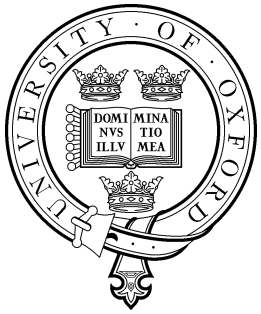
The WEAVE fibre positioner: calibration, commissioning, and first-light



Sarah Hughes
St. Cross College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2023

This thesis is dedicated to
Andrew McWilliam
for the endless support and guidance you have given me,
I couldn't have done this without you.

Acknowledgements

It's hard to believe that it's been four years since I began my DPhil in 2019. I have cherished my time in Oxford, and although the journey hasn't always been easy, I will look back on it with fond memories.

I am deeply grateful to my supervisors, Gavin Dalton and Ian Lewis, and to David Terrett for welcoming me to the WEAVE-Oxford group and for being incredible mentors throughout my DPhil. Their guidance, stimulating research discussions, and kindness have been invaluable as I progress in my career. I would also like to thank Ellen, Shoko, and Scott for their advice and experiences, which have been immensely helpful to me.

I owe a special debt of gratitude to my uncle Andrew McWilliam for his constant support and inspiration in my journey to become an astronomer.

I am also thankful to my family and friends for their unwavering support. Particularly my father, for reminding me that McDonald's is always hiring. As well as my colleague Jakob, for saying that McDonald's also has standards. To my fellow comrades Peter, Amy, Joe, Jakob, Tobias, Gustavo, Annabella, and Susanna, thank you for making my time in Oxford so enjoyable, despite the ups and downs we've had to endure. I am also grateful to Lauren and Ed for their advice as I progressed in my DPhil, drawing from their own experiences.

Special thanks go to Yorkshire Tea, Custard Creams, and random YouTube videos for getting me through this.

To Mum, Dad, and Siân, I am truly grateful for your unwavering encouragement, belief in my abilities, and continuous support throughout my DPhil. I hope you know how much you all mean to me. Lastly, Tom, you've been the best partner I could ever ask for, and I am incredibly lucky to have you by my side. I will always remember Oxford as our first home together.

Statement of Originality

I carried out the work in this thesis at the Department of Astrophysics, University of Oxford between October 2019 and April 2023 under the supervision of Prof. Gavin Dalton. My research was funded by STFC and St. Cross College, Oxford. I declare that no part of this thesis has been submitted in support of another degree, diploma or other qualification at the University of Oxford or any other university. Except where otherwise stated, the work in this thesis is mine.

Sarah Hughes

I suppose I'll have to add the force of gravity to my list of enemies.

$\sim$ *Lemony Snicket* $\sim$

Abstract

This thesis details the final assembly and calibration procedures for the WEAVE fibre positioner, from October 2019 to first-light in September 2022. The positioner was successfully mounted to the William Herschel Telescope (WHT) in May 2022. I present the first-light results of WEAVE for each of its three observing modes and discuss whether the positioner has met its design expectations.

Chapter one describes the evolution of astronomical spectroscopy and the scientific motivations behind the construction of WEAVE. This is followed by the instrument design specification and sub-systems required to meet them. Chapter two will focus on the positioner in detail and include the evolution of our quality control methods.

An account of the calibration procedures, in both the lab environment and on the elevation testing rig, is given in chapter three. The test rig revealed flexure behaviour in the positioner's structure, which was modified to correct for this. The calibration work was repeated with improved techniques, to predict changes in the positioner's metrology with elevation. This chapter then presents the results of the fibre movement testing, including the current placement timings

To maximise the scientific output from WEAVE, I applied a series of adjustments to a subset of target fields, identifying methods that will improve the survey output using the fibre assignment program, Configure, and highlight potential selection effects.

Chapter five presents the first-light results and the steps taken to achieve this. It required extensive measurements of the optical distortion map across the FoV, which were carefully mapped using astrometry of the night-sky across our focal plane. I then compare these results with the observational requirements from WEAVE's conception, to define its overall success. Chapter 6 provides an overview of the work completed in this thesis and describes potential upgrades to the instrument in the future.

Contents

Glossary **1**

1 A brief history of Astronomical Spectroscopy **4**
1.1 The basic principles 5
1.2 Fibre Spectroscopy 9
1.2.1 Integral Field Units 11
1.3 Automated Positioning Methods 13
1.4 The dawn of WEAVE 17
1.4.1 Motivation and Science objectives 18
1.5 Systems overview 21
1.5.1 The prime focus corrector 23
1.5.2 The rotator 25
1.5.3 The positioner 26
1.5.4 The optical fibres 26
1.5.5 The spectrograph 28
1.6 This Research 30

2 The WEAVE fibre positioner **33**
2.1 Structural overview 33
2.2 Cartesian robots 36
2.2.1 The Focal Plane Imager 38
2.2.2 Positioning requirements 41
2.3 The retractors 44
2.3.1 Assembly 44
2.3.2 Quality control 47
2.4 Slitlets 50
2.4.1 Assembly 50
2.4.2 Mapping 52

3 Calibration and Metrology **54**
3.1 Zenith calibrations . . . 55
3.1.1 Tumbler metrology . . . 55
3.1.2 Field plate metrology . . . 58
3.1.3 Retractor tier heights . . . 61
3.1.4 Fibre imaging and initial placement testing . . . 65
3.2 Calibration at elevation . . . 70
3.2.1 Flexure . . . 71
3.2.2 Robot calibration . . . 76
3.2.3 Movement modelling . . . 80
3.3 Fibre placement offsets . . . 89

4 Field configuration optimisation **99**
4.1 Fibre Allocation . . . 100
4.1.1 Configure . . . 102
4.1.2 Completeness . . . 105
4.1.3 Assignment cases . . . 109
4.1.4 Selection effects . . . 121
4.1.5 Overlapping fields . . . 123
4.2 Conclusion . . . 127

5 Commissioning and First-light **129**
5.1 Astrometry . . . 129
5.1.1 Image quality . . . 130
5.1.2 Telescope guiding and tracking . . . 132
5.1.3 Distortion maps . . . 135
5.1.4 Astrometric measurements . . . 137
5.2 Fibre mapping . . . 144
5.2.1 Slitlet mapping . . . 144
5.2.2 IFU mapping . . . 145
5.2.3 MOS observations . . . 151
5.3 First-light data . . . 154
5.4 Specification comparison . . . 158

6 Conclusion **162**
6.1 Further work . . . 167
6.2 The future of WEAVE . . . 168

A The WEAVE science surveys in detail **172**
A.1 Galactic Archaeology and Evolution 172
A.1.1 Stellar, Circumstellar and Interstellar physics 174
A.1.2 White Dwarf Stars 175
A.1.3 Apertif 176
A.1.4 Galaxy clusters 177
A.1.5 STePS 178
A.1.6 LOFAR 179
A.1.7 Quasi-Stellar Objects 180

B Flexure defocus model **181**

Bibliography **184**

Glossary

2dF Two-degree field of view multi-object spectrograph for the AAT.

AAT Anglo-Australian Telescope.

ADC Atmospheric Dispersion corrector.

AGN Active Galactic Nuclei.

Autofib-2 Automatic fibre positioner for the WHT.

CAD Computer Aided Design.

CASU Cambridge Astronomical Survey Unit.

CPS Core Processing System.

DIMM Differential Image Motion Monitor.

DTI Dial Test Indicator.

ESA European Space Agency.

ETSRC European Telescope Strategy Review Committee.

FCA Field Configuration algorithm.

FEA Finite Element Analysis.

FoV Field of View.

FPI Focal Plane Imager.

FRD Focal Ratio Degradation.

FTS Focus Translation System.

FWHM Full Width Half-Maximum.

GA Galactic Archaeology.

GHRIL Ground-based High Resolution Imaging Laboratory enclosure.

GTC Gran Telescopio Canarias.

HETDEX Hobby-Eberly Telescope Dark Energy Experiment.

HR High-Resolution.

IFU Integral Field Unit.

ING Issac Newton Group.

ISM Interstellar Medium.

LDSS Low Dispersion Survey Spectrograph.

LIFU Large Integral Field Unit.

LOFAR Low Frequency Array radio telescope.

LR Low-Resolution.

MaNGA Mapping Nearby Galaxies at Apache Point Observatory.

MC Monte Carlo.

mIFU mini Integral Field Unit.

ML Machine Learning.

MOS Multi-Object Spectroscopy.

PA Position Angle.

PFC Prime Focus Corrector.

PMAS Potsdam Multi-Aperture Spectrophotometer.

POS The WEAVE fibre positioner system.

PSF Point Spread Function.

QSO Quasi-stellar object.

RAS Royal Astronomical Society.

RMS Root Mean Squared.

RSD Retractor Support Disc.

RSS Residual Sum of Squares.

RV Radial velocity.

SA Simulated Annealing.

SAMI Sydney-AAO Multi-Object Integral-field Spectrograph.

SCIP Stellar, Circumstellar and Interstellar Physics survey.

SMR Spherically Mounted Retroreflector.

SNR Signal-to-Noise Ratio.

SOFA Standards of Fundamental Astronomy.

SPA Science Processing Archive.

SPIE International Society for Optics and Photonics.

STePS Stellar Populations at Intermediate Redshifts Survey.

TSP Travelling Salesman Problem.

VPH Volume Phase Holographic.

WD White Dwarf.

WEAVE WHT Enhanced Area Velocity Explorer.

WHT William Herschel Telescope.

WL WEAVE-LOFAR.

WRS WEAVE Rotator System.

ZD Zenith Distance.

Chapter 1

A brief history of Astronomical Spectroscopy

> This would have been a lot easier if we did this a billion years ago, everything would be much closer together.
>
> G. Dalton

Modern instrumentation has broadened the range and depth that astronomers can observe, allowing the structure of objects such as galaxies, clusters, and black holes to be investigated in greater detail. Spectroscopic instruments play a significant role in these observations. By studying the intricate details of light from astronomical objects, astronomers can obtain a wealth of information about their environment, formation, and evolution; gradually piecing together data to improve our understanding of the universe.

The William Herschel Telescope (WHT) Extended Area Velocity Explorer (WEAVE) is the new multi-object spectrograph installed on the 4.2 m WHT, with a multiplex of 960 individual optical fibres that are deployed over its 2° Field of View (FoV). The terminology used to describe WEAVE will be defined in this section and a detailed overview of the instrument is provided in chapter two.

This chapter will provide an overview of astronomical spectroscopy and its methodology, alongside the scientific motivation for WEAVE and the history of its conception.

1.1 The basic principles

Spectroscopy is the process of dispersing the light emitted by an object to produce a spectrum, which is a plot of the light's intensity as a function of wavelength.

Historically, spectra are produced via slit spectroscopy (Hearnshaw, 2009). Incoming light from an object is passed through a slit and often, but not always (see the Faint Object Spectrograph on the WHT; (Wynne, 1982)), a collimator to ensure the rays of light are parallel. The light is then passed through a dispersive element, typically a prism or a grating, before being focused onto a detector via a camera to produce a spectrum. In its simplest form, this requires only two lenses and either a grating or a glass prism, shown in 1.1. The long slit aperture produces a 1D spectrum, with a spectral resolution determined by the convolution of the slit width and the intrinsic dispersion of the prism or grating. The light is isolated along the direction of the slit and requires precise target acquisition to position the object appropriately.

By analysing the spectrum emitted by an object, astronomers can deduce many of its characteristics. For example, the chemical composition can be inferred using the unique spectral signature that each atomic element possesses. These properties are calculated using features known as spectral lines, which are classed as either emission or absorption lines.

Absorption lines occur when a travelling photon is absorbed by an electron orbiting in an atomic shell, causing it to gain energy and move into an excited state. Only photons with specific energies will be absorbed, as there are fixed quantum states an electron may occupy and still remain bound to the nucleus. This energy corresponds to the transition energy between states. In astronomy, absorption lines occur in when light passes through cold, dense regions of gas, such as interstellar clouds. Emission lines can be thought of as the inverse of this process. An atom in an excited state will spontaneously release a photon as the electron drops from its excited state to the ground state. They typically occur in the Interstellar Medium (ISM). An example of

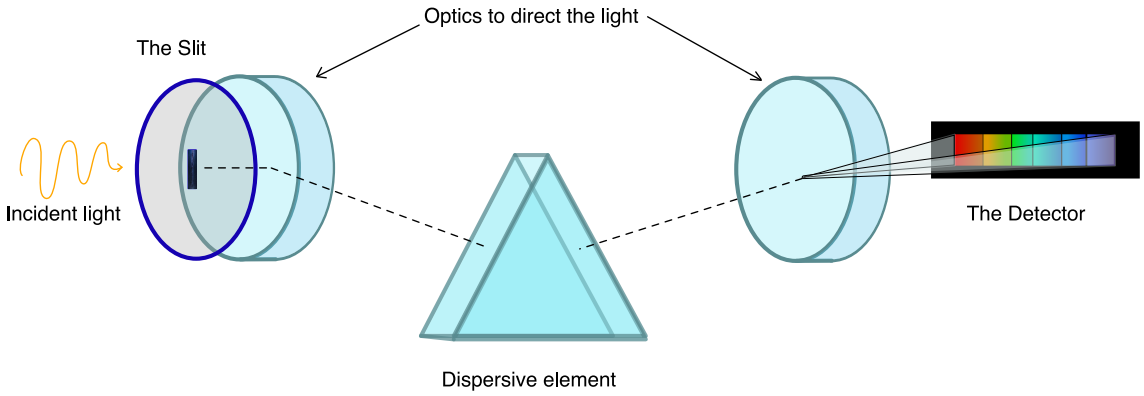


Figure 1.1: A diagram demonstrating the fundamental components of a slit spectrograph. It uses lenses to direct light passing through a slit to a dispersive element in order to produce a spectrum, which is then directed onto a detector.

a stellar spectrum, containing both types of spectral lines, can be seen in Figure 1.2. The lines are produced in the stars outer atmosphere.

The change in energy ΔE of the state of an atom is related to the wavelength of light λ emitted, where

$$\Delta E = \frac{hc}{\lambda}, \tag{1.1}$$

h is Planck's constant, c is the speed of light, and λ is the wavelength of the light. This means that the wavelength of an absorption or emission line along a spectrum can be used to identify which element it originated from (Appenzeller, 2013).

On cosmological scales, common lines in a spectrum are shifted by a similar amount in wavelength. This is called the Doppler effect and is caused by the relative movement between the observer and the object as the light travels towards us (Andrade, 1959). To calculate this shift, astronomers use well-known line structures, such as hydrogen, and compare their location to where they occur in the laboratory rest frame. Laboratory spectra are well-classified for a large catalogue of elements.

The Doppler effect is defined as

$$f = \left(1 + \frac{\Delta v}{c}\right) f_0, \tag{1.2}$$

where f is the observed frequency, f_0 is the emitted frequency, Δv is the relative velocity of the receiver to the source, and c is the speed of light. The radial velocity v_r is the velocity along the line of sight from the observer, where

$$\frac{v_r}{c} = \frac{\Delta \lambda}{\lambda_0} \tag{1.3}$$

and $\Delta\lambda$ is the difference between the observed wavelength and the rest-frame wavelength λ. This effect is used to measure the distance of astronomical objects, known as the redshift z, which is defined as

$$z = \frac{\Delta \lambda}{\lambda_0}. \tag{1.4}$$

The lines present within a spectrum are dependent on the wavelength range being observed. Therefore, the structures astronomers use to measure z depend on the instrumentation and the cosmological objects themselves.

The spectral resolution R of an instrument defines its ability to distinguish spectral lines that are spatially close, where

$$R = \frac{\lambda}{\Delta \lambda} \tag{1.5}$$

and $\Delta\lambda$ is the smallest difference that can be distinguished at wavelength λ. In practice, this is normally measured using the full-width at half maximum (FWHM) of the spectral line itself (Appenzeller, 2013).

In astronomy, an important quantity measured using a spectrum is the metallicity Z, which provides a direct link to the object's age. It is frequently used across a broad

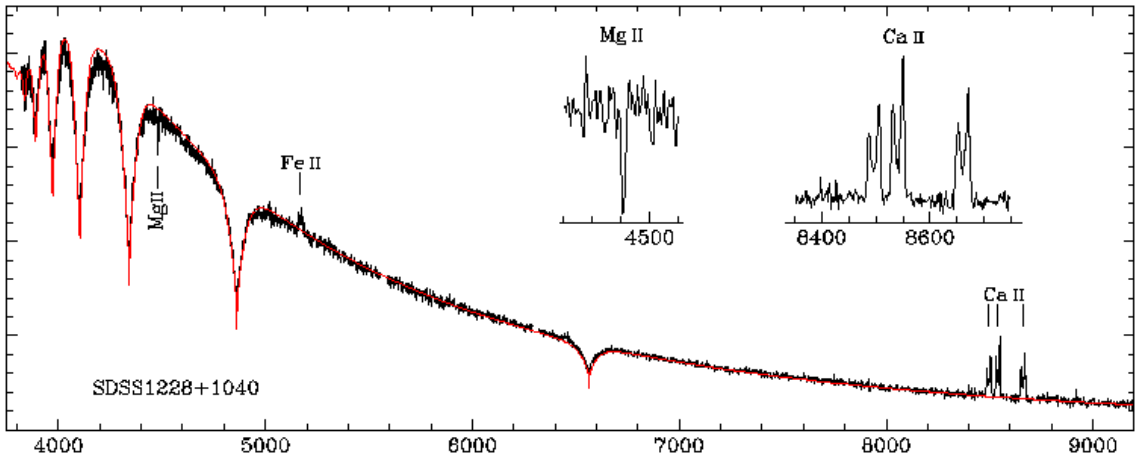


Figure 1.2: An example of a spectrum observed by SDSS, which was produced by a white dwarf star accreting planetary debris (Gansicke et al., 2007).

range of research areas for example, in White Dwarf (WD) stars, additional metals in their spectrum indicate accreting planetary debris. Astronomer's class any element heavier than Hydrogen (H) or Helium (He) as a metal. The metal content of a star is defined as the sum of the mass ratios for all elements in the star other than Hydrogen and Helium m_i,

$$Z = \sum_{i>\mathrm{He}} \frac{m_i}{M}, \tag{1.6}$$

where M is the total mass of the star. It is useful to consider the mass fraction of a star for each component:

$$X + Y + Z = 1, \tag{1.7}$$

where X is the mass fraction of hydrogen and Y is the mass fraction of Helium.

Spectroscopy is an indispensable tool that can provide astronomers with a wealth of information that is essential to expanding our knowledge of the universe. From the formation of stars and galaxies to their behaviour and evolution, it is a diverse area of research that extends far beyond astronomy and will continue to be relevant for many years to come. Further details on the spectral analysis of astronomical objects and the history of its development can be found in (Massey and Hanson, 2013) and (Hearnshaw, 1990).

1.2 Fibre Spectroscopy

As technology has rapidly advanced in recent decades, so have the instruments astronomers use to observe. Improving our understanding of the universe requires extensive observations of the night sky, and the focus has continually shifted towards observing a greater number of astronomical objects.

Increasing the sample size of surveys has a number of statistical benefits, such as minimising observational biases, reducing the errors associated with a particular measurement, and studying a range of parameter distributions. Multi-Object Spectroscopy (MOS) is the process of observing the spectra of more than one object simultaneously, typically using optical fibres or multiple slits. The components of a simple multi-object spectrograph can be seen in figure 1.3.

Optical fibres use total internal reflection to transmit light along its length. This occurs when a wavefront is completely reflected at the boundary to a second medium. Figure 1.4 demonstrates the typical structure of an optical fibre. Fibres are capable of transmitting large quantities of information across great distances with very little information loss, which is why they are used in modern communication systems, and are beneficial to astronomers. Further details on the physics of their light transmission can be found in (Hecht, 1998).

Optical fibres can be used to isolate the incoming light from objects, and by deploying fibres across the field of view, spectra can be observed for multiple objects in the same exposure. The first use of optical fibres in astronomy was the MEDUSA instrument in 1979 for the Steward Observatory, where the fibres were inserted into the focal surface (Hill, 1988). This method is referred to as the "plug-plate method" and is still used for instruments such as SDSS.

For SDSS, optical fibres are manually plugged into holes drilled into the surface of the focal plane to align with the targets in the night sky. This method requires on-site facilities to prepare the plates each night however, they are unusable for subsequent

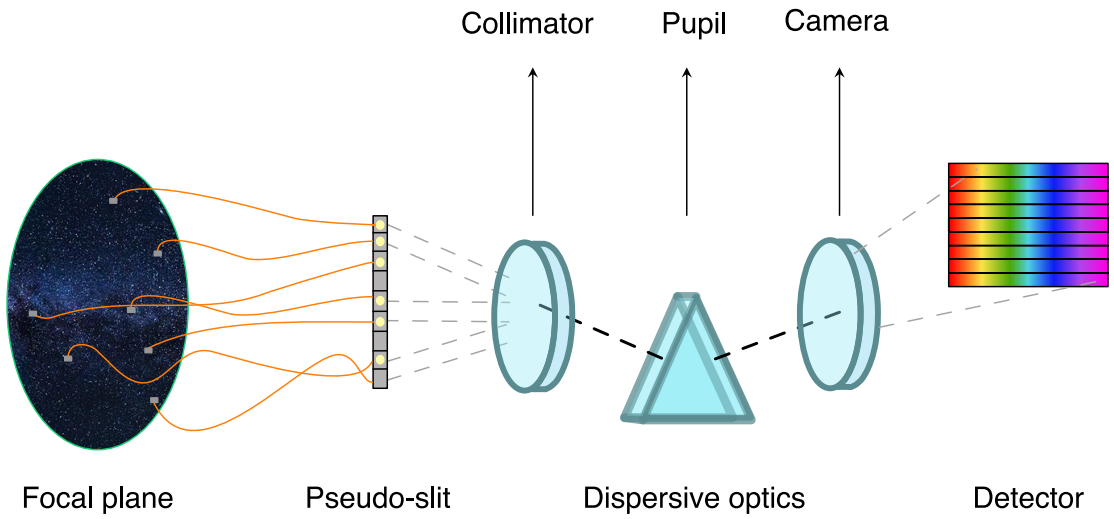


Figure 1.3: A diagram outlining the components of a multi-object spectrograph. Individual optical fibres direct light from the focal plane and their ends are arranged along a slit. Each fibre produces a spectrum on the detector after its light is passed through a dispersive element, typically a grating.

observations as the plate was drilled according to the specific temperature and seeing conditions.

The plug-plate method was developed in-parallel with the Low Dispersion Survey Spectrograph (LDSS) for the WHT; a spectrograph which used custom masks with multiple apertures to observe up to 100 objects simultaneously (Allington-Smith et al., 1990). Unfortunately, it suffers from many of the same disadvantages as the plug-plate method in terms of material waste and limited efficiency.

The use of optical fibres in instrumentation has many advantages in astronomy. Notably, the fibres enable the detectors to be positioned off of the telescope. This has major implications in terms of stability when observing, as it impacts the telescope mass balance and reduces mechanical flexing of the top-end structure.The fibres transmission efficiency is quite high and theoretically could transmit up to 90% of the incoming light (Parry, 2006) however, there are some drawbacks with the use of optical fibres. An important consideration is their length, as increasing the length of the fibres will decrease the throughput in the UV wavelength range (Lewis et al.,

2002). Another factor to consider is the focal ratio degradation (FRD). As an optical fibre is subject to exterior effects such as bending and compressing, the transmitted cone of light from the end of the optical fibre becomes larger than the input angle of the incoming light (Parry, 2006). This leads to information loss which will reduce the final throughput and the quality of data received from observations. It is essential to minimise the FRD as much as possible.

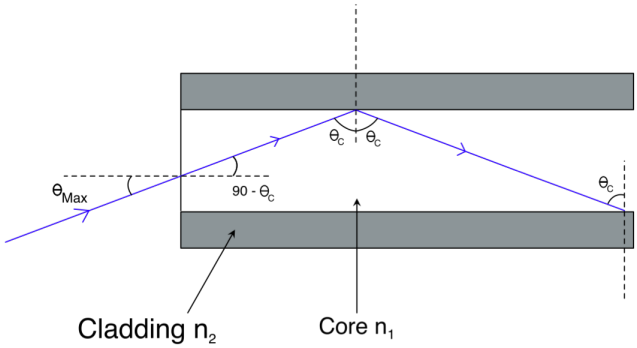


Figure 1.4: A diagram showing the propagation of light down an optical fibre. The light undergoes total internal reflection by a critical angle θ_C at the boundary of the fibre. The core of the fibre has a refractive index of n_1 and the protective cladding on the outside has a refractive index of n_2.

