and
	the black line with shaded region shows the best-fit linear model
	with $\pm 2\sigma$ slope uncertainty. The linear regression shows a strong
	correlation in which eccentricity increases as the surface brightness
	threshold increases
4.22	Parameter space plots for the nebulae detected in the FLASHES
	Sample. $R_{eff} = \sqrt{Area/\pi}$ is the effective size, e is the eccentricity
	$(0 \le e < 1)$ , $z_{Ly\alpha}$ is the redshift of the nebular Ly $\alpha$ emission, $\Delta v_{QSO}$
	and $\Delta v_{peak}$ are the velocity offsets with respect to the systemic red-
	shift and peak of QSO Ly $\alpha$ emission, respectively, $\sigma_G$ is the standard
	deviation (in km $s^{-1}$ ) of the best-fit single-peaked Gaussian line pro-
	file, and $M_i$ is the absolute i-band magnitude of the quasar. Contours
	in each plot show Gaussian Kernel Density Estimates of the 2D dis-
	tribution. Black and red lines show linear regression models with
	p < 0.05 and $p < 0.01$ , respectively. The r-values of these linear
	regressions are shown on the relevant tile. Faint grey lines indicate
	linear regression models with $p \ge 0.05$ (i.e., no correlation clearly
	indicated)
4.23	Measured asymmetry (i.e. the y:x aspect ratio, determined by a 1D
	Gaussian fit to the collapsed PSF along each axis) of the PSF in the
	final pNB images, shown in Figure 4.6

- 4.24 A comparison of the RMS line-of-sight velocity detected in the FLASHES pilot survey with gravitational motions in an NFW halo. The dashed blue line and blue shaded region represent the average and  $\pm 1\sigma$  spread in the line-of-sight RMS velocities of FLASHES pilot nebulae, respectively. The solid black curve shows the (maximum) RMS line-of-sight velocity of an NFW halo as a function of halo mass following Munari et al. (2013) ( $\sigma_{1D} = 0.68v_{200}$ , where  $v_{200}$  is the circular velocity at the virial radius). The grey shaded region indicates the halo masses of high-luminosity QSOs (HLQSOs) at a redshift of z = 2.7 in Trainor and Steidel (2012).
- 4.25 Top: log likelihood that the flat model is more appropriate than a sheared model,  $\text{Log}_{10}(\text{p}_{\text{const}})$ . Bottom: average projected specific angular momentum,  $\langle \vec{j} \rangle_f$ , versus effective radius,  $\text{R}_{\text{eff}}$ . The size of the circular markers is proportional to  $-\text{Log}_{10}(\text{p}_{\text{const}})$  (i.e. the relative likelihood of the shear model). Each target is also shown as a red square of fixed size, for clarity. The horizontal line represents the minimum resolvable angular momentum discussed in the text. . . 144
- 5.1 Variance and covariance calibration for FLASHES deep data. Left: the initial, global scaling of the variance cube. The grey shaded region shows the histogram of SNR values in the background regions of the cube. The green region shows the subset of values not associated with large systematics. The dashed black line shows a normal distribution with a standard deviation of  $\sigma = 1$ , while the solid black line (overlapping with the dashed line) shows the best-fit model to the green data. Right: the ratio between measured and propagated noise after binning with box kernels of area K. The histogram on the right shows the error on the model, which falls comfortably within a Luminosities and upper limits on luminosity for Ly $\alpha$ , NV, CIV, and 5.2 HeII for the eight targets in the sample. Error bars shown as  $\pm 2\sigma$ . . 160 Pseudo-narrowband images and  $F_{\lambda}^{R50}$  spectra for the first four deep 5.3 targets (IDs 1, 2, 5 and 7). Each pseudo-narrowband is a  $100 \times 100$

5.5	Circularly averaged radial surface brightness profiles for each targeted
	emission line. Surface brightness units, labelled [cgs] to fit on the
	plot, are the standard erg s <sup>-1</sup> cm <sup>-2</sup> arcsec <sup>-2</sup>
5.6	Velocity maps for the two targets with extended emission detected
	for all four emission lines
5.7	Dispersion maps for the two targets with extended emission detected
	for all four emission lines
6.1	A view of FIREBall-2 during its 2018 flight, taken by photographer
	Mouser Williams
6.2	Schematics of the optical design of FIREBall-1 and FIREBall2.
	FIREBall-2 added a field corrector, changed the spectrograph de-
	sign from a 'folded Offner spectrograph' to an aspherized Schmidt
	spectrograph, and changed the detector from a micro-channel plate
	to a delta-doped, electron-multiplying CCD
6.3	The broader communications system context for FIREBall-2, includ-
	ing the CSBF-owned downrange station and plane, used as backups
	in the case of loss of communications. Boxes show components of
	the system, while arrows show communications connections $170$
6.4	Simple schematic of the different communication layers in the TM/TC
	system. The left-hand side shows the ground-side system while the
	right-hand side shows the flight-side system
6.5	The J2 Connector on the CIP. The columns are labelled from right
	to left with letters, while the rows are labelled with numbers. Every
	other column has only even or odd rows, such that column 'A' has
	pins A0, A2, A4 (etc.) and column B has pins B1, B3, B5 174
6.6	The back of the SR-04 multiplexer used to combine the FIREBall-
	2 TM/TC channels (during a test in 2017- the blue cable was not
	used for flight, for obvious reasons). From left to right, the visible
	interfaces are: (i) the power supply; (ii) the network management
	port, used to configure the device; (iii) the network composite port,
	which carries the multiplexed signal, and (iv) the four individual 'data
	ports' for the channels that are to be multiplexed

6.7	Representation of the ground-side TC process as a DFA. The letter	
	'p' is the set of hexadecimal characters reserved for port identifier	
	characters, and the letter 'W' denotes the MUX package header/footer	
	character. Red lines indicate state transitions associated with errors.	
	The text below each state describes it in plain terms	30
6.8	Representation of the flight-side TC process as a DFA. The letter 'p'	
	is the set of hexadecimal characters reserved for port identifiers, and	
	the letter 'W' denotes the MUX package header/footer character. Red	
	lines indicate state transitions associated with errors. The text beside	
	each end state describes it in plain terms.	32
6.9	Flight-side TC process hardware (not including flight-side multi-	
	plexer). The DB9 connectors labelled 'TX' and 'RX' connect to the	
	transmitting and receiving lines of the flight-side MUX composite	
	port. The small chips in center frame are RS232-TTL converters.	
	The microprocessor with a red shield in the top right is the Rugge-	
	duino which handles uplinked datawords. The Ruggeduino in the	
	bottom right, programmed by CNES team members, handles dis-	
	crete command uplinks	33
6.10	A more detailed overview of the TM/TC subsystem, showing the	
	main devices, connections, and speeds. The boxes labelled 'TTL' are	
	the TTL-RS232 converter chips. Yellow boxes show the baud rates	
	of each connection	35
6.11	Schematic of the environmental testing setup in Palestine, TX. The	
	dashed line represents the boundary of the vacuum chamber. The grey	
	boxes shows the CSBF equipment. The parallel data-word (DW) and	
	HH connections on the CIP are highlighted. The blue boxes highlight	
	the new hardware, not tested during the 2009 flight	36
6.12	Pictures of the TM/TC equipment environmental testing setup in	
	Palestine, TX. The top panel shows the outside of the vacuum cham-	
	ber. The bottom panel shows the multipler (right), the CIP (top-left),	
	and the Ruggeduino connected inside the chamber	38

xxiv

# LIST OF TABLES

Number	r	Page
2.1	Sequence of weight tests performed for KCWI red channel compati-	
	bility testing. The right-most two columns show the average offsets	
	for each test along each image axis in units of pixels	. 43
2.2	Readout times for KCWI measured during detector verification, in	
	summarized form. $t_r^{quad}$ is the readout time using all four amplifiers,	
	$t_r^{dual}$ is the readout time using two amplifiers, and $t_r^{single}$ is the readout	
	time using a single amplifier	. 51
3.1	A summary of FITS Header keywords for world coordinate systems.	
	Correcting or adjusting a world-coordinate system is most easily done	
	by updating these values, as opposed to shifting or rotating the data	
	itself	. 60
4.1	Summary of FLASHES Pilot Observations	. 97
4.2	Targets included from Trainor and Steidel (2012)	. 100
4.3	pNB parameters for the FLASHES Pilot Survey	. 104
4.4	FLASHES Pilot Survey CGM properties	. 124
4.5	Distributions of Measured Sizes	. 127
4.6	Properties of Stacked Ly $\alpha$ Profiles	. 132
4.7	Comparison of Eccentricity Distributions	. 139
5.1	Summary of the observational data for ten FLASHES targets. For	
	consistency with O'Sullivan et al., 2020, the IDs in the left-most	
	column are the original target IDs from the FLASHES pilot survey.	
	The effective radius ( $R_{eff} \equiv \sqrt{Area/\pi}$ ) size and luminosity of Ly $\alpha$	
	measured in the pilot survey are presented in the fourth and fifth	
	columns. The column labelled 'Slicer' indicates the image slicer	
	(M=Medium, L=Large) used for the observations. All observations	
	used the low-resolution grating (BL).	. 153
5.2	Integrated luminosities and upper limits for Ly $\alpha$ , CIV, NV, and HeII	
	emission around the 8 deep targets. Values are given as $L_{42} \equiv (L/10^{43})$	erg $s^{-1}$ ]).
	For non-detections, upper limits $(2\sigma)$ are derived from the propagated	
	uncertainty on the luminosity measurement.	. 156

5.3	Ratios of CIV, NV, and HeII to Ly $\alpha$ in terms of integrated luminosity.
	Upper limits (2 $\sigma$ ) are provided where there is a Ly $\alpha$ detection but no
	detection of the other line. No ratios or upper limits are provided for
	target 39, which has no Ly $\alpha$ detection
5.4	Sizes of 3D detections for each target and line, in units of pkpc. $R_{max}$
	and $R_{rms}$ are the maximum and flux-weighted root-mean-square radial
	extent of the nebular emission with respect to the QSO
6.1	An illustration of the uplink process. The MUX receives two con-
	secutive messages from Port 1 ('Hello' and 'Hi') followed by a
	message from Port 3 ('Bye'). Each MUX packet is bounded by
	the header/footer characters '{},' within which the first byte spec-
	ifies the transmitting port. The central column shows the data-
	words/telecommands transmitted by the uplink script after minimiz-
	ing overhead, while the right-hand column shows the response of the
	flight-side TC process upon receiving each dataword
6.2	The itinerary developed for integration and testing of the communica-
	tions system in Fort Sumner, NM, for the 2018 flight of FIREBall-2.
	A spreadsheet containing test details and inventories was attached to
	each step, to that ensure the full system had been thoroughly verified. 189
A.1	Multi-band Photometric Data (AB Magnitudes) for the FLASHES
	Pilot Sample         203

Chapter 1

## INTRODUCTION



"What are galaxies, and why do they have the shapes and sizes we observe?" - Blumenthal et al. 1984

The above quote is a favorite of mine because of its breadth and simple curiosity. It is easy to get so bogged down in the details of a particular problem that bigger questions seem to get lost in the distance. The introduction of a PhD thesis seems an ideal place to step back and think about the bigger questions. What are galaxies? When and how did they form? What drives their evolution?

These - the broadest questions - are ones that we ultimately answer collectively. Any holistic theory of galaxy formation and evolution is the emergent result of many smaller questions answered over time. The focus of this thesis is on the role played by the gaseous environments of galaxies. In particular, it focused on observations of the extremely faint gas in galaxy environments and the instrumentation required to conduct them.

In this introductory chapter, I will outline the bigger picture surrounding this thesis: our understanding of galaxies, their environments, and their co-evolution. However, the concept of 'galaxies' (plural) is quite a modern one and that of 'galaxy environments' even more so. It therefore seems fitting to establish not just the modern scientific context, but the historical context as well. As such, I hope the reader will indulge me in a brief detour through a history, not of galaxies themselves, but of our understanding of them.



Figure 1.1: The Milky Way, seen here in a long-exposure image, has been visible to the naked eye throughout history.

## 1.1 A Human History of Galaxies

## **Ancient Times**

The nature of the Milky Way has captivated people since ancient times. Of course, the earliest explanations follow mythological, rather than scientific, themes. The English term 'Milky Way' is a direct translation of the Greek galaxias kyklos which is also the origin of the modern English word 'galaxy.' In Greek mythology, the band across the sky was the result of milk that was spilled as a goddess fed the baby Heracles. In Ancient China, the Milky Way was referred to as the 'Silver River' (which is still the literal translation of the modern Chinese term). The Silver River was drawn by a goddess using a hair-pin, to separate Earth and Heaven so that a lower goddess could not travel there to meet a lover. The Cherokees in North America called it gi li' ut sun stan un yi,' which translates to 'the place where the dog ran.' This comes from a legend in which a giant spirit dog stole cornmeal from a village, was scared away by the villagers, and spilled the cornmeal across the sky as he ran. There are many different names for the Milky Way in Irish mythology, but one of the most interesting is Scríob Chlann Uisneach, or 'the track of the children of Uisneach.' In this legend, the Milky Way connects the graves of two lovers who a bitter enemy had tried to bury as far apart as possible.

Interestingly, it was suggested by ancient philosophers Anaxagoras and Democritus in the 4th-5th Century BC that the Milky Way was made up of stars. However, they were correct for the wrong reasons and flaws in their logic were later pointed out by Aristotle. As such, this idea did not become a widely held view and would not become an accepted explanation until approximately two thousand years later.

From antiquity to pre-modern times, the Milky Way and the stars were commonly considered fixed objects on a great celestial sphere. Often, this sphere represented the heavenly realms or a divide between Earthly and heavenly realms. The state of divinity associated with such celestial objects placed them outside the reach of us mortals and outside the physical rules which govern us. It will not be until the enlightenment era and the scientific revolution that this conceptual boundary is erased and the same physical laws that govern Earthly objects are applied to the Milky Way.

## **Medieval Times**

Although the divine status of the Milky Way remained intact until the scientific revolution, there are several notable and prescient claims from Middle Eastern and Muslim astronomers during the Islamic golden age (i.e. 9th - 13th century). The astronomer Abu Rayhan al-Biruni claimed that the Milky Way was "collection of countless fragments of the nature of nebulous stars." Ibn Bâjja, born in Spain and also known by the name Avempace, considered the Milky Way to be made up of the light of many tightly-packed stars. Nasir al-Din al-Tusi similarly claimed that the Milky Way was made up of many tightly-packed stars, giving a cloudy appearance overall.

It seems clear that the idea of the Milky Way being made up of stars had re-emerged by this stage, but it was not an easy claim to verify without sufficiently powerful telescopes. As we will see momentarily, many of the main paradigm shifts towards the modern picture of galaxies were the results of new observational capabilities; resolving stars within the Milky Way, resolving the proper motion of stars, revealing faint nebulae and external galaxies - these new observations drove the debates that led us to where we are today.

### **Pre-Modern Times**

In 1610, Galileo Galilei confirmed observationally that the Milky Way was indeed made of individual stars using one of his telescopes. In 1718, Edmund Halley discovered the proper motion of three bright stars. The realization that the stars were not fixed objects on a heavenly sphere was profound. It was a major step towards bringing the stars down from a status of distant, quasi-divine objects to yet another part of a universe governed by physical laws. This discovery also marks the beginning of a transition away from the ancient paradigm in which the Milky



Figure 1.2: A wood engraving showing Thomas Wright's model of the Milky Way. An observer near the star at the center, labelled A, would see the plane of stars the same way we see the Milky Way - as a band across the sky.

Way was just a heavenly backdrop, towards one in which the Milky Way is an assembly of stars, gas, and clouds in our local environment. Not long after, Halley also measured the Earth-Sun distance, proving it to be an order of magnitude larger than previously expected. The trend of new measurements yielding unimaginable distances would be repeated many times over the next two centuries, pushing the conceptual boundary of the universe first to include the entire Milky Way as a system of stars, and then far beyond.

The idea of the Milky Way as a large system of stars orbiting a common center of gravity first came from Thomas Wright, inspired both by Halley's observations of proper motion and a religious belief in a 'Divine Center' of the universe. Wright's model, shown in Figure 1.2, consisted of a spherical shell of stars within which we were embedded. When flaws were pointed out in this model's predicted view of the night sky, he revised it so that the shell is either exceedingly thin (locally approximated as a plane) or that the morphology of the Milky Way had an annular



Figure 1.3: Left: the initial sketch of 'spiral nebula' M51 by William Parsons in 1845 based on his observations through his 72-inch telescope. Right: a modern image of M51 taken with the Hubble Space Telescope.

shape (with the 'Divine Center' inhabiting the empty central region) rather than a spherical one. He also suggested that some nebulae were distant systems with their own divine centers.

It was Immanuel Kant, inspired by reading Wright's work as well as work by Emanuel Swedenborg, who later suggested our current model of the Milky Way: a flat, rotating disk of stars. Kant also proposed the nebular hypothesis for the formation of stars and planets and advanced Wright's notion in claiming that some nebulae may be distant galaxies themselves, similar to the Milky Way.

The debate over nebulae (in particular, nebulae later resolved to have a spiral structure) raged throughout the late 1700s, 1800s, and early 1900s. Over this time, the term 'Island Universe' came to be used in reference to galaxies, and the prevailing notion of the Milky Way was as such: a self-contained system following Kant's disk model.

## **Early-Mid 20th Century**

In 1845, William Parsons discovered that Messier 51 had a spiral structure using his six-foot reflecting telescope "the Leviathan of Parsontown." Figure 1.3 shows his sketch of the so-called 'spiral nebula' beside a modern image of M51. The discovery of these nebulae with spiral structures triggered an intense debate, culminating in the famous 'Great Debate' between Heber Curtis and Harlow Shapley. After observing a nova within Andromeda, Curtis studied records for a total of 12 such novae and found them to be ten magnitudes fainter than those recorded on the Milky Way, on average. The rate of novae within the Andromeda 'nebula' was also highly anomalous compared to the rest of the Milky Way. He noted the dark lanes within

Andromeda, comparing them to dust lanes within the Milky Way, and concluded that Kant was correct: Andromeda was itself a distant galaxy.

Shapley argued that the implied distance to Andromeda was simply too large to be believable, and was supported by a much more convincing argument from Adriaan van Maanen. Van Maanen, based on his own observations, estimated that a full rotation of the Pinwheel Galaxy (Messier 101) happened on a timescale of just years. If true, the spatial scales implied by the extragalactic argument would mean that the Pinwheel galaxy rotated faster than the speed of light. Curtis conceded this point on the assumption that van Maanen's observations were correct. Of course, it would turn out that they were not. Edwin Hubble settled the debate using the 2.5 m Hooker telescope at Mount Wilson. He was able to resolve individual stars within Andromeda, which enabled him to obtain a distance measurement using Cepheid variables. After this, the debate was largely settled. The Milky Way was just one galaxy and there were many more beyond.

#### Mid-Late 20th Century

In the mid-to-late 1900s, two new areas of study brought galaxy environments into focus. The first was the discovery that the mass budget of the universe is dominated by a mysterious form of invisible matter called dark matter, and that this dark matter therefore determines the gravitational environment within which galaxies form and evolve. The second was the growing field of observations of the intergalactic and circumgalactic media: the reservoirs of baryonic material which supply fuel for galaxy growth. Far from the image of isolation conjured up by the term *Island Universes*, galaxies came to be seen as immersed in complex environments which shape their evolution.

Evidence that galaxies were 'missing' some mass came as early as the 1930s, when Jan Oort performed a re-analysis of stellar motions in the galactic plane observed by Jacobus Kapteyn. Oort concluded that there must be two 'dark stars' for every star observed to explain the motions. In 1938, Fritz Zwicky studied the motions of galaxies in the Coma cluster and found that the cluster had significantly more mass than was observable. Horace Babcok (1939) and Jan Oort (1940) also provided further evidence that there was some missing mass, based on observations of individual galaxies. However, it was not until the 1980s that dark matter became a widely studied problem.

Following work in the 1970s by Vera Rubin and Ken Ford (Rubin and Ford, 1970),

among others, dark matter became an established component in the discussion of galaxies and their environments. In the 1970s and early 1980s, it was hypothesized that dark matter could be made up of relativistic neutrinos (Cowsik and McClelland, 1973; Zeldovich, Einasto, and Shandarin, 1982). This was later termed 'hot' dark matter or HDM. However, it was later shown using numerical simulations that these energetic particles, which travel enormous distances before interacting, would smooth out density fluctuations on scales smaller than those of massive galaxy clusters (S. D. M. White, C. S. Frenk, and Davis, 1983). An alternative model in which dark matter in the early universe was made up of massive, slower-moving particles were presented around the same time (Peebles, 1982; Bond and Szalay, 1983). These were collectively referred to as 'cold' dark matter models, or CDM. A review article by Blumenthal et al., 1984 concluded that models and simulations with cold dark matter provided the strongest agreement with the galaxies we observe today. The CDM paradigm is still a standard part of our modern theory of galaxy formation.

While the exotic gravitational environments of galaxy environments were debated in the context of dark matter models, the baryonic environments were also starting to be revealed. Gunn and Peterson, 1965 constrained the density of neutral hydrogen in intergalactic space (the intergalactic medium or IGM) by studying absorption in the spectrum of a high-redshift quasar. A few years later, Lynds (1971) also studied absorption systems in a QSO spectrum to the same effect. A survey of galaxy redshifts by De Lapparent (1986) began to reveal the large-scale structure of the cosmic web, and - a year later - Hogan & Weimann would suggest using long-slit spectroscopy to image emission from the same systems which cause the Lymanalpha forest. These studies, alongside others, were the first in a new field focused on constraining the morphology, density, and other properties of the intergalactic medium. Theoretical work along these lines was further enabled by the advent of computers and numerical simulations (e.g Barnes et al., 1985), which provided an essential means of testing our theoretical models against observational data. Even more so today, numerical work forms the foundation of our understanding of the cosmic web, galaxy accretion, mergers, and other processes which are complex or chaotic enough that analytical models are insufficient. By the start of the 21st century, the role of galaxy environments and the 'diffuse universe' was understood to be one of the major puzzle pieces in the larger theory of galaxy formation.

Of particular interest in this discussion of galaxy environments was the mechanism

by which baryonic material from the IGM accretes onto dark matter halos. In the late 1970s and 1980s, as the dark matter paradigm was emerging, the classical picture of galaxy formation invoked a process typically called 'hot mode accretion' today (Rees and Ostriker, 1977; Silk, 1977; S. D. M. White and Rees, 1978; Fall and Efstathiou, 1980). In hot mode accretion, infalling gas forms a shock-front which heats it to extremely high temperatures ( $T \simeq 10^6$ K), exerting an outward pressure. This tenuous, hot gas takes a long time to cool down again and coalesce towards the center.

At the same time, there were suggestions that, under certain conditions, infalling gas may not be significantly heated at all (Binney, 1977). This would result in the deposition of large amounts of cool gas (ready to form stars) and the conservation of much of the angular momentum of infalling gas. This process is nicknamed 'cold-flow' or 'cold-mode' accretion. The discovery of galaxies rapidly forming stars at high redshift (Steidel et al., 1996) presented difficulties for a universal picture of slow, hot-mode accretion and highlighted the need for such an alternative. The prevalence of cold and hot mode accretion is still an active area of research today, and will be discussed some more in the following sections.

### **1.2 Galaxy Formation in the ACDM Paradigm**

Today, we understand much of what galaxies are and how they form. We know that a galaxy is a gravitationally bound collection of gas, dust. and stars in a gravitational well dominated by dark matter. This last detail, that the gravitational field is dominated by dark matter, is an important one. The current standard cosmology is one in which cold dark matter dominates the mass budget of the universe, and the overall mass-energy budget is dominated by a cosmological constant. This is referred to as the  $\Lambda$ CDM paradigm.

The latest measured cosmological parameters come from the ninth-year results of cosmic background microwave (CMB) observations by the Wilkinson Microwave Anisotropy Probe (WMAP9 - Hinshaw et al., 2013). These WMAP9 parameters are  $\Omega_b = 0.0463 \pm 0.0024$  (baryonic density),  $\Omega_c = 0.233 \pm 0.023$  (cold dark matter density), and  $\Omega_{\Lambda} = 0.721 \pm 0.025$  (cosmological constant / dark energy). This is the cosmology that is used throughout this thesis unless otherwise stated.

The large-scale distribution of dark matter determines the gravitational potential within which baryonic structures such as galaxies, clusters of galaxies, and gaseous filaments form. Figure 1.4 shows the large-scale distribution of dark matter in the



Figure 1.4: Snapshot from the Illustris TNG simulation (Nelson et al., 2019), showing the large scale structure of dark matter in the background. The panels on the left show zoom-ins of interacting galaxies. The panel in the top right shows the simulated visible baryonic matter while the two panels in the lower-right show zoom-ins on the visible material in two massive halos.

Illustris TNG simulation (Nelson et al., 2019). The filamentary structure, visible in simulations like this and in the real spatial distribution of galaxies (see Figure 1.5) is often referred to as 'the cosmic web' for its similarity to a spider's web.

On smaller scales, we refer to the gravitationally-bound, quasi-spherical distributions of dark matter that contain galaxies and, on larger scales, clusters of galaxies as dark matter 'halos.' The gravitational potential of galactic dark matter halos determines



Figure 1.5: A two-dimensional slice of the universe from the 2dF Survey (Colless, 1999). Each dot is a galaxy, shown in a plane of redshift (distance from the center) and right-ascension (angle). The 2dF Survey provided one of the first direct observations of the large-scale structure of the universe. The filamentary structure of the cosmic web can be seen, yet mapping the intergalactic and circumgalactic gas on this scale remains a distant goal.

the kinematics of the CGM and infalling IGM. As such, it is worth spending a little time discussing them.

Dark matter halos, and later the galaxies within them, form from the growth of small density perturbations in the early universe. The imprints of these early density fluctuations can be see as perturbations in temperature maps of the CMB. They are extremely small; the temperature of the CMB is  $T_{CMB} \simeq 2.725 \pm 0.0001$  K in every direction. However small, these over-dense regions exert a gravitational force on nearby matter. This results in a positive feedback loop: nearby material is pulled in, which increases the over-density, which increases the gravitational force on nearby material, leading to collapse.

Cold dark matter is collisionless, only interacting with itself and other matter gravitationally. As a result of this, dark matter structures collapse faster than their baryonic counterparts, for which internal pressure pushes outwards against gravity. To characterize the local overdensity and trace the growth of these collapsing dark matter structures, we use the parameter  $\delta \equiv (\rho - \bar{\rho})/\bar{\rho}$ , where  $\rho$  and  $\bar{\rho}$  are the local and average global mean density, respectively. A given spherical shell of matter initially expands with the Hubble flow (i.e. the expansion of the universe) until it reaches a particular over-density ( $\delta \simeq 1.686$ , depending on the cosmology), at
which point it begins to collapse inward.

We consider a region to have fully collapsed when it is virialized; i.e. its gravitational and kinetic energies are in equilibrium. Such a collapsed, virialized dark matter structure is what we refer to as a dark matter halo. For a flat universe with  $\Omega_m + \Omega_{\Lambda} =$ 1, the overdensity of a virialized halo, for which we use the capital letter  $\Delta$ , is given by:

$$\Delta_{vir} = \frac{18\pi^2 + 82x - 39x^2}{1+x} \tag{1.1}$$

where  $x \equiv \Omega_m - 1$  (Mo, van den Bosch, and S. White, 2010). As such, by the time a region has virialized, the overdensity has grown to  $\delta \sim \Delta_{vir} \sim 178$ .

The mass, radius, and circular velocity associated with such a virialized halo are usually referred to as the virial mass, virial radius, and virial velocity, respectively. By equating the average density of a region of radius *R* enclosing mass *M* ( $\rho = 3M/4\pi R^3$ ) with the density in terms of  $\Delta_{vir}$  ( $\rho = \Delta_{vir}\Omega_m\bar{\rho} = \Delta_{vir}\Omega_m(3H^2/8\pi G)$ ), we can obtain the expression for the virial radius below:

$$R_{vir} \simeq 163h^{-1} \text{kpc} \left[ \frac{M_{vir}}{10^{12}h^{-1} M_{\odot}} \right]^{1/3} \left[ \frac{\Delta_{vir}}{200} \right]^{-1/3} \Omega_{m,0}^{-1/3} (1+z)^{-1}$$
(1.2)

The radius of a halo is sometimes estimated by setting  $\Delta_{vir} = 200$  in Equation 1.2. This radius and the associated mass are then denoted  $R_{200}$  and  $M_{200}$ , respectively. Material within  $r \leq R_{200}$  is broadly considered to be the CGM while anything beyond is the IGM. Of course, in reality, there is no hard boundary. The  $R_{200}$  radius simply provides a useful reference.

As a final note on individual halos, the radial density distribution of dark matter halos are usually described by the profile put forward by Navarro, C. Frenk, and White, 1997. Aptly named a Navarro-Frenk-White (NFW) profile, it follows the form:

$$\rho(r) = \frac{\rho_0}{\frac{r}{R_s} \left(1 + \frac{r}{R_s}\right)^2}$$
(1.3)

where  $\rho_0$  is central density and  $R_s$  is a scale radius. The NFW profile and the broader approximation of dark matter halos as spherically symmetric systems are both idealized forms. In practice, we know that structure formation happens in a hierarchical manner. Over time, smaller halos merge and are subsumed by larger (i.e. cluster-sized) halos. It is a messy process, but these approximations are useful nonetheless.

Much of the debate on the topic of galaxy formation today is focused on the behaviour of baryonic material within the cosmological density field determined by dark matter. How does gas accrete into dark matter halos? How does the CGM evolve over time and what is its role in producing the galaxies we see today? What feedback processes dominate the evolution of the CGM? As such, let us turn now to a discussion of the CGM and IGM.

### 1.3 The Circumgalactic and Intergalactic Medium

The combined mass from the stars and gas we measure in galaxies today makes up less than half of the total baryon mass expected to have formed following the Big Bang (Mo, van den Bosch, and S. White, 2010). The remaining baryonic mass is distributed around and between galaxies, in the CGM and IGM. To understand the big picture of these gaseous reservoirs and their evolution, it helps to start at the beginning.

In its earliest moments, the universe was a hot, dense soup of fundamental particles, too energetic to be bound together into larger particles. As the universe expanded, it cooled. As the temperature dropped, it crossed a number of thresholds at which binding energies between certain particles became significant relative to their thermal energy, causing those particles to bind together.

By the time the universe was a few minutes old, the temperature had dropped enough  $(T \sim 10^9 \text{ K})$  that the nuclear binding energy between protons and neutrons became significant, causing them to bind together into atomic nuclei. The nuclei formed in this process (called 'primordial nucleosynthesis') were mostly hydrogen (~ 75%) and helium (~ 25%) with trace amounts of lithium and some heavier elements. Later still, when the universe reached a temperature of  $T \sim 4000 \text{ K}$  (at a redshift  $z \sim 1100$ ), this fluid of free electrons and atomic nuclei combined to form the first atoms. At this point, the universe was thus both highly uniform and neutral.

As discussed in the previous section, small perturbations in this near-uniform density field began a process of a run-away gravitational collapse, leading to the filamentary structure of the cosmic web seen in Figure 1.4. Baryonic material followed this collapse. Cool gas, either flowing directly into halos from the IGM or coalescing from a hot CGM, began to form stars and galaxies.

Radiation from these sources then began to re-ionize the gas around them. Evidence from absorption in very high-redshift quasars (Becker et al., 2001) tells us that this process of reionization was largely complete by z = 6, about one billion years

after the big bang. Gunn and Peterson (1965), mentioned earlier, had also used absorption in a high-redshift QSO spectrum to constrain the density of neutral gas in the IGM at z = 2. They found that "only one part in  $5 \times 10^6$  of the total mass at that time could have been in the form of intergalactic neutral hydrogen." While the onset of reionization is challenging to determine, cosmological models indicate that reionization occurs over the redshift range  $z \sim 6 - 20$  (Gnedin and Ostriker, 1997; Gnedin, 2000; Ciardi, Ferrara, and S. D. M. White, 2003; Somerville, Bullock, and Livio, 2003).

From the perspective of someone concerned with this large-scale history, galaxies are the visible by-products of a much larger, much less visible process. The bright objects we see in the distant universe are the result of extremely faint gas reacting to the gravitational influence of invisible material. In turn, as the epoch of reionization shows us, these bright sources become the driving factors in the global properties of this intergalactic gas. As such, the phases, distribution, and evolution of the IGM (and CGM) represent important puzzle pieces in any comprehensive picture of how the universe around us came to be.

In general, the IGM is usually divided up into a few temperature-density regimes. Davé et al. (2010) provides a relatively recent definition of the boundaries between these regimes. For temperature, the division is between warm/hot gas (i.e.  $T > 10^5 K$ ) and cold gas ( $T < 10^5 K$ ). For density, the axis is divided into two regimes: overdensities corresponding to collapsed halos  $\delta \ge \delta_{th}$  and overdensities corresponding to unbound gas or gas in the general IGM  $\delta < \delta th$  ( $\delta_{th} \approx 120$  at z = 0). Letting the temperature and density thresholds be denoted  $T_{th}$  and  $\rho_{th}$ , respectively, this breaks the IGM down into four main categories: (i) the warm-hot intergalactic medium (WHIM), with  $T > T_{th}$  and  $\rho < \rho_{th}$ , (ii) the cool component of the intergalactic medium with  $T \le T_{th}K$  and  $\rho < \rho_{th}$ , called the 'diffuse' component, (iii) the hot gas bound within halo, with  $T > T_{th}$  and  $\rho > \rho_{th}$ , and (iv) the cool, condensed gas within halos, with  $T < T_{th}$  and  $\rho > \rho_{th}$ .

The first two components constitute the IGM, while the latter two constitute the CGM. As a general trend over the redshift range z = 4 to z = 0, the IGM moves from being dominated by the cool, diffuse component towards having more hot gas and more gas bound in halos, as the hierarchical growth of dark matter halos creates larger potential wells and feedback from galaxies heats up the surrounding gas. The cool, condensed phase of the CGM remains a small fraction of the total baryonic content throughout, but peaks in its relative share around z = 2 - 3. This phase is



Figure 1.6: Cartoon from Tumlinson, Peeples, and Werk, 2017 showing the different components of the CGM, including accreting gas from the IGM. The mix of pristine inflows from the IGM, outflows, recycled/re-accreting gas from outflows, and diffuse gas in the halo make the CGM a complex environment.

still of critical importance to galaxy formation and evolution, as it provides galaxies with the fuel needed to grow and form stars.

# 1.4 The Co-Evolution of Galaxy and Environment

## The 10-100 Kpc Scale: Inflows, Feedback, and Mergers

There are four main processes which directly determine the evolution of the shape, size, and color of galaxies; galactic mergers, accretion of new material from the cosmic web, galaxy outflows, and the formation of stars from gas that is within the galaxy. Historically speaking, and given our human bias towards optical emission, the stellar population of a galaxy is its most defining feature - controlling its apparent color and brightness (in the absence of an AGN). I will focus here on these first three processes; mergers, inflows, and outflows.

Bluck, 2012 and Conselice et al., 2013 study the role of mergers in galaxy assembly

and find that, between a redshift of z = 1.5 and z = 3, they increase both the stellar mass and gas mass of galaxies by  $56\pm 15\%$ . This calculation includes minor mergers down to a mass ratio of 1:100. Most of the gas mass is accreted from mergers with lower-mass galaxies, while most of the accreted stellar mass comes from mergers with high-mass galaxies. On average, from z = 0 - 3, major and minor mergers are of roughly equal importance in their contributions to galaxy assembly.

High mass  $(M > 10^{11} M_{\odot})$  galaxies in the redshift range z = 1.5 - 3 form stars at a relatively constant rate, approximately doubling their stellar mass over this period (Conselice, 2012). To sustain this star formation, more gas is needed than can be obtained purely through mergers. Conselice, 2012 also show that the gas mass in massive galaxies grows by  $70 \pm 22\%$  over this period, a net accretion rate of  $61M_{\odot}yr^{-1}$ . When outflows are taken into account, the inferred gross gas accretion rate is enormous:  $96 \pm 19M_{\odot}yr^{-1}$ . Dekel et al., 2009 and C. Faucher-Giguère, Kereš, and Ma, 2011 also find accretion rates consistent with this, on the order of  $100M_{\odot}yr^{-1}$ . Given the major role of gas accretion in fueling star formation rates and galaxy growth, obtaining a handle on the reservoirs of cool around galaxies at high redshift can lend valuable insight into their evolution.

The interaction between galaxy and environment goes two ways. Radiative output from newly-formed stars ionizes and heats the surrounding gas. Galactic winds, driven by supernovae and stellar winds, inject enriched material and mechanical energy into the surrounding CGM. These processes have the effect of disrupting the cool gas in the CGM, interrupting the accretion that fueled them in the first place. For this reason, we refer to them as 'feedback' (i.e. negative feedback). Outflows and radiation from AGN are another source of feedback which enriches and energizes the CGM. Understanding the balance between accretion and feedback is one of the central dynamics in the co-evolution of galaxy and environment.

#### The Mpc Scale: Clusters, Filaments

On mega-parsec scales, the environments of galaxies are characterized by their galactic neighbours, cluster-sized dark-matter haloes within which their galactic halo is embedded, and the filaments of the IGM. Clearly, this large scale environment in turn determines the smaller scale environment: IGM gas flows into the CGM, distant satellites may become merging satellites over time, and radiation from nearby QSOs or galaxies may boost the ionizing background.

At redshifts 0.5 < z < 2, massive galaxies tend to be the most clustered, on scales of



Figure 1.7: From Thoul and Weinberg, 1995: the cooling rate of primordial gas (76% hydrogen, 24% helium, by mass). The long dashed lines show the cooling due to recombination. The short dashed lines show the cooling due to *collisional* excitation, and the dotted line represents collisional ionization. These mechanisms become important at  $T \sim 10^4$  K. The thin solid line, representing radiation from free-free electron interactions (i.e. Bremsstrahlung) begins to dominate at  $T \gtrsim 10^{5.5}$  K.

 $10-15h^{-1}$  Mpc (Hartley et al., 2010). The pertinent question here is: to what extent do these larger-scale environments influence galaxy evolution? Grützbauch et al., 2011 found that the effects of environmental density are strongest at low redshifts (z < 1) - e.g. galaxy color is correlated with environmental density - but relatively weak at higher redshifts. The galaxy halo mass and stellar mass appear to strongly dominate their properties.

### 1.5 Observational Signatures of the CGM and IGM

A logical way to identify useful tracers of IGM (and CGM) gas is to look at how it cools. Figure 1.7 is from Thoul and Weinberg, 1995 and shows the cooling rates of pristine IGM (i.e. primordial) gas. In the pristine IGM, emission driven by collisional excitation and ionization dominates for the temperature range of  $10^4 - 10^5.5$  K. A more recent and general overview of cooling in the diffuse



Figure 1.8: From Bertone and Schaye, 2012, 'How the diffuse universe cools.' Dashed lines show contributions from emission lines and solid lines show cooling from continuum emission. The colors indicate different portions of the electromagnetic spectrum. At large temperatures ( $T \ge 10^6$  K), x-ray emission continuum from Bremsstrahlung radiation dominates the cooling process. At lower temperatures  $10^4 \le T \le 10^{5.5}$  K, UV emission from both lines and continuum dominate the cooling process. IR line emission becomes significant at  $T \sim 10^3$  K.

universe (i.e.  $n_{HI} < 0.1 cm^{-3}$ ) is presented in Bertone, Aguirre, and Schaye, 2013. Figure 1.8 shows a figure from this paper, which presents a broad overview of the cooling radiation from diffuse gas in the universe. Cooling radiation from diffuse gas at T ~  $10^4 - 10^6$  K is dominated by UV line and continuum emission.

The exact contributions of different atomic species and different lines to this overall cooling rate depends on redshift. Figure 1.9, also taken from Bertone, Aguirre, and Schaye, 2013, shows the average emission per unit volume (erg s<sup>-1</sup> Mpc<sup>-3</sup>) as a function of redshift for different atomic species. Not surprisingly, hydrogen dominates the line emission contribution at all redshifts by almost an order of magnitude. The ubiquity of hydrogen throughout the IGM, combined with its dominant contributions to cooling at cool-warm temperatures, make it a clear target for any studies aiming to map the IGM.

Consisting of just a single proton and electron, it is also the simplest element. The lowest energy state of the hydrogen atom (i.e the quantum n = 1, l = 0, or '1s' state) has an energy of E = -13.6 eV. At a redshift of z = 2, atomic line emission represents ~ 80% of the cooling radiation produced by the IGM, 20% of which is



Figure 1.9: From Bertone, Aguirre, and Schaye, 2013, average emission per unit volume from different atomic species in the diffuse universe across redshifts z = 0-8.

carried by a hydrogen Lyman-alpha line alone (Bertone and Schaye, 2012). The hydrogen Lyman-alpha (henceforth Ly $\alpha$ ) line arises from the transition from the 2p (n = 2, l = 1) to the 1s (n = 1, l = 0) state. It serves as the primary tracer of CGM and IGM gas for work throughout this thesis. However, before diving into Ly $\alpha$ , let us briefly discuss another hydrogen line which is of relevance to cosmology and tracing large scale structure.

This ground state is split into two hyperfine levels based on whether the quantum spin parameters of the proton and electron are aligned or not. The difference in energy between these two hyperfine states is  $\Delta E \approx 5.87 \ \mu eV$ , corresponding to a wavelength of  $\lambda \sim 21$  cm. This emission line is ubiquitous wherever there is cool, atomic hydrogen. The 21cm line has historically been used to great effect to map the interstellar medium (ISM) in our own galaxy (van de Hulst, Muller, and Oort, 1954; Henderson, Jackson, and Kerr, 1982; Hartmann and Burton, 1997; Levine, Blitz, and Heiles, 2006). More recently, it has even been used to obtain constraints on the CGM around nearby galaxies (de Blok et al., 2014; Koribalski et al., 2018; Das et al., 2020). The 21cm line also has the potential to prove the IGM in the early

universe (Pritchard and Loeb, 2012), and indeed, there has been recent progress on making detections of the 21cm line at very high redshifts (Bowman et al., 2018; Price et al., 2018; Singh et al., 2018; Eastwood et al., 2019). However, the typical beam size of these observations is too large, and the signal too faint, to map CGM structures at intermediate redshifts ( $z \sim 2-4$ ). In addition to this, 21cm emission is produced by gas in the temperature range  $T = 10^2 - 10^3$  K, but the cool-warm gas of the IGM and CGM is at  $T \sim 10^4$  K. For tracing these structures, Ly $\alpha$  is the ideal candidate.

The energy difference of the hydrogen  $2p \rightarrow 1s$  transition is  $\Delta E = 10.2$  eV. A Ly $\alpha$  thus has a wavelength  $\lambda \approx 1215.67$  Å. Conveniently, Ly $\alpha$  at  $z \sim 2-4$  is also redshifted to visible wavelengths, enabling us to observe these rest-frame UV emission lines with ground-based, optical instruments. Ly $\alpha$  absorption profiles have been the basis of IGM studies for decades, going back to Gunn and Peterson, 1965, and observations of Ly $\alpha$  emission from the CGM are at the core of the work in this thesis.

There are three mechanisms that can produce  $Ly\alpha$  photons: (i) collisional excitation, (ii) recombination, and (iii) photo-excitation. In collisional excitation, a free electron collides with a hydrogen atom, excites the bound atom and loses kinetic energy. Energetic enough collisions can impart enough energy to free the bound electron, leading to collisional ionization. Intuitively enough, collisional excitation becomes more important at higher temperatures and densities. Photo-excitation is the excitation of the electron by absorption of an incident photon. In recombination, a free electron becomes bound by a proton, emitting a photon in the process. The term 'fluorescence' is often used to refer to the situation in which the amount of  $Ly\alpha$  emission from recombination is boosted by the presence of an external ionizing source such as a QSO. This represents an important emission mechanism for the  $Ly\alpha$  observations of the CGM today (Borisova et al., 2016; Arrigoni Battaia et al., 2019; O'Sullivan et al., 2020; Cai et al., 2019), and as such, warrants some more detailed discussion.

Recombination is typically broken into two cases; A and B. In Case A recombination, the electron is allowed to recombine into any state (i.e. any *n*), including the ground state, and all radiative transitions following this initial recombination are allowed (i.e. the electron can cascade down to the ground state in any way permitted by selection rules). In Case B recombination, both recombination directly to the ground state and higher order Lyman-series photons (e.g.  $Ly\beta$ ,  $n = 3 \rightarrow n = 1$ ) are prohibited. This



Figure 1.10: From Dijkstra, 2017. The top panel shows the probability of recombination events producing a Ly $\alpha$  photon as a function of temperature for both Case A and Case B recombination. The bottom panel shows the recombination rate for each Case as a function of temperature.

reflects the condition that most astrophysical sources will be optically thick to the photons produced by these transitions, such that they are 'immediately' reabsorbed and end up producing other photons instead. In this sense, Case A reflects an assumption of optically thin clouds while Case B reflects an assumption of optically thick clouds.

Figure 1.10, from Mark Djikstra's informative lecture series on Ly $\alpha$  radiative transfer, shows the probability that recombination will produce a Ly $\alpha$  photon under both cases, as well as the recombination rate, as a function of temperature. For Case B, recombination has a high chance (> 50%) of producing Ly $\alpha$  photons across a very wide range of temperatures, though the emissivity drops by two order of magnitudes over the range  $T = 10^2 - 10^5$  K.

Discerning the relative contributions of recombination, collisional excitation, and photo-excitation in real CGM and IGM emission is a complex task, usually requiring numerical codes to model the radiative transfer of Ly $\alpha$  photons. It also depends strongly on the assumptions about the mechanisms powering the emission. Generally speaking, there are three powering mechanisms for Ly $\alpha$  emission in the IGM and CGM. First, there is collisional excitation and ionization. This process is driven

by the thermal energy of the gas itself, which - for the IGM and CGM - is widely driven by gravitational collapse into filaments or halos. Second, there is ionization by the ambient UV background (UVB) from galaxies and quasars (i.e. the aggregate of many such distant sources). Third, there is photo-ionization and photo-excitation by proximate sources such as the central galaxy or QSO in the CGM.

Understanding the context of emission helps us determine the powering mechanism of any observed emission. For example, it can be reasonably assumed that bright  $Ly\alpha$  in the immediate vicinity of a QSO is driven by the intense ionizing radiation of the QSO. On the other hand, faint  $Ly\alpha$  emission from the general IGM, in the absence of any such sources, is more likely to be driven by gravitational cooling via collisional processes. In practical terms, however, it can be difficult to entirely rule out contributions from unseen sources such as obscured QSOs or in-situ star formation.

While there are now many observations of the CGM around high-redshift QSOs (see Chapter 4), there were not quite so many at the onset of this thesis project. Simulations, and a small number of observations, formed the basis for the expectations of how bright Ly $\alpha$  emission at z = 2 - 4 would be. C.-A. Faucher-Giguère et al., 2010 studied simulations of Ly $\alpha$  emission from cooling radiation in massive halos, finding - with some variation depending on the model - surface brightnesses on the order of  $SB_{Ly\alpha} \sim 10^{-19} - 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ . Kollmeier et al., 2010 predicted that an optically thick cloud in the IGM, powered only by photoionization by UVB radiation, would reach maximum surface brightnesses of  $SB_{L\nu\alpha} \sim 1.4 \times 10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ . Recognizing that this would take extremely long exposure times on even 10m class telescopes, they looked towards the boosted radiation fields in the vicinity of quasars. With a flux 1500 times brighter than the UVB, the authors conclude that fluorescence should be visible in a few hours on the same telescopes. By 2012, there was already small number of observations of z = 2 - 4 QSOs revealing extended Ly $\alpha$  with surface-brightnesses of  $SB_{L\nu\alpha} \sim 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$  (Stockton et al., 2002; North et al., 2012).

