Development and Demonstration of an On-Detector Technique to Limit the Impact of Atmospheric Emission Lines on Near-Infrared Spectra

by

Theodore Augustus Grosson
B.Sc., William Marsh Rice University, 2021

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astronomy

Development and Demonstration of an On-Detector Technique to Limit the Impact of Atmospheric Emission Lines on Near-Infrared Spectra

by

Theodore Augustus Grosson
B.Sc., William Marsh Rice University, 2021

Supervisory Committee

Dr. Alan McConnachie, Co-Supervisor
(Department of Physics and Astronomy)

Dr. Kim Venn, Co-Supervisor
(Department of Physics and Astronomy)

Abstract

Observations in the near-infrared using large ground-based telescopes are limited by bright atmospheric emission lines, particularly the OH Meinel bands. These lines can saturate a spectrograph on the order of minutes, resulting in the loss of information at wavelengths containing the lines. OH lines also vary on the scale of minutes, so observations longer than this timescale cannot capture this variability. Both of these properties necessitate the use of short exposure times in order to perform accurate sky subtraction. To observe faint science targets, several short exposures must be coadded instead of taking a single long exposure. Because each exposure includes its own independent read noise, this results in an increase in the total noise of the coadded image. In this thesis I present a new method to achieve longer exposure times in near-infrared spectra without the saturation of these lines, while still preserving information about their variability so that sky subtraction can still be applied. This is accomplished by periodically resetting the pixels on an H2RG detector that contain bright lines while the rest of the detector continues integrating. This method is demonstrated on the McKellar Spectrograph, where we reset the emission lines from an arc lamp while still recording their flux. I show that, when comparing the resulting spectrum and its signal-to-noise to a more conventional observing mode, the only measurable systematic difference is a result of our imperfect setup and can be removed with a standard nonlinearity correction. This method does not have the drawbacks of other measures to mitigate the effects of OH lines, such as short exposure times or completely removing the information at the relevant wavelengths, and as such shows promise for potential future use at observatories. We advocate demonstrating this method on sky spectra at existing high-quality facilities in order to test its feasibility for use in sky subtraction schemes for premier modern spectrographs.

Table of Contents

Supervisory Committee ii

Abstract iii

Table of Contents iv

List of Tables vi

List of Figures vii

Acknowledgements ix

1 Introduction 1

1.1 Infrared Astronomical Spectroscopy 1
1.2 OH Emission Lines 3
1.3 Current Observing Strategies 5
    1.3.1 Sky Subtraction 5
    1.3.2 Sky Avoidance 8
1.4 Infrared Detectors 8
    1.4.1 Guide Windows 9
1.5 On-DEtector Sky line Suppression with infrared Arrays (ODESSA) 10
1.6 Outline 11

2 Experimental Development 13

2.1 Experimental Requirements 13
2.2 Hardware Development 14
    2.2.1 McKellar Spectrograph 15
    2.2.2 Detector 19
    2.2.3 Filters 20
    2.2.4 Cryostat 23

2.3 Firmware Development . . . 25
    2.3.1 Readout Philosophy . . . 25
    2.3.2 Changes to the ODGW Readout Mode . . . 27
2.4 Software Development . . . 29
    2.4.1 Image Acquisition . . . 29
    2.4.2 Image Calibration . . . 29
    2.4.3 Calculating Flux . . . 32
    2.4.4 Extracting Spectra . . . 34
    2.4.5 Coadding Images . . . 35
2.5 Detector Characterization . . . 35
    2.5.1 Electron Conversion Gain . . . 35
    2.5.2 Readout Noise . . . 38
    2.5.3 Temperature Dependence . . . 39
    2.5.4 Persistence . . . 43
2.6 Summary . . . 44

3 Demonstration and Analysis of On-Detector Line Suppression 45
3.1 Operational and Demonstration Requirements . . . 45
3.2 Engineering Data Collection . . . 47
    3.2.1 December 2023: Physical Setup Test . . . 47
    3.2.2 January 2024: Preliminary Demonstration . . . 49
    3.2.3 February 2024: Dark Testing . . . 53
3.3 Final Data Collection . . . 55
3.4 Results . . . 58
    3.4.1 Signal-to-Noise . . . 59
    3.4.2 Systematics Introduced by Readout Mode . . . 62
    3.4.3 Other Limitations . . . 66
3.5 Summary . . . 67

4 Conclusions and Next Steps 68
4.1 Comparison to Existing Strategies . . . 68
4.2 Future Work . . . 71
4.3 Summary . . . 72

Bibliography 73

List of Tables

Table 2.1 Gratings at the McKellar Spectrograph. 16

Table 3.1 Observations for January 2024. 52
Table 3.2 Observations for February 2024. 54
Table 3.3 Observations for April 2024. 56

List of Figures

1.1 Some science cases for NIR spectroscopy. 2
1.2 OH atmospheric emission lines. 4
1.3 GNIRS long slit spectrum. 6
1.4 Sky fiber configurations. 7
1.5 ODESSA process. 11

2.1 Layout of the McKellar Spectrograph. 15
2.2 The 32-in arm of the McKellar Spectrograph. 17
2.3 McKellar grating efficiencies. 18
2.4 Observed night sky spectrum at the McKellar Spectrograph. 19
2.5 Argon emission lines. 20
2.6 The H2RG mounted in its puck. 21
2.7 Argon emission lines and nominal/acceptable FWHM of the narrowband filter. 22
2.8 Filter transmission spectra. 23
2.9 Schematic of the cryostat. 24
2.10 Detector installed in the cryostat. 24
2.11 My hardware contributions. 25
2.12 Illustration of ODGW readout. 26
2.13 Clocking scheme for ODGW mode. 27
2.14 $1/f$ noise drift in reference pixels over time. 30
2.15 Reset and flux maps of a dark image, and their associated histograms. . . . 31
2.16 First and last iterations of reset clipping procedure. 32
2.17 Bad pixel mask. 33
2.18 Comparison of different ramp-fitting methods. 34
2.19 Region of dark image used for PTCs. 36
2.20 Sample PTCs for four different superpixels. 37
2.21 Histogram of measured superpixel gains. 37

2.22 Histogram of measured RMS read noise. 38
2.23 Image quality vs. detector temperature. 39
2.24 Cold plate temperature vs. pressure of LN2 tank. 40
2.25 Emission of 300K and and 85K blackbodies. 41
2.26 Detector and radiation shield temperatures and median background flux over time. 42
2.27 Median background flux vs. radiation shield temperature. 42
2.28 Persistence prevention experiment. 43
2.29 Measured persistence current. 44

3.1 ODESSA set up on the McKellar Spectrograph. 48
3.2 The entrance slit setup. 50
3.3 Images resulting from different diffusers. 51
3.4 Resetting only half the 1108 nm line. 51
3.5 Locations of windows and regions used in dark testing. 53
3.6 Relative and absolute change in median background flux. 54
3.7 Final spectra. 58
3.8 Comparison of spectra using no-window mode and BRR mode. 59
3.9 Same as Figure 3.8, but instead comparing the two portions of the half-window-reset mode. 60
3.10 Final SNR. 60
3.11 Same as Figure 3.8, but instead comparing SNR. 61
3.12 Comparison of SNR in the 1108 nm line. 61
3.13 Typical pixel nonlinearity. 62
3.14 (Lack of) nonlinearity in a pixel which is periodically reset. 63
3.15 The reset anomaly. 64
3.16 Change in background in dark images for January and April data. 65

4.1 Simulated observation of a H = 20 mags A0 star with GNIRS at 1.6 µm. . . 70

Acknowledgements

My most heartfelt thanks go to:

Dr. Alan McConnachie, for finding a space for me to pursue my interests, encouraging me in both the successes and failures, pushing me beyond my former boundaries, and providing everything that I could ask from an advisor.

Dr. Kim Venn, for the continual support and encouragement, and for the effort in helping me find new opportunities.

Ed Chapin, Tim Hardy, Rick Murowinski, and countless others at HAA for their expertise, and for taking the time to teach me invaluable new skills.

NTCO, especially Kim, for the professional opportunities I would never have experienced without its existence.

Family and friends, near and far, for taking care of me in the high moments and the low, and always being there when I need them.

I acknowledge and respect the Lək̓ʷəŋən-speaking peoples on whose traditional territories I work and live, and the Songhees, Esquimalt, and WSÁNEĆ peoples whose historical relationships with the land continue to this day.

Chapter 1

Introduction

1.1 Infrared Astronomical Spectroscopy

Astronomical observations are completely dependent on the tools and methods used to carry them out. The goal of astronomical instrumentation development is to continuously improve the tools being used and the methodologies being applied, reducing the random and systematic uncertainties in the data we receive to refine our understanding of the physical processes which govern the objects we observe. Of the many different ways of using light to reveal this information, one which has seen significant development in recent years is infrared (IR) spectroscopy. The *James Webb Space Telescope* (*JWST*) and *Euclid* have recently begun providing observations with the Near Infrared Spectrograph (NIRSpec) (Jakobsen et al., 2022) and Near-Infrared Spectrometer and Photometer (NISP) (Euclid Collaboration et al., 2024), and several new ground-based infrared spectrographs are in various stages of development, including the Gemini Infrared Multi-Object Spectrograph (GIRMOS) (Sivanandam et al., 2018), the Maunakea Spectroscopic Explorer (MSE) (McConnachie et al., 2016), the Very Large Telescope (VLT) Multi-Object Optical and Near-infrared Spectrograph (MOONS) (Taylor et al., 2018), the Extremely Large Telescope (ELT) High Angular Resolution Monolithic Optical and Near-infrared Integral field spectrograph (HARMONI) (Thatte et al., 2010), and the Thirty Meter Telescope (TMT) InfraRed Imaging Spectrograph (IRIS) (Larkin et al., 2020).

These instruments are designed for a wide variety of scientific use cases, everywhere from the distant universe to the solar neighbourhood. Studies of the high-redshift universe necessarily require infrared observations, both for understanding the spectra of early galaxies and measuring accurate redshifts (Barton et al., 2010; Jakobsen et al., 2022; Birkmann et al., 2022; Euclid Collaboration et al., 2024). The ELT and TMT especially intend to break down the kinematic and chemical structure of galaxies at the peak epoch of star

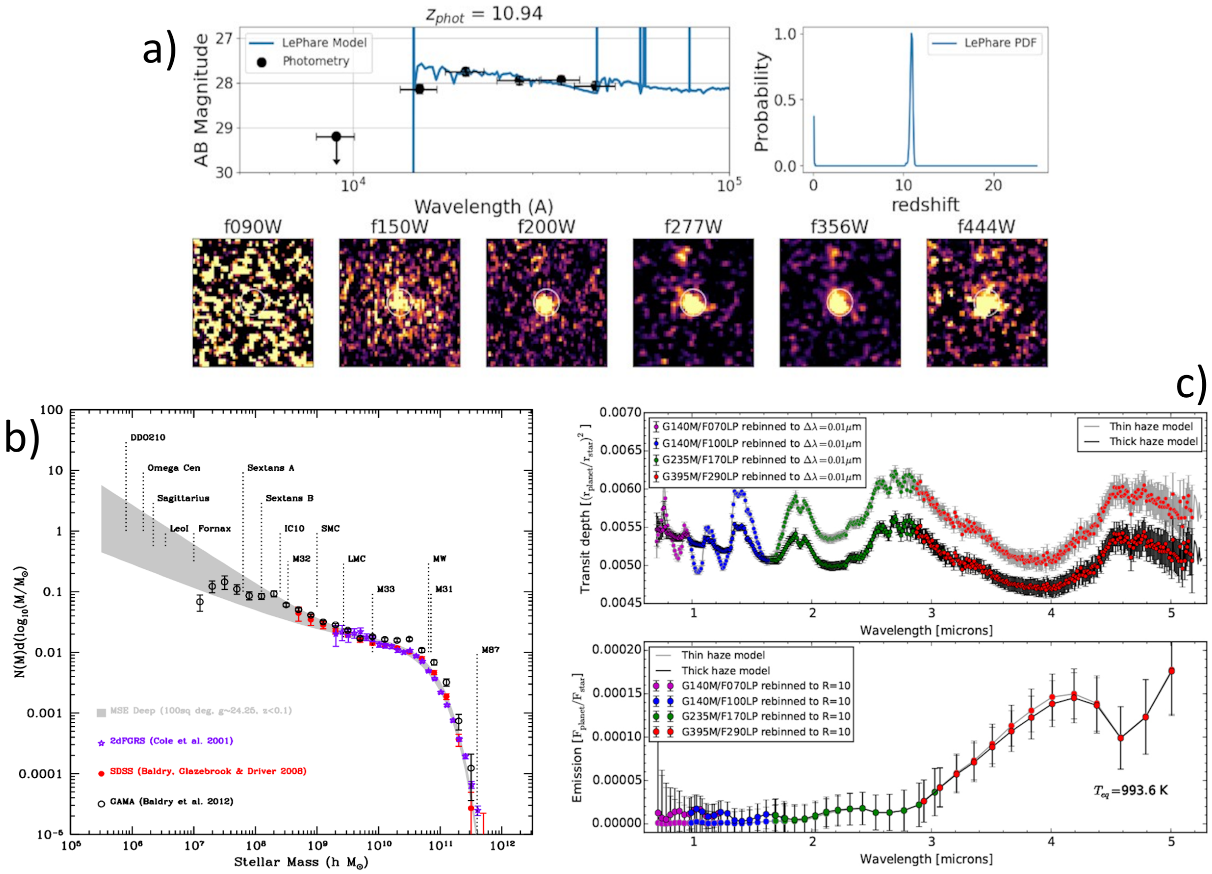

Figure 1.1: Some science cases for NIR spectroscopy. a) *JWST* photometric redshift identification of a galaxy near $z = 11$. Spectroscopy is necessary to verify the redshift. b) Stellar mass function of galaxies in the nearby universe. MSE will extend the low-mass end of observations. c) Simulated *JWST* transit and emission spectroscopy of HAT-P-26b. Images from Adams et al. (2023); McConnachie et al. (2016); Birkmann et al. (2022), respectively.

formation ($z \lesssim 5$) (Barton et al., 2010; Thatte et al., 2010). In the local universe, spectroscopy can tell the story of galactic evolution by revealing the chemistry of resolved stellar populations (Barton et al., 2010; McConnachie et al., 2016; Taylor et al., 2018). Within the Milky Way, high-resolution spectra provide precise stellar parameters for individual stars, as well as the properties of young stars in dust-filled environments (McConnachie et al., 2016; Taylor et al., 2018; Böker et al., 2022). Stellar radial velocities, transits, and direct spectroscopy also enable the discovery and characterization of individual exoplanets (Barton et al., 2010; Thatte et al., 2010; Birkmann et al., 2022).

The majority of these instruments are being built for large, ground-based telescopes with apertures ranging from 8.1 m (Gemini) to 39.3 m (ELT). The advantages of these over space-based telescopes include the relaxed constraints on physical space for the telescope, the relative flexibility and maintainability of the instruments, and the shorter timescales required for development (Böker et al., 2022). These advantages allow ground-based telescopes to be much larger and implement more advanced technology than space-based counterparts. This is a driving factor for the TMT and ELT, for example, whose extremely large apertures allow for the high spatial resolution necessary to achieve their science goals (Barton et al., 2010; Thatte et al., 2010).

While ground-based telescopes are often more versatile than space-based telescopes, they have some disadvantages. One major disadvantage is the necessity of looking through the earth's atmosphere. Turbulence from the atmosphere distorts the light passing through, preventing large telescopes from reaching their diffraction-limited resolutions if no correction is applied. In fact, most of the instruments mentioned rely heavily on new, advanced adaptive optics facilities that are being built alongside the science instruments (Allington-Smith et al., 2006; Thatte et al., 2010; Larkin et al., 2020). It is this key technology which allows ground-based telescopes to take full advantage of the high-resolution diffraction limit offered by large apertures.

In addition, ground-based telescopes face higher backgrounds than space-based telescopes, originating from warmer telescope components as well as contributions from the atmosphere itself. Some of this background can be mitigated by placing the instruments within cryogenic systems, but the various background sources imposed by the atmosphere are difficult to remove. One source of light which is especially relevant to the instruments mentioned is the infrared emission of OH. This molecule covers the near-infrared with hundreds of bright emission lines, making it difficult to distinguish target spectra within the sky-dominated images. The purpose of this thesis is to develop and demonstrate a new strategy to reduce the OH lines' impact, allowing higher-precision observations for accomplishing the science goals outlined here.

1.2 OH Emission Lines

OH emission lines, shown in Figure 1.2, were initially identified by Meinel (1950) as rotation-vibration bands of the OH molecule in the upper atmosphere. When they are formed from ozone and atomic hydrogen, OH molecules are given enough energy to be excited into elevated rotational and vibrational energy levels. Decays between these levels result in the dense emission spectrum observed (Herzberg, 1951). The smaller energy differences of rotational transitions (compared to vibrational transitions) result in the distinct

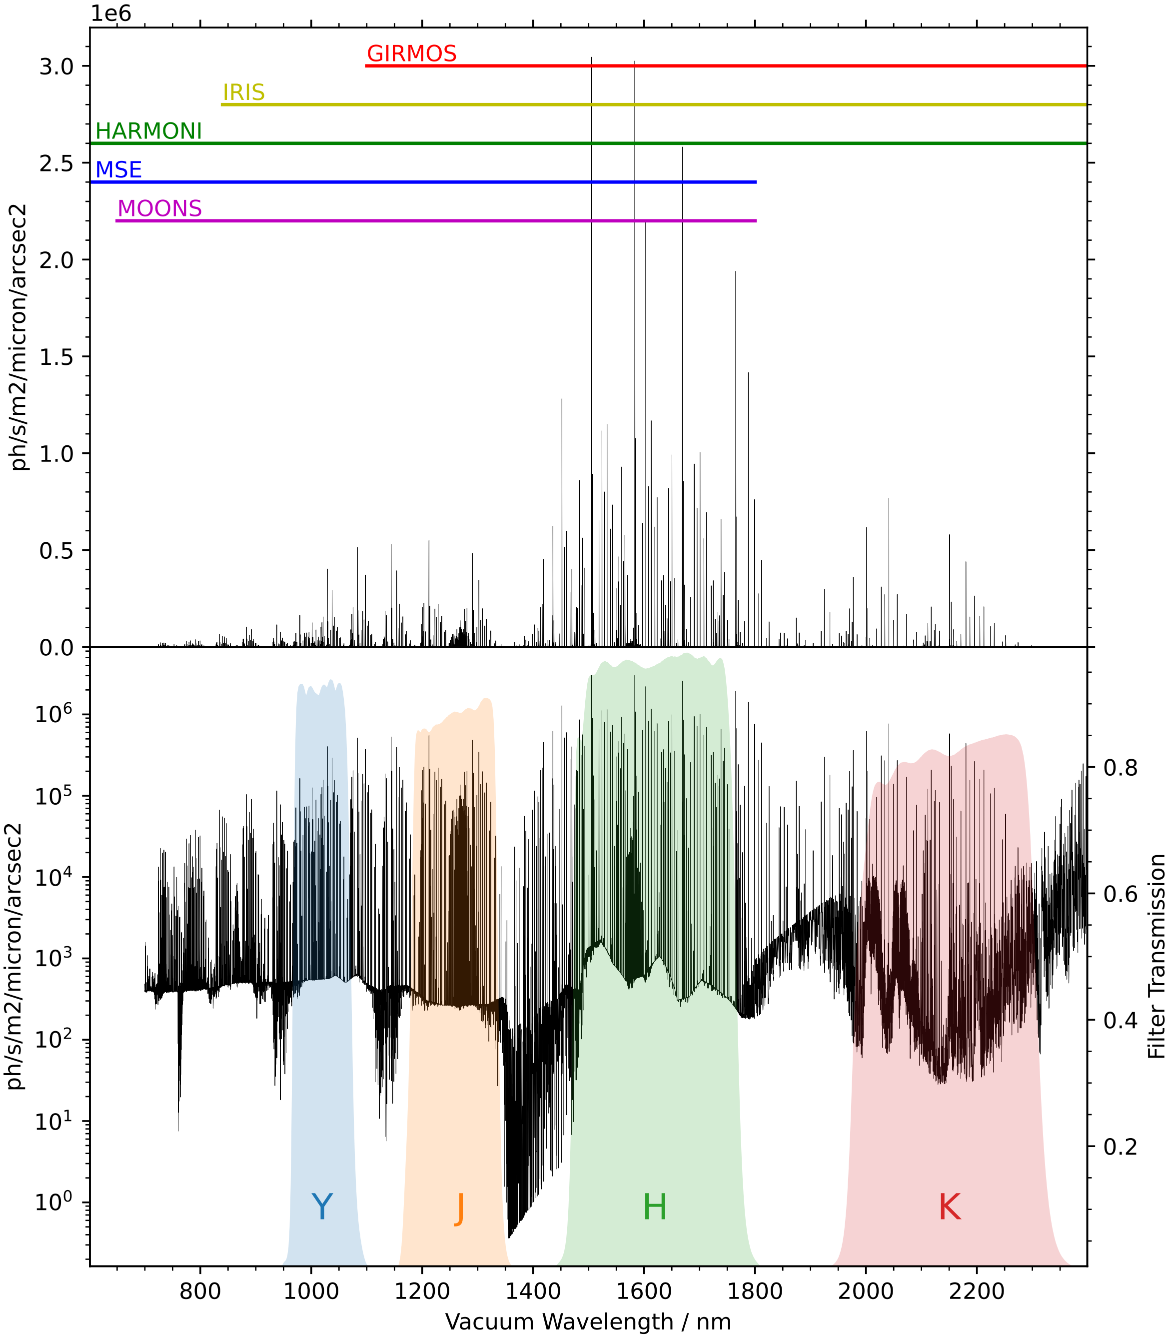

Figure 1.2: OH atmospheric emission lines, calculated with ESO Skycalc[1]. The top panel shows the wavelength ranges of selected NIR spectrographs. The bottom panel shows the total sky background as seen at the VLT along with the transmission of selected VLT HAWK-I filters[2] representative of typical NIR filters.

---

[1] `https://www.eso.org/observing/etc/skycalc`

[2] `http://svo2.cab.inta-csic.es/svo/theory/fps/`

bands seen in the spectrum; a band generally consists of one vibrational transition with a number of rotational transitions altering the total change in energy between states (Rousselot et al., 2000).

The OH bands span a large portion of the near-infrared, from 0.7–2.4 µm, with the brightest around 1.6 µm in the H-band, well within the wavelength ranges of many of the upcoming NIR spectrographs. The larger telescopes plan to target sources in the H-band fainter than 25 mags, whereas the sky continuum between the OH lines is around 19–20 mags in J- and H-bands, and the lines themselves up to 10^3 times brighter still (Puech et al., 2014; Oliva & Origlia, 1992).

In addition to their high intrinsic brightness, the emission lines vary in brightness over time and across the sky. Gravity waves propagating through the atmosphere produce periodicity of 5–15 minutes with spatial wavelengths ranging from 5–40° and variations in amplitude around 15%. However, the relative strengths of individual lines will not be impacted identically by gravity waves. Line strengths will typically increase or decrease at the same time, but line ratios do not in general remain constant. Lines lying within the same vibrational band, though, do tend to be fairly well-correlated with each other (Ramsay et al., 1992; Davies, 2007; Ellis & Bland-Hawthorn, 2008). Additional low-frequency variation arises from diurnal changes in the atmospheric composition, leading to the total sky brightness decreasing by as much as 50% over the course of a night (Ramsay et al., 1992; Moreels et al., 2008; High et al., 2010). The behaviour of the OH emission paints a complicated picture that necessitates observing strategies which account for the lines' variability.

1.3 Current Observing Strategies

1.3.1 Sky Subtraction

Though OH emission lines can be very bright, they are also very narrow, with intrinsic widths around 0.1 Å (Maihara et al., 1993b; Ellis & Bland-Hawthorn, 2008). This allows spectrographs with sufficient resolution to observe target spectra in the (relatively) faint inter-line continuum. However, if we hope to obtain any useful information from faint targets, it is necessary to come up with precise sky subtraction strategies to remove the sky spectrum. For targets around H=25, over 100 times fainter than the sky continuum, the sky subtraction needs to be accurate to within 1% (Puech et al., 2014).

Slit spectrographs have a relatively straightforward method of performing sky subtraction. Placing the target in the centre of a long slit will result in the sky spectrum being

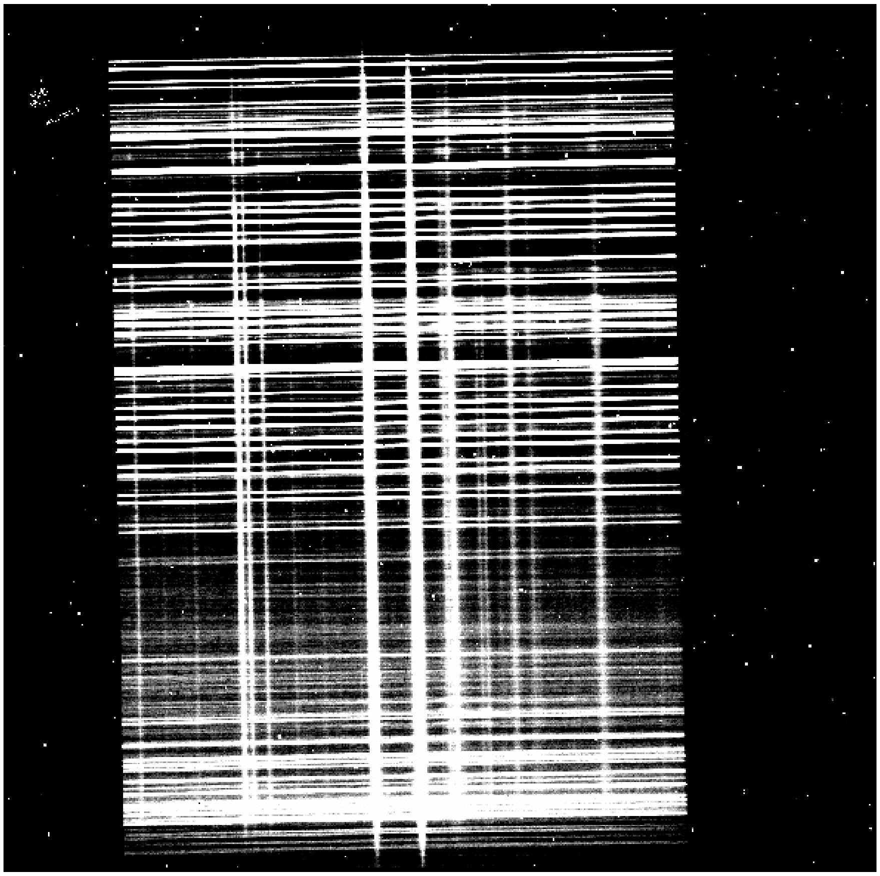

Figure 1.3: Long slit K-band spectrum taken with the Gemini Near-InfraRed Spectrograph (GNIRS). The vertical bands are stellar spectra, and the horizontal bands are sky lines. The sky lines are consistent across the full extent of the slit, allowing straightforward sky subtraction. Image from Gemini Observatory[3].

sampled on either side of the target spectrum, as seen in Figure 1.3, providing simultaneous sky data for free with the target. Sky subtraction can ignore the variation of the OH lines in this case, since the changes affect both the target and sky spectrum the same (Davies, 2007).