1.2.1 Integral Field Units

Integral Field units (IFUs) are used to produce spatially resolved spectra over a 2D field of view, with data taken over a single pointing simultaneously. This is used to produce a 3D datacube of a scalar quantity proportional to the flux density, as a function of the spatial coordinates and wavelength (Allington-Smith, 2006). To put it simply, spectral data is taken across an object's spatial dimension. The xy plane preserves the structure of the object and the z axis is the spectrum at each position, known as a Spaxel. It builds upon the detail of spectroscopy to conduct regional studies of objects such as galaxies and proto-stellar discs.

IFUs use two main structures, either lenslets or image slicers. A diagram of

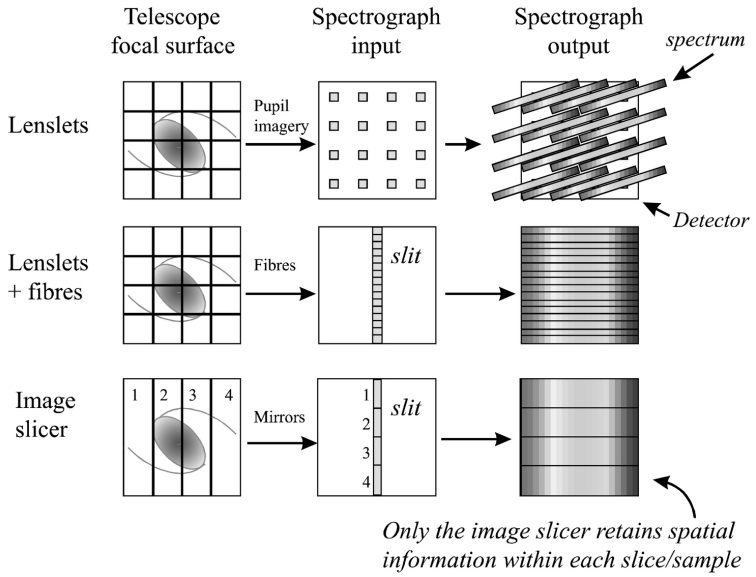


Figure 1.5: A summary of the main structures used in Integral field units from Allington-Smith and Content (1998).

these methods can be seen in figure 1.5. Lenslets use multiple apertures to isolate the incoming light. This arrangement produces a spectrum for each Spaxel over the detector, which will overlap without adequate spacing. To reduce the amount of empty space across the field of view, the dispersion angle can be adjusted however, this results in a longer spectrum which may also lead to overlapping. Therefore, lenslets are suited to situations where a long spectrum is not necessary e.g. high-resolution spectroscopy over a short wavelength range or a low resolution across a longer wavelength range.

The IFU observing modes of WEAVE provides an example of combining optical fibres with lenslet arrays to produce fibre reformatters. They are a bundle of closely packed optical fibres, typically in a hexagonal arrangement. A diagram of WEAVE's basic IFU structure is shown in Figure 1.6. In highly clustered regions, IFUs can be used as an alternative when the placement of optical fibres is not compact enough and cannot cover the target density. The fibres are reformatted into a pseudoslit to produce a long-slit spectrum that fits within the detector dimensions.

Image slicers use reflective optics to divide the image incident on the telescope

into several narrow slices. The slices are set to varying angles so that the diverging beams will exit along a unique path. The beams are then reflected by an additional offset mirror to produce a field with all the slices arranged from end-to-end at the entrance to the spectrograph. This design maximises the number of spectral and spatial resolution elements without increasing the size of the spectrograph, and can also be used in cryogenic systems. This technique was first implemented at MPIE as a 3D spectrograph (Krabbe et al., 1997).

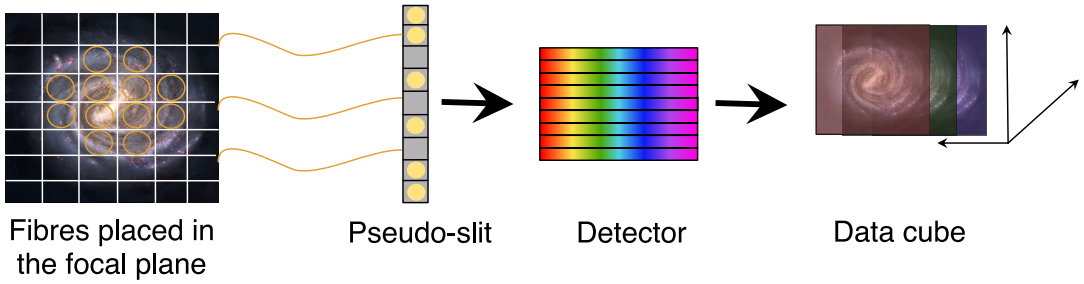


Figure 1.6: A diagram of the basic IFU components. The fibres are placed in a bundle over the target, which is then arranged in a slitlet and passed through an image slicer to produce a 3D data cube. The xy plane retains the spatial information, while the z axis changes as a function of wavelength.

Examples of recent IFU surveys include the Mapping Nearby Galaxies at Apache Point Observatory (MaNGA) survey, and the Sydney-AAO Multi-object Integral-field Spectrograph (SAMI) survey, which have been highly successful. Further details can be found in Bundy et al. (2015) and Bryant et al. (2015), and several upcoming instruments are taking advantage of this observing mode, such as Hector (Bryant et al., 2016), and VIRUS2 (Hill et al., 2020).

1.3 Automated Positioning Methods

With the continual push to observe an increasing number of objects per observation comes the need for automation. Whilst the plug-plate method has worked well for

many years, it is not a sustainable practice; a team of dedicated staff can only assemble so many observations manually, which places a hard limit on the instrument itself. The first automated MOS instrument was the MX Spectrograph (Hill and Lesser, 1986), which used 32 probes driven by stepper motors to place its 64 optical fibres. Automated fibre positioners have advanced significantly in recent years and a description of the techniques currently used in MOS instruments is summarised below. Further details on each of these methods are presented in (Gilbert, 2016).

- **Pick-and-Place positioning:** One or more robots pick up individual fibres using a gripper system and move along gantries to place them in their desired position in the focal plane. The movement of the fibres is dependent on their arrangement and choice of robot axes (Parry and Sharples, 1988).

- **Tilting Spine positioning:** Uses Echidna technology to place multiple fibres simultaneously (Gillingham et al., 2000) (Moore et al., 2003). Each fibre is held in a rigid tube called a 'spine', which can freely tilt around a ball joint at the end. An actuator that can move the spine in discrete steps is attached to the ball.

- **Radial arm positioning:** Known as *r-theta* positioners, they use multiple rotary motor arms arranged in an eccentric axis. Over a circular patrol area, every point can be accessed by at least one optical fibre. This method also allows for simultaneous fibre placement (Hu et al., 2003).

- **Roaming positioners:** Uses miniature self-mobile piezoelectric devices called 'Starbugs' to attach the fibres in the field of view (Gilbert et al., 2012). Suitable for extremely large focal planes and allows for large freedom of movement in fibre positions.

Automating the placement of optical fibres has significant advantages over manual methods. In addition to improving the efficiency of the instrument, automation

allows fields to be configured for prevailing conditions and improves the flexibility of observations by configuring fields shortly before they take place. The exact timing is dependent on the structure of the instruments focal plane. The choice of positioning technology used is determined by the specific requirements of the focal plane assembly, and each mechanism has its own complications and drawbacks. Key factors when making this decision include the target density and placement buffer around the fibres, configuration timing restrictions, and the focal surface structure.

It is important to note that pick-and-place positioning is a sequential process, whereas the alternatives are intrinsically parallel. This impacts the rate at which a field may be configured. In principle, parallel processes also have the functionality to update their positions within the focal plane to compensate for effects such as differential image motion.

There are two instruments which are highly relevant to this research; the previous WHT instrument Autofib-2 (Parry and Sharples, 1988), and 2dF, which was mounted to the Anglo-Australian Telescope (AAT) and began its observations in 1997 (Lewis et al., 2002). These instruments paved the way for WEAVE's development, and WEAVE itself incorporates many of the same features.

AutoFib-2 was installed at the top end of the WHT in 1995 and follows on from AutoFib which was installed at the AAT in 1990 (Gray and Taylor, 1990). AutoFib-2 has a 1° field of view, a multiplex of 126 optical fibres, and a resolution of 8000. Each fibre was positioned using a single pick-and-place robotic arm, which had a pneumatic gripper and a robot-vision system to monitor the fibre during its placement. The optical fibres were arranged around the edge of the field plate in bundles of 10 fibres which had to be positioned to ± 20 μm of its target in the focal plane.

The AAT is a 3.9-meter equatorial telescope, where the instrument 2dF was mounted. 2dF had a 2° field of view due to its prime focus corrector and a resolution between 500 and 2000 over a wavelength range of 400 and 1100 nm. It has a

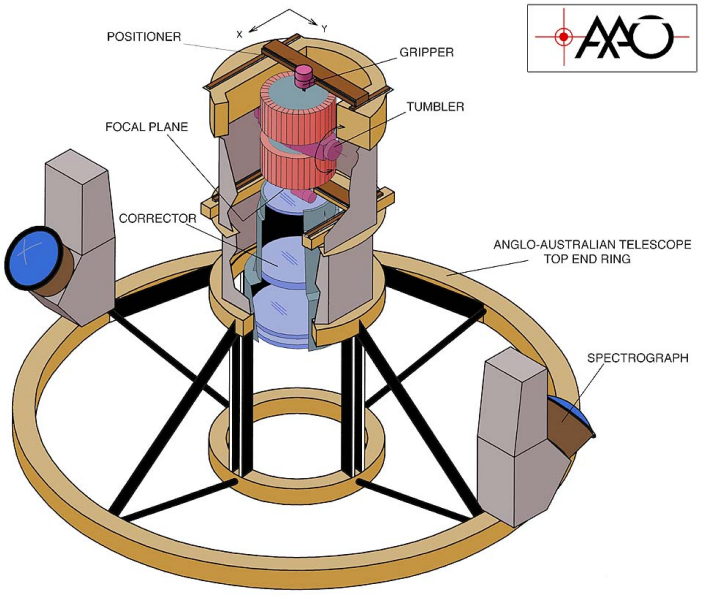


Figure 1.7: A schematic diagram of the original layout of 2dF on the AAT (Lewis et al., 2002).

multiplex of 400 optical fibres that are stored in retractors around the edge of the field plate. 2dF operates on a tumbler system with two field plates, one on each side. While one plate is being used for observations, the other side is being configured by a single robot that is positioning each group of optical fibres. It has a positioning tolerance of ± 20 μm and the robotic arm could place a fibre to an accuracy of 11 μm. 2dF is a prime-focus instrument, and originally had two spectrographs located on platforms on either side. These spectrographs were eventually replaced by AAOmega, located in further part of the dome using a long fibre run (Smith et al., 2004). With the original spectrographs on nearby platforms, this layout reduced the length of the fibres to maximise the UV throughput. A schematic diagram of 2dF is shown in figure 1.7. It was designed to conduct the largest galaxy redshift survey at the time, with ~ 2500 galaxies being observed per night and over 225,000 objects observed between the start of its surveys in 1997 and the final data release in 2003.

1.4 The dawn of WEAVE

In 2009, the Isaac Newton Group (ING) began planning their future development strategy, based on the recommendations highlighted in the road mapping report for European astronomy (Lagrange and Andersen, 2006). Another report was then released by European Telescope Strategy Review Committee (ETSRC) (Drew et al., 2009), conveying the community benefits of a high-multiplex, wide-field spectrograph in the northern hemisphere, with the 4.2 m William Herschel Telescope (WHT) at Observatorio del Roque de los Muchachos, La Palma being identified as one of the preferred sites

At a two-day meeting of the Royal Astronomical Society (RAS) in March 2010, astronomers from the ING community came together to discuss the future development of the WHT for the next decade, and the scientific opportunities available. The first day was dedicated to new science topics whilst the second day focused on instrumentation. This led to the realisation that one could build a relatively low-cost instrument with commercially available equipment that covered the majority of the desired science.

In 2013, at a meeting of the International Society for Optics and Photonics (SPIE), a new instrument for the WHT was proposed by Professor Gavin Dalton. It was decided that this instrument should be almost entirely dedicated to follow-up surveys to support the needs of recent missions such as GAIA (Gaia Collaboration et al., 2016) and the Low Frequency Array (LOFAR) (van Haarlem, M. P. et al., 2013), with a wavelength range and resolution defined to maximise the science output when combining the survey data.

This concept evolved into WEAVE; a high-multiplex, fibre-fed instrument, covering a 366-959 nm wavelength range with a 2° field of view (FoV) and two resolution modes ($R \sim 20,000$ over several windows or $R \sim 5,000$ with continuous coverage). Over its 5-year mission, WEAVE will meet its observational goals by obtaining new

spectroscopic data for a variety of astronomical objects such as galaxy clusters, the interstellar medium, white dwarf stars, and stars within the Milky-Way, providing ~ 12 million spectra over its lifetime. Further details on these surveys can be found in appendix A although they will be summarised in section 1.4.1 below.

1.4.1 Motivation and Science objectives

The *Gaia* mission launched in December 2013, with a five-year survey to produce astrometric measurements of 1 billion stars in our galaxy (Gaia Collaboration et al., 2016). *Gaia* has been a huge success and is the first instrument to couple photometric data with six-dimensional phase space information to understand the kinematic and spatial distribution of stars within the Milky Way (Lindegren et al., 2007). Its operation has recently been extended until December 2025 due to its good condition and availability of fuel.

Alongside astrometric observations, *Gaia* is capable of observing spectroscopic data at a resolution of $R \sim 11,500$ over a wavelength range of 845 to 872 nm, which includes the strong calcium triplet at $\sim$ 860 nm. This means that *Gaia* can only obtain radial velocity (RV) measurements for the brightest stars in its survey, with a magnitude $G \leq 16.5$[1] (Lindegren et al., 2007). This limited wavelength range means that abundance measurements for elements in astrophysical processes, such as slow and rapid neutron capture (s- and r-process), are missed and *Gaia* must be supported by ground-based spectroscopic surveys to achieve important scientific objectives. Most notably, studying the assembly and evolution of our galaxy by tracing the stellar metallicity content.

The WEAVE Galactic Archaeology (GA) survey will measure approximately 2 million stars in both the low and high-resolution modes with magnitude ranges of $16 < G < 20.7$ and $12 < G < 16$ respectively. These observations will cover almost

[1]At the time of writing this, it is magnitude $G \leq 14.5$ but will extend to ~ 16 by the end of *Gaia*'s mission.

all of the stellar populations within our galaxy that are accessible from the Northern hemisphere to conduct chemical-dynamic labelling of Milky-Way structures with *Gaia* (Jin et al., 2022). The GA survey is split into four sub-sections based on the resolution mode and structure observed; the galactic halo, disk dynamics, chemical abundances of stellar populations, and open clusters.

LOFAR is the new radio-interferometer for the northern hemisphere, with a baseline of 100 km in the Netherlands and 1500 km across Europe. It has a frequency range of 10-240 MHz. LOFAR was designed to conduct low-frequency surveys with a tiered structure across the northern hemisphere, to study the evolution of galaxies, their clusters, and the active galactic nuclei (AGN) that drive them. Its recent second data release has proven to be very successful, with ~ 4 million sources being identified over a 5700 deg^2 area (Jin et al., 2022).

However, additional multiwavelength data is needed to distinguish between radio sources and classify them. For LOFAR, WEAVE will be its primary source. Along with accurate redshift measurements, the spectroscopic data from WEAVE will also be used in the classification of AGN accretion modes, to estimate metallicities and velocity dispersions of objects, and calculate virial black hole masses (Jin et al., 2022) which will extend the scientific potential of LOFAR significantly.

The WEAVE-LOFAR (WL) survey makes use of all observing modes, and is divided into three tiers; wide tier, mid-tier, and deep tier. Its target selection is based solely on the flux limit and has an observational footprint that overlaps with the Hobby-Eberly telescope dark energy experiment (HETDEX) (Hill et al., 2008). As the WL survey has a higher resolution and longer wavelength coverage than HETDEX, it will be used to confirm the nature of some of the HETDEX sources and will help measure environments surrounding LOFAR sources. At this stage, there is no confirmation of the overlap of WEAVE with the Dark Energy Spectroscopic Instrument (DESI) survey (Aghamousa et al., 2016), however this will soon become

directly measurable in the coming weeks.

Apertif is the new focal-plane array receiver that has been installed on the Westerbork Synthesis Radio Telescope (WRST) (Verheijen et al. (2008a); Van Cappellen et al. (2022)). This upgraded system allows wide-field radio surveys to be conducted. From July 2019 to February 2022, Apertif observed a significant proportion of the Northern hemisphere at 1.4 GHz in order to survey the 21 cm neutral atomic hydrogen line out to $z \sim 0.2$. Apertif used a two-tier approach for their HI observations; a medium-deep tier survey covering an area of 150 deg^2, and a wide-area survey covering ~ 2200 deg^2. These surveys are crucial for understanding the evolution and morphology of galactic gas content. By cross-referencing Apertif observations with spatially resolved spectra from WEAVE's Large Integral Field Unit (LIFU), properties of the cold gas content of galaxies can be compared with spatially resolved properties such as ionised gas.

The requirements of WEAVE which are deemed desirable for the *Gaia*, LOFAR, and Apertif follow up surveys (Verheijen et al., 2008b) are described in the survey overview paper (Jin et al., 2022) and are listed in table 1.1. The corresponding observational requirements as a result of these science drivers are summarised in table 1.2. These three missions are the main driving force behind WEAVE; however, by building a flexible and versatile instrument its science case has rapidly expanded and will continue to do so over its survey lifetime.

WEAVE will take 1-hour exposures, with a detector read-out time of < 3 minutes. These overhead timings do not include the grid calibrations (see section 5.1) but do include the fibre configurations, telescope offsetting, and acquisition. These timing requirements will be evaluated during commissioning and discussed in chapter 3.

Instrument	Desirable requirements
Gaia	$R = 5000$ for radial velocities at $17 \leq V \leq 20$ mag $R = 20,000$ for stellar abundances at $12 \leq V \leq 17$ mag 10^7 stars over 10^4 continuous square degrees
LOFAR	Continuous wavelength coverage over 370-980nm $V \leq 21.5$ mag at $S/N = 5$ in continuum for redshifts 10^7 galaxies over 10^4 square degrees
Apertif	Large IFU and mini-IFUs for spatially resolved spectra of gas-rich galaxies 10^4 galaxies over 10^4 continuous square degrees

Table 1.1: The desirable requirements taken from (Jin et al., 2022) for the follow-up of the *Gaia* mission (Gaia Collaboration et al., 2016), LOFAR (van Haarlem, M. P. et al., 2013), and Apertif (Verheijen et al., 2008b) surveys

1.5 Systems overview

There are five main independent sub-systems of WEAVE; the prime focus system, the rotator system (WRS), the positioner, the optical fibres, and the spectrograph. WEAVE has three observing modes; the individual multi-object spectrograph fibres (MOS), the mini integral field units (mIFUs), and the large integral field unit (LIFU). A schematic diagram of the final arrangement of WEAVE at the top end of the WHT is shown in figure 1.8. This thesis will focus particularly on the positioner and the optical fibres. A detailed overview of the positioner and its final assembly process is given in chapter 2.

The design and operation of WEAVE are very similar to 2dF. The positioner uses a tumbler structure with two circular field plates, known as plate A and plate B. The tumbler can rotate 180° to place each field plate in either the observing position, which lies in the focal plane of the corrector, or the configuring position, which is in the plane of the two positioning robots. Each field plate has an independent set of fibres, stored in a set of retractors around its edge. This allows one field to be configured by the robots whilst the other field plate is observing and minimises the overhead time between subsequent exposures. The fibre cables from the positioner

Type	Requirement
Observational efficiency	The observational overheads for a 1 hour exposure should be <15 minutes. The desired overhead time is <3 minutes.
Target Packing (MOS and mIFU)	The separation between any pair of MOS targets must be ≤ 60" for a single pointing. The desired separation is ≤ 20".
Target Packing (IFU)	WEAVE will acquire individual targets centrally onto the FOV of an IFU to an accuracy of $\pm(1, 1)$ spatial elements. The desired accuracy is $\pm(0.1, 0.1)$.
Positioning Accuracy (MOS)	WEAVE shall be able to position each MOS optical fibre to within a positioning accuracy of 0.2" of the given astrometric position of each of its selected targets.
Guiding	The guiding accuracy shall be $< 20\%$ of the on-sky aperture size at the centre of the field for integrations of up to 1 hour.

Table 1.2: A summary of the observational requirements defined by the WEAVE collaboration, further details are provided in (Trager et al., 2013) and (Jin et al., 2022).

lead down towards the spectrograph on the Nasmyth platform, having been secured in the WRS cable wrap and along the telescope's prime focus spiders.

Table 1.3 is a summary of WEAVE's design specifications, with a 2% tolerance on this wavelength coverage in the low-resolution mode, and a tolerance of 1 nm in the high-resolution mode. It is an essential requirement that the spectral resolution is stable to less than 5% during consecutive 1-hour exposures and up to 10% over the duration of the main surveys ($\sim$ 5 year period of observing). Ideally, it will be less than 2%. These are driven by stellar radial velocity requirements from Gaia and to ensure that observations of non-binary and multiple star systems are repeatable.

An important definition used to quantify an instrument is its f-number, which is defined as the ratio between the focal length and the diameter of the aperture. The optics of this new top-end produce a plate scale of 56 μm/arcsec, with an incoming f/2.75 beam transforming to f/3.2 via the lenslets on the MOS fibre prisms. This value places a necessary condition on the fibre positioning tolerance to meet the

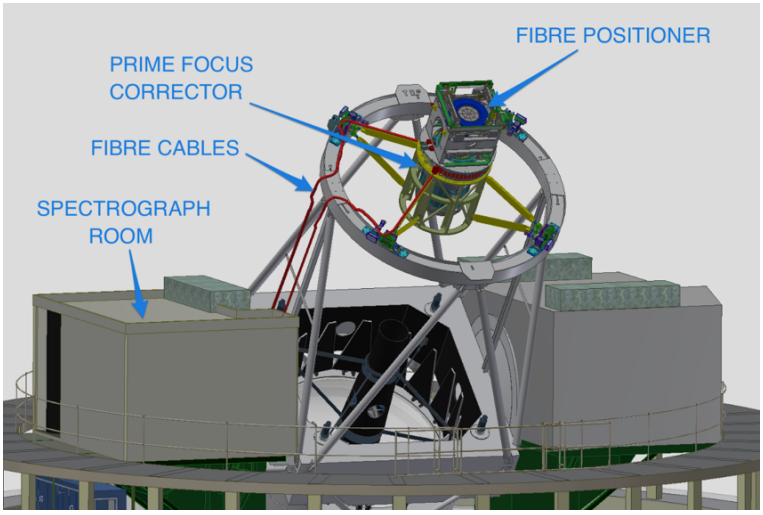


Figure 1.8: A diagram demonstrating the final arrangement of WEAVE on the top-end of the WHT, located at its prime focus.

Telescope	WHT, 4.2m ø
Field of view	2° ø
Number of fibres	960 (Plate A)/ 940 (Plate B)
Fibre size (at f/3.2)	1.3" (85μm)
Number of mIFUs, size	20 x 11" x 12" (1.3" spaxels)
LIFU size	1.3' x 1.5' (2.6" spaxels)
Low-resolution mode resolution	5750 (3000-7500)
Low resolution mode wavelength coverage (Å)	3660-9590
High-resolution mode resolution	21000 (13000-25000)
High-resolution mode wavelength coverage (Å)	4040-4650, 4730-5450, 5950-6850

Table 1.3: A summary of the design specifications for WEAVE (Dalton et al., 2014).

survey requirements stated in table 1.1. Breakdown of this tolerance is described in chapter 2.

1.5.1 The prime focus corrector

The prime focus system consists of a six-lens wide-field corrector, which contains an atmospheric dispersion compensator (ADC), an instrument rotator (see section 1.5.2), and a new top-end structure with focus/tilt adjustment. This is aligned to the back distance of the positioner. The ADC components comprise two air-spaced prismatic doublets which counter-rotate to correct for dispersion along the direction

of the parallactic angle. The corrector is located in front of the positioner. They are bolted together with the WRS to ensure that the optical fibre prisms will lie in the corrector's focal plane. The optical layout and arrangement of the lenses are shown in figure 1.9.