Although radiative transfer codes are required for a rigorous treatment, toy models can still be informative. Hennawi and Prochaska, 2013 presented simple models which can be used to estimate the column density of hydrogen from Ly $\alpha$ surface brightness measurements around a bright source under the assumption of optically thin (column density  $N_H << 10^{17} cm^{-2}$ ) or optically thick emission  $(N_H >> 10^{17} cm^{-2})$ . The source is assumed to have a spectrum following a power law in frequency with  $L_{\nu} = L_{\nu_{LL}} (\nu/\nu_{LL})^{-\alpha_Q}$ , where  $\nu_{LL}$  is the frequency of the Lyman limit ( $\lambda_{LL} = 912$  Å) and  $L_{\nu_{LL}}$  is the source luminosity at that frequency. For the optically thin case, the model of the observed surface brightness (in cgs units, i.e. erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup>) is written as:

SB [cgs] 
$$\simeq 8 \times 10^{-19} \left(\frac{1+z}{3}\right)^{-4} \left(\frac{f_{\rm C}}{1.0}\right) \left(\frac{n_{\rm H}}{0.1 \text{ cm}^{-3}}\right) \left(\frac{N_{\rm H}}{10^{20} \text{ cm}^{-2}}\right)$$
 (1.4)

where  $f_C$  is the line-of-sight covering factor of the gas and  $N_H$  is the hydrogen column density. In the optically thick limit, the model is:

SB [cgs] 
$$\simeq 6 \times 10^{-17} \left(\frac{1+z}{3}\right)^{-4} \left(\frac{f_{\rm C}}{1.0}\right) \left(\frac{R}{100 \,\rm kpc}\right) \left(\frac{L_{\nu_{\rm LL}}}{10^{30} \,\rm erg \, s^{-1} \rm Hz^{-1}}\right)^{-1}$$
 (1.5)

where R is the projected distance from the source. In the optically thin case, neither the source luminosity nor the distance to the source are factors. This is because it is assumed that the radiation field is sufficiently intense to fully ionize all of the optically thin gas. In the optically thick case, however, the gas is dense enough that it becomes 'self-shielding': only a thin layer on the surface of the cloud becomes ionized. In this picture, the incident flux does matter but only the covering factor (not the density) of the gas matters. This represents the fact that the gas in this case is acting like a mirror, reflecting (fluorescing) the incident radiation.

An important detail to note about the Ly $\alpha$  transition is that it is a *resonant* line. This refers to the fact that the Ly $\alpha$  cross-section for absorption  $\lambda \sim \lambda_{\alpha} \sim 1215.67$  Å is enormous. Dijkstra, 2017 shows that the Ly $\alpha$  cross-section is as follows:

$$\sigma_{\alpha} \sim 5.9 \times 10^{-14} (\text{T}/10^4 \text{ K})^{-1/2} \text{ cm}^{-2}$$
 (1.6)

For comparison, the Thomson cross section for the scattering of free electons is  $\sigma_{\rm T} = 6.65 \times 10^{-25}$  cm<sup>2</sup>, eleven orders of magnitude smaller. The large cross-section of the Ly $\alpha$  transition has some important implications. First, it means that the CGM and IGM become optically thick to Ly $\alpha$  at even low densities of neutral hydrogen. The sharp rise in Ly $\alpha$  optical depth at  $z \sim 6$  (discussed earlier) still corresponds only to a neutral fraction of  $x_{HI} \sim 10^{-3}$ . This is partially why it is so difficult to probe the early reionization epoch; the universe becomes optically thick to Ly $\alpha$  while the IGM is still highly ionized.

The other big implication of this large cross-section is on the spectral shape of Ly $\alpha$  emission. Ly $\alpha$  photons produced in the middle of a dense cloud (say, by collisional



Figure 1.11: From Bertone and Schaye, 2012, simulated emission levels of different UV emission lines in the z = 2 - 5 CGM and IGM. The second column shows hydrogen Ly $\alpha$  under the assumption that the gas becomes self-shielding.

excitation) will not be able to escape the cloud. Instead, they will random-walk in frequency space by a large number of scatterings - where the velocity of the atoms impart a Doppler shift on the photon. Only when the photon is scattered to high or low enough frequencies will the cloud become optically thin to it. Because of this, the characteristic profile of Ly $\alpha$  from a dense cloud in the CGM or IGM is that of a double peaked profile. This complex radiative transfer makes it difficult to interpret spectra and velocity maps from optically thick Ly $\alpha$  sources.

To measure the ionization state, metallicity, density, and kinematics of the CGM/IGM,  $Ly\alpha$  alone is not sufficient. Emission lines from other atomic species are needed for this. Figure 1.9 shows the relative emission levels of different elements (averaged over different ions), with carbon, helium, oxygen and silicon representing some of

the brightest. Specifically, among the brightest UV metal lines are CIV  $\lambda$ 1549, HeII  $\lambda$ 1640, CIII  $\lambda$ 977, and OVI  $\lambda$ 1031. Figure 1.11, from Bertone and Schaye, 2012, shows simulated emission levels for some of these.

The surface brightness levels for these metal lines are all an order of magnitude or more fainter than that of  $Ly\alpha$ . Given the difficulty in even detecting  $Ly\alpha$  from the CGM and IGM, observations of these metal lines are seriously challenging to obtain. However, the wealth of information they offer about the properties of the CGM makes them important to try and observe nonetheless.

### **1.6 Integral Field Spectroscopy as an Exploratory Tool**

Integral field spectroscopy is, in simple terms, simultaneous imaging and spectroscopy. The data produced by an integral field spectrograph essentially contains an image in which every pixel has a full spectrum. This makes them powerful exploratory tools for emission where the morphology and exact wavelength is uncertain. The adjustable spectral range of an IFS also enables flexibility in terms of both cosmological redshift and targeted emission lines. The flexibility of the data in having both spatial and spectral resolution allows signals to be optimally extracted, integrating only over the wavelengths or on-sky regions containing emission.

As is always the case, these advantages - extremely useful in exploring the CGM and IGM - come at a cost. IFS designs, by their nature, split the incoming light into more 'buckets' and require more complex optics, reducing the total throughput of the instrument. Given that the CGM and IGM are already extremely faint sources, this presents a challenge. However, with a large enough telescope, good enough seeing conditions, and long exposure times, it is one that can be overcome. The advantages and disadvantages of IFS designs, compared to other methods, will be discussed in more detail at the beginning of Chapter 2, rather than go into the topic at length here.

Integral field spectrographs (IFS), or integral field unit (IFU) spectrographs, are relatively modern instruments. The first IFS was proposed by G. Courtes in 1982 as the TIGER instrument which saw first light at the European Southern Observatory (ESO) in 1987 (Courtes et al., 1988; Bacon et al., 1995). The first IFU installed at Keck Observatory was the instrument OSIRIS (Larkin et al., 2003), which saw first light in 2005. By 2014, the two most sensitive IFUs in the world were the Palomar Cosmic Web Imager (PCWI) (Matuszewski et al., 2010) on the 5 meter Hale telescope at Palomar Observatory, and the more recently installed Multi Unit

Spectrographic Explorer (MUSE) (Caillier et al., 2014) on the 8 meter Very Large Telescope at ESO. The Keck Cosmic Web Imager (KCWI) (Morrissey et al., 2018) was in the late stages of design, to be finished and installed on the summit of Mauna Kea by 2017.

With IFUs on 5-10m class telescopes, we can directly map the CGM and IGM across at arbitrary redshifts with full spatial 2D spatial coverage without the need for multiple pointings or a large set of custom narrowband filters. This provides a powerful exploratory tool for studying the CGM and IGM.

### 1.7 Outline of Thesis

This thesis is focused on integral field spectroscopy of the CGM and IGM. The Chapters are ordered such that each is a prerequisitve for, and leads into, the next.

Chapter 2 outlines instrumental work on the Keck Cosmic Web Imager (KCWI), which was later used to conduct observations of the CGM at z = 2 - 3. Chapter 3 outlines a software packag, developed to enable IFS surveys of the CGM and IGM. These tools are used conduct the survey work presented in the following two chapters.

Chapter 4 presents the first part of the Fluorescent Lyman-Alpha Structures in Highz Environments (FLASHES) Survey. This initial pilot survey targeted Ly $\alpha$  emission in the CGM surrounding z = 2-3 QSOs. Chapter 5 presents the first follow-up deep study of the FLASHES Survey, targeting metal-line emission in the CGM around a subset of the FLASHES pilot sample.

Chapter 6 presents work on the second-generation Faint Intergalactic Redshifted Emission Balloon (FIREBall-2), a high-altitude UV telescope and IFS targeting CGM emission at low redshift  $z \leq 1$ .

Finally, Chapter 7 summarizes the outcomes of this instrumental and observational work. It also reflects on the broader changes in observational studies over this time period and presents a brief outlook on the near future of observational CGM and IGM studies.

### References

Arrigoni Battaia, F. et al. (Jan. 2019). "QSO MUSEUM I: a Sample of 61 Extended Ly α Emission Nebulae Surrounding z ~ 3 Quasars." In: *MNRAS* 482.3, pp. 3162– 3205. DOI: 10.1093/mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].

- Bacon, R. et al. (Oct. 1995). "3d Spectrography At High Spatial Resolution. I. Concept And Realization Of The Integral Field Spectrograph Tiger." In: A&AS 113, p. 347.
- Barnes, J. et al. (Aug. 1985). "Cluster-cluster Clustering." In: *ApJ* 295, pp. 368–373. DOI: 10.1086/163381.
- Becker, R. H. et al. (Dec. 2001). "Evidence for Reionization at z\_6: Detection of a Gunn-Peterson Trough in a z=6.28 Quasar." In: AJ 122.6, pp. 2850–2857. DOI: 10.1086/324231. arXiv: astro-ph/0108097 [astro-ph].
- Bertone, S., A. Aguirre, and J. Schaye (May 2013). "How the Diffuse Universe Cools." In: *MNRAS* 430, pp. 3292–3313. DOI: 10.1093/mnras/stt131. arXiv: 1301.5330.
- Bertone, S. and J. Schaye (Jan. 2012). "Rest-frame Ultraviolet Line Emission From The Intergalactic Medium at  $2 \le z \le 5$ ." In: *MNRAS* 419, pp. 780–798. DOI: 10. 1111/j.1365-2966.2011.19742.x. arXiv: 1008.1791.
- Binney, J. (July 1977). "The Physics Of Dissipational Galaxy Formation." In: *ApJ* 215, pp. 483–491. DOI: 10.1086/155378.
- Bluck, A. et al. (Mar. 2012). "The Structures and Total (Minor + Major) Merger Histories of Massive Galaxies up to z ~3 in the HST GOODS NICMOS Survey: A Possible Solution to the Size Evolution Problem." In: *ApJ* 747.1, 34, p. 34. DOI: 10.1088/0004-637X/747/1/34. arXiv: 1111.5662 [astro-ph.CO].
- Blumenthal, G. R. et al. (Oct. 1984). "Formation Of Galaxies And Large-scale Structure With Cold Dark Matter." In: *Nature* 311, pp. 517–525. doi: 10.1038/311517a0.
- Bond, J. R. and A. S. Szalay (Nov. 1983). "The Collisionless Damping Of Density Fluctuations In An Expanding Universe." In: *ApJ* 274, pp. 443–468. doi: 10. 1086/161460.
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Bowman, J. et al. (Mar. 2018). "An Absorption Profile Centred At 78 Megahertz In The Sky-averaged Spectrum." In: *Nature* 555.7694, pp. 67–70. DOI: 10.1038/nature25792. arXiv: 1810.05912 [astro-ph.CO].
- Cai, Z. et al. (Sept. 2019). "Evolution of the Cool Gas in the Circumgalactic Medium (CGM) of Massive Halos A Keck Cosmic Web Imager (KCWI) Survey of Ly $\alpha$  Emission around QSOs at  $z \approx 2$ ." In: *arXiv e-prints*, arXiv:1909.11098, arXiv:1909.11098. [astro-ph.GA].
- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: *Ground-based and Airborne Instrumentation for Astronomy V.* Vol. 9147. Proc. SPIE, 91475K. DOI: 10.1117/12.2057056.

- Ciardi, B., A. Ferrara, and S. D. M. White (Sept. 2003). "Early Reionization By The First Galaxies." In: *MNRAS* 344.1, pp. L7–L11. DOI: 10.1046/j.1365-8711.2003.06976.x. arXiv: astro-ph/0302451 [astro-ph].
- Colless, M. (Jan. 1999). "First Results from the 2dF Galaxy Redshift Survey." In: *Philosophical Transactions of the Royal Society of London Series A* 357.1750, p. 105. DOI: 10.1098/rsta.1999.0317. arXiv: astro-ph/9804079 [astro-ph].
- Conselice, C. J. (Dec. 2012). "Galaxy Formation: Where Do We Stand?" In: *arXiv e-prints*, arXiv:1212.5641, arXiv:1212.5641. arXiv: 1212.5641 [astro-ph.CO].
- Conselice, C. J. et al. (July 2013). "Gas Accretion and Mergers in Massive Galaxies at z ~2." In: *The Intriguing Life of Massive Galaxies*. Ed. by Daniel Thomas, Anna Pasquali, and Ignacio Ferreras. Vol. 295. IAU Symposium, pp. 45–48. DOI: 10.1017/S1743921313004225.
- Courtes, G. et al. (Jan. 1988). "A New Device for Faint Objects High Resolution Imagery and Bidimensional Spectrography - First Observational Results with TIGER at CFHT 3.6-METER Telescope." In: *Instrumentation for Ground-Based Optical Astronomy*, p. 266.
- Cowsik, R. and J. McClelland (Feb. 1973). "Gravity of Neutrinos of Nonzero Mass in Astrophysics." In: *ApJ* 180, pp. 7–10. DOI: 10.1086/151937.
- Das, S. et al. (May 2020). "Detection of the Diffuse HI emission in the Circumgalactic Medium of NGC 891 and NGC 4565." In: *arXiv e-prints*, arXiv:2005.13684, arXiv:2005.13684. arXiv: 2005.13684 [astro-ph.GA].
- Davé, R. et al. (Nov. 2010). "The Intergalactic Medium Over The Last 10 Billion Years - I. Lyα Absorption And Physical Conditions." In: *MNRAS* 408.4, pp. 2051– 2070. DOI: 10.1111/j.1365-2966.2010.17279.x. arXiv: 1005.2421 [astro-ph.CO].
- de Blok, W. J. G. et al. (June 2014). "Halogas Observations Of Ngc 4414: Fountains, Interaction, And Ram Pressure." In: *A&A* 566, A80, A80. DOI: 10.1051/0004-6361/201322517. arXiv: 1405.2160 [astro-ph.GA].
- Dekel, A. et al. (Jan. 2009). "Cold Streams In Early Massive Hot Haloes As The Main Mode Of Galaxy Formation." In: *Nature* 457.7228, pp. 451–454. DOI: 10.1038/nature07648. arXiv: 0808.0553 [astro-ph].
- Dijkstra, M. (Apr. 2017). "Saas-Fee Lecture Notes: Physics of Lyman Alpha Radiative Transfer." In: *arXiv e-prints*, arXiv:1704.03416, arXiv:1704.03416. arXiv: 1704.03416 [astro-ph.GA].
- Eastwood, M. W. et al. (Aug. 2019). "The 21 cm Power Spectrum from the Cosmic Dawn: First Results from the OVRO-LWA." In: *AJ* 158.2, 84, p. 84. DOI: 10. 3847/1538-3881/ab2629. arXiv: 1906.08943 [astro-ph.CO].
- Fall, S. M. and G. Efstathiou (Oct. 1980). "Formation and Rotation of Disc Galaxies with Haloes." In: *MNRAS* 193, pp. 189–206. DOI: 10.1093/mnras/193.2.189.

- Faucher-Giguère, C., D. Kereš, and C. Ma (Nov. 2011). "The Baryonic Assembly Of Dark Matter Haloes." In: *MNRAS* 417.4, pp. 2982–2999. DOI: 10.1111/j.1365-2966.2011.19457.x. arXiv: 1103.0001 [astro-ph.CO].
- Faucher-Giguère, Claude-André et al. (Dec. 2010). "Lyα Cooling Emission from Galaxy Formation." In: *ApJ* 725.1, pp. 633–657. DOI: 10.1088/0004-637X/ 725/1/633. arXiv: 1005.3041 [astro-ph.CO].
- Gnedin, N. Y. (Oct. 2000). "Effect of Reionization on Structure Formation in the Universe." In: *ApJ* 542.2, pp. 535–541. DOI: 10.1086/317042. arXiv: astro-ph/0002151 [astro-ph].
- Gnedin, N. Y. and J. P. Ostriker (Sept. 1997). "Reionization of the Universe and the Early Production of Metals." In: *ApJ* 486.2, pp. 581–598. DOI: 10.1086/304548. arXiv: astro-ph/9612127 [astro-ph].
- Grützbauch, R. et al. (Dec. 2011). "The Relationship Between Star Formation Rates, Local Density And Stellar Mass Up To z ~ 3 in the GOODS NICMOS Survey." In: *MNRAS* 418.2, pp. 938–948. DOI: 10.1111/j.1365-2966.2011.19559.x. arXiv: 1108.0402 [astro-ph.C0].
- Gunn, J. E. and B. A. Peterson (Nov. 1965). "On the Density of Neutral Hydrogen in Intergalactic Space." In: *ApJ* 142, pp. 1633–1636. DOI: 10.1086/148444.
- Hartley, W. G. et al. (Sept. 2010). "The Evolution Of Galaxy Clustering Since Z = 3 Using The Ukidss Ultra Deep Survey: The Divergence Of Passive And Starforming Galaxies." In: *MNRAS* 407.2, pp. 1212–1222. DOI: 10.1111/j.1365-2966.2010.16972.x. arXiv: 1005.1180 [astro-ph.CO].
- Hartmann, D. and W. B. Burton (1997). Atlas of Galactic Neutral Hydrogen.
- Henderson, A. P., P. D. Jackson, and F. J. Kerr (Dec. 1982). "The Distribution Of Neutral Atomic Hydrogen In Our Galaxy Beyond The Solar Circle." In: *ApJ* 263, pp. 116–122. DOI: 10.1086/160486.
- Hennawi, J. F. and J. X. Prochaska (Mar. 2013). "Quasars Probing Quasars. IV. Joint Constraints on the Circumgalactic Medium from Absorption and Emission." In: *ApJ* 766.1, 58, p. 58. DOI: 10.1088/0004-637X/766/1/58. arXiv: 1303.2708 [astro-ph.CO].
- Hinshaw, G. et al. (Oct. 2013). "Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results." In: *ApJS* 208.2, 19, p. 19. DOI: 10.1088/0067-0049/208/2/19. arXiv: 1212.5226 [astro-ph.CO].
- Kollmeier, J. A. et al. (Jan. 2010). "Lyα Emission from Cosmic Structure. I. Fluorescence." In: *ApJ* 708.2, pp. 1048–1075. DOI: 10.1088/0004-637X/708/2/1048. arXiv: 0907.0704 [astro-ph.C0].
- Koribalski, B. S. et al. (Aug. 2018). "The Local Volume H I Survey (LVHIS)." In: MNRAS 478.2, pp. 1611–1648. DOI: 10.1093/mnras/sty479. arXiv: 1904. 09648 [astro-ph.GA].

- Larkin, J. et al. (Mar. 2003). "OSIRIS: Infrared Integral Field Spectrograph for the Keck Adaptive Optics System." In: *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*. Ed. by Masanori Iye and Alan F. M. Moorwood. Vol. 4841. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 1600–1610. DOI: 10.1117/12.461775.
- Levine, E. S., L. Blitz, and C. Heiles (June 2006). "The Spiral Structure of the Outer Milky Way in Hydrogen." In: *Science* 312.5781, pp. 1773–1777. DOI: 10.1126/science.1128455. arXiv: astro-ph/0605728 [astro-ph].
- Matuszewski, M. et al. (July 2010). "The Cosmic Web Imager: An Integral Field Spectrograph For The Hale Telescope At Palomar Observatory: Instrument Design And First Results." In: *Ground-based and Airborne Instrumentation for Astronomy III.* Vol. 7735. Proc. SPIE, 77350P. DOI: 10.1117/12.856644.
- Mo, H., F. C. van den Bosch, and S. White (2010). Galaxy Formation and Evolution.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- Navarro, J., C. Frenk, and White (Dec. 1997). "A Universal Density Profile from Hierarchical Clustering." In: *ApJ* 490, pp. 493–508. DOI: 10.1086/304888. eprint: astro-ph/9611107.
- Nelson, D. et al. (May 2019). "The IllustrisTNG Simulations: Public Data Release." In: *Computational Astrophysics and Cosmology* 6.1, 2, p. 2. DOI: 10.1186/ s40668-019-0028-x. arXiv: 1812.05609 [astro-ph.GA].
- North, P. L. et al. (June 2012). "Spectroscopy of Extended Lyα Envelopes Around z = 4.5 Quasars." In: A&A 542, A91, A91. DOI: 10.1051/0004-6361/201015153. arXiv: 1205.3895 [astro-ph.CO].
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 z ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.
- Peebles, P. J. E. (Dec. 1982). "Large-scale Background Temperature And Mass Fluctuations Due To Scale-invariant Primeval Perturbations." In: *ApJ* 263, pp. L1– L5. DOI: 10.1086/183911.
- Price, D. C. et al. (Aug. 2018). "Design And Characterization Of The Large-aperture Experiment To Detect The Dark Age (LEDA) Radiometer Systems." In: *MNRAS* 478.3, pp. 4193–4213. DOI: 10.1093/mnras/sty1244. arXiv: 1709.09313 [astro-ph.IM].
- Pritchard, J. and A. Loeb (Aug. 2012). "21 cm Cosmology in the 21st Century." In: *Reports on Progress in Physics* 75.8, 086901, p. 086901. doi: 10.1088/0034-4885/75/8/086901. arXiv: 1109.6012 [astro-ph.CO].
- Rees, M. J. and J. P. Ostriker (June 1977). "Cooling, Dynamics and Fragmentation of Massive Gas Clouds: Clues to the Masses and Radii of Galaxies and Clusters.." In: *MNRAS* 179, pp. 541–559. DOI: 10.1093/mnras/179.4.541.

- Rubin, V. and Jr. Ford W. (Feb. 1970). "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions." In: *ApJ* 159, p. 379. DOI: 10.1086/150317.
- Silk, J. (Feb. 1977). "On the Fragmentation of Cosmic Gas Clouds. I. The Formation of Galaxies and the First Generation of Stars." In: *ApJ* 211, pp. 638–648. DOI: 10.1086/154972.
- Singh, S. et al. (May 2018). "SARAS 2 Constraints on Global 21 cm Signals from the Epoch of Reionization." In: *ApJ* 858.1, 54, p. 54. DOI: 10.3847/1538-4357/aabae1. arXiv: 1711.11281 [astro-ph.CO].
- Somerville, R. S., J. S. Bullock, and M. Livio (Aug. 2003). "The Epoch of Reionization in Models with Reduced Small-Scale Power." In: *ApJ* 593.2, pp. 616–621. DOI: 10.1086/376686. arXiv: astro-ph/0303481 [astro-ph].
- Steidel, C. et al. (May 1996). "Spectroscopic Confirmation of a Population of Normal Star-forming Galaxies at Redshifts z > 3." In: ApJ 462, p. L17. DOI: 10.1086/310029. arXiv: astro-ph/9602024 [astro-ph].
- Stockton, A. et al. (June 2002). "The Extended Emission-Line Region of 4C 37.43." In: *ApJ* 572.2, pp. 735–745. DOI: 10.1086/340432. arXiv: astro-ph/0205040 [astro-ph].
- Thoul, A. and D. Weinberg (Apr. 1995). "Hydrodynamic Simulations of Galaxy Formation. I. Dissipation and the Maximum Mass of Galaxies." In: *ApJ* 442, p. 480. DOI: 10.1086/175455. arXiv: astro-ph/9410009 [astro-ph].
- Tumlinson, J., M. S. Peeples, and J. K. Werk (Aug. 2017). "The Circumgalactic Medium." In: ARA&A 55.1, pp. 389–432. DOI: 10.1146/annurev-astro-091916-055240. arXiv: 1709.09180 [astro-ph.GA].
- van de Hulst, H. C., C. A. Muller, and J. H. Oort (May 1954). "The Spiral Structure Of The Outer Part Of The Galactic System Derived From The Hydrogen Emission At 21 Cm Wavelength." In: *Bull. Astron. Inst. Netherlands* 12, p. 117.
- White, S. D. M., C. S. Frenk, and M. Davis (Nov. 1983). "Clustering In A Neutrinodominated Universe." In: *ApJ* 274, pp. L1–L5. DOI: 10.1086/184139.
- White, S. D. M. and M. J. Rees (May 1978). "Core Condensation in Heavy Halos: a Two-stage Theory for Galaxy Formation and Clustering." In: *MNRAS* 183, pp. 341–358. doi: 10.1093/mnras/183.3.341.
- Zeldovich, Ia. B., J. Einasto, and S. F. Shandarin (Dec. 1982). "Giant Voids In The Universe." In: *Nature* 300.5891, pp. 407–413. DOI: 10.1038/300407a0.

# Chapter 2

# THE KECK COSMIC WEB IMAGER



Figure 2.1: The Keck Cosmic Web Imager (KCWI) being installed on the Nasmyth platform in summer 2017 (Morrissey et al., 2018).

The Keck Cosmic Web Imager (KCWI) is an integral field spectrograph installed on the Keck-2 telescope at the W. M. Keck Observatory at the summit of Mauna Kea, in Hawaii. KCWI is designed to study faint, extended emission, with a paricular focus on circumgalactic and intergalactic medium. This chapter provides an overview of the Cosmic Web Imager instrument design, contributions to the development of KCWI, and a brief comparison of PCWI and KCWI performance.

# 2.1 Advantages and Limitations of Integral Field Spectroscopy

Exploring the diffuse universe is a significant observational challenge. Flexibility is needed to map the faint nebular line emission from galaxy environments and the intergalactic medium. The diffuse universe is often complex in morphology and has

line centers that vary wildly due to gas kinematics (i.e., Doppler shifts) and radiative transfer effects (primarily absorption) (Borisova et al., 2016; Arrigoni Battaia et al., 2019; O'Sullivan et al., 2020).

These properties can make exploratory work with narrow-band imaging or singleslit spectroscopy prohibitively expensive, either in terms of time or actual cost (e.g. for manufacturing custom filters). Cool-warm gas in the CGM is expected to exhibit dispersions of  $\sigma_{1D} \leq 400$  km/s, which corresponds to a full-width-athalf-maximum (FWHM) of  $\Delta \lambda_{FWHM} \leq 13 - 15$  Å for a cosmological redshift of  $z \approx 2 - 3$ . Thus, narrow filters (on the order of a few ×10 Å) are required to obtain the optimal signal-to-noise ratio.

However, recent work shows that the velocity centroids of resonant lines like Ly $\alpha$  can vary by many thousands of kilometers per second from the systemic redshifted values. For example, in the FLASHES survey (O'Sullivan et al., 2020), the distribution of velocity centers of the nebular emission lines had a standard deviation of  $\sigma_v \simeq 1000$  km/s, which corresponds to an uncertainty of  $\sigma_\lambda \simeq 35$  Å in the position of the line center.

A wider filter could be used, but the signal-to-noise ratio will drop as  $S/N \propto 1/\sqrt{\Delta \lambda_{NB}}$  due to shot noise from the additional integrated background signal (i.e. sky and continuum sources). A narrow-band filter wide enough to accommodate a  $\pm 2\sigma$  ( $\pm 30$  Å) velocity uncertainty will be a factor of about  $\sqrt{60/15} = \sqrt{8}$  worse in signal-to-noise than an ideal filter.

Long-slit or single-slit spectroscopy offers the required spectral width to deal with this, as well as the possibility of obtaining high resolution data with relatively high throughput. However, the distribution of gas in the CGM is often highly asymmetric. Moreover, when depending on a nearby QSO to illuminate the gas, the illumination itself may be asymmetric. In a 2013 paper, Bridge et al., 2013 tested this dependence with a simple toy model and found that random orientations of a single-slit with respect to a bi-polar distribution of emitting gas yielded only a 37% success rate in detecting the emission.

It is clear then that in addition to spectral flexibility, full spatial coverage is also needed if we are to be confident in the completeness of our search for emission in galaxy environments. This flexibility is the primary benefit yielded by IFS designs. However, as is always the case, this strength comes with a trade-off.

The two primary trade-offs are throughput and field of view. IFS designs require

complex optics with many elements (mirrors, gratings, etc.), each of which incurs some loss in throughput if it does not have 100% reflectance or throughput. There is also a trade-off between signal and resolution for any detector with non-negligible read noise; the more bins you divide a signal into, the lower the SNR in each bin.

However, for background-limited observations, where the read noise is not significant compared to the shot noise, these bins can be added together in post-processing without a high cost. Signals can then be optimally extracted from the data by summing only the spatial/spectra bins containing emission. The possibility of this 'optimal extraction' (e.g. Borisova et al., 2016) enables significant gains in the final signal-to-noise ratio of an extended emission region. Lastly, larger telescopes and longer exposure times can, to some degree, compensate for lower throughput.

The more inflexible trade-off with IFS designs is the field of view or, alternatively, the spatial resolution. In order to obtain a large field of view, there would either have to be many narrow slices in the image slicer or the slices would have to be much wider themselves. The instrument complexity and cost scales sensitively with the number of slices, as the image from each slice needs its own optical path to be redirected, dispersed, and focused onto the detector.

A good illustration of this complexity is the Multi-Unit Spectroscopic Explorer (MUSE - Caillier et al., 2014) on the VLT at the European Southern Observatory in Chile. MUSE, in its wide field mode, has a field of view of  $60'' \times 60''$  and a spatial resolution of  $0.2'' \times 0.2''$ . To accomplish this, the field of view is split into 24 sub-fields, each of which has its own designated IFS, camera, and detector assembly - an enormous instrument!

Another cost driver to a large field of view with high resolution is the so-called detector 'real-estate' required. Figure 2.2 shows an example PCWI image to illustrate how 3D data is laid out on a 2D detector. In this image, the wavelength dimension is along the vertical axis, with bluer wavelengths towards the bottom and redder wavelengths towards the top. The vertical spatial dimension of the image slicer (i.e. the slice number from 1-24) is reflected by the separated 2D spectra, with left to right on the detector showing slices 1-24. Finally, the horizontal dimension of the image slicer (distance along a slice) is reflected by the x-axis within each 2D slice.

For an image slicer with a 1:1 aspect ratio, a field of angular size  $\Delta \Phi$ , and angular resolution  $\delta \theta$ , there will be  $N_{\text{slice}} = (\Delta \Phi / \Delta \theta)$  2D spectra, each requiring  $N_x = 2(\Delta \Phi / \Delta \theta)$  pixels to achieve a Nyquist sampling rate. The total number of



Figure 2.2: An example PCWI image, cropped to 1/3 of the vertical/wavelength axis. The 2D spectra of the slices are arranged from left to right. Each 2D spectrum thus represents one 'pixel' in the spatial y-axis of the field of view. Within each 2D spectrum, the spatial x-axis (i.e. distance along a slice) runs horizontally, while the wavelength axis runs vertically. The quasar HS1700+6416 can be seen as a bright continuum source on slices 9-13, while emission lines from the night sky can be seen as horizontal bands in each slice.

pixels required along the x-axis will then be  $N_X = N_{slice} \times N_x = 2(\delta \Phi / \delta \theta)^2$ , ignoring any overhead for space between slices.

To cover an instantaneous bandpass  $\Lambda$ , centered on a wavelength  $\lambda_c$ , with a spectral resolution R and Nyquist sampling, the minimum number of pixels required along the other axis of the detector is  $N_Y = 2R(\Lambda/\lambda_c)$ . Thus the total detector area required to lay out the 3D data from this image slicer is  $N_{px} = 4(\delta\Phi/\delta\theta)^2R(\Lambda/\lambda_c)$ . Given the cost of science grade CCDs, this is an expensive formula for the ratio between field of view and angular resolution!

### 2.2 The Architecture of the Cosmic Web Imagers

The Palomar Cosmic Web Imager (PCWI), which is now a facility instrument on the 5-meter Hale telescope at Palomar Observatory in California, was designed in part as a prototype instrument to The Keck Cosmic Web Imager. As such, it provides a good opportunity to illustrate the architecture of KCWI in a simpler form. The architecture of the Cosmic Web Imager instruments is that of an image-slicer-based integral field spectrograph.

In an 'image slicer' architecture, a 2D field of view is split into a series of narrow slices by a reflective block which is divided into thin strips, each oriented at a slightly different angle. These narrow slices each serve the same function as the long slit in a spectrograph. The image of each slice is passed through a dispersive element

(in this case a volume-phase holographic grating) to obtain a 2D spectrum, which is then projected onto the detector. The resulting raw image is a series of 24 separated 2D spectra which can be reconstructed into three-dimensional data cubes by the instrument's data reduction pipeline.

PCWI features a single image slicer with a field of view of  $60'' \times 40''$ , made up of 24 slices with widths of ~ 2.7'' and in-slice sampling of 0.55'' per pixel. PCWI covers a wide wavelength range of 370 - 950 nm, with default instantaneous coverage of 45 nm for the standard gratings ( $R = \lambda/\Delta\lambda \approx 5000$ ) and 90 nm for the medium resolution grating ( $R \approx 2500$ ).

For a galaxy at a redshift of  $z \approx 2.7$ , a single PCWI exposure thus allows us to fully explore a physical projected area of roughly  $320 \times 480 \text{ pkpc}^2$  with a velocity width of  $\pm 15000 \text{ km/s}$ , while maintaining a spatial resolution of  $4.4 \times 20 \text{ pkpc}^2$  and a spectral resolution of  $\Delta v = 60 \text{ km/s}$ . This provides us with a powerful tool for performing exploratory observations of the CGM and IGM without requiring prior knowledge of the distribution or *exact* redshift of the emission (some reasonable estimate is required).

PCWI can reach a limiting surface brightness of ~  $5 \times 10^{-18}$  erg/s/cm<sup>2</sup>/arcsec<sup>2</sup> in a single wavelength layer in just one hour on source, sufficient to detect the brighter parts of the high-redshift CGM seen in high-redshift studies. The Cosmic Web Imagers are designed to provide background-limited observations of faint emission, meaning read noise must be low (e.g. on the order of a few electrons per pixel per readout.) Observational strategy is also an important element in minimizing the contributions from read noise; a single one-hour exposure is obviously more desirable in terms of read noise than a series of six ten-minute exposures, but there is a trade-off.

Very long exposures suffer from poorly characterized backgrounds because emission from the night sky fluctuates noticeably even on minute timescales. At Palomar in particular, the intensity of emission lines from the sky changes enough over a 20minute exposure that even subtracting a 20-minute sky frame taken immediately after will leave sharp residuals many times louder than the Poisson noise of the background signal. If we had an ideal detector (i.e. no read noise), the best strategy would be to interleave very short exposures of the source and sky fields, to enable accurate sky subtraction.

Since such a detector does not exist, the next best thing we can do to obtain excellent

sky subtraction is using a technique called "Nod-and-Shuffle" (NAS). In this technique, the detector is covered with a mask that divides it into three regions along the wavelength axis, with the central 1/3 being exposed and the top/bottom covered. The cost of doing this is that the bandwidth is cut in three, though the full spatial field of view is maintained.

The benefit is that we can now use the different detector regions to simultaneously observe the source and sky fields without reading out. This is done by first pointing at the source field and exposing the detector for a short period (e.g. 2 minutes) so that the central 1/3 of the detector accumulates photo-electrons. When the exposure is done, the telescope "nods" to the background field, the detector electronically "shuffles" the stored charge on the central 1/3 to the top 1/3 without reading out, and a second exposure of equal length is begun, accumulating charge on the now empty central region of the detector. When this exposure is complete, it nods back to the source field, shuffles the charge on the detector back down by 1/3, and continues with a second exposure on the science field. By repeating this process, very short exposures of the target and background fields can be interleaved without adding read noise.

The inclusion of a N&S detector mask and the software controls to perform this technique is one of the major design features that make the Cosmic Web Imagers uniquely empowered to target extremely faint, diffuse emission. The sky background, which determines the noise, can be measured to such accuracy that the subtraction residuals have typical magnitudes of less than one percent of the sky brightness. The N&S method does incur a  $\sqrt{2}$  cost to the statistical SNR compared to some model-based sky-subtraction approaches, but is virtually free of systematic residuals.

In summary, the Cosmic Web Imagers are slicer-based integral-field spectrographs custom-built to explore faint, diffuse emission. PCWI has a single imager slicer, manually interchangeable gratings/filters, and a N&S mask that is also inserted and retracted manually. The Keck Cosmic Web imager features several significant upgrades to the design described here, which are outlined the following section.

### 2.3 From Palomar to Keck: KCWI Upgrades

The night sky at Mauna Kea is fainter and more stable due to laminar flow of air over the gently sloping mountain (Mauna Kea is a shield volcano). In addition, at ten meters, the Keck-2 mirror is twice the diameter of the Hale Telescope at Palomar.



Figure 2.3: The integral field unit of the Palomar Cosmic Web Imager, taken from Matuszewski et al., 2010. The left panel shows a zoomed-in view of the image slicer; a 24mm × 16mm block with 1mm reflective slices. It can be seen mounted in the bottom-right corner of the right panel. Each slice is at a slightly different angle, such that the reflected image is directed to one of the 24 pupil-array mirrors seen in the right panel. These mirrors reflect the slice images to a curved mirror which collimates the beam before it is passed through the grating.

These two details alone account for significant improvements in the sensitivity of KCWI. However, the instrument design also has a number of improvements.

While PCWI is a single-channel instrument with full coverage of the optical spectrum, KCWI is split into two distinct optical paths for the blue and red portions of the optical spectrum, each with their own camera and detector. This means that the coatings on each camera and detector can be optimized for throughput in the blue and red portions of the optical spectrum.

The flexibility of KCWI as an observational tool is substantially increased. KCWI has three interchangeable image slicers with a robotic exchanger that automatically swaps them out. The gratings and filters also have automatic exchange mechanisms, which means that an observer can change the field of view, spatial resolution, spectral resolution, and bandpass in the middle of the night.

For different combinations of image slicer, grating, and detector readout mode, the spectral resolution of the instrument ranges from R = 900 to R = 20,000, and the field of view ranges from  $20'' \times 33''$  to  $20'' \times 8''$ . In addition, the eventual accessible wavelength range of KCWI will be 350nm - 1050nm. All of this provides powerful flexibility to an observer. However, it means that there is a wide array of settings that must be calibrated. Whether the instrument is set to a spectral resolution of R = 20,000 and a central wavelength of  $\lambda_c = 825$  nm or a resolution of R = 900 and a central wavelength of  $\lambda_c = 410nm$ , the internal calibration must be able to handle



Figure 2.4: The calibration targets assembly of KCWI, shown as an overhead view of the CAD model (right) and a photograph after installation. Box (a) labels the arc lamp unit, designed and built as a part of this thesis. Box (b) labels the integrating sphere (partially obscured by the continuum housing unit.) Box (c) labels the calibration target mask assembly, which moves different masks in front of the exit port of the integrating sphere. Credit for both images goes to Matt Matuszewski, the official 'ghost in the machine' of the CWI instruments.

it or the data will not be usable.

The design and manufacture of the spectral calibration arc-lamp module is the primary contribution within this thesis to the development of KCWI. The following section will discuss the constraints, requirements, design process and resulting design for this module, which is now sitting on top of Mauna Kea, calibrating data for KCWI observers.

## 2.4 Spectral Calibration Unit Design

All observational data needs to be calibrated. KCWI contains an internal calibration unit with spectral and spatial standards used to determine the three-dimensional geometry of the data in right-ascension, declination, and wavelength. The left-hand side of figure 2.4 shows the layout of the calibration unit. The calibration unit can be broken down first into two components: (i) the calibration target and (ii) the optical path which mimics the telescope beam and directs light from the calibration target onto the image slicer.

The calibration target assembly, shown on the right-hand side of Figure 2.4, consists of an integrating sphere - a diffusive element with three 38mm diameter entry/exit ports. Light is injected at two of these ports, and exits through a third port which is covered by a selectable mask apparatus. For flat-fielding, light from a continuum source is injected into one of the three ports, no mask is set on the exit port, and the diffuse emission illuminates the entire image slicer. To obtain the spatial geometry,

the continuum source is also used, but a mask consisting of evenly-spaced vertical slits is applied, resulting in a 'continuum bars' image on the detector. These bars provide the pipeline with the necessary information to reconstruct the image formed on the image slicer.

The third port on the integrating sphere is used to inject light from the spectral standards, which are hollow-cathode arc lamps containing an inert gas such as Thorium Argon. Thorium Argon is the standard source used for PCWI. However, KCWI can have both higher and lower spectral resolution than PCWI, so it is not immediately obvious that the same arc-lamp should be used. The first step in designing this module is therefore to determine which spectral standard(s) are needed to meet the instrument requirements.

### **Choice of Spectral Standards**

Wavelength calibration involves shining light from a source with a known spectrum (an 'atlas') onto the detector, then using the known template to reverse engineer where specific wavelengths fall on the detector, and what the wavelength dispersion (e.g. in AA per pixel) is throughout the 2D space. That model is then used to reproject the data onto a uniform wavelength axis as part of the transformation from 2D image coordinates to 3D data cube coordinates. In order to perform wavelength calibration, the known source must project detectable and resolvable features (i.e. emission lines) onto the detector. There must be sufficiently many features that the model is well constrained through the space. A single bright emission line in the center of the image will not provide sufficient information to constrain the wavelength dispersion as a function of wavelength.

Thorium-Argon was used as a spectral standard for the Palomar Cosmic Web Imager. However, PCWI has a minimum resolution of  $R \approx 2500$  (for the medium-resolution grating) and a maximum resolution of  $R \approx 5000$  (for the blue/red/yellow gratings.) KCWI, on the other-hand, can have a resolution as low as R = 900 and as high as R = 20,000. There is a hard constraint imposed by the lowest resolution setting; there must be enough resolvable features to successfully calibrate the data. The constraint imposed by the high-resolution mode, on the other hand, is a softer one. Resolving features will not be a challenge, but the flux per pixel will be lower. This can be compensated by an increased exposure time, which is what makes it a soft constraint, but from a practical standpoint, exposure times for obtaining an internal calibration image should ideally be in the regime of 10 - 30 seconds, not minutes.



Figure 2.5: Simulated low-resolution spectra, generated by convolving Thorium-Argon (ThAr) and Iron-Argon (FeAr) atlases taken from NOAO. This view shows the complimentary nature of the two spectra, with the blue region around  $\lambda \sim 3500 - 4200$  benefiting from distinct FeAr lines, while the more densely packed ThAr lines have blended, reducing their individual prominence.

In order to address the first constraint, spectral atlases were downloaded from the National Optical Astronomy Observatories (NOAO) website. The available spectra were Thorium-Argon, Iron-Argon, Helium-Neon-Argon and Copper-Argon. To simulate the calibration data, the spectra were convolved with a Gaussian with a full-width-at-half-maximum (FWHM)  $\Delta \lambda_{FWHM} = \lambda_c/R = 4500\text{\AA}/900$ , where  $\lambda_c = 4500\text{\AA}$  was taken as the central wavelength of the blue arm of KCWI.

Figure 2.5 shows these atlases and their simulated low-resolution KCWI spectra. It is clear that the densely-packed lines in the blue optical portion of the Thorium-Argon atlas, while used for PCWI, is not ideally suited for calibrating the lowest-resolution blue KCWI setting. As such, another source was required with more sparsely distributed features in the blue optical range. Iron-Argon appears to be a good fit for this setting and wavelength range, but is lacking in features for the KCWI-red bandpass. Together, the lamps provide good coverage of the entire KCWI bandpass.



Figure 2.6: The housing unit for the spectral standard arc lamps. Left: a view of the CAD model in SolidWorks. The only component that cannot be clearly seen in this view is the shutter, mounted on the hidden side of the front panel. Right: a photograph of the assembled unit in the lab.

## Arc Lamp Optics and Mechanical Housing

The space allocated for the spectral standard module was  $300 \times 160 \times 160$  mm<sup>3</sup>. The module would have to house two hollow-cathode arc lamps, each of which being approximately 150 mm in length, while focusing light from both into an aperture with a diameter of 38 mm. Both lamps need to be individually shutter-controlled and light-sealed from one another and the rest of the calibration unit. Although a number of orientations were tried, it proved impossible to fit both lamps in a configuration which would orient their brightest emission directly towards the aperture.

Instead, pairs of small fold mirrors had to be used to direct the light from each into the aperture of the integrating sphere. CodeV was used to determine the best layout of the optics, while SolidWorks was used to construct the mechanical housing around the apparatus for each lamp. The light from each lamp is focused using a plano-convex lens with a diameter of 60 mm and an effective focal length of 60 mm. Two enhanced Aluminum mirrors of diameter 2" and 1" in diameter, respectively, and oriented at 45 deg, were used to direct the focused light into the integrating sphere. The aperture in each half of the box was controlled by a single-bladed, bistable shutter which takes a simple 5V signal to open or close, mitigating the need for a more expensive controller. Figure 2.6 shows the arc lamp unit, as designed in SolidWorks (left) and as built (right).

### 2.5 Instrument Verification

In addition to designing and building the spectral calibration module for KCWI, this thesis contains contributions to the verification and validation of several subsystems within KCWI.

### **Red Arm Compatibility Weight Testing**

KCWI is being developed in a phased program that has the blue channel on sky at WMKO while detailed design and fabrication of the red channel is being completed. The red channel, or 'red arm', is expected to add an estimated 717 kg of mass to the blue instrument. The instrument weighs 4900 kg in the first phase of operation. Since the addition of mass will be significant, a test of the sensitivity of the instrument to the required additional weight has been executed to aid planning of the red arm integration. For this test, weights were added to the instrument in three separate load cases, interspersed with three unloaded cases, upon each of which calibration data were taken and later used to calculate the offset in alignment. We extracted multiple sections from pairs of the calibration images and used a simple 2-D cross-correlation script in Python to determine the offset between them. This test verifed that the instrument supported the combined mass of the red and blue arms, and that the addition of the red arm would result in image deflections that could be corrected without impacting image quality.

The test was conducted by taking two types of calibration images under four different weight-loading conditions. The first image type was an 'arc-bars' image, in which an arc-lamp spectrum is shone onto a 'bars' calibration mask (consisting of five vertical slits across the field of view.) This produces sources which are discrete both spatially and spectrally (i.e. in both axes of the image.) The second type of image was an arc-flat image, in which light from the arc lamp evenly illuminates the entire field of view, producing full 2D spectra on each slice. A set of images is taken before any weights are added, called 'unloaded test one' (U1). Then the first load test was performed (L1), followed by a second unloaded test (U2), the second loaded test (L2), and so on. Table 2.1 below summarizes the test sequence. Figure 2.7 shows the weights added for test L3 in panel (a) and illustrates the image offsets in panels (b) and (c) by showing a subtraction of a 'loaded' from a 'unloaded' image (L1 subtracted from U1).



Figure 2.7: Summarized figure from the verification report for Red Arm compatibility weight testing. Panel (a) shows an example of how the weight was applied to the instrument. Panel (b) shows a 'loaded vs. unloaded' subtraction of an arc-bars image, which is an arc lamp shone through a series of vertical slits, resulting in an image of 2D point sources (i.e. discrete spatially and spectrally). Panel (c) shows a similar subtraction, except with an arc-flat image instead of an arc-bars image. An offset can clearly be seen in both between the original (black) and subtracted (red) versions of the images.

Test	Weight [kg]	Weight Location/Purpose	$\Delta x [px]$	Δy [px]
U1	None	-	< 1	< 1
L1	79	Future fold-mirror 2	+6	< 1
	45	Red filter-exchanger	-	-
	136	Red ART stage (inside enclosure)	-	-
	181	Red ART stage (outside enclosure)	-	-
U2	None	-	< 1	< 1
L2	499	Outside enclosure	+8	< 1
U3	None	-	< 1	< 1
L3	227	Outside enclosure	+5	< 1

Table 2.1: Sequence of weight tests performed for KCWI red channel compatibility testing. The right-most two columns show the average offsets for each test along each image axis in units of pixels.

The requirement for this test was that the expected offsets due to the full 717kg of the red arm (extrapolated from the results of U1, L2, and L3) were sufficiently small that they could be corrected by a minor adjustment of the dichroic, which splits between the red and blue channels. The total offset was extrapolated to be 176.58 $\mu$ m at the focal plane, corresponding to a dichroic adjustment of 1.18 arcminutes. This value is within reason, so the test satisfied the requirement that the weight from the red arm could be added without impacting image quality.



Figure 2.8: Repeatability tests for the KCWI Blue Exchanger (BEX) Mechanism. The top panel shows the change in the x and y centroids over the course of the test. Each point represents the average value measured from 5 images, each containing 2500 sources. The bottom panel shows the same thing for the 1D full-width-at-half maximum of each source.

## **Blue Exchanger Mechanism**

A key aspect to the flexibility of KCWI is a mechanism that can automatically change a grating/filter from a predefined list of filters and gratings upon request from the observer. Repeatability requirements for these mechanisms are designed to ensure that spectral calibrations can be done during the daytime to reduce the instrument overhead time for observations made at night and more generally when the thermal conditions are within  $\pm 2K$ . Stability requirements for the same mechanisms ensure that imaging and spectral resolution are preserved during observations, so that photometric stability can be achieved at time scales of an hour to take advantage of the NAS capability of the instrument.