Fiber-fed spectrographs require somewhat more effort since fibers need to be optimally positioned to sample the sky spectrum near target fibers. Several strategies for fiber placement have been proposed which are capable of low-noise sky subtraction (Rodrigues et al., 2012; Flagey et al., 2016). Some of these strategies involve placing fibers adjacent to each other to mimic slit spectroscopy. This allows near-identical sampling of the sky adjacent to the target, but can suffer from differing optical paths causing different instrumental signatures on the sky and target spectra (Sharp & Parkinson, 2010). Observing methods

---

[3] https://www.gemini.edu/instrumentation/gnirs/capability

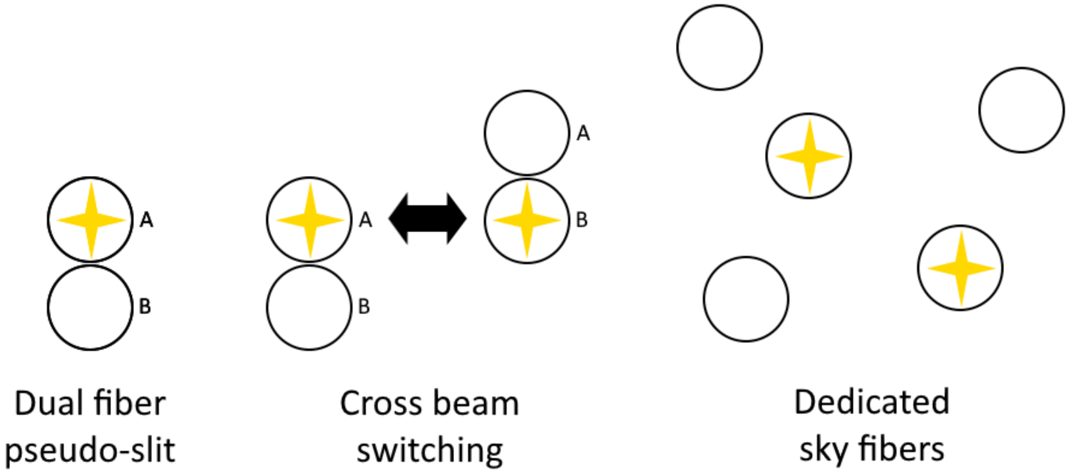

Figure 1.4: Some configurations of target (yellow star) and sky (empty) fibers. Pseudo-slit configurations mimic slit spectrographs, but each fiber has a separate instrumental PSF. Cross beam switching avoids this by switching paired fibers back and forth between the target and sky. Dedicated sky fibers require some method of reconstructing the sky at the location of the target.

such as cross beam switching solve this issue to a large extent and can provide excellent subtraction results (Rodrigues et al., 2012; Puech et al., 2014).

Some modern fiber-fed spectrographs such as MSE disfavour pseudo-slit modes, preferring instead to maximize the number of fibers on science targets (Flagey et al., 2016). This avoids the factor of two multiplexing loss associated with paired target and sky fibers, but requires more creative strategies to achieve precision sky subtraction. By placing a number of dedicated sky fibers across the field of view, one can devise methods to use those fibers to estimate the background at the location of the target. A low-order solution would be to take the average of the sky fibers and subtract directly as one might with a long slit, but this has the problems of smoothing over the OH variability and differing optical paths. Davies (2007) and Rodrigues et al. (2010), recognizing the independent variability of OH vibrational bands, improve this by independently scaling each band to best fit the observed spectrum at each spaxel. In the case of Rodrigues et al. (2010), the leftover continuum from sky fibers is also interpolated over the sky. Similarly, a PCA-based approach proposed by Kurtz & Mink (2000) and implemented by e.g. Wild & Hewett (2005); Sharp & Parkinson (2010); Soto et al. (2016) finds the basis spectra which can best subtract out the variable sky emission. These methods have been successful in producing relatively clean removal of the emission lines.

Regardless of the strategy employed, observations are fundamentally constrained by the OH emission: exposure lengths longer than a few minutes do not pick up the variation of the OH lines in the target spectra (Ramsay et al., 1992; Davies, 2007). In contrast, short exposure lengths become read noise limited, favouring longer exposures for faint targets. Even if this trade-off is made, exposure lengths for any spectrograph are limited by the bright lines anyway, which may saturate the detector within 10 minutes (Ramsay et al., 1992). Past saturation, all useful information within the affected region is lost, making subtraction of the lines impossible. Saturation also has the potential to impact nearby pixels due to bleeding, as well as subsequent observations due to persistence.

1.3.2 Sky Avoidance

In order to prevent the issues associated with saturation, there are also suggestions to avoid the wavelengths containing the emission lines from spectra altogether. Specially selected narrowband filters (FWHM < 1 nm) can probe the gaps between OH bands at the cost of removing most of the NIR spectrum, and are best suited for highly specific use cases such as the DAzLE instrument which is tuned for the detection of Lyman-α at $z > 7$ (Horton et al., 2004). High dispersion masking samples the continuum between individual lines by physically blocking the highly-dispersed light at predetermined positions, then recombining the light before sending it to the science instrument (Maihara et al., 1993a). Fiber Bragg gratings avoid the issue of grating scattering by using interference filters to remove specified wavelengths from the light prior to dispersion (Bland-Hawthorn et al., 2004). While these methods are all successful at preventing the emission lines from overwhelming light from the target, they all have the limitation of removing information about the target at each wavelength that is blocked.

1.4 Infrared Detectors

Assuming for the moment that we have a spectrograph which is capable of rapidly and simultaneously sampling target and sky spectra, we need to understand how the data are collected by detectors. The standard technology used for NIR sensing are HgCdTe hybrid detectors. In fact, all the instruments mentioned thus far which have selected a detector use the same family of detectors, the Teledyne HAWAII-xRG (Blank et al., 2011).

A key aspect of NIR detectors that distinguishes them from their ubiquitous optical counterparts, charge-coupled devices (CCDs), is the use of individual amplifiers for each pixel combined into a readout integrated circuit (ROIC), as opposed to one or a few amplifiers for the entire array. While CCDs transfer charge from one pixel to the next until they reach the readout amplifier, IR detectors measure the charge within each pixel without influencing that charge, allowing multiple samples of the same pixel. In addition, pixels can be sampled and/or reset in any order (Rieke, 2007). This allows for a much larger range of readout modes than is possible for CCDs.

For example, the standard method of measuring pixel charge in a CCD is with correlated double sampling (CDS). CDS involves measuring the charge within the pixel, flushing the charge, and measuring again; the difference between the two measured values is the total charge accumulated during the exposure (Fowler et al., 1989). Since each measurement is subject to intrinsic uncorrelated read noise from the electronics, it was realized that the non-destructive readout of IR detectors could be utilized to decrease the read noise relative to CDS. Fowler sampling, for example, consists of a series of consecutive reads at the beginning of an exposure and another series at the end. The difference of the averages of the two groups is the accumulated charge (Fowler & Gatley, 1990). Recognizing that total read noise is minimized by maximizing the number of samples, pixels can also be continuously sampled for the duration of the integration, a strategy commonly referred to as up-the-ramp (UTR) sampling (Fixsen et al., 2000). This strategy has been adopted for telescopes such as JWST, and this thesis focuses on UTR measurements.

1.4.1 Guide Windows

In our opinion, an underutilized ability of IR detectors is the ability to selectively read or reset individual pixels in arbitrary order. One implementation of this ability is "guide window" mode in HxRG detectors, in which a small subarray of pixels can be addressed independently from the full frame, allowing fast monitoring of fluence within that subarray (Blank et al., 2011). While the performance of these windows has been characterized within some specific contexts (e.g., Bezawada & Ives, 2006; Smith & Hale, 2012; Chapin et al., 2022), few science instruments have incorporated them into their observing modes. The Canada-France-Hawaii Telescope's (CFHT) Wide field Infrared Camera (WIRCam) and the Gemini South Adaptive Optics Imager (GSAOI) have implemented windows for tracking and tip/tilt corrections on guide stars, requiring minimum window speeds of a few Hz (Teeple et al., 2006; Young et al., 2012). The Carnegie Astrometric Planet Search Cameras (CAPSCam) use the reset capability to prevent saturation of bright guide stars at up to 5 Hz (Boss et al., 2009), but few other examples exist.

Perhaps a reason for this lack of enthusiasm is the known systematic effects which guide windows can produce. A well-documented issue with operating windows is the introduction of video offsets along the rows and columns of windows, though this can easily be accounted for and removed in images (Bezawada & Ives, 2006; Smith & Hale, 2012; Chapin et al., 2022). Less easy to calibrate for are the effects of multiplexer glow and self-heating. Frequent visits to small areas of the ROIC introduce additional charge to the corresponding pixels, and they also heat the area, increasing localized dark current (Tam & Hu, 1984; Smith & Hale, 2012; Regan & Bergeron, 2020). Chapin et al. (2022) note that this increased flux is also capable of spreading along the columns of windows. However, there have been few reports on how these effects impact science observations.

1.5 On-DEtector Sky line Suppression with infrared Arrays (ODESSA)

This thesis proposes a new use for guide windows in the context of OH lines in near-infrared spectra. Recognizing the advantage of increasing exposure lengths to decrease the need for stacking frames, and consequently reducing the impact of read noise on observations of faint science targets, it would be beneficial to devise a method to enable longer exposure times for near infrared spectroscopy. This method should prevent the images, or sections of the images, from being rendered useless by bright, variable emission lines. Firstly, for exposures longer than a few minutes, such a method needs to prevent the saturation of these lines, which results in the loss of all information at their locations and can cause bleeding and persistence in the detector, deteriorating the quality of observations. Secondly, it needs to be able to monitor the variations of lines on the order of minutes in order to accurately subtract them from target spectra.

The concept of On-DEtector Sky line Suppression with infrared Arrays (ODESSA) introduced here is to observe spectra with a detector capable of guide window mode. By placing a window on top of bright emission lines and periodically resetting those pixels before they have time to saturate, while continuing integration in the rest of the detector, individual sky lines at any position on the detector can be prevented from saturating while still achieving long exposure times. In addition, the charge in these window pixels is still measured between resets, so no information is lost in the wavelength region contained by the window, and the full spectrum can be reconstructed at the end of the exposure. With carefully selected reset rates, the change in flux across different sub-exposures can be measured to track variability in the emission lines. An illustration of this process using data which will be obtained in Chapter 3 is shown in Figure 1.5.

As an initial demonstration of this concept, this thesis aims to show that resetting guide

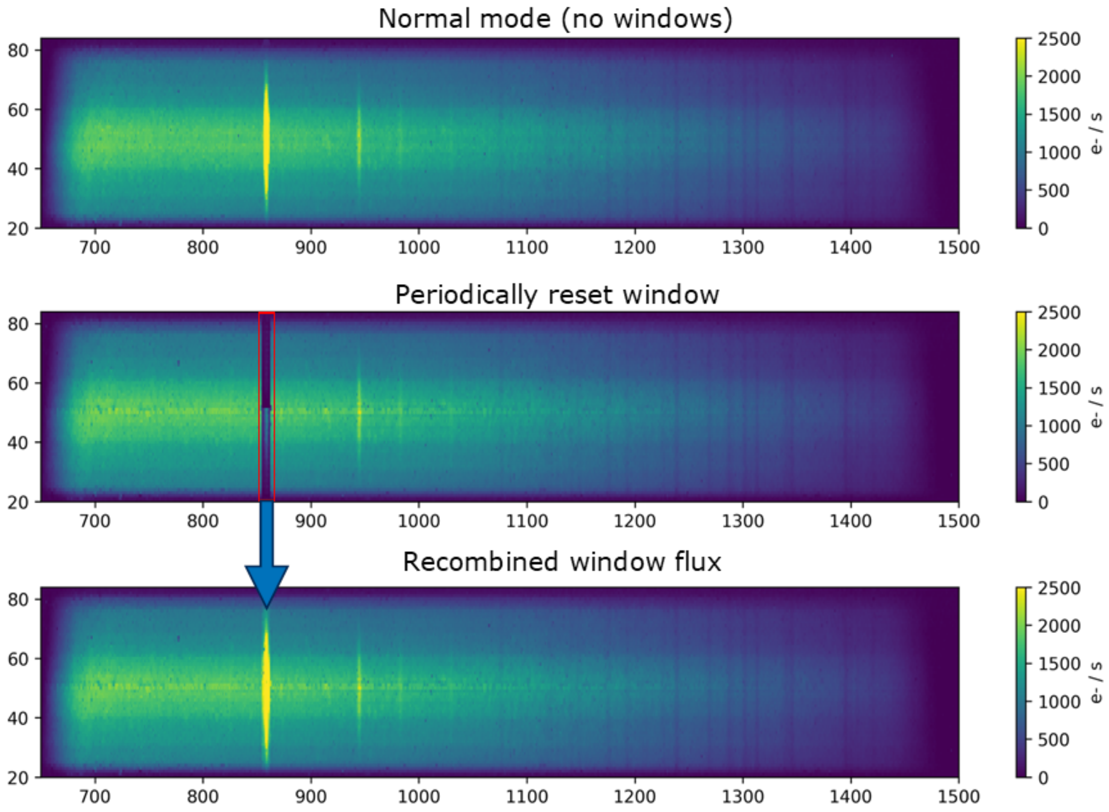

Figure 1.5: ODESSA process. The top image is a spectrum with a bright emission line observed using a normal, no-window scheme. The middle image shows the placement of a window (outlined in red) over the line in the same spectrum which is periodically reset; the data from this window is then recombined back into the image to create the bottom image. If the concept works ideally, the top and bottom images should be identical. Details on how these images were obtained will be described later in the thesis.

windows can be effectively used to prevent bright emission lines from reaching saturation without significant adverse effects on the final spectrum. If this is the case, future work will explore the concept's usefulness in performing precise sky subtraction. If successful, this method could enable significantly longer exposure times for NIR spectra, reducing uncertainty from read noise and improving the precision of observations, ultimately advancing our understanding of astrophysical phenomena.

1.6 Outline

This thesis is organized as follows: Chapter 1 has introduced the motivation behind the work and background for the major concepts. Chapter 2 presents the laboratory development necessary for ODESSA to be implemented, including hardware, firmware, and software development, as well as characterization of the physical system. Chapter 3 details the demonstration of the concept using an on-telescope setup and presents an analysis of the main results. Chapter 4 presents conclusions about the practicality of the ODESSA concept based on the experiments carried out, proposes some future improvements, and summarizes the thesis.

A major portion of this thesis pertains to work done leading up to and including a conceptual demonstration of the ODESSA system. By its nature, this included a wide variety of tasks, from image analysis software coding to mechanical engineering design. As such, this work was highly collaborative, and several tasks were carried out by colleagues at the National Research Council (NRC) Herzberg Astronomy and Astrophysics Research Centre (HAA). These include Ed Chapin, Research Council Officer, Software; Tim Hardy, Research Council Officer, Electronics; Jordan Lothrop, Mechanical Technologist; Rick Murowinski; Mustafa Wasif, co-op student; and Matthew Winn, co-op student. Throughout the thesis, I note in the text wherever tasks were performed by individuals other than myself.

Chapter 2

Experimental Development

This chapter provides a detailed account of the development necessary before a demonstration of ODESSA could be carried out. Working with the hardware available led to restrictions on the scope of the experiment, including using artificial spectra instead of sky spectra and having a bright background which needed to be worked around. Custom software was written to obtain and analyze images, and basic characterization of the detector system was carried out. Although some flaws limit the capabilities, we find that our setup is sufficient to carry out the demonstration.

2.1 Experimental Requirements

The successful demonstration of the ODESSA technique to limit atmospheric emission lines in spectra necessarily requires the use of a spectrograph mounted on a telescope. The Dominion Astrophysical Observatory's (DAO) 1.2-m telescope in Victoria, BC hosts the McKellar Spectrograph, a Coudé-focus spectrograph which has been in operation since 1962. This spectrograph is commonly used to observe medium- to high-resolution stellar spectra in visible wavelengths, but has also been used as a testbench for experiments at HAA (e.g., Andersen et al., 2008; Jackson et al., 2024, in press), making it a good candidate for our own experiment.

The detector typically used at the McKellar Spectrograph is a helium-cooled Scientific Imaging Technologies (SITe-4) CCD, which can record low-noise data at blue–visible wavelengths. However, our primary intent of developing the ODESSA method is to apply it at red–near infrared wavelengths, where OH lines are especially strong. HAA owns three "spare" Teledyne HAWAII-2RG (H2RG) HgCdTe infrared detectors, which are capable of the UTR sampling critical for the ODESSA concept and one of which was previously used for the HAA On-Detector Guide Window project of Chapin et al. (2022). Therefore, as a starting point for our demonstration, we paired one of these detectors with the McKellar Spectrograph, using it to observe at wavelengths which contain the OH lines.

The McKellar Spectrograph has a selection of mostly visible-wavelength gratings with different wavelength ranges and resolutions (Richardson, 1968). OH lines become strong at wavelengths longer than about 1 µm, and H2RG detectors are sensitive to wavelengths longer than about 800 nm, so the selection of grating and spectral lines needs to match both of these constraints, i.e., we must select sky lines which lie in the overlap between the detector sensitivity and the wavelength range of the grating we select.

Background emission at NIR wavelengths is a major source of noise. Therefore, it is important for the way in which the detector is mounted and operated on the spectrograph to minimize noise caused by the bright IR light emitted by the ∼300K Coudé room. Sensible methods to suppress the impact of this background radiation include cooling the detector in a cryostat and including filters which block longer-wavelength IR light. These must be selected to fit the constraints imposed by the physical layout of the spectrograph as well as the wavelengths we choose to observe. As will be seen in Section 2.5.3, the selections we make do not perform perfectly, so our method of taking and analyzing images also needs to account for a non-negligible amount of IR background. This includes optimization of the operation of the detector itself, the format of images taken during the experiment, and the data reduction process.

These initial considerations have already identified important areas of the experiment requiring careful thought, namely:

- the spectrograph and detector setup needed to observe the emission lines we select,
- selection of the lines themselves,
- minimization of background radiation,
- detector operation, and
- image analysis.

This chapter addresses each issue in turn.

2.2 Hardware Development

With the above constraints in mind, the final ODESSA configuration includes one of the available H2RGs in a liquid nitrogen-cooled cryostat, controlled by an Astronomical Research Cameras (ARC) Generation 4 controller and mounted on the 32-in arm of the

McKellar Spectrograph. This section describes each hardware component in detail, as well as rationale for the choices made and some measurements that demonstrate sufficient performance of the system.

2.2.1 McKellar Spectrograph

The McKellar Spectrograph, commissioned in 1962, is situated at the Coudé focus of the 1.2-m telescope at the DAO. Light from the primary mirror is redirected through a slit and into the spectrograph room as shown in the floor plan in Figure 2.1. Here, a choice of gratings, listed in Table 2.1, disperses the light to the detector. The plate holders labelled in Figure 2.1 have since been replaced with cryostat mounts for digital detectors. One of these mounts, as well as some other components of the spectrograph, are shown in the photograph in Figure 2.2. In its more than 60 years of operation, the McKellar Spectrograph has produced high-resolution stellar spectra that have been especially useful

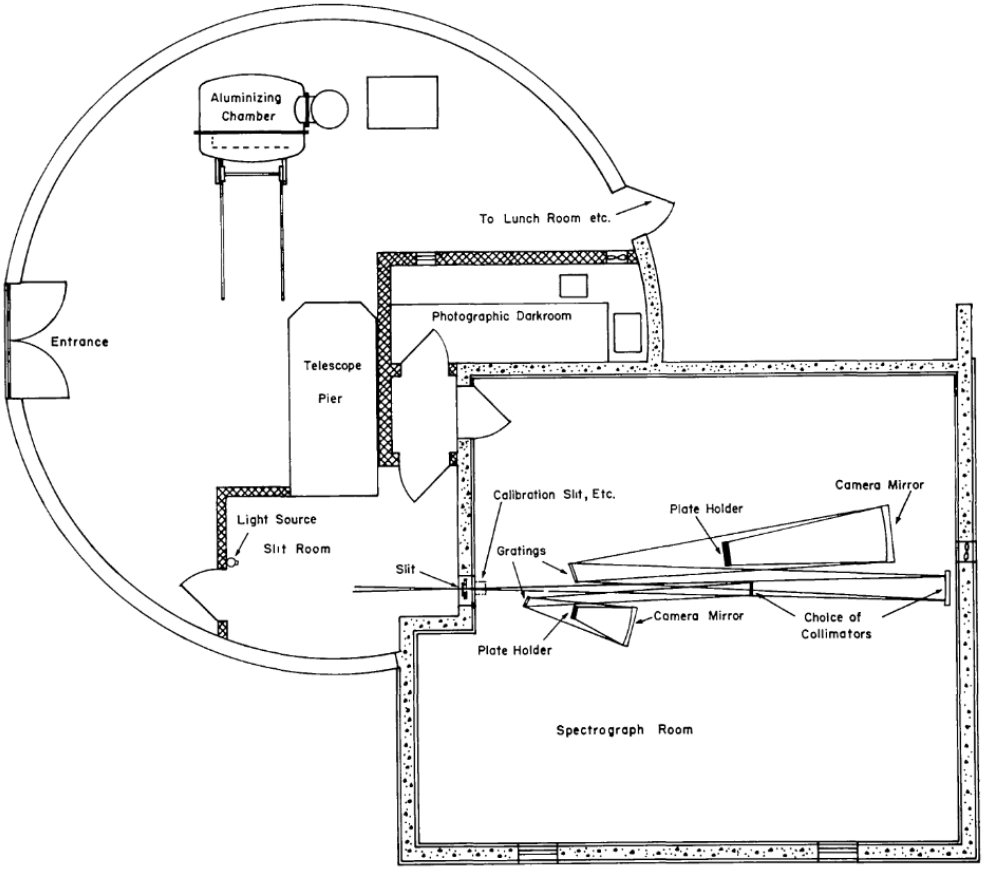

Figure 2.1: Layout of the McKellar Spectrograph. Figure from Richardson (1968).

Name	Lines/mm	Blaze (Å)	Range (Å)	Dispersion (Å/mm)
9682M	830	4060	3500-5800	2.4
9681M	830	8125	5800-9900	4.8
32121	1200	5000	3600-6500	10.1
32121H	1200	6000	5000-8500	10.1
32122	1200	3000	3200-4000	5.0
3282	830	4060	3500-6000	6.5
3281	830	8125	6000-9900	13.1
3261	600	4000	3200-5000	17.9
3261H	**600**	**7000**	**5000-9500**	**17.9**
3262	600	3500	3200-4200	9.0
3231	300	4200	3500-8000	40.9

Table 2.1: Gratings at the McKellar Spectrograph. Gratings available with the 92-in camera have names beginning with "92," and those available with the 32-in camera begin with "32." The grating we selected is in bold.

in spectral classification and radial velocity measurements (e.g., Batten, 1974; Aikman, 1976; Scarfe et al., 1990).

The McKellar Spectrograph has two camera arms: a 96-in focal length arm typically used for higher resolution spectroscopy, and a 32-in focal length arm for lower resolution spectroscopy. We prefer observations of the sky lines to have a low spectral resolution to constrain the emission lines to within a few pixels, so we use the 32-in arm. This choice allows us to use a diffraction grating with 600 lines mm^{-1} (R$\sim$20,000) and a blaze wavelength of 700 nm, the reddest available with lower resolution ("3261H").

Grating Efficiency

The available gratings operate mostly at optical wavelengths, and their performance at the reddest wavelengths is not well characterized. Before proceeding further, we wanted to ensure that the grating we selected had high enough throughput around 1 µm to provide useful data at these wavelengths. To do this, we measured the wavelength-dependent efficiency of this grating. We also made equivalent measurements for the reddest 96-in grating (830 lines mm^{-1}, blaze 813 nm).

The relative efficiencies of the gratings were measured by using the McKellar Spectrograph equipped with the SITe-4 CCD. We imaged a broadband light source located outside the entrance slit of the spectrograph through three different narrowband filters for each grating: one each centred at 700, 850, and 1000 nm. The intensity of the source was also

Figure 2.2: The 32-in arm of the McKellar Spectrograph. Light from the collimator (left) is dispersed off a diffraction grating (not shown), reflected off the 32-in-focal length mirror (centre), then redirected onto the focal plane by a flat mirror on the detector mount (right).

measured at the entrance slit using a calibrated photodiode. For each wavelength and for both gratings, the CCD image was bias-subtracted and an approximate wavelength solution was applied to the spectra using the nominal central wavelength and dispersion settings of the spectrograph. The spectra were then divided by the nominal quantum efficiency (QE) curve of the CCD, which decreased from about 80% to 20% from 700 nm to 1 μm, then divided by the exposure time of the image. Because exposure lengths were all less than 10 seconds, we assumed negligible dark current. Finally, the total counts of the observed narrowband features were divided by the intensity measured by the photodiode, and the two resulting efficiency curves were normalized to be equal to one at the measured wavelength closest to the blaze wavelength of the grating. The efficiencies are shown in Figure 2.3. Although the efficiencies of both gratings drop toward 1 μm, the drop is less than a factor of 2 over the wavelength range examined. This gave us confidence that the efficiencies around 1 μm would remain high enough to provide useful data that could successfully demonstrate the ODESSA technique.

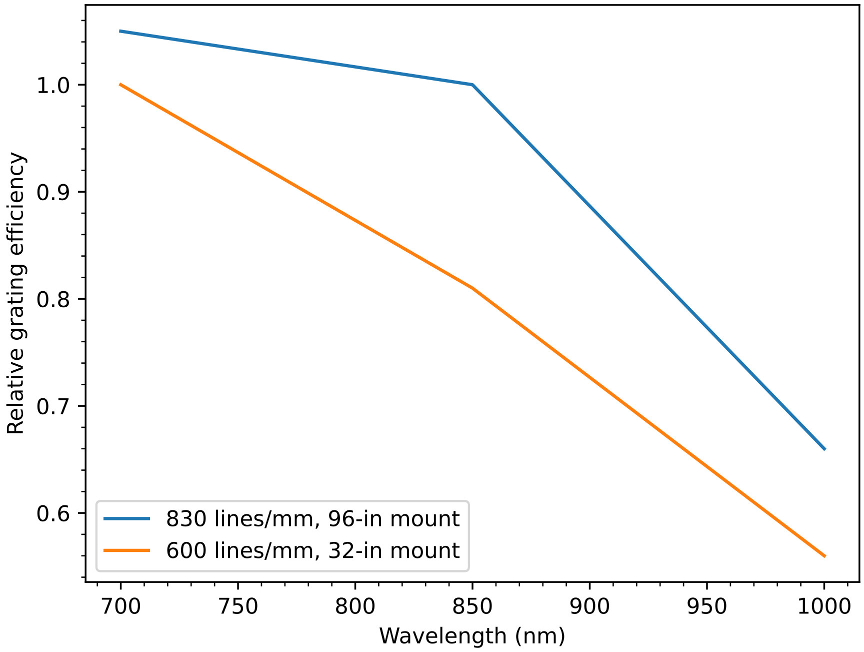

Figure 2.3: Efficiencies of two McKellar Spectrograph diffraction gratings, relative to their efficiencies at ∼blaze wavelength.

Sky Line Selection

To determine which sky lines might be best suited for a demonstration of ODESSA, we received a small allocation of telescope time during the night of 1 December 2022. We again used the SITe-4 CCD on the McKellar Spectrograph. We used the 600 line mm^{-1} grating on the 32-in arm of the spectrograph, centred at a wavelength of 935 nm, to take twelve 900-second exposures of blank sky.