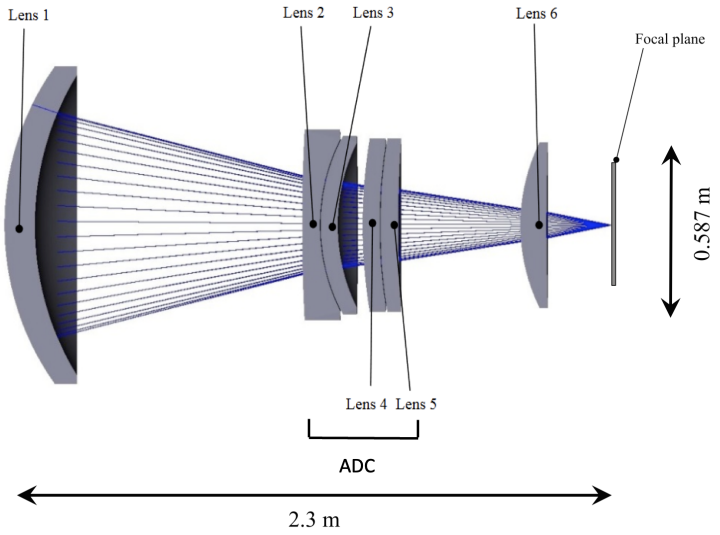


Figure 1.9: The optical layout of the Prime Focus Corrector, demonstrating the position of all six lenses (Dalton et al., 2014). The lenses are aligned to a tolerance of 50 μm (Tomás et al., 2018), and the corrector prescription delivers $< 0.6"$ polychromatic images over a flat focal surface.

A summary of the design specifications of the Corrector is shown in Table 1.4. The ADC can only compensate for the dispersion effects of the atmosphere. Other effects, such as differential refraction[2], are corrected for as much as possible in the placement of the fibres in the field.

Parameter	Design specification
Field of View	2°
Focal surface	410 mm diameter FoV
Telescope Focal Length	11.55m
F-number	f/2.75

Table 1.4: A summary of the design specifications for the Corrector (Schallig, 2019).

[2]An effect caused by changes in the air column density as a function of zenith distance. This causes a distortion across the field of view that varies in both amplitude and direction.

The prime focus system has a new top-end structure to accommodate the size and mass of the positioner. The original WHT secondary flip-ring structure was removed, and a new set of spiders has been mounted directly to the outer ring. A series of cam adjusters[3] are located inside the ring mount, which allows for fine adjustment of the focus and tilt. The combination of these structures is the focus translation system (FTS) (Dalton et al., 2020). Further details on the development and testing of the prime focus system can be found in (Dalmases et al., 2018), (Tomás et al., 2018), and (Dalton et al., 2020).

1.5.2 The rotator

The WRS is designed to compensate for the field rotation effect due to the alt-azimuth mount of the WHT and to protect the optical fibre cables as the telescope moves. In March 2019, the WRS arrived at the observatory and was successfully mounted to the positioner after its arrival in 2021. The two systems attach via an interface plate. Extreme care was taken in designing and manufacturing the interface plate to ensure that the alignment between the positioner and the WRS is correct. The WRS and positioner subsystems were aligned and attached in 2021 by the ING engineering team after both components arrived at the observatory.

The design for the WRS can be seen in figure 1.10. It is based on the rotator system produced by the company IDOM for the Gran Telescopio Canarias (GTC) and was chosen for its compact design yet improved cable accessibility. It consists of a cross-roller bearing which is driven by a frameless direct-drive motor. Further details on the design and construction of the WRS subsystem can be found in (Vicente et al., 2020).

[3]A type of hydraulic control plunger, activated with magnets. It is used to make small, timed increments.

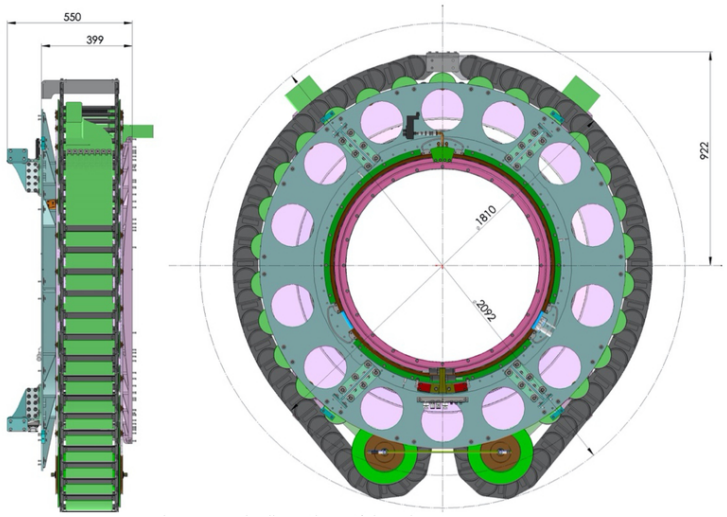


Figure 1.10: A diagram of the main components of the WEAVE rotator system (Vicente et al., 2020).

1.5.3 The positioner

As mentioned previously, the WEAVE fibre positioner (POS) uses pick-and-place technology with a central tumbler structure that closely follows 2dF. A fundamental theme of WEAVE's design is to use existing and readily available technology wherever possible, and the positioner is no exception to this. It consists of two commercial cartesian robotic arms, supplied by the company SCHUNK, that share a common x-axis to place the optical fibres, and a third independent robot to image the fibres in the focal plane. A detailed overview of the positioner and the integration of the optical fibre sub-system is described in chapter 2.

1.5.4 The optical fibres

As summarised in table 1.3, WEAVE has three observing modes. The positioner has a multiplex of 960 MOS fibres, 20 mIFUs, and one LIFU. Each mIFU contains a bundle of 37 optical fibres arranged hexagonally, which are located on plate B. The LIFU contains a hexagonal bundle of 547 optical fibres and is mounted between the two halves of the tumbler drum, which is rotated 90° for observations. The observing

mode specifications are summarised in table 1.5.

Observing mode	Fibre size	Unit type	Multiplex	Field of view	Configuration time
MOS	1.3”	Single fibre	960[a]	1.3”	$\sim$ 55 minutes
mIFU	1.3”	IFU	20	11” $\times$ 12”	$<$ 20 minutes
LIFU	2.6”	IFU	1	$1.3' \times 1.5'$	$\sim$ 1 minute

Table 1.5: The main parameters for the three observing modes of WEAVE. [a]Plate A has a multiplex of 980 optical fibres whilst Plate B has 960, this is to accommodate the guide fibres and the mIFU retractors.

The fibres are sorted using custom cables to safely route the fibres from the positioner to the spectrograph, whilst providing resistance to fatigue and shock, along with increased flexibility (Mignot et al., 2020b). At the spectrograph end, the MOS fibres are divided into groups of 24 as they line the slit, and towards the positioner, the cables gradually separate into smaller cables with groups of six fibres. The MOS optical fibres are stored around the edge of the field plate in a pulley system called a retractor. Each retractor contains six fibres, with three independent pulley systems on each side, in a tiered arrangement. The positioner contains 336 retractors in total, 168 located around each plate. The structure of the retractors and their installation is described in chapter 2.

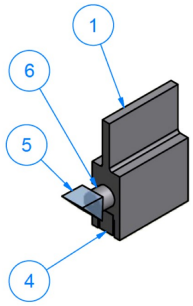


Figure 1.11: A computer-aided design (CAD) of the MOS fibre button structure, where (1) is the button vane, (5) is the prism, (4) is the magnet, and (6) is the ferrule.

On the end of each optical fibre is a magnetic stainless steel button that the fibre is passed through. The button design follows directly from 2dF and is shown in figure

1.11. Its outer structure is shaped to allow the robot gripper mechanism to grasp it, and a magnet is embedded in its base to secure it to the field plate. Each magnet is carefully glued in the same orientation in the internal slot for the magnet. The shape of the slot ensures that its magnetic field is completely contained, which prevents neighbouring fibres from snapping together when placed in close proximity during observations.

A prism is then attached to the polished fibre end to guide the incoming light along the optical fibre. Each fibre is fed through a buffer called a ferrule, which protects the fibre as it sits along the length of the button. This structure is glued in place at a fixed distance to prevent the button from pushing against the prism as the fibre is withdrawn, which could cause the prism to detach.

The prisms for the MOS fibres are approximately 1 mm in diameter, with a concave lens formed into its top surface. This is to direct incident light into the fibre efficiently, by transforming the incoming beam from f/2.75 to f/3.2. By reducing the f-number of the beam, the design for the spectrograph collimator was simplified significantly. Further details on the assembly and testing of the optical fibres can be found in (Sayède et al., 2016), and (Mignot et al., 2020a).

1.5.5 The spectrograph

The WEAVE spectrograph, constructed by NOVA in the Netherlands was successfully delivered to the WHT in August 2020. The spectrograph is contained within the ground-based high-resolution imaging laboratory (GHRIL) enclosure, which is located on one of the Nasmyth platforms of the telescope. Placing the spectrograph here improves the stability of the system as the telescope moves, and reduces restrictions on the size of the spectrograph due to weight and balance. This requires a longer fibre run and consequently introduces loss due to scattering as the light propagates. Information on the spectrograph's performance testing can be found in (Rogers et al.,

2014) and (Stuik et al., 2018).

The spectrograph has a dual-arm configuration, containing one red arm and one blue arm. This is achieved using a dichroic filter at approximately 600 nm. It has been designed to give continuous wavelength coverage ($3660 - 9590$Å) in the low-resolution mode ($R \sim 5750$) and coverage over several narrow regions (either $4040 - 4650$ or $4730 - 5450$, and $5950 - 6850$Å) in the high-resolution mode ($R \sim 21000$). The HR ranges were selected to ensure that the Hβ and Mg b lines were both covered, with the aim to also include the Fe line complex around 5270Å. Both arms of the spectrograph use volume phase holographic (VPH) gratings[4], with an exchange mechanism in each arm to switch between the low and high-resolution modes. The system has an interchangeable slit input so that the input can be switched between the MOS, mIFU, and LIFU modes with ease, whilst simultaneously back-illuminating the fibres in the configuring position. This is necessary for the robots to image and locate each fibre.

The incoming light from the slit is aligned using a collimating mirror and a series of lenses are located in each arm of the spectrograph. The collimator produces a 190 mm diameter beam, which is then passed through a dichroic filter that splits the beam and directs them into the VPH gratings for dispersion (Izazaga et al., 2018b). A diagram of the optical path through the spectrograph can be seen in figure 1.12.

The spectrograph has two cameras, each coated in an anti-reflective layer that is specific to the respective arm of the dual-beam configuration (366 to 606 nm for the blue camera and 579 to 959 nm for the red camera). Both cameras consist of 8 lenses, where one of them is aspherical, and each camera has 8k (spectral) $\times$ 6k (spatial) pixels (Dalton et al., 2014). Details of the manufacturing and assessment of the camera lenses can be found in (Izazaga et al., 2018a), and the collimating mirror

[4]A VPH grating works by varying the refractive index of the material itself, often perpendicularly to the surface fringes. They have high efficiencies and their shape can be easily customised, making them ideal for use in astronomy.

in (Izazaga et al., 2018b).

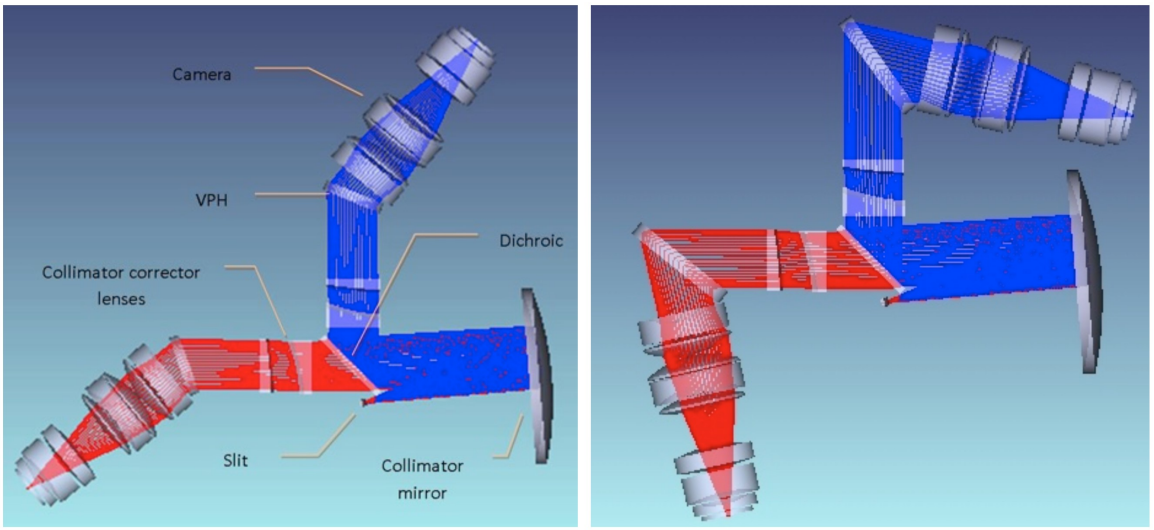


Figure 1.12: The optical arrangement of the red and blue arms of the Spectrograph for both the high (left) and low (right) resolution modes of WEAVE.

1.6 This Research

WEAVE is the new multi-object spectrograph for the 4.2 m WHT, with a multiplex of 960 individual optical fibres, 20 mini IFUs (mIFUs) and one large IFU (LIFU), which are deployed by its fibre positioner. In December 2020, the WEAVE fibre positioner was shipped from the laboratory in Oxford and following a period of repair work, metrology measurements, and calibrations, first light was successfully achieved in October 2022.

The final assembly of the positioner began with the installation of the MOS and mIFU optical fibre cables from December 2019. I and members of the WEAVE-Oxford group spent approximately one year carefully handling and feeding in the cables by-hand, whilst maintaining a record of each fibre's location within the positioner, alongside information about their cables and slitlet assignment. I then used this information to generate a series of 24 configurations for each plate, so that the trace of each fibre's spectra along the spectrograph detector could be determined on-sky.

This technique was successfully used during the commissioning period and the fibre traces have now been incorporated into the Core Processing System (CPS) of the WEAVE pipeline.

Due to the COVID-19 pandemic, I and Ian Lewis assembled the MOS and mIFU slits which interface with the WEAVE spectrograph. The entire process took two weeks, and required the careful application of index-matching gel to each slitlet, and installing them along the slit whilst preventing air bubbles occurring along the interface.

The focus of this work is on the calibration, commissioning, and first-light results of the WEAVE fibre positioner. We begin by providing a detailed overview of the positioner's structure and functionality in chapter 2. This includes a description of the retractor installations and modifications made throughout 2020, prior to the positioner's shipment.

Chapter 3 describes the calibration procedures conducted both in-lab and after the positioner's arrival at the observatory. It was first aligned and attached to the rotator subsystem before being mounted to a testing rig, which was used to simulate the movement and orientation of the positioner on the telescope.

In chapter 4, I conducted a series of simulations to optimise the fibre allocation process, using fields supplied by members of the WEAVE-LOFAR team. This analysis explored the distribution in the fibre assignment and was used to identify potential selection effects.

Chapter 5 presents the research conducted during WEAVE's commissioning period, including the astrometry of the night sky to the focal plane, and the mapping of the fibre slitlets within the spectrograph. This chapter presents the first-light results for the LIFU observing mode and will discuss the extent to which WEAVE has met its observational requirements. This is continued in chapter 6, which summarises this research and presents the future outlook for WEAVE in the coming years.

This is an exceptionally exciting time within astronomy, as many new instruments are achieving first light, with WEAVE sitting at the forefront of ground-based observations. WEAVE will provide a versatile and efficient method of observing the northern-hemisphere sky, and its spectroscopic observations will unlock new science from existing instruments such as GAIA and LOFAR.

Chapter 2

The WEAVE fibre positioner

> So you actually read, you're one of those.
>
> J. Wenninger

This chapter will provide a detailed overview of the WEAVE fibre positioner, with a focus on its structure and functionality. It includes the final assembly process and the challenges that this presented. The installation of the optical fibre systems into the positioner began in December 2019 until September 2020[1]. It was completed by myself alongside members of the WEAVE Oxford group. The positioner successfully arrived at the WHT in December 2020, after a period of in-lab operational testing.

2.1 Structural overview

The WEAVE fibre positioner uses pick-and-place technology that directly follows from 2dF. As with 2dF, WEAVE has a central tumbler structure for continuous flexible configuring of its optical fibres. The tumbler is split into two symmetrical halves separated by a central beam, which forms the rotation axis. This is shown in figure 2.1. At one end of the beam is the motor which drives the tumbler's rotation, and at

[1]This covered the first COVID-19 lockdown from March 2020, where Gavin and Ian were the only members allowed into the department until I rejoined them in July 2020.

the other end is a bearing. This allows the tumbler to rotate up to ±175°. The hollow central beam provides a path for the fibres to be stored and fed through during their installation, and the additional space surrounding the beam hosts the LIFU on one side and the MOS auto-guider and back-illumination unit on the other side.

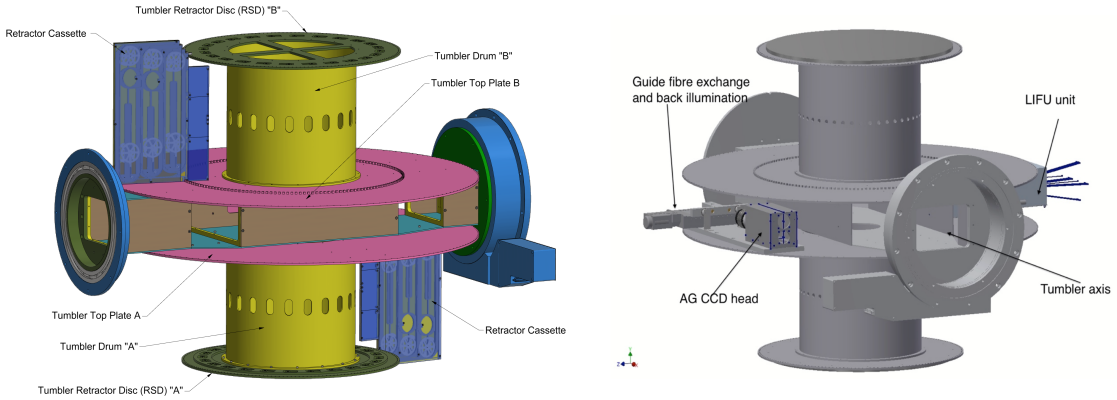


Figure 2.1: (*Left*) A diagram of the tumbler structure and its components, including the location of the retractors relative to the tumbler. (*Right*) A diagram of the tumbler structure demonstrating the position of the LIFU and the guide fibre back illumination and exchange unit.

An aluminium top plate is attached on each side of the beam's length to support a drum. The drum ends with a structure called the retractor support disk (RSD), and two invar circular field plates, known as plate A and plate B. These field plates have an active radius of 205 mm and have been ground and polished until their surface is flat to ±4 μm when bolted down. Invar is a unique material known for its low thermal expansion coefficient, which makes it an ideal alloy for manufacturing field plates. However, its crystalline structure takes an exceeding long time to settle, with its dimensions changing up to six months after production. To improve the invar's dimensional stability, it is artificially aged via careful temperature control, however this comes at the cost of increasing the coefficient of thermal expansion (Wang et al., 2013). Measurements of this surface and its implications are discussed in chapter 3.

Both field plates have a series of small holes, which have been carefully filled

with black paint to form grid spots across the surface. These grids are used to transform between the coordinate systems which exist inside the positioner (Schallig, 2019), they are the robot coordinates of the Focal Plane imager (FPI, see section 2.2.1), the cartesian robot coordinates (see section 2.2), the plate coordinates, and the on-sky coordinates. The grid spots are imaged by the upper focus of the camera attached to the positioner's robot gantries , and by the plate viewing camera of the FPI. By measuring the position of each spot across the field plate, they can be used as reference marks to remove linear affects such as scale changes, rotation, or the non-perpendicularity of the robot gantries, as well as non-linear effects such as the flexure of the gantries due to gravity. The incorporation of these grid points into the calibration software is described fully in (Schallig, 2019).

Each half of the tumbler supports 168 individual retractor units, which bolt to both the RSD and the top plate, stiffening the entire structure. In the configuring position, the field plate is accessible by the positioning robots. When observing, the field plate is 2.4 mm behind the corrector focal plane of the telescope. This 2.4 mm difference accounts for the button and prism geometry to place the entrance of the fibre at the focus. An important design specification for WEAVE is that the difference in height across the field plate must be less than ± 15 μm edge-to-edge when observing, including the relative tilt (Hughes et al., 2020). The x-axis tilt in the observing position can be compensated for by adjusting the rotation angle of the tumbler for each plate. In the configuring position, residual tilt can be partially compensated for during the placement of the fibres. Depending on the orientation of the tumbler, each plate is in either the configuring position or the observing orientation, with the exception of the LIFU observations. In this case, the tumbler is rotated to 90° so that the LIFU aperture lies in the focal plane. When the tumbler is in this position the field plates are inaccessible to the robots.

The tumbler is supported by the positioner's outer structure, which also supports

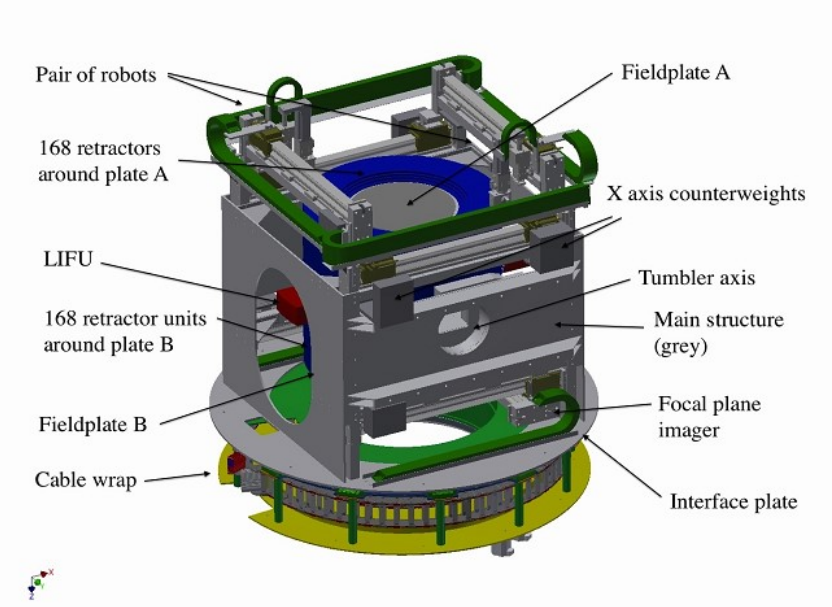


Figure 2.2: An early diagram of the Positioner, labelling the main components. Small changes have been made to the arrangement of the Positioner since this diagram was produced, such as the precise location of the x axis counterweights however, the overall structure remains unchanged.

the Cartesian robot gantries and their counterweights. An interface plate is located at the base of the positioner, which is used to align the positioner to the WRS system. This interface plate is used as an absolute reference point and is the basis of the positioner's metrology (Schallig, 2019). A diagram of the positioner's structure can be seen in figure 2.2. The positioner's mass is approximately one tonne and its dimensions are 2m x 2m x 2m including its protective covers.

2.2 Cartesian robots

WEAVE uses three commercial robotic gantries named by the WEAVE team; Nona, Morta and Decima[2]. Nona and Morta are two Cartesian robots which move along gantries in the xy plane on the configuring side of the positioner, with both robots sharing a pair of x gantries and independent y gantries. The robots have individual

[2]Nona, Morta, and Decima are names from Roman mythology. They are the three sisters who represent destiny and control the thread of life.

z and θ-axes, giving them a full kinematic range when orientating themselves over the field plate. Decima is referred to as the Focal Plane Imager (FPI) and scans the telescope focal plane. Further details on the FPI calibration is presented in section 2.2.1.