The Blue Exchanger (BEX) mechanism is composed of six axes of motion. Five of the six axes are linear stages plus one rotator stage. All stages except the two pusher stages, which are stepper motors, are controlled with servo motors. The tests consisted of generating a pseudo-random list of moves which would iterate through



Figure 2.9: Stability tests for KCWI's Blue Exchanger Mechanism. Both panels show the standard deviations of the centroid position and full-width-at-halfmaximum (FWHM) of sources in the images over a five minute period. The position appears to have a typical standard deviation of  $\sigma(X) \simeq \sigma(Y) \sim 0.03$  and the FWHM appears to have a typical standard deviation of 0.1px in both axes.

every possible configuration, pausing in between each move to take a set of calibration images. The calibration images taken were arc-bars images, which produce 2D (spatial/spectral) point sources. See panel (b) of Figure 2.7 for a zoomed-in example of five such sources. Broadly speaking, the tests were divided into (i) repeatability/hysteresis tests - checking how accurately each stage managed to return to the same position after a sequence of moves and (ii) stability tests - checking the consistency of the position and focus of the images while maintaining a single position. The repeatability was assessed by measuring the average centroid and full-width-at-half-maximum (FWHM) for each set of images taken between moves. Each one of these sets consisted of five images, each of which in turn contained approximately ~ 2500 sources. Figure 2.8 shows the repeatability measurements for the source position and focus (i.e. FWHM.) There appears to be a gradual upwards trend in the y-centroid across the ten sets of images, but the magnitude of this change is only approximately 0.5px in total over the full set of tests. None of the other measurements show any such trend.

Figure 2.9 shows the stability measurements for the centroid and FWHM of sources within each five minute period. The centroid seems stable to within approximately 0.03 pixels and the FWHMs stable to within approximately 0.1 pixels, though each pixel has a slightly elongated tail. These values fall within the acceptable range



Figure 2.10: Stability of the flux within each source over five-minute periods. The black histogram shows the standard deviation in the flux values for all measured sources as a percentage of the initial flux. The red histogram shows the expected distribution from shot noise and a read noise of 5 electons/pixel, while the blue distribution shows the same for twice the read noise.

for the test. Finally, Figure 2.10 shows the stability of the flux of each source over the same intervals. The change in flux from image to image is found to be mostly consistent with shot noise and a read noise of approximately 5 electrons per pixel. All tests were found to satisfy the stability and repeatability requirements. One requirement which was not met was the time taken for grating and filter changes, which initially had a soft requirement of  $\leq 2$  minutes, but which took ~ 3 and ~ 2.5 minutes, respectively.

### **Camera Articulation Stage Mechanism**

The KCWI Blue-arm Articulation Stage is responsible for moving the camera and detector through the spectral range of the instrument. Therefore it is important that the Articulation Stage meets the requirements set out in terms of accuracy, repeatability, and stability while station-keeping. Speed is also a relevant factor in the performance of this component, as it plays into the overall time taken for calibrations and for changing configurations during the night. For this test, in similar fashion to the Blue Exchanger Test, the articulation stage was moved between a series of angular positions (0, 25, 50, 75, and 100 degrees). A sequence of 15 moves was executed 7 times for a total of 105 moves, and throughout, the script kept track


Figure 2.11: KCWI Articulation Stage verification test results. The top left panel shows the time taken, in seconds, as a function of the move size, in degrees. The time taken at the required speed (1 deg/s) is shown as a dashed red line. The top right panel shows the net speed (deg/s) for a move versus the move size, with the minimum speed shown as a blue dashed line. The bottom panel shows the RMS deviation of the stage position, in arcseconds, as a function of position, in degrees.

of the target position, command time, and time of completion. After each move, the mechanism held its station for 180 seconds to test stability. The speed of the mechanism was found to consistently satisfy the requirement if 1 degree per second, and the RMS deviations of the stage while station-keeping were shown to be less than 0.4'' in all cases and less then 0.2'' in most.

# Nod-and-Shuffle Mechanism

The NAS observation technique is a central element to observations with KCWI, allowing for extremely accurate sky subtraction. This technique requires a mask to block the upper and lower thirds of the detector, with the center exposed to the incoming light. Since there is usually a comfortable amount of space between the edges of the inserted mask and the detector area needed for science, the science requirements are not highly sensitive to small fluctuations in the mask position.



Figure 2.12: The three types of test images taken for the NAS mask mechanism verification. In the top row, panel (a) shows the 'open' position, where the NAS mask is fully removed. Panel (b) shows the 'test' position, with the mask half inserted. Panel (c) shows the 'dark' position, with the mask almost fully inserted, such that the left edge falls on the first slice.

However, given the central importance of the mechanism and its likely frequent use, it pays to be thorough and rigorous in testing its repeatability. For this test, the NAS mask was set to a sequence of three different positions: 'open' - in which the mask is off the detector completely, 'test' - in which the mask covers just over half of the detector width, and 'dark' - in which the mask is almost fully inserted, but offset so that its left edge falls on an illuminated slice. At each position, a continuum flat image was taken. Figure 2.12 shows an example of a continuum flat image taken at each position. This set of three images was repeated for sixty iterations over the course of approximately ten hours and used to verify the repeatability and stability of the mask's position. In each image, the average position and slope (i.e. any rotation) of the left edge of the mask was measured. Figure 2.13 shows the centers and slopes from these tests. For the 'test' and 'dark' positions, respectively, the RMS value of the mask position is  $78\mu$ m and  $71\mu$ m, which is comfortably within the large margins on either side of the CCD region containing data.

## **KCWI Blue Detector Performance**

The performance of the KCWI detector is of critical significance to the scientific goals of the instrument. Alongside other factors such as throughput and quantum efficiency, low read noise and dark current are required to maintain a high level of sensitivity, while the amplifier gain (and thus full well value) determines the dynamic range of the detector. Also, in practical terms, short readout times are important to an observer, particularly for calibration sets. For this test, sets of continuum flat images were taken for a series of exposure times and used to generate photon transfer curves, allowing us to calculate the CCD gain and read noise. Figure 2.14



Figure 2.13: Results of the KCWI NAS mask mechanism test. The top panel shows the mask edge position across the iterations of removing/inserting the mask into the 'test' position. The second panel shows the slope of the mask edge for the same images. The bottom two panels are the same as the top two, except applied to the 'dark' position.



Figure 2.14: Photon transfer curves for the four KCWI amplifiers (UL = upper-left, UR = upper-right, etc.) in  $1 \times 1$  binning, slow readout mode at high gain. The horizontal dashed line represents the read noise component in each panel, while the sloped dashed line represents shot noise. Inset are the values for the read noise ( $\sigma$ ) and gain (k).

shows an example of these transfer curves for the quad amplifier mode. These tests were repeated for each amplifier mode, readout speed, and binning mode. Read noise was generally found to be approximately  $4 e^{-}/px$  for the fast readout mode, and approximately  $2.5 - 3.0 e^{-}/px$  for the slow mode. In all cases, we used KCWI's high gain mode. The data taken during these tests also allowed for readout times to be calculated. Table 2.2 below summarizes the readout times for each combination of binning and readout mode. The range in each combination is determined by the number of amplifiers used, with single-amplifier mode being the slowest and quad-amp mode being the fastest.

Long 'dark' calibration images were taken to measure the dark current and to get a measure of the rate of incidence of cosmic rays in the laboratory setting. Although the conditions were not perfectly dark, they were also used to obtain upper limits on the dark current for each amplifier. The dark was found to be  $\leq 3 e^{-1}/\text{px/hr}$  and

Binning	Mode	$t_r^{quad}$ [s]	$t_r^{dual}$ [s]	t <sub>r</sub> <sup>single</sup> [s]
$2 \times 2$	Fast	7	13	25
$2 \times 2$	Slow	27	53	106
$1 \times 1$	Fast	19	38	75
$1 \times 1$	Slow	85	170	337

Table 2.2: Readout times for KCWI measured during detector verification, in summarized form.  $t_r^{quad}$  is the readout time using all four amplifiers,  $t_r^{dual}$  is the readout time using two amplifiers, and  $t_r^{single}$  is the readout time using a single amplifier.

the rate of cosmic rays detected was approximately the 1 cr s<sup>-1</sup> cm<sup>-2</sup> expected.



### 2.6 On Sky Performance: PCWI and KCWI Comparison

Figure 2.15: Comparison of a 2hr PCWI observation (1hr source and 1hr sky in NAS mode) with a 1hr KCWI observation (standard observing mode).

While there are many more KCWI observations in this thesis, it would not be fitting to end this chapter without at least a cursory look at KCWI's on-sky performance. Figure 2.15 shows the FLASHES (O'Sullivan et al., 2020) target SDSS0958+4703 observed with both the Palomar and Cosmic Web Imagers. The PCWI data represents the result of two on-sky hours in NAS mode (i.e. one hour on source, one hour on sky) taken as a set of three 40-minute NAS sequences. The KCWI data represents one hour on target in standard mode, using a 10-minute reference sky image (and subsequent background correction) for sky subtraction. Both panels

show the same physical area and use the same color maps. It is immediately clear at a glance that the KCWI data is both higher resolution and significantly deeper. The morphology, kinematics, and spectra of both observations clearly agree, but the KCWI observation fills in much more of the fainter nebular emission (i.e.  $\leq 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ ).

# References

- Arrigoni Battaia, F. et al. (Jan. 2019). "QSO MUSEUM I: a Sample of 61 Extended Ly *α* Emission Nebulae Surrounding *z* ~ 3 Quasars." In: *MNRAS* 482.3, pp. 3162– 3205. DOI: 10.1093/mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Bridge, C. R. et al. (June 2013). "A New Population of High-z, Dusty Lyα Emitters and Blobs Discovered by WISE: Feedback Caught in the Act?" In: *ApJ* 769, 91, p. 91. DOI: 10.1088/0004-637X/769/2/91. arXiv: 1205.4030.
- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: Groundbased and Airborne Instrumentation for Astronomy V. Vol. 9147. Proc. SPIE, 91475K. DOI: 10.1117/12.2057056.
- Matuszewski, M. et al. (July 2010). "The Cosmic Web Imager: An Integral Field Spectrograph For The Hale Telescope At Palomar Observatory: Instrument Design And First Results." In: *Ground-based and Airborne Instrumentation for Astronomy III.* Vol. 7735. Proc. SPIE, 77350P. DOI: 10.1117/12.856644.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 *z* ~ 2.3 − 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.

# **CWITOOLS**

#### **3.1 Introduction**

The Palomar and Keck Cosmic Web Imagers (hereafter KCWI and PCWI) are integral field unit (IFU) spectrographs designed to study faint, extended emission (Matuszewski et al., 2010; Morrissey et al., 2018). PCWI was installed on the Hale 5m telescope at Palomar Observatory in 2009, while KCWI was installed on the Keck-2 10m telescope in the W. M. Keck Observatory in 2017. In 2014, the Multi-unit Spectroscopic Explorer (MUSE) (Caillier et al., 2014) was installed on the 8m VLT at the European Southern Observatory. This new set of instruments on 5-10m class telescopes has enabled observers to directly detect signals on the order of  $10^{-18}$  erg/s/cm<sup>2</sup>/arcsec<sup>2</sup> in less than an hour of telescope time (Martin et al., 2014a; Martin et al., 2014b). This in turn has enabled surveys of unprecedented sizes mapping the circumgalactic medium around high-redshift galaxies and quasars (Borisova et al., 2016; Arrigoni Battaia et al., 2019; O'Sullivan et al., 2020; Cai et al., 2019). As the observational field grows and sample sizes increase, data analysis becomes an increasingly prevalent issue. Here we present CWITools, a data analysis Python3 toolkit tailored to handling the data produced by the PCWI and KCWI data reduction pipelines (DRPs). This toolkit can be seen and used as a pipeline itself, taking input in the form of individual data cubes and producing scientific products such as white-light images, pseudo-narrow-band images, spectra, and velocity maps. CWITools was initially built out of necessity, as a toolkit for work on the FLASHES Survey (O'Sullivan et al., 2020); a survey of extended emission in z = 2 - 3 QSO environments. Over the past two years, it has been transformed into a publicly available, user-friendly package with help menus, documentation and application examples. This package is open source, and can be adapted to work with any three-dimensional data. However, in order to limit the scope for the purposes of testing and robustness, we focus primarily on applications involving data from PCWI and KCWI.

We begin by providing an overview of the context and architecture of CWITools, including a detailed description of the PCWI and KCWI pipelines' output. We



Figure 3.1: The architecture of CWITools, showing the broad pipeline from standard DRP results to scientific results, as well as the nature of the input and output of each module. Arrows represent different modules within the package and boxes represent the different types of data. For example, at the top, the reduction module (green arrow) takes standard DRP data cubes as input, and outputs corrected/coadded data cubes. On the right hand side are specific examples of the kinds of data products at each stage.

then describe the methodology of each broad processing step within CWITools; reduction, extraction, synthesis, modeling, and measurement. Finally, we provide an example application of CWITools in detecting nebular emission around a source at high redshift. In general, since full code-specific documentation and examples exist online for the package, we will limit the discussion here to design, methodology, application examples, and recommendations.

# 3.2 Package Architecture

CWITools is intended to provide a bridge between the output of the standard instrument data reduction pipelines (DRPs) and scientific products. In particular, PCWI and KCWI are both designed to image faint and diffuse emission, which lends a particular scientific focus to the package, though by no means an exclusive one. While there is no ubiquitous procedure which applies to all scientific projects, there are certain steps which are more or less universal to the process of extracting and measuring signals in IFU data: cropping, correcting 3D coordinate systems, masking or subtracting sources, extracting spectra or velocity maps, etc. CWITools is thus intended to provide observers with a flexible and easy-to-use set of tools with which they can customize a pipeline to suit their needs.

The standard DRPs have the goal of producing fully calibrated, three-dimensional data cubes for each individual exposure. Typically, observers reduce data on a night-by-night basis. Others may organize their data based on instrument configurations used or by target observed during a given observing run. For the purpose of scientific work, it makes the most sense to organize any subsequent analysis by science target. As such, the central element of any CWITools is a '.list' file which indicates simply: (i) where the input data for this target can be found, (ii) a set of unique IDs for files associated with this target, and (iii) where to save output and intermediate products.

The end point of CWITools is whatever scientific measurement is needed for the discussion and analysis part of a project. Whatever the scientific case - the functional architecture can be broken into the same main modular components: reduction, extraction, synthesis, modeling, and measurement. Reduction involves steps which are in essence further corrections of the data, required to compile the final observation (e.g. cropping, coadding). Extraction refers to steps which are focused on isolating a specific element or signal within the data (e.g. removing foreground or background). Synthesis is the generation of first-level scientific products from the isolated signal, such as surface brightness maps or velocity maps. Modeling is the fitting and evaluation of models, applied either to the generated scientific products or directly to the isolated 3D signal. Finally, measurement is the calculation of physical quantities from synthesized products, 3D data, or models. The distinction between synthesis and measurement in the context of this package is that the former produces 1D-3D data structures (e.g. spectra or velocity maps) upon which further measurements or analysis can be performed, while the latter produces scalar results which represent end-points in the data analysis process (e.g. luminosity). Although the term 'pipeline' implies a waterfall-like process through the above steps and the

bulk of pipelines will go through the above steps in some similar order, there is no strict one-directional flow imposed by the design of the package. There is always a point beyond which automation becomes more cumbersome than the alternative manual work required and, in the context of this package, that point is reached in determining the exact order of operations. Steps often need to be skipped, re-arranged, or repeated depending on the scientific objective and there is no 'one size fits all' data analysis pipeline. As such, while CWITools provides a number of template pipelines as examples, they are intended as templates to be modified and adapted to observers' needs.

Figure 3.1 shows the internal modular structure of CWITools, as well as the associated inputs and outputs of each stage. Each module corresponds to a Python module within the main package (e.g. cwitools.reduction), within which are functions associated with that stage. In addition to these modules, which can be imported into Python environments and used to construct a pipeline within any scripting environment, CWITools contains a library of command-line scripts which serve as wrappers to these functions. These scripts are designed as an interface for users who are less familiar with Python scripting, and simply want access to the tools. Upon installation, they are added to the user's environment as terminal commands (e.g. 'cwi\_coadd' becomes the direct command to coadd data cubes), each of which has a help menu guiding the user on its usage. Users who are not familiar with Python can then construct their data analysis pipelines using simple bash scripts containing a number of these commands. Short examples of both a Python environment pipeline and a bash script pipeline are included in the Appendix, and a full set of examples are available within the package data itself.

## 3.3 Cosmic Web Imager Data Format

The standard data reduction pipelines of PCWI and KCWI produce three-dimensional data cubes containing two spatial axes and one wavelength axis. In this section, we briefly describe the different output file types produced by the standard DRPs and the 3D coordinate systems of those data structures, so as to lay a foundation for the discussion of the methodology.

#### **Standard DRP Output**

The standard pipelines of both PCWI and KCWI first apply the usual reduction steps - bias correction, flat fielding, dark subtraction, etc. - to the raw 2D detector images.

Each 2D image contains 24 2D spectra (one for each slice in the image slicer) arranged side by side. Using a series of calibration images, the pipeline reconstructs these 2D images into 3D data cubes with two spatial axes and one wavelength axis. The first 3D data product is given extension "cube.fits," with each exposure producing a non-sky-subtracted object frame (*ocube*), a sky-subtracted intensity frame (*icube*), a frame of the sky data or sky model used (*scube*), and an associated 3D variance estimate (vcube). As the 3D data is refined through subsequent stages, the filenames are updated to reflect the stages. For example, after a slice-to-slice relativeresponse correction (stage6\_rr in the PCWI DRP), the letter 'r' is appended to each so that the files now have the extensions 'icuber.fits,' 'ocuber.fits,' 'scuber.fits,' etc. After flux calibration using a standard star, the appended letter is changed (rather than added) to 's,' so the filenames are now 'icubes.fits,' 'ocubes.fits,' etc. A full filename will include a unique identifier combined with one of these extensions. For example, the fully reduced, flux-calibrated cube for KCWI exposure number 116 might be 'kb200115\_00116\_icubes.fits.' This is of central relevance here because a core operational mode of CWITools involves providing as input a list of unique IDs (e.g. '200115\_00116' for this exposure) and a 'cube type' (e.g. icubes.fits) to work with. This allows users a simple interface with a high level of flexibility. For example, if an observer has three exposures for a certain target, they just need to store the three unique IDs for those exposures in a CWITools '.list' file, after which they can run any operation on any data product for that target by providing both the ID list and the desired cube type (e.g. coadd the intensity cubes, then coadd the sky cubes).

All flux-calibrated CWI data cubes are produced in "FLAM" units - i.e. flux per unit wavelength:  $F_{\lambda} \equiv \text{erg s}^{-1} \text{ cm}^{-2} \text{ Å}^{-1}$ . While PCWI outputs directly in  $F_{\lambda}$  units, KCWI data cubes are produced in units of "FLAM16,"  $F_{\lambda,16} = 10^{16}F_{\lambda}$ . Non-fluxcalibrated data cubes have units of 'electrons' - i.e. the number of photo-electrons measured in each voxel.

## **Coordinate Systems and Headers**

There are three coordinate systems which are of relevance when analyzing CWI data (and IFU data in general). The first is the *world coordinate system* (WCS) - which refers to the real world measurements of wavelength ( $\lambda$ ), right-ascension ( $\alpha$ ), and declination ( $\delta$ ). The second is the *image coordinate system*, referring to the axes within the data cube. Let us denote these as *x*, *y*, and *z*, where *x* and *y* are spatial axes

and z is the wavelength axis. It is important to note that the world-coordinate axes  $\alpha$  and  $\delta$  only correspond directly to x and y when (i) the position angle is a multiple of 90° and (ii) the field of view is small. In general, a one-to-one correspondence between image and WCS axes should not be taken for granted; i.e.  $\alpha \to \alpha(x, y)$  and  $\delta \to \delta(x, y)$ .

FITS image formats contain "header" objects which store meta-data about the image such as timestamps, configuration details, and exposure times. The headers also contain the necessary information to translate between the two above coordinate systems. Specifically, they contain sets of keywords to define (i) the number of axes and size of each axis, (ii) a reference point in the image for a known world coordinate (e.g. the right-ascension and declination at a given x, y), and (iii) the change in world coordinates along each image axis (e.g. the change in  $\alpha$  along the xaxis). Table 3.1 lists these keywords and their meanings. The widely used package Astropy provides a convenient way to handle coordinate systems by creating WCS objects which store this information and provide some useful functions such as mapping (x, y, z) to ( $\alpha$ ,  $\delta$ ,  $\lambda$ ) and vice versa (Astropy Collaboration, Price-Whelan, et al., 2018; Astropy Collaboration, Robitaille, et al., 2013).

As a final note on coordinate systems, a common source of confusion is the varying conventions when it comes to ordering axes and defining origins. The FITS headers for CWI data specify the axes such that the order of the axes is (1, 2, 3) = (x, y, w). However, when loading the data in a Python shell (e.g. with AstroPy or NumPy), the data structure has the order of axes reversed: (1, 2, 3) = (w, y, x). Furthermore, while the values in FITS headers are 1-indexed (i.e. the index of the first pixel is 1), data structures in Python are typically 0-indexed. This must be taken into account when handling header keywords such as CRPIX1, or converting between coordinate systems.

# 3.4 Module: Reduction

The reduction module is focused on steps for further refinement of the final observational data. This includes any steps beyond the standard date reduction pipeline which are required to create the final, fully calibrated, combined data cube for a given target. In this section, we will describe each of these steps in detail.



Figure 3.2: Automatic cropping parameters obtained by CWITools. This view is presented to the user if automatic cropping is requested. Each panel presents a summed one-dimensional profile for a different axis (x, y, z from top to bottom). Data within the cropped range, delineated by vertical red lines, is highlighted in blue, while data outside the range is black. This is most useful as a first step, from which the user can determine the best cropping parameters to suit their needs.

# Cropping

The output data cubes from both the PCWI and KCWI DRPs both require cropping along all three axes. While a user can determine the crop parameters they want to use, there are some defaults determined by the nature of the final PCWI and KCWI data cubes. The layout of the 2D spectra of the slices in detector space is such that alternating slices are staggered in wavelength. This means that the bandpass for each slice is slightly different, and the wavelength range which is common to all slices is slightly less than the instantaneous bandpass of any one slice. Only

Keyword	Description	Example
NAXIS	The number of axes	3
NAXISA	The length of image/data axis A	127
CTYPE <b>B</b>	The type of world coordinate axis <b>B</b>	RA–TAN
CNAME <b>B</b>	The name of world coordinate axis <b>B</b>	KCWI RA
CUNITB	The units of world coordinate axis <b>B</b>	deg
CRVAL <b>B</b>	The reference value for world coordinate axis <b>B</b>	255.25857
CRPIX <b>B</b>	The reference pixel along image axis <b>B</b>	32
CDA B	Change in world axis <b>B</b> per pixel of image axis <b>A</b>	$8.09 \times 10^{-5}$

Table 3.1: A summary of FITS Header keywords for world coordinate systems. Correcting or adjusting a world-coordinate system is most easily done by updating these values, as opposed to shifting or rotating the data itself.

data within this common overlap region can be reliably calibrated by the DRP. The headers of KCWI and PCWI FITS files contain the keywords 'WAVGOOD0' and 'WAVGOOD1' which indicate this range. Therefore, as a default, CWITools will crop the wavelength to this range.

Spatially, there are different reasons to crop PCWI and KCWI data. For PCWI, the x-axis (FITS axis 1, NumPy axis 2) of the data contains some buffer, going slightly beyond the edge of each slice. The exact margin can be determined by looking at a fully reduced cube from a continuum flat image, but is usually approximately about 10 pixels on either side. There is no padding or margin along the y-axis, which contains only the 24 slices. For fully reduced KCWI data, there is padding if stage 7 (differential atmospheric refraction correction) has been applied, and the amount of padding differs depending on the slicer setting, so there is no hard-coded default for the spatial padding. Instead, the 'auto-crop' mode (i.e. the default used in absence of user input) is to trim empty rows and columns. Users should be aware that this may not be sufficient to avoid edge artifacts which may be present in the data. The cropping tool has a plot functionality which can be used to view the crop settings overlaid on a profile of each axis. This is a helpful tool in selecting the best settings. As with all steps that modify existing cubes, the output is saved by default with a modified file extension. In this case, the ".fits" of the input is replaced with ".c.fits," to indicate that it has been cropped.



Figure 3.3: Automatic WCS correction using source fitting. As one option for spatial WCS correction, CWITools assumes that the initial WCS is approximately correct, then identifies and fits the nearest source to that location. The above view is the view presented to the user during this step if requested, so that the user can inspect the fit visually. The top left panel shows the full field of view, with a white cross indicating the expected location of the primary source. The white box indicates the search area, which can be adjusted by the user. The red cross indicates the fitted location of the source. The top right panel shows a zoom in on the white box. The bottom two panels show the one-dimensional PSF of the source along each axis and a simple 1D Gaussian fit to the data.

# World Coordinate System Correction

As described earlier, the world coordinate system (WCS) is the three-dimensional coordinate system of right ascension ( $\alpha$ ), declination ( $\delta$ ), and wavelength ( $\lambda$ ). FITS headers for three-dimensional data contain reserved key words which determine the translation from image coordinates (x, y, z) to world coordinates ( $\alpha$ ,  $\delta$ ,  $\lambda$ ).



Figure 3.4: Cropping and WCS correction applied to an individual cube. Left (a): a spatial 2D snapshot of the data cube prior to cropping and WCS correction. The red circle has the correct coordinates for the source, SDSS0958+4703. Right (b): the cropped and WCS-corrected cube, the red circle now aligns with the visible source.

To constrain the transformation, several pieces of information are provided. The first is the position of a specified world coordinate in image coordinates; i.e. that the 3D position ( $\alpha_0$ ,  $\delta_0$ ,  $\lambda_0$ ) is coincident with ( $x_0$ ,  $y_0$ ,  $z_0$ ). The values  $\alpha_0$ ,  $\delta_0$ , and  $\lambda_0$  are stored in 'CRVAL' keywords, where the central value for each specified world-coordinate axis is given. Since the order of the world coordinate axes is usually right-ascension, declination, and wavelength, this means the header contains 'CRVAL1= $\alpha_0$ ,' 'CRVAL2= $\delta_0$ ,' and 'CRVAL3= $\lambda_0$ .' It is important to clarify that the numbering here is for the values associated with the world-coordinate axes, not with image axes. The coincident image coordinate is stored in 'CRPIX' keywords: 'CRPIX1= $x_0$ ,' 'CRPIX2= $y_0$ ' and 'CRPIX3= $z_0$ '. Finally, the change in each world coordinate axis along each image is provided by the "CD Matrix," which is a set of keywords of the form 'CDW\_I,' which specifies the change in world-coordinate axis W per pixel along image axis I. As an example, 'CD1\_2' encodes the change in right-ascension per pixel along the *y*-axis. The units of these values are given by keywords 'CUNIT1,' 'CUNIT2,' and 'CUNIT3.'

Generally, the CD Matrix can be taken as accurate for all PCWI and KCWI data cubes. The only rectification that is usually required is an adjustment of the central reference point. This is done separately for the spatial and wavelength axes. In each case, there are two basic approaches to choose from: measuring the location of a feature with known world-coordinates or cross-correlating the input data so that they are at least aligned. The former provides a correction in absolute terms, but requires a measurable source with known coordinates, which is not always available. The latter provides a fall-back for these cases, such that the input data can be ensured to have consistent world-coordinate systems, but the absolute values may remain

inaccurate. Each of the four processes is described below. Cubes with corrected coordinates systems are saved by default with the added file extension ".wc.fits" (for WCS-corrected).

#### **Spatial Correction: Source Fitting**

The preferred way to correct the spatial axes is to measure the location of a known source within the image. This is done by first creating a white-light image from the input data (see Synthesis module in Section 3.6). The default operating assumption is that the initial WCS is approximately correct, and a  $10'' \times 10''$  box around the estimated location of the source is extracted. In the case that the initial WCS is extremely inaccurate, an initial guess of the source location can be provided, and the size of the box can be adjusted. Once the box has been extracted, 1D profiles in *x* and *y* are formed by summing along the image axes, and a 1D Moffat profile is fit to the source to obtain the best-fit center. CRPIX1 and CRPIX2 are updated to the *x* and *y* centers, respectively, and CRVAL1/CRVAL2 are updated to the known RA/DEC of the source.

### Wavelength Correction: Line Fitting

To correct the wavelength axis, a known sky-line can be fit with a simple Gaussian model. The default way to do this is with sky cubes (e.g. "scubes.fits") and known sky emission lines. CWITools package data includes a full blue-optical sky spectrum for Keck, and a preliminary set of known emission lines in both the Palomar and Keck blue-optical sky spectra. For example, there is a bright mercury line (thanks to light pollution) in the Palomar night sky, Hg I  $\lambda$ 4358.3. This line was extremely useful in correcting the Palomar data for the FLASHES Pilot Survey (O'Sullivan et al., 2020). A high-SNR sky spectrum is compiled from the input sky cube by summing over both spatial axes. As in the spatial PSF fitting, the default assumption is that the initial estimate of the sky-line is extracted from the spectrum, and a 1D Gaussian model is fitted to the data to obtain the true center. The difference between the initial WCS' estimate of the line position and the fitted position is calculated in units of pixels, and the header keyword CRPIX3 is updated accordingly so that the WCS is consistent with the measured position of the source.

#### Wavelength Correction: 1D Cross-Correlation

If no spatial source or sky line is available, cross-correlation can be used to ensure that the input data are all self-consistent and aligned, even if the absolute world-coordinate solution is not exactly known. To do this, as in the 2D cross-correlation above, one image must be picked as the reference point. For each sky cube, a 1D spectrum is then generated and each spectrum is cross-correlated with the reference spectrum. A list of relative offsets, in units of pixels, is then calculated between the spectra. The CRPIX3 header keywords in all but the reference image are updated based on the measured offsets so that they are consistent with the WCS of the reference image. Any error in the reference image's WCS thus remains in the corrected WCS.

### **Custom WCS Correction Routines**

Certain science cases, such as creating mosaics of very faint emission, satisfy none of the above requirements and require more advanced, home-made methods. As such, CWITools separates the WCS measurement and WCS correction steps. WCS measurement produces a WCS correction file which contains a table of the desired values for the CRVAL and CRPIX keywords for each input cube. If using either of the in-built methods (feature fitting or cross-correlation), this is automatically generated by the WCS measurement function. If using a more complex method, or in the case that some manual adjustment is required, a user can generate this table themselves by whatever method they desire. As long as the format of the table (which is quite simple) is correct, the table file can be fed as input into the 'apply\_wcs()' function. Because CWITools uses the WCS to automatically coadd data, this step is crucial in determining the quality of the final, coadded data.

#### Coadding

The single most ubiquitous and (relatively speaking) computationally intensive task in PCWI and KCWI data analysis is coadding data onto a common three-dimensional grid, including variance propagation. There is, of course, already a plethora of openly accessible and efficient Python code which performs some subset of this task, but there are several key issues which led us to develop an entirely custom algorithm.

First, nearly all of the widely used packages implementing coadding or drizzling



Figure 3.5: Spatial 2D slices of coadded frames using different drizzle factors. In panel (a) on the left, the coadd was performed normally, with no drizzling (i.e.  $f_{drz} = 1.0$ ), while in panel (b), a very low Drizzle factor of  $f_{drz} = 0.4$  was used. Normally, one would use a factor in the range  $f_{drz} \sim 0.6-0.8$ , but we use a low factor here to make the visual difference clear. The image in panel (b) is clearly sharper with higher frequency noise, while that in panel (a) looks smooth in comparison.

algorithms (e.g. Avila et al., 2015) are written with two-dimensional imaging data in mind. Second, knowledge of every computational step is needed in order to accurately propagate variance through the coadd process. This is trivial if the data only needs to be resampled and shifted linearly along its axes, but coadding images with arbitrary rotations makes the resampling - and thus the mathematics of error propagation - significantly more complicated. As such, in order to have an algorithm which can coadd arbitrary input in terms of position angle and spatial sampling, we have developed a custom three-dimensional coadding algorithm, including a drizzling factor. As a final note, CWITool's coadding algorithm makes use of existing 3D mask cubes produced by the PCWI and KCWI pipelines - automatically loading and using them if the user requests it. These cubes flag noisy edge pixels, pixels affected by cosmic rays, and other potentially corrupted pixels. These can then be excluded from the coadd, improving the quality of the final product.

The CWITools coadd process is split into two main steps. First, the input cubes are aligned in wavelength. At the moment, only input with a common sampling rate in wavelength is accepted, since this is by far the most common scenario, though a future update is planned to allow multiple wavelength sampling rates in the input. The minimum and maximum input wavelengths are determined ( $\lambda_{min}$  and  $\lambda_{max}$ ),

as well as the input resolution,  $\Delta\lambda$ . A new common wavelength grid is generated spanning the range  $[\lambda_{min} - \Delta\lambda, \lambda_{max} + \Delta\lambda]$  with the same resolution. Then, for each input cube, the cube is padded with zeros along the z-axis until it is the same length as the common wavelength axis, and the offset (in Angstrom) between the first index of the original input grid and the first index of the new common grid is determined. Let  $\delta\lambda_i$  be this offset. The required shift in pixels is then  $\delta z = \delta\lambda_i/\Delta\lambda$ . This is split into an integer shift,  $\delta z_{INT} = int(\delta z)$ , and a sub-pixel shift,  $\delta z_{SUB} = \delta z - \delta z_{INT}$ . The integer shift component requires no interpolation, and thus no error propagation. The data and associated variance are just rolled along the z-axis by  $\delta z_{INT}$ . The sub-pixel shift is then performed using linear interpolation, implemented as a convolution with the 1D kernel  $K_z = [\delta z_{SUB}, 1 - \delta z_{SUB}]$ . To propagate the error on this step, the variance is convolved with  $K_z^2$ . At the end of this step, the cubes are all aligned in wavelength and have the same z-axis length.

The second major step is the spaxel-by-spaxel projection of the input cubes' footprints onto a common coadd grid. The on-sky footprints of each input field of view are calculated and the footprint required to encompass all of the input data is derived. The minimum spatial sampling of the input is taken as the uniform spatial sampling of the desired output grid. This information is then used to construct a new header and empty data cube for the coadded data. The 2D (x,y) vertices of each input pixel are mapped from input image coordinates to on-sky coordinates using the input WCS (with Astropy's WCS class). The on-sky coordinates of these vertices are then mapped to output image coordinates using the newly constructed WCS. The footprint of the input pixel on the output frame is then represented as a Polygon object, using a Python package called Shapely. The coadd frame pixels within this footprint are also represented as Polygons, and the overlapping area between the input pixel and each output pixel is calculated. This step is computationally intensive, but provides a high level of flexibility and robustness to the coadd method, as the polygons are entirely flexible in shape and orientation. In particular, this allows us to implement a 'Drizzle' factor, shrinking the size of the input pixels by a certain amount (typically to 70 - 80% the original size) to increase the spatial sampling of the coadd (Avila et al., 2015). Time is also not a major constraint for the typical use-case of CWITools coadding, as it only needs to be performed a small number of times per target. That said, the process still only takes about twenty seconds to add 3-4 high-resolution KCWI cubes, including error propagation and masking. The contributions from the individual input cubes are weighted by exposure time,

*E*. Thus, if we let the index *i* iterate over the input cubes, and the indices *j*, *k* iterate over the two spatial axes, such that  $x_{ij}$  is the  $j^{th}$  x-pixel of the  $i^{th}$  cube, then the final coadded flux is given below in Equation 3.1.

$$F_{coadd}(x, y) = \frac{\sum_{i} E_i \left[ \sum_{j} \sum_{k} F_{in}(x_{ij}, y_{ik}) f(x, y, x_{ij}, y_{ik}) \right]}{\sum_{i} E_i}$$
(3.1)

Here,  $f(x, y, x_{ij}, y_{ik})$  is the fraction of the footprint of the input pixel  $(x_{ij}, y_{ik})$  that falls on the output pixel (x, y). Since the wavelength axes have been aligned, and the process here is applied at all wavelength layers, the third axis has simply been dropped from the notation. The propagated variance is then as shown in Equation 3.2.

$$V_{coadd}(x, y) = \frac{\sum_{i} E_{i}^{2} \left[ \sum_{j} \sum_{k} V_{in}(x_{ij}, y_{ik}) f^{2}(x, y, x_{ij}, y_{ik}) \right]}{(\sum_{i} E_{i})^{2}}$$
(3.2)

It should be noted that the interpolation involved in the wavelength alignment and flux redistribution introduces additional covariance in the coadded cube. While CWITools does not currently have a built-in way to calibrate covariance, a full section is dedicated to discussing this in O'Sullivan et al. (2020). Some pixels at the edge of the coadded field of view may only be partially covered by the input data after all of the input cubes have been added. A threshold is (optionally) applied to reject any edge pixels with very low coverage, and referred to in the code as 'pxthresh' - meaning pixel coverage threshold. Setting this to a high value (i.e. 0.9) will mean that only more-or-less fully covered pixels are included.

As a final step, empty rows, columns, and wavelength layers are trimmed from the coadded data. Here, a second (optional) threshold is applied based on exposure time. If the input data has any spatial dithering, some spaxels will have longer total exposure times than others. Spaxels with significantly lower exposure times will appear noisier and may complicate analysis. The parameter 'expthresh' sets the minimum exposure time (as a fraction of the maximum) for a spaxel to be included in the data. This threshold is only applied as an extension of the trimming; that is, rows and columns on the edge of the field of view with lower exposure times will be trimmed from the data. It does not remove or mask regions throughout the data arbitrarily. As an example, if three equal-length exposures are taken with a 0, -1'', +1'' dithering pattern along the x-axis, the edge -1'' regions along that axis in the



Figure 3.6: Illustration of the exposure and pixel coverage thresholds in CWITools' coadd function. Panel (a) shows an illustration of two overlapping fields of view with equal exposure time. The numbers in each area represent the local stacked exposure time relative to the maximum stacked exposure time. These are the values considered when applying the exposure threshold. Panel (b) illustrates the pixel coverage threshold. The white grid represents the coadd pixel grid, while the red rectangle represents the footprint of an input pixel.

coadd will have fractional exposure times of 1/3. Setting expthresh to 0.5 would eliminate these regions and keep only the overlapping central part of the field of view. Coadded cubes are saved by default using the same name as the ".list" file, though the output filename can be specified during usage.

#### Variance Estimation and Scaling

While KDRP and PDRP produce 3D variance estimates for each exposure, there may occasionally be a need to estimate the variance from the data itself. For example, if some procedure is performed on the data for which the error propagation is prohibitively complex or if there is some problem with the data that affects the pipeline estimate, it may be preferable to estimate the variance empirically. In these cases, the basic approach, which is broadly similar to that used in Borisova et al., 2016, is to first estimate a 2D variance map by taking the variance along the z-axis and then scale that 2D variance map to match the noise properties of each wavelength layer in the data. This is complicated by the presence of real signal in the data so, as a first pass, let us assume that the cube is dominated by noise and that the number of voxels containing such signal is negligible.

The cube is first divided into bins of size  $\Delta z$  pixels (i.e. wavelength layers). For each bin, the variance is taken along the z-axis to produce a local estimate of the (x,y) variance. Then, for each layer, the distribution of SNR values using this variance estimate is calculated. The following step relies on an assumption that the noise is Gaussian (or at least approximately Gaussian) in form as it involves the assertion that the distribution of SNR should follow a standard normal distribution ( $\mu = 0, \sigma = 1$ ). This, in turn, relies on the assumption that (i) the distribution is dominated by background pixels and (ii) the noise within the background of the data is Gaussian in nature. These are both reasonable assumptions for a long (i.e. sky-limited) exposure containing no bright sources and with only a small fraction of the voxels/spaxels occupied by real emission. This assumption obviously breaks down under different circumstances, which we will discuss shortly. If the measured distribution of SNR values in this layer has standard deviation  $\sigma_i$ , then the variance rescaling factor is  $r_{var} = 1/\sigma_i^2$ .

The assumption of a standard-normal distribution only applies for background regions which are shot-noise limited. In the case where the input data contains large regions of bright emission, these regions must be masked and excluded from the SNR distribution which is used to calculate the scaling factor. Ultimately, there must still be a reasonably large background region - confidently free of real signal to justify use of this method. Otherwise, this method should not be used to estimate the variance. If a mask is provided, two restrictions apply. The first is that within every z-axis bin, there must still be enough unmasked wavelength layers to obtain the variance along the z-axis in every spaxel. If some regions of the mask are very extended in wavelength, the  $\Delta z$  parameter should be increased to ensure that enough unmasked layers remain in each bin. This, of course, reduces the accuracy of the local 2D variance estimate, but is a necessary step. The second restriction is that every wavelength layer must still contain a sufficient number of unmasked spaxels to obtain a reliable distribution of SNR values. If this is not the case, again, this method should not be used. As a last resort, if both of the above restrictions cannot be met, a single scaling factor can be applied to the entire variance cube by combining all background voxels into a single distribution.

If a user already has a variance cube, but believes it needs to be scaled (e.g. to



Figure 3.7: Variance scaling by assertion of a standard normal distribution in 'background' regions. The grey shaded histogram shows the distribution of SNR values based on the input data and variance. Note the logarithmic scale on the y-axis. The green histogram shows the distribution of SNR values after large, contiguous 3D objects (either systematic residuals or real emission regions) have been detected and masked. The dashed black line shows a standard normal distribution, and the solid black line shows the best-fit Gaussian model used to calculate the re-scaling factor.

account for covariance introduced by coadding), then the initial variance cube can be provided and only the rescaling part of the algorithm will be applied. Estimated and scaled variance cubes are saved by default with the extension ".var.fits."

### **Slice-to-Slice Scattered Light Correction**

In image-slicer integral field units, each slice of the field of view is sent along a different optical path. Part of the standard DRP's job is thus to correct for the differing relative response (i.e. throughput) of different slices, which can be caused by dust on the slices or pupil mirrors, or edge-of-field effects. One sliceto-slice correction that falls beyond the scope of the standard KCWI and PCWI DRPs involves scattered light. Slices containing very bright sources can sometimes



Figure 3.8: Calibration of covariance in a single KCWI data cube. The panel on the left shows the ratio between the observed noise and the propagated noise assuming no covariance ( $\sigma_{obs}$  and  $\sigma_{ideal}$ ) - after the binning the data - as a function of bin size. The red line shows the best-fit two-component model, with  $\sigma_{obs}/\sigma_{ideal} = C(1 + \alpha \log(K))$  for  $K \le K_{thresh}$  and  $\sigma_{obs}/\sigma_{ideal} = \beta$  for  $K > K_{thresh}$ . The right panel shows a histogram of the fractional residuals.

contain an additional, relatively flat scattered light component across the slice. To remove this, CWITools runs through the 1D profile of each slice at each wavelength layer, estimates the background level, and subtracts it. The estimate is made by conservatively sigma-clipping the 1D profile to remove bright sources and then taking the median of the remaining pixels. This method is suitable for relatively clear fields with a single, very bright source. In fields with multiple sources, it can be difficult to obtain a reliable background estimate and this method should only be used with appropriate caution. Slice-corrected cubes are by default saved with the extension 'sc.'

# **Air-to-Vacuum and Heliocentric Corrections**

Lastly, the reduction module contains two common corrections for the wavelength axes of input data: conversion from air wavelengths to vacuum wavelengths, and a heliocentric velocity correction. For the former, CWITools uses an implementation from the package PyAstronomy to convert the wavelength axis. Because the correction from air to vacuum wavelengths depends on wavelength, the data must be interpolated onto the new wavelength grid, using either linear or cubic interpolation. Error propagation is not yet available for this function, so the variance should be rescaled or estimated anew after the application of this change. For the heliocentric correction, CWITools uses Astropy's SkyCoordinate class, and offers a choice between updating the header keywords to modify the wavelength axis, or to keep the original wavelength axis and shift the data using interpolation.

### 3.5 Module: Extraction

The extraction module could just as well be called the 'isolation' module, as the ultimate goal is to isolate a specific signal, be it a point source, extended source, continuum emission, or line emission. While specifics may vary, as always, there are a few more or less ubiquitous steps in this process. First and foremost among them is the removal of point sources by modeling of the point-spread-function. Second is the removal of any unwanted component of the emission which is slowly varying both spatially and spectrally, referred to loosely as 'background subtraction.' Masking, smoothing, and segmenting the data (into contiguous regions above a threshold) are also common steps towards isolating a signal. In this section, we describe the CWITools implementation of each of these.

# **PSF Subtraction**

PSF subtraction requires first modeling the PSF in 3D and then subtracting the model. Analytical PSF models such as Gaussian or Moffat profiles provide robustness of shape, which helps when trying to avoid overfitting, and are more well suited to fitting blended sources. However, the real instrument PSFs in PCWI and KCWI are more complex than a simple Gaussian or Moffat. As such, relying on these models for PSF subtraction leads to significant, systematic residuals. Systematic errors can be significantly worse than random error as they run the risk of creating false negatives and false positives. As such, the CWITools PSF-subtraction follows an empirical approach, some variant of which is widely used in existing observational work (Arrigoni Battaia et al., 2019; Cai et al., 2019; O'Sullivan et al., 2020).

The most common reason for performing PSF subtraction in CWI data is to disentangle point sources and extended, nebular line emission. The key property of nebular line emission that enables this particular method is that it is spectrally confined to relatively narrow portions of the overall bandwidth. The empirical approach to building a 3D PSF model involves constructing a 2D model of the PSF by summing over wavelength layers which do not contain nebular emission ('continuum wavelengths'), then scaling it to match the PSF in each wavelength layer. The benefit of this approach is that arbitrarily complex instrument PSFs can be reliably subtracted, provided the shape does not change strongly as a function of wavelength. The most significant drawback of this approach is that it struggles to handle blended PSFs of two or more sources, as adjacent sources will be included in the empirical model. It is also not well suited to separating a diffuse/resolved continuum source from an unresolved continuum source. The way to achieve these goals following an empirical approach is to use an isolated source to obtain a PSF model. While this may be added as an option in a future update of CWITools, the current version focuses on extended nebular emission.

In the current implementation of this method, the user can specify a series of wavelength regions to exclude from WL images. For each wavelength layer, a new WL image is calculate by summing over a window of width  $\delta \lambda_{WL}$  centered on the current layer. If the window is clipped on either side by the limits of the z-axis or there are masked layers within it, it is grown incrementally until the 'effective' window size (i.e. the total *useable* bandwidth) is equal to  $\delta \lambda_{WL}$ . This ensures that the number of wavelength layers summed for each WL image is consistent, ensuring that the SNR of the PSF model also remains roughly consistent (it will of course vary anyway depending on the spectrum of the continuum emission). Pixels within a radius of  $r_{min}$  (default value is typically  $r_{min} = 1''$ ) from the center of the source are used to calculate a scaling factor for the PSF at each wavelength layer. The scaled model is then subtracted from the layer out to a radius of  $r_{max}$ , which is typically set to 2-3 times the seeing (i.e.  $r_{max} \sim 5''$ ).

CWITools has two modes of PSF subtraction: 1D and 2D. In the 2D version, the above process takes place using full 2D white-light images and circular regions of radius  $r_{min}$  and  $r_{max}$ . In the 1D method, the PSF models are created, scaled, and subtracted on a slice-by-slice basis. This is motivated by the fact that, for bright sources, there can be a significant scattered light component which is slice-dependent, and thus better fit by a model for that slice alone. The same  $r_{min}$  and  $r_{max}$  are used, only now in a 1D sense and for each slice.

In either method, variance data can be provided. If it is, the variance on the PSF model is calculated and error is propagated throughout the subtraction. Final output



Figure 3.9: PSF-subtraction of a bright source to isolate extended emission. The top panel (a) shows a section of the 2D spectrum of a slice containing the bright source. The bottom panel shows the same 2D spectrum, with the same color map, after PSF subtraction. Bright, extended  $Ly\alpha$  emission can be seen clearly after the subtraction. The small, bright spots to the left (blue) side of the extended emission are systematic residuals caused by sharp  $Ly\alpha$  forest absorption features, where the PSF shape changes rapidly as a function of wavelength. The central pixels used to scale the PSF are masked, as these cannot be used to measure signal.

is saved with the extension ".ps.fits" (for PSF-Subtracted) and ".ps.var.fits" for the associated variance. Optionally, the PSF model can also be saved.

# **Background Subtraction**

As mentioned briefly above, the goal of background subtraction (BGSub) is to separate and remove any slowly varying component, spatially or spectrally. Examples of such signals include residuals left over after an imperfect sky subtraction or extended continuum emission from a (foreground) galaxy. As such, the term 'background' is used quite loosely. There are many ways to approach removing background signals. CWITools has three methods to choose from at the time of writing: (i) polynomial spectral fitting, (ii) median filtering, and (iii) simple median subtraction.