The resulting coadded bias- and dark-subtracted spectrum is shown in Figure 2.4, with sky lines cross-matched to the catalogue of Hanuschik (2003) labelled. The blue spectrum is unbinned, and the orange spectrum shows binned data. Despite a combined 3 hours of exposure time, the sky lines are barely distinguishable above the background from the telescope. While the low QE of the CCD in this wavelength range likely plays a role (a problem the H2RG would not have), the relatively small aperture of the telescope and a degraded primary mirror coating limit the amount of light which makes it from the sky to the detector. As a result, optimizing the configuration of the spectrograph is not likely to result in significantly clearer lines. This is in stark contrast to Figure 1.2, which shows

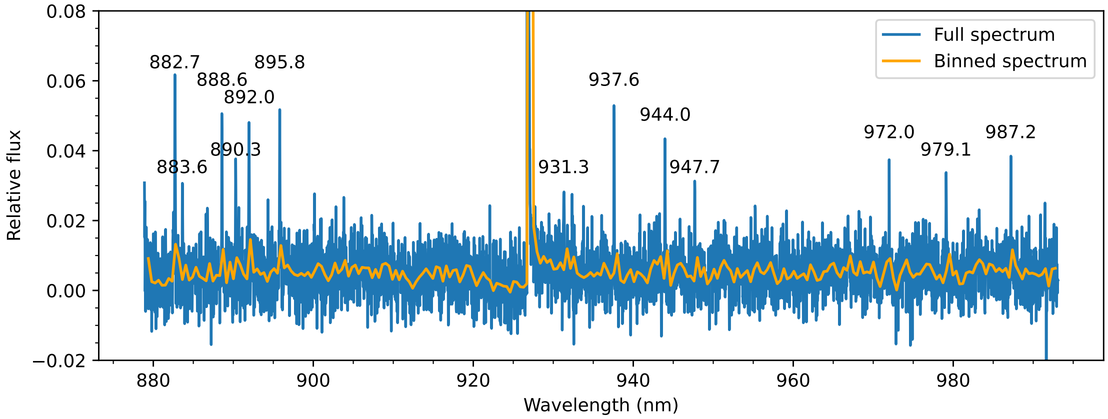

Figure 2.4: Observed night sky spectrum at the McKellar Spectrograph, with known sky lines labelled with their wavelengths. The large spike around 927 nm is a bad row on the CCD. The 15-px binned spectrum is also shown.

obvious, bright emission lines at these wavelengths at the VLT.

Given the very weak nature of the sky lines available to us, instead of attempting to observe these faint sky lines with the McKellar, we decided that a simpler and more effective route would be to use arc lamps to create artificial sky lines. ODESSA is fundamentally a method for resetting any emission line which may saturate the detector quickly, so we conclude that a demonstration with a simulated "sky" spectrum is just as useful as a demonstration with real sky spectra. The 1.2-m dome has an FeAr calibration lamp with strong Ar emission lines in the vicinity of 1 µm, as seen in Figure 2.5. Our choice of filter, presented later in Section 2.2.3, ultimately resulted in our selection of the 1108 nm and 1111 nm Ar lines to use as mock-OH lines.

2.2.2 Detector

Teledyne Technologies[1] produces a family of HgCdTe infrared detectors called HAWAII-xRG (HxRG). These are imagers that are optimized for low-light astronomical use and are ubiquitous in both ground- and space-based telescopes (Blank et al., 2011). They are also capable of a guide window mode, which will be discussed thoroughly in Section 2.3. In our setup, we use an H2RG sensitive from 0.8 to 5.3 µm, with a 2040x2040 pixel sensing area, four rows/columns of reference pixels along each edge, and an 18-µm pixel pitch.

Our choice of detector was guided by availability: HAA owns three different H2RGs which were part of the *JWST* Fine Guidance Sensor programme, including one flight spare.

---

[1] https://www.teledyne-si.com/

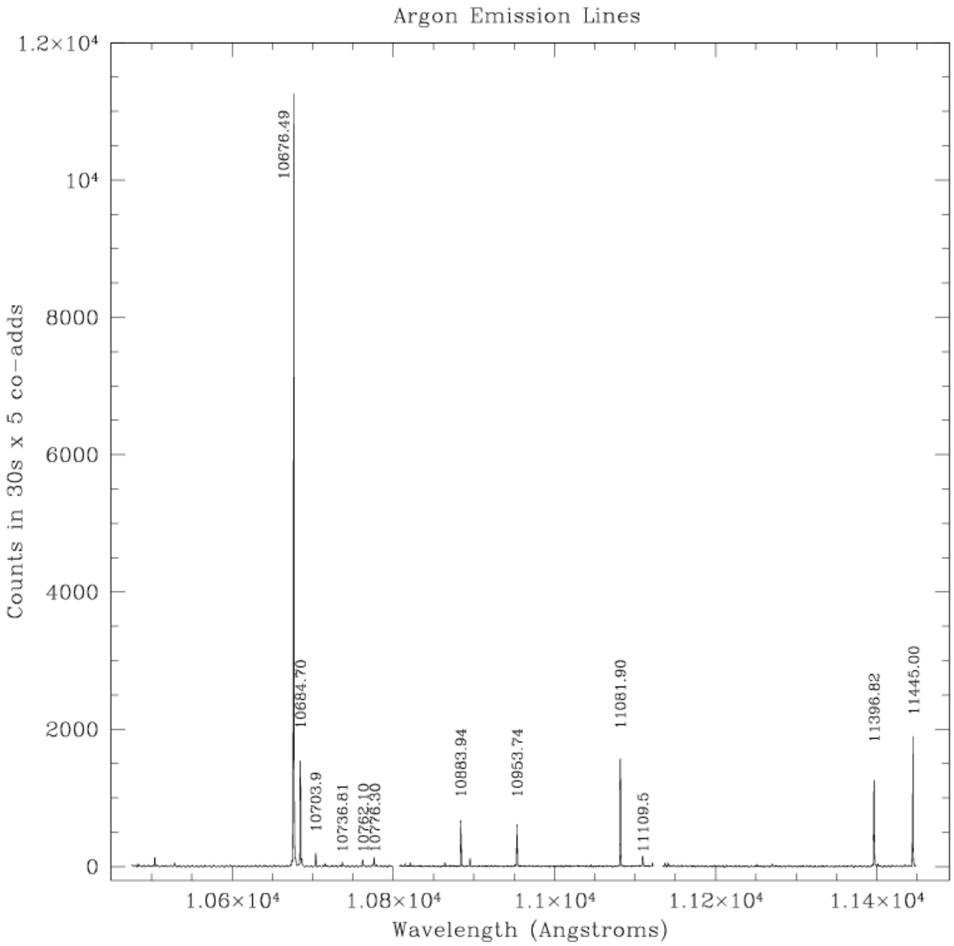

Figure 2.5: Some argon emission lines. Image from Gemini Observatory[2].

The one which was used in Chapin et al. (2022) had around 50% bad pixels at temperatures below 50K, while the other two had around 10% and < 1% bad pixels, respectively. As the third detector (the flight spare) was clearly superior to the other two, Hardy installed it in the *JWST* flight package ("puck") that came with the first detector and that would allow us to mount it in a cryostat. The detector and its puck are shown in Figure 2.6.

2.2.3 Filters

Infrared detectors are sensitive to blackbody emission at room temperature, so for our application it is necessary to include IR-blocking filters to prevent their immediate saturation

---

[2] https://www.gemini.edu/instrumentation/gnirs/calibrations

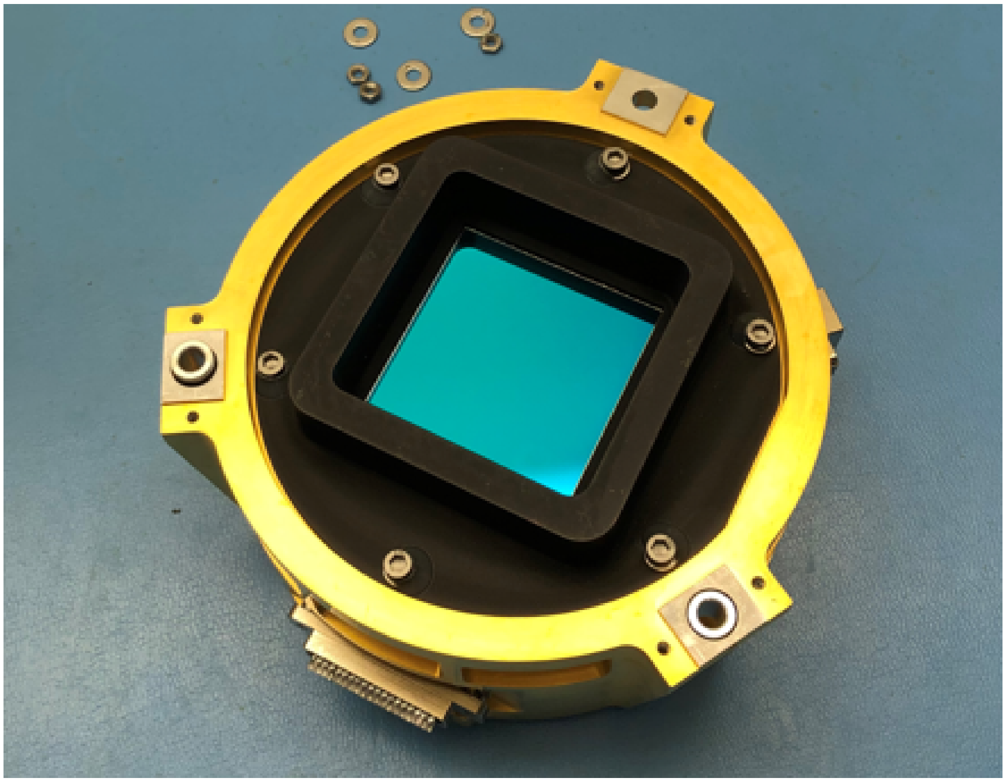

Figure 2.6: The H2RG mounted in its puck.

with IR background photons. Andover Corporation[3] offers custom filters that could meet the requirement of high transmission around the Ar emission lines we plan to observe and high blocking out to the detector cutoff of 5.3 µm. However, these are much more expensive than Andover's off-the-shelf narrowband filters, and would take much longer to be delivered. It was determined that a well-selected narrowband filter, paired with some additional filters designed for IR-blocking, could perform just as well as their custom solution, but for a much lower cost.

Of the narrowband filters offered by Andover, the most sensible for observing the Ar emission lines is a 9.5 mm-thick filter with a central wavelength of 1100 nm, FWHM of 54 nm, and a throughput outside this band below 10^{-4}. The nominal FWHM of this filter and the manufacturing tolerance is shown alongside the Ar lines in Figure 2.7. This filter has strong transmission in the region of the 1108 and 1111 nm Ar lines, so we selected this filter for the ODESSA demonstration with the aim of using these two Ar emission lines as our "OH lines".

The blocking of the narrowband filter out to 5.3 µm is not sufficient to suppress the

---

[3] https://www.andovercorp.com/

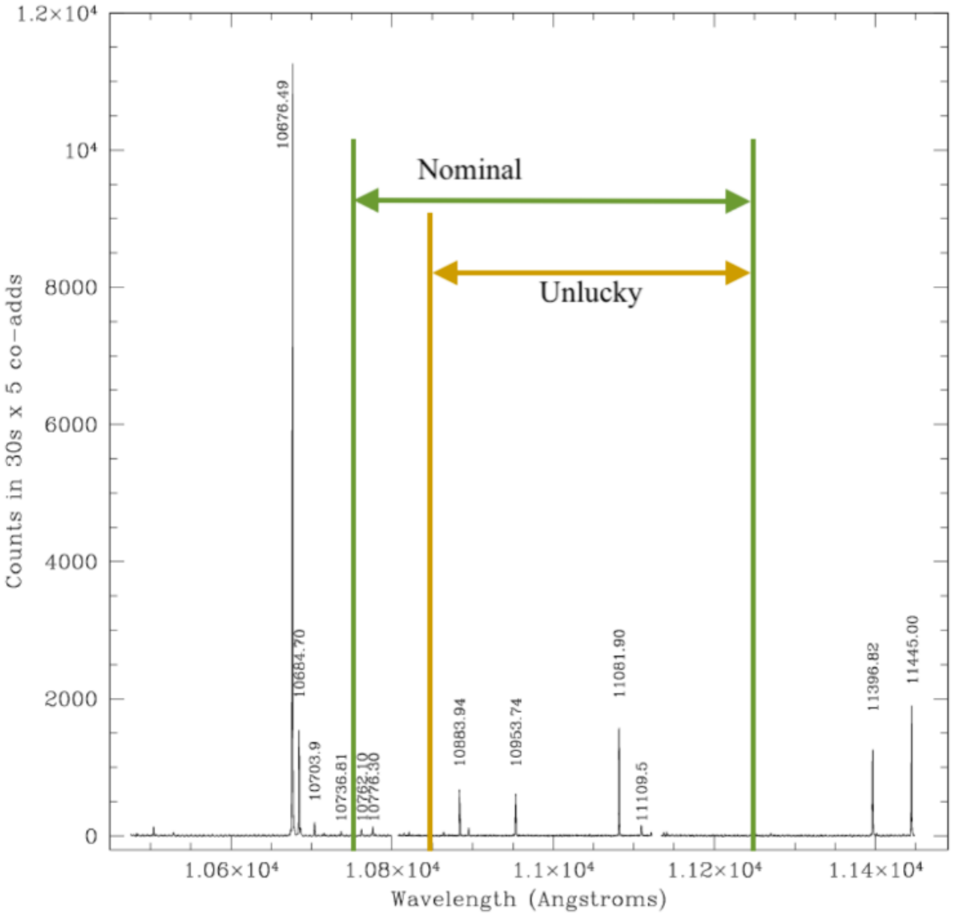

Figure 2.7: Argon emission lines and nominal/acceptable FWHM of the narrowband filter.

background blackbody radiation as much as necessary. Therefore, we included two additional spare IR-blocking filters from the HAA electronics lab, which were 1 and 2 mm thick. To estimate the total blocking power of the filter stack, we measured the transmission of each of these three filters from 1.25–6.67 μm using a Spectrum Two™ IR spectrometer at the University of Victoria.

The resulting transmission spectra of each filter, reported to a precision of 10^{-4}, are shown in Figure 2.8. We confirm that the transmission of the Andover filter is below 10^{-4} outside its bandpass. The 2-mm thick filter has a transmission spectrum consistent with that of PK-50 glass, a known IR blocking glass, and has transmission below 10^{-4} at 5 μm. The 1-mm thick filter, whose composition we do not know, has higher transmission at long wavelengths, but it was found to aid in reducing the IR background in images, so we included it as well in the filter stack. We estimate the total transmission of the combined filters is less than 10^{-9} at 5 μm.

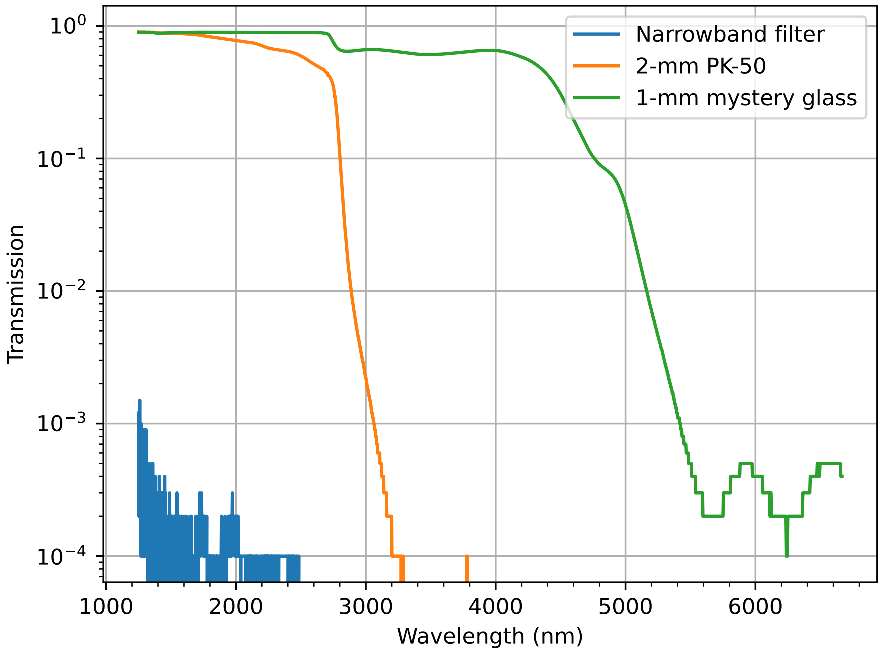

Figure 2.8: Measured filter transmission spectra, to a precision of 10^{-4}. The narrowband filter is the 1100 nm filter from Andover, and "PK-50" and "mystery glass" are IR blocking filters from the HAA electronics lab.

2.2.4 Cryostat

The detector needs to be mounted to the spectrograph using a cryostat which keeps the detector around its ideal operating temperature well below 60K (Rauscher et al., 2011). This can typically be achieved with liquid helium (^{4}He), which is how the detector was cooled in the cryostat of Chapin et al. (2022). However, this cryostat was too large to fit onto the mount of the 32-in arm of the McKellar Spectrograph. Instead, we used a smaller, liquid nitrogen (LN2)-cooled cryostat despite its lesser cooling capability.

The new cryostat uses two LN2 tanks to cool the detector and radiation shield by conduction: one is situated adjacent to the cold plate onto which the detector is strapped, and the other is strapped to the radiation shield which surrounds the detector. Windows on the front of the cryostat body and radiation shield allow light to enter. A schematic is shown in Figure 2.9.

The available space for the detector in the new cryostat is smaller than in the previously used cryostat, so the mount holding the puck to the cold plate was redesigned by Lothrop and Winn. The existing filter holder, which is necessary to hold the IR-blocking filters,

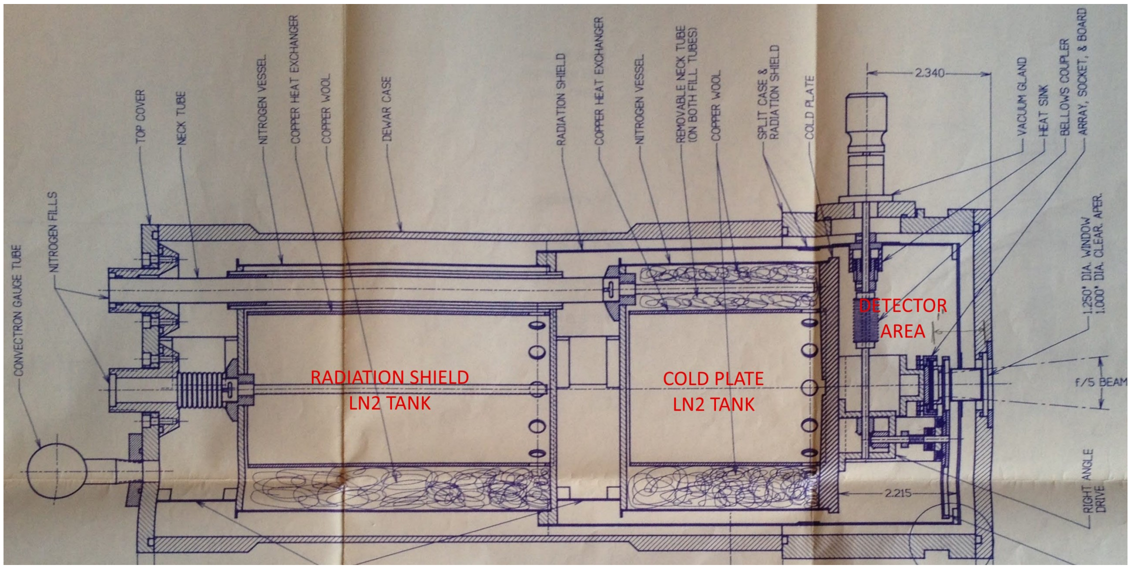

Figure 2.9: Schematic of the cryostat.

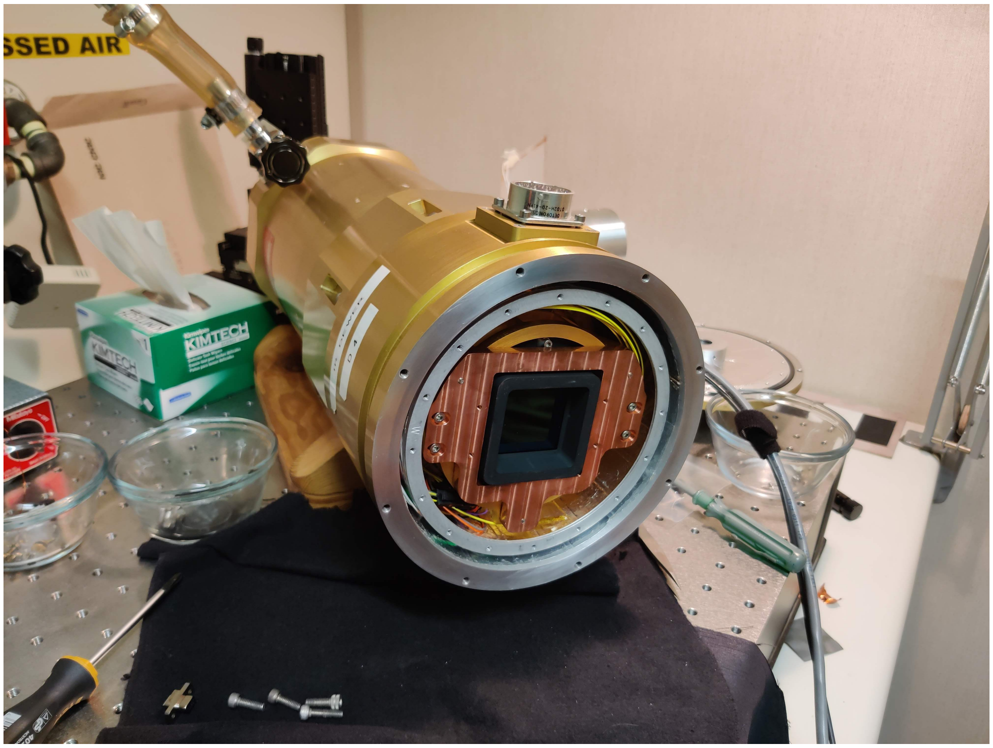

Figure 2.10: Detector installed in the cryostat.

could also no longer fit within the radiation shield. I therefore created a conceptual design for a holder which would fit through the window in the radiation shield, and this piece was finalized by Lothrop and Winn. I also resoldered the wires for the harness connecting to the puck through the cryostat exterior to a new connector which fits into the side of the cryostat. These components to which I directly contributed are shown in Figure 2.11.

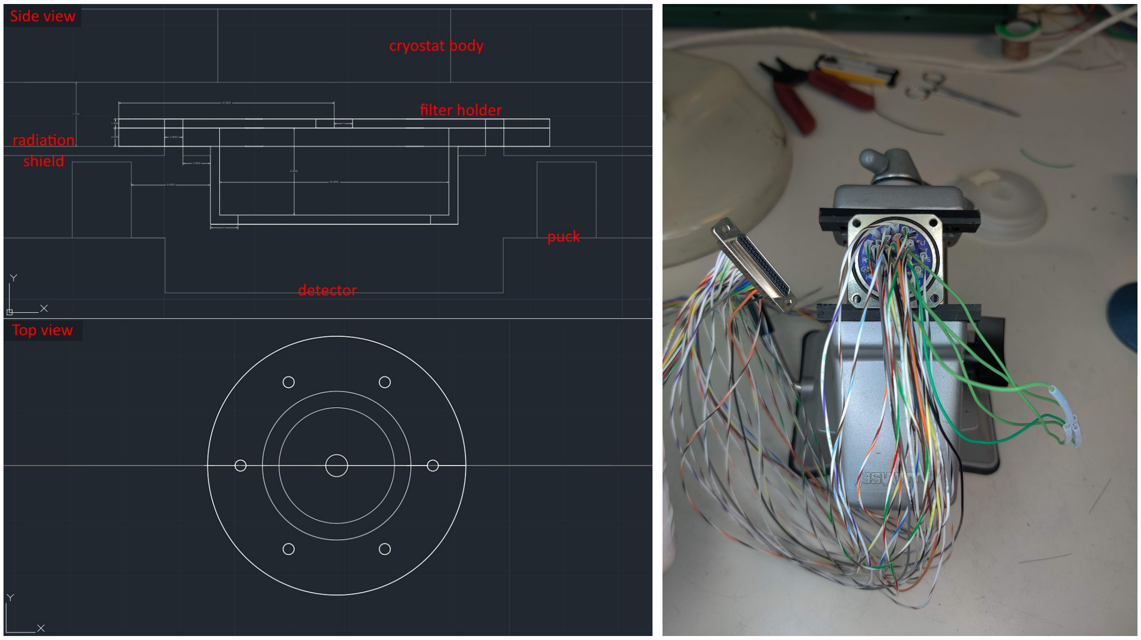

Figure 2.11: Hardware components for which I made a direct contribution. *Left*: Side and face view of preliminary design for the new filter holder. *Right*: re-soldered harness attaching the puck to the exterior of the cryostat.

2.3 Firmware Development

The code to control the readout of the detector is built off of C source code and a C++ API that came with the ARC Gen-4 controller[1] used in the experiments of Chapin et al. (2022). Modifications to the capabilities of this controller can be made easily by editing the source code and initializing it to the controller. As of the beginning of this project, Chapin et al. (2022) had created an On-Detector Guide Window (ODGW) readout mode which implemented up-the-ramp (UTR) sampling and the operation of windows during exposures. We extended this code to optimize it for ODESSA.

2.3.1 Readout Philosophy

The detector is read out by a series of operations defined by the controller. In standard UTR mode without any windows, these operations individually sample each pixel in a

---

[1] https://www.astro-cam.com/Gen4Hardware.php

row, one row at a time until the full frame is read. This can be repeated as many times as desired to add more samples up the ramp. While sampling any pixel or row, the controller can also specify that that pixel or row should be reset.

In ODGW mode with one or more windows, the controller samples an entire window after each row, resetting the window pixels if desired. It then appends the values of the entire window to the end of the row. If n windows have been defined in the exposure, the windows are rotated through so that each window is read every n-full-frame rows. An illustration is shown in Figure 2.12.

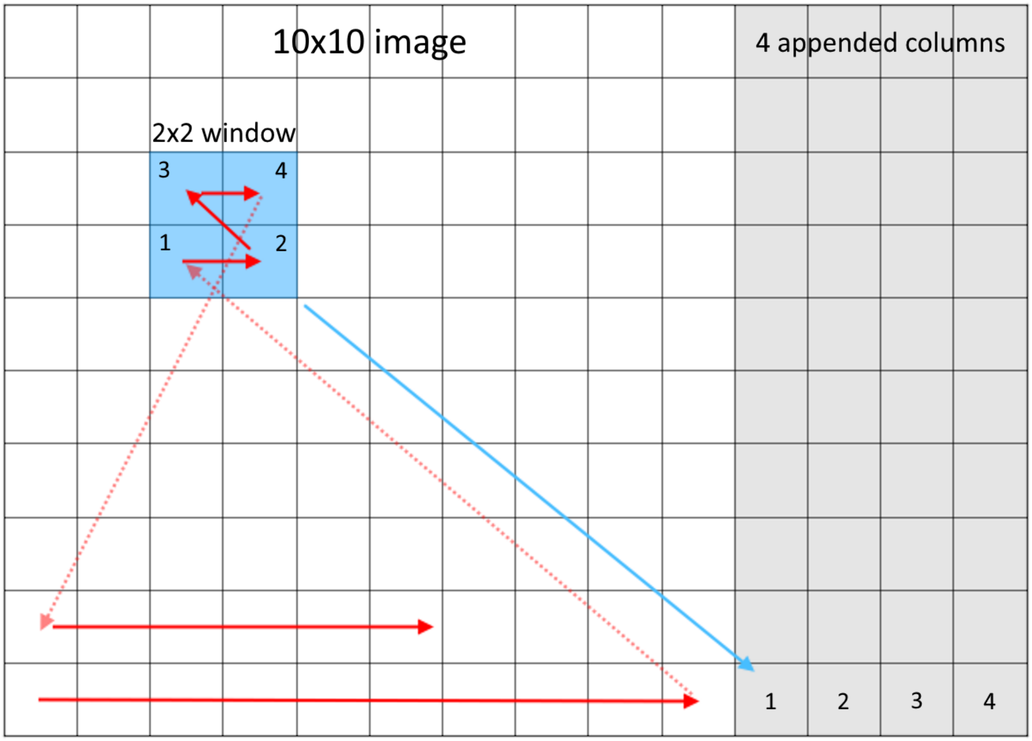

Figure 2.12: Illustration of how windows are read out in ODGW mode with a 10×10 frame and 2×2 window. Red arrows represent the order in which pixels are visited. Window pixels are appended to the end of a row according to their visit order, indicated by the numbers.