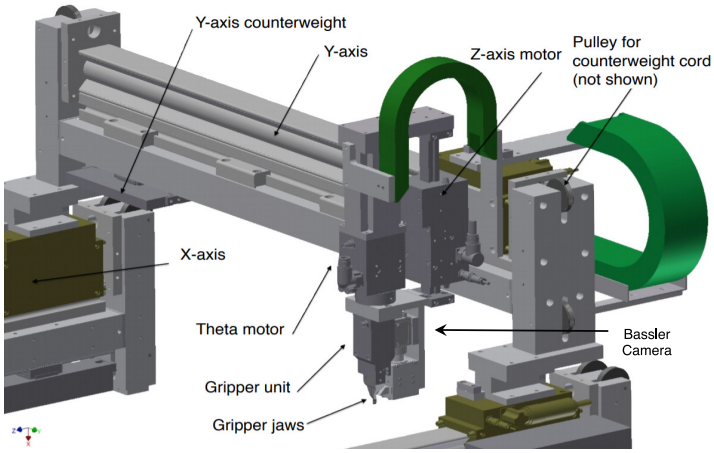


Figure 2.3: A CAD diagram of the Cartesian robotic system for WEAVE. The z-axis linear motor slide was changed in 2019 to prevent a systematic drift in positioning with the axes temperature (Schallig, 2019).

Both Nona and Morta possess a gripper mechanism and camera system, shown in figure 2.3. The profile of the gripper increases in width in order to prevent the jaws from flexing. The camera has two focus heights; an upper focus to scan the field plate grid, and a lower focus to image and centroid the fibre prisms (Gilbert, 2016). Either Nona or Morta may pick up an individual fibre from the park position, or the field plate. When configuring a field of fibres, both robots move in a combined 5-axis kinematic. They have overlapping patrol regions and proximity limits to prevent collisions. The configuration of a single WEAVE field is designed to take approximately one hour to complete and will be discussed in chapter 3.

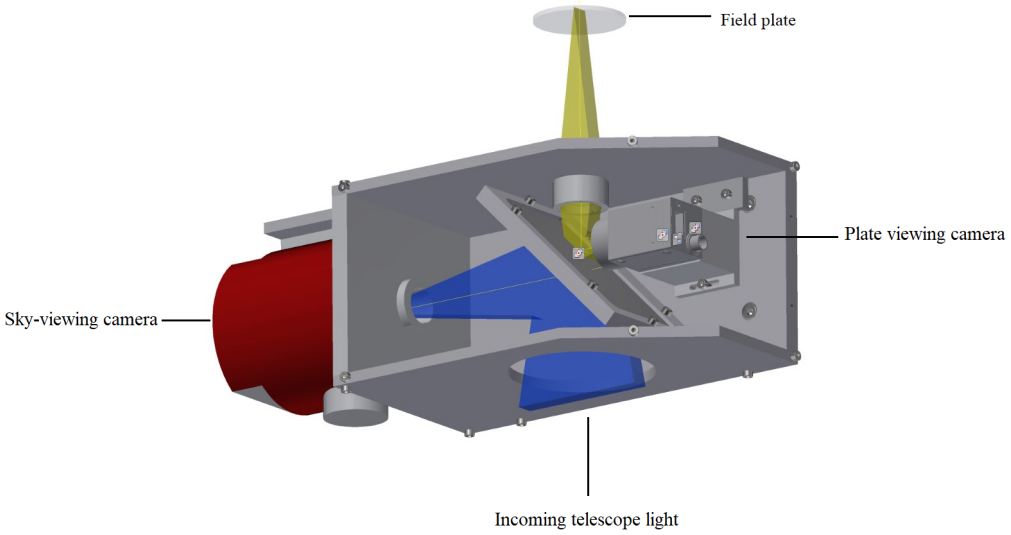


Figure 2.4: A CAD diagram of the FPI, with the incoming light path from both the sky and the focal plane.

2.2.1 The Focal Plane Imager

The FPI is mounted to the y-axis motor of the third positioner robot Decima, and can be found near the interface plate. It is designed to image both the night sky and the observing field plate so that a mapping can be generated between the two surfaces. This procedure is extremely important for the astrometry of WEAVE, which will be discussed in chapter 5. Unlike Nona and Morta, the FPI can only move in the xy-plane and consists of two cameras; a thermo-electrically cooled camera (Bigeye G-123B Cool) with a pixel size of 6.45 μm (0.113"), and a Basler camera (Hughes et al., 2020). A diagram of the FPI's structure can be seen in figure 2.4.

In September 2020, Ian Lewis and I began the procedure[3] to measure and adjust the focus of the FPIs plate viewing camera and the internal optics of the BigEye sky-viewing camera. This required the use of an optical bench to arrange the components at the correct distances before setting the focus. This setup is shown in figure 2.5 and the steps are given as follows:

1. The distance between the FPI ledge and the field plate must be carefully ar-

[3]This procedure was developed by Ian Lewis prior to the start of my DPhil.

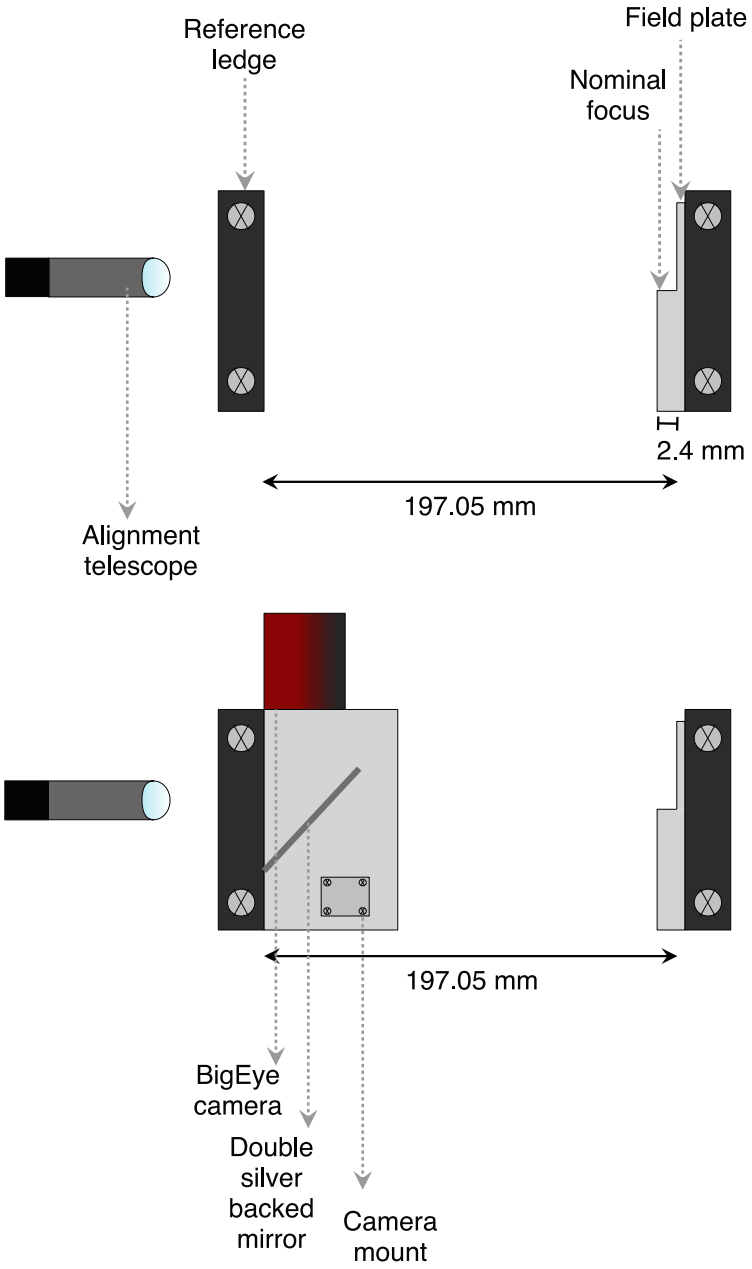


Figure 2.5: A diagram of the two main stages of the FPI focus setup. (*Top*) demonstrates two mounts set at a fixed difference corresponding to the distance between the FPI reference ledge and the field plate (Schallig, 2019). This is then used to ensure that the dummy field plate is visible yet slightly out of focus compared to the nominal focus position. (*Bottom*) The FPI placed against the reference ledge, where the internal optics can be adjusted to ensure that the BigEye camera chip is in focus and parfocal with the plate viewing camera.

ranged to be 197.05 mm apart. Two mounts were arranged on an optical bench to simulate this spacing, where one represents the position of the FPI and the other represents the dummy field plate.

2. On one mount, a reference surface was created to simulate the field plate and nominal focus position using an aluminium block, where half of the surface face protrudes by 2.4 mm. This is designed to highlight the difference in height between the nominal focus on the fibre buttons and the field plate itself, where the surface is flat to ± 10 μm. This reference surface was covered in a high-contrast target to provide visual aid as the focus of the plate-viewing camera was adjusted.

3. A small alignment telescope was placed behind the mount that represents the FPI ledge, at a distance of approximately 600 mm. The focus of the Basler camera was manually adjusted until the focal plane reference surface (the upper surface) was sharply defined whilst the half representing the field plate could still be seen slightly out of focus.

4. At this stage, it is crucial not to make any further manual adjustments to the alignment telescope focus. The FPI was placed with its lower surface against the reference mount and the alignment telescope was used to image the CCD chip of the BigEye camera, which should now be in focus.

5. To refine the focus of the image on the CCD chip, the position of the double-backed silver mirror was adjusted until the corner of the BigEye camera chip could be seen clearly with a sharp definition. This ensures that the BigEye camera is parfocal with the focal plane of the telescope.

6. With the 45° double-backed silver mirror in position, the plate viewing camera was installed and its focus was checked by imaging the further side of the ref-

erence surface, which represents the field plate. Images taken throughout this process can be seen in figure 2.6.

Setting the FPI focus was an iterative process, demonstrated by the right-hand side of figure 2.6. This image demonstrates a clear tilt in focus when imaging the grid points along the field plate, which must be removed. This was resolved by replacing the mount of the double silvered-fold mirror. Any changes to the focus must be completed by dismounting the FPI and using the same optical setup described above and cannot be performed on-sky. The quality of the images produced by the FPI will be described further in chapter 5.

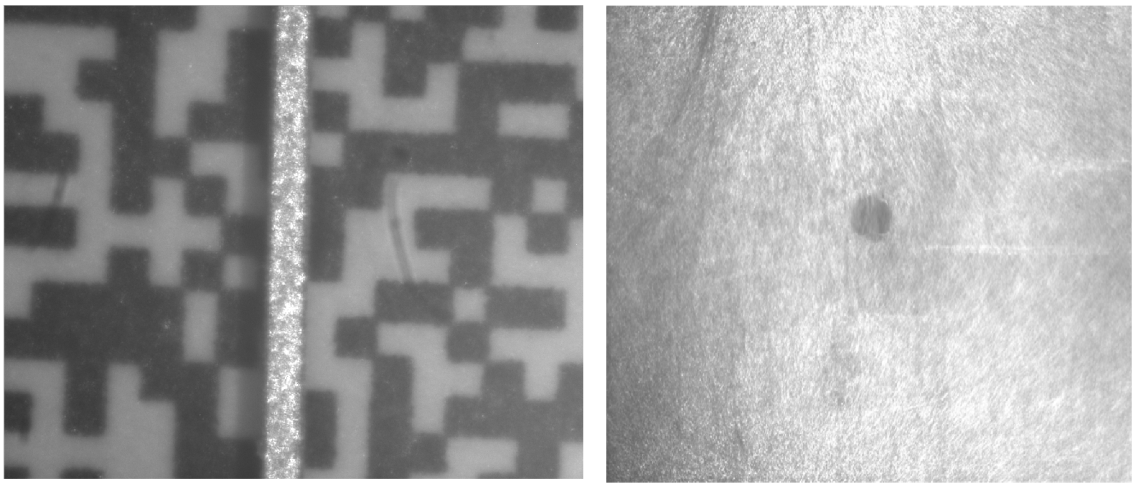


Figure 2.6: *(Left)* The surface used to adjust the focus of the FPI's BigEye camera, as seen by the alignment telescope. Each half of the image has a printed QR code as a reference surface. The surface on the left-hand side represents the field plate, which is visible yet slightly out of focus in comparison to the nominal focus on the right-hand side. The reference surface representing the nominal focus is 2.4 mm closer to mimic the location of the buttons in the focal plane. *(Right)* An image of a grid spot taken using the FPI camera. A focus gradient can be seen, which was due to the positioning of the Basler camera, as its adjustment axis was slightly tilted.

2.2.2 Positioning requirements

To ensure WEAVE could complete the necessary follow-up observations, the magnitude limitations of *Gaia* were used as a basis to define the initial positioning toler-

ances. The derivation of these requirements are presented in (Dalton et al., 2014) and (Gilbert, 2016).

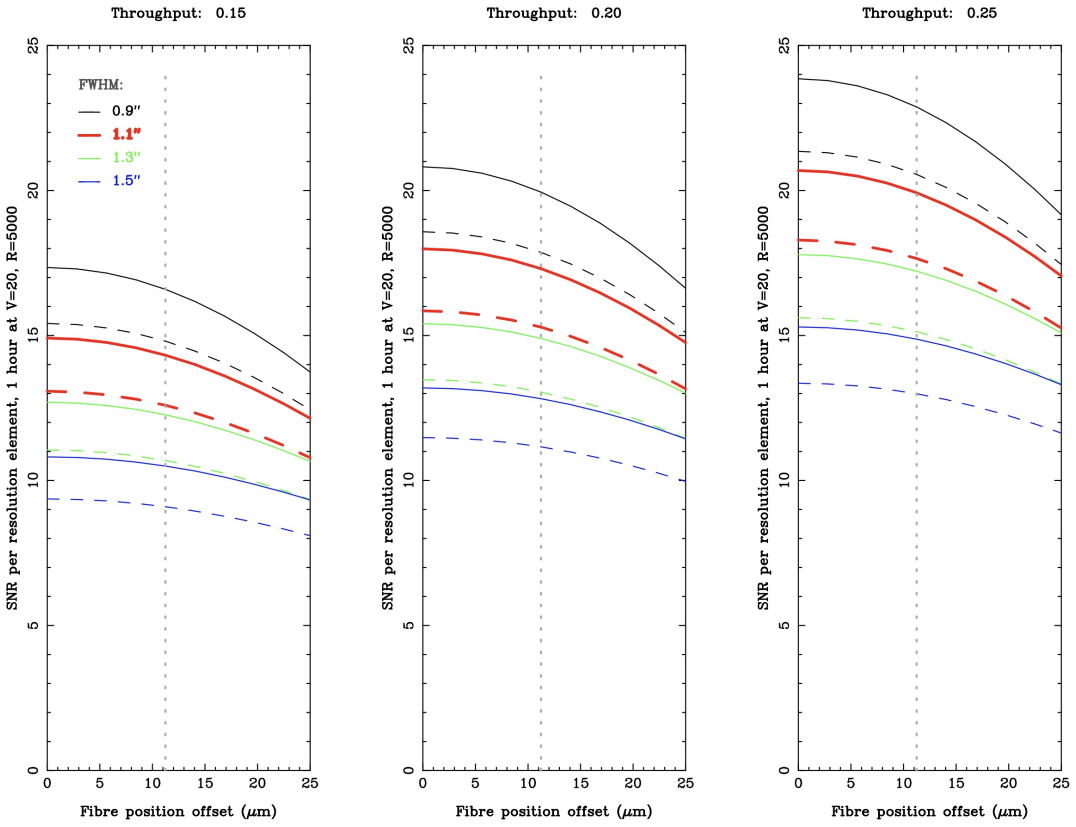


Figure 2.7: The simulated SNR for a 3-pixel wide spectral resolution element, with a nominal faint magnitude limit $V = 20$ and a 60-minute exposure time. The readout noise of the detector is taken as 2.5 e-, along with 1.2 for the airmass and a 1.3" diameter fibre core. Each sub-figure represents a different system throughput within the anticipated range; each line represents the post-PFC PSF width in different seeing conditions, solid lines being for a best-case (Gaussian) and dashed lines for the worst-case (Moffat, $\beta = 3$) PSF shapes. The bold (red) lines correspond to a reported 75th percentile seeing of 0.800 for the WHT. The dashed vertical line shows the combined maximum offset requirement of 0.2" for the PFC and positioner subsystems combined [figure credit to Gavin Dalton, with additions from James Gilbert, and appears as figure 1.6 in (Gilbert, 2016)].

With *Gaia*'s nominal faint magnitude limit ($V = 20$), the positioning accuracy is a method of increasing the signal-to-noise ratio (SNR) of the data. Combining this with the seeing conditions at the WHT, the fibre 85 μm diameter, and the simulated throughput of the optical system, the SNR per resolution element can be estimated.

This was simulated by Professor Gavin Dalton in figure 2.7.

To achieve the scientific goals defined in section 1.4.1, the offset of the optical fibres must not exceed 0.2". This is shared between the prime focus corrector and the positioner (Gilbert, 2016), and corresponds to ± 12 μm. It is expected that $\pm 8 \mu$m will be occupied by the placement of the fibres on the field plate (known as the positioning tolerance). The positioner budget is presented in table 2.1[4]. Exceeding the total misalignment contribution will result in a reduction of the observed magnitude and a lower SNR in the spectra (Schallig, 2019).

Description	RMS value in μm
Centroiding of the fibre image	$\ll 1$
Gripper release repeatability	3.6
Robot Gantry accuracy	7.1
Coordinate reference frame correction	4.0
Gantry flexure uncertainty	2.0
Total misalignment contribution from POS	9.1

Table 2.1: A breakdown of the fibre positioning budget allowed for the POS, from (Schallig, 2019) and reproduced with permission from (Middleton, 2018). The *fibmeasure* software is credited to (Gilbert, 2016).

When conducting placement testing with the prototype fibres, the positioning accuracy was $\sim \pm 20$ μm. I developed the calibration procedures with the aim of understanding the behaviour of the system during the fibre placement, so as to consistently place each fibre on the first attempt. Chapter 3 will evaluate whether these calibration procedures have been sufficient. The resulting SNR from WEAVE's first-light observations is discussed in chapter 5.

[4]It is unclear how accurate the gripper repeatability value is at the time of writing this thesis.

2.3 The retractors

As mentioned in section 1.5.4, the positioner contains a total of 336 retractors, where each retractor stores 6 optical fibres. The retractors were installed by hand over the course of one year by myself and members of the WEAVE-Oxford group. This section describes the assembly and installation process, and the ongoing quality control challenges that have presented themselves throughout commissioning.

2.3.1 Assembly

Installing each retractor into the positioner is a delicate procedure that requires two people and takes an hour on average to complete. The fibre buttons are placed onto a thin metal strip, which secures them as they are passed up through the central drum of the positioner[5]. One person feeds in the fibres through the central tumbler beam and another reaches in from the top of the positioner to draw them through the correct slot of the drum. The fibres are then laid out on a nearby workstation before being removed from the metal strip one by one and fed into the retractor mechanism. The structure and arrangement of the retractors can be seen in figure 2.8.

A single retractor contains twelve pulleys, arranged into three tiers on each side. On the right-hand side of figure 2.8, the spare length box can be seen to the right of the pulley systems. This box stores the additional length of the fibres, which allows the fibre to be re-terminated if its end becomes damaged. Each fibre from the metal strip is fed around a pair of pulleys and its button is placed onto a metal porch. The top pulley is attached to a hub wound with a constant-force spring, whilst the other spins freely about a perpendicular axis. When the fibre is withdrawn, the spring extends as its pulley moves downwards. This allows the optical fibres to be moved back and forth under constant tension. The fibre length leading out of the retractor is pressed into a tightly-fitting groove along the retractor frame and the side panels

[5]We had a very handy piece of plastic railing for this process that I named Colin.

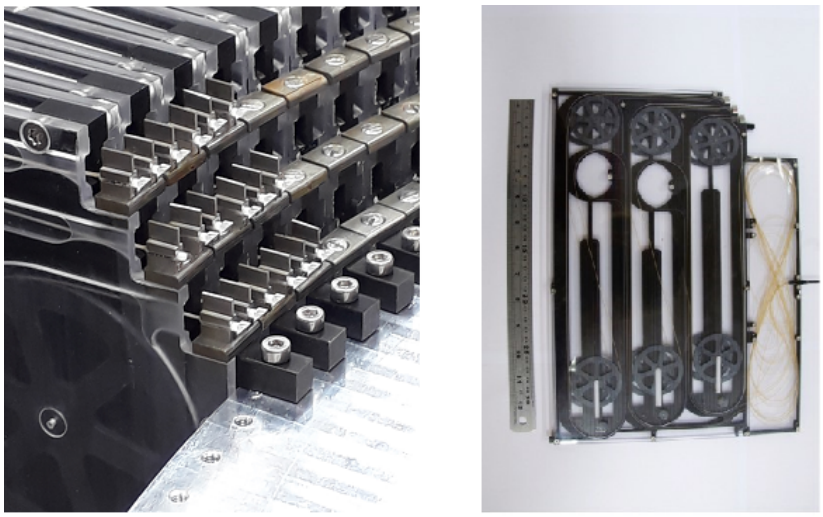


Figure 2.8: (*Left*) An image of three retractors installed into the Positioner. Each retractor contains six optical fibres, arranged across three tiers. (*Right*) An image of a single retractor, demonstrating its internal structure and the arrangement of the pulleys (credit to Ellen Schallig for this photo).

are attached once all six fibres have been integrated into the pulley systems.

The movement of each fibre is tested manually to identify potential movement problems. Examples include pulleys failing to spin, the bottom pulley becoming jammed along its track so it cannot withdraw the fibre, and a cogging movement as the pulley spins. To prevent these failure modes, the retractors components are inspected and tested extensively. The quality control evolved during the installation period, as we began to recognise the common signs that a retractor will fail. The list below summarises the potential problems and solutions that arose during their installation:

- A bent hub axle or twisted spring, which can both be replaced.
- A protruding yoke screw, which is fixed by over-driving and easing off the torsion.
- Debris or flash in slots or along the pulley itself, which can be removed.
- A ribbed surface of the slot in the outer cover, which can be sanded down.

The time taken to assemble a single retractor significantly increases if these features are present, and may take several attempts at assembly before it is acceptable. Components can be easily replaced due to the number of spares available, however this may not resolve the underlying issue. If the retractor passes the manual testing, it is bolted radially to the field plate drum at the RSD. If any issues are identified, the retractor is re-opened and inspected for the source of the problem. A difficult retractor may take approximately 3 hours to assemble after several repeated attempts. The successful installation of the plate B fibre system can be seen in figure 2.9, which took approximately six months to install. In the WEAVE design specifications, up to 5% of the MOS fibres may be disabled during survey operations however, inspections and interventions should take place at regular intervals to reduce this number.

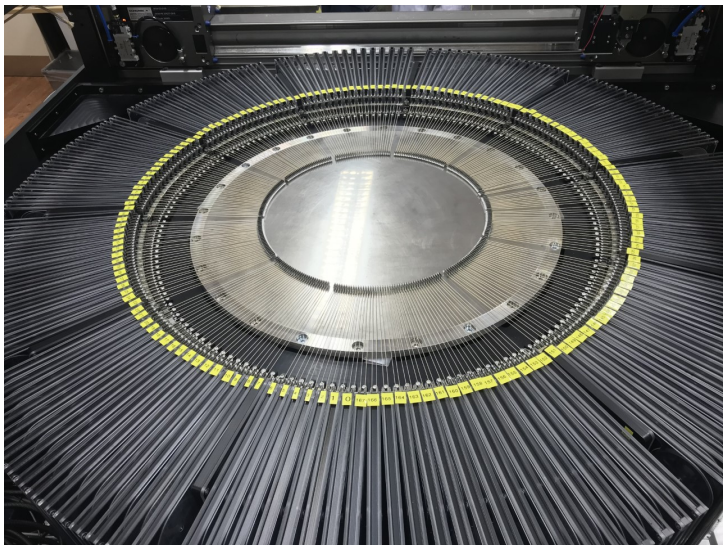


Figure 2.9: The complete installation of all 168 retractors on plate B, deployed in a ring configuration to image the prism of each fibre. The gaps correspond to the location of the mIFUs.