### **BGSub Method 1: Polynomial Fitting**

The 'polyfit' method fits a low-order polynomial (i.e. polynomial degree  $k_p \approx 1-5$ ) to the spectra in each spaxel. This method is probably the best choice for applications involving narrow-line nebular emission within data where the total bandwidth is large relative to the width of the emission. In such a scenario, the emission line features will be too small to be over-fit by such a slowly varying polynomial, and the fit will be dominated by continuum wavelengths. However, if the total bandwidth is similar to the width of the emission, even low-order polynomials will be more susceptible to over-fitting. In either case, wavelength regions known to contain emission lines can be masked by the user, ensuring that the polynomial is only fit to the background



Figure 3.10: Integrated spectrum of a region before (top) and after (bottom) PSF and polynomial background subtraction. Black vertical bands indicate the locations of noisy residuals from extremely sharp absorption features or bright sky line, both of which can be masked. Vertical red lines, from left to right, indicate the positions of emission lines: HI Ly $\alpha$   $\lambda$ 1216, NV  $\lambda\lambda$ 1239, 1243, SiIV  $\lambda$ 1394, CIV  $\lambda\lambda$ 1548, 1551, HeII  $\lambda$ 1640.

signal.

If the background polynomial model is:

$$p(k,\lambda) = \sum_{i=0}^{k} c_i \lambda^i,$$
(3.3)

then, assuming that the wavelength of each layer is a well known quantity, and the only variance comes from uncertainty in the coefficients, the variance on the background model can be written as:

$$V(k,\lambda) = \sum_{q=0}^{q=k} \left( \frac{\delta p(k,\lambda)}{\delta c_q} \cdot \delta c_q \right)^2.$$
(3.4)

The partial derivative expands to:

$$\frac{\delta p(k,\lambda)}{\delta c_q} = \sum_{i=0}^{i=k} \left( \lambda^i \frac{\delta c_i}{\delta c_q} + c_i \frac{\delta \lambda^i}{\delta c_q} \right) = \sum_{i=0}^{i=k} \lambda^i \frac{\delta c_i}{\delta c_q},$$
(3.5)

where again, wavelength is assumed to be a well known quantity. Plugging this in

to the expression for the variance, we get:

$$V(k,\lambda) = \sum_{q=0}^{q=k} \left( \sum_{i=0}^{i=k} \lambda^i \frac{\delta c_i}{\delta c_q} \cdot \delta c_q \right)^2.$$
(3.6)

The covariance matrix for the polynomial coefficients is returned by the polynomial fitting routine used (NumPy's polyfit). The off-diagonal covariance terms are typically very small compared to the diagonal terms. As such, the simplifying approximation can be made that the variables are independent, such that  $\delta c_i / \delta c_q = \delta_j^i$ , where  $\delta_j^i$  is the Kronecker delta function. This reduces the above expression to

$$V(k,\lambda) = \sum_{q=0}^{q=k} \lambda^{2q} (\delta c_q)^2 = \sum_{q=0}^{q=k} \lambda^{2q} \operatorname{Var}(c_q).$$
(3.7)

The variance estimate can be re-scaled following Section 3.4 if needed to account for added covariance in the data.

#### **BGSub Method 2: Median Filtering**

Another common method to subtract background signals is median filtering along the wavelength/z-axis. The background is estimated by means of a median filter with window size  $\Delta z \ge 5$ , in pixels, wherein the value at each z-index is replaced with the local median. As in the polyfit method, wavelength regions can be masked to exclude them from contributing to the estimate of the local median. However, the window size must be sufficiently large that at least five unmasked pixels remain in the window, or a median is not well defined. While median filtering can be useful in scenarios where a polynomial fit performs poorly, median-filtered models are less well behaved in the sense that they can produce sharp, small-scale features and change discontinuously from one index to the next. Especially in the context of masking certain wavelengths, they should be used with caution and the background model (which can be saved as an option in CWITools) should be inspected.

The variance on the mean in a given window of size  $N_z$  centered on index *i*, as a function of the existing variance estimates, is:

$$V(N_z, i) = \frac{1}{N_z^2} \sum_{j=i-\Delta z/2}^{j=i+\Delta z/2} V_j$$
(3.8)

where  $V_j$  is the input variance at index j. The ratio of the variance on the mean to the variance on the median is  $4n/(\pi(2n + 1))$  where  $N \equiv 2n + 1$  is the total sample size. This can be rewritten in terms of N as  $\pi N/(2(N - 1))$ . Thus, the variance on the *median* background model is:

$$V_{med}(N_z, i) = \frac{\pi}{2N_z(N_z - 1)} \sum_{j=i-N_z/2}^{j=i+N_z/2} V_j.$$
 (3.9)

#### **BSub Method 3: Median Subtraction**

Where the median filtering method works along the z-axis, the simpler 'medsub' method subtracts the spatial median at each wavelength layer. This is most useful in data that (i) is well flat-fielded and (ii) contains mostly background spaxels at all wavelength layers. The latter constraint can be relaxed if a mask is provided by the user to exclude non-background regions, but the remaining background region should still be sufficiently large at all wavelength layers to obtain a reliable median. Since the variance propagation on this is a simpler version of that described in the median filtering method above, it does not need to be described in detail again.

# Masking

Masking can be necessary after the subtraction of some point sources, where relatively small errors in the shape of the PSF near the core can still lead to loud residuals, or in order to mask the locations of foreground sources to exclude them from subsequent analysis. CWITools' extraction module contains a method for creating 2D binary masks using DS9 region files (extraction.get\_mask). The method takes in 2D image data to obtain the dimensions for the output mask. A user can also define a custom 1D (i.e. spectral), 2D (spatial), or 3D mask, and use the CWITools masking method to apply it to any compatible data. The mask and data can both be 3D, both 2D, 3D data with a 2D mask (in which case the mask is treated as a spatial mask and applied at all wavelength layers), or 3D data with a 1D mask (in which case the mask is treated as a spectral mask). Masked data is saved by default with a ".M.fits" file extension.

# Smoothing

While there are many existing smoothing and filtering methods available from libraries such as NumPy and SciPy, three custom methods are included in CWITools to allow for error propagation (by squaring the smoothing kernel). The main smoothing method is a generic one which applies smoothing along a requested subset of axes for 1D, 2D, or 3D input data. For 2D or 3D data, smoothing can be applied along any subset of the axes, and is implemented as a series of convolutions with 1D smoothing kernels. This is permitted by the fact that the two available kernels are Gaussian and Box kernels, both of which are separable into 1D components. Smoothing along the z-axis with a kernel of size  $N_z$  is described mathematically as:

$$F'(x, y, z) = \frac{\sum_{i=-N_z/2}^{i=+N_z/2} F(x, y, z+i) K_{1D}(i)}{\sum_i K_{1D}(i)}.$$
(3.10)

The naive error propagation (i.e. ignoring the effect of covariance) follows:

$$V'(x, y, z) = \frac{\sum_{i=-N_z/2}^{i=+N_z/2} V(x, y, z+i) K_{1D}^2(i)}{(\sum_i K_{1D}(i))^2}.$$
(3.11)

As mentioned before, O'Sullivan et al. (2020) contains a full discussion of bootstrapping variance estimates to include covariance in PCWI data, following the lead of earlier work by the Calar Alto Legacy Integral Field Area survey (Husemann et al., 2013) and the SDSS-IV MaNGA IFU Galaxy Survey (Law et al., 2016). CWITools does not yet have a built-in implementation of this method, but it is planned for a future release.

Two other methods are included as faster, stripped-down alternatives for the most common smoothing applications: spatial smoothing and wavelength smoothing. These make use of a faster convolution method ('SciPy.ndimage.convolve') and have simpler usage in that they always take 3D data and do not require the user to specify the axes.

## Segmentation

While there are any number of additional steps that may be required depending on the application, a common final step in the extraction process is segmentation: dividing the cube into contiguous 3D regions above a certain threshold. The CWITools implementation of this is built using an existing segmentation routine from the popular package 'Scikit Image' (specifically skimage.measure.label). The function allows the user to set the segmentation threshold either in absolute physical units or in terms of SNR. It also provides the ability to limit the segmentation to specific

wavelength ranges of interest and set a lower bound on the number of voxels required for a region to be included. Wavelength ranges containing common nebular emission lines can be included by providing the redshift of suspected emission and a velocity window to explore around each line. Similarly, known bright sky lines in the Palomar/Keck sky spectra can be excluded automatically. Custom ranges to include or exclude can also be provided by the user (e.g. to rule out bad wavelength regions or extend the velocity range of an extra broad line). The output of this method - a cube containing labelled regions which I call the 'object' cube as a shorthand - can be used in synthesizing results and making measurements for specific objects later on.

#### **3.6 Module: Synthesis**

The synthesis module takes 3D data products as input and returns vectors or arrays containing scientifically relevant results such as radial profiles or velocity maps. 'Object' cubes produced by segmentation can be used for many of these functions to generate such results from specific regions only.

## White-Light Images

Generation of white-light (WL) images is a straight forward process: a 2D image is formed by summing over a broad wavelength range. Two sets of wavelength regions should generally be excluded in doing so: bright sky lines and any nebular emission lines present in the data. The user can specify wavelength regions to mask, and also select an option to automatically mask some known sky lines which are built into the package (stored in the 'data/sky/' subdirectory in the installation directory). The wavelength region used should also be restricted to the 'WAVGOOD' region indicated by the header information, but usually the data will have been cropped to this range before generating WL images. The input is generally assumed to be in the standard KCWI/PCWI units of 'FLAM16' - meaning  $10^{16} \times \text{erg s}^{-1}\text{cm}^{-2}\text{Å}^{-1}$ . If this is the input unit (i.e. if the keyword 'BUNIT' is set to 'FLAM16'), the WL image is converted to surface-brightness units ( $10^{16} \times \text{erg s}^{-1}\text{cm}^{-2}\text{arcsec}^{-2}$ ) following

$$SB_{16}(x, y) = \frac{\Delta\lambda}{(\Delta\theta)^2} \sum_{z} F_{\lambda, 16}(x, y, z), \qquad (3.12)$$

where  $(\delta \theta)^2$  is the size of the input spaxels in units of square arcseconds, and  $\Delta \lambda$  is the size of each wavelength layer in units of Angstrom. Error is propagated

by summing the variance data over the same wavelength layers and squaring the multiplicative term.

$$\operatorname{Var}(\operatorname{SB}_{16})(x, y) = \left(\frac{\Delta\lambda}{(\Delta\theta)^2}\right)^2 \sum_{z} V_{\lambda, 16}(x, y, z).$$
(3.13)

#### **Pseudo Narrowband Images**

A pseudo-Narrowband (pNB) image is a Narrowband image formed by summing wavelength layers of a datacube. The synthesis module has a method for generating pNB images with or without the subtraction of a local white-light image. The user provides the central wavelength and bandwidth of the desired pNB image, and is returned the image, an estimate of variance on the image, a local white-light image, and the variance on the white-light image. The variance estimates are derived from 3D variance cubes, if provided, or by taking the variance along the z-axis in the input data, if not. Optionally, the user can request white-light subtraction, in which case the white-light image is scaled and subtracted from the pNB image. The scaling factor is calculated using a user-provided location, typically the location of a bright, central continuum source to be subtracted.

This tool is useful as an initial exploration of a data cube. By generating a series of these images at different wavelengths, one can form a channel map. Channel maps are useful tools in exploring IFU data, especially when the nature (i.e. center and width) of any emission is not already well known, and the signal may be too faint to detect easily on a voxel-by-voxel basis (see O'Sullivan et al., 2020). In the case of non-detections, this tool provides a useful way to obtain upper limits on the surface brightness of suspected emission.

A commonly studied property of both galaxies and extended emission is the circularly average radial surface brightness profile, usually as a function of some physical distance (i.e. comoving or proper kiloparsecs). CWITools provides a convenient function for the calculation of a radial surface brightness profile, where the radius can be returned in units of pixels, arcseconds, comoving kiloparsecs, or proper kiloparsecs. The radial bins are defined by four parameters: a minimum radius, a maximum radius, the number of (equally spaced) bins, and a scale setting which determines whether to make the bins of equal size in linear space or log space. The user provides a surface brightness map and central location as input. The surface brightness profile can be easily obtained using either the synthesis.pseudo\_nb method or synthesis.obj\_sb method (below) which measures the surface brightness of a defined 3D region.

# **Object Surface Brightness, Spectra and Moments**

In many IFU studies, though especially those regarding extended nebular emission, 3D object masks (generated by segmentation - see extraction.segment) are a central feature. These masks contain integer-labelled, contiguous 3D regions which we refer to here as 'objects.' Once one has a 3D region delineating the object of their study, it becomes trivial to generate useful products. CWITools has a number of functions which accept 3D object masks, along with specified object IDs (i.e. the number of the region to study). Specifically, the synthesis module contains a method for generating: (i) a surface brightness map of an object (synthesis.obj\_sb), (ii) an integrated 1D spectrum of an object (synthesis.obj\_spec), and (iii) maps of the first two z-moments (i.e. velocity and dispersion) for an object (synthesis.obj\_moments).

Object surface brightness maps are obtained by setting all non-object voxels to zero, summing the cube along the z-axis, and applying a conversion from units of  $F_{\lambda}$  (i.e. erg/s/cm2/Å) to units of surface brightness (erg/s/cm2/arcsec2). Onedimensional spectra are similarly obtained by setting all non-object voxels to zero and summing along the two spatial axes. In this case, no unit conversion is required. The user can decide whether to apply the 3D mask in full or to sum the full spectra of all spaxels within the object boundary. The latter method can be useful for the purpose of presentation in that it shows the background noise throughout the rest of the spectrum, and in recovering the full shape of emission lines, as the thresholding step used to obtain 3D masks necessarily cuts out the faint wings of a profile once they fall below the noise level. Finally, 2D maps of the first and second moments in wavelength are obtained through a straight-forward moments calculation (see measurement.first\_moment and measurement.second\_moment for details and uncertainty propagation). In the case where moments are being calculated using a 3D mask, it is assumed that all object voxels contain positive flux, so no further threshold is applied and the 'closing window' method is not used.

# 3.7 Module: Modeling

The modeling module provides wrapper functions for some common models, model fitting methods, model comparison, and some other miscellaneous useful tools.

While there is a wealth of existing modeling functionalities available from Astropy and other packages, the CWITools modeling module contains wrappers for models commonly used in PCWI/KCWI data analysis with a self-consistent syntax for use in model fitting and evaluation. At the moment, all of the models and modeling functions are one dimensional, as the main applications considered are the fitting of emission lines and surface brightness profiles. Future updates, beyond the initial release, may include additional models, such as 2D kinematic or surface brightness models.

### Wrappers for Models and Fitting

Model fitting within CWITools is done by minimizing a residual sum of squares (RSS) function using SciPy's implementation of differential evolution as the default optimizer. Differential evolution is a stochastic method of optimization, which is less susceptible to local minima than standard gradient descent methods. This can be of importance in fitting models in the presence of significant noise, where the cost function is not smooth.

Differential Evolution (in scipy.optimize), like many other available optimization methods, finds the minimum of an objective function of the form f(p, [args]), where p is the vector of free parameters to be optimized and [args] is a sequence of any additional, fixed parameters required for the function. Since CWITools uses a least-squares approach, the objective function is one which computes the residual sum-of-squares (RSS) for any given model, set of model parameters, and input data. This flexible RSS method (modeling.rss\_func) has the form rss\_func(p, f, x, y), where p is the vector of model parameters, f is the model function, and x and y are the input data. The model function itself must be of the form f(p, x). For convenience, CWITools has a number of common functions written in this form, with more to be added later. The current library of functions includes a 1D Voigt, Gaussian, and Moffat profiles for line-profile or PSF fitting. For the fitting of radial profiles, 1D Sersic, Exponential, and Power-law profiles are included.

#### **Model Comparison**

Model comparison can be performed using one of two information criteria: the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Both the AIC and BIC indicate the *relative* likelihood that a given model is the best representation of the observed data out of all models considered, taking into account
the degrees of freedom of the models. The AIC is calculated as

$$AIC = 2k + n\ln(RSS), \qquad (3.14)$$

where k is the number of parameters, n is the number of data points used in fitting, and RSS is the residual sum of squares. A lower score is better for both the AIC and BIC. The AIC also has a correction for small sample sizes, denoted as AICc:

AICc = AIC + 
$$\frac{2k^2 + 2k}{n - k - 1}$$
. (3.15)

The correction term vanishes as *n* approaches infinity. This term is always included in the CWITools implementation of the AIC. The BIC is calculated as:

BIC = 
$$k \ln(n) + n \ln\left(\frac{\text{RSS}}{n}\right)$$
. (3.16)

A set of AIC or BIC values can be converted into a set of weights indicating the relative likelihood of each model. Following Wagenmakers and Farrell, 2004, these weights can be calculated as:

$$w_i = \frac{\exp\left(-\frac{1}{2}\Delta_i(\text{BIC})\right)}{\Sigma_j \exp\left(-\frac{1}{2}\Delta_j(\text{BIC})\right)}.$$
(3.17)

It is important to note that these weights indicate *relative* likelihoods, with respect only to the other models considered. The scientific significance of such relative likelihoods therefore depends strongly on the total set of models considered. As a random example, it would be misleading to claim an absorption line is Gaussian in shape if the only models considered were a Gaussian model and a flat continuum model, but if other common line shapes (e.g. Lorentzian, Voigt) were considered and a Gaussian still had the lowest AIC/BIC, then it may be a reasonable claim.

## 3.8 Module: Measurement

# **First and Second Moments**

In context of IFU data analysis, moments are typically calculated along the z-axis to derive velocity and dispersion. For that reason, we focus here on moments calculated along the z axis in wavelength units. The measurement module contains implementations of these moments calculations that accept a 1D wavelength axis and spectrum as input. The synthesis module contains a method which builds on this to create two-dimensional moment maps for 3D objects. The first moment of a 1D spectrum with wavelength  $\lambda$  and flux *F* is calculated as follows:

$$\mu_1 = \frac{\sum_i \lambda_k F_k}{\sum_k F_k}.$$
(3.18)



Figure 3.11: z-Moment maps, shown in units of km  $^{-1}$ , calculated using CWITools. From left to right: first moment, propagated error on the first moment, second moment, and propagated error on the second moment.

As a short-hand, let us refer to the numerator here as  $N_1$  and the denominator as D. The error on this calculation can then be shown to be

$$\sigma(\mu_1) = \sqrt{\sum_{i} \frac{(\lambda_i D - N_1)^2}{D^4} V_i},$$
(3.19)

where  $V_i = \sigma^2(F_i)$  is the variance on the flux. The second moment is calculated as follows:

$$\mu_2 = \sqrt{\frac{\sum_i (\lambda_k - \mu_1)^2 F_k}{\sum_k F_k}}.$$
(3.20)

Let us refer to the numerator this time as  $N_2$ . The denominator is the same as in Equation 3.8. The error on this calculation can be shown to be:

$$\sigma(\mu_2) = \frac{1}{2\mu_2} \sqrt{\sum_i \frac{((\lambda_i - \mu_1)^2 D - N_2)^2}{D^4} V_i}.$$
(3.21)

It should be noted that the numerators here will have slightly higher variance due to the covariance between adjacent pixels. This can be taken into account following Section 3.4.

### **Integrated Luminosity**

The integrated luminosity can be calculated for 1D, 2D, or 3D data and an optional object mask of the same dimensions. If no object mask is provided, all of the input data is summed. If the input is two-dimensional, the input units are assumed to be those of surface brightness and the total flux,  $F_{tot}$  is calculated following:

$$F_{\text{tot}} = (\Delta \theta)^2 \Sigma_{\text{x}} \Sigma_{\text{y}} \text{SB}(\text{x}, \text{y}) M(\text{x}, \text{y})$$
(3.22)



Figure 3.12: Left: an illustration of the three characteristic radii provided by the measurements module. From largest to smallest, they are: (i) the maximum radius ( $R_{max}$  - dotted circle), (ii) the effective radius ( $R_{eff}$  - dashed circle), and (iii) the flux-weighted RMS radius ( $R_{rms}$  - solid circle). Right: a radial profile generated by the synthesis module for the same nebula. The same three radii are shown on this axis.

where  $\Delta \theta$  is the angular area in units of  $\operatorname{arcsec}^2$  and M(x, y) is the 2D binary mask. If the input is 3D, the input unit is assumed to be units of  $F_{\lambda}$ , and the luminosity is calculated as:

$$F_{\text{tot}} = (\Delta \lambda) \sum_{x} \sum_{y} \sum_{z} F_{\lambda}(x, y, z) M(x, y, z), \qquad (3.23)$$

where  $\Delta \lambda$  is the size of the wavelength layers, usually in Angstrom. Finally, if the input is one dimensional, the luminosity is calculated as:

$$F_{\text{tot}} = (\Delta \lambda) \sum_{z} F_{\lambda}(z) M(z). \qquad (3.24)$$

The total flux is then converted to luminosity using the luminosity distance,  $D_L(z)$ :

$$L_{\text{tot}} = 4\pi D_{\text{L}}^2(z) F_{\text{tot}}$$
(3.25)

The error on the luminosity is obtained by summing the variance and squaring the multiplicative terms. The usual caveat regarding covariance (Section 3.4) applies here, and should be considered before using the estimated error.

# **Characteristic Sizes**

CWITools provides a number of ways to measure the size of a 2D or 3D object. First and foremost among them is area, either units of square pixels or square arcseconds. Several characteristic radii are also defined, as in O'Sullivan et al., 2020; effective radius (Reff), maximum radius (Rmax), and RMS radius (Rrms), defined as follows:

$$R_{\rm eff} = \sqrt{A/\pi} \tag{3.26}$$

$$\mathbf{R}_{\max} = \max[|r_{ij} - r_c| \times M_{ij}] \tag{3.27}$$

$$\mathbf{R}_{\rm rms} = \sqrt{\frac{\sum_i \sum_j |r_{ij} - r_c|^2 F_{ij} M_{ij}}{\sum_i \sum_j F_{ij} M_{ij}}}$$
(3.28)

where the indices *i* and *j* iterate over the spatial axes,  $r_{ij}$  is the 2D vector from the image origin to the point  $(x_i, y_j)$ ,  $r_c$  is the vector from the origin to the flux-weighted centroid of the nebula,  $F_{ij}$  is the total flux or surface-brightness at position  $(x_i, y_j)$ , and  $M_{ij}$  is the 2D binary object mask, such that  $M_{ij} = 1$  for object spaxels. If 3D data and object masks are provided,  $F_{ij}$  and  $M_{ij}$  are formed by taking the sum and maximum value along the z-axis, respectively. Each measurement, taken alone, serves as a useful reference for a different aspect of the object size, but lacks any information about shape. R<sub>eff</sub> serves as a convenient proxy for total measured size, R<sub>max</sub> describes the maximum extent of the nebula in any direction from its flux-weighted centroid, and R<sub>rms</sub> provides a characteristic scale at which the emission is concentrated.

# **Asymmetry and Eccentricity**

Another commonly used measurement in describing the morphology of an object is its asymmetry or eccentricity. The asymmetry parameter,  $\alpha$  reflects the minor axis to major axis ratio of the emission. The calculation of this parameter, used in O'Sullivan et al., 2020, is taken from Arrigoni Battaia et al. (2016), who in turn based it on work by Stoughton et al., 2002. It derives the parameter  $\alpha$  from second order spatial moments defined as:

$$M_{xx} = \left\langle \frac{(x - x_c)^2}{r^2} \right\rangle_f,\tag{3.29}$$

$$M_{yy} = \left\langle \frac{(y - y_c)^2}{r^2} \right\rangle_f,$$
 (3.30)

and

$$M_{xy} = \left\langle \frac{(y - y_{\rm c})(x - x_{\rm c})}{r^2} \right\rangle_f,$$
 (3.31)

where  $(x_c, y_c)$  is the flux-weighted centroid of the emission. As an aside, CWITools' measurement module contains a wrapper to enable arbitrary moment calculations of the form  $M(xx, yy, p, q, f) = \langle xx^p yy^q \rangle_f$ , where xx and yy are 2D mesh-grids

of x-position and y-position, and f is a 2D grid of flux-like weights. These second order moments are used to define the terms

$$Q \equiv M_{xx} - M_{yy} \tag{3.32}$$

$$U \equiv 2M_{xy} \tag{3.33}$$

which are then used to derive the asymmetry

$$\alpha = b/a = \frac{1 - \sqrt{Q^2 + U^2}}{1 + \sqrt{Q^2 + U^2}}.$$
(3.34)

The 'elliptical' eccentricity, which is another representation of the same thing, is defined as

$$e \equiv \sqrt{1 - \alpha^2}.\tag{3.35}$$

Both of the above functions accept 2D or 3D data as input. If an object mask is not provided, all input data is used in the above calculation. This is generally only recommended if the input data already contains only the isolated 2D or 3D signal (i.e. an object mask has been applied by some means).

### **Specific Angular Momentum**

One commonly studied phenomenon, in terms of kinematics, is the presence of structured kinematic shear. Shear-like features can arise from a number of different physical phenomena including inflows, outflows, and galactic rotation. Specific angular momentum is typically defined as angular momentum per unit mass, and has a dimensionality of square distance per unit time (e.g.  $m^2s^{-1}$ ). Here, we use the flux of an object as a proxy for mass, and calculate the weighted average over the entire object. We also bear in mind that, in IFU data, we do not have full 3D vectors for the distance; rather, we have only projected radius and line-of-sight velocity. Thus to be more precise in our wording, the measurement is of the flux-weighted average *projected specific* angular momentum of an object. This is defined as:

$$\langle \vec{j} \rangle_f = \frac{\sum_x \sum_y F(x, y) \vec{R}_{\perp}(x, y) \times \vec{v}_z(x, y)}{\sum_x \sum_y F(x, y)}$$
(3.36)

where  $\vec{R}_{\perp}(x, y)$  is the projected radius, in pkpc, from the flux-weighted centroid of the nebula to the point (x, y), F(x, y) is the flux at that point, and  $\vec{v}_z(x, y)$  is the line-of-sight velocity, in km s<sup>-1</sup>. The units of the measurement are thus pkpckms<sup>-1</sup>. This measurement provides a useful, quantitative insight into whether the kinematics of an object appear to be dominated by a structured velocity gradient (which aligns with the distribution of flux to produce a large value) or whether the kinematics appear to be primarily noisy or chaotic.

#### **3.9** Summary and Discussion

We have presented here a configurable pipeline for the extraction, modeling and measurement of signals in three-dimensional integral field spectroscopy data, and shown its application on extracting extended Ly $\alpha$  emission from a high-redshift QSO. CWITools provides a comprehensive and flexible suite of tools for correcting, coadding, and analyzing KCWI and PCWI data cubes. Here, we discuss briefly the extension of CWITools to other instruments and general future work.

#### **Extension to Other IFUs**

CWITools, as discussed, is built specifically with the Cosmic Web Imager instruments at Keck and Palomar observatory in mind. That being said, there is much that IFS data has in common, regardless of which instrument produces it. Certain header keywords are FITS standards, and others are not. Certain methods (e.g. polynomial subtraction) may apply universally, while others (e.g. analytical PSF modeling) may depend on instrument specifics. It is thus worth a quick look at the areas in which CWITools becomes instrument-specific, and how it could be adapted to other instruments.

The majority of methods in CWITools can be applied to any input containing the same data structure (a wavelength axis and two spatial axes). However, there are a few areas in which the specifics of the CWI format feature strongly. First and foremost among these is the nature of the file-types saved by the standard DRP. In KDRP and PDRP, for a single exposure, a separate data cube is saved for intensity, background, variance, and masks. As a counter example, the data cubes produced by MUSE store the variance associated with an intensity cube as a second HDU within the same file. This represents the biggest challenge in using CWITools for other instruments. However, in lieu of a package update with flexible input in every method, it can be patched by writing a tool to convert data from one instrument to the format of another. For example, a MUSE data cube could be loaded, written to separate intensity and variance files, and then used as input to CWITools. As such, this problem is easy to solve with a single function that converts data formats. This method can also be applied to data structures and header keywords; a specific function can be written to re-order axes and re-name keywords. Once this function is written, it is trivial to apply it to data and then continue using CWITools as on CWI data.

The next category of incompatibilities arises from package data that is specific to

one instrument. For example, CWITools stores a list of known bright sky-lines at Palomar and Keck observatories. If certain tools were to be applied to MUSE data, a list of bright sky lines at European Southern Observatory would need to be added. CWITools also contains information about the different gratings and slicers available to PCWI and KCWI. These are used to estimate the spatial and spectral resolution. Again, this information would need to be added to make the same function available to another instrument. However, this amounts to simple data entry and is not a major obstacle.

On the surface, it seems almost trivial to address these concerns and adapt CWITools to other instruments and it may indeed be so. However, there is always the risk of so-called 'unknown unknowns.' Are there assumptions being made that may not apply to all instruments? For example, what if the spaxels in a certain IFS are not spatially adjacent and have gaps between them (say, due to a sparse lenslet array)? The assumption that spaxels are always spatially adjacent is present in the current implementations of PSF modeling and measuring the radial extent of objects. To adjust for this, all such methods would need to be updated to work in world coordinates rather than image coordinates, spatially. The possibility for such unanticipated, fundamental inconsistencies is a major reason for the focused development of CWITools. Extension to other instruments would likely have to be rolled out one at a time, with some period of testing and development required. Having said that, CWITools has intentionally been built in a modular way to allow for this possibility.

# **Future Development**

The development of this toolkit was motivated in large part by survey work. As discussed briefly in the introduction, the latest generation of IFS instruments on 5-10m class telescopes has enabled surveys of 'relatively large' samples of tens of targets. IFS data analysis can be extremely complex and time consuming. If such surveys are to grow in size and remain scientifically productive, data analysis pipelines will be needed. Much of CWITools was developed from scratch because there were no easily available tools to perform the desired tasks. However, it has been developed as an open source package and we extend an open invitation to contributors so that this particular wheel does not need to be reinvented. As such, we encourage anyone in the community who would like to see certain features to contact the authors and get involved.

## References

- Arrigoni Battaia, F. et al. (Sept. 2016). "The Stacked Ly $\alpha$  Emission Profile from the Circum-Galactic Medium of  $z \sim 2$  Quasars." In: ApJ 829, 3, p. 3. DOI: 10.3847/0004-637X/829/1/3. arXiv: 1604.02942.
- (Jan. 2019). "QSO MUSEUM I: a Sample of 61 Extended Ly α Emission Nebulae Surrounding z ~ 3 Quasars." In: MNRAS 482.3, pp. 3162–3205. DOI: 10.1093/ mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].
- Astropy Collaboration, A. M. Price-Whelan, et al. (Sept. 2018). "The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package." In: AJ 156.3, 123, p. 123. DOI: 10.3847/1538-3881/aabc4f. arXiv: 1801.02634 [astro-ph.IM].
- Astropy Collaboration, T. Robitaille, et al. (Oct. 2013). "Astropy: A Community Python Package for Astronomy." In: *A&A* 558, A33, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM].
- Avila, R. J. et al. (Sept. 2015). "DrizzlePac 2.0 Introducing New Features." In: Astronomical Data Analysis Software an Systems XXIV (ADASS XXIV). Ed. by A. R. Taylor and E. Rosolowsky. Vol. 495. Astronomical Society of the Pacific Conference Series, p. 281. arXiv: 1411.5605 [astro-ph.IM].
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Cai, Z. et al. (Sept. 2019). "Evolution of the Cool Gas in the Circumgalactic Medium (CGM) of Massive Halos A Keck Cosmic Web Imager (KCWI) Survey of Ly $\alpha$  Emission around QSOs at  $z \approx 2$ ." In: *arXiv e-prints*, arXiv:1909.11098, arXiv:1909.11098. [astro-ph.GA].
- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: Groundbased and Airborne Instrumentation for Astronomy V. Vol. 9147. Proc. SPIE, 91475K. doi: 10.1117/12.2057056.
- Husemann, B. et al. (Jan. 2013). "CALIFA, the Calar Alto Legacy Integral Field Area survey. II. First public data release." In: A&A 549, A87, A87. DOI: 10. 1051/0004-6361/201220582. arXiv: 1210.8150 [astro-ph.CO].
- Law, D. et al. (Oct. 2016). "The Data Reduction Pipeline for the SDSS-IV MaNGA IFU Galaxy Survey." In: *AJ* 152.4, 83, p. 83. DOI: 10.3847/0004-6256/152/4/83. arXiv: 1607.08619 [astro-ph.IM].
- Martin, D. C. et al. (May 2014a). "Intergalactic Medium Emission Observations with the Cosmic Web Imager. I. The Circum-QSO Medium of QSO 1549+19, and Evidence for a Filamentary Gas Inflow." In: *ApJ* 786, 106, p. 106. DOI: 10.1088/0004-637X/786/2/106. arXiv: 1402.4816.

- Martin, D. C. et al. (May 2014b). "Intergalactic Medium Emission Observations with the Cosmic Web Imager. II. Discovery of Extended, Kinematically Linked Emission around SSA22 Lyα Blob 2." In: *ApJ* 786, 107, p. 107. DOI: 10.1088/0004-637X/786/2/107. arXiv: 1402.4809.
- Matuszewski, M. et al. (July 2010). "The Cosmic Web Imager: An Integral Field Spectrograph For The Hale Telescope At Palomar Observatory: Instrument Design And First Results." In: *Ground-based and Airborne Instrumentation for Astronomy III.* Vol. 7735. Proc. SPIE, 77350P. DOI: 10.1117/12.856644.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 *z* ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.
- Stoughton, C. et al. (Jan. 2002). "Sloan Digital Sky Survey: Early Data Release." In: *AJ* 123.1, pp. 485–548. DOI: 10.1086/324741.
- Wagenmakers, E. and S. Farrell (2004). "AIC Model Selection Using Akaike Weights." In: *Psychon. Bull. Rev.* 11.1, pp. 192–196.

# Chapter 4

# THE FLASHES SURVEY I: PILOT STUDY

# O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 *z* ~ 2.3 − 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. doi: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.

#### 4.1 Overview

We present the pilot study of the Fluorescent Lyman-Alpha Structures in High-z Environments (FLASHES) Survey, the largest integral-field spectroscopy survey to date of the circumgalactic medium at z = 2.3 - 3.1. We observed 48 quasar fields with the Palomar Cosmic Web Imager (Matuszewski et al., 2010) to an average (2 $\sigma$ ) limiting surface brightness of  $6 \times 10^{-18}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup> (in a 1" aperture and ~ 20Å bandwidth). Extended HI Lyman- $\alpha$  emission is discovered around 37/48 of the observed quasars, ranging in projected radius from 14 - 55proper kiloparsecs (pkpc), with one nebula exceeding 100 pkpc in effective diameter. The dimming-adjusted circularly averaged surface brightness profile peaks at  $1 \times$  $10^{-15}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup> at R<sub>1</sub> ~ 20 pkpc and integrated luminosities range from  $0.4 - 9.4 \times 10^{43}$  erg s<sup>-1</sup>. The emission appears to have an eccentric morphology and an average covering factor of  $\sim 30 - 40\%$  at small radii. On average, the nebular spectra are redshifted with respect to both the systemic redshift and Ly $\alpha$  peak of the quasar spectrum. The integrated spectra of the nebulae mostly have single or double-peaked profiles with global dispersions ranging from 143 - 708 km s<sup>-1</sup>, though the individual Gaussian components of lines with complex shapes mostly have dispersions  $\leq 400$  km s<sup>-1</sup>, and the flux-weighted velocity centroids of the lines vary by thousands of km s<sup>-1</sup> with respect to the QSO redshifts. Finally, the rootmean-square velocities of the nebulae are found to be consistent with those expected from gravitational motions in dark matter halos of mass  $\text{Log}_{10}(M_h[M_{\odot}]) \simeq 12.2^{+0.7}_{-1.2}$ . We compare these results to existing surveys at higher and lower redshift.

# 4.2 Introduction

To understand the evolution of galaxies and their properties, it is critical to understand their environments. Our current picture of galaxy formation takes place in a universe dominated by cold dark matter (Blumenthal et al., 1984). In this picture, dark matter structures collapse in a hierarchical manner, dragging with them the baryonic material that eventually forms and fuels galaxies. A key element of this framework is the interplay between galaxies and their environments; galaxies form and evolve through a series of interactions with both the circumgalactic and intergalactic medium (CGM and IGM; e.g., Bond, Kofman, and Pogosyan 1996; Fukugita, Hogan, and Peebles 1998). A long history of accretion, outflows, and merger events underlies the properties of galaxies that we observe today (e.g, Kereš et al. 2005; Dekel et al. 2009; Fumagalli, O'Meara, and Prochaska 2011; Correa et al. 2015).

With the development of highly sensitive integral field spectrographs, there is now the opportunity to contribute substantial direct observational evidence to the discussion around high-redshift galaxy environments, which has so far taken place largely in the realms of theory and simulation. The sensitivity, spatial resolution, and spectral flexibility of these new instruments enable exploratory surveys which map the density, morphology, composition, and kinematics of the CGM. Several integral-field spectroscopy (IFS) studies focused on individual targets have produced remarkable insights into galactic environments at high redshift ( $z \ge 2$ ). Umehata et al. (2019) reported the discovery of giant Ly $\alpha$  filaments, spanning more than a megaparsec, embedded in a z = 3.1 protocluster. Several kinematic studies of extended nebulae around QSOs have revealed evidence for intergalactic gas spiraling into dark matter halos (Martin et al., 2019; Martin et al., 2016; Arrigoni Battaia et al., 2018). A number of studies over the past 5-6 years have revealed giant Ly $\alpha$ nebulae around individual high-redshift galaxies and QSOs (e.g. Cai et al., 2018; Martin et al., 2014a; Martin et al., 2014b) as well as connecting pairs of QSOs (Arrigoni Battaia et al., 2019a). Multi-phase observations of similar nebulae have also begun to emerge (Cantalupo et al., 2019; Marques-Chaves et al., 2019). However, to fully characterize the morphology, composition, and dynamics of the CGM, large samples with multi-wavelength observations are needed.

Christensen et al. (2006) and Herenz et al. (2015) provide some of the first examples of IFU surveys of high-redshift QSO environments. These studies focused on extended Ly $\alpha$  with sample sizes of seven and five, detecting extended emission in 4/7 and 1/5 targets, respectively. More recently, teams using the Multi-Unit Spectroscopic Explorer (MUSE) on the Very Large Telescope (Caillier et al., 2014)



Figure 4.1: IFS surveys of extended emission around high redshift galaxies. Surveys are shown as stacked histograms representing the number of targets in each.

have produced surveys of Ly $\alpha$  emission around quasars and galaxies at  $z \gtrsim 3$  with sample sizes on the order of tens of targets. Borisova et al. (2016) (hereafter B16) studied 17 bright radio-quiet quasars (and 2 radio-loud) at  $z \sim 3.5$ , finding ubiquitous "giant" Ly $\alpha$  nebulae on scales larger than 100 pkpc, with clear asymmetries and a circularly-averaged radial profile following power laws. Arrigoni Battaia et al. (2019b) (hereafter A19) studied 61 QSOs with a median redshift of 3.17, finding Ly $\alpha$  nebulae extending on the order of tens of kpc around their quasars.

The nebulae they discover have some spread in their degree of spatial symmetry, and they find that their radial profiles are best fit by an exponential profile with a scale length  $r_H \sim 15$  pkpc. They compare this to a narrow-band study at  $z \sim 2$ (Arrigoni Battaia et al., 2016), but with the actual centroid of Ly $\alpha$  emission varying by thousands of km s<sup>-1</sup> from systemic QSO redshifts, it is not clear how reliable narrow-band imaging is without prior knowledge of the emission wavelength. Wisotzki et al. (2016) performed an ultra-deep exposure of the *Hubble* Deep Field South with MUSE, reaching a (1 $\sigma$ ) limiting surface brightness of 1 × 10<sup>-19</sup> erg s<sup>-1</sup>cm<sup>-2</sup> arcsec<sup>-2</sup>. They report detections of extended Ly $\alpha$  halos around 21 of the 26 total z = 3 - 6 galaxies in their sample, on spatial scales of  $R_{\perp} \sim O(10$  pkpc). The remaining 5 non-detections represent the faintest galaxies in the sample, and thus are thought to be a matter of insufficient sensitivity, making the overall result



Figure 4.2: The FLASHES pilot sample in redshift (z) vs. absolute i-band magnitude ( $M_i$ ). Circles indicate targets for which the value of  $M_i$  is estimated from the given apparent magnitude, while diamonds indicate those for which a value of  $M_i$  was provided in the SDSS DR12Q. The colorbar indicates the WISE infrared W2-W3 (4.6 – 11.6µm) color. The blue histograms represent the distributions of z and  $M_i$  in the SDSS DR12Q in the same redshift range.

consistent with the ubiquitous  $Ly\alpha$  halos reported in B16. More recently, Cai et al. (2019) observed 16 QSOs with redshifts z = 2.1 - 2.3 using the Keck Cosmic Web Imager (KCWI) (Morrissey et al., 2018) and report extended emission around all of them, although 2/16 of the nebulae are reported to have projected sizes smaller than 50 pkpc. The authors find that the nebulae are more asymmetric and lower in surface brightness than the z > 3 MUSE studies.

We have utilized the Palomar Cosmic Web Imager (PCWI) (Matuszewski et al., 2010) to conduct a pilot study of the gaseous environments of quasars spanning a redshift range of z = 2.3 - 3.1, filling a gap in existing observations (see Figure 4.1). This survey, which we call FLASHES (Fluorescent Lyman-Alpha Structures

in High-z Environments), consists of a broad pilot survey component presented in this paper, and a follow-up deep survey to be presented in a future paper. The pilot survey aims to map  $Ly\alpha$  emission from the CGM around the full sample of 48 quasars at redshifts  $2.3 \le z \le 3.1$  to a  $2\sigma$  surface-brightness limit of ~  $5 \times 10^{-18}$  erg  $s^{-1}cm^{-2}arcsec^{-2}$  in a ~ 1 arcsec<sup>2</sup> aperture for a pseudo-Narrow-band (pNB) image with a typical bandwidth of 20Å (a limit of ~  $6 \times 10^{-18}$  erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup> was achieved in the final pNB images). Based on existing observational work (e.g., A19, B16), this is expected to be sufficient to map CGM Ly $\alpha$  emission within a 50 – 100 proper kiloparsecs (pkpc) of the quasars, and enable us to constrain the morphology and kinematics of the CGM in this redshift range. In addition, we search for the presence of the gaseous filaments that are theorized to feed gas from the cosmic web into dark matter halos. Recent observations have offered tantalizing direct evidence supporting cold-flow accretion from multiple filaments forming 'cold inflow disks' (Martin et al., 2015; Martin et al., 2016; Martin et al., 2019). A larger sample of observations will allow us to test the validity of such models and their utility in constraining the gas dynamics associated with cold-flow accretion. A subset of these targets will be followed up with deep KCWI exposures for the latter component of the survey, targeting  $Ly\alpha$  emission at surface brightness levels an order of magnitude fainter than in the pilot survey, as well as targeting emission from metals such as  $\lambda$ CIV1549 and  $\lambda$ HeII1640 which probe the multi-phase structure of the CGM.

In this paper, we focus exclusively on the FLASHES pilot survey. In Section 4.3, we describe the survey methodology, target selection, and choice of observables. In Section 4.4, we present a summary of the observations and data. In Section 4.5, we describe the data reduction with the standard PCWI pipeline and a newly developed Python3 package, 'CWITools,' and the data analysis required to extract and characterize the nebular emission. In Section 4.6, we present the core observational results: emission maps, kinematic maps, spectra, symmetries and radial profiles. Finally, in Section 4.7, we discuss the implication of our results, sensitivity limits, and comparisons to existing work, before summarizing our findings in Section 4.8. For calculations of the luminosity distance and physical plate scales (pkpc per pixel) throughout the paper, we use a  $\Lambda$ CDM cosmology with  $H_0 = 70$ kms<sup>-1</sup>Mpc<sup>-1</sup>,  $\Omega_b = 0.3$ , and  $\Omega_{\Lambda} = 0.7$ .

ID	Target Name	Farget NameCoordinatesZQSOM		M <sub>i</sub>	Seeing	Clouds	Exp	$SB_{2\sigma}{}^{g}$
				AB mag	arcsec		min	$10^{-18}$ cgs
1	HS1700+6416 <sup>a</sup>	17:01:01.00 +64:12:09.10	$2.7375 \pm 0.0010$	-31.01	1.3	CR	222	0.48
2	SDSS1112+1521	11:12:52.45 +15:21:23.50	$2.7898 \pm 0.0003$	-28.38	1.1	CR	60	1.07
3	SDSS0834+1238	08:34:08.63 +12:38:36.54	$2.7465 \pm 0.0013$	-28.29	1	CR-P	60	1.06
4	SDSS1011+2941 <sup>a</sup>	10:11:56.00 +29:41:42.00	2.6400	-30.11	1.3	CR	60	1.27
5	SDSS0735+3744	07:35:35.44 +37:44:50.42	$2.7514 \pm 0.0003$	-27.90	1.3	CR	60	0.94
6	SDSS0103+1316 <sup>a</sup>	01:03:11.27 +13:16:17.70	$2.6985 \pm 0.0010$	-29.93	1.9	CR	56	1.08
7	SDSS0958+4703	09:58:45.42 +47:03:24.43	$2.4907 \pm 0.0003$	-28.30	2.1	CR	60	1.61
8	SDSS0132+3326 <sup>d</sup>	01:32:44.60 +33:26:55.42	$2.4205 \pm 0.0003$	-26.84	1.4	CR-F	60	0.81
9	SDSS0837+1459	08:37:12.89 +14:59:17.38	$2.5100 \pm 0.0004$	-28.27	2	CR	48	1.36
10	SDSS2241+1225	22:41:45.11 +12:25:57.24	$2.6222 \pm 0.0006$	-28.28	1.7	CR	60	1.23
11	SDSS1626+4858 <sup>e</sup>	16:25:59.89 +48:58:17.49	$2.7347 \pm 0.0007$	-28.49	1.1	CR-P	54	1.03
12	SDSS2328+0443 <sup>c</sup>	23:28:28.48 +04:43:46.84	$2.5681 \pm 0.0001$	-24.55	1.7	CR-P	60	1.17
13	SDSS1002+2008	10:02:55.43 +20:08:02.56	$2.6555 \pm 0.0006$	-27.25	1.2	CR	60	1.62
14	SDSS1218+2414	12:18:10.98 +24:14:10.90	$2.3752 \pm 0.0008$	-28.93	1.5	CR-P	40	1.84
15	SDSS0108+1635 <sup>a</sup>	01:08:06.40 +16:35:50.00	$2.6399 \pm 0.0003$	-29.32	1.6	Р	70	1.29
16	SDSS0753+4030	07:53:26.11 +30:40:38.63	$2.9304 \pm 0.0004$	-28.72	1.5	TC	60	1.06
17	SDSS0057+0346	00:57:37.78 +03:46:45.03	$2.4365 \pm 0.0005$	-27.83	2.1	CR	60	1.43
18	SDSS0012+3344	00:12:15.26 +33:44:00.33	$2.4502 \pm 0.0003$	-27.72	2	CR-P	60	1.96
19	SDSS0013+1630 <sup>d</sup>	00:13:55.86 +16:30:51.78	$2.5907 \pm 0.0002$	-27.93	1.6	CR	60	0.89
20	SDSS0006+1614	00:06:39.47 +16:14:59.30	$2.4216 \pm 0.0005$	-27.85	1.8	TC	60	1.29
21	SDSS0730+4340	07:30:02.80 +43:40:03.04	$2.9367 \pm 0.0005$	-27.39	1.8	CR	60	0.64
22	SDSS0822+1626	08:22:00.22 +16:26:52.87	$2.4541 \pm 0.0005$	-28.06	1.4	CR-P	60	1.39
23	SDSS1428+2336	14:28:10.96 +23:36:40.21	$2.7792 \pm 0.0004$	-27.83	1.2	CR-F	60	0.98
24	SDSS0851+3148 <sup>c</sup>	08:51:24.79 +31:48:55.72	$2.6384 \pm 0.0005$	-24.50	1.3	CR-P	60	1.36
25	SDSS0214+1912 <sup><i>a</i></sup>	02:14:29.71 +19:12:37.40	2.4710	-28.83	1.8	CR-P	70	1.19
26	SDSS0015+2927	00:15:53.14 +29:27:21.45	$3.0755 \pm 0.0003$	-28.46	1.3	CR	60	0.72
		Continued on next page		1	1	1		·

 Table 4.1: Summary of FLASHES Pilot Observations

ID	Target Name	Coordinates z <sub>QSO</sub>		M <sub>i</sub>	Seeing	Clouds	Exp	$\mathrm{SB}_{2\sigma}{}^{g}$
				AB mag	arcsec		min	$10^{-18}$ cgs
27	SDSS2339+1901 <sup>b</sup>	23:39:44.60 +19:01:52.00	2.6200	-29.72	1.8	CR-P	60	1.54
28	SDSS0300+0222 <sup>b</sup>	03:00:46.02 +02:22:45.24	2.5240	-28.73	2	CR	68	1.26
29	SDSS0639+3819	06:39:01.60 +38:19:15.24	$2.5393 \pm 0.0003$	-26.55	1.3	CR	60	1.87
30	SDSS2338+1504 <sup>e</sup>	23:38:23.16 +15:04:45.22	$2.4121 \pm 0.0009$	-28.45	1.8	CR-P	56	1.26
31	SDSS0321+4132	03:21:08.45 +41:32:20.86	$2.4457 \pm 0.0007$	-29.97	1.8	TC	70	1.10
32	SDSS0211+3117	02:11:39.25 +31:17:24.67	$2.7854 \pm 0.0005$	-27.45	1.7	CR-P	60	1.12
33	SDSS0118+1950	01:18:39.93 +19:50:27.86	$2.7780 \pm 0.0002$	-28.24	1.1	CR	60	1.10
34	SDSS0144+0838	01:44:14.08 +08:38:20.40	$2.4307 \pm 0.0008$	-27.69	1.8	TC	60	1.10
35	SDSS1532+3059	15:32:58.24 +30:59:06.59	$2.5492 \pm 0.0004$	-28.92	1.3	CR-P	40	1.43
36	SDSS2151+0921	21:51:55.30 +09:21:14.07	$2.4493 \pm 0.0005$	-27.63	1.3	CR-P	56	1.22
37	SDSS0303+3838	03:03:09.16 +38:38:57.20	$2.7989 \pm 0.0004$	-28.03	1.1	CR	60	1.52
38	SDSS0126+1559	01:26:36.12 +15:59:29.94	$2.6969 \pm 0.0004$	-27.39	1.5	Р	64	1.45
39	SDSS2234+2637 <sup>c</sup>	22:34:53.07 +26:37:25.00	$2.7774 \pm 0.0009$	-25.15	1.5	CR-P	60	1.09
40	SDSS2259+2326	22:59:04.02 +23:26:43.91	$2.4622 \pm 0.0012$	-28.05	1.3	CR	60	1.41
41	SDSS0041+1925	00:41:09.83 +19:25:19.85	$2.7096 \pm 0.0007$	-26.66	1.5	TC	60	1.10
42	SDSS1552+1757	15:52:00.50 +17:57:22.70	$2.7034 \pm 0.0003$	-24.82	1.2	CR	80	1.22
43	SDSS0205+1902 <sup>b</sup>	02:05:27.51 +19:02:29.10	2.7030	-29.73	1.6	TC	70	1.25
44	SDSS1258+2123 <sup>c</sup>	12:58:11.25 +21:23:59.70	$2.6245 \pm 0.0003$	-24.62	1.3	CR	70	1.13
45	SDSS2340+2418 <sup><i>c</i>,<i>d</i></sup>	23:40:39.74 +24:18:59.15	$2.3513 \pm 0.0007$	-25.37	1.4	CR-F	60	0.50
46	SDSS0107+1104 <sup><i>c</i>,<i>d</i></sup>	01:07:14.66 +11:04:46.10	$2.5369 \pm 0.0010$	-25.31	1.4	CR	60	1.00
47	SDSS2350+3135 <sup>c</sup>	23:50:36.46 +31:35:05.02	$2.8285 \pm 0.0020$	-25.52	1.5	CR	60	1.43
48	SDSS0137+2405 <sup>c</sup>	01:37:58.65 +24:05:41.01	$2.4398 \pm 0.0012$	-24.42	1.8	CR-P	60	1.97

(a) Literature target; (b) Target selected from SIMBAD to fill observing schedule; (c) Dust-obscured targets

(d) Observed without Nod-and-Shuffle technique; (e) Radio-loud QSO

(g) Limiting surface brightness in a 1  $\operatorname{arcsec}^2$  aperture in a single cube layer.