More specifically, each of these operations consists of a combination of electronic pulses sent to the detector. Pixels are selected by horizontal and vertical "clock" pulses (HCLK and VCLK, respectively), and READEN and RESETEN pulses activate reading and resetting of the pixel being visited, respectively. FSYNC pulses reset the sample position to the starting corner of the array. The general clocking scheme is shown in Figure 2.13.

The readout scheme is programmed directly into the controller, but there are certain

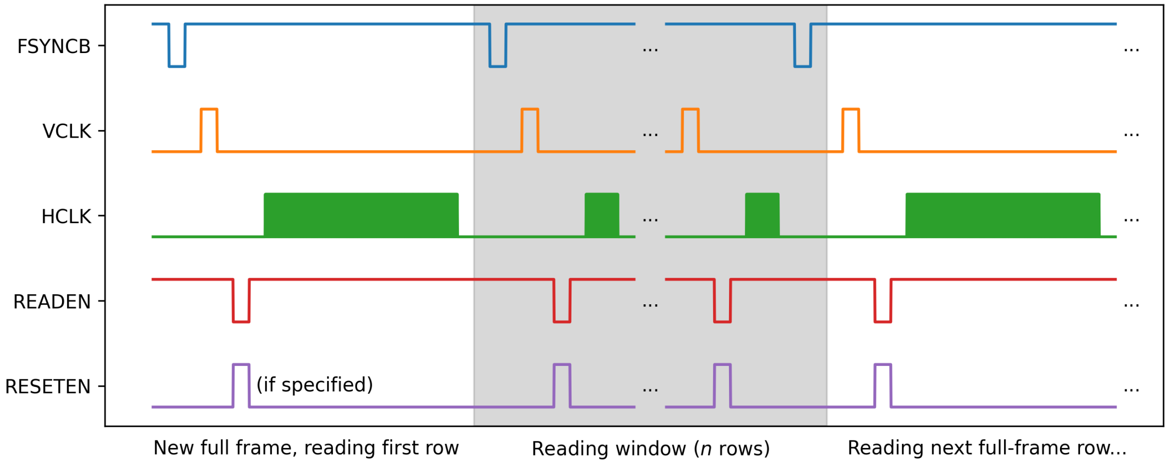

Figure 2.13: Clocking scheme for ODGW mode. Adapted from Figure 2 from Chapin et al. (2022).

variables that are set in exposure configuration files at runtime. Configurable parameters include the number of UTR reads to include after an initial reset, how often to reset the full frame, how often to reset the window(s), the size of the window(s), and the location of the window(s). In the case of multiple windows, each window can have an independent reset rate and location, but they must all be the same size.

2.3.2 Changes to the ODGW Readout Mode

Some changes to the initial ODGW mode were necessary for the demonstration of resetting emission lines. In the remainder of the thesis, I refer to the updated readout configuration as “bright region reset” (BRR) mode. Firstly, in case of exposures with many samples up the ramp, Wasif and Chapin adapted the firmware so that each frame is immediately written to its own FITS file, instead of waiting until the entire exposure is finished and writing to a single data cube. This ensures that potentially very large images do not have to be stored in memory for the entire exposure, and it also allows near real-time monitoring of the images as they are being exposed. I created a Python script that displays each frame as it is written to a FITS file and automatically stitches each file into one three-dimensional FITS file at the end of the exposure.

To ensure consistency across images regardless of the presence and size of windows, it is desirable to have each frame take the exact same amount of time. Initially, this was not the case, as reading out a window added an additional amount of time to each row proportional to the number of pixels in the window. To fix this in the BRR mode, we define a maximum window size of 32 rows × 16 columns, and introduce delays to the clocking scheme to increase the effective size of the window to this maximum size. For example, the clocking of an 8 × 8 window would consist of 8 HCLKs, 8 delays of the same length as the HCLK pulse, and a VCLK, repeated 8 times; then 24 iterations of 16 HCLK delays and a VCLK delay.

There is a small discrepancy between the theoretical time it should take to read one frame based on the lengths of the waveforms being sent and the actual measured time. This discrepancy is the same across all exposures, about 2% longer to read out than expected, though the cause is unknown. To prevent any potential issues this discrepancy could cause, I added timers into the code to measure the actual time taken to read out each row and frame; these measured times are reported in the FITS headers alongside the expected values.

Due to the significant IR background present in our images, we wanted to make the period between each UTR sample as short as possible so that as many samples as possible could be measured before the detector saturated. Since data is transferred from the detector through only a single video channel (a limitation imposed by available hardware), it takes around 40 seconds to read out a single 2048×2048 pixel frame—a factor of a few times longer than it takes to reach saturation. In the ODESSA demonstrations, we only observe a single spectral order which spans fewer than 100 rows on the detector. We decided to skip over most of the other rows, which do not contain any useful information, saving the time it would take to otherwise read out each of these unnecessary rows. This is done by adding an appropriate number of VCLK pulses after each frame's initial FSYNC pulse, reading the number of rows desired, then skipping straight to the next frame's FSYNC. The number of rows to skip as well of the number of rows to read are specified in the exposure configuration files. For ODESSA, we generally skipped over the first ∼1000 rows and read out 100 rows, resulting in a frame period of 2.7 s. This was generally short enough to sample up the ramp around three times before saturation due to the background.

The core of the ODESSA concept is the ability to reset windows at some ideal rate: resetting too slowly results in saturation and loss of any subsequent information, while resetting too quickly results in ramps whose flux cannot be measured as precisely due to too few UTR samples between resets. Depending on the flux of the emission lines we wish to reset, it may be optimal to have different reset rates for each line. The initial ODGW mode did not allow this ability, so I enabled it for BRR. This involved propagating a user-configurable setting into the firmware logic which decides whether a pixel is reset each time it is visited.

2.4 Software Development

Similarly to the firmware, data reduction software was created initially for the ODGW project, and was developed further to suit the needs of ODESSA. The reduction process is detailed here.

2.4.1 Image Acquisition

Individual exposures are configured by creating a JSON file containing the configurable exposure parameters. These include the name stem of the output files, the size and location of the full frame, video channel information, the size and location of windows if they exist, the number of full frame and window reads and resets to perform, and the number of times to repeat the exposure. This file is read by a C++ script which initializes the controller with the binary file created by compiling the firmware, executes the necessary commands on the controller, and writes the output data to FITS files with appropriate header entries. The raw data returned by the controller is inverted—i.e., higher electron counts corresponds to lower analog-to-digital units (ADU)—and pairs of columns are swapped, so these are both corrected before writing to FITS. Each UTR frame is written to its own FITS file as soon as the detector is read out.

To simplify and speed up the image taking process, the C++ executable is wrapped in a Python script which can be quickly edited and called from the command line. After asynchronously calling the C++ script, the Python script listens for the newly generated FITS files and displays them as soon as they are written. Configurable parameters in the script allow this image to be reset- and dark-subtracted if desired, allowing continuous monitoring of the accumulated charge in the detector. Once the image sequence has finished, the script stitches each frame in an exposure into a single three-dimensional UTR FITS file for that exposure.

2.4.2 Image Calibration

Electronic systems have long been known to exhibit noise following a $1/f$ power spectrum ("$1/f$ noise") as a result of thermal fluctuations within semiconductors (Voss & Clarke, 1976). In IR detectors, this can be seen as a fluctuation in the values of pixels over time, as seen for reference pixels in Figure 2.14. To remove the $1/f$ noise variation, we calculate the residuals of the best-fitting line for each pixel with respect to the read number. We then take the median of these residuals for the eight reference pixels in each row of a frame

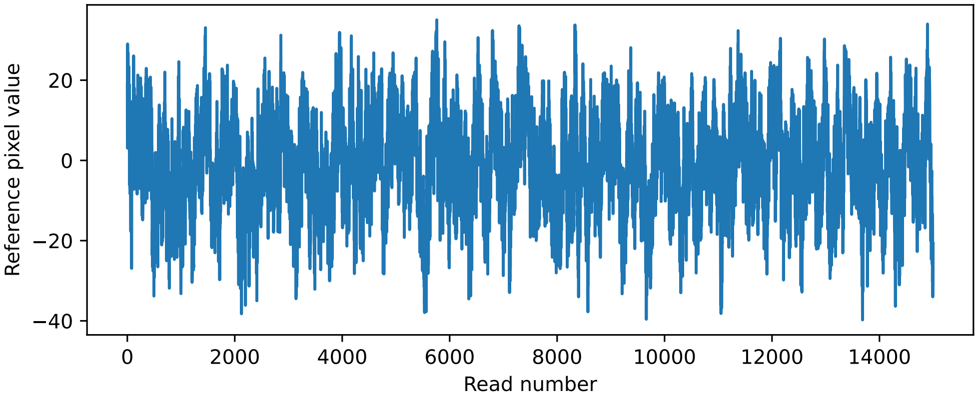

Figure 2.14: $1/f$ noise drift in reference pixels over time. The range of this graph is approximately 5 minutes.

(four on either side of the active pixels) and subtract the median from the corresponding row in each frame.

While an ideal detector would produce a ramp with constant slope for a stable source, the measured number of counts for actual detectors tends to roll off while approaching full well capacity (Blank et al., 2011; Plazas et al., 2017; Artigau et al., 2018). We attempted to create a linearity correction as is typical for this stage of the calibration process (Bushouse et al., 2022), but we were hindered by the presence of the so-called "reset anomaly" which affects H2RGs (Bezawada & Ives, 2006; Rauscher et al., 2007; Langevin et al., 2022). At present, we forgo the linearization step, and comment on the implications of this decision in Section 3.4.2.

Each exposure needs a matching dark exposure with the exact same readout configuration to account for any additional signal produced by operating windows during readout. The number, location, and size of the windows needs to be identical, as well as the number of UTR full frame reads. Master dark images are created for each configuration by taking the median of a series of reference-subtracted dark data cubes with window reads still appended to each row. The resulting master dark is subtracted directly from a reference-subtracted data cube. Although typically dark images are used to measure dark current (thermally-freed electrons which are collected in pixels), the dominant source of signal in our dark images is from external thermal IR emission. These sources behave in roughly the same manner, though dark current is much smaller, so I continue to refer to these images as "dark images."

Bad pixel masks are created from several dark exposures. For each of these darks, the reference pixels are subtracted, and the flux and reset are measured from the ramps using the method in Section 2.4.3. Bad pixels are identified using an iterative sigma-clipping procedure on both the reset and the flux images. The raw reset and flux images, as well as histograms of their values, are shown in Figure 2.15.

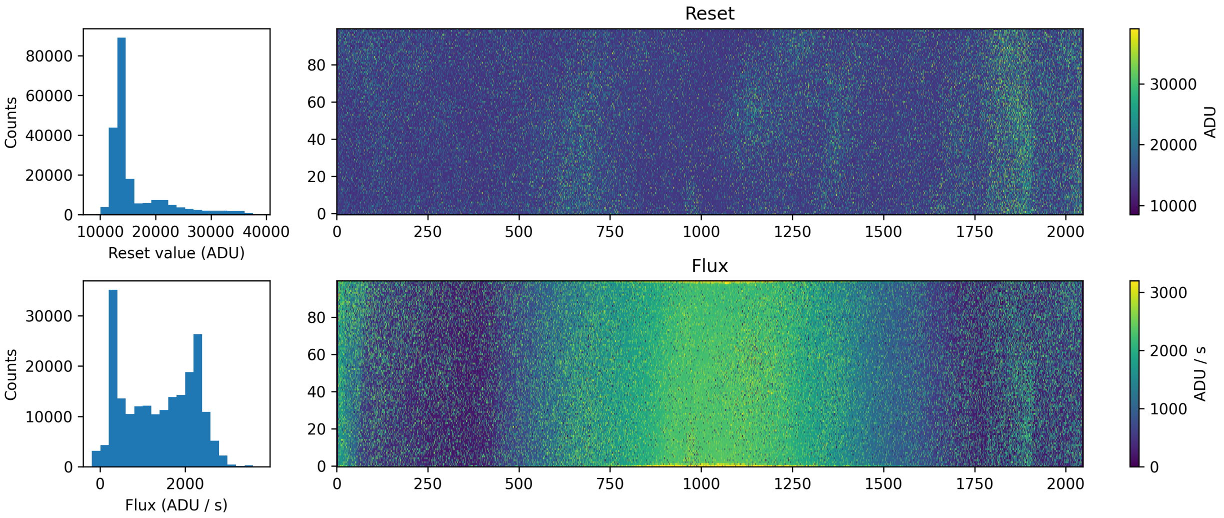

Figure 2.15: Reset and flux maps of a dark image, and their associated histograms.

Due to the long tail of the reset histogram, we found that a reliable way to end the iterative clipping process is to identify when the mean of the non-masked pixels' reset values is sufficiently close to the mean of the best-fitting Gaussian to the histogram's peak. We assume the latter to be the "true" reset value for good pixels. For example, Figure 2.16 shows the first and last iterations on the dark image shown. If the two Gaussians are farther apart than a threshold value, then pixels farther than $3\sigma_{\mathrm{pop}}$ from the population mean are masked. We arbitrarily set the threshold value to be $0.05\sigma_{\mathrm{tot}}$, where $\sigma_{\mathrm{tot}}^2 = \sigma_{\mathrm{pop}}^2 + \sigma_{\mathrm{hist}}^2$ is the combined population and histogram standard deviations. This process is then repeated until the two Gaussians are within the $0.05\sigma_{\mathrm{tot}}$ threshold of each other.

Most bad pixels are caught by the filter on the reset image, but some pixels which have typical reset behaviour exhibit atypical flux behaviour. The bright IR background pattern on our detector means we cannot sigma clip on the raw flux image. Instead, we create a smoothed flux map by convolving the masked flux map with a 2D top hat function of radius 10 pixels, and take the residual of the flux map minus the smoothed map. We expect the mean of the residual to be zero, so we iteratively clip at 3σ until the mean is within the standard error of the mean $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$, where n is the number of remaining good pixels.

To create a final bad pixel mask, we use the above procedure to identify bad pixels

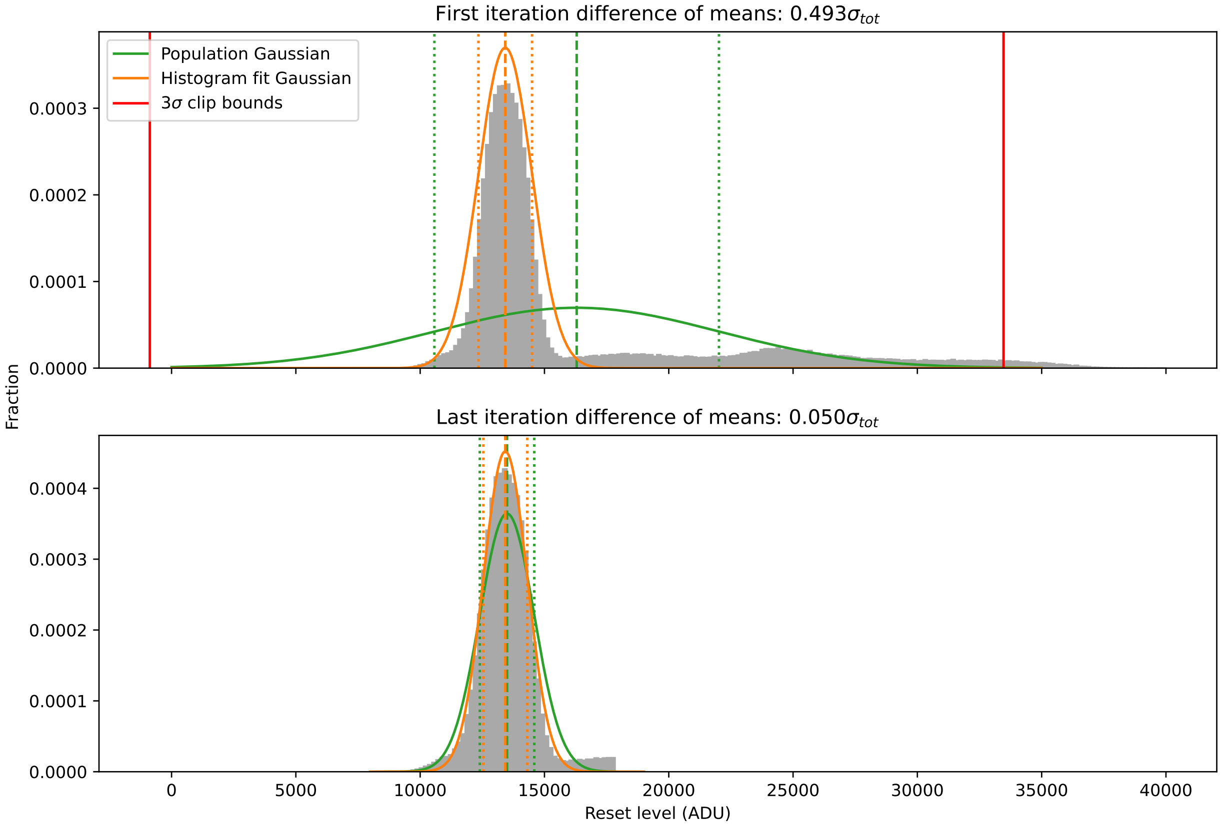

Figure 2.16: First and last iterations of reset clipping procedure. The mean and standard deviations of the two Gaussians are represented by vertical dashed and dotted lines, respectively. For each iteration, pixels with reset values outside the red lines are added to the bad pixel mask.

on a series of dark images. In the final mask, a pixel is marked as bad if that pixel was masked in at least half of the input dark images. We find that the total fraction of bad pixels on this detector when operating around 65K is about 27%, though that number is somewhat lower on the portion of the detector that receives light. A comparison of a masked and unmasked dark image is shown in Figure 2.17. Bad pixel masks are applied to dark-subtracted images, and the bad pixels are disregarded for the remainder of the analysis.

2.4.3 Calculating Flux

After basic image calibration, we want to convert the data cubes into two-dimensional images showing the flux and uncertainty per pixel in electrons / second. We do this by fitting a slope to UTR reads of each pixel, then using the electron conversion gain of

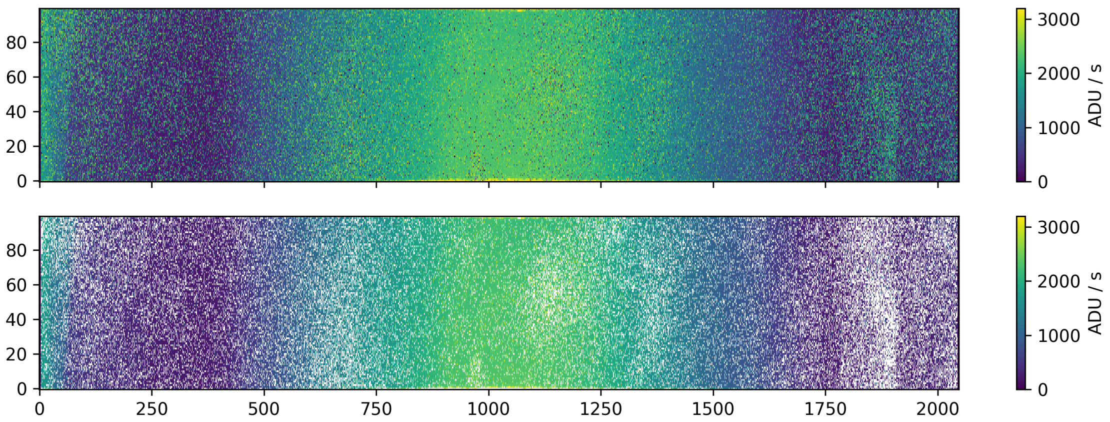

Figure 2.17: Dark image before and after applying the bad pixel mask. Bad pixels are marked in white.

the detector and the read interval to convert the slope from ADU / read to e$^-$ / s. We first separate the full frame and window components of each exposure and treat them individually.

To measure the slope and uncertainty of each pixel's ramp, we follow the method of Rauscher et al. (2007) and Robberto (2009), which was developed for the H2RG detectors of the *JWST* Near-Infrared Spectrograph (NIRSpec). This method uses the slope found by an ordinary least squares (OLS) regression,

$$b = \frac{\Sigma_i (x_i - \bar{x})(y_i - \bar{y})}{\Sigma_i (x_i - \bar{x})^2} \times \frac{egain}{t} \ (\mathrm{e}^-/\mathrm{s}),$$

where x_i are the read numbers, y_i are the measured counts, $egain$ is the electron conversion gain in e$^-$ / ADU, and t is the time between each read in seconds. The variance of the flux is derived by considering the Poisson and read noise introduced to every measurement, then correcting for the fact that the measurements are not independent samples: the electron counts are cumulative, so consecutive samples will tend to be closer together than a random sampling would expect. The resulting derived equation for the variance of the flux, slightly simplified for our case, is

$$\sigma_b^2 = \frac{6}{5}\frac{n^2+1}{n(n^2-1)}\frac{b}{t} + 12\frac{1}{n(n^2-1)}\frac{\sigma_{RN}^2}{t^2},$$

where σ_{RN} is the RMS read noise in e$^-$.

This uncertainty is similar to that which is found by an OLS regression, though more robust against correlated noise. Compared to correlated double sampling (CDS), which is typical for detectors which are incapable of nondestructive sampling such as CCDs, the uncertainty used here is much lower due to the greater number of samples. Similarly, the uncertainty of a full frame pixel's slope is much larger than that of a window pixel due the smaller number of samples. This allows the signal-to-noise ratio (SNR) within a window to be significantly higher than outside it. A comparison of these different fitting methods is shown in Figure 2.18.

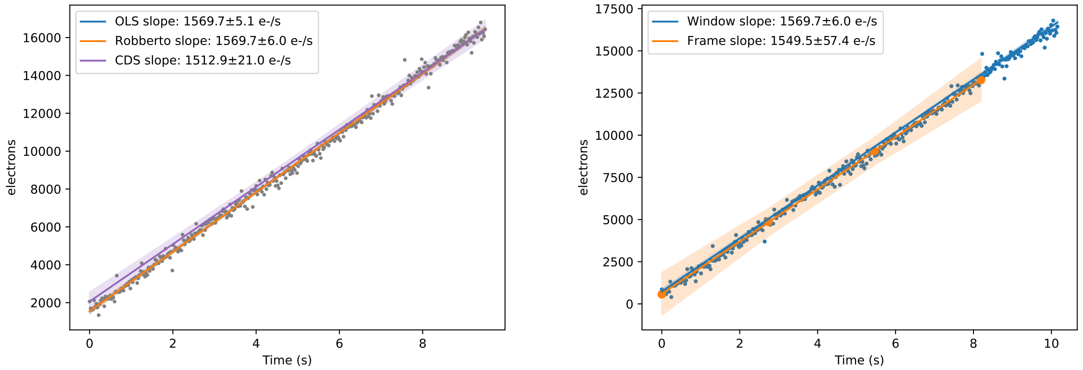

Figure 2.18: Comparison of different ramp-fitting methods. *Left*: Different methods of fitting a window ramp. *Right*: Robberto method of fitting a window and full frame ramp of the same pixel.

In the case of a window which has been reset during the exposure, a new ramp is fit after every reset. A total flux and uncertainty for the pixel is found by finding the weighted mean of each individual fitted flux, where the weight is the inverse variance uncertainty of that pixel:

$$b_{\text{tot}} = \frac{\Sigma_i \frac{b_i}{\sigma^2_{b,i}}}{\Sigma_i \frac{1}{\sigma^2_{b,i}}}, \text{ and } \sigma^2_{b,tot} = \frac{1}{\Sigma_i \frac{1}{\sigma^2_{b,i}}}.$$

Since the flux of the full frame and window(s) of an exposure have been calculated separately, we then recombine the window flux into its appropriate location in the full frame, creating a single flux image using all the data collected.

2.4.4 Extracting Spectra

To convert 2D flux images into 1D spectra, we use a weighted mean extraction with the weights based on "flat" images of an incandescent lamp. After fluxes for the spectra have been calculated and median combined to create a master flat, we define a region of pixels on the detector with which to determine the weights. This region includes the entire extent of the detector which receives light from the spectrograph. Weights for each row within this region are determined by taking the median of the pixels in that row, and weights for rows outside the region are set to 0. Per-column 1D spectra and uncertainties are then extracted by taking the weighted mean of a flux or uncertainty (variance) image.

Wavelength solutions are found using master arc images created in the same way as the master flats. The spectra of the master arcs are extracted as above, then a Gaussian is fit to each known emission line. For ODESSA, we use the 1108.19 nm and 1110.95 nm Ar lines from an FeAr arc lamp. The pixel location of the mean of each Gaussian is used to fit a linear wavelength—pixel relationship, from which a wavelength is determined for each value in the 1D spectra.

2.4.5 Coadding Images

After a group of like images has been extracted into 1D spectra, they are resampled onto a common wavelength grid corresponding to the range and twice the pixel resolution of the detector: 1101—1128 nm in 0.016 nm increments. This is done using the `specutils` Python package's `LinearInterpolatedResampler`, which resamples both the flux and uncertainty of a given spectrum onto a new grid. Finally, spectra are coadded using an inverse variance weighted mean and normalized by their mean. The final coadded error spectra represent the total RMS noise per wavelength resolution element.

2.5 Detector Characterization

In order for the above reduction method to produce meaningful values, it is necessary to have a strong understanding of the nature of the imaging system. To achieve this, we carried out a series of characterization tests of the detector in its cryostat.

2.5.1 Electron Conversion Gain

After electrons are collected in the pixels of a detector, they are converted into units commonly called analog-to-digital units (ADUs). To obtain images in units of Poisson-statistics-obeying electrons, astronomers often multiply ADUs by the detector's electron conversion gain. The typical method of measuring the gain for a detector is the photon transfer curve (PTC) method of Janesick et al. (1985). The PTC method takes advantage of the Poisson nature of the collected electrons by measuring the mean and variance of a series of flat field images—the slope of the variance-vs-mean curve is equal to the gain.

While this is a simple enough procedure for CCDs, CMOS and IR detectors' independent pixels means that each pixel has its own gain (Le Graët et al., 2022). In order to reliably measure the mean and variance of a single pixel over a range of illuminations, one would need to obtain hundreds of flat field images. In addition, the correlated noise of the UTR samples means that ramps cannot be used as a shortcut to obtain many illuminations at once. Because of the difficulty of obtaining a proper PTC dataset, as well as our lack of a decent flat field setup, we use several approximations to measure our detector's electron conversion gain.

Although we were unable to obtain true flat field images, the smoothly-varying IR background in our dark images is bright enough to obtain images at a few different illuminations up to saturation. We take 50 UTR dark exposures which, in this setup, each contain a reset frame plus three UTR frames before saturation in the brightest part of the image. To avoid the correlation between samples in the same exposure, exposures are divided evenly so that only one of the first, second, or third frames is used from each exposure. Useful PTCs cannot be created from pixels that receive no light, so we use only the pixels in the region marked in Figure 2.19, corresponding roughly to the FWHM of the background flux.

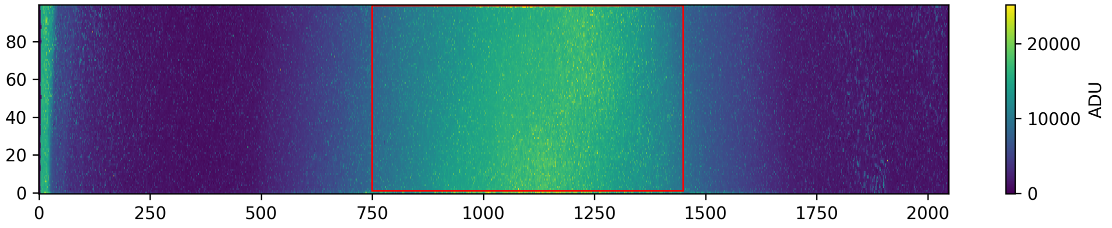

Figure 2.19: Region of dark image used for PTCs.