Midway through the retractor installations for plate A, a complication was discovered. When successive retractors were bolted to the tumbler, the left-hand tier 1 fibres would consistently stall when withdrawn. This was not observed during the manual testing of individual retractors and was caused by their compact arrangement around the edge of the field plate. Neighbouring units were applying pressure at their

contact point, which coincided with the slot for the left-hand tier 1 pulley hub. This additional pressure prevented the pulleys from spinning freely, causing the system to lock in place. The solution was to reduce the retractor's thickness by machining their covers down by ~ 100 μm over the contact region. This involved retrofitting approximately half of the plate A retractors, which were removed, modified, and replaced. This successfully resolved the behaviour and the installation process could continue.

2.3.2 Quality control

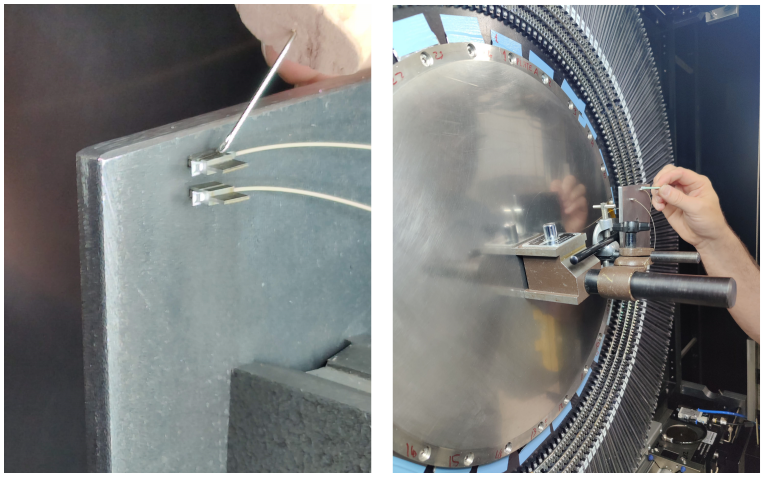


Figure 2.10: (*right*) The repair process developed after the positioner was successfully mounted to the WHT. A pair of magnetic clamps holding small metal plates are secured to the circular field plate, to secure the fibres at an appropriate height. (*left*) A close-up image of the in situ repair process, as the ferrule is cleaned of previous adhesive and superglue is fed into the slot after the prism has been re-positioned.

It is crucial that the retractors operate smoothly during observations. The highest risk failure mode is when a retractor stalls and be cannot be retracted. This can lead to loops of fibre around the edge of the field plate caused by the jaws leaving the fibre behind. The gripper units may then catch on these loops as the robots move, damaging the nearby fibres. Additional failure modes such as the prism being removed or the ferrule slipping back through the button can be fixed in situ. As configuration testing began, an increasing number of ferrules that had slipped back were observed,

causing the fibres to be preemptively disabled to prevent the prism from detaching. This repair work can be seen in figure 2.10 and was completed in several stages over the course of 2022 and early 2023.

Despite the care taken to when assembling and installing the retractors, extensive repair work was required on WEAVE's fibre sub-system after its arrival at the WHT. Figure 2.11 shows the location of the fibres systems that have been repaired across both plates as of January 2023. This is a cumulative list that does not represent the current state of the fibre system, as many of these fibres have been brought back into operations.

As of February 2023, there have been 197 retractors repairs for MOS A and 80 for MOS B thus far. The disparity between them is due to the ongoing discovery of loose ferrules on MOS A, which comprises 108 of these entries. This is the result of quality control differences between the two companies that manufactured the optical fibres for WEAVE. MOS A contains the majority of the cables that require additional maintenance. Additional complications such as magnets and prisms detaching, and stalled retractors are also highlighted in figure 2.11 however, these occur on a less frequent basis.

The in situ repairs are an ongoing process, and there has only been one instance of a repeated failure mode following the restoration of a fibre. Ideally, the maintenance procedures will decrease as these quality control issues are found and full operations begin. The majority of the remaining repairs involve either fibre splicing or prism reattachment, which both require an optical bench to be completed. This exact setup is currently being tested by the ING team and will become a part of the routine maintenance procedures, although the exact schedule has not been determined yet. Based on lifetime testing of the fibre, is was expected that the retractors will last for the 5-year survey period without failing, however these tests do not account for the range of conditions and orientations that the retractors experience. More information

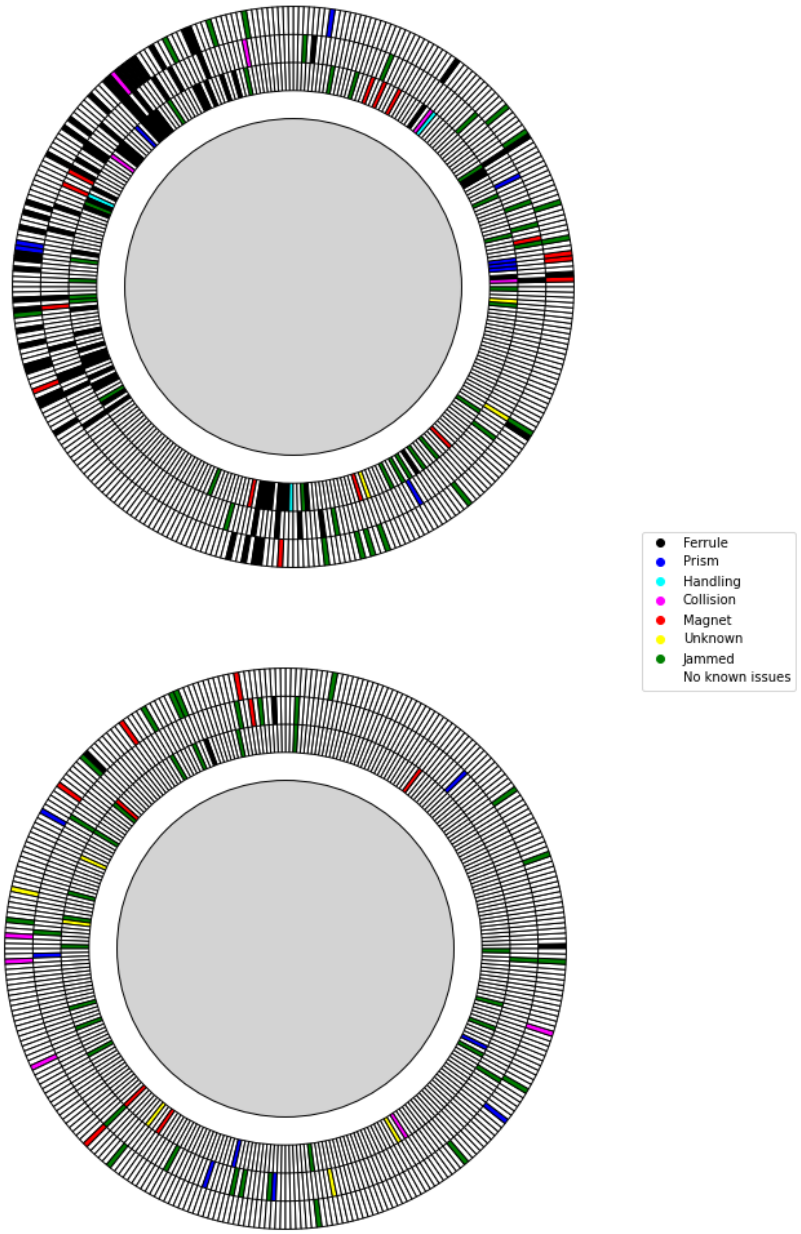


Figure 2.11: The cumulative distribution of disabled fibres for (*top*) MOS A, and (*bottom*) MOS B. These diagrams display the main failure modes since the fibres were first installed and occurred as part of the configuration testing, which is currently ongoing. The majority of these fibres have been repaired and the list of fibres currently disabled is significantly less. Routine maintenance procedures are being scheduled regularly as commissioning continues.

on the lifetime testing of the retractors can be found in Gilbert (2016).

2.4 Slitlets

From the slit to the focal plane, each fibre follows a complex construction. Inside the positioner, the fibre bundles become separate strands with their own button, prism, and pulley sub-system. At the spectrograph interface, the optical fibres are ordered linearly. This optimises the detector's spatial dimension. The fibres are grouped into structures called slitlets. A slitlet is a metal casing that secures the fibre ends on the focal surface through the use of small grooves and adhesives. The ordering of the fibres must be record the ordering of the fibres so that the spectrum on the detector can be matched to the correct observation. This section will describe the assembly and mapping process of the MOS and mIFU slits.



Figure 2.12: The MOS plate B slit, with $\sim$ 75% of all slitlets successfully installed. Index-matching gel has been applied to the surface of the slitlets to prevent any reflection at the boundary between the slitlets and the slit window. The fibres are back-illuminated to identify any air bubbles in the gel across their surface (Hughes et al., 2020).

2.4.1 Assembly

Each slitlet is secured into a single slit structure, which is a curved comb-shaped object with an array of "v" shaped grooves to support the fibres. The slit is curved in both the x and the z directions. The z curvature allows the collimator design to

relax away from a flat input focal surface, simplifying the optics. The x-curvature corrects for distortion arising from the grating for the LR mode, and ensures that all fibres have the same wavelength solution in the LR (i.e. the same lines at the same x-pixel for all fibres). In the HR, this cannot be met, and so the wavelength range changes slightly along the slit. This means that the focus of the same spectral line is curved along the slit.

The three observing modes have independent slits and contain a varying number of slitlets; 40 for both of the MOS with 24 fibres per slitlet, 20 for the mIFU with 37 fibres per slitlet, and 29 for the LIFU with 21 fibres per slitlet. The MOS and mIFU slitlets were assembled by myself and Ian Lewis in the Oxford laboratory, using instructions developed by Ellen Schallig at NOVA for the LIFU assembly. Figure 2.12 demonstrates the arrangement of the slitlets for MOS B.

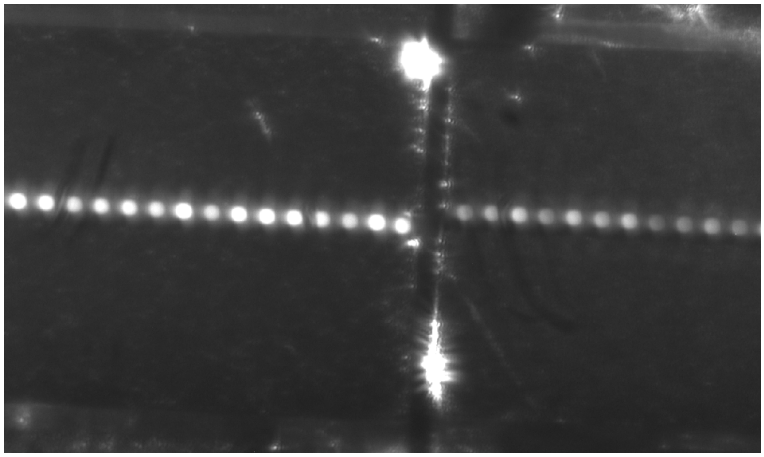


Figure 2.13: An image of the mIFU slitlet taken using a small alignment telescope with a Basssler camera attached. These images were used to identify possible air bubbles in the index-matching gel over the fibres, which could introduce reflection and refraction at the interface to the spectrograph. The small circles represent the back-illuminated fibres along the slitlet.

To prevent reflection and refraction at this interface, the slitlets must be installed using an index-matching gel and sit flush with the front surface of the slit. Failure to do so will introduce air bubbles into the system. After securing a slitlet between the teeth of the slit, a Basler camera was used to image the slitlet's position. Attached to

the camera was a small alignment telescope[6] which could focus to a distance of 600 mm. These images allowed us to identify any air bubbles covering the surface of the fibres, shown in figure 2.13. If air bubbles were present, the slitlet was removed from the slit structure and cleaned, before applying a new layer of gel and re-inserting it into position. It took approximately 2 weeks to complete the assembly of the three slits.

2.4.2 Mapping

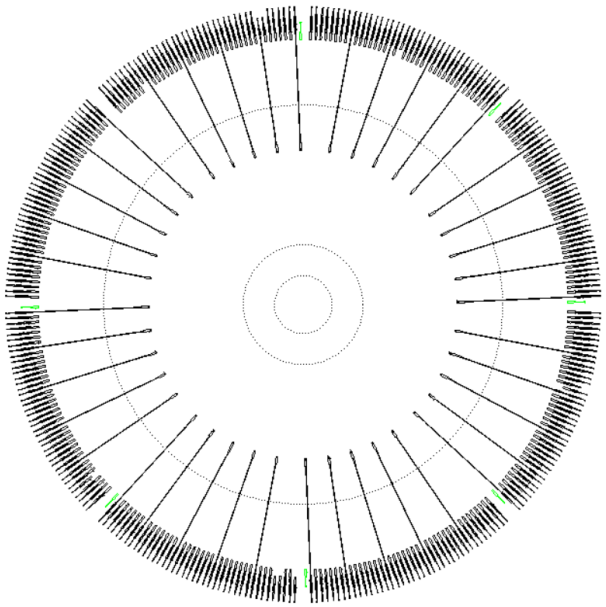


Figure 2.14: A diagram of a single slitmap configuration for MOS A. One fibre from each slitlet has been positioned radially, and exposures will be taken using the WHT's calibration unit to illuminate the deployed fibres. This will identify their position along the spectrograph detector.

There are several labelling conventions within the optical fibre sub-system that were carefully recorded to track their relative position within WEAVE. The first numbers were those used by the manufacturers (SEDI and FTO), which were recorded alongside the location of the fibres within the positioner itself and its new fibre ID. The range of fibre IDs that belong to each slitlet is recorded, but not their relative

[6]This is the same alignment telescope supplied by NOVA and used in the FPI focus adjustment.

position along the spectrograph slit, as the location of each fibre along its slitlet is unknown.

The back-illumination unit is designed to shine a light down all the fibres of a single observing mode simultaneously; either MOS A, MOS B or the mIFUs. To identify where each fibre ID is located along the slit, an alternative method must be used. I designed a sequence of configurations to be deployed, where a single fibre from each slitlet is placed radially on the field plate from its park position. The calibration unit installed on the telescope is then be turned on to illuminate these fibres individually, whilst the parked fibres remain dark[7].

This required a total of 24 configurations per plate to cover every fibre along the slitlet. An example of a single slitmap configuration can be seen in figure 2.14. This information, along with the positioner fibre IDs, has been integrated into the data processing software as part of the commissioning process.

[7]There is some light transmitted by the parked fibres, however it does not affect the identification of the fibres position

Chapter 3

Calibration and Metrology

> Is this data weird or do I just not know anything?
>
> P. Watson

Chapter 3 will discuss the metrology and calibration measurements that were taken prior to the shipping of the positioner to the observatory, and the additional procedures completed after its arrival.

In early 2021, it was quickly realised that the positioner was damaged during transportation. Many small bolts attached to the outer structure had been sheared off, which implied the positioner experienced a considerable shock. This was exacerbated by the movement of the counterweights. The WHT engineering team spent several months carefully disassembling and reassembling the tumbler with the fibres still installed to repair and maintain the positioner. To verify the gravitational stability of the tumbler when mounted to the telescope, flexure measurements were taken across a series of elevations and rotation angles. This led to mechanical interventions and further calibration measurements for each position.

3.1 Zenith calibrations

In the Oxford laboratory, the positioner was built from the interface plate upwards. This orientation can be seen in figure 2.2 and is referred to as the zenith position. It was not possible to change the positioner's orientation until it reached the observatory, due to limited space and the absence of suitable equipment to control its inclination. As a result, the in-lab calibration procedures were developed in the zenith position before being repeated on the testing rig at the observatory.

3.1.1 Tumbler metrology

A crucial requirement for WEAVE is that the field plates are co-focal when observing. The height change across both field plates must not deviate by more than 15 μm along its x and y-axes so that they lie within the focal plane of the prime focus subsystem. This requirement corresponds to the depth of focus at f/2.7 and exceeding this budget will cause fibres near the edge of the field plate to be out of focus during observations.

This calibration process began by measuring the position of the field plates when they are in the observing position. A dial gauge was mounted to a calibration bracket attached to a ledge on the FPI gantry (Schallig, 2019), which can be moved by hand to measure a sparse grid of points across their surface. A diagram of this setup is shown in figure 3.1. The gauge has been zeroed against the calibration bracket. The FPI gantry flexure was measured to be ± 0.02 mm from its mean position across the field, which is significantly smaller than any potential seeing effects. This was measured by Ellen Schallig and is described further in Hughes et al. (2020).

The tilt along the x and y direction of the field plate is calculated using simple trigonometry, where

$$\tan(\theta) = \frac{\delta h}{\delta y}, \tag{3.1}$$

and δh is the change in height along the plate axis δy. The height at the centre of

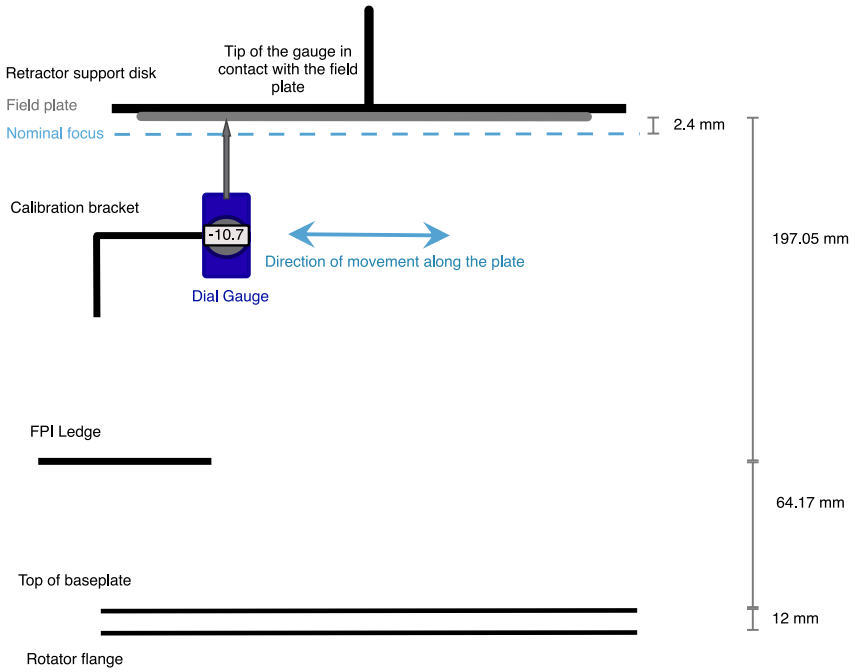


Figure 3.1: Arrangement of the dial gauge with respect to the field plates and the tumbler. The dial gauge is attached to the calibration bracket, which can be moved by hand to specific x, and y-positions along the gantries. Height measurements can then be made along the field plates to identify any tilt that must be corrected.

the field plate gives an estimate of the central tumbler beam's position. Plate A and plate B were found to have a tilt of 500 μm and 200 μm respectively. This forms a cone-like shape about the tumbler's axis which is described in figure 3.2.

However, discrepancies between the centre of the field plate in the observing position and the configuring position implied that the tumbler itself is tilted relative to the axis of rotation. Measurements of the grid spot focus following the retractor installations implied a vertical shift of approximately 1 mm. These changes are due to a shift in height across the central beam as it is bolted to the bearing and the motor along with a downward movement of the structure due to the additional weight of the retractors.

After inserting shims beneath each field plate and lowering the tumbler axis by 80 μm, subsequent dial gauge measurements gave a tilt of 4 μm and 16 μm across the x

and y directions for plate A and a tilt of 11 μm and 30 μm for x and y across plate B, with an average height relative to the interface plate of -9.933 mm and -9.9971 mm respectively. These values are outside the allowed tolerance therefore further adjustment is needed.

It is important to note that in the configuring position, there is no absolute reference point like the calibration bracket. All measurements taken at this side of the positioner are relative to the position of the robots, with gantries that flex as the robots move. This is of little consequence for this set of measurements, so long as both plates are proven to be in tolerance when in the observing position using the calibration bracket. Additional measurements affected by the absence of this absolute reference will be discussed later in this chapter.

The freedom of movement of the bearing and motor bolts within their housing suggests that only a maximum shift of $\pm 0.03^\circ$ is possible, so we assume that the effect of the tilt of the positioning robot plane relative to the FPI gantry is -0.08°. As a result, the central tumbler beam had been shifted by 200 µm from its position prior to the retractor installations and needed to be adjusted.

To compensate for the vertical shift in the tumbler's position, the axis only needed to be lowered by a further 80 µm when including the tilt across each field plate. The arrangement measured is shown in figure 3.2. Lowering the populated tumbler was a difficult process. With the aid of gravity, the structure was moved downward very slowly by loosening and re-tightening the bolts around the tumbler beam in a controlled procedure.

Having adjusted and fine-tuned the position of the central beam, the remaining tilts across both field plates in the observing position must still be corrected. The measurements taken after the vertical shift of the tumbler gave a change in height of 500 μm across plate A and 200 μm across plate B. Shims were inserted beneath each of them, with a total thickness of 700 µm. These shims were not initially distributed

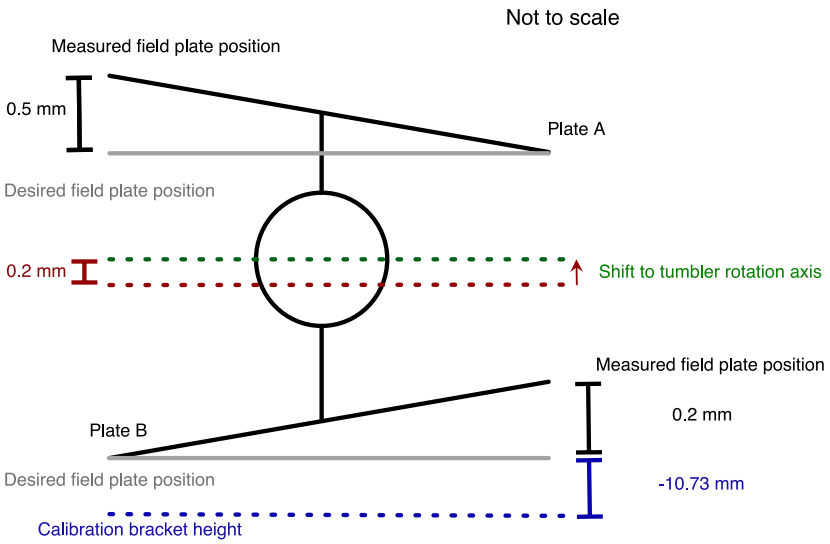


Figure 3.2: A diagram demonstrating the arrangement of the tumbler after the fibre installations were completed. There is an obvious cone structure formed by the two field plates which can only be corrected for using positive shims. The central tumbler beam has also shifted by 200 µm from the previous measurement and must be shifted to ensure the plates are co-focal in the observing position (Hughes et al., 2020).

evenly and are discussed further in section 3.2. However, they significantly reduced the tilt in the observing position, so that the y-axis tilt across plate A and plate B was 15 μm and 10 μm prior to the positioner being shipped. It was expected that these values would produce a tilt of $\sim$ 100 μm across the field plates when in the configuring position.

3.1.2 Field plate metrology

As mentioned previously, it is expected the field plates will exhibit their own independent surface structure due to the effects of surface polishing and flexure. This flexure is introduced when the field plate is bolted down at its edge to the RSD, which is not a flat surface (Schallig, 2019). The final surface profile must be mapped to ensure that each fibre can be placed within the 10 μm error budget, as the tilt and displacement of the field plates relative to the robot gantries make a contribution. Failure to do so

will cause some fibres to be dropped onto the field plate at an inappropriate height.

Following the work completed in section 3.1.1, the field plates now lie parfocal when in the observing position. However, this is not the case when the plates are being configured. As there is no absolute reference point for this side of the positioner, all measurements of the field plates surface profile are relative to the flexure of the robot gentries themselves and cannot be decoupled. It is not possible to complete these measurements with the field plates in the observing position.