#### 4.3 Survey Methodology

# **Choice of Observables**

At a redshift of z = 2, Bertone, Aguirre, and Schaye (2013) estimate that 80% of the energy emitted by the diffuse material of the CGM/IGM is carried by emission lines, with the remaining 20% in continuum emission. The Hydrogen Lyman series - and primarily Ly $\alpha$  - is the main contributor to this, carrying 20% of the line emission energy budget. Metal lines serve as better tracers for a wider range of over-densities or temperatures. They are typically an order of magnitude fainter than Ly $\alpha$  and depend strongly on gas metallicity and phase (Bertone and Schaye, 2012). The ubiquity and brightness of Ly $\alpha$  make it a clear choice for the pilot survey's goal of detecting and mapping the cool-warm phase of the CGM. With the Ly $\alpha$  line, we can constrain the morphology, density and baryonic mass of detected nebulae. Targets of interest can then be followed up in the deep study component of the survey, targeting metal lines such as HeII and CIV, in order to get a more complete picture of the multi-phase CGM.

# **Target Selection**

The FLASHES sample is primarily selected from SDSS DR12Q - the QSO Catalog from the 12th Data Release of the Sloan Digital Sky Survey (Alam et al., 2015). Targets were chosen within the redshift range of  $z \simeq 2.3$  and  $z \simeq 3.1$  based on the observability of  $Ly\alpha$  given the wavelength range accessible to the medium resolution grating of PCWI. An effort was made to select targets evenly across this redshift range though operational constraints such as the number of required instrument settings on a single night or the times at which various targets were observable from Palomar at low airmass, limited this effort. An effort was also made to select a range of absolute i-Band (rest-frame optical) magnitudes, as opposed to focusing on the brightest quasars, in order to explore any dependence of the nebular emission on QSO brightness. However, the FLASHES sample is still somewhat biased towards brighter QSOs when compared to the full distribution in the SDSS DR12Q. The distributions of the pilot sample in redshift and absolute i-band magnitude, alongside the distribution of these values in the SDSS DR12Q, is shown in Figure 4.2. A WISE color cut of W2[4.8 $\mu$ m] – W3[11.6 $\mu$ m] > 4.8 was used to identify heavily dust-obscured targets within the SDSS DR12Q which were expected to exhibit extended Ly $\alpha$  emission, as discussed in Bridge et al. (2013). The FLASHES pilot sample includes 9 of these dust-obscured targets, indicated in Figure 4.2 by the colorbar. We note that these 9 targets are not classical 'Type II' QSOs, as their spectra do contain broad line emission despite exhibiting heavily suppressed continuum emission. Over the course of the pilot survey, we included four additional targets from Trainor and Steidel (2012) which were known to exhibit extended emission, but lacked any IFS observations. Table 4.2 shows the names of these targets in both papers, for reference.

ID	FLASHES Name	Name in Source
1	HS1700+6416	HS1700+6416
4	SDSS1011+2941	Q1009+29 (CSO 38)
6	SDSS0103+1316	Q0100+13 (PHL 957)
15	SDSS0108+1635	HS0105+1619

Table 4.2: Targets included from Trainor and Steidel (2012)

In this work, a distinction is made between the total detection rate and the 'blind' detection rate, which excludes these five targets. Three targets were selected by searching the SIMBAD Astronomical Database based on coordinates and redshift to fill gaps in our observing schedule where no suitable targets were available from the SDSS DR12Q. Finally, two soft constraints were applied in our selection. First, targets with few obscuring foreground stars and galaxies were preferred, as blended and nearby sources can make the data analysis step of isolating the nebular emission prohibitively complicated. Second, where radio data was available, radio-quiet sources were preferred. One of the goals of FLASHES is to study gas dynamics and cold inflows from the cosmic web, and the presence of jets associated with radio-loud quasars would complicate this analysis. Of the 48 pilot targets, only two are detected in radio and classify as radio-loud using the criterion  $R = f_v^{1.4\text{GHz}}/f_v^{4400\text{\AA}} \gtrsim 10$  (Kellermann et al., 1989). Table 4.1 provides a breakdown of all of the pilot survey targets, coordinates, and sources.

The FLASHES target selection is multi-pronged, and there are biases in the methodology towards radio-quiet quasars with fields relatively clear of nearby/foreground sources. Any biases in the SDSS DR12Q will also be inherited. As such, the authors caution that while this is the first large sample of its kind in this redshift range, the results of this work should not be quickly or trivially extrapolated to the wider galaxy population.

#### 4.4 Observations and Ancillary Data

We observed 48 QSO fields between 2015 and 2018 on the 5-meter Hale telescope at Palomar using PCWI. PCWI is an image-slicer IFS mounted at the Cassegrain focus of the 5-meter Hale telescope at Palomar Observatory. The instrument field of view is  $60'' \times 40''$  (approximately  $480 \times 320$  pkpc<sup>2</sup> at  $z \sim 2 - 3$ ). The longer axis is composed of 24 slices with a width of  $\sim 2.5''$  and an in-slice pixel size of ~ 0.55". Typical seeing at Palomar is ~ 1.5" full-width-at-half-maximum, so individual exposures are slit-limited along the y-axis and seeing-limited along the x-axis. Exposures can be dithered to increase the sampling rate along the y-axis. Gratings and filters are interchangeable on PCWI. Our pilot observations used the medium resolution Richardson grating, which has a resolution of  $R \simeq 2500$  and operates over a bandpass of 400 - 600 nm. With a spectral plate scale of 0.55 Å/px, the minimum resolution element  $\Delta \lambda \sim 2 \text{\AA}$  is sampled above the Nyquist rate. For all observations, we use a filter with a bandpass of 350 - 580 nm. For 44/48 of the targets, Nod-and-Shuffle (N&S) mode of PCWI was used (see Matuszewski et al., 2010 for details). In short, N&S allows for highly accurate sky subtraction, almost entirely free of systematic residuals, at the cost of bandwidth and statistical noise. The standard pilot observation consists of three 40 minute N&S observations (20 minutes on sky, 20 minutes on source), stacked for a total of 1 hour on source and 1 hour on sky. Seeing conditions at Palomar are generally such that the full-width at half-maximum (FWHM) of a point source is 1 - 2''. To increase the spatial sampling, the second and third N&S observations are dithered by  $\pm 1''$ perpendicular to the direction of the slices. N&S mode was not used for four targets in this sample (see Table 4.1). This was done on one observing run in the interest of spending more telescope time on source rather than on sky, but the increase in systematic sky residuals was not deemed worth it for future observations. For these, an A-B-B-A pattern was used to alternate between 20-minute science frames and 10-minute sky frames. Lastly, one target (HS1700+6416) has a significantly longer total exposure time, as it was one of the earliest targets to be observed. However, as it still represents an initial exploration, it is included in the Pilot sample.

The goal for each target was 60 minutes on source and 60 minutes on sky. For four targets, we obtained 56, 54, 48, and 40 minutes in total due to time lost to poor weather. Some fields are not centered exactly QSO, due in part to guiding constraints (the guider and instrument fields of view have a fixed offset and orientation) and in part to position foreground sources in such a way that they could be

masked/subtracted. Because of this, a small number of the fields shown in Figure 4.6 have blank areas on one side. Multi-wavelength ancillary data were obtained for each target when available. Near- and far-UV data were obtained from GALEX (Bianchi et al., 2011). Photometric u, g, r, i and z-band magnitudes were obtained from the Sloan Digital Sky Survey's Photometric Catalog's 12th data release (SDSS DR12 - Alam et al., 2015). 2MASS J, H and K-band magnitudes as well as WISE  $3.35\mu$ m,  $4.6\mu$ m,  $11.6\mu$ m, and  $22.1\mu$ m magnitudes were obtained from the AllWISE Data Release (Cutri et al., 2013). Finally, 1.4GHz radio fluxes were obtained from the FIRST Survey (Helfand, R. L. White, and Becker, 2015). All magnitudes and fluxes were converted to AB magnitudes for consistency. These data are presented in Appendix A.

# 4.5 Data Reduction and Analysis

## **Standard Pipeline Reduction**

Initial data reduction is performed using the standard PCWI Data Reduction Pipeline<sup>1</sup>, which converts raw, 2D science frames into flux-calibrated, three-dimensional cubes with real-world coordinate systems in RA, DEC and wavelength. A detailed description of PCWI calibration products, with useful reference images, is available in Matuszewski et al. (2010).

All frames are initially cosmic-ray subtracted and bias subtracted. As PCWI is a Cassegrain-mounted instrument, there are sometimes slight offsets in the data due to gravitational flexure. These are corrected using a 2D cross-correlation method before the construction of 3D data products. The pipeline then maps from the 2D space of raw images to the 3D image coordinates (x, y, z) and onsky/wavelength coordinates  $(\alpha, \delta, \lambda)$  using a 'continuum bars' image and an 'arc-flat' image, which have known spatial and spectral features, respectively. The uneven illumination of the image slicer is then corrected for in two steps - first correcting the profile within each slice, and then correcting the slice-to-slice variation. Finally, a spectrophotometric standard star observation is used to convert detector counts to physical flux units. The final product of this pipeline is a three-dimensional, flux calibrated data cube for each individual exposure. For the four targets observed without N&S mode, sky subtraction was performed by extracting 2D sky spectra from the adjacent sky frames and scaling them on a slice-by-slice basis.

<sup>&</sup>lt;sup>1</sup>PCWI DRP: https://github.com/scizen9/pderp

# **Cube Correction and Coadding**

The large volume of data in this survey and complex nature of the 3D IFS data required the development of a toolkit for common reduction and analysis functions. CWITools<sup>2</sup> is a Python3 toolkit written specifically for the analysis of PCWI and KCWI data. It is available publicly on GitHub and will be presented in more detail in a future paper.

Before coadding, individual exposure cubes are first corrected by adjusting their world-coordinate system and trimming them. The RA/DEC coordinate system is corrected for each frame using the known location of a visible source in the field (typically the target QSO, though occasionally an adjacent star). The actual position of the source is measured in image coordinates, and then the coordinate system is updated such that that location accurately points to the known RA/DEC. This does not correct for any errors in rotation, though these are expected to be negligible. In a similar way, the wavelength axis is corrected using the positions of known sky emission lines. Finally, the cube is trimmed to only the wavelength range which is shared by all slices (as each slice has a slightly different bandpass), and edge pixels are trimmed off the spatial axes. The corrected and cropped input cubes are then coadded.

CWITools uses a custom-built method for this which calculates the footprint of each input pixel on the coadd frame, and distributes flux onto the coadd grid accordingly. The on-sky footprint of each input frame is calculated, and a new world-coordinate-system representing the coadd frame is constructed so that it encompasses all input data and has an aspect ratio of 1:1. The wavelength axes of the input cubes are first aligned using linear interpolation to perform any sub-pixel shifts (variance is propagated by convolving with the square of the convolution matrix used to shift the data). With the cubes aligned in wavelength, the problem of coadding then becomes two-dimensional. To calculate the footprint of each input pixel on the coadd grid, the vertices of each input pixel is represented as a vector of four  $(x_i^v, y_i^v)$  coordinates, where the *i* subscript denotes the input coordinate system and the superscript v denotes that they represent the pixel vertices (not the center). These vertices are then transformed into a vector of on-sky coordinates (i.e. a vector of  $\alpha^{\nu}$ ,  $\delta^{\nu}$  coordinates) and from there into a vector of coadd frame coordinates  $(x_o^v, y_o^v)$ . A polygon representing the footprint of the input pixel is then created in coadd coordinates, and the overlapping area with each pixel in the coadd

<sup>&</sup>lt;sup>2</sup>CWITools: https://github.com/dbosul/CWITools

grid is calculated. The flux from the input pixel is then redistributed accordingly, following:

$$F_{coadd}(x, y) = \sum_{x_i} \sum_{y_i} F_{input}(x_i, y_i) f(x, y, x_i, y_i)$$
(4.1)

where  $f(x, y, x_i, y_i)$  is the fraction of the footprint of the input pixel  $(x_i, y_i)$  that falls on the output pixel (x, y). Because the pixels are represented as flexible polygons, this method allows for the input of frames with arbitrary spatial resolution and position-angle. It also allows a variance estimate to be propagated following:

$$V_{coadd}(x, y) = \sum_{x_i} \sum_{y_i} V_{input}(x_i, y_i) f^2(x, y, x_i, y_i).$$
(4.2)

However, dividing up the pixel this way is implicitly performing linear interpolation, which introduces covariance between the pixels. We discuss how this is handled in Section 4.5. The final pixel size in the coadded image has a 1:1 aspect ratio with the same plate scale as the x-axis of the input data ( $\simeq 0.55$  arcsec px<sup>-1</sup>).

ID	$\lambda_{ m c}$	$\delta\lambda$	δ	$\mathrm{SB}_{2\sigma}{}^a$
	Å	Å	$\rm km~s^{-1}$	$10^{-18}$ cgs
1	4555	21	1383	2.92
2	4605	13	846	4.44
3	4569	25	1641	6.02
4	4437	20	1352	5.49
5	4559	19	1250	4.39
6	4522	16	1061	4.84
7	4240	20	1415	10.43
8	4160	20	1442	4.85
9	4283	29	2031	7.98
10	4430	23	1557	6.43
11	4522	29	1923	5.78
12	4313	15	1043	5.4

Table 4.3: pNB parameters for the FLASHES Pilot Survey

Continued on next page

ID	$\lambda_{ m c}$	δλ	δ	$\mathrm{SB}_{2\sigma}{}^a$
	Å	Å	km s <sup>-1</sup>	$10^{-18}$ cgs
13	4456	12	807	5.44
14	4128	27	1962	9.52
15	4441	15	1013	5.92
16	4786	34	2131	6.48
17	4205	27	1926	7.99
18	4193	10	715	6.15
19	4365	18	1237	4.81
20	4189	17	1217	5.55
21	4806	21	1310	2.94
22	4237	26	1840	8.43
23	4611	15	975	4.22
24	4442	27	1823	7.79
25	4240	32	2264	7.22
26	4980	10	602	2.76
27	4407	27	1837	9.11
28	4295	13	908	4.83
29	4304	25	1742	9.28
30	4166	17	1224	5.95
31	4208	20	1425	5.56
32	4615	15	975	4.14
33	4590	19	1241	4.93
34	4211	18	1282	5
35	4351	13	896	4.4
36	4188	15	1074	5.2
37	4620	12	779	4.39
38	4488	15	1002	6.13
39	4587	23	1504	5.89
40	4211	23	1638	6.63
41	4497	23	1534	6.66
42	4497	23	1534	6.36
43	4537	25	1653	6.04
44	4396	19	1296	6.15
15	4066	20	1475	4.85

Continued on next page

ID	$\lambda_{ m c}$	δλ	δ	$\mathrm{SB}_{2\sigma}{}^a$
	Å	Å	$km s^{-1}$	$10^{-18}$ cgs
46	4308	19	1323	4.67
47	4650	24	1548	8.09
48	4194	16	1144	7.82

In this section, we describe the steps taken to extract extended  $Ly\alpha$  emission in the CGM and produce scientific products from the data. We initially search for extended emission using a two-dimensional channel map method, which trades spectral resolution for an increased signal-to-noise ratio (SNR). Once emission is identified, we then analyze it in three dimensions to obtain kinematics and spectra.

## **Generation of Pseudo-Narrow-Band Images**

In order to identify extended emission, an initial exploration of the cubes is performed using pNB images, which are narrow-band images formed by collapsing wavelength layers of the data cube. For the purpose of studying extended emission, continuum emission must be subtracted. For each pNB image, a white-light (WL) image is formed by summing ~ 50Å on either side of the current pNB bandpass. Pixels within a circular region of radius ~ 1.5" around the QSO are then used to calculate a set of scaling factors between the images. The scaling factors are sigmaclipped at  $\pm 3\sigma$  and the resulting mean is taken as the global scaling factor for the WL image. The WL image is then scaled and subtracted from the pNB image in a circular region out to ~ 5". Nearby continuum sources are identified using the SDSS catalog from the built-in catalog function of the SAOImage DS9 tool (Joye and Mandel, 2003). These sources are masked and excluded from all subsequent analysis in this work. The masks are shown as black regions in the pNB panels of Figure 4.6.

Variance images are also produced for the pNB images, using the propagated error on the coadded cubes as input. To prioritize the extraction of faint emission on large spatial scales, the data is smoothed with a simple  $5 \times 5$  (pixels) box kernel. This smoothing increases the covariance in the data. In the next section, we describe an empirical variance calibration method which scales the propagated variance estimates to account for covariance.



Figure 4.3: PSF subtraction performed for off-center (i.e. continuum) wavelengths for the first five targets. The columns, from left to right, show: the scaled white-light image, the off-center pNB image, the subtracted off-center pNB image, the Ly $\alpha$  pNB image (for comparison), and the subtracted Ly $\alpha$  pNB image.

# **Covariance in the pNB Images**

IFS data contains covariance between adjacent pixels from resampling onto a regular wavelength or spatial axis, distributing flux onto a new pixel grid when coadding, and any subsequent smoothing or binning steps. This complicates efforts to use standard error analysis when smoothing or summing flux from data cubes. It is extremely complex, and typically beyond the scope of standard data reduction pipelines, to



Figure 4.4: Calibration of variance measurement in FLASHES pilot data. Black crosses indicate individual calibration measurements. The solid black curve indicates the averaged profile, while the grey-shaded region represents the  $\pm 1\sigma$  uncertainty. The solid red curve indicates the functional fit to  $\sigma_{meas}/\sigma_{nocov} = (1 + \alpha_v Log(N_k))$ , with  $\alpha_v = 0.79$ , and the horizontal dashed red line indicates the approximate asymptote for the relationship at large  $N_k$  ( $\beta_v \simeq 2.6$ ).

analytically determine the exact form of this covariance. Because of this, the variance produced by such pipelines underestimate the true noise of the data. As a first step, following the approach of similar studies in the field (Borisova et al., 2016; Arrigoni Battaia et al., 2019b), we re-scale the propagated variance on each coadded cube to match the global noise properties of the cube. This is done comparing the variance with the distribution of voxel values in the cube background (i.e. after masking sources and emission-lines). The variance rescaling factor at this stage is approximately ~ 1.5 for all cubes. However, we also have to take into account the covariance added by smoothing. For smoothing with a uniform box kernel K(m, l) of side  $N_k$  and where K = 1 for all m, l, the propagated variance assuming independent variables is:

$$V'_{nocov}(x,y) = \frac{\sum_{m=-N_k+1}^{m=N_k-1} \sum_{l=-N_k+1}^{l=N_k-1} V(x+m,y+l)}{N_k^2}.$$
(4.3)

The numerator here is a binning operation, while the denominator is a fixed normalization factor which does not add to the covariance. As such, to account for the covariance introduced by this smoothing operation, we adopt the methodology used by the data reduction pipelines for the Calar Alto Legacy Integral Field Area survey (Husemann et al., 2013) and the SDSS-IV MaNGA IFU Galaxy Survey (Law et al., 2016). These pipelines estimate the covariance in spatially summed/binned spectra using an empirical calibration of variance as a function of the kernel size. This is done by binning data by  $N_k \times N_k$  pixels, measuring the noise in the binned signal ( $\sigma_{meas}$ ), and comparing it to the error derived under the assumption of no covariance ( $\sigma_{nocov}$ ). The relationship between these variables is then fit with a functional form following  $\sigma_{meas} = (1 + \alpha_v Log(N_k^2))\sigma_{nocov}$ , and used to calibrate future variance estimates. For large  $N_k$ , where most of the data under the kernel is uncorrelated, this functional form beaks down and instead follows a simple scaling form of  $\sigma_{meas} \simeq \beta_v \sigma_{nocov}$ .

To perform this calibration, we generate a set of pNB images at continuum wavelengths, such that they contain no extended emission and contain mostly empty background after WL subtraction. From our 48 coadded data cubes, we can generate ~ 440 such pNB images. We then measure the noise after smoothing these images with box kernels of size  $N_k = 1, 3, 5, 7$ , and 9, and compare it to the noise estimate from the simply propagated variance. When measuring the noise of the smoothed image, we use only values for which the convolution with the smoothing kernel did not rely on zero padding (this would underestimate the noise). We then fit the above functional form to find  $\alpha_v = 0.79$  and  $\beta_v \approx 2.6$ . Figure 4.4 shows the result of this calibration. The inset shows the error on the calibration itself: the model estimates the variance to within  $\pm 18\%$  ( $\pm 2\sigma$ ). Since we smooth our pNB images with a 5 × 5 box kernel, we rescale the variance by a factor of  $f_{var} = (2.11)^2 \approx 4.45$ . We also scale the variance following this form when calculating the integrated SNR of an extended region.



Figure 4.5: Example of the optimization of pNB image parameters (wavelength center and bandpass) for target HS1700+6416 (ID 1). The colormap in both panels shows the integrated SNR in a circular region of radius 75pkpc around the QSO. The left panel shows the initial coarse grid, which searches a wide velocity range of  $\pm 10000$  km s<sup>-1</sup>. The red lines indicate the  $\pm 1000$  km s<sup>-1</sup> zoom-in for the high resolution grid, shown on the right.

# **Optimizing the pNB Image Parameters**

To identify extended emission in the cubes, a 2D approach using pNB images is adopted. This approach is chosen over a 3D voxel-by-voxel search so that the signal can be integrated along the wavelength axis. The basic approach is to generate sets of pNB images with varying combinations of wavelength center and width and measure the integrated SNR in the vicinity of the QSO (within a projected radius of 75 pkpc). The pNB parameters which optimize the signal in the vicinity of the QSO is then chosen. This is done in two steps: first using a coarse grid with a large range of wavelength/center combinations to find an approximate central wavelength, and then using a fine grid with a smaller range. The pNB centers in the coarse grid range over  $\pm 10000$  km s<sup>-1</sup> in velocity from the peak of Ly $\alpha$  emission in the QSO spectrum, in steps of  $\delta v_{center} = 1000 \text{ km s}^{-1}$ . This large range is motivated by the findings of previous works such as Borisova et al. (2016) that the centroid of Ly $\alpha$ emission can vary by thousands of kilometers per second from the QSO redshift. The pNB velocity width for the coarse grid is varied over a range of 500 - 2500 km s<sup>-1</sup> in steps of  $\delta v_{width} = 500 \text{ km s}^{-1}$ . This process runs semi-automatically, as regular visual inspection of the pNB images is needed to ensure that the integrated SNR is not influenced by systematic errors (under-subtraction of the QSO, bright sky line residuals, etc).

Once an approximate wavelength center is identified from the first step, a higher resolution grid is generated to refine the best center/width. This time, the pNB velocity centers range over a smaller range of  $\pm 1000$  km s<sup>-1</sup> with a step-size of 100 km s<sup>-1</sup>, and the bandwidth ranges from 500 - 2500 km s<sup>-1</sup> in steps of 250 km s<sup>-1</sup>. This grid is used to identify the optimal center and width for the pNB image for each target. Figure 4.5 shows an example of the coarse and fine grids generated during this process for one target (ID 3 / HS1700+6416). For targets in which there is no clear choice of center/width, the default setting is to be centered on the peak of QSO Ly $\alpha$  emission with a velocity width of 1500 km s<sup>-1</sup>. Table 4.3 shows the final parameters for each pNB image alongside the  $2\sigma$  limiting surface brightness achieved in a 1 arcsec<sup>2</sup> aperture.

# **Extracting Emission from pNB Images**

When the optimal center and width of the pNB are identified, the final data products are produced, including the WL image, non-subtracted pNB image, subtracted pNB image, source mask, variance map, and SNR map. These are used to identify regions of extended emission. The data is initially segmented by a threshold of SNR  $\geq 2\sigma$ . The integrated SNR of each region is then calculated (taking covariance into account after summing under the region) and an integrated SNR threshold of SNR<sub>int</sub>  $\geq 4.5\sigma$ is applied. The search is limited to a 250 pkpc<sup>2</sup> box around the QSO. If no regions are found of a sufficient SNR, the target is counted as a non-detection. If there are detected regions, the total integrated SNR of all regions is measured and used to determine the order of the targets (from highest to lowest).

# **Characterizing 2D Morphology**

In order to highlight different characteristics, we measure the size of the nebulae in three ways. First, we use the maximum extent of the nebula from its flux-weighted centroid,  $R_{max}$ . Secondly, we define an *effective* radius to be the radius of an equivalent circular area, i.e.,  $R_{eff} = \sqrt{Area/\pi}$ . We emphasize that  $R_{eff}$  is not a true radius, but a characteristic scale. Finally, we measure the flux-weighted root-mean-square radius,  $R_{rms} = \sqrt{\langle R^2 \rangle_f}$ , using the flux values under the 2D nebular mask. While  $R_{max}$  and  $R_{eff}$  give a sense of the maximum and average extent of the nebula, respectively,  $R_{rms}$  size gives a sense of the characteristic scale at which most of the emission is concentrated.

Beyond measurement of size, the 2D morphology is characterized by three parameters: eccentricity (i.e. asymmetry), displacement, and covering factor. To quantify the symmetry of the nebulae and allow for direct comparison with existing literature, we adopt the same measurement of spatial symmetry as presented in A19. This parameter,  $\alpha$ , is derived from the second-order spatial moments and reflects the ratio of the semi-minor axis (*b*) to the semi-major axis (*a*) of the emission (i.e.,  $\alpha = b/a$ ). We then convert it to an elliptical eccentricity parameter (e), which we find to be more intuitive, following:

$$e = \sqrt{1 - b^2/a^2} = \sqrt{1 - a^2}.$$
 (4.4)

The displacement, which we denote  $d_{QSO}$ , is the projected physical distance (in proper kiloparsecs) between the flux-weighted centroid of the nebular emission (under the mask) and the quasar.

## **Radial Profiles**

Radial surface brightness profiles are measured from a minimum projected radius of 18 pkpc to a maximum radius of 150 pkpc in logarithmic bins of 0.1 dex. All of the detected emission in this sample falls within this range. The 2D object mask is not applied when calculating the circularly averaged surface-brightness profile, but the locations of known and subtracted continuum sources are masked. For non-detections, the wavelength of any CGM Ly $\alpha$  emission is not known, so it may not be contained in the bandpass of the pNB image. For this reason, the averaged radial profile including non-detections may slightly underestimate the true radial profile. The covering factor is calculated using the same radial bins, and defined as the fraction of pixels in each annular region above an SNR of  $2\sigma$ .

# Luminosities

The integrated luminosity of each nebula is calculated following

$$L_{tot} = \frac{(\delta\theta)^2}{4\pi D_L^2(z)} \Sigma_x \Sigma_y SB(x, y) M(x, y)$$
(4.5)

where SB(x, y) is the 2D surface-brightness map in units of erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>, M(x, y) is the binary 2D mask defined earlier,  $\delta\theta$  is the angular size of a pixel, and  $D_L(z)$  is the luminosity distance at the redshift of the target. We note that integrated luminosities are sensitive to the surface-brightness threshold used to

define M(x, y). Any comparison to luminosities reported in other works should consider the differences in cosmic dimming-adjusted surface brightness limits.

## **Point-Source Subtraction in 3D**

CWITools performs 3D point-spread function (PSF) subtraction in a similar fashion to Borisova et al. (2016), which is a basic extrapolation of the pNB method, described earlier, to 3D. A white-light image is formed by summing all of the wavelength layers of the cube, which is then used to identify the positions of any point sources. For each point source above a certain SNR threshold, the following routine is repeated: for each wavelength layer in the cube, a broad-band (i.e., white-light) image centered on the current wavelength layer is formed by summing over a large spectral range (~ 100Å). This image is then scaled and subtracted from the wavelength layer using the method described in Section 4.5. The underlying assumption of this technique is that the shape of the PSF will be dominated by white-light, not nebular emission. In the case of obscured quasars with faint continuum or quasars with particularly bright extended emission, the wavelength range containing nebular emission may be masked to prevent it being used for the white-light image. A small inner radius roughly equal to the seeing ( $\sim 1''$ ) is used to calculate the scaling factors, and the scaled WL image is subtracted out to a larger radius, typically a few times the seeing ( $\sim 5''$ ). Once this PSF subtraction is completed for all detectable point sources, any remaining continuum or scattered light is subtracted (if necessary) using a low-order (k = 1 or 2) polynomial fit to the spectrum in each spaxel. If strong nebular emission is identified, it can be masked during this fitting process to avoid over-fitting. Finally, the PSF cores of bright sources that have been subtracted are masked to prevent noisy residuals influencing any measurements later on. As with the 2D pNB images, the positions of known continuum sources are identified and masked using sources from the 12th SDSS Data Release (Alam et al., 2015).

#### **Integrated Nebular Spectra and Line-Fitting**

To create an approximate 3D mask encompassing the emission, the spatial object mask, M(x, y) is extended along the wavelength axis over the same range as was used to form the final pNB image. Nebular spectra are obtained by summing over the spatial axes under the 3D mask. The spectra are fit with both a simple Gaussian model, a model consisting of multiple (1 - 4) Gaussian components, and a simple linear model. To determine which model best represents the data, we calculate the



Figure 4.6: FLASHES Pilot survey observations (ID 1-6). Each tile is  $250 \times 250 \text{ pkpc}^2$  in size, centered on the QSO. The left-most four columns show a whitelight image,  $\text{Ly}\alpha$  surface brightness, velocity, and dispersion. Surface brightness is in cgs units, erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>. The black bar in the top white-light image shows 100 pkpc and the black square shows the box kernel used to smooth the WL and pNB data. Foreground sources in each field have been masked, with the masked regions shown in black. The rightmost column shows integrated nebular spectra (black) and scaled QSO spectra (grey). The spectra are summed over the object masks and shown in units of  $10^{-17}$  erg s<sup>-1</sup> cm<sup>-2</sup> Å<sup>-1</sup>. Spectra are shown rest-frame wavelength, according to the systemic QSO redshift. Blue lines indicate the fluxweighted centers of nebular emission, while black lines indicate the peak of QSO emission. (Caption continues in next Figure)



Figure 4.7: (continued) A very bright mercury sky emission line (Hg  $\lambda$ 4358.3) is masked in some spectra and shown here as a vertical red band wherever it appears. Empty regions (shown in white) in the pNB images are outside the field of view. An ellipse representing the FWHM of the QSO's PSF is shown in each tile. Red ellipses are used for smoothed PSF (in the WL and pNB images) while black ellipses are used for the unsmoothed PSF (moment maps).



Figure 4.8: (continued)



Figure 4.9: (continued)



Figure 4.10: (continued)


Figure 4.11: (continued)



Figure 4.12: (continued)

Bayesian Information Criterion (BIC) for each, following

$$BIC = n \ln(RSS/n) + k \ln n \tag{4.6}$$

where n is the number of independent variables (i.e., length of the spectrum), k is the number of free parameters in the model, and *RSS* is the residual sum of squares of the model. Lower BIC values indicate a better representation of the data. Weights representing the *relative* likelihood of a set of models can be derived from the BIC values as:

$$w_i = \frac{\exp\left(-\frac{1}{2}\Delta_i(BIC)\right)}{\sum_j \exp\left(-\frac{1}{2}\Delta_j(BIC)\right)}$$
(4.7)

where  $\Delta_i(BIC)$  is the difference between the *i*th BIC value and the minimum BIC value of the set (Wagenmakers and Farrell, 2004). The value  $0 \le w_i \le 1$  corresponds to the relative likelihood that the *i*th model is the best representation of the data (among those considered). The linear model is included as a simpler alternative in order to validate the single-component Gaussian models; if only Gaussian models were considered, the BIC would still indicate a single-component Gaussian as the best fit for pure noise. This multiple-component fit provides an important piece of contextual information when interpreting the global dispersions of the nebulae - as complex line shapes with multiple components can appear quite broad when viewed as a single Gaussian, or otherwise treated as a single kinematic component (e.g., by calculating the second moment).

To create a stacked Ly $\alpha$  spectrum from the individual detections, spectra are linearly interpolated onto a rest-frame wavelength grid ranging from 1200 Å to 1230 Å with a sampling rate of 0.14 Å/px (approximately the PCWI sampling around Ly $\alpha$  at z = 3). As there is more than one measure of redshift (e.g. the flux-weighted center of emission vs. the systemic QSO redshift), we create four versions of the stacked spectrum, each using a different central wavelength (i.e., redshift) to convert to restframe units: (i) the systemic QSO redshift given in DR12Q, (ii) the flux-weighted center of Ly $\alpha$  emission, (iii) the peak of Ly $\alpha$  emission in the QSO spectrum, and (iv) the HeII  $\lambda$ 1640 redshift from DR12Q. The stacked spectra are simple averages of the individual detections, though non-detections are necessarily excluded as no integrated nebular spectrum can be measured.

#### **2D Moment Maps**

Two-dimensional first and second flux-weighted z-moment maps are calculated as:

$$\mu_{\lambda,1} = \frac{\sum_k \lambda(k) I_k}{\sum_k I_k} \tag{4.8}$$

$$\mu_{\lambda,2} = \sqrt{\frac{\sum_{k} (\lambda_k - \mu_1)^2 I_k}{\sum_{k} I_k}}$$
(4.9)

where  $I_k$  is the  $k^{th}$  wavelength layer of the intensity cube and  $\lambda_k$  is the wavelength at that layer. I and  $\mu$  are both two-dimensional arrays with the spatial indices (i, j)omitted for simplicity (i.e.,  $I_k = I_{ijk}$ ). No smoothing is applied to the 3D data prior to calculating the moments.

As a statistical moment is not well defined for a distribution with negative weights, some non-negative threshold must be applied to the spectra before calculating the first or second moment. For bright signals, a high SNR threshold can be applied which rejects virtually all noise while also retaining enough signal for an accurate measurement. However, for fainter signals, it can be challenging to find a threshold which satisfies both of these requirements. A simple positive threshold (i.e.,  $F_{\lambda} > 0$ ) can be applied, positive fluctuations in the background noise will then bias the calculation. For the calculation of the first moment,  $\mu_1$ , an iterative approach can be used to overcome this. The effect of evenly distributed noise (in a well background-subtracted signal) will be to bias the result towards the center of whichever wavelength window is used. If the wavelength window is centered on the true first moment, then this biasing effect will be negligible. As such, if we perform this calculation iteratively, updating the center of the window each time to the new value of  $\mu_1$ , the window center will eventually converge on the true value. If the size of the wavelength window used for the calculation is also reduced as the solution converges, this further mitigates any biasing effect from unevenly distributed noise. We use this method to determine the first moment (i.e. velocity center) of the spectra in each spaxel, with a starting window size of 25Å (to fully explore the range used for the pNB images), reduced in steps of  $\Delta \lambda = 1$ Å until a minimum window size of 10Å is reached.

For the second moment, a convergent method cannot be used to the same effect, as the influence of normally distributed noise on the second moment is to unilaterally increase its value. Instead, we apply a basic non-negative threshold and treat the derived values as upper limits. The spatially resolved maps still provide our only insights into the 2D distribution of the second moment, and as such are valuable despite this limitation. We can rely on line-fitting of the integrated nebular spectrum (see Section 4.5) for more robust measurements of the global dispersions of the nebulae.

Once the moments are calculated,  $Ly\alpha$  velocity and dispersion maps can be derived as:

$$\mathbf{v}(\mathbf{i},\mathbf{j}) = \left(\frac{\mu_{\lambda,1}(\mathbf{i},\mathbf{j}) - \lambda_0}{\lambda_0}\right)\mathbf{c}$$
(4.10)

$$\sigma_{\rm v}({\rm i},{\rm j}) = \left(\frac{\mu_{\lambda,2}({\rm i},{\rm j})}{\lambda_0}\right) {\rm c} \tag{4.11}$$

where  $\lambda_0$  is the flux-weighted average wavelength of the integrated nebular spectrum,  $\mu_{\lambda,1}(i, j)$  is the first moment in wavelength (Eq. 4.8) at that position, and *c* is the speed of light in vacuum. For each nebula, we also calculate the flux-weighted, one dimensional root-mean-square velocity along the line of sight,  $v_{rms} = \sqrt{\langle v^2 \rangle_f}$ . To be clear, this is the root-mean-square of velocities in individual spaxels *relative* to the flux-weighted average velocity of the nebula. Finally, we measure the offset between the flux-weighted average velocity of the nebula and three key wavelengths: the wavelength of Ly $\alpha$  at the systemic redshift of the QSO ( $\lambda_{\alpha,QSO}$ ), the wavelength of the peak of Ly $\alpha$  emission in the QSO spectrum ( $\lambda_{\alpha,peak}$ ), and the wavelength of Ly $\alpha$  at the HeII  $\lambda$ 1640 redshift of the QSO ( $\lambda_{\alpha,HeII}$ ):

$$\Delta v_{\rm QSO} = \left(\frac{\lambda_{\alpha,\rm QSO} - \lambda_0}{\lambda_0}\right)c \tag{4.12}$$

$$\Delta v_{\text{peak}} = \left(\frac{\lambda_{\alpha,\text{peak}} - \lambda_0}{\lambda_0}\right) c \tag{4.13}$$

$$\Delta v_{\text{HeII}} = \left(\frac{\lambda_{\alpha,\text{HeII}} - \lambda_0}{\lambda_0}\right) c. \tag{4.14}$$

 Table 4.4: FLASHES Pilot Survey CGM properties

ID	$L_{43}^{a}$	R <sub>eff</sub>	R <sub>rms</sub>	R <sub>max</sub>	d <sub>QSO</sub>	е	n <sub>reg</sub> <sup>b</sup>	ZLyα	$\Delta v_{QSO}$	$\Delta v_{peak}$	$\sigma_{\rm v}{}^c$	$N_G^d$
	erg s <sup>-1</sup>	pkpc	pkpc	pkpc	pkpc	(0-1)		$(\pm 2\sigma)$	km s <sup>-1</sup>	km s <sup>-1</sup>	km s <sup>-1</sup>	(0-4)
1	4.3	41.4	28.3	63.0	21.8	0.80	2	$2.746 \pm 0.0008$	+692	-119	395	2
2	9.4	54.8	47.2	110.2	9.0	0.67	3	$2.788 \pm 0.0010$	-178	-171	325	2
3	8.0	45.1	40.7	83.8	41.0	0.82	1	$2.759 \pm 0.0010$	+992	+1624	459	2
4	6.3	43.0	37.9	69.4	10.5	0.88	1	$2.651 \pm 0.0012$	+871	+2179	473	3
5	3.4	30.2	22.4	45.5	19.5	0.87	1	$2.747 \pm 0.0014$	-393	+93	263	1
6	5.4	38.5	32.9	55.9	8.2	0.79	1	2.721±0.0012	+1813	-65	265	2
7	6.1	42.5	45.4	93.1	7.0	0.56	2	$2.487 \pm 0.0012$	-351	+459	707	2
8	2.5	30.8	25.8	53.3	21.4	0.84	1	$2.425 \pm 0.0010$	+378	+852	288	3
9	3.9	33.9	39.7	77.2	7.7	0.72	2	$2.520 \pm 0.0014$	+871	+708	560	1
10	3.2	28.8	28.1	68.0	11.9	0.82	1	$2.642 \pm 0.0016$	+1644	-110	474	2
11	3.0	28.2	22.8	45.7	25.5	0.86	1	$2.714 \pm 0.0018$	-1646	+928	664	2
12	2.0	25.4	17.0	31.1	27.0	0.60	1	$2.551 \pm 0.0016$	-1433	-738	277	1
13	2.1	28.7	23.8	56.3	34.1	0.73	1	$2.663 \pm 0.0016$	+601	-60	202	1
14	2.6	27.4	58.6	119.8	20.9	0.96	3	$2.394 \pm 0.0022$	+1694	+1132	654	2
15	2.4	25.6	22.9	54.0	19.5	0.78	2	$2.652 \pm 0.0016$	+1025	+59	337	1
16	5.5	31.0	36.8	72.5	10.2	0.69	4	$2.935 \pm 0.0016$	+354	+360	627	2
17	3.6	31.4	36.3	76.1	17.1	0.72	3	$2.454 \pm 0.0012$	+1553	+389	571	4
18	1.8	28.4	48.9	90.3	13.8	0.97	3	$2.448 \pm 0.0018$	-216	-192	214	1
19	1.3	26.2	32.7	52.1	9.8	0.99	2	$2.595 \pm 0.0018$	+332	+719	137	1
20	1.5	27.7	30.9	49.1	60.7	0.70	2	$2.443 \pm 0.0022$	+1871	+944	430	1
21	1.2	25.3	34.1	54.2	28.1	0.96	2	$2.950 \pm 0.0018$	+1029	+406	374	2
22	1.8	23.4	16.3	31.6	33.9	0.51	1	$2.486 \pm 0.0016$	+2753	+888	424	2
23	1.0	18.7	15.5	36.5	22.1	0.93	1	$2.792 \pm 0.0018$	+1031	-50	390	2
24	2.9	26.0	26.6	45.4	5.9	0.64	1	$2.650 \pm 0.0012$	+909	+564	608	2
25	2.7	25.1	20.9	42.8	30.5	0.87	1	$2.487 \pm 0.0018$	+1367	+418	707	2
Continued on next page												

ID	$L_{43}^{a}$	R <sub>eff</sub>	R <sub>rms</sub>	R <sub>max</sub>	d <sub>QSO</sub>	е	n <sub>reg</sub> <sup>b</sup>	z <sub>Lyα</sub>	$\Delta v_{QSO}$	$\Delta v_{peak}$	$\sigma_{\rm v}{}^c$	N <sub>G</sub> <sup>d</sup>
	erg s <sup>-1</sup>	pkpc	pkpc	pkpc	pkpc	(0-1)		$(\pm 2\sigma)$	km s <sup>-1</sup>	km s <sup>-1</sup>	km s <sup>-1</sup>	(0-4)
26	0.7	19.2	57.6	79.0	4.3	~1.00	2	$3.097 \pm 0.0022$	+1553	+1280	240	1
27	2.1	16.6	12.3	27.4	21.0	0.67	1	$2.623 \pm 0.0020$	+281	+317	408	1
28	1.1	21.8	17.5	36.0	27.6	0.71	1	$2.533 \pm 0.0020$	+781	+273	419	1
29	2.4	21.8	30.1	59.1	15.7	0.94	2	$2.537 \pm 0.0020$	-158	+53	488	2
30	1.1	19.9	23.3	41.9	20.0	0.98	1	$2.427 \pm 0.0030$	+1327	+674	360	1
31	0.7	16.9	13.0	25.3	23.7	0.91	1	$2.462 \pm 0.0022$	+1377	-554	499	2
32	0.7	17.6	14.0	25.9	79.4	0.89	1	$2.793 \pm 0.0022$	+583	+605	390	3
33	1.2	18.2	21.8	40.5	22.6	0.97	1	$2.771 \pm 0.0022$	-571	+46	196	1
34	0.7	18.1	12.9	24.2	34.8	0.76	1	$2.461 \pm 0.0016$	+2659	+1309	427	2
35	0.7	18.6	19.1	37.0	21.8	0.92	1	$2.575 \pm 0.0022$	+2179	-153	276	1
36	0.6	16.2	11.7	19.0	24.3	0.75	1	$2.444 \pm 0.0032$	-455	-348	143	2
37	0.4	13.6	15.8	28.3	71.9	0.98	1	$2.800 \pm 0.0020$	+67	+613	194	1

(a)  $L_{43} = L/10^{43} \text{ erg s}^{-1}$ 

(b)  $n_{reg}$  is the number of distinct spatial components.

(c)  $\sigma_v$ ; the global dispersion as determined by a single-component Gaussian fit.

(d)  $N_G$  is the best-fit number of Gaussian components according to the Bayesian Information Criterion.



Figure 4.13: Cumulative distributions of the sizes of the detected nebulae in the FLASHES Pilot sample, as measured using: effective radius ( $R_{eff}$ ), maximum radial extent ( $R_{max}$ ), and flux-weighted RMS radius ( $R_{rms}$ ).

## 4.6 Results

In this section, we present the 2D morphologies, eccentricities, radial profiles, kinematic properties, and integrated spectra of the nebulae detected in the FLASHES Pilot sample. For the survey as a whole, we present an averaged radial profile, covering factors, and distributions of kinematics. In order to provide a more complete physical picture of each QSO environment, the basic observational data (surface brightness, velocity, dispersion, and integrated spectra) are displayed side-by-side in the extended Figure 4.6 for each target.

#### Size and Luminosity

The leftmost column of Figure 4.6 shows the pNB images generated following Section 4.5. An average limiting surface brightness of ~  $6 \times 10^{-18}$  erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup> (in a 1 arcsec<sup>2</sup> aperture) was achieved. The individual limiting surface brightnesses are presented alongside the observational details in Table 4.1 We detect nebulae (i.e. regions of emission with SNR<sub>int</sub> > 4.5) around 37 of the 48 objects in our sample. Of these, only one has an effective diameter  $D_{eff} = 2R_{eff} \ge 100$  pkpc. Excluding the four targets obtained from literature, which were previously known to contain extended emission, we find a detection rate of 33/44. The nebulae are found to have projected radii on the order of tens of proper kiloparsecs, with  $R_{eff} \simeq 13 - 55$  pkpc and  $R_{rms} \simeq 12 - 59$  pkpc. The maximum radial extent of the nebulae are found to span a much larger range than the effective radii ( $R_{max} \simeq 19 - 120$  pkpc) indicating

	Min(R)	Max(R)	Mean(R)	Median(R)	$\sigma(\mathbf{R})$
	pkpc	pkpc	pkpc	pkpc	pkpc
R <sub>max</sub>	19	120	42	42	32
R <sub>eff</sub>	14	55	21	23	14
R <sub>rms</sub>	12	59	22	22	16

 Table 4.5: Distributions of Measured Sizes

some degree of asymmetry. We plot the cumulative distribution functions for each measurement in Figure 4.13. Table 4.5 below summarizes the distributions of these three parameters. The integrated luminosities range from  $L_{min} = 0.4 \times 10^{43}$  erg s<sup>-1</sup> to  $L_{max} = 9.4 \times 10^{43}$  erg s<sup>-1</sup>, with mean  $L_{avg} = 2.7 \times 10^{43}$  erg s<sup>-1</sup> and standard deviation  $\sigma_L = \pm 2.13 \times 10^{43}$  erg s<sup>-1</sup>.