Due to our limited sample size, we make the assumption that the all the individual gains across the detector are approximately equal. However, the varying brightness of the background flux means we cannot take the mean and variance across the entire detector. Instead, we group the detector into 4 × 4 "superpixels," and take the mean and variance of pairs of images within those superpixels, ignoring bad pixels. PTCs are then created for each superpixel, and the slope of the PTC is measured using OLS. Sample PTCs and their best-fitting slopes are shown in Figure 2.20, and a histogram of all measured PTC slopes is shown in Figure 2.21. The inverse of the median of each measured slope is taken to be the average gain of the detector, and the 1σ quantiles are the uncertainty. We measure an average gain of 3.0 ± 0.6 e^{-}/ADU, which is consistent with the expected gain for

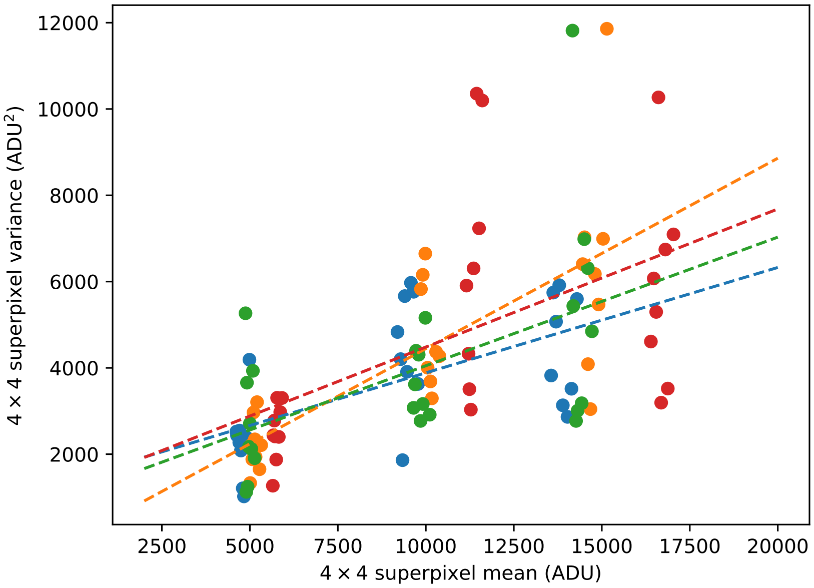

Figure 2.20: Sample PTCs for four different superpixels, distinguished by colour. Each point represents a pair of dark images, and dashed lines are the best-fitting PTCs for their respective colours.

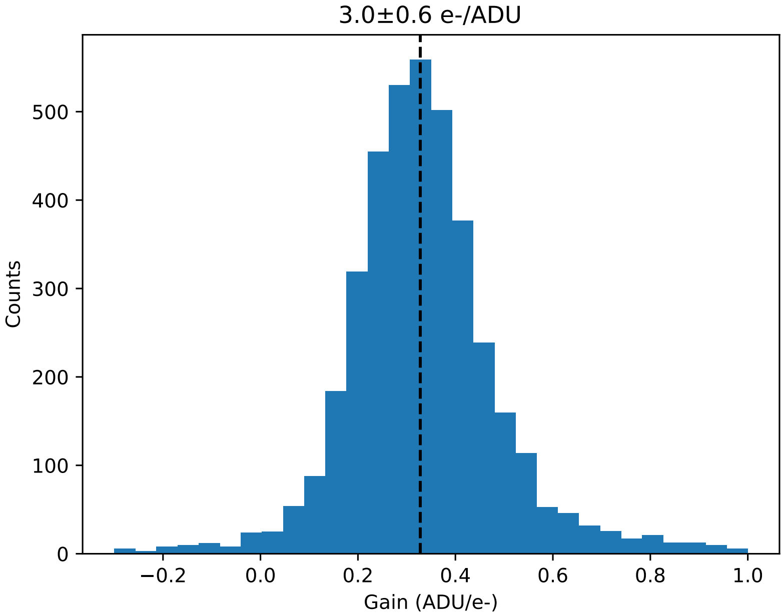

Figure 2.21: Histogram of measured superpixel gains.

HxRG detectors (Hilbert & Rest, 2016; Ives et al., 2020; Le Graët et al., 2022).

2.5.2 Readout Noise

The PTC method can also be used to measure readout noise, but it requires finely sampled data in the low-flux regime. We therefore estimate the readout noise of the detector by finding the variance of the reset values of each pixel using the same data as in Section 2.5.1. A histogram of the measured RMS read noise values is shown in Figure 2.22, and the median read noise across the detector is 153 e^-. This method is sensitive to changes to the reset level across exposures, as seen in Chapin et al. (2022). However, in order to separate read noise variance from dark current or IR background variance, it would be necessary to take images with much shorter exposure times than we were able to achieve. The measured read noise here is thus likely an overestimate of the true read noise of the detector, which Teledyne reports should be well below 100 e^- (Blank et al., 2011).

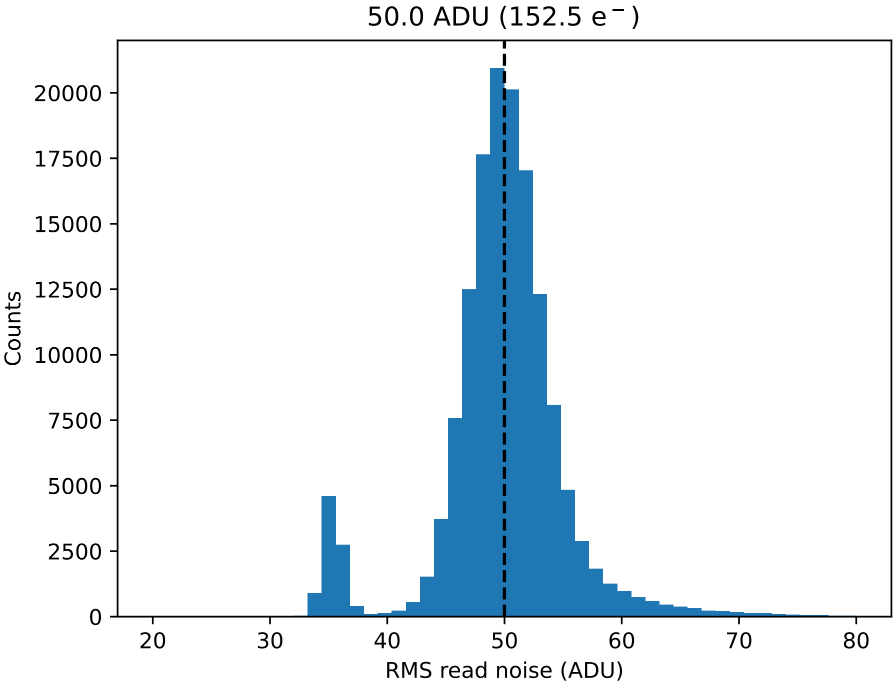

Figure 2.22: Histogram of measured RMS read noise.

2.5.3 Temperature Dependence

Detector Temperature

The quality of image depends on the temperature of the detector. While the detector was still in the ^{4}He-cooled cryostat, we took some dark images at various temperatures to compare the effects. Figure 2.23 shows that the number of hot pixels increases dramatically with higher temperatures, in agreement with Rauscher et al. (2011) and Regan & Bergeron (2020). Between 70 and 80K, the detector becomes dominated by bad pixels which drastically reduce the amount of useful data contained in an image.

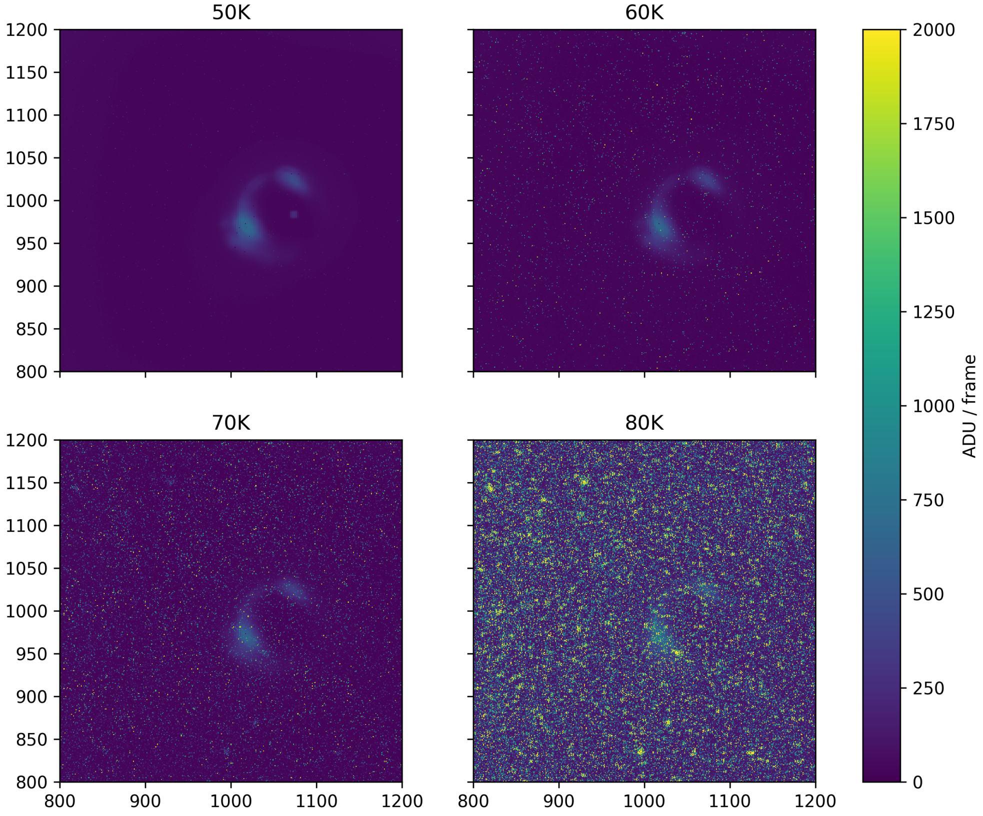

Figure 2.23: Dark images of the central region of the detector at different temperatures.

Since the cryostat that we use is cooled by 77K LN2, it is necessary to find a way to make the detector colder than the cryostat would typically be able to achieve in order to obtain usable images. To do this, we attach a vacuum pump to the LN2 tank that is adjacent to the cold plate on which the detector sits. By lowering the pressure in the tank, the boiling point of the LN2 decreases, so higher-energy molecules evaporate more easily, leaving behind lower-energy molecules as liquid in the tank. We find that by pumping the

LN2 tank to a few Torr, we can achieve a cold plate temperature near 50K, as seen in Figure 2.24. The temperature of the detector is somewhat warmer, reaching around 62K.

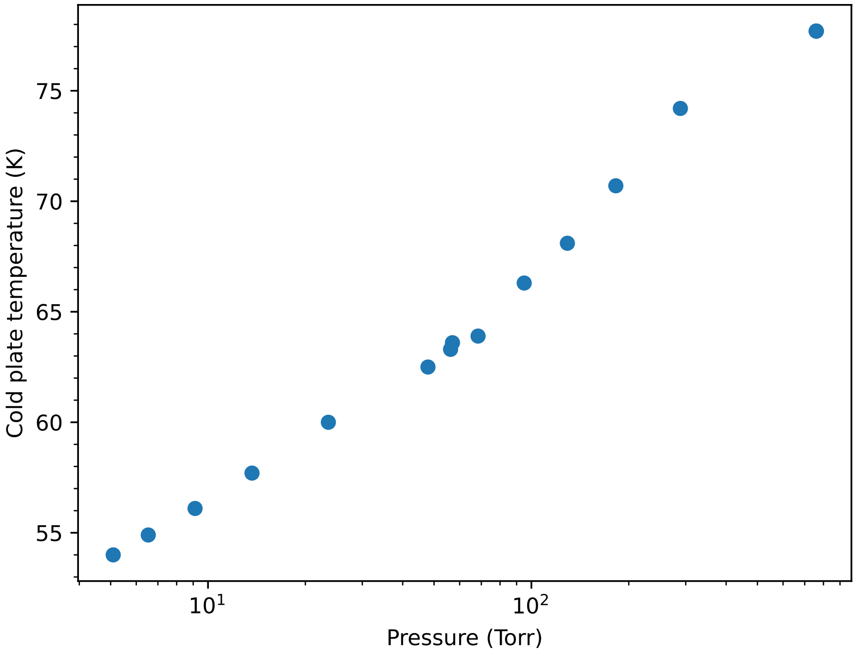

Figure 2.24: Cold plate temperature vs. pressure of LN2 tank.

Radiation Shield Temperature

Because our detector is sensitive to radiation out to 5 μm, it is able to detect blackbody radiation from room temperature objects which is not completely blocked by the filters. In addition, slightly warmer components within the cryostat can emit a non-negligible amount of radiation at 5 μm, so even when the radiation shield and filters are cold, they contribute to the background seen by the detector. A typical background for our detector can be seen in Figures 2.15 and 2.19. As mentioned in Section 2.2.3, stacking a narrowband filter plus two IR-blocking filters in the cryostat is successful at reducing the external IR background to a manageable level. However, this background still varies with the temperature of the radiation shield, most likely due to blackbody emission of the filters, which only reach a temperature of around 85K. An estimate of the blackbody radiation seen from sources external and internal to the cryostat is shown in Figure 2.25. A significant challenge posed by our choice of cryostat was reducing the background to a level which would not

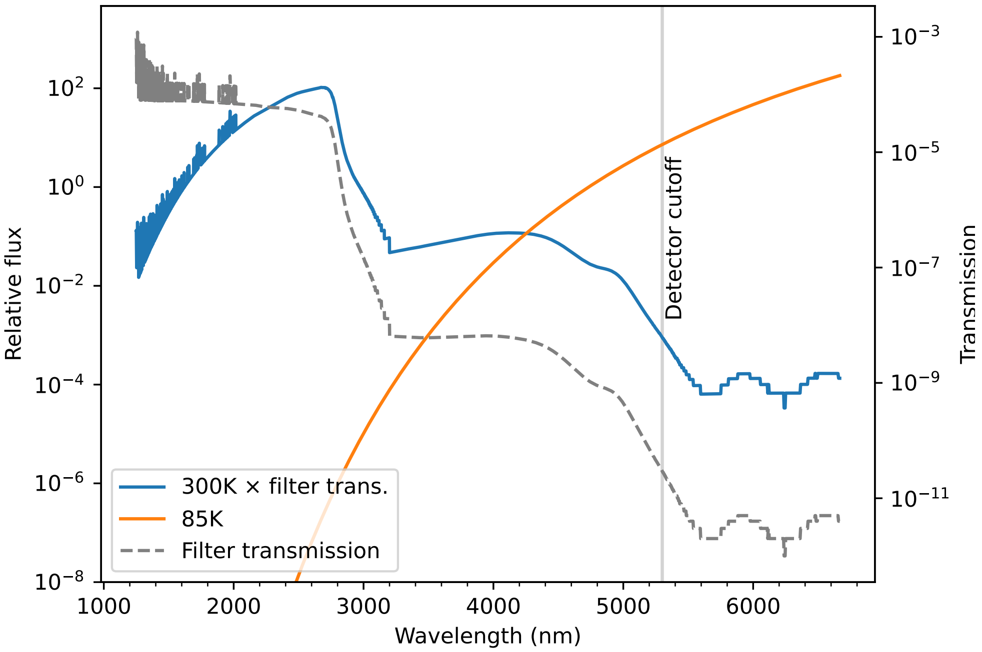

Figure 2.25: Emission of 300K and and 85K blackbodies, as seen by our detector. The 300K emission coming from outside the cryostat is multiplied by the total transmission of the filter stack, while the 85K emission comes from the filters themselves. Blackbody fluxes assume the same emitting area and perfect emissivity.

immediately saturate the detector.

We used a series of dark images taken as the radiation shield was cooling to quantify the effect of the radiation shield temperature on the background. Figure 2.26 shows how the temperatures of the detector and radiation shield and the median background flux in the image vary over time. The flux closely follows the temperature of the radiation shield with a lag of around 16 minutes; Figure 2.27 shows the relationship between the two after shifting the radiation shield temperature by this amount. At operating temperatures, there is a strong correlation between the two, and the turnoff at higher temperatures is likely due to the high detector temperature increasing the number of bad pixels. To mitigate this effect, we ensure that the radiation shield is always at as low a temperature as possible when taking images.

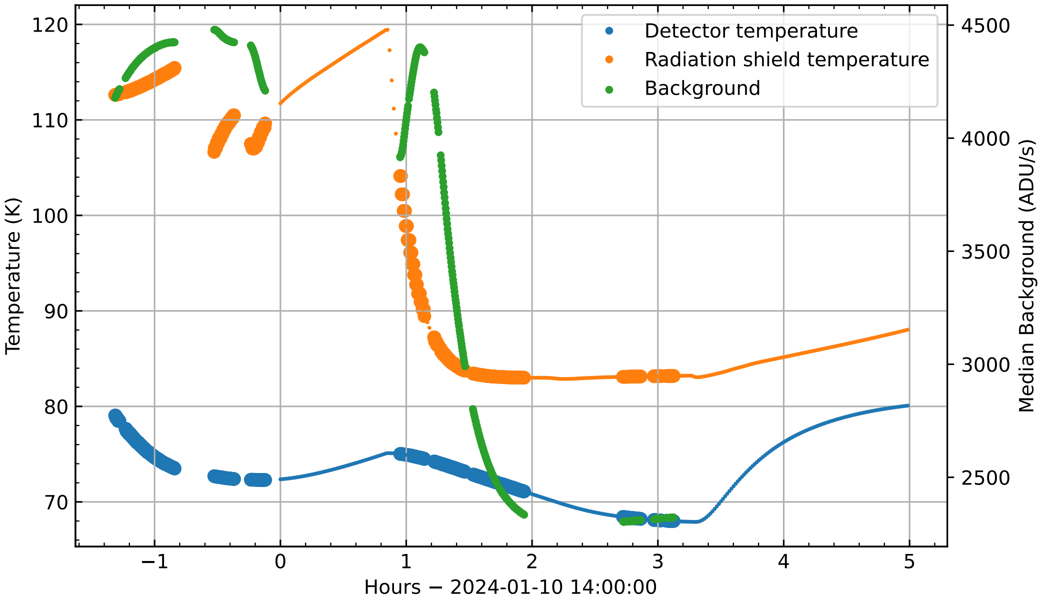

Figure 2.26: Detector and radiation shield temperatures and median background flux over time as the cryostat cooled. Larger temperature points are measurements that have corresponding images for background measurement. The background follows the radiation shield temperature fairly closely with a lag of 16 minutes.

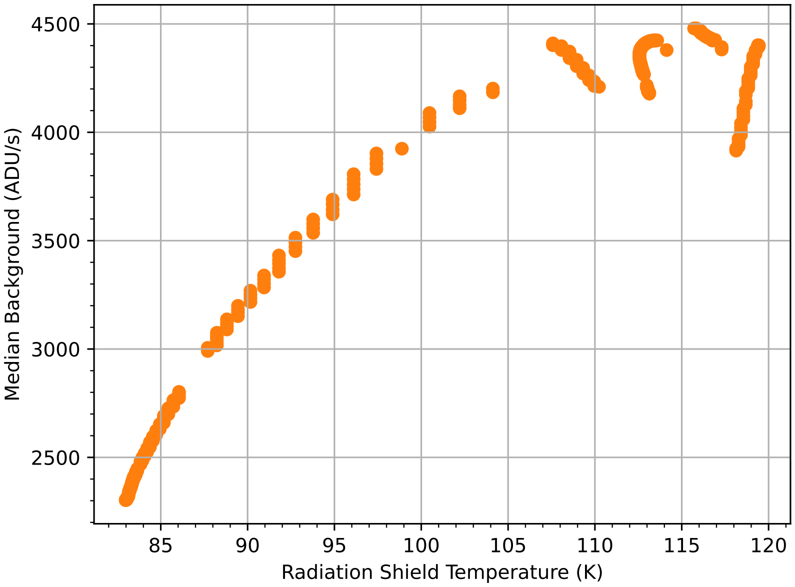

Figure 2.27: Median background flux vs. radiation shield temperature, after offsetting the temperature by 16 mins to account for the lag between temperature and background.

2.5.4 Persistence

One of the potential benefits of resetting bright emission lines during exposures is the suppression of persistence current. If left to saturate, bright sources can leave charge embedded within "traps" in the detector. The charge is then slowly released, resulting in an additional apparent flux in subsequent images lasting potentially for hours (Smith et al., 2008; Ives et al., 2020). Preventing the traps from filling in the first place via resetting of bright sources could thus be a strategy to mitigate persistence current (Bezawada & Ives, 2006).

To test the strength of persistence current in our detector, and the extent to which windows can prevent it, we carried out an experiment while the detector was still in the ^{4}He-cooled cryostat comparing images with and without windows. The procedure consisted of placing two 8×8 windows in different locations on the detector: one in the bright part of a spot caused by illuminating an LED, and one far from the spot. An initial dark image is taken, with constant reading of both windows, then the LED is flashed to saturate a small area of the detector, and the electrons are allowed to "soak" in the detector to fill any vacant charge traps. A second dark is then taken, and the flux measured in the difference between the two dark images should consist solely of persistence current. This sequence is repeated once without resetting the windows, and once holding the windows in constant reset while the LED is on. An illustration is shown in Figure 2.28.

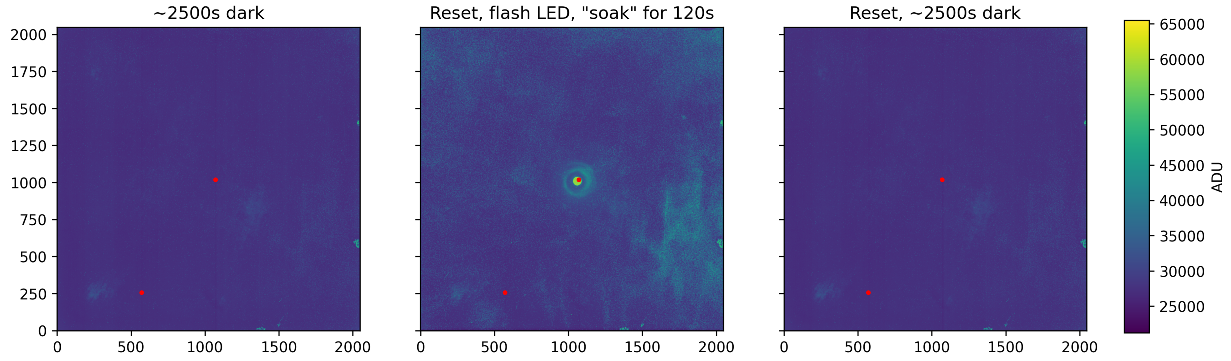

Figure 2.28: The image sequence of the persistence measurement experiment. Window locations are shown in red.

We found that the obvious persistence current which appears in the illuminated window which is not reset in Figure 2.29 is suppressed in the reset illuminated window. The unilluminated window, which does not exhibit any persistence current, remains unaffected as expected. The illuminated window appears to in fact be overcorrected by the window

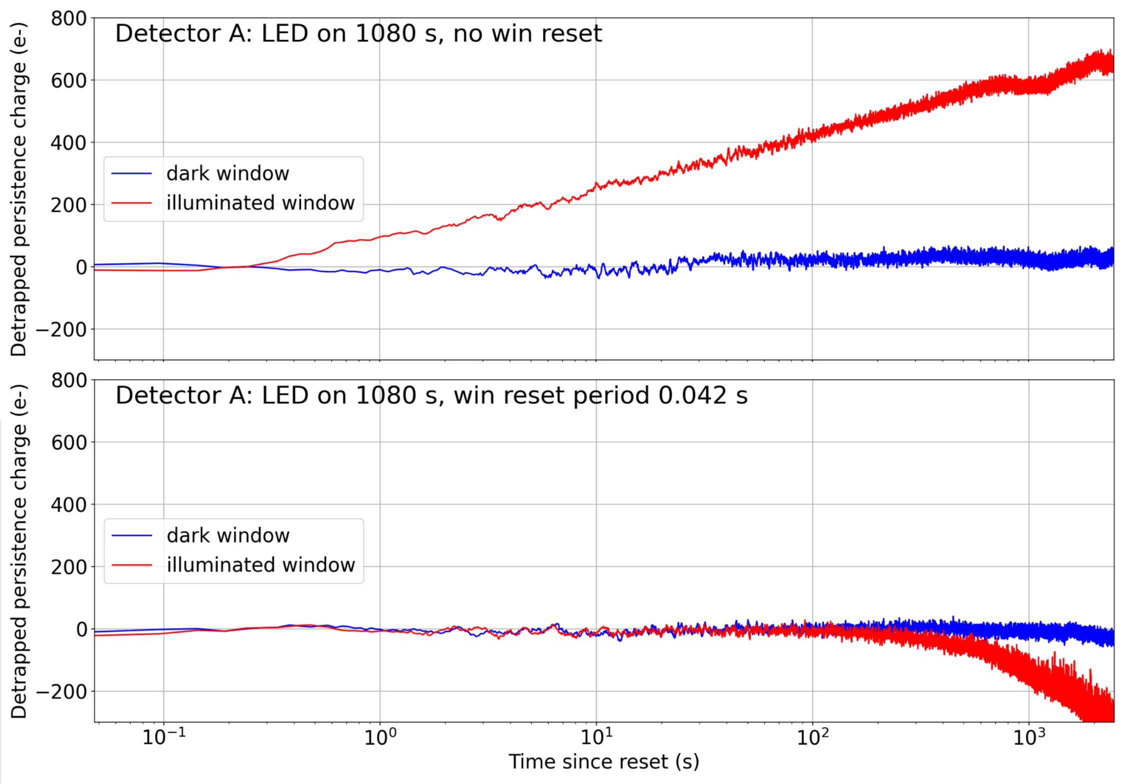

Figure 2.29: Measured persistence current without (top) and with (bottom) window resetting. The clear persistence current is removed, but is replaced by an anomalous negative apparent current.

resetting procedure, as seen by the anomalous negative apparent current. While the cause of this behaviour remains unclear, the results are more thoroughly discussed by Chapin et al. (2024, in press).

2.6 Summary

Significant effort went into building the system that would demonstrate ODESSA. We adapted a detector used in a previous experiment for use on the McKellar Spectrograph, and we ensured that both were set up so that we could easily image spectra of emission lines in the near infrared. We also built the firmware that is the core of the ODESSA concept and made a software pipeline to turn images from the detector into useful data products. Limitations in the setup resulted in a significant IR background in images, and also led us to use artificial spectra in lieu of actual sky spectra. Finally, we measured some characteristics of the system, finding a typical electron conversion gain, a quite high read noise value, and a strong dependence of image quality on temperature.

Chapter 3

Demonstration and Analysis of On-Detector Line Suppression

This chapter describes the on-telescope demonstration of ODESSA. This was carried out at the McKellar Spectrograph by observing a composite spectrum of an arc and an incandescent lamp, which represent a source spectrum contaminated with bright sky lines. Several data runs were necessary to identify the optimal methodology for the demonstration, including understanding some of the sources of systematic uncertainties such as varying background. The final dataset allows for the comparison of spectra obtained by resetting emission lines against those obtained without the use of windows. The results show that the ODESSA process does not have any major impacts on the final reconstructed spectra; it in fact improves the signal-to-noise for the spectral regions contained by windows. Residual differences in the line intensities of standard and reconstructed spectra can be attributed to correctable processes, and oddities introduced by window operation ultimately do not produce any detectable effect.

3.1 Operational and Demonstration Requirements

A successful demonstration of the ODESSA concept needs to show that periodically resetting windows on a detector and subsequently reconstructing the full spectrum does not significantly degrade the quality of the spectrum compared to one obtained using a more standard readout mode. In particular, we want to demonstrate whether systematic effects caused by windows significantly affect the features of the final 1D coadded spectra. If they do not, then this bodes well for the future application of this methodology to science observations that are otherwise negatively impacted by sky emission lines.