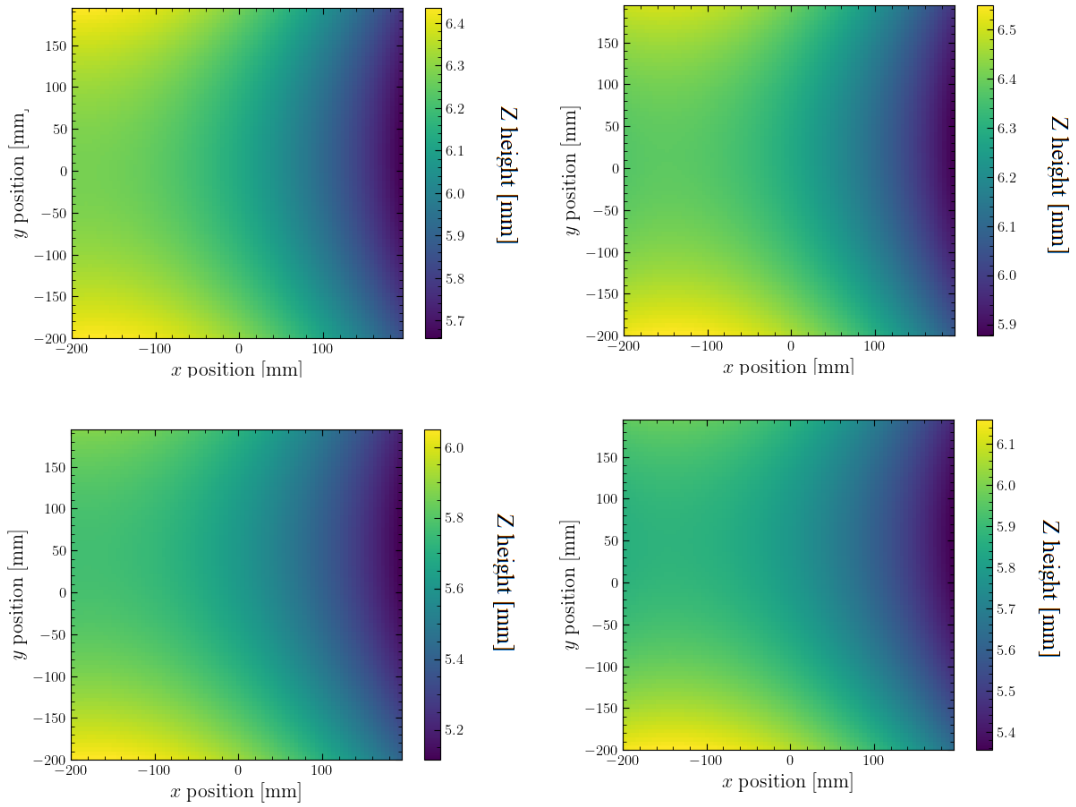


Figure 3.3: A contour map of the z-height measurements for field plate A (*top*) and B (*bottom*) using the robots Nona (*left*) and Morta (*right*). All measurements were made using gauge blocks at the same set of coordinates across the plates (Hughes et al., 2020).

The first set of field plate maps were taken with the positioner in the zenith position, using gauge blocks between the surface of the plate and the closed robot gripper jaws. Gauge blocks are small metal strips with a precisely measured thickness.

They can be stacked together to accurately measure small distances by checking if they fit and were used in this instance to protect both the robot jaws and the plate itself from damage when measuring the distance between them. These measurements were taken at a series of points in robot coordinates across the field plates, whilst the z-axis was kept at a height of 32.5 mm. These maps are shown in figure 3.3. The surface profiles of the field plates are represented as contour maps across the effective field radius in figure 3.3. A linear least squares solution to a second order polynomial was calculated using the set of gauge block measurements. Given a scalar observation (the dependent variable y_i) at each set of coordinates, the data can be described in pairs $(\overrightarrow{x_i}, y_i)$.

To predict the scalar quantity for additional coordinates, a linear function that relates the input with our scalar quantity must be estimated, such that

$$y_i \approx \overrightarrow{w}^\top \overrightarrow{x_i}, \tag{3.2}$$

where $\overrightarrow{w}^\top$ is the weight vector that minimises the squared prediction error. This is defined as

$$\text{Squared error} = \sum_{i=1}^{N} (y_i - \overrightarrow{x_i} \cdot \overrightarrow{w})^2. \tag{3.3}$$

The solution is given by

$$\overrightarrow{w} = \left(X^\top X\right)^{-1} \left(X^\top Y\right), \tag{3.4}$$

where X and Y are the vertical stacks of $\overrightarrow{x_i}$ and y_i for compact notation. This derivation can be found in (Pillow, 2018) and is outside the scope of this thesis.

The solution was used to interpolate across a finer grid of points to predict the plate height at the location of each target over the active radius of the field plate and generate a smooth continuous map. The maps displayed a maximum shift in height of 700 μm for plate A and 800 μm for plate B. However, the use of the gauge blocks does introduce additional errors into the profile of the map, on the order of a few microns.

The choice in height using the gauge blocks is subjective, as the person measuring the profile determines which block represents the maximum height difference. If care is not taken, an upward force can be exerted on the jaws themselves as the block is inserted. To minimise this contribution, I took each set of measurements for the initial maps. Verification and repeated sets were taken by either myself or Prof. Gavin Dalton.

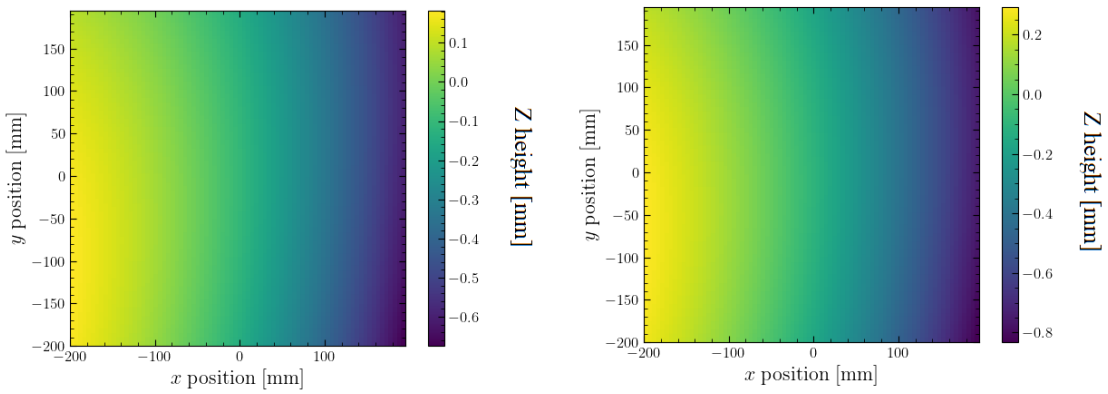


Figure 3.4: The difference in the z-height contour maps between Nona and Morta for field plate A (*left*) and B (*right*).

Although, it is not possible to completely decouple the flexure of the gantries, the differences in their behaviour across each plate can be isolated. Figure 3.4 shows the difference between Nona and Morta in the surface profiles of plate A and plate B independently. Both contour maps of figure 3.4 display the same directional tilt, although they have different magnitudes. These are 70 μm and 100 μm across plate A and B respectively. This change in magnitude is due to the relative tilt between the field plate and both robot axes, which is dependent on the plate being configured.

3.1.3 Retractor tier heights

As the retractors are bolted to both the RSD and the tumbler's top plate, as with the field plates it is expected that there will be a variation in the z height of the

retractor porches as a function of position. These fluctuations must be mapped to reduce errors in the placement of the fibres to and from their park position and to account for the relative tilt between the field plates and the RSD caused by the new shims.

Using gauge blocks to measure the height of each tier was not possible due to the fibre's compact arrangement, as it carried a high risk of damaging the neighbouring fibres. To minimise this risk, only measurements of the tier 3 park heights were taken using the robots directly. With the jaws of the gripper unit closed, the z-axis of the robot can be manually lowered over the park coordinates of a specific tier 3 fibre. The park heights for tiers 1 and 2 could not be measured due to the risk of collision when lowering the z-axis manually.

As the fibres have been installed by hand, their exact position in the system is currently unknown. Whilst they have been assembled to rest safely on their retractor porch, they are not in the optimal position respective to them. The nominal park position of every fibre is when the button vane is radial to the centre of the field and aligned to its pivot position, shown in figure 3.5.

To adjust the position of each fibre, one robot (either Nona or Morta) was slowly driven over the fibre's nominal park position and lowered to the default pick-up height with the jaws open. The initial positioning of the robot over the fibre was shifted in an iterative process, as the gripper jaws would frequently land on the button vane. This was performed by driving the robot to the pick-up xy coordinates and lowering the z axis to a few millimeters above the nominal pick-up height. Its position with respect to the button was inspected visually, as shown in figure 3.5, and adjustments can then be made to the robots xy position before the height is lowered and the jaws surround the button vane. Once over the fibre, an image is taken using the gripper camera and centroiding software developed by (Gilbert, 2016) was used locate the position of the fibre relative to the gripper jaws. A high-level software script, independent to the

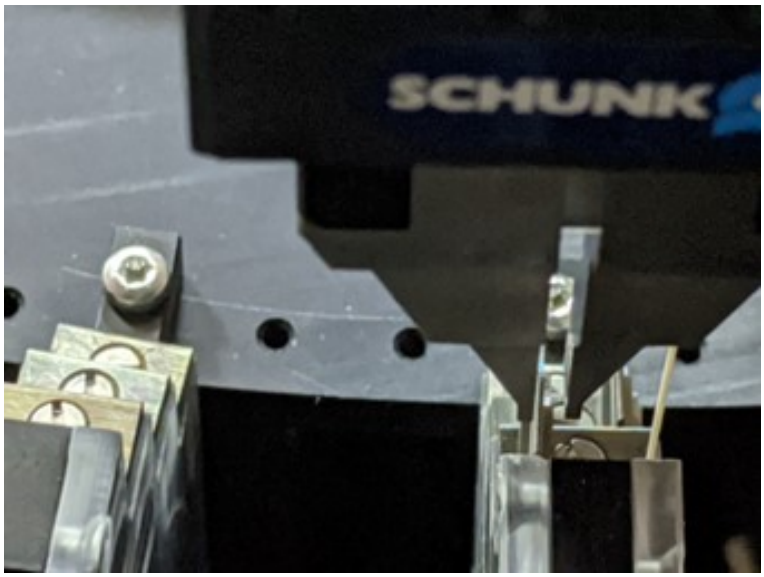


Figure 3.5: An image of the gripper jaws relative to the vane of the fibre button. It highlights the perspective used to determine any adjustments to the robots position when aligning the fibres on their pivot position when resting on the retractor porch. The jaws should be centred on the button to move down to the nominal pick-up height and image the fibre. This is the first step in the procedure for picking up a fibre that has not been moved yet.

configuring software, was then used to carefully tweak the fibre's position with the gripper jaws closed, so that it sat within a few microns of its optimal park coordinates. Following this procedure, the retractor tier height measurements can begin.

The method of measuring the tier heights is completed in a series of stages. First, the two tier 3 fibres resting on the porch must be positioned on the field plate at a $\pm 14^{\circ}$ angle from their pivot point. This ensures that the buttons and fibre peak are clear of the porch. Either Nona or Morta is then driven over the fibre's park coordinates, where the gripper jaws are then closed, and the brakes are switched off so that the z-axis can be gently lowered by hand. The data for the fibre park height values is recorded directly from the z-axis encoders when the jaws are resting on the porch. Measurements were taken at every tier 3 park position and can be seen in figure 3.6.

Whilst this procedure cannot be completed for tier's 1 and 2, the tier 3 measure-

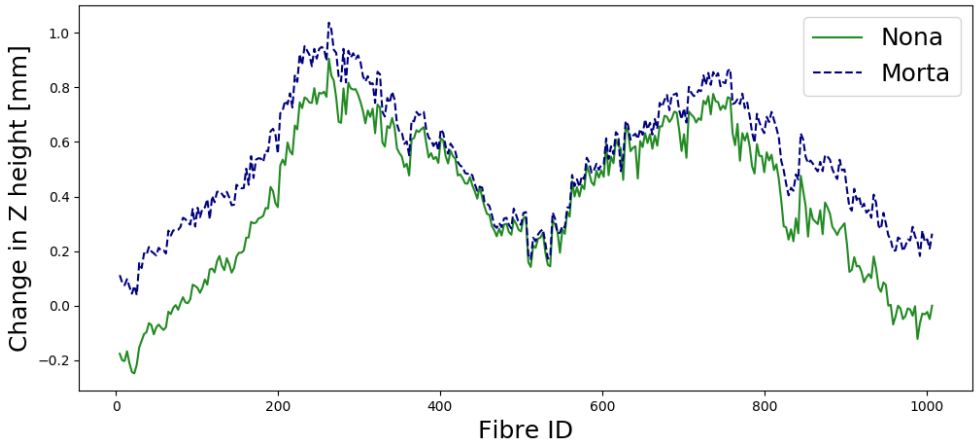


Figure 3.6: The deviation from the nominal z park height for every tier 3 fibre retractor porch on plate A measured using both Nona and Morta (Hughes et al., 2020).

ments can be extrapolated to predict their height. The central frame of each retractor has been manufactured with a vertical distance of 10 mm between subsequent tiers, which is applied to the corresponding tier 3 height measurement. However, this does not account for the variation in height caused by the tilt across the individual porches or in the manufacturing of the retractor frames.

The variation in z-height as a function of fibre ID along plate A can be seen in figure 3.6 for both Nona and Morta. The height of the retractors exhibits a periodic motion for both robots with a notable offset. This cyclic behaviour is due to the differences between the x-stages as the robots move around the edge of the RSD. The offset between Nona and Morta is maximised at the furthest points across the field along the y-direction and due to differences in the rest height of the y-stages across each x-axis.

This behaviour can be accounted for by applying an offset value for each fibre per robot. This value is then applied when calculating the appropriate z-axis height when picking up and placing the fibres from their park position. These offsets have been incorporated into the positioner's software.

From figure 3.6, the porch-to-porch differences are of the order of $\sim$ 10 μm by comparing subsequent measurements. This can be compensated for by over-driving the robots in the z direction when parking the tier 1 and 2 fibres. Extreme cases often correspond to a proud porch screw that secures the porch to the retractor frame causing the gripper unit to partially rest on the screw when it is lowered. This can be rectified by re-adjusting the screw and repeating the measurement.

3.1.4 Fibre imaging and initial placement testing

Orientating the fibres to the correct park position was a lengthy process that took place over several days per plate. This is due to the manual adjustment of the robot's position relative to the button vane as the z axis was lowered to its pick-up height. During the in-lab testing phase, all initial fibre movements, including the park position alignment and the park position alignment were completed using Nona. This was necessary for consistent testing as further calibration was required at this stage for Morta's θ-axis and gripper camera focus.

Each fibre has a local patrol region where it may be positioned on the field plate. This region is a radial segment which emanates from the retractor pivot point with a maximum angle of $\pm 14^{\circ}$ and a radius defined by its tier. Tiers 1, 2, and 3 have a maximum radial extent of 200 mm, 190 mm, and 180 mm respectively[1]. A fibre may be placed anywhere within this region, provided that it does not violate the rules defined in (Terrett et al., 2014). These restrictions protect the fibres from collisions and becoming entangled. The allocation of the fibres to targets in the field of view using the software *Configure* will be discussed in chapter 4.

Figure 3.7 presents a flow diagram outlining the anatomy of a single fibre move, independent of its destination. It should be noted that the study of the fibre movement sequence is currently an ongoing process alongside commissioning, as fields are

[1]These restrictions are subject to change, based on the retractors performance during commissioning and the possible need for additional safety measures.

processed by the positioner, it is crucial to identify unexpected behaviour which could be problematic.

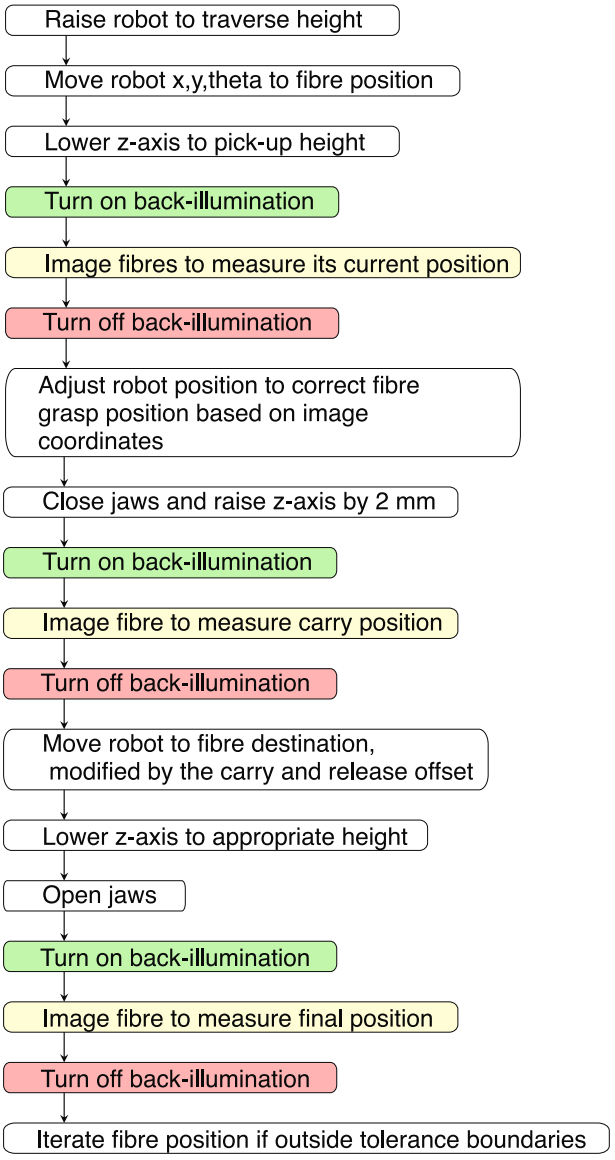


Figure 3.7: A flow diagram describing the anatomy of a fibre move, which summarises the main steps that occur when either Nona or Morta attempts to pick-up and move a fibre to a designated destination.

The initial radial fibre movement tests were conducted using the high-level software to observe and modify subsections of the kinematic movement. Particularly for tier 1 fibres, as their porches are approximately in line with the field plate height.

Originally, the robot z-axis moved independently, whilst the kinematics of the x, y, and θ axes proceeded simultaneously. The kinematics were reconfigured to carry out a synchronised 5-axis movement, as it was found that the independent$z$-axis movement made the largest contribution to the fibre placement timing.

In the case of the tier 1 fibres, a two-stage move was required when travelling from park to prevent collisions between the robot and positioned fibres. The robot moves with a reduced speed to a midway point along its trajectory. When this height above the plate is reached, the tier 1 fibre can then be placed in the nominal manner. This process is reversed when the fibre is being parked from the field and has been tested on all tier 1 fibres. The testing of the movement sequence is ongoing as configurations are conducted as part of commissioning.

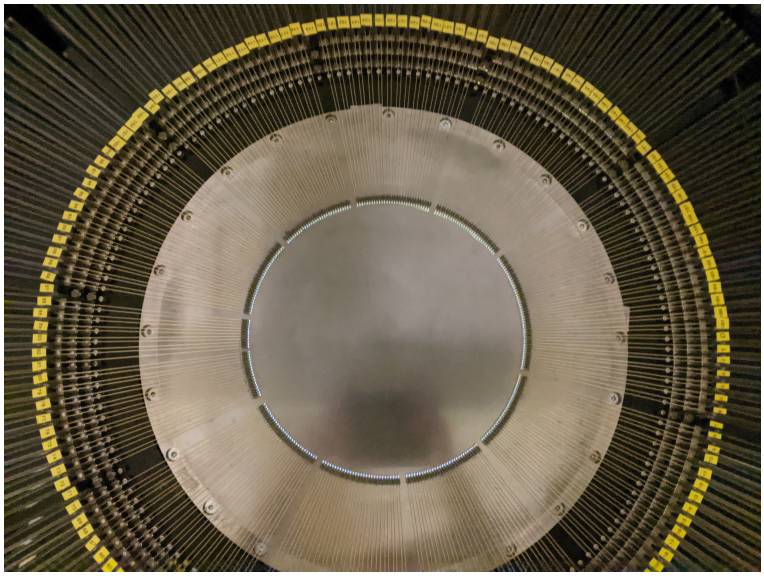


Figure 3.8: The ring configuration used to image the prism of each fibre on plate B and the first test of the final configuration software that will be used during observations. Every fibre in tier 3 has successfully been placed radially on the plate, with gaps corresponding to the location of the mIFUs.

By default, the field configuration software uses a standard outline for each fibre based on the design specifications, with an additional 50 μm buffer around its edge. This outline is used to calculate potential collisions during the target allocation process, but does not account for variations in the alignment between the prism and the

button front-edge due to the assembly process.

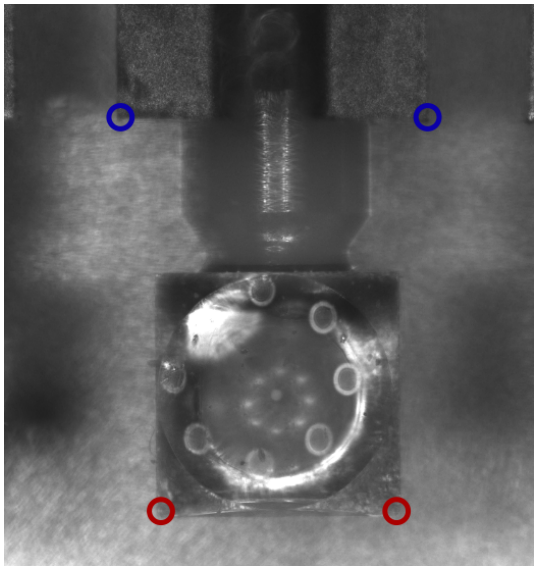


Figure 3.9: An image of the prism attached to an optical fibre, taken with Nona's gripper camera, and represented as a CAD diagram in figure 1.11. This image will be used to define the boundary of the prism within the configuration software, highlighted by the red and blue circles which define the position of the boundary vertices. The structure at the top is the front of the button (*blue circles*), the square piece is the prism (*red circles*), and the section that connects them is the fibre ferrule. The rings in the image around the centroid are reflections of the floodlights, which are LEDs arranged around the edge of the camera aperture.

To update these outlines, a new definition must be measured using direct imaging. Images were taken of each prism using Nona's gripper camera. A single tier of fibres was deployed radially to form a ring configuration to optimise this process, which is shown in figure 3.8. The imaging process is independent of the configuration deployed and must be repeated following the repair of any fibre.

A benefit of this process is that prisms which have been damaged or severely misaligned are easily identified from the images, and can be preemptively disabled before the prisms become detached or before failures occur due to centroiding problems.

To meet the target separation requirements for configurations, the limitations of the optical fibre crossings must be well understood. The previous movement tests have not included any crossings as the fibres were deployed radially or with opposing

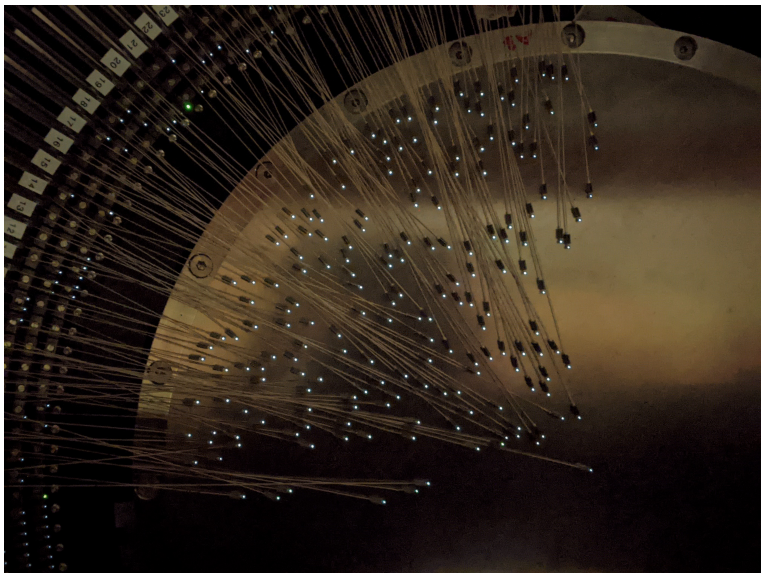


Figure 3.10: The first configuration to test the crossing of optical fibres, initially completed using a list of individual movements, before being repeated by the final configuration software. The field contains one octant of the plate A fibres across all three tiers.

angle pairs. To incorporate crossings when the fibres are deployed, the order they are placed in must be carefully documented. If the robot attempts to move a fibre which is underneath another, it may cause a significant amount of collateral damage to nearby fibres in the field.