## **2D Morphology**

From a quick glance at the pNB images in Figure 4.6, it is clear that there is quite a spread in the spatial symmetry of the nebulae. As discussed in Section 4.5, we quantify this using the eccentricity parameter,  $0 < e \leq 1$ . The value for each target is presented in Table 4.4. The detected nebulae are found to exhibit eccentric morphologies, ranging from a minimum of e = 0.51 to a maximum of  $e \sim 1$ , with a mean of 0.82 and a standard deviation  $\sigma_e = 0.13$ . A number of targets with  $e \simeq 0.9 - 1.0$  appear to be the result of two or more co-linear patches of emission (IDs 14, 18, 19, and 26). To provide some context, we present the number of distinct spatial components in each object mask alongside the eccentricity in Table 4.4. It is important to remember that what is being measured here is the collective eccentricity of the detected regions, and that - with deeper sensitivity fainter emission filling the space between and around these regions might be detected, which would lower the eccentricity. We explore the relationship between the surface brightness threshold and measured eccentricity in more detail in Section 4.7.

The distance between flux-weighted center of mass of the detections and the QSO has a mean value of  $d_{QSO,avg} = 18$  pkpc and also a spread of  $\sigma(d_{QSO}) = 18$  pkpc. Of the 37 detections, 34 have centers of mass within 50 pkpc of the QSO, while three (IDs 20, 32, 37) have large displacements. While ID 20 appears to have some connection to the QSO, IDs 32 and 37 appear similar in nature to the displaced emission seen in A19's target 25.

## **Radial Profiles**



Figure 4.14: Circularly averaged radial profiles of the detected CGM, centered on the QSOs. The top panel shows three examples: a bright detection (ID2), an intermediate detection (ID 22) and a non-detection (ID 48). The middle panel shows the averaged profiles in observed surface brightness, with a Sérsic fit and a power-law fit. The bottom panel shows the average of the profiles after scaling each by  $(1 + z)^4$  to correct for cosmological surface brightness dimming, with the same fits. The x-axis is shown in log-scale.

Figure 4.14 shows the average radial surface-brightness profiles of the FLASHES Pilot survey observations. The average radial surface brightness profile peaks at around  $SB_{max}^{obs} \simeq 6 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2} (SB_{max}^{(1+z)^4} \simeq 1 \times 10^{-15} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}).$  On average, the bulk of emission appears to fall within 50 pkpc of the quasar. We



Figure 4.15: Covering fraction of  $SNR \ge 2\sigma$  as a function of radius. The average for all fields in the sample is shown with black diamond markers, while the average of all detections is shown with blue circle markers.

fit two models to each profile: a power law model with the form  $I(R) = I_0(R/R_e)^{\alpha}$ and a Sérsic profile. For both observed and adjusted profiles, the emission appears to be best described by a profile with Sérsic index  $n \approx 0.5 - 0.6$  and half-light radius  $R_e \approx 24$  pkpc. The observed profile has intensity (surface brightness) at the half-light radius  $I_e = 5 \times 10^{-18}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>, while the adjusted profile has  $I_e^{(1+z)^4} \approx 0.8 \times 10^{-15}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>. We note that an exponential profile is a Sérsic wih n = 1, and that a parameter space of n = 0.1 - 6.0 was explored during the fitting process using a stochastic optimizer (differential evolution from SciPy - (Virtanen et al., 2019; Storn and Price, 1997)) which is less susceptible to local minima than standard gradient descent algorithms.

Figure 4.15 shows the covering factor as a function of projected radius for the same two sample-wide averages. There is a stark contrast between the peak value of  $\sim 30\%$  for the sample-wide average and the near unity covering factor reported by A19. Even among the detections, the average covering factor does not exceed 40%. We discuss these findings further in Section 4.7.

#### 2D Kinematic (Moment) Maps

The second column from the left in Figure 4.6 shows 2D Ly $\alpha$  velocity (first wavelength moment) maps generated as discussed in Section 4.5. The majority of velocities fall within ±300 km s<sup>-1</sup> of the flux-weighted mean velocity of each nebula. The vast majority of the targets do not exhibit any clear kinematic structure. However, two targets (ID 4 and ID 7) stand out from the rest of the sample in this regard. The Eastern side of the extended emission around target 4 appears to be mostly blueshifted, while the Western side appears to be mostly redshifted. For target 7, the South/South-East side of the nebula appears to be broadly redshifted while the North/North-West side is mostly blueshifted. Determining the significance of such structures is non-trivial given the spectral resolution and spatial covariance in the data. We thus present a full discussion on tests for kinematic coherence in the data in Section 4.7.

The third column from the left in Figure 4.6 shows two-dimensional maps of the second wavelength/velocity moment (i.e. velocity dispersion). What appears immediately obvious is that the average dispersions of the nebulae vary significantly, over a range of  $\sim 200 - 400$  km s<sup>-1</sup>. Within the individual nebulae, it is difficult to recognize any clear patterns. It is worth repeating here (as discussed in Section 4.5) that these dispersions are upper limits and are influenced by the size of the wavelength window used to calculate them. To obtain more accurate dispersion maps, deep observations are required as they will allow line fitting techniques to be used on a spaxel-by-spaxel basis.

Figure 4.17 shows the distributions of three velocity offsets,  $\Delta v_{QSO}$ ,  $\Delta v_{peak}$ , and  $\Delta v_{HeII}$ . The distribution of velocity offsets with respect to the systemic redshift ( $\Delta v_{QSO}$ ) is spread over a wide range, from  $\Delta v_{QSO}^{min} = -1647 \text{ km s}^{-1}$  to  $\Delta v_{QSO}^{max} = +2754 \text{ km s}^{-1}$  with a median of  $\Delta v_{QSO}^{med} = +871 \text{ km s}^{-1}$  and a standard deviation of  $\sigma(\Delta v_{QSO}) = 994 \text{ km s}^{-1}$ . The distribution of offsets with respect to the peak of Ly $\alpha$  emission in the QSO spectrum is more concentrated, ranging from  $\Delta v_{peak}^{med} = -738 \text{ km s}^{-1}$  to  $\Delta v_{peak}^{max} = +2179 \text{ km s}^{-1}$  with a median value of  $\Delta v_{peak}^{med} = +390 \text{ km s}^{-1}$  and a standard deviation of  $\sigma(\Delta v_{peak}) = 606 \text{ km s}^{-1}$ . Finally, the spread in velocity with respect to  $z_{HeII}$  is the widest of all, ranging from  $\Delta v_{HeII}^{min} = -1090 \text{ km s}^{-1}$  to  $\Delta v_{HeII}^{max} = +3709 \text{ km s}^{-1}$  with a standard deviation of  $\sigma(\Delta v_{HeII}) = 1130 \text{ km s}^{-1}$ . The median of distribution is also significantly red-



Figure 4.16: Top panel: global dispersions of the detected nebulae as measured from a single-component Gaussian fit. Bottom panel: dispersions of individual features when fitting spectra with a composite model of 1-4 Gaussian components.

shifted ( $\Delta v_{\text{HeII}}^{\text{med}}$  = +1195 km s<sup>-1</sup>).

The rightmost column in Figure 4.6 shows the integrated nebular spectra, extracted from the data cubes by first applying the 2D emission mask and summing over the spatial axes. Fits to the data indicate that 33/37 of the profiles can be decently described by a one- or two-component Gaussian fit, with four targets exhibiting a more complex line structure. We note that, as these spectra are spatially integrated, the line shape may be a result of the superposition of spatially separated components as well as being influenced by Ly $\alpha$  radiative transfer within a single, unresolved emitter. Given that the global dispersion will be heavily influenced by the presence of multiple kinematic components, we present two sets of measurements in Figure 4.16 in order to distinguish between the extrinsic (i.e., superposition of spatially separated components) and the intrinsic (i.e., line broadening) dispersion. The former is measured as the width of single-component Gaussian fits (top panel). These dispersions range from  $\sigma_v^{min} = 143 \text{ km s}^{-1}$  to  $\sigma_v^{max} = 708 \text{ km s}^{-1}$ , with a mean of  $\sigma_v^{avg} = 399 \text{ km s}^{-1}$  and a  $1\sigma$  spread in this distribution of 154 km s<sup>-1</sup>. The latter



Figure 4.17: Distributions of CGM Ly $\alpha$  velocity offsets with respect to different redshifts. The top panel shows velocity with respect to the best-fit SDSS/DR12Q QSO redshift. The middle panel shows velocity with respect to the peak of Ly $\alpha$  emission in the QSO spectrum. The bottom panel shows velocity offset with respect to the HeII  $\lambda$ 1640 redshift from SDSS.

Redshift	$F_{\lambda}{}^{a}$	Vavg	$\sigma_{ m v}$			
	$erg s^{-1}cm^{-2}Å^{-1}$	$km s^{-1}$	km s <sup>-1</sup>			
ZLyα	3.0	8	430			
Zpeak	1.4	+311	721			
Z <sub>QSO</sub>	0.9	+754	1049			
Z <sub>HeII</sub>	0.8	+1367	1035			
(a) Amplitude of Gaussian fit.						

Table 4.6: Properties of Stacked Ly $\alpha$  Profiles

is indicated by the dispersions of the individual Gaussian components wherever a multi-Gaussian (i.e. 1-4 Gaussian components) is the best-fit model. With few exceptions, these dispersions are found to be < 400 km/s. The single-component dispersion and the best-fit number of peaks are presented in Table 4.4.

# Stacked Lya Profiles

Figure 4.18 shows stacked Ly $\alpha$  profiles of the detected CGM emission in the pilot sample, converted to rest-frame units using (i) the redshift of the CGM Ly $\alpha$  emission



Figure 4.18: Stacked Ly $\alpha$  profiles of the CGM detections in the FLASHES pilot survey. Different colors indicate different redshifts used to convert from observed to rest-frame wavelengths: the redshift of the CGM Ly $\alpha$  emission itself, the redshift of the peak of Ly $\alpha$  emission in the QSO (blue), the QSO's systemic redshift from DR12Q (green), and the HeII  $\lambda$ 1640 redshift from SDSS (red).

in each field  $(z_{Ly\alpha})$ , (ii) the redshift corresponding to the peak of  $Ly\alpha$  emission in the QSO spectra  $(z_{peak})$ , (iii) the SDSS/DR12Q best-fit systemic redshift of the QSO  $(z_{QSO})$ , and (iv) the redshift of HeII emission in the QSO spectrum  $(z_{HeII})$ . The averaged line profiles exhibit typical Gaussian shapes, with widths reflecting the velocity distributions of the emission relative to each redshift. Table 4.6 presents the amplitude, mean and standard deviation of each stacked profile. With the exception of the  $z_{Ly\alpha}$ -aligned profile, all of the stacked spectra have a clear redward bias.

## 4.7 Discussion

## From Non-Detections to Giant Lya Nebulae

Borisova et al. (2016) report ubiquitous giant nebulae ( $R_{max} \ge 50$  pkpc) in their sample of 19 quasars at  $z \sim 3.5$ , with a limiting sensitivity of  $10^{-18}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>



Figure 4.19: Comparison of  $L_{Ly\alpha}/R_{max}^2$  as a function of  $R_{max}$  for different surveys. The top panel shows the comparison for sizes in proper kiloparsecs, while the bottom panel shows the same comparison for comoving kiloparsecs. The quantity  $L_{Ly\alpha}/R_{max}^2$  should depend only on the intrinsic radial surface brightness profile of the emission, so comparing nebula of equal size under this metric provides an equitable comparison of the average surface brightness of *detected* regions.

in a 1 arcsec<sup>2</sup> aperture in a 1.25Å layer. Arrigoni Battaia et al. (2019b) report ubiquitous nebulae on scales of tens to hundreds of pkpc around their sample of 61  $z \sim 3.1$  quasars with similar sensitivity to B16. Cai et al. (2019) report nebulae with projected diameters greater than 50 pkpc for 14/16  $z \simeq 2.1 - 2.3$  QSOs, again at comparable sensitivity but at significantly lower redshift. Our work now reveals nebulae around 37/48  $z \simeq 2.3 - 3.1$  quasars on spatial scales of tens of pkpc. Because our detection method used wavelength integrated data, there is no perfect one-to-



Figure 4.20: Cumulative distribution of nebular eccentricities for the FLASHES survey, A19, and C19. The solid black line shows the distribution as measured from the FLASHES  $1\sigma$  contour object masks. The dashed black line shows those measured from  $2\sigma$  contour object masks. The red line shows the distribution of values presented in A19. The green dashed line shows the distribution from C19.

one sensitivity comparison with the above surveys. The average dimming-adjusted radial profile measured here appears to be almost an order of magnitude fainter than that reported in A19, with a peak brightness of  $SB_{max}^{adj} \simeq 10^{-15}$  compared to  $10^{-14}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>. C19 reports a median surface brightness profile which is also significantly fainter than both A19 and B16, albeit slightly brighter than our average.

From this picture, it appears possible that there is some redshift evolution from  $z \sim 3$  to  $z \sim 2$  towards lower average Ly $\alpha$  surface brightness in the vicinity of QSOs. However, when comparing averaged radial surface brightness profiles, there is a degeneracy between the covering fraction of emitting gas and the average surface brightness of that gas; a faint nebula covering a large area factor may have the same circularly averaged radial surface-brightness profile as a small but bright nebula. Assuming that the luminosity, L grows approximately as  $L(R) \propto R^2$ , where R is the radius, the quantity  $L(R)/R^2$  should depend only on the intrinsic radial surface brightness profile of the emitting gas. Comparing this quantity for nebulae of similar

size then provides a comparison of the average intrinsic surface brightness within the nebular region, which can be used to distinguish between the two above scenarios.

In Figure 4.19, we compare the detected emission from A19, B16, C19, and this paper in the parameter space of  $L(R)/R^2$  vs. R, where we have used  $R_{max}$  as a proxy for size because it is readily available in all studies. We perform this comparison both for sizes measured in pkpc and comoving kiloparsecs (ckpc). No obvious overall difference emerges between the studies. From this comparison, we find that there is no systematic difference in the average intrinsic surface brightness of the detected regions at different redshifts; i.e. nebulae of similar size have similar average brightness. This implies that the driving factor between the fainter circularly averaged profiles in this work (and C19) is the lower covering fraction of detected gas, rather than globally fainter emission. Although they overlap with the other surveys, the average surface brightness measured by B16 does appear systematically higher than the other surveys, possibly because their sample focused on brighter QSOs (although we measure no significant relationship between QSO magnitude and Ly $\alpha$  luminosity here).

## Asymmetry of the Ly $\alpha$ Emission

It is clear from visual inspection of the pNB images in Figure 4.6 alone that there is a pronounced degree of asymmetry in many of the detected nebulae. The distribution of values of the eccentricity parameter (*e*) supports this impression, with a mean value of 0.82. Figure 4.20 compares the cumulative distributions of e for the FLASHES pilot sample with those presented in C19 and A19 (none were presented in B16). We use the two-sample Kolmogorov-Smirnov (K-S) test to compare the distributions of *e*, and find that we can reject the null hypothesis (that the two samples are from the same underlying distribution) when comparing to A19 (p < 0.002 - see Table 4.7 for exact values). However, when comparing to C19 using the two sample K-S test, we cannot reject the null hypothesis ( $p \approx 0.44$ ). Table 4.7 summarizes the results of the K-S tests. The means of the distributions for A19, C19, and this work, respectively, are  $e_{A19} = 0.69$ ,  $e_{C19} = 0.82$ , and  $e_F = 0.82$ , with 1 $\sigma$  spreads in each distribution of  $\sigma(e_{A19}) = 0.15$ ,  $\sigma(e_{C19}) = 0.1$ , and  $\sigma(e_F) = 0.13$ .

It is important to note that there are significant differences in extraction technique and sensitivity between our work, A19, and C19. Changes in morphology may



Figure 4.21: Change in eccentricity as a function of the increase in limiting surface brightness. Contours show a Gaussian kernel density estimate and the black line with shaded region shows the best-fit linear model with  $\pm 2\sigma$  slope uncertainty. The linear regression shows a strong correlation in which eccentricity increases as the surface brightness threshold increases.

equally be a result of changing sensitivity limits as of intrinsic differences in CGM properties. The different extraction techniques, in particular, make a one-to-one comparison of limiting sensitivity very difficult. We can, however, test whether the eccentricity itself depends on the limiting surface brightness used within our sample. In Figure 4.21, we show the distribution of eccentricities calculated for different SNR isophotes ( $SNR_{iso} = 1.5, 2.0, 2.5, and 3.0$ ) for data within 100 pkpc of the QSO. Eccentricities are only calculated if there are at least 10 spaxels within the isophotal threshold. The top panel shows the same data with the SNR isophotes converted to absolute surface brightness isophotes. Linear regression to the data in both panels does not indicate a significant correlation. This, combined with the fact that C19 report similar eccentricities having higher sensitivity than both the FLASHES Pilot survey and AB19, indicates that limiting sensitivity is at least not the primary driver of the increased eccentricity. As the Ly $\alpha$  emission we are observing is likely powered by ionizing emission from the QSO, both the illumination and intrinsic distribution of gas play important roles in determining the morphology of the detected nebulae. These findings, combined with the finding from the previous section - that a lower covering factor seems to be driving the reduced average



Figure 4.22: Parameter space plots for the nebulae detected in the FLASHES Sample.  $R_{eff} = \sqrt{Area/\pi}$  is the effective size, e is the eccentricity ( $0 \le e < 1$ ),  $z_{Ly\alpha}$  is the redshift of the nebular Ly $\alpha$  emission,  $\Delta v_{QSO}$  and  $\Delta v_{peak}$  are the velocity offsets with respect to the systemic redshift and peak of QSO Ly $\alpha$  emission, respectively,  $\sigma_G$  is the standard deviation (in km s<sup>-1</sup>) of the best-fit single-peaked Gaussian line profile, and M<sub>i</sub> is the absolute i-band magnitude of the quasar. Contours in each plot show Gaussian Kernel Density Estimates of the 2D distribution. Black and red lines show linear regression models with p < 0.05 and p < 0.01, respectively. The r-values of these linear regressions are shown on the relevant tile. Faint grey lines indicate linear regression models with  $p \ge 0.05$  (i.e., no correlation clearly indicated).

Test	K-S Statistic	p-value
FLASHES v. A19	0.377	$1.9 \times 10^{-3}$
FLASHES v. C19	0.245	$4.4 \times 10^{-1}$

Table 4.7: Comparison of Eccentricity Distributions

surface-brightness - paint a picture of a  $z \sim 2 - 3$  CGM that is increasingly patchy and asymmetric at lower redshifts.

## **PSF** Asymmetry



Figure 4.23: Measured asymmetry (i.e. the y:x aspect ratio, determined by a 1D Gaussian fit to the collapsed PSF along each axis) of the PSF in the final pNB images, shown in Figure 4.6.

# **Relationships between Global Nebular Properties**

In Figure 4.22, we present a corner plot comparing some key measured properties of the detected nebulae. For each comparison, we test for any relationship between the parameters using a simple linear regression. If the result appears significant (i.e. has a p-value < 0.05), we plot the best-fit line and show the r-value of the linear regression, indicating the strength of the correlation  $(-1 \le r \le +1)$ . The strongest correlation found is no surprise - being between effective radius and luminosity. Visual inspection of this tile confirms a roughly quadratic relationship, as can be expected for these parameters. Eccentricity appears inversely related to luminosity, which can be explained if smaller detections tend to be more eccentric (see Section 4.7). We find no significant correlation between the absolute i-band magnitude of the QSO,  $M_i$ , and the effective size. A weak correlation is found between the velocity offset from the systemic redshift ( $\Delta v_{QSO}$ ) and the global dispersion as measured from a Gaussian fit ( $\sigma_G$ ). It is not immediately obvious what might cause such a relationship, though it seems plausible that the dispersion and local absorption are both influenced by certain global properties of the surrounding CGM (such as the average temperature of Ly $\alpha$  emitting/absorbing gas). The correlation is not strong enough to motivate a thorough study here, but presents an interesting element to test with the more sensitive deep survey data. Beyond these few instances, there appear to be no significant correlations between any of the other measured nebular properties.

## Kinematics of the Ly $\alpha$ Emission

The flux-weighted centroid of the Ly $\alpha$  emission measured in our sample varies by many hundreds of km s<sup>-1</sup> from the systemic redshift of the QSO ( $\sigma(\Delta v_z) = 994$  km s<sup>-1</sup>) and from the peak of QSO Ly $\alpha$  emission ( $\sigma(\Delta v_{peak}) = 606$  km s<sup>-1</sup>). The spread with respect to the SDSS HeII  $\lambda 1640$  redshift is even more significant, with  $\sigma(\Delta v_{HeII}) = 1130$  km s<sup>-1</sup>. This spread, comparable to that reported in A19, highlights the challenge faced by narrow-band imaging searches for Ly $\alpha$  emission from the CGM around specific targets. All three velocity offset distributions, shown in Figure 4.17, present a clear bias towards the red. Some of this effect may be attributable to the re-absorption of blueshifted emission (i.e. rest-frame  $\lambda \le 1216$  Å) in the intervening IGM. However, it could also indicate that the majority of detections feature outflowing gas with a red-dominated line profile (e.g Gronke, Bull, and Dijkstra, 2015).

The average dispersions of the nebulae, shown in the third column of Figure 4.6, appear to be in agreement with the finding of A19, in that nearly all targets have mean dispersions  $\sigma_{avg} \leq 400 \text{ km s}^{-1}$ . As we note in Section 4.5, the statistical second moments here provide upper limits in the presence of strong noise. However, this finding is supported by our line-fitting analysis of the integrated nebular spectra. The global dispersions of the integrated spectra have a mean of  $\sigma_v = 399 \text{ km s}^{-1}$  even when modelled using a single Gaussian component. In the bottom panel of the same figure, we see that the vast majority of global dispersions above 400 km s<sup>-1</sup>] disappear when multiple Gaussian components are allowed, indicating that these line-widths are the result of complex line shapes, attributable in part to both the



Figure 4.24: A comparison of the RMS line-of-sight velocity detected in the FLASHES pilot survey with gravitational motions in an NFW halo. The dashed blue line and blue shaded region represent the average and  $\pm 1\sigma$  spread in the line-of-sight RMS velocities of FLASHES pilot nebulae, respectively. The solid black curve shows the (maximum) RMS line-of-sight velocity of an NFW halo as a function of halo mass following Munari et al. (2013) ( $\sigma_{1D} = 0.68v_{200}$ , where  $v_{200}$  is the circular velocity at the virial radius). The grey shaded region indicates the halo masses of high-luminosity QSOs (HLQSOs) at a redshift of z = 2.7 in Trainor and Steidel (2012).

superposition of spatially distinct kinematic components and intrinsically complex spectra (i.e. within a single spaxel).

Approximately one third (15/37) of the detected nebulae appear to be best fit by a single peak, while a plurality (17/37) seem to be best described by a two-component fit, and the remaining few (5/37) have more complex line shapes with three or more

components. We note that these best-fit measurements, determined using the BIC, only represent the relative likelihood of the models considered, and the presence of considerable noise and occasional systematics such as bright sky-line residuals should be taken into account when interpreting these results. For example, for target 24, a bright sky line (Hg  $\lambda$ 4358.3) coincided almost exactly with the position of the redshifted Ly $\alpha$  line. A small wavelength region around this line had to be masked before analyzing the data, so the line complexity here is likely artificial.

In Figure 4.24, we compare the measured RMS velocities to the line-of-sight RMS velocity (v<sub>RMS.1D</sub>) expected for Nevarro-Frenk-White (NFW) halos (Navarro, Frenk, and White, 1997) with concentration parameter c = 3.5. We measure the RMS velocity of each nebula detected in the sample and find the average value to be  $v_{RMS,avg} = 208 \pm 128$  km s<sup>-1</sup>, which corresponds to the values expected from a halo mass range of  $\text{Log}_{10}(M_h[M_{\odot}]) = 12.2^{+0.7}_{-1.2}$ . Trainor and Steidel (2012) measured the halo masses for a sample of high-luminosity QSOs at a redshift of  $z \simeq 2.7$ , and found the range to be  $Log_{10}(M_h[M_{\odot}]) = 12.3 \pm 0.5$ . An analysis of the clustering of  $z \sim 1.5$  QSOs in the 2dF QSO Redshift Survey by da Ângela et al. (2008) found that QSOs tend to inhabit  $M_h \simeq 3 \times 10^{12} h^{-1} M_{\odot}$ , regardless of luminosity or redshift, while M. White et al. (2012) studied the clustering of  $2.2 \le z \le 2.8$ QSOs in the Baryonic Oscillation Spectroscopic Survey and found their results to be consistent with QSO host halo masses of  $M_h \simeq 10^{12} h^{-1} M_{\odot}$ . We thus find that the RMS velocity values among the FLASHES pilot detections are broadly consistent with those expected from gravitational motions in the host dark matter halos of QSOs at their redshift (median redshift  $z \simeq 2.63$ ). It is important to note that there are many more effects contributing to the observed Ly $\alpha$  kinematics beyond gravitational motions in an ideal NFW halo; e.g. outflows, mergers, AGN feedback, and radiative transfer. This comparison was performed to test for any clear inconsistency between the expected and measured kinematics. The fact that the results appear to be consistent with halo motions only tells us that we cannot directly rule out an interpretation of the moment maps as reflecting physical motions, not that this is the most appropriate interpretation. The FLASHES deep survey will provide us with an opportunity to perform more detailed modeling of kinematics, including radiative transfer effects.

#### Coherence in the Ly $\alpha$ Moment Maps

As we mentioned in Section 4.7, two targets - IDs 4 and 7 - appear to exhibit some coherent kinematic structure, with regions that are systematically redshifted or blueshifted with respect to their flux weighted center. We test for the presence of systematic structure in two ways: first by measuring the specific projected angular momentum of each nebula and second by performing a simple comparison of 2D kinematic models with and without shear.

## **Specific Projected Angular Momentum**

We define the flux-weighted average of the projected specific angular momentum for each nebula as:

$$\langle \vec{j} \rangle_f = \frac{\sum_x \sum_y F(x, y) \vec{R}_{\perp}(x, y) \times \vec{v}_z(x, y)}{\sum_x \sum_y F(x, y)}$$
(4.15)

where  $\vec{R}_{\perp}(x, y)$  is the projected radius, in pkpc, from the flux-weighted centroid of the nebula to the point (x, y), F(x, y) is the flux at that point, and  $\vec{v}_z(x, y)$  is the line-of-sight velocity, in km s<sup>-1</sup>.

To determine whether a given measured value of  $\langle \vec{j} \rangle_f$  is significant, we estimate a 'minimum' value; i.e. that measured from two adjacent but spatially and kinematically distinct regions, A and B. The effective area of the seeing 'disk' in an individual exposure is approximately  $\theta_{slit} \times \theta_{seeing}$ , where  $\theta_{slit} \simeq 2''.5$  is the angular width of a slit and  $\theta_{seeing} \ 1 - 2''$  is the typical seeing (FWHM) at Palomar. As such, let us consider two adjacent  $3'' \times 3''$  regions ( $6 \times 6 \text{ px}^2$ ). Our typical error on the average velocity in a region of this size is  $\delta v_{reg} \sim 20 \text{ km s}^{-1}$ , taking covariance from binning into account (see Section 4.5). Let us consider the average velocities of region A and B to be  $-2\delta v_{reg}$  and  $+2\delta v_{reg}$ , respectively, such that they are kinematically separated. Finally, assuming a physical plate scale of  $\delta R_{\perp}/\delta x = 8 \text{ pkpc px}^{-1}$ , which is typical for our redshift range, we get  $|\langle \vec{j} \rangle_f|_{min} = 1783 \text{ pkpc km s}^{-1}$ .

## Flat vs. Sheared Model Comparison

We perform a comparison of two basic models using the BIC (see Section 4.5); a flat model,  $v(x, y) = v_0$ , and a model with linear terms in x and y,  $v(x, y) = v_0 + Ax + By$ .



Figure 4.25: Top: log likelihood that the flat model is more appropriate than a sheared model,  $\text{Log}_{10}(\text{p}_{\text{const}})$ . Bottom: average projected specific angular momentum,  $\langle \vec{j} \rangle_f$ , versus effective radius,  $\text{R}_{\text{eff}}$ . The size of the circular markers is proportional to  $-\text{Log}_{10}(\text{p}_{\text{const}})$  (i.e. the relative likelihood of the shear model). Each target is also shown as a red square of fixed size, for clarity. The horizontal line represents the minimum resolvable angular momentum discussed in the text.

This provides a qualitative test as to whether the moment map is flat or has any spatial dependence, to first order. As before, we use the BIC values to estimate the likelihood of each model.  $p_{flat}$  represents the likelihood that the flat model is more appropriate, while  $p_{shear} = 1 - p_{flat}$  indicates the likelihood that the shear model is more appropriate. For the majority of fields, the result is clearly in favor of the flat model ( $p_{flat} > 0.05 - 27/37$ ) or only weakly indicative of the sheared model ( $p_{flat} > 0.01 - 30/37$ ). A small number of targets indicate some significant likelihood that the shear model better represents the data ( $0.01 > p_{flat} > 10^{-5}$ , 5/37). However, for targets 4 and 7, there is a vanishing probability that the flat model is better ( $p_{flat} \sim 10^{-15}$  and  $p_{flat} \sim 10^{-21}$ ).

Figure 4.25 shows the detected nebulae in the parameter space of  $R_{eff}$  versus  $|\langle \vec{j} \rangle_f|$ . The size of the markers is indicative of the likelihood of a shear model being correct (size =  $-Log_{10}(p_{flat})$ ). From this combined perspective, it is clear that targets 4 and 7 represent two targets which are (i) among the largest detections, (ii) have significant projected specific angular momentum, and (iii) have strong indications from the BIC values that the velocity map is sheared rather than flat. We thus conclude that there is strong evidence of coherent kinematics in these two fields, though we leave the physical interpretation and modeling of this effect as a topic for the deep component of the FLASHES Survey.

## 4.8 Conclusions

We have conducted the first large IFS survey targeting the z = 2.3 - 3.1 CGM in emission. We observed 48 quasar fields over a four-year period using PCWI on the Hale 5m telescope at Palomar Observatory. We find that:

- I Of the 48 quasars observed (to an average limiting surface brightness of  $\sim 6 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$  in a 1" aperture), 37 exhibit extended Ly $\alpha$  emission on a wide range of scales, varying in flux-weighted radius over  $R_{rms} = 12 59$ kpc and in maximum (radial) extent over  $R_{max} = 19 120$  pkpc. The average flux-weighted projected radius of the nebulae is  $R_{rms}^{avg} = 22$  pkpc and the spread in the distribution of these sizes is  $\sigma(R_{rms}) = 16$  pkpc. The reported sizes are smaller than those in A19 or B16 by about  $\Delta R_{max} \sim 30$  pkpc, and comparable to those in C19.
- II The average radial profiles peak at  $SB_{max}^{obs} = 6 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ in observed surface brightness and  $SB_{max}^{adj} = 1 \times 10^{-15} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ when adjusted for cosmological surface brightness dimming.
- III The integrated nebular luminosities range from  $L_{min} = 0.4 \times 10^{43} \text{ erg s}^{-1}$  to  $L_{max} = 9.4 \times 10^{43} \text{ erg s}^{-1}$ .

- IV The nebulae are asymmetric on average, with measured eccentricities ranging from e = 0.51 to  $e \sim 1.0$ , and a sample-wide mean eccentricity of  $e_{avg} = 0.82$ . We find that the nebulae have a slightly higher eccentricity on average than those found by A19 around  $z \gtrsim 3$  quasars, but the same mean value as those reported around  $z \simeq 2.1 - 2.3$  QSOS by C19.
- V The S/N  $\ge 2\sigma$  covering factor profiles peak at  $f_c \simeq 30\%$  at small radii for the sample-wide average when non-detections are included and ~ 40% when they are excluded.
- VI The flux-weighted average velocity of the nebulae varies by thousands of km s<sup>-1</sup> with respect to the systemic QSO redshift ( $\sigma(\Delta v_{QSO}) = 994 \text{ km s}^{-1}$ ) and has a redshifted bias ( $\Delta v_{QSO,med} = +871 \text{ km s}^{-1}$ ). The flux-weighted average velocity of the nebulae also varies significantly with respect to the Ly $\alpha$  peak of the QSO spectrum, albeit by a smaller but considerable amount ( $\sigma(\Delta v_{peak}) = 606 \text{ km s}^{-1}$ ) and has a lesser but still redshifted bias with a value of  $\Delta v_{peak,med} = +390 \text{ km s}^{-1}$ .
- VII Most of the integrated nebular emission line profiles are either single-peaked (15/37) or double-peaked (17/37) with a few nebulae exhibiting more complex line shapes.
- VIII Global dispersions for the nebulae range from 143 708 km/s, with a mean of 399 km s<sup>-1</sup> and standard deviation of 155 km/s. The average RMS line-of-sight velocity is found to be  $v_{RMS,avg} = 208 \pm 128$  km s<sup>-1</sup>, consistent with that expected from QSO host halos with a mass range of  $Log_{10}(M_h[M_{\odot}]) = 12.2^{+0.7}_{-1.2}$ .

## 4.9 Acknowledgements

This work was supported by the National Science Foundation (NSF Award Number 1716907). The authors would like to thank Chuck Steidel (California Institute of Technology) and Ryan Trainor (Franklin & Marshall College) for their insightful discussions and help with target selection, as well as Sebastiano Cantalupo (ETH Zürich) and Heather Knutson (Caltech). We would also like to thank the staff at Palomar Observatory for their continuous support over nearly seventy nights of observations since 2014. Finally, we would like to thank our reviewer for their time, effort and thoughtful feedback.

## References

- Alam, S. et al. (July 2015). "The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III." In: *ApJS* 219.1, 12, p. 12. DOI: 10.1088/0067-0049/219/1/12. arXiv: 1501.00963 [astro-ph.IM].
- Arrigoni Battaia, F. et al. (Sept. 2016). "The Stacked Ly $\alpha$  Emission Profile from the Circum-Galactic Medium of  $z \sim 2$  Quasars." In: ApJ 829, 3, p. 3. DOI: 10.3847/0004-637X/829/1/3. arXiv: 1604.02942.
- (Jan. 2018). "Inspiraling Halo Accretion Mapped in Ly  $\alpha$  Emission around a  $z \sim 3$  Quasar." In: *MNRAS* 473.3, pp. 3907–3940. DOI: 10.1093/mnras/stx2465. arXiv: 1709.08228 [astro-ph.GA].
- (Nov. 2019a). "Discovery of Intergalactic Bridges Connecting Two Faint  $z \sim 3$  Quasars." In: A&A 631, A18, A18. DOI: 10.1051/0004-6361/201936211. arXiv: 1909.00829 [astro-ph.GA].
- (Jan. 2019b). "QSO MUSEUM I: a Sample of 61 Extended Ly  $\alpha$  Emission Nebulae Surrounding  $z \sim 3$  Quasars." In: *MNRAS* 482.3, pp. 3162–3205. DOI: 10.1093/mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].
- Bertone, S., A. Aguirre, and J. Schaye (May 2013). "How the Diffuse Universe Cools." In: *MNRAS* 430, pp. 3292–3313. DOI: 10.1093/mnras/stt131. arXiv: 1301.5330.
- Bertone, S. and J. Schaye (Jan. 2012). "Rest-frame Ultraviolet Line Emission From The Intergalactic Medium at  $2 \le z \le 5$ ." In: *MNRAS* 419, pp. 780–798. DOI: 10. 1111/j.1365-2966.2011.19742.x. arXiv: 1008.1791.
- Bianchi, L. et al. (Sept. 2011). "GALEX Catalogs Of Uv Sources: Statistical Properties And Sample Science Applications: Hot White Dwarfs In The Milky Way." In: *Ap&SS* 335.1, pp. 161–169. DOI: 10.1007/s10509-010-0581-x.
- Blumenthal, G. R. et al. (Oct. 1984). "Formation Of Galaxies And Large-scale Structure With Cold Dark Matter." In: *Nature* 311, pp. 517–525. doi: 10.1038/311517a0.
- Bond, J. R., L. Kofman, and D. Pogosyan (May 1996). "How Filaments Of Galaxies Are Woven Into The Cosmic Web." In: *Nature* 380.6575, pp. 603–606. DOI: 10.1038/380603a0. arXiv: astro-ph/9512141 [astro-ph].
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Bridge, C. R. et al. (June 2013). "A New Population of High-z, Dusty Lyα Emitters and Blobs Discovered by WISE: Feedback Caught in the Act?" In: *ApJ* 769, 91, p. 91. DOI: 10.1088/0004-637X/769/2/91. arXiv: 1205.4030.

- Cai, Z. et al. (July 2018). "Keck/Palomar Cosmic Web Imagers Reveal an Enormous Ly $\alpha$  Nebula in an Extremely Overdense Quasi-stellar Object Pair Field at z = 2.45." In: *ApJ* 861.1, L3, p. L3. DOI: 10.3847/2041-8213/aacce6. arXiv: 1803.10781 [astro-ph.GA].
- (Sept. 2019). "Evolution of the Cool Gas in the Circumgalactic Medium (CGM) of Massive Halos A Keck Cosmic Web Imager (KCWI) Survey of Ly $\alpha$  Emission around QSOs at  $z \approx 2$ ." In: *arXiv e-prints*, arXiv:1909.11098, arXiv:1909.11098. arXiv: 1909.11098 [astro-ph.GA].
- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: Groundbased and Airborne Instrumentation for Astronomy V. Vol. 9147. Proc. SPIE, 91475K. DOI: 10.1117/12.2057056.
- Cantalupo, S. et al. (Mar. 2019). "The Large and Small Scale Properties of the Intergalactic Gas in the Slug Ly  $\alpha$  Nebula Revealed by MUSE HeII Emission Observations." In: *MNRAS* 483.4, pp. 5188–5204. DOI: 10.1093/mnras/sty3481. arXiv: 1811.11783 [astro-ph.GA].
- Christensen, L. et al. (Dec. 2006). "Extended Lyman-*α* Emission Around Bright Quasars." In: *A&A* 459, pp. 717–729. DOI: 10.1051/0004-6361:20065318. eprint: astro-ph/0603835.
- Correa, C. A. et al. (June 2015). "The Accretion History Of Dark Matter Haloes -II. The Connections With The Mass Power Spectrum And The Density Profile." In: *MNRAS* 450.2, pp. 1521–1537. DOI: 10.1093/mnras/stv697. arXiv: 1501. 04382 [astro-ph.CO].
- Cutri, R. M. et al. (Nov. 2013). "VizieR Online Data Catalog: AllWISE Data Release (Cutri+ 2013)." In: *VizieR Online Data Catalog* 2328.
- da Ângela, J. et al. (Jan. 2008). "The 2dF-SDSS LRG and QSO Survey: QSO Clustering and the L-z Degeneracy." In: *MNRAS* 383.2, pp. 565–580. DOI: 10.1111/ j.1365-2966.2007.12552.x. arXiv: astro-ph/0612401 [astro-ph].
- Dekel, A. et al. (Jan. 2009). "Cold Streams In Early Massive Hot Haloes As The Main Mode Of Galaxy Formation." In: *Nature* 457.7228, pp. 451–454. DOI: 10.1038/nature07648. arXiv: 0808.0553 [astro-ph].
- Fukugita, M., C. J. Hogan, and P. J. E. Peebles (Aug. 1998). "The Cosmic Baryon Budget." In: *ApJ* 503.2, pp. 518–530. DOI: 10.1086/306025. arXiv: astro-ph/9712020 [astro-ph].
- Fumagalli, M., J. M. O'Meara, and J. X. Prochaska (Dec. 2011). "Detection of Pristine Gas Two Billion Years After the Big Bang." In: *Science* 334.6060, p. 1245. DOI: 10.1126/science.1213581. arXiv: 1111.2334 [astro-ph.CO].
- Gronke, M., P. Bull, and M. Dijkstra (Oct. 2015). "A Systematic Study of Lymanα Transfer through Outflowing Shells: Model Parameter Estimation." In: *ApJ* 812.2, 123, p. 123. DOI: 10.1088/0004-637X/812/2/123. arXiv: 1506.03836 [astro-ph.GA].

- Helfand, D. J., R. L. White, and R. H. Becker (Mar. 2015). "The Last of FIRST: The Final Catalog and Source Identifications." In: *ApJ* 801.1, 26, p. 26. DOI: 10.1088/0004-637X/801/1/26. arXiv: 1501.01555 [astro-ph.GA].
- Herenz, E. C. et al. (May 2015). "Where is the Fuzz? Undetected Lyman α Nebulae Around Quasars at z ~ 2.3." In: A&A 576, A115, A115. DOI: 10.1051/0004-6361/201425580. arXiv: 1502.05132.
- Husemann, B. et al. (Jan. 2013). "CALIFA, the Calar Alto Legacy Integral Field Area survey. II. First public data release." In: A&A 549, A87, A87. DOI: 10. 1051/0004-6361/201220582. arXiv: 1210.8150 [astro-ph.CO].
- Joye, W. A. and E. Mandel (2003). "New Features of SAOImage DS9." In: Astronomical Data Analysis Software and Systems XII. Ed. by H. E. Payne, R. I. Jedrzejewski, and R. N. Hook. Vol. 295. Astronomical Society of the Pacific Conference Series, p. 489.
- Kellermann, K. I. et al. (Oct. 1989). "VLA Observations of Objects in the Palomar Bright Quasar Survey." In: *AJ* 98, p. 1195. DOI: 10.1086/115207.
- Kereš, D. et al. (Oct. 2005). "How Do Galaxies Get Their Gas?" In: *MNRAS* 363, pp. 2–28. doi: 10.1111/j.1365-2966.2005.09451.x. eprint: astro-ph/0407095.
- Law, D. et al. (Oct. 2016). "The Data Reduction Pipeline for the SDSS-IV MaNGA IFU Galaxy Survey." In: *AJ* 152.4, 83, p. 83. DOI: 10.3847/0004-6256/152/ 4/83. arXiv: 1607.08619 [astro-ph.IM].
- Marques-Chaves, R. et al. (Sept. 2019). "Discovery of a Giant and Luminous Ly $\alpha$ , CIV, and HeII nebula at z = 3.326 with Extreme Emission Line Ratios." In: A&A 629, A23, A23. DOI: 10.1051/0004-6361/201936013. arXiv: 1907.08486 [astro-ph.GA].
- Martin, D. C. et al. (May 2014a). "Intergalactic Medium Emission Observations with the Cosmic Web Imager. I. The Circum-QSO Medium of QSO 1549+19, and Evidence for a Filamentary Gas Inflow." In: *ApJ* 786, 106, p. 106. DOI: 10.1088/0004-637X/786/2/106. arXiv: 1402.4816.
- (May 2014b). "Intergalactic Medium Emission Observations with the Cosmic Web Imager. II. Discovery of Extended, Kinematically Linked Emission around SSA22 Lyα Blob 2." In: *ApJ* 786, 107, p. 107. DOI: 10.1088/0004-637X/786/2/107. arXiv: 1402.4809.
- (Aug. 2015). "A Giant Protogalactic Disk linked to the Cosmic Web." In: *Nature* 524, pp. 192–195. DOI: 10.1038/nature14616.
- (June 2016). "A Newly Forming Cold Flow Protogalactic Disk, a Signature of Cold Accretion from the Cosmic Web." In: *ApJ* 824, L5, p. L5. DOI: 10.3847/2041-8205/824/1/L5.

- Martin, D. C. et al. (July 2019). "Multi-filament Gas Inflows Fuelling Young Starforming Galaxies." In: *Nat. Astron.* 3, pp. 822–831. DOI: 10.1038/s41550-019-0791-2. arXiv: 1904.11465.
- Matuszewski, M. et al. (July 2010). "The Cosmic Web Imager: An Integral Field Spectrograph For The Hale Telescope At Palomar Observatory: Instrument Design And First Results." In: *Ground-based and Airborne Instrumentation for Astronomy III*. Vol. 7735. Proc. SPIE, 77350P. DOI: 10.1117/12.856644.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- Munari, E. et al. (Apr. 2013). "The Relation Between Velocity Dispersion And Mass In Simulated Clusters Of Galaxies: Dependence On The Tracer And The Baryonic Physics." In: *MNRAS* 430.4, pp. 2638–2649. DOI: 10.1093/mnras/stt049. arXiv: 1301.1682 [astro-ph.C0].
- Navarro, J., C. Frenk, and White (Dec. 1997). "A Universal Density Profile from Hierarchical Clustering." In: *ApJ* 490, pp. 493–508. DOI: 10.1086/304888. eprint: astro-ph/9611107.
- Storn, R. and K. Price (Dec. 1997). "Differential Evolution A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces". In: *Journal of Global Optimization* 11.4, pp. 341–359. ISSN: 1573-2916. DOI: 10.1023/A: 1008202821328. URL: https://doi.org/10.1023/A:1008202821328.
- Trainor, R. F. and C. C. Steidel (June 2012). "The Halo Masses and Galaxy Environments of Hyperluminous QSOs at z ~= 2.7 in the Keck Baryonic Structure Survey." In: *ApJ* 752, 39, p. 39. DOI: 10.1088/0004-637X/752/1/39. arXiv: 1204.3636.
- Umehata, H. et al. (Oct. 2019). "Gas Filaments Of The Cosmic Web Located Around Active Galaxies In A Protocluster." In: *Science* 366.6461, pp. 97–100. ISSN: 1095-9203. DOI: 10.1126/science.aaw5949. URL: http://dx.doi. org/10.1126/science.aaw5949.
- Virtanen, P. et al. (July 2019). "SciPy 1.0–Fundamental Algorithms for Scientific Computing in Python." In: *arXiv e-prints*, arXiv:1907.10121, arXiv:1907.10121. arXiv: 1907.10121 [cs.MS].
- Wagenmakers, E. and S. Farrell (2004). "AIC Model Selection Using Akaike Weights." In: *Psychon. Bull. Rev.* 11.1, pp. 192–196.
- White, M. et al. (Aug. 2012). "The Clustering Of Intermediate-redshift Quasars As Measured By The Baryon Oscillation Spectroscopic Survey." In: *MNRAS* 424.2, pp. 933–950. DOI: 10.1111/j.1365-2966.2012.21251.x. arXiv: 1203.5306 [astro-ph.CO].
- Wisotzki, L. et al. (Mar. 2016). "Extended Lyman  $\alpha$  Haloes Around Individual High-redshift Galaxies Revealed By MUSE." In: A&A 587, A98, A98. DOI: 10. 1051/0004-6361/201527384. arXiv: 1509.05143.

## Chapter 5

# FLASHES II: THE MULTIPHASE CGM AT z = 2.3 - 3.0

#### 5.1 Introduction

Each galaxy we observe today is the result of billions of years of - among other things - gas accretion, feedback, and mergers. Understanding the baryonic environments of galaxies in the early universe is thus of central importance to developing any holistic theory of galaxy formation. The development of several highly sensitive integral field spectrographs on 5-10m class telescopes ((Matuszewski et al., 2010; Caillier et al., 2014; Morrissey et al., 2018) has enabled surveys directly mapping emission from the circumgalactic medium at high redshift ( $z \ge 2$ ) (Borisova et al., 2016; Cai et al., 2019; Wisotzki et al., 2016; Arrigoni Battaia et al., 2018). Among these was the FLASHES Pilot Survey (O'Sullivan et al., 2020), which presented integral field spectroscopy observations of emission from the circumgalactic medium (CGM) around 48 z = 2.3 - 3.0 QSOs. Here, we present the first follow-up study, using the newly installed Keck Cosmic Web Imager (KCWI - Morrissey et al., 2018) to reveal the multi-phase CGM around as subset of these targets.