As a consequence of these considerations, it is desirable to ensure that the data that we analyze have high signal-to-noise. This prevents the data from being dominated by Poisson fluctuations in the signal, and allows us to meaningfully compare the noise properties of spectra obtained using the ODESSA technique and more standard techniques. To this end, we carried out several engineering runs designed to finalize requirements for the demonstration prior to the final data collection run.

We first need to identify the data that must be collected with the spectrograph. This involves understanding how the spectrum will appear in images using the setup defined in Section 2.2, including the usable range of the detector; the size, location, and focus of spectral features; and the light sources producing the spectra. Since we decided to use mock sky lines due to the faintness of OH lines observed with the telescope, we need a light source which reasonably replicates a spectrum that might be observed on-sky. For this, it is useful to have two separate sources: one simulating sky lines, and one simulating the "true" spectrum of some astronomical source without atmospheric contamination. Such mock spectra are not readily available at the DAO, so it is necessary to come up with our own.

Once we understand the light sources and necessary data, we need to understand systematic effects resulting from our setup. A major source of uncertainty is the bright IR background present in previous test images; therefore we need to characterize the behaviour of this background in operating conditions. As was discovered in Section 2.5.3, the background varies significantly over time, so in order to take several hours of combined exposures, we must continuously monitor the background by interleaving dark images with our spectrum exposures.

An additional source of uncertainty is the scale of the systematic effects that windows introduce into images. As has been well documented, operating windows on a HxRG detector can introduce voltage offsets, self-heating, and crosstalk effects (Bezawada & Ives, 2006; Smith & Hale, 2012; Regan & Bergeron, 2020; Chapin et al., 2022). Therefore we need to characterize what window-related effects occur as a result of our mode of operation.

Once all of these characteristics are better understood, we can use them to define how to optimally carry out the ODESSA demonstration. This chapter describes the tests which allowed for these characterizations and which guided our decisions on how best to collect the final data. It also details the results of the main on-telescope tests, including comparison between the spectra obtained using the different methods, their signal-to-noise ratios (SNR), systematic changes introduced by the BRR readout mode, and other limitations of our setup.

For each of the data runs that were carried out, the data collection process generally consisted of:

1. Setting up the equipment at the McKellar Spectrograph and remotely connecting to the computer that runs our image-taking code through the 1.2-m dome control room (except for the dark testing run in Section 3.2.3);
2. Filling the two cryostat tanks with LN2, attaching vacuum pumps to the cryostat body and to the cold plate tank, and waiting for the detector temperature to reach ∼65K before beginning images;
3. Aligning and focusing the spectrum if necessary;
4. Taking the necessary exposures to achieve the goals of the run and storing the resulting images and temperature logs on the Canadian Advanced Network For Astronomy Research (CANFAR) platform[1].

3.2 Engineering Data Collection

Three engineering runs were carried out with the intention of fulfilling the requirements outlined above. Each of these runs are detailed here.

3.2.1 December 2023: Physical Setup Test

December 2023 was the first time we were able to set up our equipment on the McKellar Spectrograph. Because we only had use of the spectrograph during the day and had to allow regular nighttime observing, we used this as an engineering run to test the processes needed to set up the system and obtain the data we want. Primarily, this included finding physical places to put all the different components of our system in the Coudé room, ensuring that the cryostat LN2 tanks could be filled in its upward-facing orientation, finding the correct configurations necessary to align and focus the spectra, and finding a way to create the light source that could mimic an astronomical target with sky lines superimposed.

Most of the components of ODESSA are shown in Figure 3.1. The ARC controller and temperature controller, which we use as a temperature logger for the detector and radiation shield, are connected to both the cryostat and the computer running the imaging code. We remotely connect to this computer so that all imaging can be carried out from the control room. The cryostat pressure gauge measures the pressure within the body of the cryostat, which needs to be at high vacuum to maximize thermal insulation from the outer shell of the cryostat, so the cryostat body is connected to a continuously-running

---

[1] https://www.canfar.net/

Figure 3.1: ODESSA set up in the Coudé room of the McKellar Spectrograph, with major components labelled. The ARC controller and temperature controller connect directly to both the cryostat and the computer (out of view). The two vacuum pumps connect to different parts of the cryostat, where they are both monitored by pressure gauges. The fill tube of the LN2 transfer dewar is used for both of the cryostat's LN2 tanks.

turbo pump. The cryostat cold plate LN2 tank is evacuated with a scroll pump for the procedure discussed in Section 2.5.3; this pressure is monitored by a second pressure gauge. The LN2 transfer dewar can be connected to both of the cryostat's LN2 tanks to fill them, and we found no problems with doing so while the cryostat was secured to its mount.

Alignment of the spectrograph was made easier by the Python script described in Section 2.4; the script immediately displays images once they are taken and can copy them to a shared directory, so that they can be accessed by the software typically used for alignment at the McKellar Spectrograph. With this script in place, alignment consisted of finding the precise position of the cryostat on its mount which would result in a spectrum at best focus and aligned with the detector pixels; i.e., the spectrum should take up as few pixel rows as possible, and the centre of the spectrum should fall along a single row for its entire length to simplify spectral extraction. We also selected a grating angle which would place the Ar lines of interest in the centre of the image. These positions were noted to speed up future alignment. We also found that, to maximize the number of UTR samples before saturation while still containing the full height of the spectrum, it was ideal to read 100 rows near the centre of the detector, skipping over the top and bottom portions of the detector. Thus a full image has a size of 100 rows × 2048 columns. We use this image size for the remainder of our testing.

The final goal of this run was to create a light source which mimics some astronomical spectrum contaminated by sky lines. To do this, we needed to place both an arc lamp and an incandescent lamp in the optical path of the spectrograph. While the McKellar Spectrograph is equipped with both lamps for calibration purposes, it does not have a way to place both sources in the beam at the same time. Therefore, instead of using the McKellar incandescent lamp, we attached a clip-on lamp to a tripod standing just outside the entrance slit, aiming it directly at the slit while the arc lamp was activated. This does not place the incandescent light directly in the optical path, so we experimented with different diffusers placed against the entrance slit. These diffusers needed to be strong enough to redirect the incadescent light into the beam, but not so strong that the arc light is washed away. We found that the McKellar diffusers which could attach directly to the slit completely removed the arc lines from the images, while the neutral density filters (NDFs) did not diffuse enough incandescent light into the beam. We were able to produce good results using an improvised setup involving two layers of a plastic bag—we dub this a PIE (Plastic-bag Improvised Elastic-secured) diffuser. A photograph of this setup is shown in Figure 3.2. A comparison of the images resulting from the different types of diffusers is shown in Figure 3.3; the superiority of the PIE diffuser is obvious. As such, we used this setup for the remaining McKellar data runs.

3.2.2 January 2024: Preliminary Demonstration

January 2024 was the first of a few data runs during which our use of the McKellar Spectrograph was uninterrupted by nighttime observing, allowing us to have a consistent setup across each day of each run. During this run, we wanted to characterize the IR background of the system as well as perform an initial demonstration of the ODESSA concept. After setting up and cooling down the cryostat, we found that a poor cryostat vacuum was preventing the detector and radiation shield from reaching as low of a temperature as we desired. Reattaching the turbo pump to the cryostat halfway through the day allowed the temperature to settle at lower values. While this meant that we were unable to take as many exposures at operating temperature as we hoped for, it did allow us to characterize

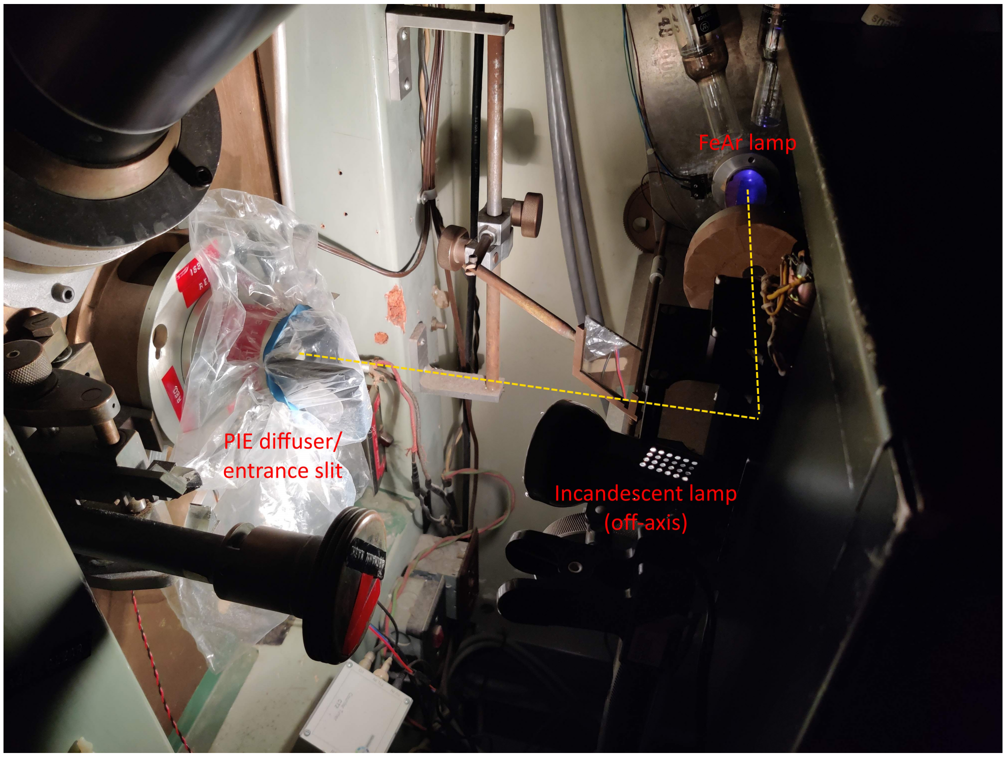

Figure 3.2: The entrance slit setup used for ODESSA. The PIE diffuser is on the left, covering the entrance slit. The clip-on incandescent lamp is in the center, and the arc lamp is in the top right. The arc light is redirected into the beam by a flat mirror. The optical path is indicated by the yellow dashed line.

the response of the background to high temperatures. While the temperature settled, we took several sequences of varying numbers of dark images, each image being composed of a reset followed by three UTR samples (∼8 sec exposure, as long as possible before saturation). A summary of each sequence is shown in Table 3.1. As discussed in Section 2.5.3, we found a strong dependence of the background on the radiation shield temperature.

During this run, we also tested the general procedure for the ODESSA concept. Knowing that the IR background can change rapidly, we preceded each sequence of spectral observations with a sequence of dark images with the same window configurations, creating one master dark image for each sequence. Rapidly changing background can also create uncertainty as to the source of differences between images that use and do not use a window; for example, it may not be possible to distinguish whether an increase in noise is caused by operating a window or by a changing background in the system. We therefore decided to make two different comparisons of datasets. First, we periodically reset an entire spectral line and compare the resulting image to one with no windows. Second, we

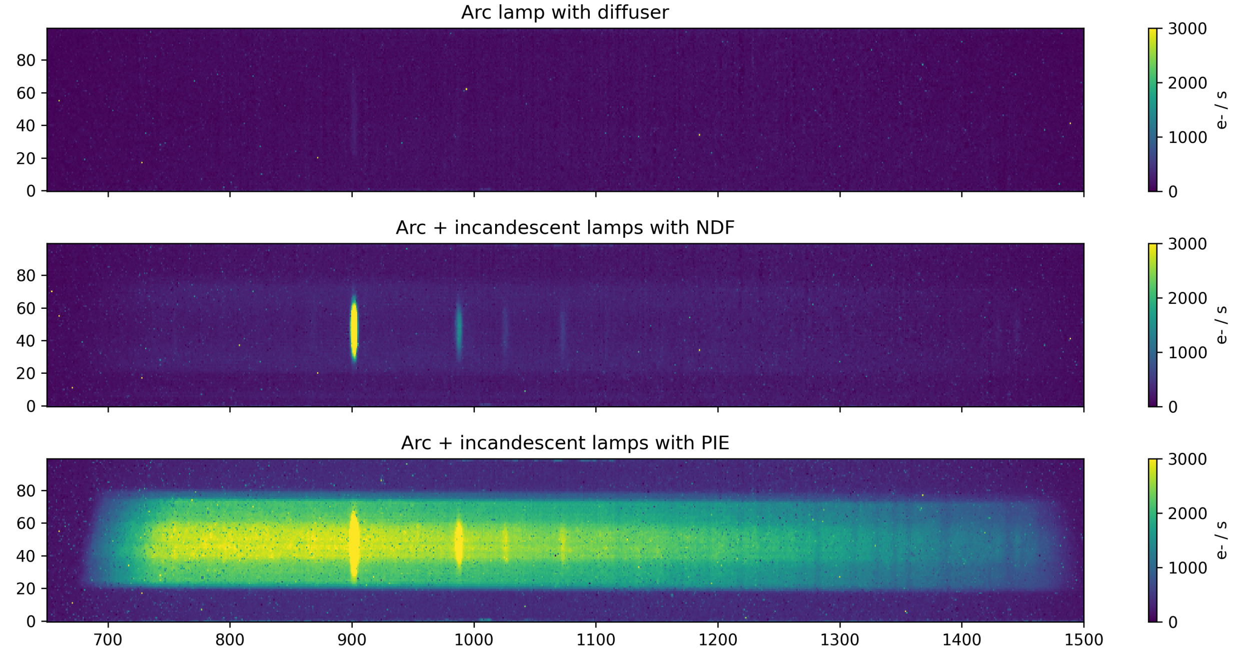

Figure 3.3: Images resulting from different diffusers. Even the lightest diffuser, as in the top image, almost completely obscures the arc lines. The neutral density filter used for the middle image has an optical density of 0.3 (transmittance ~0.5), but does not diffuse enough incandescent light. The two-layer PIE diffuser produced a clear two-component spectrum.

reset only half of a spectral line, and separately analyze the top and bottom halves. An illustration of this second technique is shown in Figure 3.4. This approach has the benefit of comparing the BRR mode to regular techniques at *exactly* the same instant under the same conditions.

Finally, an undiagnosed bug caused the ARC controller to freeze when running too large of a window, so we use two adjacent windows of the same size to reset the entire

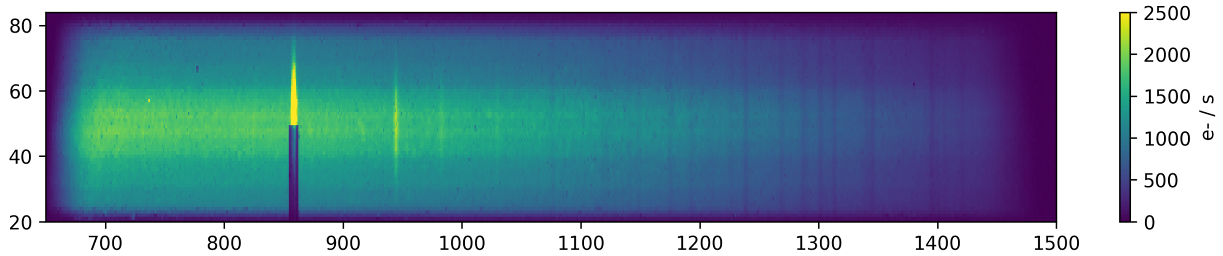

Figure 3.4: Resetting only half the 1108 nm line so that the top and bottom halves of the spectrum can be compared directly.

line at once. This means that each window is only sampled at half the frequency as when operating only one window (i.e., 50 times per full frame, about 18 Hz), but this is still plenty of UTR samples to precisely measure the window flux. The window size we use, which contains the entire half-emission line, is 32 × 12 px.

The observing sequences are summarized in Table 3.1. Although this data collection run was useful in testing procedures for ODESSA, the detector temperature was too high to obtain as ideal of a dataset as we wanted for a full demonstration. We instead used these data to develop the reduction software described in Section 2.4, leaving the actual demonstration for a later run.

Sequence start (PST)	Sequence type	Images in sequence	Det./RS temp. (K)	Windows	Notes
			2024-01-10		
12:41	dark	10	79/113	–	
12:46	dark	100	78/113	–	
13:28	dark	41	73/107	–	Refilled RS LN2
13:46	dark	30	72/107	–	Refilled RS LN2
14:57	dark	50	75/104	–	Attached turbo pump to cryostat
15:13	dark	10	74/87	–	
15:16	dark	50	74/86	–	
15:32	dark	50	73/83	–	
15:45	dark	50	72/83	–	
16:43	dark	37	68/83	–	
16:58	dark	15	68/83	–	
17:04	dark	15	68/83	–	
			2024-01-11		
14:40	dark	10	76/86	–	
14:44	spectrum	10	76/86	–	
14:56	dark	7	75/86	[1]	
15:03	spectrum	10	75/86	"	
15:10	dark	10	75/86	[1,2]	
15:14	spectrum	10	74/87	"	

[1] Over lower half of 1108 nm line, resetting every 50 rows

[2] Over upper half of 1108 nm line, resetting every 50 rows

Table 3.1: Observations for January 2024. Sequence type “spectrum” refers to an image containing both the arc and incandescent light. Occasionally, sequences would be cut short by a bad image which required the ARC controller to be reset.

3.2.3 February 2024: Dark Testing

The January data were useful in characterizing the IR background for high temperatures, but we also wanted to observe the behaviour of the background at operating temperatures, and whether operating a window had any effect on the background. To do this, we blocked the exterior window of the cryostat with copper tape, sealing out most of the 300K blackbody radiation from the room and leaving just the radiation produced within the cryostat.

We took several sequences of dark images within the dark room of the electronics lab at HAA, each image consisting of a reset and four UTR reads. Each sequence either had no window, one 32 row × 12 column window only being read, or one 32×12 window being reset every 50 rows. This window was placed on either a bright, medium, or faint region of the background feature, as shown in Figure 3.5. An electrical grounding issue resulted in the detector temperature gauge not returning any data, but the radiation shield temperature varied by < 0.1K, so we assume the detector temperature also remained stable between 60 and 70K. The sequences are summarized in Table 3.2.

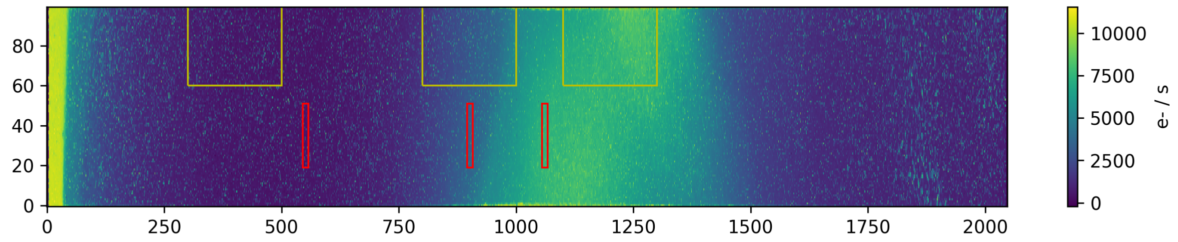

Figure 3.5: Locations of windows and regions used in dark testing. Windows are outlined in red, and regions used to measure the background flux are outlined in yellow.

The data were converted into flux images using the procedure defined in Section 2.4 (albeit without the dark-subtraction step). We ended up not using the window reads, but instead observed the change in background over time in three 40 × 200-px regions of the detector: one centred over an area with high background flux, one with low background flux, and one with an intermediate flux. The regions are shown in Figure 3.5. To avoid potential interference from windows, which are known to affect the rows and columns in which they are operating (Bezawada & Ives, 2006; Smith & Hale, 2012; Chapin et al., 2022), we sample the background only in the top 40 rows of the image. The intermediate-background region unavoidably had to overlap some of the window columns, but enough columns were sampled to minimize any window column effects. The measured change in background flux for each region in each exposure is shown in Figure 3.6.

Sequence start (PST)	Sequence type	Images in sequence	Det./RS temp. (K)	Windows
2024-02-15				
13:18	dark	20	–/82	–
13:26	dark	20	–/82	bright, read only
13:32	dark	20	–/82	medium, read only
13:38	dark	20	–/82	faint, read only
13:46	dark	20	–/82	bright, reset every 50 rows
13:52	dark	20	–/82	medium, reset every 50 rows
13:58	dark	20	–/82	faint, reset every 50 rows

Table 3.2: Observations for February 2024.

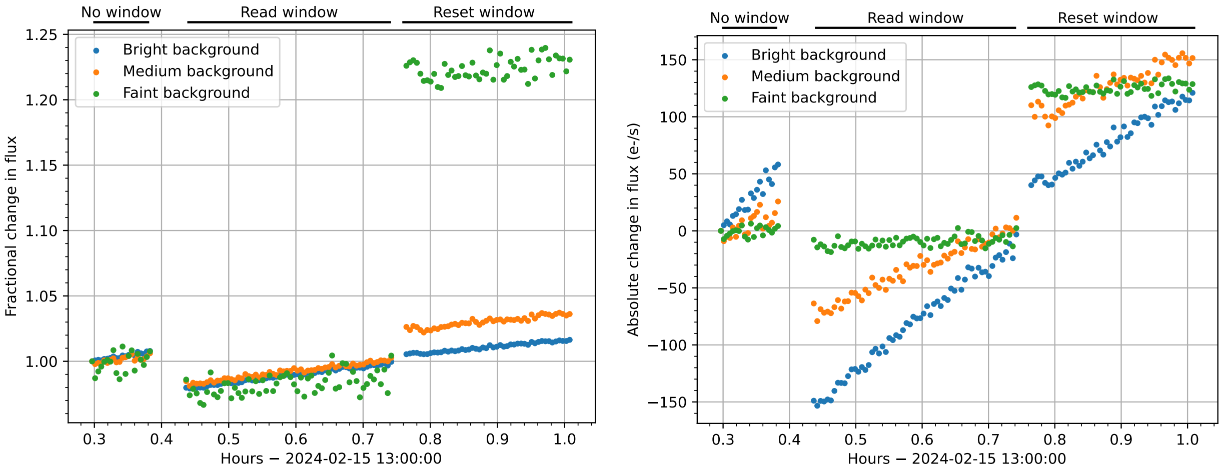

Figure 3.6: Relative and absolute change in median background flux during data collection for three different regions of the detector. The bright, medium, and faint regions had initial median fluxes of 7400, 4200, and 600 e$^-$/s, respectively.

Each region exhibits a flux which increases at a constant relative rate over time when continuously taking exposures, about 6.6% / hr or 4.6% over the full duration of the data collection. The consistent relative rate across the detector suggests that the source of the changing flux is likely changing thermal emissions within the cryostat. The measured radiation shield temperature increased from 81.54K to 81.60K over the course of the exposures; integrating blackbody Planck spectra from 1 µm to the detector cutoff of 5.3 µm at these temperatures, assuming perfect emissivity and detector response, results in a total flux increase of 1.7% due to change in temperature. Other sources of varying thermal radiation which may account for the remainder of the increasing flux could be the inner body of the cryostat or the copper tape covering the external window, though we have no temperature measurements to confirm either of these.

The effect of the drifting background flux on our data reinforces the need for mitigation by interleaving dark images throughout the data collection process and creating a new master dark image as often as possible. Additional features in the background flux seen in this data run and others are further discussed in Section 3.4.2.

3.3 Final Data Collection

Our final time slot at the McKellar Spectrograph was dedicated completely to obtaining the main images needed for the ODESSA demonstration. In order to build up as large a statistical sample as possible, we spent two full days collecting the image sequences listed in Table 3.3. While we employed the same strategy as Section 3.2.2 of interleaving dark with light images, we also included a greater variety of images and window configurations.

As in January, we included images with no window, a window resetting the lower half of the 1108 nm Ar line, and two windows resetting the full line. In addition, we included equivalent images in which the window was being read only, so that we might be able to distinguish between any features in the images caused by resetting a window and only reading it, such as seen in Section 3.2.3. We also included a sequence demonstrating reset of both arc lines at once. For each of these window configurations, we took dark and arc plus incandescent exposures as in January, as well as arc-only and incandescent-only exposures. As well as being used for calculating spectrum wavelength solutions, the arc-only images can be used to represent an artificial sky spectrum, while the incandescent-only images can represent the "true" spectrum of our artificial target without sky contamination. In the interest of time, we only took incandescent-only exposures for images without any windows.

For all exposures, we continued to include a reset followed by three UTR reads, with a 100×2048-pixel frame. We achieved better focus during this run than in January, so the Ar lines occupied fewer columns on the detector for these images. Therefore, we were able to decrease the window sizes to 32×8 pixels.

After cooling down and aligning the spectrograph, we found that the arbitrary offset of the analog-to-digital converters (ADCs) was drifting over time, causing stored pixel values to overflow or underflow their 16-bit unsigned integer storage. Because the drift was rapid enough to significantly change over the course of a single image, the saved UTR images would not have reliable values, so we discarded sequences in which this occurred. The

Sequence start (PDT)	Sequence type	Images in sequence	Det./RS temp. (K)	Windows	Notes
			2024-04-17		
11:44	dark	–	62/87	–	Discarded due to ADC offset drift
11:50	spectrum	–	62/87	–	
12:02	arc	–	62/88	–	
13:20	dark	20	63/85	[1], read only	Refilled RS LN2
13:26	spectrum	20	"	"	
13:32	arc	19	"	"	
13:39	dark	19	"	[1], reset[a]	
13:47	spectrum	20	"	"	
13:54	arc	20	"	"	
14:01	dark	20	"	[1,2], read only	
14:07	spectrum	20	"	"	
14:13	arc	20	"	"	
14:20	dark	20	"	[1,2], reset[a]	
14:26	spectrum	20	"	"	
14:32	arc	20	"	"	
14:40	dark	20	"	–	
14:46	flat	20	"	–	
14:54	spectrum	20	"	–	
15:01	arc	20	"	–	
15:07	dark	20	"	[1], read only	
15:16	spectrum	20	64/85	"	
15:22	arc	20	64/86	"	
15:29	dark	19	"	[1], reset[a]	
15:35	spectrum	20	"	"	
15:41	arc	20	"	"	
15:48	dark	20	"	[1,2], read only	
15:54	spectrum	20	"	"	
16:01	arc	19	"	"	
16:07	dark	20	"	[1,2], reset[a]	
16:13	spectrum	20	"	"	
16:19	arc	20	"	"	

[1] Over lower half of 1108 nm line

[2] Over upper half of 1108 nm line

[a] Every 50 rows

Table 3.3: Observations for April 2024. Sequence type "spectrum" refers to an image containing both the arc and incandescent light; "arc" is arc only, and "flat" is incandescent only. Eight sequences had a corrupted image which was discarded. Table continues on next page.

Sequence start (PDT)	Sequence type	Images in sequence	Det./RS temp. (K)	Windows	Notes
			2024-04-18		
10:26	dark	20	62/85	–	
10:33	flat	20	"	–	
10:39	spectrum	20	"	–	
10:45	arc	20	"	–	
10:52	dark	20	"	[1], read only	
10:59	spectrum	20	"	"	
11:05	arc	20	"	"	
11:11	dark	20	"	[1], reset[a]	
11:20	spectrum	20	"	"	
11:26	arc	20	"	"	
11:32	dark	20	"	[1,2], read only	
11:38	spectrum	20	"	"	
11:44	arc	20	"	"	
11:50	dark	19	"	[1,2], reset[a]	
11:56	spectrum	19	"	"	
12:02	arc	20	"	"	
13:38	dark	20	63/86	–	New masters
13:43	flat	20	"	–	
13:49	spectrum	20	"	–	
13:55	arc	20	"	–	
14:01	dark	20	"	[1], read only	
14:07	spectrum	20	"	"	
14:13	arc	20	"	"	
14:19	dark	20	"	[1], reset[a]	
14:25	spectrum	20	"	"	
14:30	arc	20	"	"	
14:36	dark	19	"	[1,2], read only	
14:45	spectrum	19	"	"	
14:51	arc	20	64/86	"	
14:57	dark	20	"	[1,2], reset[a]	
15:03	spectrum	20	"	"	
15:08	arc	20	"	"	
15:24	dark	20	"	[1,2], reset[a]; [3,4], reset[b]	
15:30	arc	20	"	"	
15:36	spectrum	20	"	"	

[1] Over lower half of 1108 nm line

[2] Over upper half of 1108 nm line

[3] Over lower half of 1111 nm line

[4] Over upper half of 1111 nm line

[a] Every 50 rows

[b] Every 100 rows

drift stopped after the first three sequences and was not observed again. The cause of this behaviour remains undetermined.