The configuration software *Delta* (Terrett et al., 2014) is designed to track this ordering and prevent the robot from causing damage by restricting the possible moves it may make. The performance of *Delta* had been simulated to test its capabilities however, it has not previously been used to configure a field with real fibres. To minimise the risk to the fibres themselves, the fibre movement order was generated by *Delta* but the first fibre crossings were initially executed using the higher level software. After carefully observing the positioner's behaviour for the first few moves, an octant of the field plate, shown in figure 3.10, was configured using the complete operation of *Delta*'s capabilities. The order of the fibres was closely controlled to prevent entanglements and ensure the fibre placement could be easily reversed. Plate A was used during the in-lab testing process, and this test has been successfully

completed on plate B at the observatory.

The tier 1 fibres were positioned from park first, as their crossings would occur at a lower height closer to the field plate edge. The deployment of this octant did not reveal any potential limitations as a result of the fibre crossings. However, full field configurations may reveal additional restrictions that must be implemented.

3.2 Calibration at elevation

In 2021, the Issac Newton Group (ING) engineering team spent several months carefully detaching and reassembling the tumbler structure as a result of obvious damage to the counterweight system caused to the positioner during its shipping, which resulted in several sheared bolts that needed to be extracted and repaired, and the misalignment of the tumbler structure itself. This was identified by a significant defocus when conducting scans of the field plate grid. Figure 3.11 shows a CAD diagram outlining the disassembly process.

Eventually, a decision was made to redesign and install a new robot counterweight system, due to concerns regarding the gravitational load distribution on the counterweight rails themselves. The details of this work are outside the scope of this thesis, however the reformation of the tumbler structure had a significant impact on the positioner metrology.

Following the repair work, the positioner was then aligned and bolted to the WRS system via the interface plate, and both systems were placed on a testing rig that could manually change elevation. Combined with the rotation of the WRS, the orientation of the positioner when mounted to the telescope could now be simulated.

An important orientation of the positioner when it is at an elevation angle of 90°. This is referred to as Access Park 3 (AP3) and is the only position where maintenance procedures for WEAVE can be completed once mounted to the telescope. With

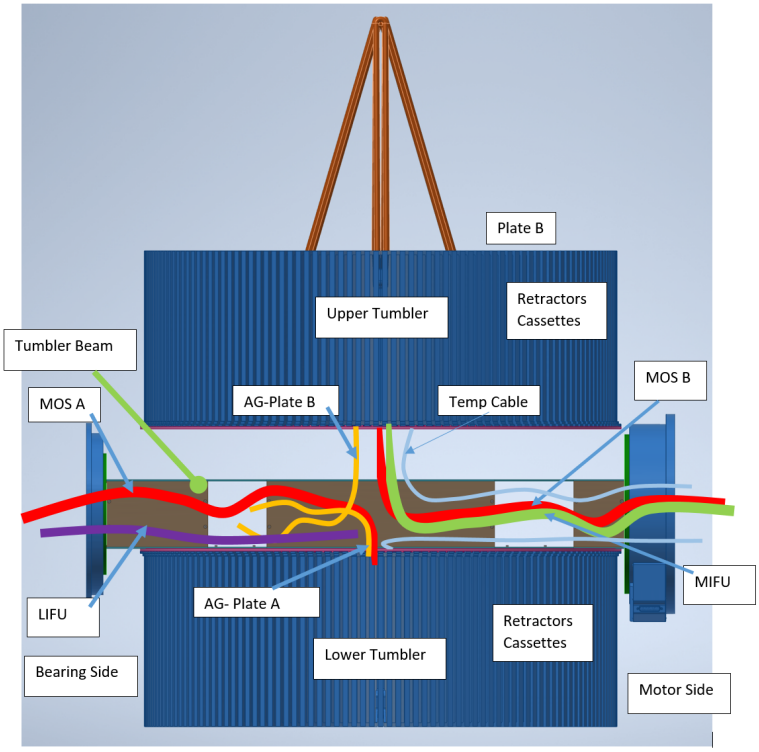


Figure 3.11: A diagram highlighting the method used during the disassembly of the positioner's tumbler in January 2021. Straps were placed around the retractor support disc (RSD) to carefully lift the top-half of the tumbler to allow the ING engineering team access to the central beam.

the testing rig in this position, retractor maintenance was conducted to evaluate the procedure's feasibility. Extensive inspections and interventions were required for many of the retractors on both field plates, as the shock caused a significant disturbance to their internal components.

3.2.1 Flexure

This section describes the tumbler flexure measurements completed by Kevin Dee in 2021, which are summarised in (Hughes et al., 2022a), and the implications this has on the focal plane. Further details on this work are described in appendix B.

The WHT has a limited range in its focus adjustment, designed to account for the effect of thermal expansion and flexure of the top-end structure. However, there are

additional requirements in the focus tolerance, driven by the allowed defocus across the button of each fibre in the focal plane. These are:

- The tilt across the field plate when observing must not exceed $\pm 0.0045°$.
- The back focal distance between the interface plate and the observing field plate must be 273.2 ± 0.025 mm.

To stay within the focus budget of the WHT, the flexure behaviour of the positioner-WRS system must be evaluated. During WEAVE's design phase, a Finite Element Analysis (FEA) model was created to model the flexure resulting from gravitational loading. For the tumbler system, the model consisted of a series of solid bodies for the drums, central beam, RSD, and field plates to predict their combined flexure but not the flexure of the individual components. Direct measurements of the flexure began in 2021 using the elevation testing rig, this setup shown in figure 3.12.



Figure 3.12: (*Left*) The arrangement of a single digital Dial Test Indicato (DTI) when measuring the shift in position of the observing field plate with respect to the WRS interface plate, as the positioner's orientation is changed. It is attached to an angle iron spanning the width of the interface. The tilt across the field plate was measured using two DTIs and the distance to the interface was measured with a depth micrometer, which were zeroed at ZD 0°. (*Right*) The positioner mounted to the WRS, on a testing rig that allows the system to change elevation.

In this section, the elevation angle of the positioner is referred to as its Zenith Distance (ZD), as the testing rig moves from 0°, where the telescope is pointing at zenith, down to 90°, where it is pointing at the horizon. It is important to make this distinction, as during commissioning the telescope position is referred to in terms of the on-sky position angle (PA).

To establish an accurate reference point for future metrology measurements, the position of the field plates relative to the WRS interface plate was determined using a laser tracker. This approach provided an absolute reference point, instead of relying on the FPI calibration ledge as a substitute reference for the WRS interface. The initial measurements were therefore critical for ensuring precision and consistency in subsequent metrology evaluations. These results were obtained using four magnetic Spherically Mounted Retroreflector (SMR) nests, with the interface plate being used as an absolute reference. The nests were attached to the observing field plate along the x and y directions to calculate the tilt with respect to the focal plane at ZD 0°. The new measurements are summarised in table 3.1, and compared to the 273.2 ± 0.025 mm back focal distance requirement. It was clear that the tumblers present position was out of specification.

	Plate A	Plate B
x direction tilt	0.009°	0.046°
y direction tilt	0.024°	0.135°
Movement of the centre position	0.12 mm	2.293 mm

Table 3.1: The results of the tilt and shift in position of each field plate's centre when in the observing orientation, relative to the WRS interface.

To study the tumbler's metrology with elevation, two digital Dial Test Indicator (DTI) gauges were mounted to an iron bar. This was placed between the interface plate and the observing field plate, and spans the width of the rig. The setup is shown in figure 3.12. The DTI's were arranged along the iron bar to place the tips of each gauge at opposing ends of the field plate. A set of measurements were then

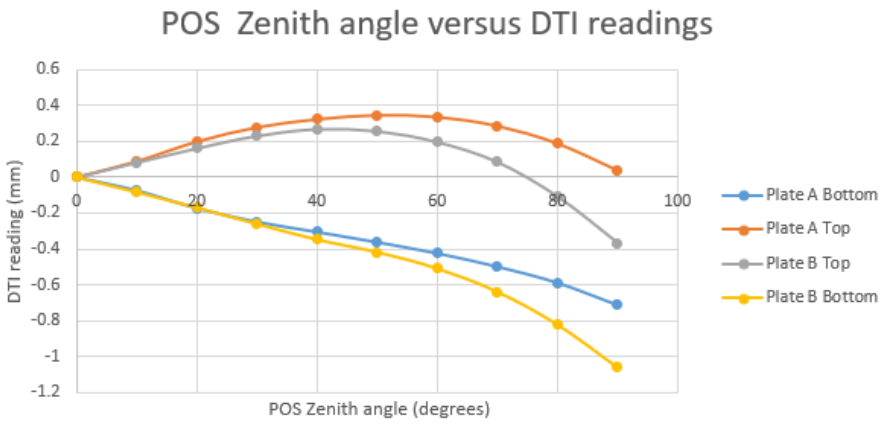


Figure 3.13: The displacement as a function of elevation, measured using two DTI's at both the top and bottom positions on the field plate. The bottom measurements show a continual increase in the displacement, with a maximum value of $\sim$ 1000 μm at ZD 50°. A curved profile can be observed for the top position, where the displacement increases to a maximum value of $\sim$ 400 μm before moving back towards its original value [image credit to Kevin Dee; used with permission].

taken for the x and y axis tilts as a function of elevation angle, shown in figure 3.13. The maximum displacement occurred at ZD 50° with 200 μm and 1000 μm being observed for plate A and B respectively.

The magnitude of these displacements meant that mechanical intervention was required to prevent the tumbler from shifting the focal plane during observations. Additional testing with the laser tracker revealed that the primary source of the flexure was likely caused by the central beam and the attached drums. Aluminium flat plates and reinforcement brackets were inserted into the central beam to improve its stiffness, which may have been reduced when further bolt holes were drilled during reassembly. A clamp was then placed around one end of the beam on the bearing side to prevent it from twisting as the elevation angle changed. These interventions are shown in figure 3.14.

The additions to the central beam's structure significantly reduced the magnitude of the displacement from $\sim$ 1000 μm to 400 μm for plate B, however there is still a distinctive curved profile for plate A which must be removed. This profile was

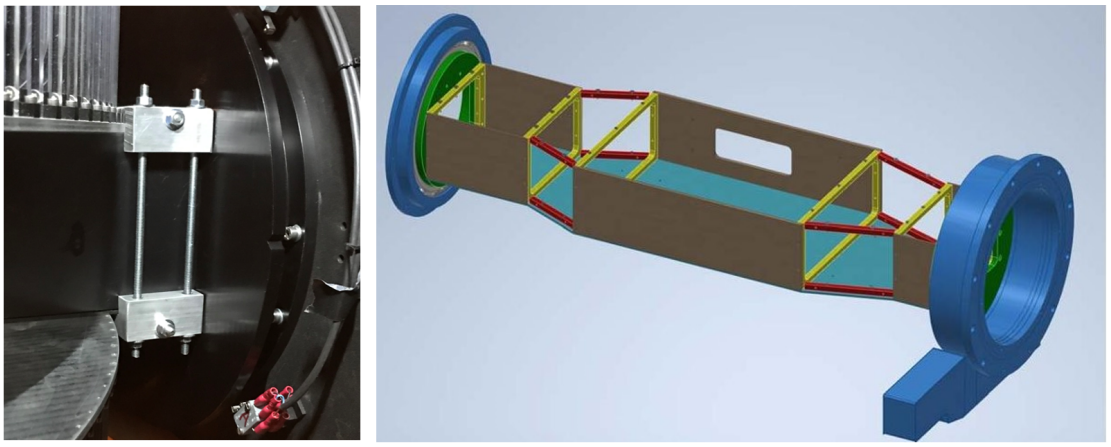


Figure 3.14: (*left*) The clamp inserted around the bearing side of the tumbler beam, to prevent it from twisting as the positioner's orientation was changed. (*right*) A schematic diagram showing the reinforcement bars, which hold in place an aluminium flat along the top of the tumbler beam. These additional structures increase the beams rigidity, which has been weakened by the bolt holes made in the side panels.

caused by the flexing of the tumbler top plates and the RSD relative to one another. Four jacks were inserted between them to reduce the curvature and magnitude of the flexure further and these results are shown in figure 3.15. There is now a maximum displacement of 200 μm across both plates. The jacks were then replaced with a central reinforcement strut as a permanent solution, pictured in figure 3.16, which gave reproducible results that are the same order of magnitude.

Whilst the magnitude of the flexure is outside the tolerances defined by the focal adjustment for the WHT, further analysis of the field plates final focus positions revealed that the flexure results do not impose any restrictions on the orientations which may be used during the WEAVE survey observations. This analysis was performed by Kevin Dee and is described in appendix B. In essence, a well sampled optical system has a point-spread function (PSF) which scales with the apertures f-number^2. However, as WEAVE does not sample the entire focal plane, the tolerance to the defocus is reduced. This tolerance is quantified as the residual sum of squares (RSS) of the defocus, which can then be used to model the resulting PSF as a Moffat function.

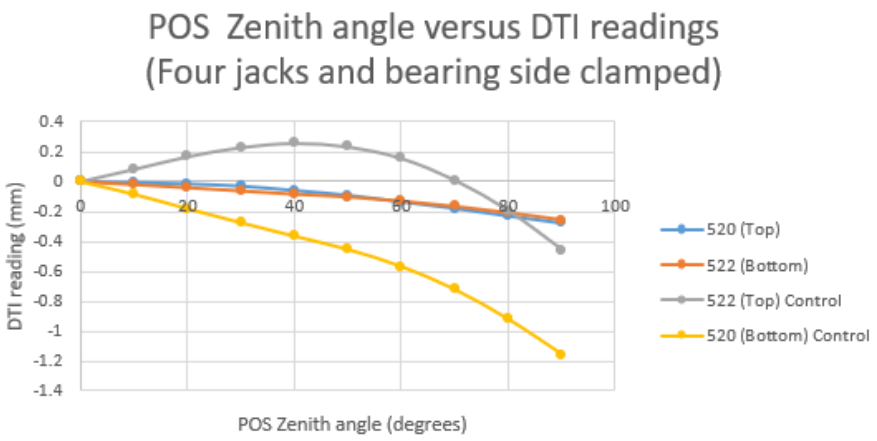


Figure 3.15: Shows the displacement as a function of elevation, when both the clamp, stiffening panels, and four additional jacks are used to hold the tumblers structure. Here the term "Control" refers to the measurements taken prior to any components installed to reduce the flexure. The additional components significantly reduces the magnitude of the displacement, and removes the curved profiles observed for plate A [image credit to Kevin Dee; used with permission].

The least favourable outcome to the remaining flexure is a PSF with a Full-Width Half Maximum (FWHM) of 1.2" at ± 150 μm defocus. For the vast majority of the observing positions, both plates display a tilt within a tolerance of ± 14 μm. A limit of $\pm 155°$ was imposed to prevent the WRS from reaching the limits of its mechanical travel range, which is $\pm 160°$. However, the effect of the flexure behaviour will need careful consideration, particularly as we receive results from commissioning and first-light. At the time of writing, no additional restrictions have been implemented. This is a consequence of the limited FEA, which assumes that the modelled components are rigid solids, when clearly this is not the case.

3.2.2 Robot calibration

After conducting simple operational tests with the positioner at an inclination, it became clear that the calibration procedures conducted in the zenith position were no longer adequate and that the increased gravitational load had introduced unexpected

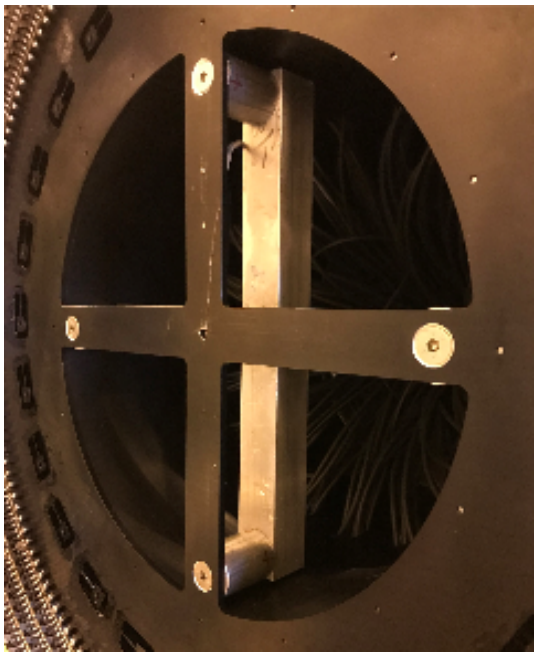


Figure 3.16: An image of the strut installed between the tumblers top-plate and the retractor support disc prior to the field plate reinstallation. It was designed to replace the four jacks used to reduce the flexure of the top-plates with respect to the tumbler drums.

system changes. This was first evident when correcting the fibre's resting position in park using the gripper jaws and during radial fibre movement tests, as the robot often hit the button vane when lowering the z axis. Therefore, additional measurements were required to complete the final calibrations before the positioner could begin observations.

The fibre movements and calibration procedures are all made relative to the fibre's position when viewed by the gripper camera. Therefore, it is vital that the position of the fibre, as measured from its perspective, is accurate and precise. This process began by calculating the new focus height and rotation centre for all three cameras on each robot.

The first step is to set the zero point of the encoders for both Nona and Morta's x and y axes, which has been chosen to be (0,0) on plate A. Each robot must be driven to the centre grid spot on the field plate. To make this process easier, a fibre

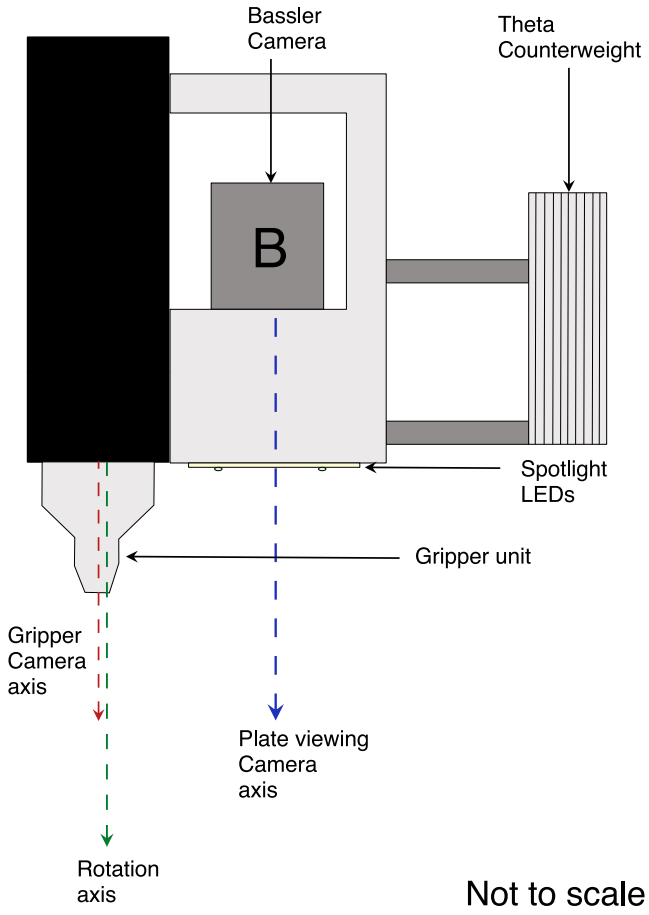


Figure 3.17: A schematic diagram of the two camera focii on each Cartesian robot gantry in relation to its centre of rotation.

can be manually placed over the centre spot before driving the robot to place the fibre within its field of view, as it is only the relative measurements that matter. The centroiding software will calculate the fibre's position relative to the robot, which is then adjusted until the robot is in position and the encoders are re-zeroed.

After setting the encoders, the plate scale and axis of rotation must then be measured for the Bassler camera on both Nona and Morta at its upper and lower focus. A diagram highlighting the position of the axis of rotation is shown in figure 3.17. This process has been continually refined, since the in-lab attempts did not measure this value to a high enough accuracy. The height z-axis of the robot is lowered until the fibre is in focus within the gripper cameras field of view. Setting

the rotation centre must be completed whilst the robot is in the θ orientation used when the robot itself is parked. This is because the calculated rotation centre is not self-consistent when performed at alternative θ angles due to the relative tilt of the robot axes.

To begin the calibration procedure for the robots rotation centre, a fibre must first be placed on the field plate. The robot is then manually driven to a location where the fibre is in the centre of the gripper camera's field of view, or as close as possible to it. An image is taken with the gripper camera focus to measure its position, followed by an image taken with the camera's upper focus to locate the fibre within its field of view. This should be different to that of the gripper-viewing camera focus. The robot's position is adjusted to centre the fibre within the camera's upper focus. A sequence of four moves is then carried out by the robot to image positions around the fibre in a square formation using the upper focus. The position of the fibre is measured whilst the robot is stationary at each location and the difference between them is calculated. This is an in an iterative process, as the distances between each point are used to set the plate scale until the measurements for all four positions agree. After the plate scale is set, the rotation centre is adjusted, in camera pixels, until the measured position of the fibre agrees with the measured position calculated using images taken by the gripper camera focus.

After setting the camera parameters, several scans of the field plate were taken. As mentioned in section 2.1, these grid scans use the centroiding software to remove the offset, rotation, scaling, and non-perpendicularity of the axes when placing the fibres on the field plate (Schallig, 2019). These scans were taken ensure that the measurements of the fibre position are consistent, with all three camera systems agreeing on the location of the fibre.

It is essential that the robot's optical system agrees with the fibre's position for both foci. The gripper camera's focus, for which the entire fibre placement sequence

depends on, must be within ± 1 μm. For the upper focus, this tolerance is increased to ± 5 μm as it only requires the fibre to be somewhere within the its field of view to be a functional part of the positioner's software. It will only be used when searching for misplaced fibres.

Following the positioner's installation to the WHT, a small shift has been observed in the rotation axis for Morta's gripper camera focus on the order of a few pixels when imaging fibres at opposing sides of the field plate. As part of the WEAVE commissioning period, the position of the cameras rotation centre is now being measured as a function of fibre ID and the positioner's orientation to ensure that the calibration does not vary.

3.2.3 Movement modelling

The final results from section 3.2.1 display a linear trend in the tumbler flexure, with a maximum displacement of 200 μm across an angular range of ZD 0° to ZD 90°. This residual flexure will impact the metrology of both field plates and the retractor heights, which must be remeasured following the reassembly of the tumbler. This section describes the updated measurement procedures and models used to estimate the changes in the retractor and field plate metrology with elevation.

To be consistent with the flexure adjustments and the WRS alignment, a new set of base maps was measured at ZD 40°. Due to the restricted access whilst the positioner is on the testing rig, it was no longer possible to use the previous gauge block method to measure the contour maps. Instead, DTIs were used to measure the surface profile of the field plates in the configuring position, by mounting them in the place of the new theta counterweight system. The setup of the DTIs as they measure the surface profile of the field plate is shown in figure 3.18. The gauges were zeroed with the robot driven to the centre of the field plate (0,0) whilst the z-axis was at height of 37 mm.

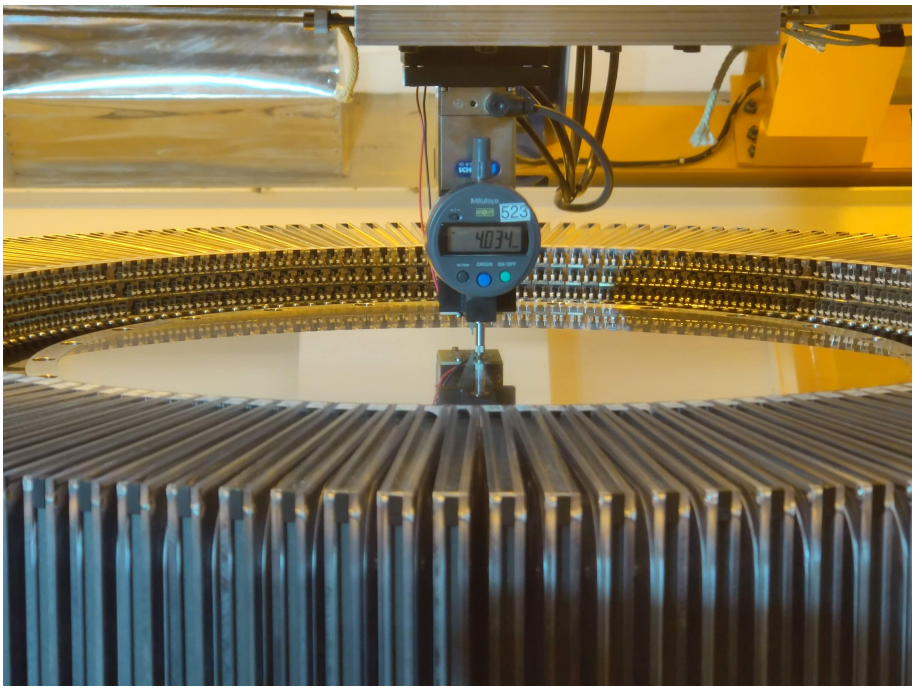


Figure 3.18: The DTI gauge used to measure the z-height of the field plates, for a range of positioner orientations. The gauge has been mounted in place of the theta counterweight system, meaning that every measurement must be made at a fixed orientation. The gauge is zeroed at coordinates (0, 0, 37) mm on the field plate.