From the several large studies that have emerged targeting HI Lyman- $\alpha$  (henceforth  $Ly\alpha$ ), a number of details about the high-redshift CGM now seem clear. First, that large, extended Ly $\alpha$  nebulae with radial extents on the order of tens of proper kiloparsecs, and radial profiles approximately - though not exactly - following an exponential decline, are near ubiquitous. Second, that the morphology of these nebulae is often strongly asymmetric and eccentric. Third, that the nebulae tend to peak in observed surface brightness on the order of  $10^{-17}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>, with typical integrated luminosities on the order of  $Log(L [erg s^{-1}]) \sim 42 - 43$ . Finally, while the kinematics of these  $Ly\alpha$  nebulae appear mostly chaotic, they occasionally do exhibit some kind of coherent structure on large scales and, among these, some appear to exhibit signs of rotation. It is important to note that the necessary reliance on Ly $\alpha$  as a bright tracer of this gas has also left some gaps in our knowledge. What is the typical composition of this gas, or at the very least what are the relative emissive levels of various atomic species in the CGM? Do the morphologies and kinematics of other atomic species agree with those of the giant  $Ly\alpha$  nebulae? What is the CGM gas mass? How much of the gas we are seeing

is inflowing versus outflowing? In order to address these and related questions, we need to push beyond  $Ly\alpha$  and measure the various atomic species of the multi-phase CGM.

In this chapter, we present preliminary results from the multi-phase follow-up observations of the CGM around 8 QSOs selected from the FLASHES Pilot sample and covering rest-frame UV emission lines from Ly $\alpha \lambda 1216$  through to HeII  $\lambda 1640$ . In Section 5.2 we present a summary of the observations. In Section 5.3, we present the data reduction and analysis performed to measure the extended emission. In Sections 5.4 and 5.5, we present and then summarize the current observational results.

# 5.2 Observations

Targets were selected from the FLASHES pilot sample for observations with the Keck Cosmic Web Imager (KCWI) (Morrissey et al., 2018). The blue low-resolution (BL) grating was used so that the rest-frame range of  $1216\text{\AA} - 1640\text{\AA}$  could be covered in a single exposure. The targeted exposure time for each field was between 2 - 4 hours, as permitted by weather and target visibility. As result of poor weather, one target (ID 5) fell short of the 2h objective. Another target (ID 5) was allowed some extra time to fill a small gap in the observing schedule, resulting in a total of 5.4h. Observations were conducted in September 2017, February and December 2018, and May 2019.

Table 5.1 summarizes the observations, including the size and luminosity of Ly $\alpha$  emission detected in the pilot survey. Rather than following up only on the brightest Ly $\alpha$  targets, 5/8 of the targets were selected based on the merit of a bright detection in the pilot survey and 3/8 were selected because they represented non-detections of Ly $\alpha$  in the pilot survey. The motivation here was to explore the relationship between metal-line emission and Ly $\alpha$  emission in both extremes.

## 5.3 Data Reduction and Analysis

Bias correction, cosmic-ray removal, flat-fielding, cube reconstruction, differential atmospheric refraction (DAR) correction, and flux calibration were performed with the standard KCWI Data Reduction Pipeline (DRP). Further corrections and coadding were performed using CWITools (see Chapter 3 or O'Sullivan et al., 2020). Specifically, the world-coordinate system (WCS) was corrected by measuring the position of the known QSOs and updating the FITS header to align that position with the QSO's right-ascension and declination. The individual exposure cubes

ID	Coordinates	zqso	$L_{Ly\alpha}$	R <sub>eff</sub>	Slicer	$t_{exp}$
			$10^{43} \text{ erg s}^{-1}$	[pkpc]		[h]
1	17:01:01.00 +64:12:09.10	2.737	4.3	41.4	Μ	2.6
2	11:12:52.45 +15:21:23.50	2.790	9.4	54.8	Μ	5.4
5	07:35:35.44 +37:44:50.42	2.751	3.4	30.2	M	1.6
7	09:58:45.42 +47:03:24.43	2.491	6.1	42.5	Μ	2
9	08:37:12.89 +14:59:17.38	2.510	3.9	33.9	L	4.3
39	22:34:53.07 +26:37:25.00	2.777	-	-	Μ	2.5
42	15:52:00.50 +17:57:22.70	2.580	-	-	Μ	2
43	02:05:27.51 +19:02:29.10	2.703	-	-	M	4

Table 5.1: Summary of the observational data for ten FLASHES targets. For consistency with O'Sullivan et al., 2020, the IDs in the left-most column are the original target IDs from the FLASHES pilot survey. The effective radius ( $R_{eff} \equiv \sqrt{Area/\pi}$ ) size and luminosity of Ly $\alpha$  measured in the pilot survey are presented in the fourth and fifth columns. The column labelled 'Slicer' indicates the image slicer (M=Medium, L=Large) used for the observations. All observations used the low-resolution grating (BL).

were then coadded (weighted by exposure time) using a drizzle factor of  $d_F = 1.0$  (i.e. no drizzling) for medium slicer data and  $d_f = 0.8$ .

PSF subtraction was performed using CWITools. To summarize the method briefly here, an empirical PSF model for each wavelength layer is created by summing over a large bandpass on either side of that wavelength layer and then scaling the resulting image using pixels within  $r \sim 1''$ . After PSF-subtraction, background subtraction was performed with a polynomial model of order k = 5. Finally, the core of the PSF was masked.

Initial exploration of the data was performed by summing over the circular region within a projected radius of  $R_{\perp} = 50$  pkpc. For future reference, we label this spectrum  $F_{\lambda}^{R50} \equiv F(\lambda, R_{\perp} \le 50 \text{ pkpc})$  [erg s<sup>-1</sup> cm<sup>-2</sup> Å<sup>-1</sup>]. The spectrum was used to correct or verify the redshift of any identifiable CGM Ly $\alpha$  emission,  $z_{Ly\alpha}$ , as well as any metal-line emission (which should more accurately reflect the systemic QSO redshift),  $z_{QSO}$ .

With the redshifts identified, pseudo-narrowband images with velocity widths of  $\Delta v = \pm 750 \text{ km s}^{-1}$  were formed around the expected wavelengths for Ly $\alpha$  (1216 Å), NV (1240 Å), CIV (1549 Å), and HeII (1640 Å). These images, and the associates portions of the  $F_{\lambda}^{R50}$  spectra are shown in Figures 5.3 and 5.4. These narrow-band images are then used to obtain radial surface brightness profiles and luminosities of each line, which provide the baseline results and upper limits in the absence of a



Figure 5.1: Variance and covariance calibration for FLASHES deep data. Left: the initial, global scaling of the variance cube. The grey shaded region shows the histogram of SNR values in the background regions of the cube. The green region shows the subset of values not associated with large systematics. The dashed black line shows a normal distribution with a standard deviation of  $\sigma = 1$ , while the solid black line (overlapping with the dashed line) shows the best-fit model to the green data. Right: the ratio between measured and propagated noise after binning with box kernels of area K. The histogram on the right shows the error on the model, which falls comfortably within a  $\pm 20\%$  margin.

successful voxel-by-voxel 3D detection.

Voxel-by-voxel detection is performed by applying an initial SNR threshold of SNR  $\geq 2.5$  to the PSF-subtracted, background-subtracted data cubes to create 3D binary masks. The integrated SNR of each contiguous 3D region in these masks is then calculated and a second threshold of SNR<sub>int</sub>  $\geq 5$  is applied. The central wavelengths of the remaining 3D objects are then cross-checked with the expected wavelengths for the redshifted Ly $\alpha$ , CIV, NV, and HeII lines, with candidates being considered within a velocity window of  $\delta v = \pm 2000 \text{ km s}^{-1}$ .

If a detection is made using the 3D method, velocity and dispersion maps are generated for the object. If the spatial extent of an object is greater than  $R_{\perp} = 50$  pkpc, the integrated luminosity is measured by summing all of the flux under the 3D mask. Otherwise, the integrated luminosity is measured by summing over the inner 50 pkpc in the pseudo-narrowband images.

To perform this 3D segmentation and to obtain uncertainties on measured properties, accurate estimates of the variance are needed. An initial estimate of the variance on the final data cubes is obtained by dividing each cube into bins along the z-axis,
creating local 2D variance maps for each bin, and then scaling the local 2D map to match the global noise properties of each wavelength layer. The variance estimate is scaled globally to match the noise properties of the data.

Rescaling of the variance is done by calculating the distribution of SNR values for background voxels in the cube. As our observations are sky/background limited, the SNR in blank sky regions should follow a standard normal distribution. These background regions are identified by excluding regions of 50 or more connected voxels with  $|\text{SNR}| \ge 3$ , which represent either emission regions or systematic residuals (e.g. from sky subtraction near bright sky lines). The left panel of Figure 5.1 shows this variance scaling for target 2. The initial estimate is typically very close to the true value, so the variance only needs to be scaled by a small factor 0.95 < f < 1.05.

To calculate the integrated SNR of objects, as well as the variance on summed quantities, the effect of covariance also needs to be calibrated. This is done following a similar approach to O'Sullivan et al., 2020. The data is binned by increasingly large box kernels and, for each bin size, the ratio between the actual noise and ideally propagated (i.e. assuming no covariance) noise is measured. Then, a function is fit to describe the dependence of this ratio on kernel size. This function is of the form

$$\sigma_{\text{meas}}/\sigma_{\text{ideal}} = \eta (1 + \alpha \ln(K))$$
 for  $K \le T$  (5.1)

$$= \eta (1 + \alpha \ln(T)) \equiv \beta \qquad \text{for } K > T \qquad (5.2)$$

where  $\sigma_{\text{meas}}$  is the measured noise after binning,  $\sigma_{\text{ideal}}$  is the ideally propagated noise,  $\eta$  is a normalizing factor (to allow for inaccuracy in the initial variance estimate),  $\alpha$ parameterizes the dependence on the kernel area, K, and T is a threshold in kernel area above which the model becomes flat. This threshold reflects the maximum size of covariant structures in the data. The right-hand side of Figure 5.1 shows the covariance curve for target 2.

### 5.4 Preliminary Results

Figures 5.3 and 5.4 show the pseudo-narrowband images and spectra obtained for each of the targets. Figure 5.2 and Table 5.2 present the measured luminosities and  $2\sigma$  upper limits for extended Ly $\alpha$ , NV, CIV, and HeII in each of the eight target fields. Extended Ly $\alpha$ , NV, CIV, and HeII emission is detected in 3/8 targets. Two targets have just Ly $\alpha$  and CIV detections. Two targets have just Ly $\alpha$  detections. One target (ID 39) remains a non-detection for all lines, including Ly $\alpha$ .

Table 5.3 presents the ratios and upper limits on the ratios between the luminosities

ID	$L_{42}(Ly\alpha)$	$L_{42}(NV)$	$L_{42}(CIV)$	L <sub>42</sub> (HeII)
7	$129.9 \pm 1.7$	$2.6 \pm 0.5$	$14.5 \pm 0.6$	$7.1 \pm 0.5$
2	$152.1 \pm 1.2$	$4.4 \pm 1.1$	$23.3 \pm 1.0$	$8.0 \pm 1.9$
1	$37.6 \pm 3.0$	< 5.7	< 5.8	< 6.1
5	$72.8 \pm 1.1$	$4.5 \pm 1.1$	$7.4 \pm 1.1$	$9.3 \pm 1.2$
9	$50.5 \pm 0.3$	< 0.6	$2.0 \pm 0.4$	
39	< 1.4	< 1.5	< 1.7	< 1.9
43	$19.7 \pm 0.9$	< 1.2	< 1.3	< 1.6
42	$39.7 \pm 0.7$	< 1.3	$8.3 \pm 0.6$	< 1.5

Table 5.2: Integrated luminosities and upper limits for Ly $\alpha$ , CIV, NV, and HeII emission around the 8 deep targets. Values are given as  $L_{42} \equiv (L/10^{43} [\text{erg s}^{-1}])$ . For non-detections, upper limits  $(2\sigma)$  are derived from the propagated uncertainty on the luminosity measurement.

ID	$NV/Ly\alpha$	$CIV/Ly\alpha$	HeII/Ly $\alpha$
7	$0.020 \pm 0.004$	$0.112 \pm 0.005$	$0.055 \pm 0.004$
2	$0.029 \pm 0.007$	$0.153 \pm 0.007$	$0.053 \pm 0.013$
1	< 0.151	< 0.154	< 0.162
5	$0.062 \pm 0.015$	$0.101 \pm 0.016$	$0.127 \pm 0.016$
9	< 0.013	$0.040 \pm 0.007$	
43	< 0.063	< 0.067	< 0.084
42	< 0.032	$0.208 \pm 0.017$	< 0.039

Table 5.3: Ratios of CIV, NV, and HeII to  $Ly\alpha$  in terms of integrated luminosity. Upper limits  $(2\sigma)$  are provided where there is a  $Ly\alpha$  detection but no detection of the other line. No ratios or upper limits are provided for target 39, which has no  $Ly\alpha$  detection.

of the Ly $\alpha$  emission and the three targeted metal emission lines. Having no detection in Ly $\alpha$ , target 39 is omitted from these results. The ratios between NV and Ly $\alpha$ range from NV/Ly $\alpha < 0.013$  to NV/Ly $\alpha \simeq 0.06$ , the ratios between CIV and Ly $\alpha$ range from CIV/Ly $\alpha \simeq 0.04 - 0.2$ , and the ratios between CIV and Ly $\alpha$  range from HeII/Ly $\alpha < 0.039$  to HeII/Ly $\alpha \simeq 0.127$ .

Table 5.4 presents two measurements of size for each 3D detection: the fluxweighted root-mean-square distance from the QSO,  $R_{rms}$  and the maximum radial extent of the nebulae,  $R_{max}$ . Ly $\alpha$  emission is measured with maximum extents spanning  $R_{max} \simeq 33 - 127$  pkpc and flux-weighted radii spanning  $R_{max} \simeq 16 - 45$  pkpc. For target 7, the metal emission is concentrated on the same spatial scale as the Ly $\alpha$  emission, with RMS radii of  $R_{rms} \simeq 21 - 22$  pkpc for NV, CIV, and HeII and  $R_{rms} \simeq 25$  pkpc for Ly $\alpha$ . The maximum extent of the detected metal line emission,

ID	$R_{max}^{Ly\alpha}$	$R_{rms}^{Ly\alpha}$	R <sub>max</sub> <sup>NV</sup>	R <sub>rms</sub> <sup>NV</sup>	R <sub>max</sub> <sup>CIV</sup>	R <sup>CIV</sup> <sub>rms</sub>	R <sup>HeII</sup> <sub>max</sub>	R <sup>HeII</sup> rms
7	65.73	24.72	35.79	21.13	39.8	21.99	46.16	21.84
2	90.63	23.5	23.91	15.65	35.64	16.25	23.54	15.35
1	54.02	28.35						
5	37.76	17.09					19.81	14.17
9	126.9	45.07						
43	32.75	17.57			16.21	11.23		
42	33.35	16.47			17.46	12.96		

Table 5.4: Sizes of 3D detections for each target and line, in units of pkpc.  $R_{max}$  and  $R_{rms}$  are the maximum and flux-weighted root-mean-square radial extent of the nebular emission with respect to the QSO.

however, is approximately 50 - 70% that of Ly $\alpha$ . Qualitatively, the same result (similar R<sub>rms</sub> and smaller R<sub>max</sub>) applies to targets 5, 42, and 43.

Figure 5.5 presents the circularly averaged surface brightness profiles for each line and each target. The Ly $\alpha$  profiles peak for the brightest nebulae at an average brightness of a few ×10<sup>-17</sup> erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup>, while the non-detection in target 39 has upper limits at ~ 1 × 10<sup>-18</sup> erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup>. The radial profiles for NV, CIV, and HeII all appear fainter than their Ly $\alpha$  counterparts by an order of magnitude or more.

For the two targets with 3D detections in all four lines, first and second moment maps were generated following the same methodology described in Chapter 3. Figure 5.3 and Figure 5.4 show the first z-moment (velocity) and second z-moment (dispersion), respectively. Although a more detailed analysis has not yet been performed, the CIV and HeII velocity maps appear broadly consistent with the Ly $\alpha$  maps. The second moment maps for NV show dispersions of  $\sigma_v^{NV} \ge 400 \text{ km s}^{-1}$ , though this is likely a result of the double-peaked line shape rather than the intrinsic line width. The dispersions for HeII show values  $\sigma_v^{\text{HeII}} \le 100 \text{ km s}^{-1}$ . However, this is narrower than a resolution element of the observations (which have R = 900), and as such is a systematic effect caused by the artificially narrow velocity width of the detected object (due to SNR thresholding).

### 5.5 Summary and Ongoing Work

So far, the first deep study of the FLASHES survey has yielded the following results:

1. Of 8 targeted QSO fields (5 pilot survey detections and 3 non-detections), 7/8 contain Ly $\alpha$  emission. Three fields contain extended NV, CIV, and HeII emis-

sion in addition to extended Ly $\alpha$  nebulae. Another two contain detected CIV and Ly $\alpha$ , but no NV or HeII. Ly $\alpha$  alone is detected in two more and one field contains no detections whatsoever. If the remaining non-detection contains an extended Ly $\alpha$  nebula, its luminosity is L(Ly $\alpha$ ) < 1.4 × 10<sup>42</sup> erg s<sup>-1</sup>.

- 2. The detected luminosities in NV range from  $L(NV) = 2.6 4.5 \times 10^{42} \text{ erg s}^{-1}$ , while upper limits of  $L(NV) \leq 1 \times 10^{42} \text{ erg s}^{-1}$  are placed on most other fields. The range of measured CIV luminosities is  $L(CIV) = 2.0 - 23.3 \times 10^{42} \text{ erg s}^{-1}$ , with upper limits of  $L(CIV) \leq 1.5 \times 10^{42} \text{ erg s}^{-1}$  placed on other fields. The range of measured HeII luminosities is  $L(HeII) = 7.1 - 9.3 \times 10^{42} \text{ erg s}^{-1}$ , with upper limits of  $L(HeII) \leq 1 - 2 \times 10^{42} \text{ erg s}^{-1}$  placed on other fields.
- 3. Kinematic maps are obtained for all four lines for two targets. The velocity maps show broad agreement across the lines, though it is difficult to compare for the small detections in NV and HeII. Velocities are found to be within a few hundred km s<sup>-1</sup> for all maps. Dispersions for LyA and CIV are mostly in the range of  $\sigma_v = 200 400$  km s<sup>-1</sup>. The NV and HeII velocity maps are affected systematics and more modeling work is needed to measure these dispersions carefully.

At the time of writing this thesis, this deep study is a work in progress. A number of tasks remain to finalize the data analysis and enable a more detailed discussion of the results:

- 1. Perform a final round of analysis attempting to improve the PSF and background subtraction, to ensure no detections have been missed due to systematic errors.
- 2. Use adaptive kernel smoothing to optimally extract signal in all cubes.
- 3. Leverage the HeII and CIV observations to constrain the gas mass following Hennawi and Prochaska, 2013.
- 4. Adopt a modeling approach to estimate metallicity and gas temperatures from the observations.

### References

Arrigoni Battaia, F. et al. (Jan. 2018). "Inspiraling Halo Accretion Mapped in Ly  $\alpha$  Emission around a  $z \sim 3$  Quasar." In: *MNRAS* 473.3, pp. 3907–3940. doi: 10.1093/mnras/stx2465. arXiv: 1709.08228 [astro-ph.GA].

- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Cai, Z. et al. (Sept. 2019). "Evolution of the Cool Gas in the Circumgalactic Medium (CGM) of Massive Halos A Keck Cosmic Web Imager (KCWI) Survey of Ly $\alpha$  Emission around QSOs at  $z \approx 2$ ." In: *arXiv e-prints*, arXiv:1909.11098, arXiv:1909.11098. [astro-ph.GA].
- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: Groundbased and Airborne Instrumentation for Astronomy V. Vol. 9147. Proc. SPIE, 91475K. doi: 10.1117/12.2057056.
- Hennawi, J. F. and J. X. Prochaska (Mar. 2013). "Quasars Probing Quasars. IV. Joint Constraints on the Circumgalactic Medium from Absorption and Emission." In: *ApJ* 766.1, 58, p. 58. DOI: 10.1088/0004-637X/766/1/58. arXiv: 1303.2708 [astro-ph.CO].
- Matuszewski, M. et al. (July 2010). "The Cosmic Web Imager: An Integral Field Spectrograph For The Hale Telescope At Palomar Observatory: Instrument Design And First Results." In: *Ground-based and Airborne Instrumentation for Astronomy III.* Vol. 7735. Proc. SPIE, 77350P. DOI: 10.1117/12.856644.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 z ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.
- Wisotzki, L. et al. (Mar. 2016). "Extended Lyman  $\alpha$  Haloes Around Individual High-redshift Galaxies Revealed By MUSE." In: A&A 587, A98, A98. DOI: 10. 1051/0004-6361/201527384. arXiv: 1509.05143.



Figure 5.2: Luminosities and upper limits on luminosity for Ly $\alpha$ , NV, CIV, and HeII for the eight targets in the sample. Error bars shown as  $\pm 2\sigma$ .



Figure 5.3: Pseudo-narrowband images and  $F_{\lambda}^{R50}$  spectra for the first four deep targets (IDs 1, 2, 5 and 7). Each pseudo-narrowband is a  $100 \times 100$  pkpc<sup>2</sup> square.



Figure 5.4: Pseudo-narrowband images and  $F_{\lambda}^{R50}$  spectra for the last four deep targets (IDs 9, 39, 43 and 42). Each pseudo-narrowband is a  $100 \times 100$  pkpc<sup>2</sup> square.



Figure 5.5: Circularly averaged radial surface brightness profiles for each targeted emission line. Surface brightness units, labelled [cgs] to fit on the plot, are the standard erg s<sup>-1</sup>cm<sup>-2</sup>arcsec<sup>-2</sup>.



Figure 5.6: Velocity maps for the two targets with extended emission detected for all four emission lines.



Figure 5.7: Dispersion maps for the two targets with extended emission detected for all four emission lines.

# Chapter 6

# FIREBALL-2 - THE FAINT INTERGALACTIC REDSHIFTED EMISSION BALLOON



Figure 6.1: A view of FIREBall-2 during its 2018 flight, taken by photographer Mouser Williams.

# 6.1 Constraining the Low Redshift CGM

To understand the evolution of the CGM and IGM, as well as their role in galaxy formation and evolution, we need observations across a wide range of redshifts. The low redshift universe (i.e.  $z \leq 1$ ) is of particular interest when studying galaxy evolution. This might be in part for the obvious reason that the galaxies and environments we see there are more immediate predecessors to those we see today, but there are a number of other reasons to study the low-redshift CGM and IGM. We know from galaxy population studies that there is an evolution over cosmic time from blue, star-forming galaxies to 'red and dead' ellipticals. Much debate has prevailed over the dominant cause of this transition. To what extent is it caused by feedback from star formation or AGN? What are the effects of this feedback on environmental properties? How does the hierarchical growth of more and more

massive dark matter halos impact cold flow accretion? To answer these questions, we need to reveal the environments of low redshift galaxies, study their kinematics and composition, and cross-correlate these properties with the properties of the galaxies themselves. Such a cross-correlation is also harder to do at high-redshift, where the most commonly studied emission is boosted by QSO fluorescence, but the physical plate-scale in proper kiloparsecs per arcsecond makes separation of the QSO and host galaxy untenable. As such, studying the CGM at low redshift is of central importance to the development of any broad theory of galaxy formation and evolution.

However, while cosmological redshift at  $z \sim 2-3$  places many of the brightest emission lines in the diffuse universe (e.g. Ly $\alpha \lambda 1216$ , CIV  $\lambda 1549$ ) squarely in the optical portion of the electromagnetic spectrum, low-redshift studies must target these lines at ultra-violet wavelengths. This presents a significant instrumental challenge on two fronts. First and foremost, standard science-grade CCDs do not perform well at ultra-violet wavelengths. This is because UV emission does not penetrate far into the silicon wafer, so the electron-hole pair formed by a UV photon is formed near to the surface where the semi-conductor's electric field is weak, and the pair simply recombines rather than being picked up as a photo-electron. While there are alternative designs to CCDs, the quantum efficiency of UV detectors has historically been extremely low compared to their optical counterparts. GALEX, a highly successful UV mission which mapped two thirds of the sky at near- and far-UV wavelengths, did so with a quantum efficiency of about 5%. Understandably, this presents a problem for the scientific goal of studying some of the faintest UV emission in the universe. The second challenge involved in studying UV emission comes from the Earth's atmosphere, which absorbs almost all UV light. The only way to circumvent this is to get above the atmosphere. To do so, one needs a sounding rocket, high-altitude balloon, or satellite. Satellites are significantly more expensive than either of the other two options, and usually only hardware with a high technology readiness level, meaning either sounding rockets or balloons must be used to prove that the technology works first. Sounding rockets provide only brief windows of time at high altitude (i.e. on the order of minutes, rather than hours) which means that they are not well suited to a study that requires long exposures. Thus, to develop a UV detector for the study of the faint, diffuse universe, a highaltitude balloon is needed.

#### 6.2 The History and Architecture of FIREBall

FIREBall - the Faint Intergalactic Redshifted Emission Balloon - is a high-altitude balloon-born telescope with a UV integral field unit at its prime focus. While motivated by the goal of mapping the diffuse universe at low redshift, it also broadly serves as a unique test-bed for the next generation of UV observational capabilities. FIREBall has had two generations. FIREBall-1 flew in 2007 (Tuttle et al., 2010) and 2009 (Milliard et al., 2010), while FIREBall-2 flew in 2018 (Hamden et al., 2019) and has another launch scheduled for 2020. The first flight of each generation (i.e. 2007 and 2018) served as engineering demonstrations because of complications during launch. In 2007, a pivot mechanism failed during launch and reduced the telescope's ability to maintain stable pointing for exposures longer than a minute. Several bright stars were observed instead, so the flight still yielded a successful test of the instrument. In 2018, a tear in the balloon reduced the flight time to just forty minutes above minimum altitude. The best launch opportunity also happened to coincide with a full moon, and the deflated balloon reflected moonlight into the instrument via an unexpected path. Thus, the flight was only a fraction of the intended time and the background level was orders of magnitude higher. Despite these obstacles,  $Ly\alpha$  was still detected from one bright galaxy and the flight served to verify the performance of the entire payload. The 2007 and 2018 flights demonstrate the high risk that comes with the high gains of balloon-borne science. The entire system is assembled and tested in a time-span of a couple of months, the winds and weather in a 2-3 week launch window need to be ideal, and there are multiple critical risk areas (balloon, pivot, pointing system, spectrograph, detector, communications), most of which are novel and custom-built for one specific flight. The 2009 flight demonstrates the rewards of such immense effort - FIREBall-1 obtained the strongest constraints on the CGM and IGM available from any instrument at the time. Hopefully, the 2020 flight will do the same in providing the first deep, direct measurements of UV emission from the low redshift CGM and IGM.

# **Optical System**

The architectures of both generations of FIREBall have much in common. Figure 6.2 shows an overview of both missions and the differences between them. Both generations share the same basic structure. A 1.2 meter altitude-azimuth sidereostat reflects incoming light upwards onto a 1 meter parabolic primary mirror. The entrance to the spectrograph is placed at the prime focus, in addition to some optical



Figure 6.2: Schematics of the optical design of FIREBall-1 and FIREBall2. FIREBall-2 added a field corrector, changed the spectrograph design from a 'folded Offner spectrograph' to an aspherized Schmidt spectrograph, and changed the detector from a micro-channel plate to a delta-doped, electron-multiplying CCD.

element which redirects visible light into the guider camera. The spectrograph's internal optics are reflective, to avoid the heavy losses in the UV caused by traditional lens materials. Despite the shared overall structure, the designs of FIREBall-1 (FB-1) and FIREBall-2 (FB-2) are very distinct.

The integral field unit in FB-1 is fiber-based - with the focal plane consisting of a hexagonal bundle of optical fibers which project the input light onto the detector

in pseudo-slits which are then dispersed. A dichroic is placed before the entrance to the spectrograph which redirects visible light to the guider camera. While the dichroic has a transmission of  $\geq 90\%$  at  $\lambda \sim 2200$ Å, the optical fibers are a source of considerable losses in this design, at a transmission of 70%. In FB-2, the fiberbased entrance to the integral field unit is replaced with slit-mask design. A series of pre-made masks, designed for specific scientific targets, are installed on a carousel which can be used to interchange them at the focal plane. Each mask consists of two layers: an upper layer with a reflective surface directs visible light back and into the guider camera, while the lower layer contains the slit mask for the science targets.

Another major change from FB-1 to FB-2 is the detector being used. While FB-1 used a similar microchannel plane detector to the one used in GALEX (Martin et al., 2005), with a quantum efficiency of QE  $\simeq$  5%, FB-2 uses a 'delta-doped electronmultiplying CCD,' newly developed at the Jet Propulsion Laboratory in Pasadena, CA. While standard optical CCDs suffer losses due to recombination of the electronhole pairs formed near the surface of the detector (see Section 1), these CCDs have a Boron coating which is on the order of a single layer of atoms thick (hence the term 'delta-doped' - as in a delta function). This layer carries the electric field, strengthening it near the surface of the detector and leading to the capture of more photo-electrons. Where microchannel plates have  $QE_{mcp} \simeq 5\%$  in the UV, these  $\delta$ -doped EMCCDs have a UV quantum efficiency of QE<sub> $\delta d$ </sub>  $\simeq 50 - 70\%$ , an order of magnitude improvement. The CCDs also have an electron-multiplying mode, in which an extremely high gain is applied and a single event (i.e. photo-electron) cascades through the readout process to become tens of thousands of electrons. The gain is stochastic in nature, so this mode does not provide an accurate count of how many electrons per pixel were measured. Instead, a high threshold is applied (e.g.  $t > \gtrsim 10000e^{-}$ ) to generate a binary mask which represents the pixels which contained an electron before readout. As read noise only adds a few electrons per pixel during the readout process, this effectively eliminates read noise entirely, leaving only clock-induced-charge (CIC) - which is extremely low - as a source of false-positives. This 'photon-counting' mode only works well in the regime where the expected value of photons per pixel is approximately one within the set exposure time. As such, it might not be well suited to measuring bright stars, but it is extremely well suited to measuring the diffuse universe.



Figure 6.3: The broader communications system context for FIREBall-2, including the CSBF-owned downrange station and plane, used as backups in the case of loss of communications. Boxes show components of the system, while arrows show communications connections.

While the above two changes constitute the most significant upgrades from the optical design of FB-1 to FB2, there were a number of smaller but still significant changes. A field correcting relay of two spherical mirrors was added between the prime focus and the spectrograph. This addition corrected the field out to a  $30' \times 30'$  field of view, though only  $30'' \times 12''$  is used by the detector. The grating was also changed from a "folded Offner spectrograph" (Tuttle et al., 2010) to an aspherized Schmidt spectrograph (Hamden et al., 2019).

# **Communications System**

In addition to the telescope and instrument, balloon-borne missions require complex pointing controls systems and remote communication systems. The French National Center for Space Studies (CNES - Centre Nationale d'Etudes Spatiales) was responsible for the former, and its inner workings are beyond the scope of this thesis. Suffice to say that it is difficult to maintain arcsecond pointing while swinging from

170

a giant balloon at 128,000 ft. The communications system, on the other hand, is a topic of focus within this thesis. As such, it is worth spending some time on the architectural decisions related to the communication system.

The Columbia Scientific Balloon Facility (CSBF), located in Palestine, Texas and with a base in Fort Sumner, New Mexico, is the NASA facility which manages their overall scientific balloon program. CSBF handles the actual balloon deployment and launch, as well as the radio communication between the ground and payload during flight. As such, the communications system of any balloon-borne (NASA/CSBF) project consists of three main 'vertical' (i.e. from ground to balloon) layers: (i) the ground-side communications linking between control stations and the CSBF system, (ii) the CSBF-controlled radio uplink to their hardware on the payload, and (iii) the flight-side communications between the CSBF hardware and instrumentation or other devices. In addition to this breakdown, there are a few main subcategories of communication to be considered. The bulk of the communication system is the so-called telemetry: downlinked house-keeping data with status updates and feedback from flight-hardware; and tele-commanding: uplinked commands issued by team members at various control stations on the ground. For short-hand, the telemetry/tele-commanding system is called the TM/TC system. In addition to the TM/TC system, there is a separate downlink with a large bandwidth (1MBps) for downloading science data (i.e. FITS format images) and a video downlink for streaming the guider video to a monitor the team can watch during flight. Gillian Kyne, previously at Caltech and now at JPL, worked primarily on detector development but also led the development of the science downlink. Our collaborators from Columbia University, who built the guider system, oversaw the video downlink.

The most important architectural decision for the communications system is which CSBF flight-hardware will be used, as this serves as the bottle-neck through which all control of the telescope and instruments is funneled. There are essentially three choices here: the science instrument package (SIP), the 'mini'-SIP, and the consolidated instrument package (CIP), which is usually the choice for short-duration flights. The SIP and mini-SIP are the typical choices for long-duration balloons with flights of a few days or more. FIREBall-1 flew the mini-SIP, which provides a direct serial link to the payload over which commands can be sent. The payload of FIREBall-2 had been redesigned using the CIP, which is 1/3 the size and signifi-



Figure 6.4: Simple schematic of the different communication layers in the TM/TC system. The left-hand side shows the ground-side system while the right-hand side shows the flight-side system.

cantly lighter than the mini-SIP, but does not provide a serial link over which to issue commands. Instead, the CIP provides a functionality to upload two-byte datawords one at a time, which are represented on the flight-side as a parallel signal on a series of 16 pins, rather than as a serial output. Had we been able to design the system from scratch, the mini-SIP would have been the better choice on the whole, but as it was, the payload had already been designed with this decision made, and it could not be easily reverted due to the increased mass and footprint of the mini-SIP. While it allowed considerable savings in weight and space - two precious commodities in ballooning - this decision meant that all telescope and instrument commanding had to be broken into a series of two-byte datawords, each sent through the CSBF system and parsed from parallel back to serial before being recombined into the full command and forwarded to the flight devices.

There are a few more highly-connected and performance-sensitive decisions to be made once the flight hardware has been chosen. What control stations need to be connected, and how many connections does each need? What bandwidths should be allocated to each device? How should the commanding and telemetry on multiple channels be combined (i.e. what multiplexing process should be used)? Importantly, for the CIP-based architecture, how can the multiplexed commands be compressed to squeeze through an effective uplink bandwidth of  $\sim 60 - 70$  bits per second?

### 6.3 Telemetry and Telecommand System Overview

The full communications system for FIREBall-2 consists of three major components: (i) a data downlink for images taken by the instrument, (ii) a video downlink for the guider, and (iii) the telemetry and telecommand subsystem, abbreviated as TM/TC. The TM/TC subsystem encapsulates everything used for controlling and receiving feedback from the main flight-side devices. Figure 6.4 shows a simple schematic of the different layers involved in this process. In short, there are four layers, which have components on both the ground and flight side. First and foremost is the control stations and their respective devices on the flight side (e.g. the ground-station controlling the detector and the on-board detector computer). This 'layer' represents the sending and receiving ends of the TM/TC process. Next is the multiplexing layer, which combines the outgoing signals from the various stations and devices into a single communications channel and separates it back into multiple channels at the receiving end. Before being sent to the ground-side equipment for radio communication, this signal needs to be processed and formatted correctly. It also needs to be parsed after being received by the radio equipment on the flight-side. These processing steps - handling the interface with the radio equipment - are referred to as TC Processes because they mostly relate to the uplink (i.e. telecommand) direction. Finally, there is the radio communication step itself, which is fully controlled by the Columbia Scientific Balloon Facility.

The TM/TC system has four main external interfaces which constrain its design: (i) FIREBall-2 ground-side control stations, (ii) CSBF ground-side equipment (called the GSE), (iii) CSBF flight-side equipment (the consolidated instrument package or CIP), and (iv) FIREBall-2 flight-side equipment. The FIREBall-2 devices and operating stations determine the requirements of the system in terms of speed and stability, while the CSBF equipment determines some of the strongest design constraints. As such, it seems like a good idea to cover these two layers first to set up the design of the system.

### 6.4 CSBF Radio System and Interface

The radio uplink provided by CSBF and its interfaces provide the hardest constraints on the design of the communications system. This is true in part because we have no design control over the system, but also because the choice of which CSBF system to use is one of the key architectural decisions of the mission. CSBF offers a few alternative pieces of hardware, which vary significantly in the type of interface provided, available bandwidths, size, and weight. The decision was made early in the development of FIREBall-2 to use CSBF's consolidated instrument package (CIP), because it is about one third the size and weight of the feasible alternative, the mini-SIP (Support Instrument Package). However, the CIP comes with a significant bottle-neck on the uplink bandwidth, as we will discuss shortly.



Figure 6.5: The J2 Connector on the CIP. The columns are labelled from right to left with letters, while the rows are labelled with numbers. Every other column has only even or odd rows, such that column 'A' has pins A0, A2, A4 (etc.) and column B has pins B1, B3, B5.

For the TM/TC downlink, we utilized a voltage controlled oscillator transmitting on the standard Inter-Range Instrumentation Group channel HH with a frequency response of 100 kHz. CSBF recommended the use of biphase encoding, which uses two bits for every signal bit to achieve better stability, so the maximum signal frequency that could be carried on this downlink was 50,000 baud.

For the uplink, the CIP can receive two types of commands: discrete and data-word commands. A discrete command sends a pulse to a single pin on the CIP, which is useful as a control for power switches or reset buttons. A data-word command sends a two-byte packet of data to the CIP, which appears in parallel on a series of pins. All of this output appears on the flight-side J2 connector on the CIP, a picture of which is shown in Figure 6.5. The output pins on the J2 connector are open collector outputs with external pull-up resistors, such that they default to a signal of 1 (or HIGH) when there is no input and drop to 0 (LOW) when there is input on the pins. On the ground side, the GSE requires a serial input at 1200 baud for issuing commands. Commands are sent to the GSE by the ground-side FIREBall-2 equipment, where they are parsed and trigger radio uplinks to the CIP. Commands must have the following syntax:

# S[yyyy] [xx][q]S[yyyy] [xx][q]S[yyyy] [xx][q][CR][LF]

'S' is just the ascii character S (x53 - Hexadecimal value 53). The phrase '[yyyy]' is the four-character command, either the hex representation of the dataword to be

sent or the pin number to be toggled for a discrete command. The '[**xx**] term is the address to send to (an internal setting determined by CSBF and given to each team). '[**q**]' determines the type of command being sent: 'K' (x4B) for discrete commands and 'W' (x57) for datawords. '[**CR**]' (x0D) is the carriage return character and '[**LF**]'(x0A) is the line-feed character. So, for example, if our address was '07' and we wanted to send the character-string 'Hi,' we would first convert 'Hi' into its hexadecimal representation: H = x48 and i = x69. Then our command would be:

### S4869 07WS4869 07WS4869 07W[CR][LF],

To send a discrete command to pin E13 on the J2 connector, we would check the CIP manual<sup>1</sup> for the command which toggles this pin, find that it is Hexadecimal 22, and send:

# S0022 07KS0022 07KS0022 07K[CR][LF].

Although the connection to the GSE is at 1200 Bd, the true bottle-neck in the uplink process comes from the rate at which the GSE can (or should) issue radio commands. While the rate can be adjusted to some extent, the speed is approximately 0.25s per command. In terms of sending serial data as two-byte data-words, this corresponds to an effective speed of 64 Bd, before taking any overhead bytes into account. This provides the most limiting constraint on the uplink speed throughout the system. As we will discuss shortly, this slow uplink speed proved to be one of the major operational challenges of the 2018 campaign. While the design of FIREBall-2 may be too far advanced to make as central a change as replacing the CIP with a much larger and heavier alternative, it should be seriously considered for FIREBall-3, should there be one.

#### 6.5 Ground Control Stations and Flight Devices

Three main flight-side systems need to be controlled: the detector computer, the guider computer, and the CNES gondola-pointing equipment.

#### Guider System - GDR1 and GDR2

When the balloon is in flight and the telescope is pointed at one of the target fields, the guidance system and gondola-pointing system work together to keep the telescope accurately pointed. The guiding software has a flight-side component and

<sup>&</sup>lt;sup>1</sup>Available at www.csbf.nasa.gov

a ground-side component. The flight-side software on the guider on-board computer (GOBC) automatically reads in the guider images, tracks the positions of multiple guide stars, and sends relevant data to the ground station. The ground-side guider software consists of a graphical user interface which displays the telemetry data and enables the user to send commands to the flight-side adjusting certain parameters such as the exposure time for the guider camera.

The guider telemetry includes one-dimensional profiles of the guide stars and associated model parameters, as well as other general house-keeping data. Given that the telescope will essentially need to guide for the entire flight, this means that the guider channel requires a reliable and constant telemetry connection with a moderate bandwidth. Through testing different connection speeds, a rate of 9600 baud was found to be sufficient to handle this link without degradation of performance.

In the event of a critical failure, such as a power failure or software crash, a backup connection was needed to restart the flight-side software. The serial port used for this purpose needed to be modified to allow direct access to a terminal where the user could issue commands. Since the main guider connection could not be used for both purposes, a second serial connection was established. This connection could function at a very low bandwidth, since it was only required to enable infrequent and short text commands and required little to no telemetry. We determined that a connection of 4800 Bd was comfortably sufficient for the needs of this backup channel. For the rest of this chapter, we will refer to the main guider connection (which connects the ground and flight software) as GDR1 and the backup terminal connection as GDR2. In terms of hardware interfaces, both were USB inputs on the ground station computer while GDR1 was a DB9 connection and GDR2 a USB connection to the flange of the GOBC.

#### **CNES Gondola Pointing System - NPT**

The main mechanical controls for the gondola are owned and operated by the French contingent of the FIREBall team from the Centre Nationale d'Etudes Spatiales (CNES). Throughout this chapter, we will use the French acronym for this system, NPT (nacelle pointage télescope). The NPT system is responsible for stabilizing the telescope's pointing, which is no easy task for a five ton payload hanging from a balloon at 128,000 ft. Of the three major systems, NPT had both the most demanding communications requirements and the highest risk associated with failure. Telemetry data on this channel sends critical information about all of the

mechanical components on the gondola. This was also the channel with the highest flux of uplinked commands, as every change in pointing or target acquisition needed to be issued through it. The commands themselves contained high-precision target destinations for various mechanisms as well as checksums (to avoid commands being executed in error). The density of information in these commands made them quite long (relative to the uplink bandwidth - e.g. 50 bytes) and incompressible. Primarily because of the telemetry data, the NPT was allocated the maximum bandwidth possible on the HH channel (38, 400 baud).

Although the NPT channel demanded high throughput and reliability, its interfaces were relatively simple compared to other devices. Both the ground station and on-board system used a straight-forward serial DB9 connection.

# **Detector - DET**

The detector on-board computer (DOBC) is responsible for image acquisition, monitoring the temperature and pressure of the instrument, and managing the science data downlink. Communication is done through a serial connection which is configured to allow terminal access to the DOBC. Although it is of central importance to the scientific mission, the demands of the detector communications channel (labelled DET) are relatively low. Commands issued via the uplink are executed as terminal commands in a linux environment. As such, shorthand aliases for the most common commands were made to limit the number of bytes being sent. The telemetry data on the DET channel only needs to pass along the plain-text output of the terminal. The only feedback sought after on this channel is: (i) confirmation that a command was received/executed and (ii) the occasional display of a few lines of a text file (for reading the most recent temperature/pressure). As such, a low downlink speed does not significantly impact performance on this channel. Similar to GDR2, a speed of 4800 Bd was deemed sufficient for the DET channel. As we will see in the following section, the multiplexer port speeds do not operate in a 'zero sum' way. Each port uses only the bandwidth required, with the *maximum* set by the channel speed, and the rest of the available bandwidth is dynamically allocated to each port as needed.

The hardware interface for the DET channel consisted of a DB9 connection on the ground-side control station computer and a circular Milspec connector attached to the flange of the DOBC enclosure on the flight-side.

### 6.6 Multiplexing Layer

At its most basic, multiplexing several signals into one is a straight-forward process. If port A sends a string of bytes encoding the message 'Hello' and port B sends a message with 'World,' then a single channel could combine them into the string "{AHello}{BWorld}" where the curly brackets represent the bounds of a message and the first byte of each message is reserved to identify the port it came from. A multiplexing device would simply be one that takes these two inputs and outputs the combined string, to be passed to a similar device on the other end of the chain which understands the syntax and splits the data streams back into separate streams for ports A and B. However, as with many things, there is a significant difference between a simple concept and an effective implementation. The most important non-trivial element is the optimization of throughput/efficiency for multiple channels which collectively exceed or even just approach the total bandwidth of the combined channel. In this case, a good multiplexer has to make intelligent decisions about the size of packets and the rate at which packets can be sent from each port. It also needs to be able to handle technical details such as different hardware or software flow-control settings, baud rates, or encodings on the input channels. The commercial multiplexers used for FIREBall-1 were retained for FIREBall-2, and as such built into the design from an early stage. These were Data Comm for Business (DCB) SR-04 Statistical Multiplexers. A commercial multiplexer was not strictly required for the uplink side, as the TC Process (which is described in the next section) manually parses the multiplexer output anyway. However, multiplexing is an important process in maintaining stable telemetry, given the constant flux of data from both the GDR1 and NPT channels. If the FIREBall-2 architecture was rebuilt using the mini-SIP or SIP instead of the CIP, the multiplexers would then be similarly useful for uplink.

The SR-04 multiplexers take four data inputs via RJ45 (i.e. standard Ethernet cable) ports called 'Data Port 1,' 'Data Port 2,' and so on, as seen in Figure 6.6. As discussed earlier, the four input channels were DET, GDR1, GDR2, and NPT. The combined data stream is sent and received via the 'network composite' port, which also takes an RJ45 input. The baud rate for this port determines the maximum (instantaneous) combined bandwidth. We set the composite port speed to 38, 400 Bd, as the next highest value was 57, 600 Bd, beyond the capacity of the HH channel when using biphase encoding. The SR-04 multiplexers allocate bandwidth intelligently, such that the sum of the individual bandwidths may exceed the composite port speed and the multiplexer will allocate the available bandwidth as needed. As



Figure 6.6: The back of the SR-04 multiplexer used to combine the FIREBall-2 TM/TC channels (during a test in 2017- the blue cable was not used for flight, for obvious reasons). From left to right, the visible interfaces are: (i) the power supply; (ii) the network management port, used to configure the device; (iii) the network composite port, which carries the multiplexed signal, and (iv) the four individual 'data ports' for the channels that are to be multiplexed.

such, the NPT channel could be set to 38,400 Bd, the other channels set to their lowest functionally required speeds, and the multiplexer would allocate as much of the 38,400 Bd as possible to the NPT channel depending on availability. This was important to do, because 38,400 Bd was close to the limit at which the CNES software would suffer performance degradation.

To interface with the separated radio uplink and downlink channels, the composite outputs from both multiplexers were split into a TX (transmitting) and an RX (receiving) line. Connecting multiplexer TX and RX lines to the radio equipment requires an interfacing layer to adjust the signal levels and format the data correctly. This layer is called the TC Process layer because it is mostly of importance for the uplink direction.

# 6.7 TC Process Layer

The TC process is the interfacing layer between the multiplexers and the CSBF equipment. The CSBF HH downlink is a simple process: it takes a serial signal with transistor-transistor logic (TTL) levels (i.e. 0-5V) as input on the flight side, and outputs this signal as a serial signal with RS232 levels on the ground side. The only real interface handling required for the downlink is the conversion of the multiplexer output (in RS232) to TTL levels on the flight-side. Most of the work is on the uplink side, formatting the data with the GSE syntax and then parsing the parallel output of the CIP to turn it back into a serial signal.



Figure 6.7: Representation of the ground-side TC process as a DFA. The letter 'p' is the set of hexadecimal characters reserved for port identifier characters, and the letter 'W' denotes the MUX package header/footer character. Red lines indicate state transitions associated with errors. The text below each state describes it in plain terms.

On the ground-side, this means splitting the TX line from the multiplexer into twobyte packets, wrapping each packet in the CSBF commanding syntax, and adjusting the baud rate. In physical terms, the RJ45 output of the multiplexer is split into TX and RX on separate DB9 connectors. The connector carrying the TX line is passed through a DB9-to-USB adapter and plugged into a USB port on a dedicated TC process computer. There, a Python script (the 'uplink script') reads the serial input on the port (at 38,400 Bd), prepares the commands for the GSE, and sends them out via a second serial port (connected to the GSE) at 1200 Bd. There is also a downlink script, but this just transparently passes the HH data through the TC process computer so that it can be monitored and logged.

On the flight side, the uplinked dataword is displayed in parallel (i.e. on 16 separate pins) as a 1 ms low-going pulse. Two 'strobe' pins also pulse, one going low and one going high, to signal that a dataword is present. The flight side TC process is thus responsible for parsing the output from the CIP and forwarding it as a serial signal to the flight-side multiplexer. This is done using a microprocessor called a Ruggeduino, a ruggedized version of the popular Arduino microprocessors, built to be generally robust and withstand a large temperature range. Pins A3 through C7 on the J2 connector are used for the 16 bits of the dataword (A3-C3) and strobes (C5 and C7). The Ruggeduino reads in the parallel signal whenever both strobes pulse, converts it to hexadecimal characters, and then forwards them to the multiplexer.