Each of the consecutive three or four sequences of images with the same window configuration has its own master dark image and bad pixel map. Three master arc and flat images were created, one for 17 Apr images, and one each for the first and second halves of 18 Apr images. The data were reduced and spectra were extracted using the procedures defined in Section 2.4. The spectra were combined to produce one coadded spectrum per window configuration, for a total of eight final spectra which are analyzed in the next section.

3.4 Results

Each of these eight spectra is shown in Figure 3.7. In this section, I refer to each spectrum by its number shown here here.

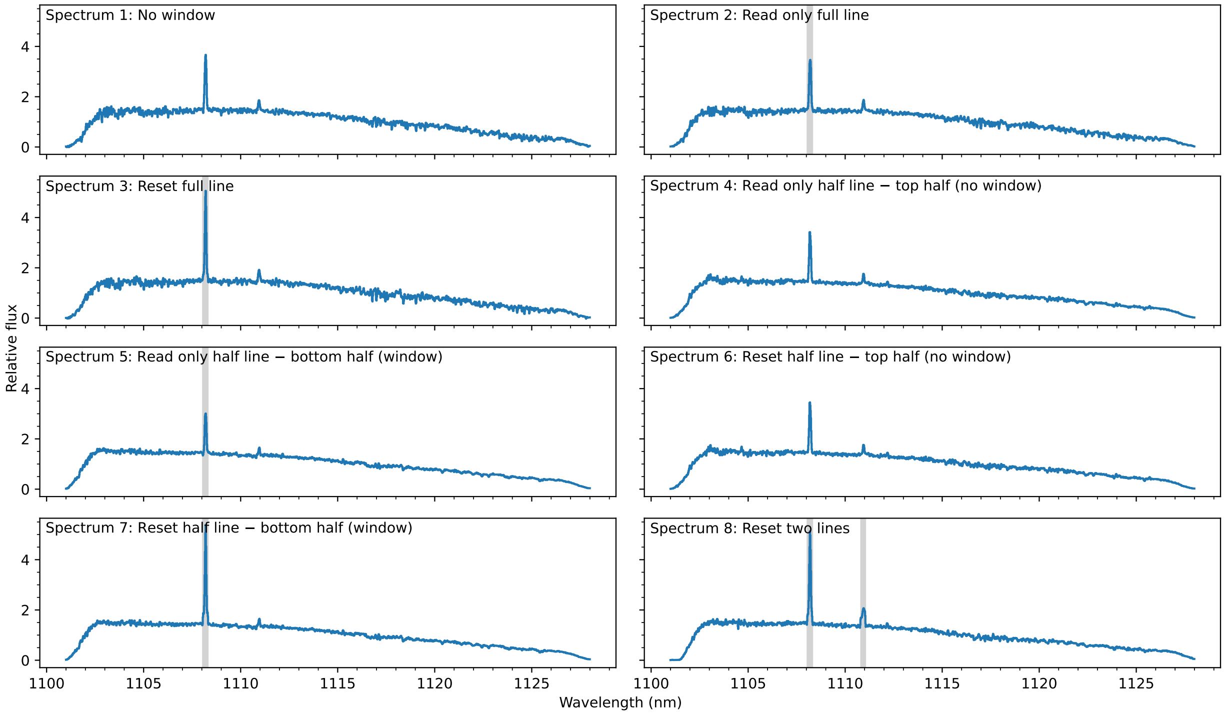

Figure 3.7: The eight spectra produced. Regions which are covered by a window are shaded in grey in their respective spectra.

Spectra 1 and 3 are shown side-by-side in Figure 3.8. Immediately we can see that the overall shape of the spectrum is unaffected by the use of the BRR mode. To aid in uncovering any differences, I also show the ratio of the two modes (BRR / no-window) in

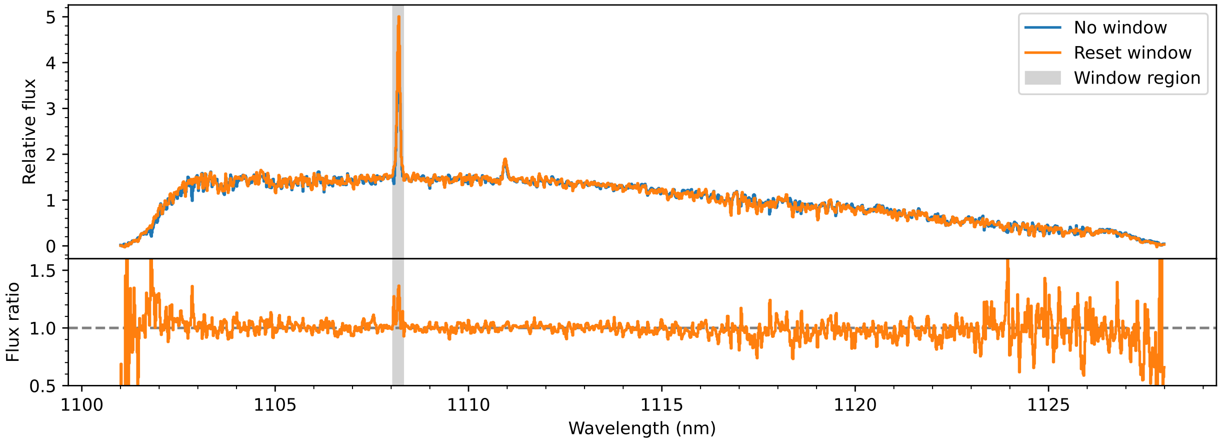

Figure 3.8: Comparison of spectra using no-window mode and BRR mode. The bottom panel shows the ratio of the two (BRR / no-window mode).

the lower panel of this plot. The ratio reveals that the BRR mode results in a slightly brighter emission line, for reasons discussed in Section 3.4.2. On either end of the spectrum, deviations of the ratio from unity are caused by slight shifts in the edge of the illuminated region of the detector between image sequences. These shifts are physical shifts of the detector in the optical path, so they can be safely ignored for this experiment.

Similarly, we can compare the two portions of the half-window-reset configuration described above to directly test the two modes on the same spectrum, without having to consider variations between exposures. The spectra are shown in Figure 3.9. As before, the only significant difference between the two spectra is the increased flux in the window being reset. A comparison between BRR mode and a window which is read but not reset yields the same results, as does a comparison against the exposures that use four windows to reset both emission lines. The consistency between the spectra in each exposure mode (barring the increased flux in reset windows) suggests that BRR mode could indeed be useful in actual observations.

3.4.1 Signal-to-Noise

While a visual comparison of the spectra is very intuitive, a quantification of the differences between them requires looking at the signal to noise as a function of wavelength in each mode. This is achieved by dividing the final spectra by their respective error spectra as created by the process defined in Section 2.4. The signal to noise for each window configuration is shown in Figure 3.10, and a comparison between Spectra 1 and 3 is shown in Figure 3.11. The SNR remains largely the same outside the window between modes; the

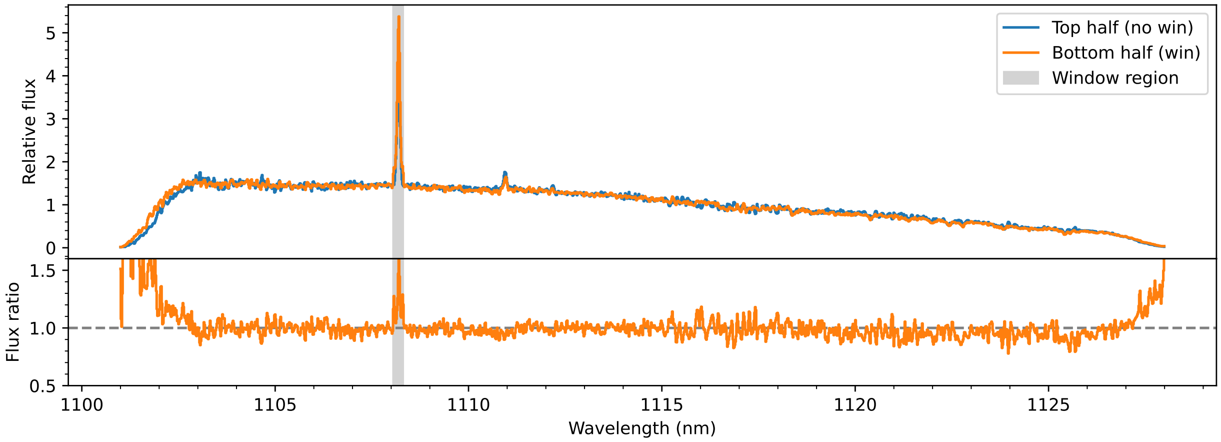

Figure 3.9: Same as Figure 3.8, but instead comparing the two portions of the half-window-reset mode.

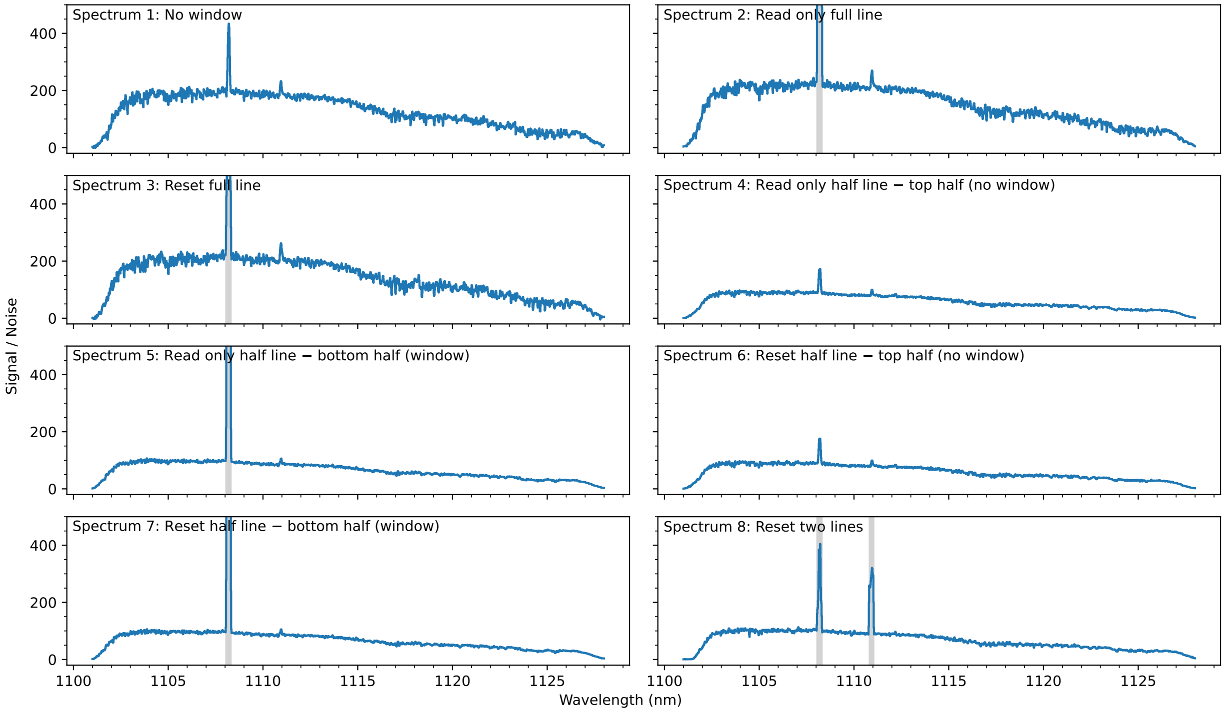

Figure 3.10: Each configuration's signal to noise. Regions which are covered by a window are shaded in grey in their respective spectra.

slight decrease in SNR for the no-window mode is a result of discarding the first sequence of images.

The operation of windows during an exposure significantly increases the SNR within the window, as seen in Figure 3.12. This change is due to the number of UTR samples to

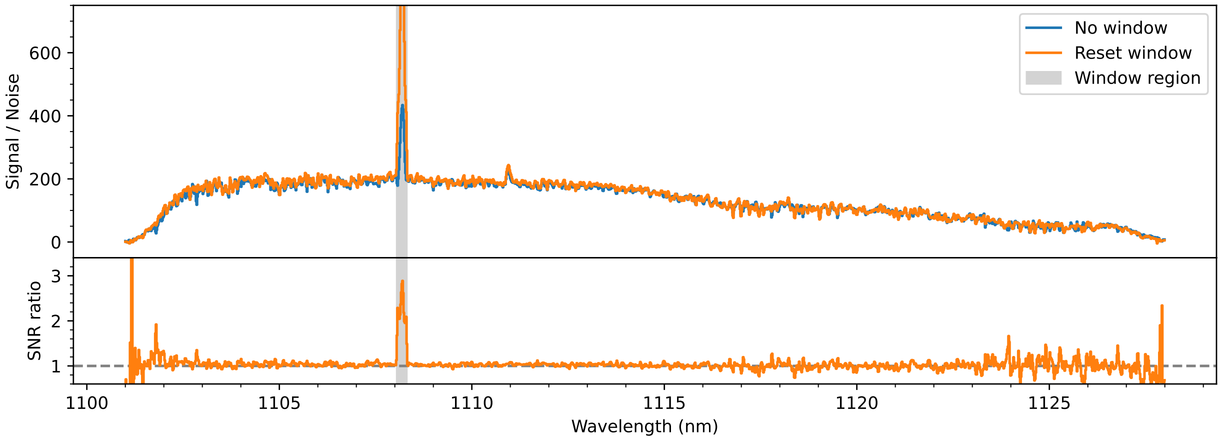

Figure 3.11: Same as Figure 3.8, but instead comparing SNR.

which a flux is fit. Without a window, the variance of the slope is large due to a limited number of samples, while reading a window dramatically increases the number of samples (by a factor of up to 100 in this case). Resetting the window periodically (as is necessary for bright sky lines) increases the uncertainty somewhat, due to fitting slopes to several shorter ramps instead of one long ramp. However, even with this increase in noise, the BRR mode still has a SNR significantly higher than the no-window mode, demonstrating the ability of the BRR mode to improve monitoring of quickly varying sky lines.

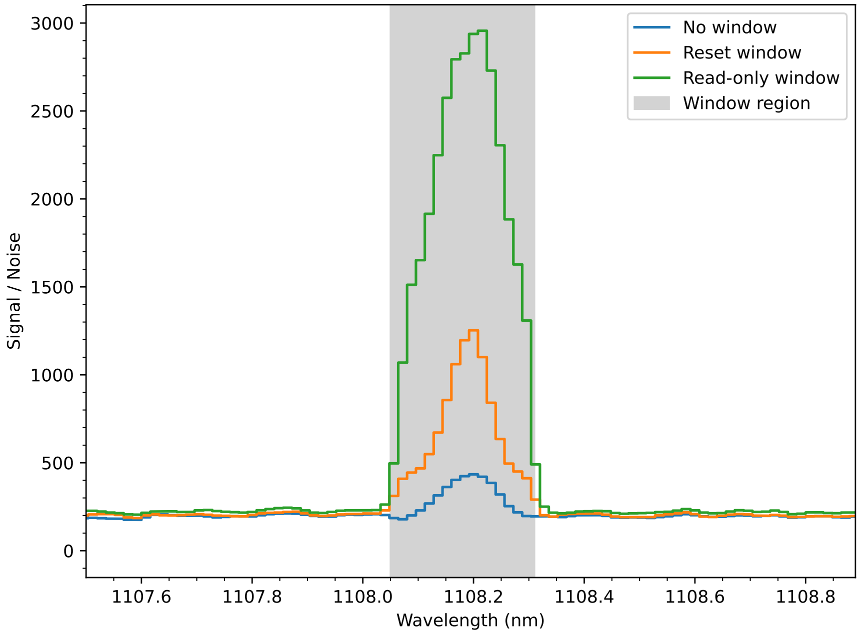

Figure 3.12: Comparison of SNR in the 1108 nm line.

3.4.2 Systematics Introduced by Readout Mode

Nonlinearity

Figure 3.8 showed that the measured flux within a window which is being reset is higher than the measured flux without any resets. The most likely cause for this effect is pixel nonlinearity. As charge accumulates within a pixel during an exposure, that charge effectively decreases the bias voltage, resulting in a lower measured charge within the pixel (Plazas et al., 2017). Nonlinearity therefore becomes more significant at higher fluence.

As mentioned in Section 2.4, a typical data reduction process would include a pixel linearization step, which we did not include. Figure 3.13 shows a pixel with high flux exhibiting significant nonlinearity. We attempted a nonlinearity correction by fitting a second-order polynomial to each pixel's ADU ramp in a median dark image. The quadratic term of this polynomial is then added back into reference-subtracted ramps, so that the data follow the linear-only terms. Figure 3.13 shows the best-fitting slope to the ramp

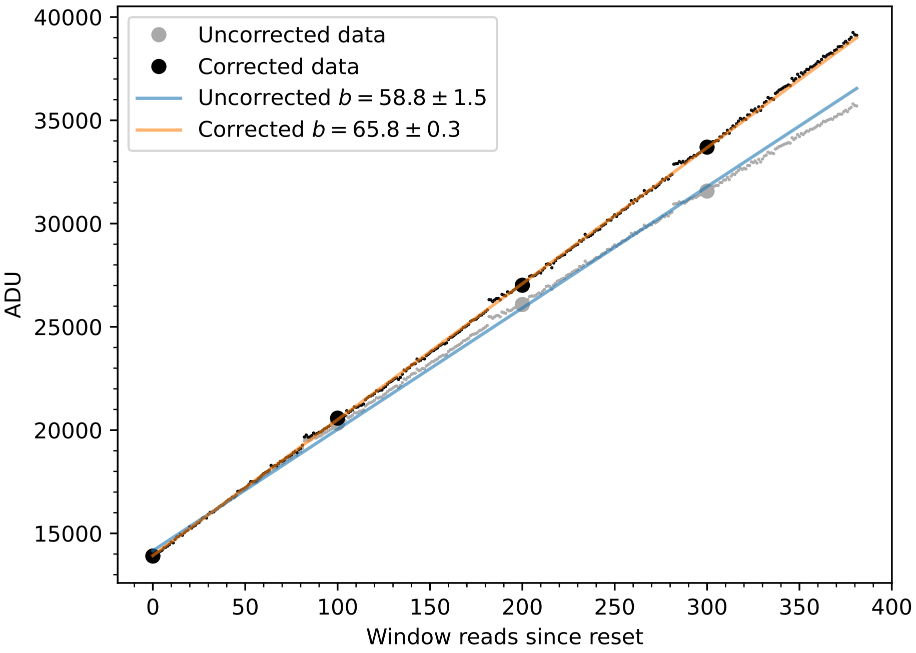

Figure 3.13: Typical pixel nonlinearity. Large points are full frame reads, and small points are window reads for the same pixel. The best-fitting OLS slopes and uncertainties for the linearity-uncorrected and corrected full-frame ramps in ADU / window period are shown as well.

both before and after linearization; the corrected slope is 12% higher than the uncorrected slope, indicating that a linearization step is necessary for accurate results. Averaged across the window, the corrected flux is 18% higher than the uncorrected flux, which is similar to the increased flux seen in the window region in Figure 3.8.

However, the same pixel experiences less nonlinearity when it is periodically reset by the BRR mode, as seen in Figure 3.14. Because of the resetting, the number of electrons in the pixel never reaches the nonlinear regime, so each of the seven measured ramps, when combined, result in a flux which is already quite close to the ideal, linearized flux. Therefore, we expect windows in BRR mode to have a higher measured flux than the same image in no-window mode, as is seen in the results above.

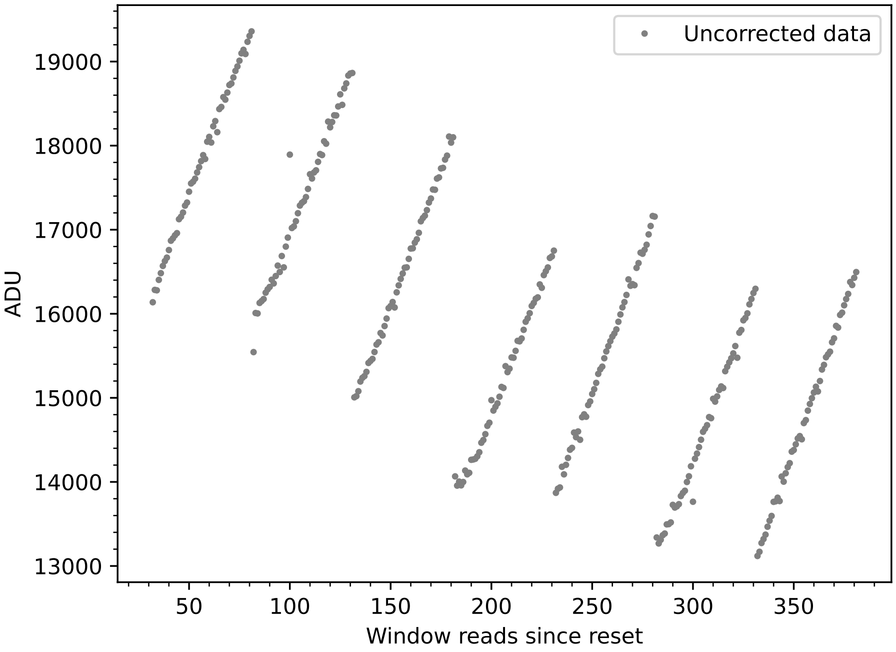

Figure 3.14: (Lack of) nonlinearity in a pixel which is periodically reset. The fluence in the pixel never reaches the nonlinear regime.

Unfortunately, we cannot uniformly apply this nonlinearity correction across the entire detector. As a result of the high background flux, for most pixels in our data we can only fit the four full-frame UTR reads for each image. This becomes a problem due to a feature dubbed the "reset anomaly", which is shown in Figure 3.15 (Bezawada & Ives, 2006; Rauscher et al., 2007; Langevin et al., 2022). The anomaly is characterized by an abrupt increase in a pixel's apparent flux directly after being reset, before flattening to a more

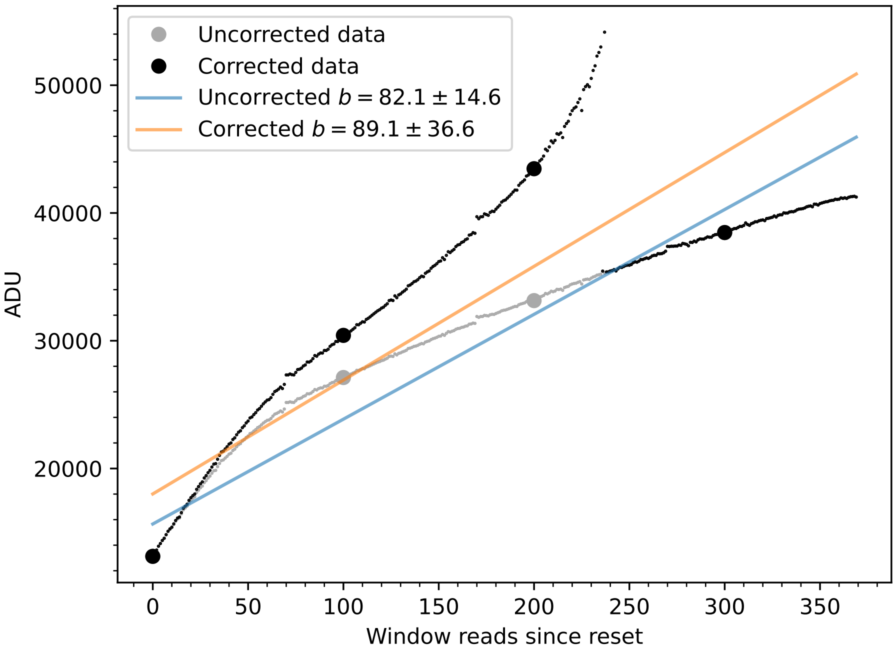

Figure 3.15: Example of the reset anomaly in one ramp, and the naive linearity correction applied to it.

linear rate. The 2nd-order polynomial fit to pixels with the reset anomaly overestimates the nonlinearity correction, producing wildly inaccurate "corrected" ADU values.

Bezawada & Ives (2006) and Rauscher et al. (2007) speculate that the reset anomaly is due to settling charges within the ROIC after a pixel is reset, and Rauscher et al. (2007) demonstrate that the anomaly is repeatable and well-fit by a four-parameter model including a linear and an exponential term. Our restriction to only four samples per ramp makes this model of little use to us, though, and the model would also need an additional parameter to account for nonlinearity at high count levels. However, an implementation of the ODESSA concept with lower IR background would have the ability to take exposures many times longer than we have here, allowing us to fit a nonlinearity correction. We thus expect that the inability to correct for nonlinearity and the reset anomaly is unique to our specific experiment.

Qualitatively, the increased flux within windows can be explained by the nonlinearity described above, and the effect is on the same order as the observed change. The data collected for this experiment, though, are insufficient to properly correct for nonlinearity, so future work needs to confirm whether nonlinearity is the sole cause.

Background flux

Several authors have already noted that the operation of windows can affect the pixels along the windows' rows and columns. Reset level offsets are obvious along these rows and columns, and offsets in flux can also be seen to a lesser extent, though Smith & Hale (2012) and Chapin et al. (2022) emphasize that these offsets are all repeatable and can be calibrated out. Smith & Hale (2012) and Regan & Bergeron (2020) also demonstrate that high-cadence windows can exhibit significant self-heating, which produces a localized glow affecting pixels near the windows (Tam & Hu, 1984).

However, Figure 3.6 shows that, during the operation of windows, there is an offset in flux even in pixels not along rows or columns containing a window. We find that the offsets are not localized to pixels near windows, nor are they repeatable across data runs. While Figure 3.6 showed a decrease in flux across the detector while reading windows and an increase while resetting them in February, Figure 3.16 shows that, in January data, flux increases for all images with windows; and in April data, flux remains the same while reading windows and *decreases* when resetting them. This behaviour is difficult to explain by any sort of multiplexer glow mechanism.

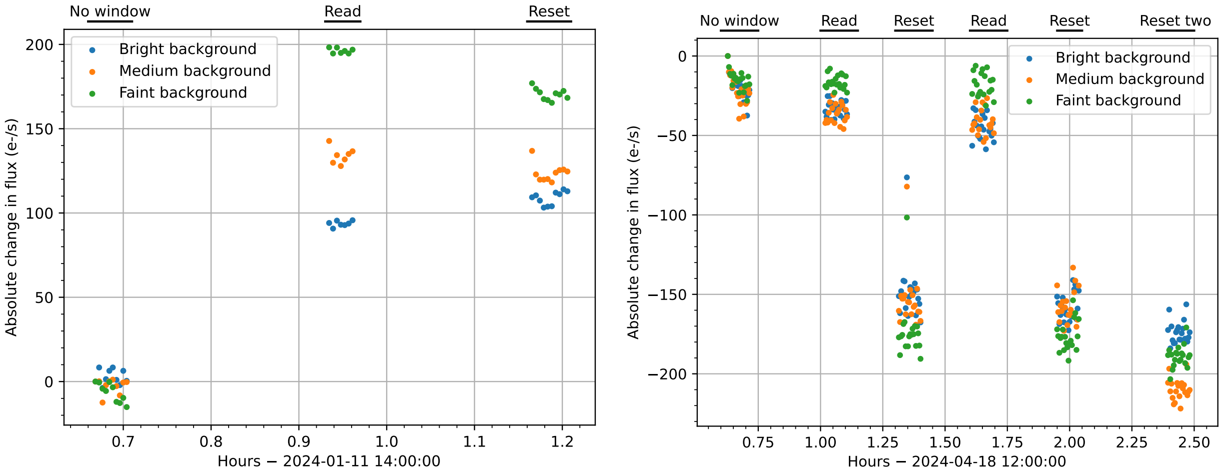

Figure 3.16: Change in background dark images for January and April data in different regions of the detector with different window configurations. Backgrounds are measured in the same way as Figure 3.6.