The simplest version of the uplink script would be to take every pair of bytes from the multiplexer and wrap them in the command syntax. The flight-side,

MUX Packet	TCs	Flight-side Response
{1Hello}		
	{1	Set active port to 1
	He	Send {1He}
	11	Send {111}
	0}	Send {10}
{1Hi}		
	Hi	Send {1Hi}
{3Bye}		
	{3	Set active port to 3
	By	Send {3By}
	e}	Send {3e}

Table 6.1: An illustration of the uplink process. The MUX receives two consecutive messages from Port 1 ('Hello' and 'Hi') followed by a message from Port 3 ('Bye'). Each MUX packet is bounded by the header/footer characters '{},' within which the first byte specifies the transmitting port. The central column shows the data-words/telecommands transmitted by the uplink script after minimizing overhead, while the right-hand column shows the response of the flight-side TC process upon receiving each dataword.

in turn, would also transparently pass every two bytes it receives. However, as mentioned in Section 6.6 the multiplexer adds overhead data to separate data from each transmitting port and identify them. Because of the slow speed of the uplink, the amount of overhead data needs to be minimized. We implemented three main mechanisms to minimize the overhead on the uplink channel. First, consecutive individual packets from the same port would be concatenated until a timeout of 1 second was reached, a maximum length was reached, or a packet was sent from another port. This effectively meant that every full command issued had at most three bytes of overhead (header, footer and port number). Second, the port number was only uplinked if it was different from the previous port to communicate. To enable this, the flight-side TC process kept track of the most recent port to communicate. Every data-word received would be immediately transmitted to the 'active port' until a dataword with the syntax header, port number was received, at which point the active port would be updated. Finally, header and footer bytes were removed from the uplinked message and re-inserted by the flight-side TC process. The only caveat to this last part is that, for commands with uneven numbers of characters, the last transmitted byte would be a footer character. Table 6.1 shows an example of this process playing out, which is often the easiest way to get a clear understanding of the process.



Figure 6.8: Representation of the flight-side TC process as a DFA. The letter 'p' is the set of hexadecimal characters reserved for port identifiers, and the letter 'W' denotes the MUX package header/footer character. Red lines indicate state transitions associated with errors. The text beside each end state describes it in plain terms.

To ensure there were no corner cases in the uplink scripts which could result in the TC process becoming frozen or acting incorrectly, we developed the ground-side and flight-side scripts using deterministic finite automata (DFA) as a tool for more formal verification of the software. DFAs represent possible states and state transitions for a process based on a string of input or events. In this case, they provided a useful way to represent all possible states of the uplink TC processes, which could help me ensure in a deterministic manner that there were no 'gaps' in the software where a command could be lost or corrupted. Figure 6.7 shows the DFA representing the packet-collection for the ground-side uplink process. Figure 6.8 shows the same for the flight-side TC processes.

In addition to the uplink and downlink scripts, we developed software to monitor the communications channels and create logs. This software enabled me to keep track of when a command had finished uplinking and help the team in avoiding a 'traffic collision' of coincident commands. The uplink script was able to handle this



Figure 6.9: Flight-side TC process hardware (not including flight-side multiplexer). The DB9 connectors labelled 'TX' and 'RX' connect to the transmitting and receiving lines of the flight-side MUX composite port. The small chips in center frame are RS232-TTL converters. The microprocessor with a red shield in the top right is the Ruggeduino which handles uplinked datawords. The Ruggeduino in the bottom right, programmed by CNES team members, handles discrete command uplinks.

if it did happen, but it would cause confusion in keeping track of our activity. For monitoring the telemetry, the downlink log file was tracked in four separate terminals (one tracking each device). This allowed the team to see the telemetry data in real time. Although the byte-representation of the data is not human-readable, it is still clear when there is noise on the channel because the terminal floods suddenly with non-standard characters. It also helps to see when the telemetry from CNES or the Guider software is active, as a diagnostic to determine whether the on-board systems are powered on and transmitting data.

As a final note on the TC process, there are two main failure modes associated with the uplink process which can lead to successfully transmitted but inaccurate commands. While there is always the broader risk of critical failures (e.g. power failure) and continuous, noise-inducing failures (e.g. a faulty serial connection to the MUX), the insidious nature of these failure modes makes them particularly dangerous. For example, a small slew command could be changed to a very large one if one of the bits encoding the target destination becomes corrupted and is still transmitted.

The first is a mode in which an individual dataword is corrupted or lost during uplink. This could be caused by radio interference, an unstable connection on any of the sixteen dataword pins, or a systematic failure in the flight-side TC process software. The second mode is one in which the delay between uplinked datawords causes a buffering failure. The way to address the latter is to use a buffering algorithm with a long wait time, or one which waits for a specific footer before interpreting the received bytes.

Fortunately, the two high-risk channels - GDR1 and NPT - had robust buffering mechanisms on their flight-side computers which were not affected by the slow uplink. They also performed their own internal checksum validations. These channels are high risk because they control the pointing of the telescope. The DET and GDR2 channels are terminal-based access points. This means that buffering is not required, as the computers will not interpret the typed command until the return character is sent. It also means that commands can be validated by eye, as the telemetry data on these ports echoes the received bytes and shows the typed command as if it were in a local terminal. An incorrect terminal time is inconvenient but not likely to cause a critical failure. As such, additional error-checking was not a high priority for the TC process, especially given the risk associated with added complexity.

#### 6.8 Integration and Testing

Figure 6.10 shows a detailed outline of the full TM/TC system. This is a mission critical system; if it does not work, the payload is a dead weight in the air. As such, it is extremely important that the various mechanisms are properly tested. FIREBall-2 is an international collaboration, with different subsystems being owned, developed, and tested across Caltech, JPL, Columbia University, LAM, and CNES. As such, the majority of system-level testing and integration can only take place in the few months before launch when the team comes together. We will discuss these tests shortly. First, though, there was a list of smaller tests that needed to be carried out as the communications system was being developed.



Figure 6.10: A more detailed overview of the TM/TC subsystem, showing the main devices, connections, and speeds. The boxes labelled 'TTL' are the TTL-RS232 converter chips. Yellow boxes show the baud rates of each connection.

Most of the hardware used in the communications system for the 2018 flight was inherited from a previous 2009 flight, and had been successfully flight tested. The 2009 flight, however, used a different central piece of CSBF equipment (the mini-SIP), and a different setup for the rest of the TM/TC system. As such, new software and hardware had to be introduced. We have already discussed some of the development of the software in Section 6.7. However, it also needed to be tested under flight-like conditions, and the hardware (i.e. Ruggeduino and the new RS232-TTL chips) needed to be environmentally tested.

# Vacuum Chamber Testing at Caltech

The flight altitude of FB-2 is approximately 128,000 feet, at which atmospheric pressure is on the order of  $10^{-2}$  atm. The first series of tests conducted on the Ruggeduino were vaccum chamber tests, conducted on the Caltech campus. The goal of this was to test for critical failure of the device from extended periods at low pressure. The microprocessor was first programmed to communicate contin-



Figure 6.11: Schematic of the environmental testing setup in Palestine, TX. The dashed line represents the boundary of the vacuum chamber. The grey boxes shows the CSBF equipment. The parallel data-word (DW) and HH connections on the CIP are highlighted. The blue boxes highlight the new hardware, not tested during the 2009 flight.

uously for a period of 24 hours at a pressure of 0.001 atm (760 milli-torr). In this first round of testing, simple four-bit datawords were sent via a feedthrough from an Arduino outside the vacuum chamber and echoed back by the Ruggeduino inside the chamber. The Arduino forwarded the responses to a laptop, which saved them in a log file. We later processed the log file to check for consistency with the input. We then repeated the test sending large amounts of serial data via the serial Tx/Rx pins of the Ruggeduino. Both tests were passed with no critical failures.

### **Environmental Testing in Palestine, TX**

In 2017, we travelled to the CSBF headquarters in Palestine, Texas to conduct environmental testing on the new flight hardware and perform the first integration tests of our TM/TC system with the CSBF hardware. Figure 6.11 shows a schematic of the test setup. The flight-side hardware (Ruggeduino, MUX, and RS232-TTL converters) were placed inside a pressure and temperature controlled chamber and connected to the CIP as they would be in flight. The CIP was also connected to a receiver and transmitter outside the vacuum chamber via the same feedthrough.

A script generated simulated multiplexer data on a dedicated test computer, which was fed via a USB cable into a separate port, to simulate the setup used during flight. The uplink TC process, described above, then parsed the incoming data and sent commands to the GSE. The Ruggeduino was programmed to behave as if during flight, parsing the incoming datawords when the strobes pulsed, and sending them on towards the MUX via the RS232-TTL converters. Data Port 1 on the MUX was connected to the testing computer via a feedthrough in the vacuum chamber, so that the data received on the 'flight side' could be monitored and data could be sent over the downlink.

Four tests were conducted over six hours with this set-up. First, the temperature was dropped to -40C. A two-hour test was conducted at 600 mbar (0.5 atm,) at which point the pressure was dropped to full vacuum and the test continued for another hour. After this, the pressure was raised back to 600 mbar. The temperature was raised to +50 for a second two hour test and a final hour long test at full vacuum. Overall, the test was a success and the uplink/downlink worked without failure. However, there was inevitably some noise as the test was being conducted using radio communication indoors. This was controlled for by connecting the CIP with the GSE through a wired connection, at which point the issue disappeared.

### Integration and Testing in Fort Sumner, NM

As mentioned earlier, the vast majority of system-level integration and testing for the communications system took place during the campaign months leading up to the 2018 flight. This is because the communications system needed the CSBF equipment and the presence of each team (detector, guider, CNES) to be exhaustively tested. Given the relatively short time frame, a detailed list of tests was drawn up to ensure the system had been fully verified before launch. This testing process, the stages for which are laid out in Table 6.2, begins with basic inventory checks and goes thorough testing of each component (and small assembly) in the communications system before moving on to the final integration tests, in which the control stations and flight devices are included.

The final two stages of pre-flight testing are system-wide indoor tests, in which the telescope itself remains in the lab and the team runs through exercises, and outdoor sky tests in which the telescope is suspended from a crane over the course of a night and target fields are actually acquired. For the former, indoor interference made the



Figure 6.12: Pictures of the TM/TC equipment environmental testing setup in Palestine, TX. The top panel shows the outside of the vacuum chamber. The bottom panel shows the multipler (right), the CIP (top-left), and the Ruggeduino connected inside the chamber.

full radio system too noisy to use. As such, we often used a wired connection to the gondola, but left the rest of the communications system intact. For the on-sky tests, the conditions for radio communication were also not ideal (the line-of-sight to the radio tower was not always clear), but were much better. These tests required that the full communications system be in place.

1. Inventory Check		
	1.1	Hardware Inventory
	1.2	Software Inventory
	1.3	External Interfaces
2. Unit/Module Tests		
	2.1	Serial Cables Inspection
	2.2	Serial Ports Inspection
	2.3	MUX Box Inspection
	2.4	MUX-MUX Wired Test
	2.5	Ground-Flight System Direct Test
	2.6	FLT TC Process Test
	2.7	GND TC Process Test
	2.8	CIP Isolation Test
3. Integration Tests		
	3.1	Single device, wired MUX
	3.2	Multiple device, wired MUX
	3.3	Single device, MUX with CIP
	3.4	Multiple device, MUX with CIP
	3.5	Full indoor system test
	3.6	Full sky Test

Table 6.2: The itinerary developed for integration and testing of the communications system in Fort Sumner, NM, for the 2018 flight of FIREBall-2. A spreadsheet containing test details and inventories was attached to each step, to that ensure the full system had been thoroughly verified.

# 6.9 2018 Flight

FIREBall-2 was launched at 10:20 AM MDT on the  $22^{nd}$  of September 2018 from the CSBF base in Fort Sumner, New Mexico. Hoisted by 1.1 million cubic meters of helium, the payload reached the intended altitude of 39km (~ 128,000ft) approximately five hours after launch. After three hours at altitude, however, the balloon began to sink lower. Sadly, this premature descent coincided with sunset, so the mission only obtained 45 minutes of science observations while losing altitude. The planned flight had included six hours of science observations. Every kilometer of lost altitude at 30 - 40km corresponds to a reduction in throughput of 5%.

In addition to these difficulties, the rarity of launch opportunities during the previous weeks (due to bad weather) had forced us to accept a launch one day before a full moon. While this already would have meant an increased background, the shape of the deflated balloon conspired to compound this problem by focusing moonlight directly into the instrument. The resulting background level was more than a

thousand times higher than expected and prevented the use of photon-counting observations, which require an extremely low flux.

Despite these extremely unfortunate circumstances, the 2018 flight of FIREBall-2 proved to be a successful engineering flight for the payload. The team was in full control of an EMCCD, UV spectrograph-carrying 1m telescope at  $\sim$  39km. Picouet et al., 2020 (submitted to JATIS - the Journal of Astronomical Telescopes, Instruments and Systems), details the performance of the instrument during this flight and the expected improvements, which we will summarize here.

The  $5\sigma$  sensitivity limit in terms of the Ly $\alpha$  luminosity of a 7" extended source in a 30 minute observation window is found to be  $L_{5\sigma}^{Ly\alpha} \sim 1.5 \times 10^{43}$  erg s<sup>-1</sup>. This would be sufficient to detect many of the nebulae reported in Borisova et al., 2016, Arrigoni Battaia et al., 2019 and O'Sullivan et al., 2020, but whether similarly bright nebulae still exist at  $z \sim 0.7$  is still unknown.

The surface brightness limit at a distance 7" away from a central source is found to be  $1.7 \times 10^{-16}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>. Taking into account the  $(1+z)^4$  redshift dimming, this corresponds to an equivalent limit of ~ 9 × 10<sup>-18</sup> erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup> at z = 2.5, which would be sufficient to detect the brightest CGM regions reported in the above surveys (which tend to peak at ~ 5 × 10<sup>-17</sup> erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>).

### 6.10 Future Work

FIREBall-2 will be back in 2021 with another flight out of Fort Sumner, NM. With the balloon deployment issue (which affected two flights in 2018) solved and, ideally, a darker sky, this next flight has a great chance of yielding the first detections of extended Ly $\alpha$  around  $z \sim 0.7$  galaxies.

Vincent Picouet (Picouet et al. 2020, submitted to JATIS) has done the hard work of characterizing the 2018 flight performance and anticipating the performance of the 2021 flight, so we encourage the reader to read that paper, rather than repeat the findings here. For now, let it suffice to say that only minor changes are expected (including some baffling to block rogue moonlight), and the major benefit is expected from a more successful flight and observing conditions.

### References

Arrigoni Battaia, F. et al. (Jan. 2019). "QSO MUSEUM I: a Sample of 61 Extended Ly *α* Emission Nebulae Surrounding *z* ~ 3 Quasars." In: *MNRAS* 482.3, pp. 3162–3205. DOI: 10.1093/mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. DOI: 10.3847/0004-637X/831/1/39.
- Hamden, E. et al. (May 2019). "FIREBall-2: Advancing TRL While Doing Proof-of-concept Astrophysics on a Suborbital Platform." In: *Proc. SPIE*. Vol. 10982. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 1098220. DOI: 10.1117/12.2518711. arXiv: 1905.00433.
- Martin, D. C. et al. (Jan. 2005). "The Galaxy Evolution Explorer: A Space Ultraviolet Survey Mission." In: *ApJ* 619.1, pp. L1–L6. doi: 10.1086/426387. arXiv: astro-ph/0411302 [astro-ph].
- Milliard, B. et al. (2010). "FIREBALL: the Faint Intergalactic Medium Redshifted Emission Balloon: Overview and First Science Flight Results". In: *Proc. SPIE*. Vol. 7732. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 773205. DOI: 10.1117/12.857850.
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 *z* ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.
- Tuttle, Sarah E. et al. (2010). "FIREBALL: the First Ultraviolet Fiber-Fed Spectrograph". In: *Proc. SPIE*. Vol. 7732. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 773227. DOI: 10.1117/12.857877.

### Chapter 7

## SUMMARY AND OUTLOOK

This final chapter will briefly summarize the contributions from this thesis and then present some thoughts on the near future of observational CGM studies. In particular, this latter part will focus on: the bottle-neck to IFS survey work presented by data analysis, the potential of meta-analyses of the growing field of observations, and the instrumental challenge of resolving galaxy properties. First, though, a summary.

### 7.1 Summary

Our observational picture of high-redshift galaxy environments has become clearer in the past six years. Prior to 2015, fully-resolved, direct detections of CGM emission were few and far between, with only a few small-sample studies published (e.g. Christensen et al., 2006). Today, by a simple counting of the survey work by Herenz et al. (2015), Wisotzki et al. (2016), Borisova et al. (2016), Arrigoni Battaia et al. (2019), Cai et al. (2019), Farina et al. (2019), and this thesis (O'Sullivan et al., 2020), we now collectively have 201 deep IFS observations of the CGM around quasars and galaxies at z = 2 - 7. These data place us in a new observational era, in which we have statistically significant samples capable of constraining the mean morphological and kinematic properties of the CGM at high redshift, as well as its evolution over this period.

#### The FLASHES Survey I: Ly $\alpha$ Pilot Study

The pilot study component of the FLASHES Survey (O'Sullivan et al., 2020) represents 48 of the above 201 deep IFS observations, and the only large sample covering the redshift range z = 2.3 - 3.0. While the detailed list of conclusions is available in Chapter 4, here is a short summary of the main observational results from the pilot survey:

1. Extended Ly $\alpha$  nebulae are a common feature of z = 2.3 - 3.1 QSOs, though the average surface brightness appears lower than those reported at higher redshifts and - possibly as a result of this - 11/48 target fields contained no significant detections.

- 2. The average observed radial surface brightness profiles of the nebulae are best-fit by a Sérsic profile with index n = 0.5 0.6 and half-light radius  $R_e \simeq 24$  pkpc. Though this is the average size, the flux-weighted average radii spread over a relatively wide range, from  $R_{\rm rms} = 12 59$  pkpc.
- 3. The stacked, circularly-averaged radial surface brightness profile of these z = 2.3 3.0 nebulae appears several times fainter than those reported at higher redshift (Borisova et al., 2016; Arrigoni Battaia et al., 2019). However, a more detailed analysis showed that the average intrinsic brightness *within* the detected regions is comparable, implying that the majority of this difference arises from a lower covering factor rather than a globally fainter emission. Combined with the fact that the FLASHES Pilot detections are also more eccentric on average, and similar findings by Cai et al., 2019 for a sample of z = 2.25 QSOs, this paints a picture that the neutral CGM may be growing clumpier at lower redshifts.
- 4. The centroid of  $Ly\alpha$  emission in these nebulae varies by thousands of kilometers per second with respect to the QSO redshift and is, on average, biased towards redder wavelengths by +871 km s<sup>-1</sup>.
- 5. With the medium resolution (R = 2500) PCWI data, the line profiles appear to be mostly single or double-peaked.
- 6. The RMS velocities of the nebulae appear broadly consistent with those expected from gravitational motions in dark matter halos in the mass range  $\text{Log}_{10}(\text{M}_{h}[\text{M}_{\odot}]) = 12.2^{+0.7}_{-1.2}$ , expected to host  $z \simeq 2.7$  QSOs.
- 7. The majority of the kinematic maps do not show strong signs of coherence, but two in particular exhibit a high flux-weighted specific angular momentum. Martin, O'Sullivan, et al., 2019b performed detail modeling of the kinematics of one of these targets, finding strong evidence for the presence of cold-flow accretion along filaments connected to the QSO.

### The FLASHES Survey II: Deep Study of CGM Metal Line Emission

With this recent rapid growth in observations of extended Ly $\alpha$  nebulae at high redshift, the natural next step is to look beyond hydrogen. UV metal-line emission from atomic species such as NV, CIV, and HeII can help us constrain the mass and metallicity of the CGM, probe wider ranges in the temperature-density phase space, and provide additional handles on gas kinematics.

However, these lines remain extremely challenging to detect. Only recently have any such detections been made. The most notable is probably the study using archival MUSE data by Guo et al., 2020. The authors re-analyzed data for 80  $z \sim 3$ QSOs, looking both to make individual detections and derive a statistical result from stacking the data. They directly detect metal-line emission in about 20% of the sample. From the statistical profiles, they detect CIV  $\lambda$ 1549, HeII  $\lambda$ 1640, and CIII]  $\lambda$ 1909 out to a radius of 45 pkpc and conclude that CIV and HeII have much steeper radial declines than Ly $\alpha$ .

The first deep study of the FLASHES Survey (O'Sullivan et al. in prep / Chapter 5 of this thesis) targeted a sub-sample of 10 QSOs from the pilot study with deep KCWI observations and wavelength coverage sufficient to detect all of the above lines. From this study, we have so far learned:

- 1. Of 8 targeted QSO fields (5 pilot survey detections and 3 non-detections), 7/8 contain Ly $\alpha$  emission. Three fields contain extended NV, CIV, and HeII emission in addition to extended Ly $\alpha$  nebulae. Another two contain detected CIV and Ly $\alpha$ , but no NV or HeII. Ly $\alpha$  alone is detected in two more and one field contains no detections whatsoever. If the remaining non-detection contains an extended Ly $\alpha$  nebula, its luminosity is L(Ly $\alpha$ ) < 1.4 × 10<sup>42</sup> erg s<sup>-1</sup>.
- 2. The detected luminosities in NV range from  $L(NV) = 2.6 4.5 \times 10^{42} \text{ erg s}^{-1}$ , while upper limits of  $L(NV) \leq 1 \times 10^{42} \text{ erg s}^{-1}$  are placed on most other fields. CIV luminosities are found to be  $L(CIV) = 2.0 - 23.3 \times 10^{42} \text{ erg s}^{-1}$ , with upper limits of  $L(CIV) \leq 1.5 \times 10^{42} \text{ erg s}^{-1}$  placed on other fields. The range of measured HeII luminosities is  $L(HeII) = 7.1 - 9.3 \times 10^{42} \text{ erg s}^{-1}$ , with upper limits of  $L(HeII) \leq 1 - 2 \times 10^{42} \text{ erg s}^{-1}$  placed on other fields.
- 3. Kinematic maps are obtained for all four lines for two targets. The velocity maps show broad agreement across the lines, though it is difficult to compare for the small detections in NV and HeII. Velocities are found to be within a few hundred km s<sup>-1</sup> for all maps. Dispersions for LyA and CIV are mostly in the range of  $\sigma_v = 200 400$  km s<sup>-1</sup>. The NV and HeII velocity maps are affected systematics and more modeling work is needed to measure these dispersions carefully.

## The Keck Cosmic Web Imager

The relatively sudden surge in spatially and spectrally resolved CGM observations was empowered by both new instrumentation and new software. The Palomar Cosmic Web Imager, installed in 2009, empowered O'Sullivan et al., 2020. The Multi-Unit Spectroscopic Explorer (MUSE) (Caillier et al., 2014), installed in 2014, empowered Borisova et al. (2016), Arrigoni Battaia et al. (2019), and Farina et al. (2019).

Chapter 2 of this thesis outlined contributions to the development and testing of the Keck Cosmic Web Imager (KCWI - Morrissey et al., 2018). KCWI was installed in 2017 and has since empowered the aforementioned CGM survey by Cai et al., 2019, as well as the FLASHES Deep Survey (O'Sullivan et al., in prep), and a plethora of other studies (Simha et al., 2020; Gannon et al., 2020; Bresolin et al., 2020; Burchett et al., 2020; Kalita and Ebeling, 2019; Martin, O'Sullivan, et al., 2019a; van Dokkum et al., 2019; Law et al., 2018).

## **CWITools**

The data analysis required to conduct IFS surveys is complex and time-intensive. In Borisova et al. (2016), the authors present a software package called 'CubEx', written in FORTRAN for the extraction of faint signals from MUSE data. The need for this kind of tool is evidenced by the use of CubEx in subsequent surveys such as Arrigoni Battaia et al. (2019) and Cai et al. (2019).

Chapter 3 of this thesis presented a software package, written in Python and specifically designed for the Cosmic Web Imager instruments, called CWITools (O'Sullivan and Chen, in prep). CWITools was developed out of necessity for the FLASHES Survey (Chapters 4 and 5). The full release is scheduled to coincide roughly with the completion of this PhD thesis, though a number of colleagues and collaborators have already benefited from the availability of the beta version of CWITools. It is my hope that the development of CWITools will empower many future PCWI and KCWI studies.

## FIREBall-2

Studying Ly $\alpha$  and bright metal lines such as CIV  $\lambda 1549$  in the low-redshift ( $z \leq 1$ ) CGM in detail is a major technological challenge. The Faint Intergalactic Medium Redshifted Emission Balloon (FIREBall-2) is trying to observe some of the faintest gas in the universe with integral field spectroscopy in a part of electromagnetic spectrum where throughput and quantum efficiency do not come easily. The rewards for success are significant - the first observations of the CGM the low-z universe but the risk is high. Not only does every part of the instrument-telescope system need to work, but the weather needs to be clear, the high-altitude wind speeds need to be low, the sky needs to be dark, the balloon and attachment need to function without problem, and the launch needs to go off smoothly.

Chapter 6 presented contributions to the development of new UV capabilities as part of the wonderful team sending FIREBall-2 (Hamden et al., 2019) to the edge of the atmosphere. This contribution was a successful implementation of the commanding and telemetry system which enabled us to collect at least some observational data and manage the payload through an unexpectedly early descent caused by a tear in the balloon.

Despite this unfortunate outcome, FIREBall-2's 2018 flight was an engineering success for the telescope and instrument. Every system - from communication to guidance to image acquisition and download - worked as planned. The team controlled a fully functional UV telescope and instrument at ~ 120,000 ft with a novel, UV-capable, photon-counting CCD. Despite the fact that the moonlight and deflated balloon conspired to raise the background level by two orders of magnitude, FIREBall-2 detected continuum emission from three of the brightest target galaxies and obtained a limiting surface brightness of  $1 \times 10^{-16}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup> for extended emission at a distance of 7" from the galaxy centers (Picouet et al. 2020, submitted to JATIS). Taking into account planned improvements and the hope for better observing conditions, the predicted performance of the 2021 flight is that this limiting surface brightness will be improved to  $1.4 - 1.8 \times 10^{-17}$  erg s<sup>-1</sup> cm<sup>-2</sup> arcsec<sup>-2</sup>.

### 7.2 Outlook

With the contributions from this thesis thus summarized, let us end with a brief outlook on a few topics regarding the future of CGM and IGM observations.

### 7.3 Resolving Galaxy and Environment Simultaneously

To fully understand the relationship between galaxy and environment, we need to study the properties of both. Unfortunately, one of the major limitations of the current field of IFU CGM observations is the inability to resolve the host galaxy, small satellites (e.g. if there is an interacting pair), and their properties.

Given that this is a thesis concerned largely with the instrumentation required to study the CGM, it seems appropriate to compare an ideal observation with current and planned facilities. Putting aside any restrictions for the moment, the ideal observation to study galaxy and CGM would have both spatial and resolution, so that the kinematics of both (for various emission lines) could be studied, and in particular, it would have the following properties:

- 1. Field of view: O(100 pkpc), to cover the CGM.
- 2. Angular resolution: O(1 pkpc), to resolve host galaxies at high redshift.
- 3. Spectral resolution:  $O(10 \text{ km s}^{-1})$ , to resolve galaxy and CGM kinematics.
- 4. Wavelength coverage: O(100 nm), to probe different emission lines.

First, let us justify each of these briefly. The virial radii of massive ( $M_h \sim 10^{12} M_{\odot}$ ) dark matter halos are typically on the order of one hundred proper kiloparsecs ( $R_{vir} \sim 163h^{-1}$  pkpc, Equation 1.2). Thus, a field of view covering  $300 \times 300$  pkpc<sup>2</sup> would be sufficient to cover the CGM around a central source. High-redshift galaxies are known to be more compact than their modern counterparts with half-light radii of just a few kiloparsecs (Bouwens et al., 2004). As such, kiloparsec resolution would be needed to resolve the host galaxy. Gravitational motions in galactic dark matter halos are typically on the order of 100 - 300 km s<sup>-1</sup>. To resolve such velocities well, a resolution element  $\Delta v < 100$  km s<sup>-1</sup> is needed. Finally, if we want to probe bright UV emission lines, we will need wavelength coverage on the order of 100 nm (e.g. Ly $\alpha \lambda 1216$  and CIV  $\lambda 1640$  at a redshift of z = 2.5 are separated by 148 nm).

One proper kiloparsec at a redshift of z = 2 has an angular size of  $\delta\theta \sim 0.1''$ . This resolution either requires adaptive optics on a ground-based telescope or a spacebased telescope with diameter of  $D \gtrsim 1.5$  m. MUSE, on the VLT, is designed to work with adaptive optics, achieving a spatial resolution on the order of 0.06-0.08'', albeit with a Strehl ratio of only 10%. The field of view in this mode is  $\sim 7.4'' \times 7.4''$ , a physical size of  $\sim 60 \times 60$  pkpc at z = 2.5. Thus, the required angular resolution is feasible with existing instrumentation, but to achieve it in combination with the desired field of view, a new instrument is required.

To resolve a velocity width of  $\Delta v \leq 100 \text{ km s}^{-1}$ , a spectral resolution of  $R \gtrsim 3000$  is needed. KCWI, in a configuration with the low-resolution blue grating (BL) and small image slicer, can achieve a resolution of  $R \sim 3600 \ (\Delta v \sim 75 \text{ km s}^{-1})$  and an instantaneous bandpass of  $\Delta \lambda \sim 200 \text{ nm}$ . Again, this shows that this particular aspect is within reason for existing instrumentation, but the challenge remains to

achieve all of the desired requirements simultaneously. KCWI does not have adaptive optics and, even if it did, neither KCWI nor MUSE could simultaneously achieve the desired angular resolution and field of view.

So, what about future instrumentation? The Thirty Meter Telescope (TMT) looms on the horizon for ground-based observations and the concept for the Large UV / Optical / IR (LUVOIR) Surveyor is being developed. Will either of these host instruments capable of this ideal observation? As it stands, the answer is that it is not likely, but there is some potential.

TMT is planned to come online with an IR adaptive optics system, the Narrow Field Infrared Adaptive Optics System (NFIRAOS). This will provide diffractionlimited resolution in the J, H, and K bands ( $0.8 - 2.4\mu$ m) across a field of view of  $34'' \times 34''$ , or  $280 \times 280$  pkpc<sup>2</sup>. Among the first light instruments is an IR integral field spectrograph - the Infrared Imaging Spectrograph (IRIS) - which takes broadband images of the full AO field of view and offers integral field spectroscopy of the very inner region (maximum  $2.25'' \times 4.4''$  with one of the slicer IFU options). Despite the small field of view, this could provide IFS observations of H $\alpha$  (656.28 nm) emission at redshifts z = 0.5 - 2.5, resolving galaxies. Used in combination with MUSE/KCWI Ly $\alpha$  observations or Keck-OSIRIS H $\alpha$  observations, this could offer a way to simultaneously observe environment and galaxy connection.

LUVOIR, still in the concept development phase, has a multi-object spectrograph (LUVOIR Multi-Object Spectrograph - LUMOS) planned, but this would not provide the same contiguous spatial coverage as an IFS. There is a flexible IFS planned as part of the Extreme Coronagraph for Living Planetary Systems (ECLIPS) - though whether or not it could be used for CGM imaging around a bright central source, rather than imaging planets around a nearby star, is unclear to me.

#### Meta-analyses of IFU CGM Studies

Since the early 2000s, a large number of IFS studies have reported the surface brightness, morphology, and kinematics of the CGM around QSOs (and other galaxies) across a wide redshift range of  $z \simeq 2-7$  (see the summary at the beginning of this chapter for a list).

A number of these papers have dedicated a section of their analysis to comparing and contrasting with other studies at different redshifts, but there is great potential for a more detailed study which aggregates all of this data and carefully controls for the varying selection, sensitivity, and methodology. This would allow the researcher to leverage hundreds of observations in exploring the redshift evolution of CGM morphology and brightness, as well as dependence on QSO magnitude, IR color, or other environmental properties (e.g. local clustering).

In O'Sullivan et al., 2020, we reported a difference in average surface brightness of CGM emission compared to higher redshift studies by Borisova et al., 2016 and Arrigoni Battaia et al., 2019. I found that this was primarily explained by a lower covering factor, as the detected regions themselves were of a similar intrinsic brightness to those reported at higher redshift.

While we are still confident in this conclusion, the comparison was limited to two other surveys and it was a small part of the larger project. It would be great to see a longer, more detailed discussion of this and other trends that perhaps the authors of each individual survey did not have the time to detect or notice. A preliminary study along these lines could be a good early PhD student project, as it involves a lot of reading, requires no new observations, and has the potential for very interesting take-aways.

#### **Future FLASHES Projects**

The core observational work of this thesis is the FLASHES Survey; a multi-year survey of the CGM around  $z \simeq 2.3 - 3.1$  QSOs. This thesis has presented the pilot survey, targeting Ly $\alpha$  with moderately deep PCWI observations, and the first follow-up study, targeting metal line emission in addition to Ly $\alpha$  with deep KCWI observations. These were the highest priority projects to complete, but the FLASHES database contains yet-to-be-published observations and has the potential for a number of additional studies.

First, we obtained deep (1-3hr) high-resolution ( $R \sim 9000$ ) KCWI observations of three QSOs with extended emission. The goal here is to study the spectral shape of the Ly $\alpha$  emission in detail and study the effect of spectral resolution on kinematic maps. Second, we obtained moderately deep ( $\leq 1hr$ ) observations of Ly $\alpha$  for the majority of the 11 non-detections in the FLASHES Pilot study.

As the ubiquity of extended Ly $\alpha$  halos around high-redshift QSOs becomes more and more established, the few exceptions to the rule become more interesting. However, the upper limits provided by the pilot survey were not very restrictive. In Chapter 5, we included three non-detections from the pilot sample. In 2/3 of the fields, we detect significantly extended Ly $\alpha$  and even CIV emission. However, in one field (SDSS2234+2637), we find no extended emission whatsoever, despite an exposure time of 2.5 hr on KCWI.

Against the backdrop of surveys reporting nearly ubiquitous Ly $\alpha$  halos around highredshift galaxies and simulations which predict as much, the convincing absence of such emission seems more curious than its presence. What is different about the QSO SDSS2234+2637 that makes it devoid of emission? It happens to be one of the faintest QSOs we have observed, but O'Sullivan et al. (2020) found no significant relationship between QSO magnitude and Ly $\alpha$  luminosity. Perhaps it is a result of absorption in the foreground. Either way, a study of these follow-up observations of pilot sample non-detections could yield interesting insight into the range of brightnesses and sizes of Ly $\alpha$  emission around high-z QSOs.

Finally, though this project is not one based on a specific subset of observations, all observations in the FLASHES database could be used to search for coincident emission and absorption features. Imaging absorbers in emission could lend insight into the nature of damped  $Ly\alpha$  absorbers and Lyman limit systems.

#### References

- Arrigoni Battaia, F. et al. (Jan. 2019). "QSO MUSEUM I: a Sample of 61 Extended Ly α Emission Nebulae Surrounding z ~ 3 Quasars." In: *MNRAS* 482.3, pp. 3162– 3205. DOI: 10.1093/mnras/sty2827. arXiv: 1808.10857 [astro-ph.GA].
- Borisova, E. et al. (Nov. 2016). "Ubiquitous Giant Ly $\alpha$  Nebulae around the Brightest Quasars at  $z \sim 3.5$  Revealed with MUSE." In: *ApJ* 831, 39, p. 39. doi: 10.3847/0004-637X/831/1/39.
- Bouwens, R. J. et al. (Aug. 2004). "Galaxy Size Evolution at High Redshift and Surface Brightness Selection Effects: Constraints from the Hubble Ultra Deep Field." In: *ApJ* 611.1, pp. L1–L4. DOI: 10.1086/423786. arXiv: astro-ph/0406562 [astro-ph].
- Bresolin, F. et al. (May 2020). "Internal Kinematics of Giant H II Regions in M101 with the Keck Cosmic Web Imager." In: *MNRAS* 495.4, pp. 4347–4365. DOI: 10.1093/mnras/staa1472. arXiv: 2005.10369 [astro-ph.GA].
- Burchett, J. et al. (May 2020). "Circumgalactic Mg II Emission from an Isotropic Starburst Galaxy Outflow Mapped by KCWI." In: *arXiv e-prints*, arXiv:2005.03017, arXiv:2005.03017 [astro-ph.GA].
- Cai, Z. et al. (Sept. 2019). "Evolution of the Cool Gas in the Circumgalactic Medium (CGM) of Massive Halos A Keck Cosmic Web Imager (KCWI) Survey of Ly $\alpha$  Emission around QSOs at  $z \approx 2$ ." In: *arXiv e-prints*, arXiv:1909.11098, arXiv:1909.11098. [astro-ph.GA].

- Caillier, P. et al. (July 2014). "MUSE from Europe to the Chilean Sky." In: Groundbased and Airborne Instrumentation for Astronomy V. Vol. 9147. Proc. SPIE, 91475K. doi: 10.1117/12.2057056.
- Christensen, L. et al. (Dec. 2006). "Extended Lyman- $\alpha$  Emission Around Bright Quasars." In: A&A 459, pp. 717–729. DOI: 10.1051/0004-6361:20065318. eprint: astro-ph/0603835.
- Farina, E. P. et al. (Dec. 2019). "The REQUIEM Survey. I. A Search for Extended Lyα Nebular Emission Around 31 z > 5.7 Quasars." In: *ApJ* 887.2, 196, p. 196. DOI: 10.3847/1538-4357/ab5847.
- Gannon, J. et al. (May 2020). "On the Stellar Kinematics and Mass of the Virgo Ultradiffuse Galaxy VCC 1287." In: *MNRAS* 495.3, pp. 2582–2598. DOI: 10. 1093/mnras/staa1282. arXiv: 2005.03041 [astro-ph.GA].
- Guo, Yucheng et al. (Jan. 2020). "Metal enrichment in the circumgalactic medium and Ly $\alpha$  haloes around quasars at z ~3". In: *arXiv e-prints*, arXiv:2001.05473, arXiv:2001.05473 [astro-ph.GA].
- Hamden, E. et al. (May 2019). "FIREBall-2: Advancing TRL While Doing Proof-of-concept Astrophysics on a Suborbital Platform." In: *Proc. SPIE*. Vol. 10982. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 1098220. DOI: 10.1117/12.2518711. arXiv: 1905.00433.
- Herenz, E. C. et al. (May 2015). "Where is the Fuzz? Undetected Lyman α Nebulae Around Quasars at z ~ 2.3." In: A&A 576, A115, A115. DOI: 10.1051/0004-6361/201425580. arXiv: 1502.05132.
- Kalita, Boris S. and Harald Ebeling (Dec. 2019). "Jellyfish: Resolving the Kinematics of Extreme Ram-pressure Stripping at z~0.3." In: *ApJ* 887.2, 158, p. 158. DOI: 10.3847/1538-4357/ab5184. arXiv: 1910.11898 [astro-ph.GA].
- Law, D. R. et al. (Oct. 2018). "Imaging Spectroscopy of Ionized Gaseous Nebulae around Optically Faint AGNs at Redshift *z* ~ 2." In: *ApJ* 866.2, 119, p. 119. doi: 10.3847/1538-4357/aae156. arXiv: 1809.04089 [astro-ph.GA].
- Martin, D. C., D. O'Sullivan, et al. (July 2019a). "Multi-filament Gas Inflows Fuelling Young Star-Forming Galaxies." In: *Nature Astronomy* 3, pp. 822–831. DOI: 10.1038/s41550-019-0791-2. arXiv: 1904.11465 [astro-ph.GA].
- Martin, D. C., D. O'Sullivan, et al. (July 2019b). "Multi-filament Gas Inflows Fuelling Young Star-forming Galaxies." In: *Nat. Astron.* 3, pp. 822–831. doi: 10.1038/s41550-019-0791-2. arXiv: 1904.11465.
- Morrissey, P. et al. (Sept. 2018). "The Keck Cosmic Web Imager Integral Field Spectrograph." In: *ApJ* 864.1, p. 93. DOI: 10.3847/1538-4357/aad597.
- O'Sullivan, D. et al. (May 2020). "The FLASHES Survey I: Integral Field Spectroscopy of the CGM around 48 z ~ 2.3 3.1 QSOs." In: *ApJ* 894.1, 3, p. 3. DOI: 10.3847/1538-4357/ab838c. arXiv: 1911.10740.

- Simha, S. et al. (May 2020). "Disentangling the Cosmic Web Towards FRB 190608." In: *arXiv e-prints*, arXiv:2005.13157, arXiv:2005.13157. arXiv: 2005.13157 [astro-ph.GA].
- van Dokkum, P. et al. (Aug. 2019). "Spatially Resolved Stellar Kinematics of the Ultra-diffuse Galaxy Dragonfly 44. I. Observations, Kinematics, and Cold Dark Matter Halo Fits." In: *ApJ* 880.2, 91, p. 91. DOI: 10.3847/1538-4357/ab2914. arXiv: 1904.04838 [astro-ph.GA].
- Wisotzki, L. et al. (Mar. 2016). "Extended Lyman  $\alpha$  Haloes Around Individual High-redshift Galaxies Revealed By MUSE." In: A&A 587, A98, A98. DOI: 10. 1051/0004-6361/201527384. arXiv: 1509.05143.

## Appendix A

# EXTENDED PHOTOMETRY DATA FOR FLASHES TARGETS

Target	GALEX <sup>a</sup>		SDSS DR12 <sup>b</sup>					/	2MASS <sup>a</sup>	2	WISE <sup>d</sup>				FIRST <sup>e</sup>
SDSS/HS	FUV	NUV	u	g	r	i	Z	Н	J	K	W1	W2	W3	W4	1.4GHz
1700+6416	18.99	18.77	16.74	16.05	15.94	15.84	15.77	16.64	17.53	16.80	15.70	15.49	13.88	13.09	
0006+1614		22.86	19.18	18.33	18.13	18.10	17.84	15.22	16.15	15.78	18.00	17.42	16.00	15.38	
0012+3344		20.81	18.97	18.32	18.27	18.26	17.97	17.37	18.18	17.46	17.74	17.40	15.61	14.53	
0013+1630			18.93	18.33	18.26	18.17	17.93	16.93	18.22	17.55	17.26	17.14	15.74	14.83	
0015+2927			19.31	18.15	17.99	18.01	17.90	17.23	18.29	17.46	17.59	17.38	16.01	15.46	
0041+1925			20.95	19.86	19.70	19.50	19.32	17.22	18.20	17.70	19.34	19.18	17.35	15.46	
0057+0346		20.65	18.84	18.18	18.13	18.06	17.84				17.62	17.34	16.19	15.30	
0103+1316			17.32	16.57	16.37	16.27	16.16	17.00	18.37	17.21	16.37	16.00	14.02	13.18	
0107+1104			21.51	20.96	20.66	20.68	20.39	15.85	16.84	16.38	19.98	20.22	16.88	15.62	
0108+1635			18.1	17.19	17.00	16.87	16.67				16.56	16.30	14.75	13.91	
0118+1950			19.11	18.14	17.99	18.01	17.89	16.04	17.18	16.73	17.84	17.64	16.12	15.21	
0126+1559			19.77	19.00	18.82	18.81	18.60	16.77	18.10	17.50	18.62	18.37	16.78	15.66	
0132+3326			19.73	19.10	19.18	19.10	18.77	•••			18.00	18.02	17.32	15.27	
0137+2405			24.93	22.23	21.80	21.67	22.06	16.55	17.75	17.54	20.30	20.00	16.98	15.06	
0144+0838			18.92	18.38	18.26	18.27	18.09				18.25	17.46	15.66	14.92	
0205+1902			18.31	17.45	17.27	17.07	16.90	16.64	17.53	16.80	16.75	16.43	14.58	13.84	
0211+3117			19.71	19.00	18.86	18.86	18.79	17.30	18.31	17.75	18.89	18.54	16.64	15.14	
0214+1912			18.77	17.97	17.91	17.74	17.39	16.64	17.53	16.80	16.79	16.41	14.97	13.98	
0300+0222		22.04	18.63	18.04	17.95	17.89	17.61	16.64	17.53	16.80	17.50	17.13	15.43	14.45	
0303+3838		•••	20.52	19.24	18.96	18.87	18.70	16.64	17.53	16.80	18.44	17.98	16.17	15.26	
Continued on next page															

# Table A.1: Multi-band Photometric Data (AB Magnitudes)

for the FLASHES Pilot Sample

Target	Target GALEX <sup>a</sup>			SE	OSS DR1	$2^b$		,	2MASS <sup>a</sup>	2	WISE <sup>d</sup>				FIRST <sup>e</sup>
SDSS/HS	FUV	NUV	u	g	r	i	Z	Н	J	Κ	W1	W2	W3	W4	1.4GHz
0321+4132			18.16	17.22	16.75	16.59	16.31	16.64	17.53	16.80	16.08	15.71	14.31	13.59	
0639+3819			21.43	20.36	20.34	20.09	19.69	16.64	17.53	16.80	19.18	19.33	16.92	15.32	
0730+4340		22.43	20.52	19.19	19.06	19.00	18.87	16.20	17.17	16.79	18.67	18.41	16.89	15.70	
0735+3744			20.32	18.68	18.56	18.34	18.13				17.86	17.59	15.78	15.40	
0822+1626	19.79	19.45	18.36	17.88	17.88	17.90	17.67	16.48	17.69	16.39	17.60	17.24	15.63	14.93	
0834+1238			18.95	18.17	18.02	17.94	17.82	17.09	18.00	17.34	17.78	17.40	15.49	14.28	
0837+1459			18.4	17.74	17.74	17.72	17.44				17.24	16.86	15.17	14.41	
0851+3148			22.58	21.32	21.64	21.60	21.47	15.46	16.68	15.82	20.66	18.96	14.97	13.41	
0958+4703	20.99	21.63	18.5	17.73	17.73	17.65	17.35		•••		17.31	17.19	15.88	14.84	
1002+2008			20.01	19.09	18.94	18.85	18.64		•••		18.50	18.16	15.29	13.31	
1011+2941			16.76	16.17	16.09	16.02	15.90				15.87	15.64	14.21	13.45	
1112+1521			19.58	18.10	17.96	17.82	17.58		•••		17.23	17.09	16.43	15.12	
1218+2414			17.46	16.91	16.97	16.94	16.72	16.34	17.68	16.51	16.70	16.40	14.59	13.85	
1258+2123			22.27	21.15	21.33	21.50	20.88				20.54	19.46	15.81	14.54	
1428+2336			20.11	18.82	18.58	18.44	18.39	16.78	17.82	16.87	18.26	17.86	16.06	15.22	
1532+3059			17.9	17.25	17.17	17.14	16.98				16.86	16.55	15.20	14.61	
1552+1757			23.78	21.55	21.31	21.31	20.76	15.28	16.26	15.73	19.06	19.03	17.58	15.19	
1625+4858		22.18	19.52	18.09	17.94	17.63	17.41	15.79	16.57	16.15	17.25	17.11	15.94	15.32	12.86
1625+4858		22.18	19.52	18.09	17.94	17.63	17.41	15.79	16.57	16.15	17.25	17.11	15.94	15.32	12.86
2151+0921			18.96	18.42	18.38	18.36	18.10	16.77	17.67	17.47	18.21	18.01	16.77	15.38	
2234+2637			23.59	22.03	21.50	21.00	20.41	16.04	17.17	16.18	20.35	20.21	16.87	15.59	
2241+1225			18.73	18.05	17.93	17.84	17.70				17.60	17.19	15.53	15.05	
2259+2326			19.02	18.26	18.11	17.99	17.65				17.33	16.96	15.40	14.49	
2328+0443			22.67	20.78	21.14	21.55	20.76	16.26	17.33	16.83	19.99	19.17	16.00	14.66	
2338+1504	21.3	21.66	18.19	17.68	17.63	17.50	17.22				16.99	16.69	15.49	14.98	12.27
2339+1901		22.3	18.12	17.20	17.12	17.00	16.59	16.64	17.53	16.80	16.04	15.89	14.96	14.24	•••
2340+2418			21.13	20.69	20.56	20.53	20.09	16.90	18.02	16.91	19.71	20.07	17.05	14.98	•••
2350+3135			22.94	21.02	20.67	20.82	20.65				19.92	20.56	17.34	15.33	•••
Continued on next page															

Target	GAL	.EX <sup>a</sup>		SD	$2^b$		2MASS <sup>c</sup>				FIRST <sup>e</sup>				
SDSS/HS	FUV	NUV	u	g	r	i	Z	Н	J	Κ	W1	W2	W3	W4	1.4GHz
(a) GALEX DR5 (Bianchi et al., 2011)															
(b) SDSS DR12 (Alam et al., 2015)															
(c) 2MASS Catalog (Skrutskie et al., 2006)															
(d) AllWISE Catalog (Cutri et al., 2013)															
(e) FIRST Survey (Helfand, White, and Becker, 2015)															