The change in flux additionally cannot be explained by varying reset levels along a nonlinear pixel ramp—there is no association at all between image flux and reset offsets—and blackbody temperature variations should cause smooth changes in flux which are proportionally uniform as seen in Figure 3.6.

The ultimate cause of this behaviour remains unclear, although it accompanies other odd characteristics associated only with the H2RG detectors and ARC controller at HAA (Chapin et al., 2022, 2024, in press). A potential cause could be shifting voltage levels within the ROIC, which could be associated with various grounding issues that were encountered over the course of the experiment. Bugs in the controller could also have an effect, as other groups investigating HxRG guide windows who have not seen this issue have not used ARC Gen-4 controllers. Regardless of the cause, the flux changes do not cause a noticeable effect in our final results, and we conclude that the magnitude of the effect is beneath the current noise level.

3.4.3 Other Limitations

Other limitations of this experiment are related to non-ideal hardware. Our selection of cryostat, which followed from the constraints of the McKellar Spectrograph, resulted in an operating temperature which was higher than ideal, leading to degraded detector performance and a considerable IR background.

As was seen in Section 2.5.3, the number of bad pixels on the detector increases quickly with temperature, with around 27% of pixels being masked in the final data runs. The loss of more than a quarter of the data collected makes identification of structured systematic effects more difficult, since such structures may be hidden behind densely-masked regions of images. Such large numbers of bad pixels also leads to the masking procedure inevitably missing some pixels, making any differences between noise caused by poor image quality and noise caused by the use of BRR mode less clear. Although no major reductions of image quality were found when using the BRR mode, the relatively poor initial quality of the images should be kept in mind.

Additionally, the warm temperature of the cryostat and imperfect IR blocking led to high background in all images. Aside from the additional major contribution of shot noise to the data, this required us to limit our exposures to $<$10 seconds each. On-sky observations, on the other hand, can reach exposure lengths of up to 10 minutes before saturation. A more comprehensive experiment might have used exposure lengths beyond what would be typical for this wavelength range, demonstrating directly the ability of BRR mode to allow long exposure times in spite of bright emission lines.

3.5 Summary

We have shown that resetting bright emission lines to prevent their saturation is a sensible concept for NIR observations that does not result in any significant degradation in the spectra that are obtained. The spectra produced by the BRR mode are nearly identical to those produced by a standard UTR imaging mode, with the added benefit of increased signal-to-noise within windows. The differences that we do identify can be attributed to the limitations of our setup, specifically regarding nonlinearity and the high IR background that is part of our images.

Chapter 4

Conclusions and Next Steps

The previous chapters have detailed the procedure necessary to carry out on-detector prevention of the saturation of atmospheric OH emission lines. Using an H2RG detector and custom software with the McKellar Spectrograph, we completed a successful demonstration of the use of BRR mode to achieve this. Compared to a more conventional observing method, the BRR mode does not introduce any measurable systematic changes outside the windows, and the increased flux within the window is attributable to correctable non-linearity effects specific to our setup. Therefore, resetting regions of infrared detectors shows promise as a method to limit the impact of bright OH lines on spectra. The work detailed here justifies further development and testing to quantify the full value of implementing the BRR method at major observatories. In this chapter, I discuss how the BRR method compares to some other methods and explore some potential avenues for further development.

4.1 Comparison to Existing Strategies

Some existing strategies involve avoiding OH emission lines completely. In comparison to narrowband filters, high dispersion masking, and fiber Bragg gratings, BRR mode immediately allows observations to use information at wavelengths covered by the emission lines. Despite allowing these bright lines to fall on the detector, BRR mode still allows longer exposures by resetting the regions at risk of saturation, reducing the read noise introduced by stacking many short exposures. The preservation of the observed spectrum outside the windows in our experiments suggests that, even if nothing were gained from the information in the windows, BRR mode would perform just as well as any of these hardware-based solutions. BRR mode also has the advantage of not requiring any custom-built optical components; it can be implemented in software and firmware for any detector-controller combination capable of guide window mode (which most upcoming IR detectors are).

To compare the utility of the BRR mode in improving existing, or developing new, methods of sky subtraction for NIR spectra, it will be necessary to test this concept further. As discussed in the previous chapter, some limitations in our setup prevent us from taking images that are directly comparable to on-sky images from premier observatories located at excellent sites such as Maunakea. The bright IR background from the Coudé room and the large number of bad pixels on the detector need to be overcome to further reduce uncertainty of the impact of windows. Implementing this concept on a premier spectrograph with simultaneous target and sky spectra would make it possible to devise an experiment which directly demonstrates the ability of the BRR mode to improve the achievable quality of sky-subtracted spectra. The end goal would be to compare signal-to-noise and systematic uncertainties of the final sky-subtracted spectra using BRR mode to more conventional methods.

It is unlikely that the effect of using windows in science observations should be very large. The odd behaviour related to the windows noted here and in Chapin et al. (2022, 2024, in press), such as measured flux offsets, seems to be limited to the detectors and controller being used at HAA. Systematic effects which have been reported elsewhere, such as row and column count offsets and self-heating, are more repeatable and can be completely removed with calibration images (Bezawada & Ives, 2006; Smith & Hale, 2012). It is not critical to BRR mode to visit windows at a very high cadence, so well-chosen read and reset rates may even reduce self-heating effects below what has been observed previously. This also alleviates the concerns of Baril et al. (2006), who discourage the use of windows for guiding due to the extra time it takes to interleave windows with the full frame.

As a simple example of what BRR might look like on a major telescope, I use the Gemini Near-InfraRed Spectrograph (GNIRS) integration time calculator[1] to estimate the observed spectrum of a science target contaminated by the sky background. Figure 4.1 shows the spectra of an H-magnitude 20 A0-type star and the estimated background for a 5-minute long-slit exposure in the H-band. This exposure assumes “good” image quality, “clear” cloud cover, and “low” water vapour as defined by the calculator. The calculator assumes a constant sky background brightness in the H-band of ∼14 mag / arcsec^2. In 5 minutes, the brightest OH lines reach the nonlinear regime of the detector (65,000 e$^-$) and are approaching full well (90,000 e$^-$). In a single exposure, the target spectrum continuum

---

[1] `https://www.gemini.edu/instrumentation/gnirs/exposure-time-estimation`

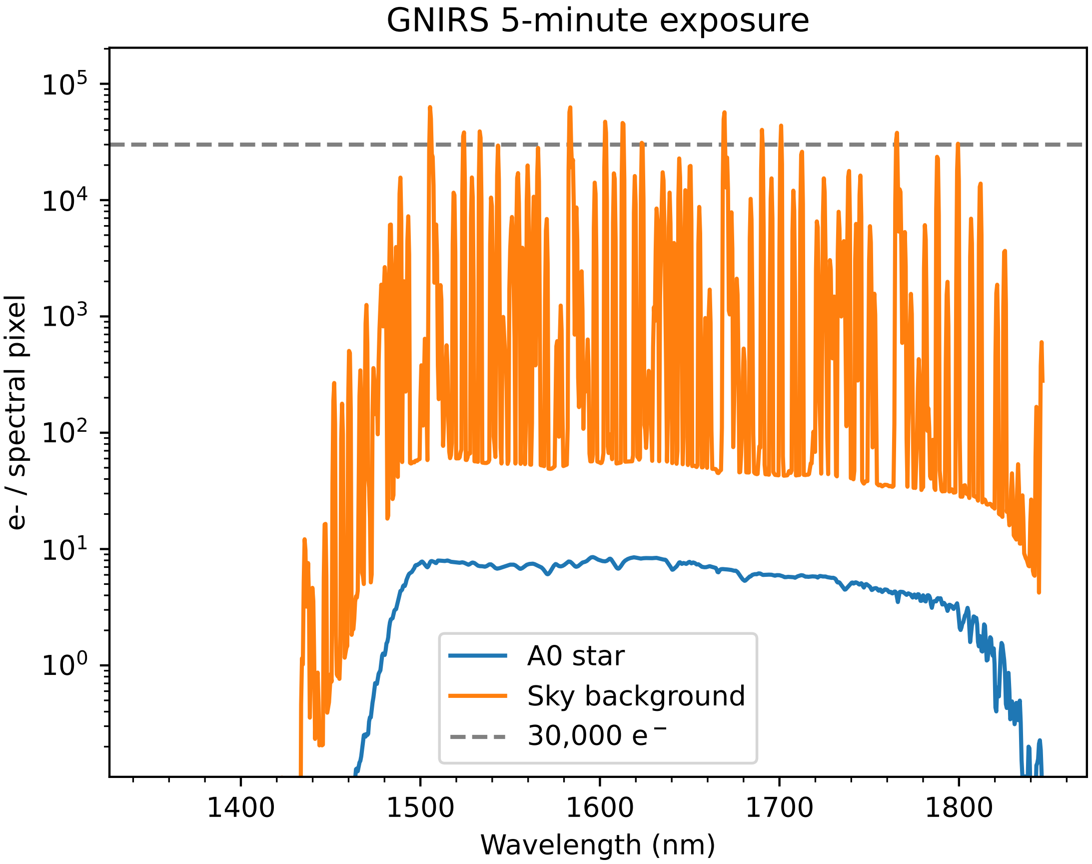

Figure 4.1: Simulated observation of a H = 20 mags A0 star with GNIRS at 1.6 µm. Placing windows over the 11 brightest lines would allow double the exposure time while keeping the entire detector within the linear regime. The dashed grey line represents approximately half the linear well depth; pixels exceeding this threshold must be reset to double the exposure length.

lies at around 3 e^- per pixel, and read noise for faint target mode is 7 e^-.

To double the exposure length on this detector without letting any OH lines saturate—thus decreasing read noise of coadded images by a factor of $\sqrt{2}$—one would need to place windows on the 11 lines brighter than 30,000 e^-, resetting each of them at least once during the 10-minute exposure. To reach even longer exposure times and lower readout noise, additional windows can be defined on the next brightest lines, which would be reset at any point before saturation. With read noise on this detector on the same order as the signal itself, observations could significantly benefit from the noise reduction offered by long exposures in BRR mode. Sky subtraction quality should be the same or better than before, as no information is lost from saturation, and the variability of lines can even be monitored at a more frequent rate.

4.2 Future Work

Future work on this project should focus on correcting for the limitations identified previously. Currently, the most significant uncertainty in our results is whether pixel nonlinearity is solely responsible for the discrepancies between spectra within windows. Therefore, we need to implement a correction for the nonlinearity to determine whether this is the case.

As previously mentioned, the reset anomaly currently prevents a satisfactory linearity correction for our data. However, Rauscher et al. (2007) find that the reset anomaly is repeatable for individual pixels, so a useful strategy may be to search for old data from the same detector which contain more UTR samples per exposure. If sufficient data exist, we could fit a correction to the old data and apply it to the data used in this thesis.

Implementing a nonlinearity correction would leave the bright IR background as the largest limiting factor. To reduce the impact of the IR background on our images, it will be necessary to explore alternative hardware setups. The cryostat used by Chapin et al. (2022) did not suffer the same background issue as the cryostat used here, and the temperature was more well-controlled then. An in-house solution could be to refit the detector to the old cryostat to mitigate temperature-related issues, though a new method of observing "sky lines" would have to be devised due to the old cryostat's inability to be mounted on the McKellar Spectrograph.

A more comprehensive solution would be to implement the necessary readout software onto an existing spectrograph which is optimized for observations in the NIR. In addition to the use of a higher-quality instrument, this would have the advantage of being able to conduct true sky observations with which we could directly compare different sky subtraction schemes using simultaneous target and sky observations. While many of the instruments mentioned in Chapter 1 are still under development, some operational NIR spectrographs already using HxRG detectors include NIFS and IGRINS-2 at Gemini (McGregor et al., 2003; Lee et al., 2022), OSIRIS and MOSFIRE at Keck (Larkin et al., 2006; McLean et al., 2012), SPIRou at CFHT (Donati et al., 2020), CRIRES and KMOS at VLT (Kaeufl et al., 2004; Sharples et al., 2013), MOIRCS at Subaru (Ichikawa et al., 2006), and NIRPS at La Silla Observatory (Wildi et al., 2017). These spectrographs are all installed at sufficiently high-quality sites to enable the on-sky demonstration we intend to carry out. Investigation will be necessary to determine which instruments are most suitable for the BRR mode, as well as the prerequisites for implementing BRR on facilities such as these. In the absence of on-sky observations, simulated comparison of various sky subtraction methods could reveal how much improvement could be realistically expected from the BRR mode. This would in turn allow a more detailed risk-benefit analysis for implementation at observatories.

4.3 Summary

In this thesis, I have described a new method for dealing with bright sky lines in infrared spectroscopy which has the potential to improve signal to noise of faint targets and improve sky subtraction. In order for observations to not be limited by bright OH emission lines from the atmosphere, it is necessary to devise ways to remove them from spectra. Existing methods require either short exposures, which become dominated by read noise, or removing all information at the wavelengths of the emission lines. The method introduced here avoids both of these shortfalls by resetting the pixels of the detector containing the lines before they are able to saturate.

Several colleagues at HAA and I designed an experimental setup to demonstrate this method. We assembled the hardware, which resulted in restrictions on the experiment such as a bright background and limited exposure times. I adapted existing software to optimize it for the BRR observing mode, allowing customized detector control as well as a data reduction pipeline designed for this experiment. I carried out several on-telescope tests demonstrating both the effectiveness of the BRR method as well as the extent to which the operation of windows affects the final data products.

I show that the BRR mode does not have a significant impact on observations apart from a feature which I argue is straightforward to correct, given that it results from our specific, imperfect setup. As a result, the method demonstrated here is worthy of further exploration, and it could potentially lead to the reduction of noise and systematics on next-generation spectrographs. Implementing BRR on modern, large telescopes could thus enable key science on the faintest objects observable.

Bibliography

Adams, N. J., Conselice, C. J., Ferreira, L., et al. 2023, MNRAS, 518, 4755, doi: 10.1093/mnras/stac3347

Aikman, G. C. L. 1976, Publications of the Dominion Astrophysical Observatory Victoria, 14, 379. https://ui.adsabs.harvard.edu/abs/1976PDAO...14..379A

Allington-Smith, J. R., Content, R., Dubbeldam, C. M., Robertson, D. J., & Preuss, W. 2006, MNRAS, 371, 380, doi: 10.1111/j.1365-2966.2006.10666.x

Andersen, D. R., Fischer, M., Conan, R., Fletcher, M., & Véran, J.-P. 2008, Proc. SPIE, 7015, Adaptive Optics Systems, 70150H, doi: 10.1117/12.789792

Artigau, É., Doyon, R., Hernandez, O., et al. 2018, Proc. SPIE, 10709, High Energy, Optical, and Infrared Detectors for Astronomy, 107091P, doi: 10.1117/12.2314475

Baril, M. R., Ward, J., Teeple, D., et al. 2006, Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, 62690Z, doi: 10.1117/12.672113

Barton, E. J., Larkin, J. E., Moore, A. M., et al. 2010, Proc. SPIE, 7735, Ground-Based and Airborne Instrumentation for Astronomy, 77355M, doi: 10.1117/12.856521

Batten, A. H. 1974, Publications of the Dominion Astrophysical Observatory Victoria, 14, 191. https://ui.adsabs.harvard.edu/abs/1974PDAO...14..191B

Bezawada, N., & Ives, D. 2006, Proc. SPIE, 6276, High Energy, Optical, and Infrared Detectors for Astronomy, 62760O, doi: 10.1117/12.670244

Birkmann, S. M., Ferruit, P., Giardino, G., et al. 2022, A&A, 661, A83, doi: 10.1051/0004-6361/202142592

Bland-Hawthorn, J., Englund, M., & Edvell, G. 2004, Opt. Express, 12, 5902, doi: 10.1364/OPEX.12.005902

Blank, R., Anglin, S., Beletic, J. W., et al. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 437, Solar Polarization 6, ed. J. R. Kuhn, D. M. Harrington, H. Lin, S. V. Berdyugina, J. Trujillo-Bueno, S. L. Keil, & T. Rimmele, 383

Böker, T., Arribas, S., Lützgendorf, N., et al. 2022, A&A, 661, A82, doi: 10.1051/0004-6361/202142589

Boss, A. P., Weinberger, A. J., Anglada-Escudé, G., et al. 2009, PASP, 121, 1218, doi: 10.1086/647960

Bushouse, H., Eisenhamer, J., Dencheva, N., et al. 2022, JWST Calibration Pipeline, Zenodo, doi: 10.5281/ZENODO.7038885

Chapin, E., Hardy, T., Grosson, T., & Lothrop, J. 2024, in press, Proc. SPIE, 13103, X-Ray, Optical, and Infrared Detectors for Astronomy

Chapin, E. L., Dunn, J., Hardy, T., Hubner, O., & Lothrop, J. 2022, Proc. SPIE, 12191, X-Ray, Optical, and Infrared Detectors for Astronomy, 121910R, doi: 10.1117/12.2627979

Davies, R. I. 2007, MNRAS, 375, 1099, doi: 10.1111/j.1365-2966.2006.11383.x

Donati, J. F., Kouach, D., Moutou, C., et al. 2020, MNRAS, 498, 5684, doi: 10.1093/mnras/staa2569

Ellis, S. C., & Bland-Hawthorn, J. 2008, MNRAS, 386, 47, doi: 10.1111/j.1365-2966.2008.13021.x

Euclid Collaboration, Jahnke, K., Gillard, W., et al. 2024, Euclid. III. The NISP Instrument, arXiv e-prints, doi: 10.48550/arXiv.2405.13493

Fixsen, D. J., Offenberg, J. D., Hanisch, R. J., et al. 2000, PASP, 112, 1350, doi: 10.1086/316626

Flagey, N., McConnachie, A., Szeto, K., Murowinski, R., & Mignot, S. 2016, Proc. SPIE, 9910, Observatory Operations: Strategies, Processes, and Systems, 99101F, doi: 10.1117/12.2230932

Fowler, A., Joyce, R., Gatley, I., Gates, J., & Herring, J. 1989, Proc. SPIE, 1107, Infrared Detectors, Focal Plane Arrays, and Imaging Sensors, 22, doi: 10.1117/12.960647

Fowler, A. M., & Gatley, I. 1990, ApJ, 353, L33, doi: 10.1086/185701

Hanuschik, R. W. 2003, A&A, 407, 1157, doi: 10.1051/0004-6361:20030885

Herzberg, G. 1951, JRASC, 45, 100. https://ui.adsabs.harvard.edu/abs/1951JRASC..45..100H

High, F. W., Stubbs, C. W., Stalder, B., Gilmore, D. K., & Tonry, J. L. 2010, PASP, 122, 722, doi: 10.1086/653715

Hilbert, B., & Rest, A. 2016, NIRCam Detector Gain Values in CV2, Tech. Rep. JWST-STScI-004622, Space Telescope Science Institute. https://ui.adsabs.harvard.edu/abs/2016jwst.rept.4622H

Horton, A., Parry, I., Bland-Hawthorn, J., et al. 2004, Proc. SPIE, 5492, Ground-Based Instrumentation for Astronomy, 1022, doi: 10.1117/12.551591

Ichikawa, T., Suzuki, R., Tokoku, C., et al. 2006, Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, 626916, doi: 10.1117/12.670078

Ives, D., Alvarez, D., Bezawada, N., George, E., & Serra, B. 2020, Proc. SPIE, 11454, X-Ray, Optical, and Infrared Detectors for Astronomy, 114541N, doi: 10.1117/12.2562408

Jackson, K., Véran, J.-P., van Kooten, M. A. M., et al. 2024, in press, Proc. SPIE, 13097, Adaptive Optics Systems

Jakobsen, P., Ferruit, P., Alves de Oliveira, C., et al. 2022, A&A, 661, A80, doi: 10.1051/0004-6361/202142663

Janesick, J., Klaasen, K., & Elliott, T. 1985, Proc. SPIE, 570, Solid State Imaging Arrays, 7, doi: 10.1117/12.950297

Kaeufl, H.-U., Ballester, P., Biereichel, P., et al. 2004, Proc. SPIE, 5492, Ground-Based Instrumentation for Astronomy, 1218, doi: 10.1117/12.551480

Kurtz, M. J., & Mink, D. J. 2000, ApJ, 533, L183, doi: 10.1086/312620

Langevin, Y., Carlier, V., Hannou, C., et al. 2022, Proc. SPIE, 12180, Optical, Infrared, and Millimeter Wave, 1218037, doi: 10.1117/12.2629310

Larkin, J., Barczys, M., Krabbe, A., et al. 2006, Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, 62691A, doi: 10.1117/12.672061

Larkin, J. E., Wright, S. A., Chisholm, E. M., et al. 2020, Proc. SPIE, 11447, Ground-Based and Airborne Instrumentation for Astronomy, 114471Y, doi: 10.1117/12.2561184

Le Graët, J., Secroun, A., Barbier, R., et al. 2022, Proc. SPIE, 12191, X-Ray, Optical, and Infrared Detectors for Astronomy, 121911M, doi: 10.1117/12.2628974

Lee, H.-I., Ramos, F., Prado, P., et al. 2022, Proc. SPIE, 12189, Software and Cyberinfrastructure for Astronomy, 121891Y, doi: 10.1117/12.2628852

Maihara, T., Iwamuro, F., Hall, D. N., et al. 1993a, Proc. SPIE, 1946, Infrared Detectors and Instrumentation, 581, doi: 10.1117/12.158709

Maihara, T., Iwamuro, F., Yamashita, T., et al. 1993b, PASP, 105, 940, doi: 10.1086/133259

McConnachie, A., Babusiaux, C., Balogh, M., et al. 2016, The Detailed Science Case for the Maunakea Spectroscopic Explorer: The Composition and Dynamics of the Faint Universe, arXiv e-prints, doi: 10.48550/arXiv.1606.00043

McGregor, P. J., Hart, J., Conroy, P. G., et al. 2003, Proc. SPIE, 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, 1581, doi: 10.1117/12.459448

McLean, I. S., Steidel, C. C., Epps, H. W., et al. 2012, Proc. SPIE, 8446, Ground-Based and Airborne Instrumentation for Astronomy, 84460J, doi: 10.1117/12.924794

Meinel, I. A. B. 1950, ApJ, 111, 555, doi: 10.1086/145296

Moreels, G., Clairemidi, J., Faivre, M., et al. 2008, ExA, 22, 87, doi: 10.1007/s10686-008-9089-6

Oliva, E., & Origlia, L. 1992, A&A, 254, 466. https://ui.adsabs.harvard.edu/abs/1992A&A...254..466O

Plazas, A. A., Shapiro, C., Smith, R., Rhodes, J., & Huff, E. 2017, JInst, 12, C04009, doi: 10.1088/1748-0221/12/04/C04009

Puech, M., Rodrigues, M., Yang, Y., et al. 2014, Proc. SPIE, 9147, Ground-Based and Airborne Instrumentation for Astronomy, 91476L, doi: 10.1117/12.2055137

Ramsay, S. K., Mountain, C. M., & Geballe, T. R. 1992, MNRAS, 259, 751, doi: 10.1093/mnras/259.4.751

Rauscher, B. J., Lindler, D. J., Mott, D. B., et al. 2011, PASP, 123, 953, doi: 10.1086/661663

Rauscher, B. J., Fox, O., Ferruit, P., et al. 2007, PASP, 119, 768, doi: 10.1086/520887

Regan, M. W., & Bergeron, L. E. 2020, JATIS, 6, 1, doi: 10.1117/1.JATIS.6.1.016001

Richardson, E. H. 1968, JRASC, 62, 313. https://ui.adsabs.harvard.edu/abs/1968JRASC..62..313R

Rieke, G. H. 2007, Annual Review of Astronomy and Astrophysics, 45, 77, doi: 10.1146/annurev.astro.44.051905.092436

Robberto, M. 2009, Derivation of Correct Noise Equation for General MULTIACCUM Readout, Tech. Rep. JWST-STScI-001853, SM-12, Space Telescope Science Institute. https://ui.adsabs.harvard.edu/abs/2009jwst.rept.1853R

Rodrigues, M., Flores, H., Puech, M., Yang, Y., & Royer, F. 2010, Proc. SPIE, 7735, Ground-Based and Airborne Instrumentation for Astronomy, 77356L, doi: 10.1117/12.856956

Rodrigues, M., Cirasuolo, M., Hammer, F., et al. 2012, Proc. SPIE, 8450, Modern Technologies in Space- and Ground-based Telescopes and Instrumentation, 84503H, doi: 10.1117/12.924878

Rousselot, P., Lidman, C., Cuby, J. G., Moreels, G., & Monnet, G. 2000, A&A, 354, 1134. https://ui.adsabs.harvard.edu/abs/2000A&A...354.1134R

Scarfe, C. D., Batten, A. H., & Fletcher, J. M. 1990, Publications of the Dominion Astrophysical Observatory Victoria, 18, 21. https://ui.adsabs.harvard.edu/abs/1990PDAO...18...21S

Sharp, R., & Parkinson, H. 2010, MNRAS, 408, 2495, doi: 10.1111/j.1365-2966.2010.17298.x

Sharples, R., Bender, R., Agudo Berbel, A., et al. 2013, The Messenger, 151, 21. https://ui.adsabs.harvard.edu/abs/2013Msngr.151...21S

Sivanandam, S., Chapman, S., Simard, L., et al. 2018, Proc. SPIE, 10702, Ground-Based and Airborne Instrumentation for Astronomy, 107021J, doi: 10.1117/12.2313924

Smith, R. M., & Hale, D. 2012, Proc. SPIE, 8453, High Energy, Optical, and Infrared Detectors for Astronomy, 84530Y, doi: 10.1117/12.927148

Smith, R. M., Zavodny, M., Rahmer, G., & Bonati, M. 2008, Proc. SPIE, 7021, High Energy, Optical, and Infrared Detectors for Astronomy, 70210J, doi: 10.1117/12.789372

Soto, K. T., Lilly, S. J., Bacon, R., Richard, J., & Conseil, S. 2016, MNRAS, 458, 3210, doi: 10.1093/mnras/stw474

Tam, S., & Hu, C. 1984, IEEE Transactions on Electron Devices, 31, 1264, doi: 10.1109/T-ED.1984.21698

Taylor, W., Cirasuolo, M., Afonso, J., et al. 2018, Proc. SPIE, 10702, Ground-Based and Airborne Instrumentation for Astronomy, 107021G, doi: 10.1117/12.2313403

Teeple, D., Riopel, M., Baril, M., et al. 2006, Proc. SPIE, 6274, Advanced Software and Control for Astronomy, 62740I, doi: 10.1117/12.670355

Thatte, N., Tecza, M., Clarke, F., et al. 2010, Proc. SPIE, 7735, Ground-Based and Airborne Instrumentation for Astronomy, 77352I, doi: 10.1117/12.857445

Voss, R. F., & Clarke, J. 1976, PhRvL, 36, 42, doi: 10.1103/PhysRevLett.36.42

Wild, V., & Hewett, P. C. 2005, MNRAS, 358, 1083, doi: 10.1111/j.1365-2966.2005.08844.x

Wildi, F., Blind, N., Reshetov, V., et al. 2017, Proc. SPIE, 10400, Techniques and Instrumentation for Detection of Exoplanets, 1040018, doi: 10.1117/12.2275660

Young, P. J., McGregor, P., van Harmelen, J., & Neichel, B. 2012, Proc. SPIE, 8451, Software and Cyberinfrastructure for Astronomy, 845124, doi: 10.1117/12.925273