The location of the dial gauge introduces an offset from the gripper's position by ±55 mm along the x-axis. Therefore, the robot's range of movement must be restricted to prevent collisions between the robot gripper unit and the retractors, whilst still covering the active radius of the field plates. To produce a complete map across both plates, each robot must make measurements at two different rotation angles. The majority of the coordinates will be measured at a rotation angle of $-90°$, and $+90°$ for the remaining points.

At the centre of the field plate, the rotational offset between ±90° was measured to transform all of the plate height values into the same reference frame. This method was verified by measuring a smaller range of points that could be accessed by each robot at for both theta angles. The offset between each point was consistent with the offset at (0,0) to within 10 μm, and could therefore be applied as a global offset between the two rotational frames (Hughes et al., 2022a). The updated DTI base z-

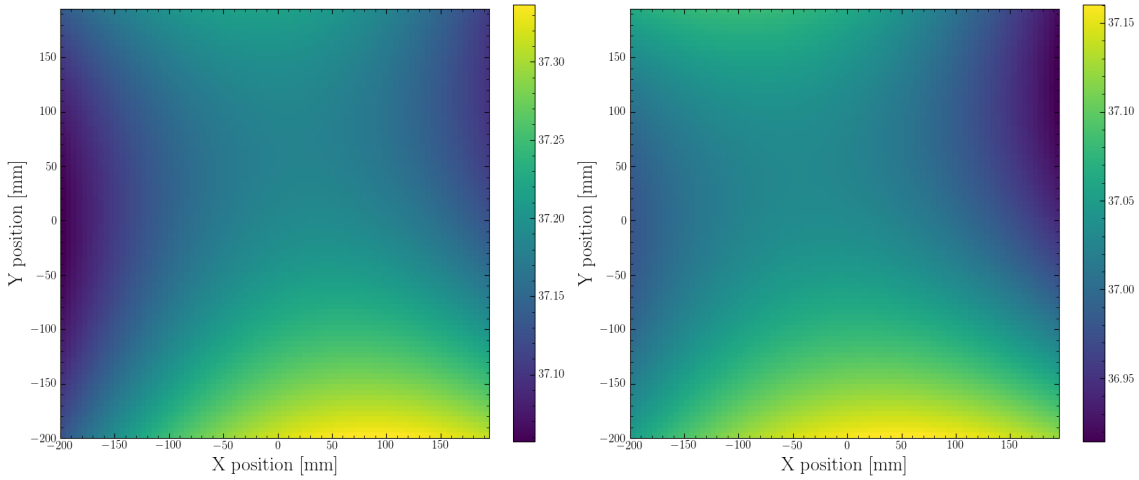


Figure 3.19: The field plate z-height as a function of position for (left) Morta and (right) Nona. All measurements were made using a DTI gauge attached to each robot, for a consistent set of coordinates across the plate. These measurements are relative to the flexure of the robot gantries (Hughes et al., 2022a).

maps are shown in figure 3.19. These contour maps were produced using the previous method described in section 3.1.2 and show a significant change in surface profile of the field plates. This results from the new distribution of shims inserted beneath each field plate, the reinforcement strut, and central beam stiffenings. Across both field plates, the maps display a maximum height difference of $\sim$ 200 μm along the x-direction in the configuration position.

The method of measuring the retractor park height remained unchanged, as it was still possible to manually lower the robot's z-axis with the restricted access on the testing rig. A DTI could not be used at ZD 40° due to the restricted range and additional load to the z-axis. A complete set of updated retractor height measurements can be seen in figure 3.20. The retractors clearly display the same periodic behaviour seen in section 3.1.3, with an increase in magnitude of the peaks. This is expected due to the reinforcement of the tumbler top plate and RSD flexure, in addition to increased gravitational loading of the robot gantries at ZD 40°. It is not expected that the extraction and repair of the retractor systems would introduce a significant changes to this distribution, except between neighbouring retractors measurements.

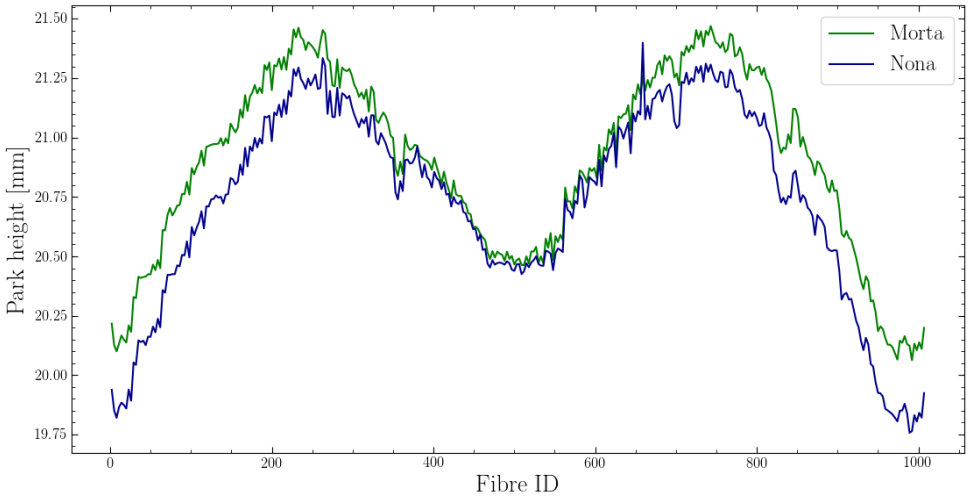


Figure 3.20: The complete set of tier 3 park height measurements for plate A at ZD 40°, for both Nona and Morta (Hughes et al., 2022a).

These measurements were used to update the applied offset when picking and placing the fibres in park.

These base measurements for the field plate profile and retractor height metrology are used as the foundation for predicting the positioners metrology as a function of its orientation. The use of the DTIs allows contour maps to be easily measured for a broad range of positioner elevation and WRS rotation angles. To simulate the orientations which will be used during observations, contour maps were taken across elevation angles of ZD 0°, 15°, 30°, 40°, 50°, and 65°, and rotation angles +120°, +60°, 0°, −60°, and −120°. The differences between these maps and the base map taken at ZD 40° rotation angle 0° can be seen in figure 3.21. As expected, the notable cases coincide with the largest shift in both elevation and rotation angle with a maximum difference of 100 μm across the field plate maps. These differences must be accounted for during the placement of the fibres as observations take place to meet the required positioning tolerances.

This process must also be repeated for the retractor tier heights. Unfortunately, they cannot be easily measured as a function of position, as the access becomes in-

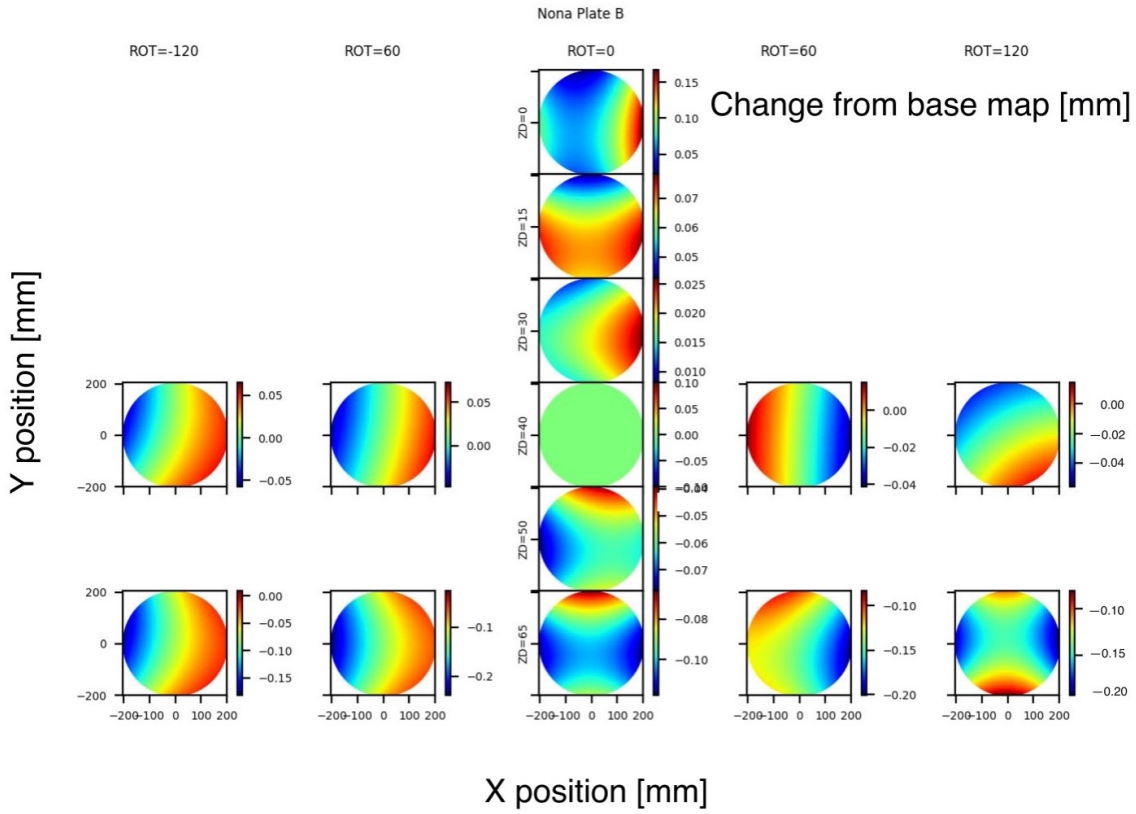


Figure 3.21: The plate B Nona z-height maps, for a set of elevation and rotation angles that span the range of the orientations which will be used on-sky. These plots show the change in the field plate surface profile relative to the base map measured at ZD 40°. The individual plots are the difference between the plate maps at rotation angles +120°, +60°, −60°, and −120° and ZD angles 0°, 15°, 30°, 40°, 50°, and 65°. The colour represents the change in z-height of the plate from the map at ZD 40° and rotation angle 0° in mm (Hughes et al., 2022a).

creasingly restricted with the positioner's elevation angle. Measuring every individual retractor park heights is also a lengthy process which would significantly slow the progression of commissioning if a complete set was taken for each of the corresponding map orientations. Therefore, I devised a model to predict their evolution as a function of the positioner's orientation by extrapolating the available z-maps.

First, I define a function that describes the base plate map $f(x, y)$ at ZD 40° rotation 0°, along with another function to describe the complete set of retractor park heights $M(X, Y)$. The plate map function $f(x, y)$ is then extrapolated beyond the 205 mm active radius of the field plates to 289.75 mm, which gives an estimate of

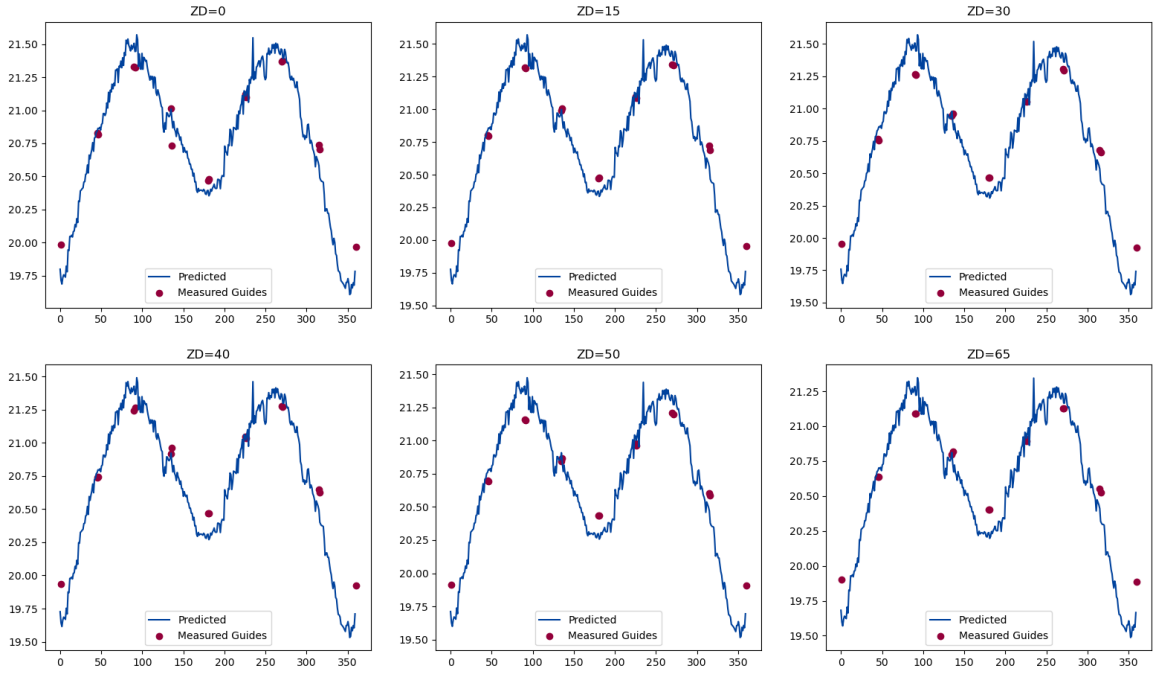


Figure 3.22: Plots of predicted park heights for Nona on plate A, calculated using the plate maps measured at each ZD angle, the offset between the park heights, and the base plate map at ZD 40°. These predictions span the ZD angles 0°,15°,30°,40°,50°, and 65°. The red points represent the tier 3 guide retractor heights, which were measured directly using the robot gripper jaws for comparison.

the plate map value at each park position. The difference between the tier 3 retractor height and the extrapolated plate height is then calculated for every park position and defined as the function

$$\Delta M\left(X,Y\right) = f\left(X,Y\right) - M\left(X,Y\right). \tag{3.5}$$

Once calculated ΔM can then be applied to the the new set of measured plate maps, which have been extrapolated in the same manner as $f\left(x,y\right)$. The result is a complete set of predictive park height measurements $F\left(X,Y\right)$, where

$$F\left(X,Y\right) = f_{\text{new}}\left(X,Y\right) + \Delta M\left(X,Y\right) \tag{3.6}$$

and $f_{\text{new}}\left(X,Y\right)$ is the extrapolated plate map for the each orientation.

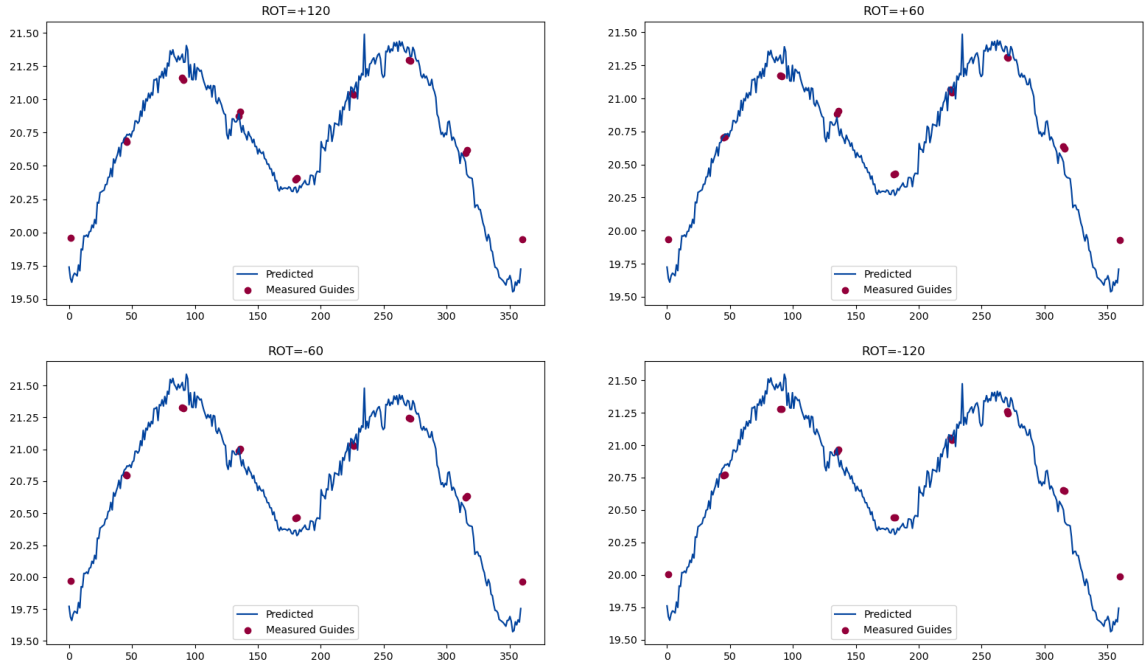


Figure 3.23: Plots of the predicted park heights for Nona on plate A, calculated using the plate maps measured at each rotation angle, the offset between the park heights, and the base plate map at ZD 40° rotation angle 0. These predictions were calculated for rotation angles +120°,+60°, −60°, and −120°. The red points represent the tier 3 guide retractor heights, which were measured directly using the robot gripper jaws for comparison.

When modelling a positioner's metrology as a function of position, the effect of elevation and rotation angle are evaluated independently. Their predictions are then applied as a series of corrections to the plate maps and park heights in the positioner's software as the fibres are configured.

To verify our model predictions, direct park height measurements were taken across every guide fibre retractor for the complete set of orientations using the robot gripper jaws. Both plate A and B have eight guide fibre retractors which are evenly spaced along the perimeter of the field plate. The guide fibre retractors on plate A do not contain any additional MOS fibres, meaning that only tier 1 is populated. This rapidly improves the time taken between measurements, as there are no fibres to be moved. This is not the case for plate B, where the guide fibre retractors contain MOS fibres to compensate for the space taken by the mIFUs. However, only sixteen MOS

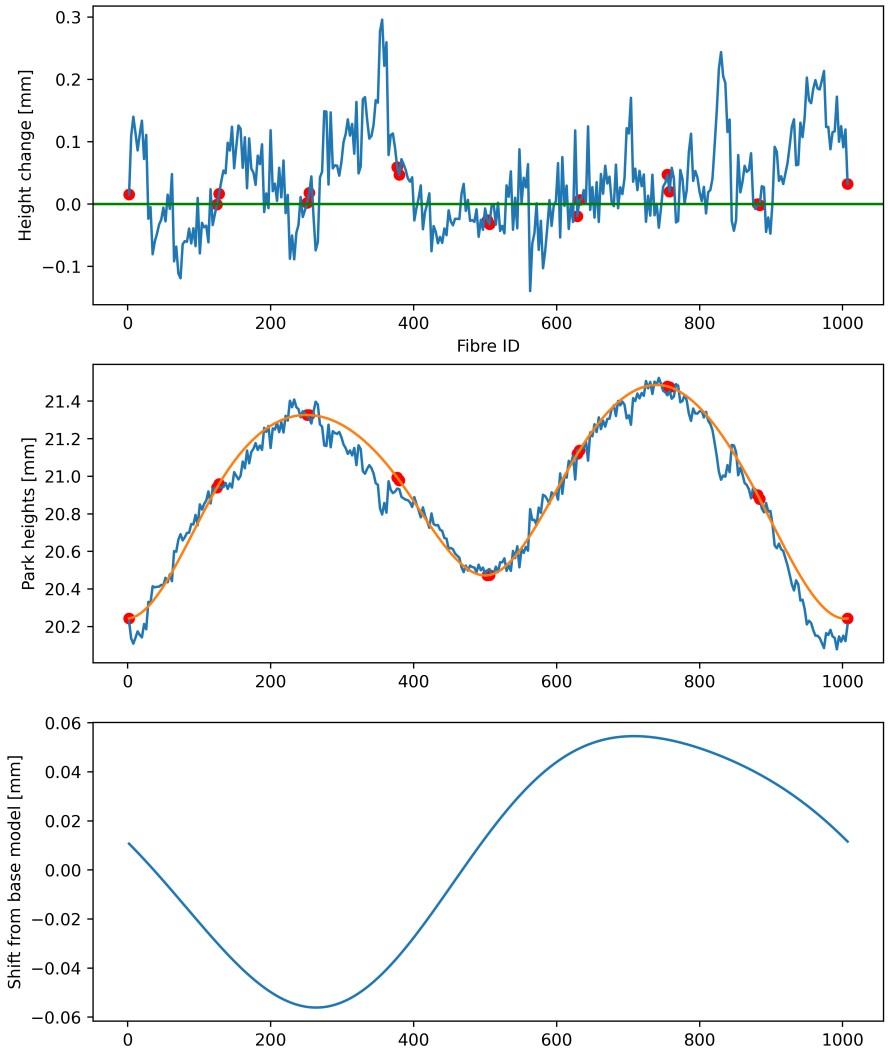


Figure 3.24: (*Middle*) A plot of the predicted park heights at ZD 40° and rotation +60° represented by the blue line, with the measured guide retractor heights displayed as red points. The orange line represents a simple spline fit to the guide retractor heights over the range of the predicted heights for a smooth comparison. (*Top*) A plot of the residuals of the spline fit and the predicted park heights, shown by the blue line. Here the red points represent the residuals between the measured guide retractor heights and the spline fit with the zero point emphasised by the green line. (*Bottom*) The difference between the predictive map at ZD 40° rotation +60° and the base model prediction generated at ZD 40° rotation 0°. This plot depicts the change in the extrapolated plate map between positioner orientations.

fibres must be deployed on plate B and can be done in a single configuration as there are no fibre crossings. Figure 3.22 shows the complete set of predictive park heights for each elevation angle, alongside the direct measurements of the guide retractor

park heights, marked by the red points. This is repeated for the set of rotation angles in figure 3.23.

Comparing the models prediction at the guide fibre park positions against the direct measurements demonstrates that there is an uncertainty on the scale of ~ 40 μm. A plot of the residuals for the guide retractors measurements can be seen in figure 3.24. Initially, a spline function was fitted to these residuals as their behaviour appeared to be periodic however, its behaviour is not physically motivated and applying these residuals as an additional offset may introduce greater uncertainties rather than reducing their magnitude. Instead, the spline fit was applied to the predicted park heights including the measured guide retractors to describe the distribution as a smooth continuous function without the variation between retractor pairs. This was used as a baseline to redefine the residuals and analyse their distribution. This is shown in figure 3.24, and includes the transformation of the predictive model as a result of rotation.

It appears that the majority of the residuals for the direct measurements and the predictive map are at a greater height than the spline model, which will not cause the fibres to be parked incorrectly. It is also not clear how representative the guide retractor measurements are of the complete set of residuals. Taking this into consideration, along with the magnitude of the residuals, it was decided that the models will be implemented within the positioner's software. In addition, the robots will be over-driven by ~ 100 μm when placing the fibres back in park to account for the ~ 40 μm uncertainty. This overdriving should not exert excessive force on the gripper jaws which might cause them to slip down the button vane.

3.3 Fibre placement offsets

There are three types of fibre movements that the positioner carries out; from the park position to the field plate, from one position on the field plate to another, and from the field plate back to the fibre's park position. These movements are referred to as *Unpark*, *Move*, and *Park* respectively. The positioner's software is designed to image the fibre at several stages along its trajectory, including its final placement position. If the fibre is not placed within ± 10 μm of its target, the fibre is picked back up and placed back down using an offset calculated by the positioner software (Terrett et al., 2014). This offset is based on its previous placement and its current position relative to the robot gripper jaws as it is being held. This is an iterative process designed to place each fibre within the positioning tolerance.

Ultimately, the primary goal of the positioner is to precisely place each optical fibre in the field of view to a tolerance of ± 10 μm on its first attempt, in a short enough time frame to allow the entire field to be successfully configured within one hour. This is no small feat. Since the first prototype fibres were installed in 2019, extensive testing has taken place to achieve this requirement.

The ability to place a fibre within tolerance on the first iteration is dictated by the offsets which are introduced during each stage of the fibre's movement cycle. They can be described in three main components; the grasp offset, the release offset, and the carry offset. It should be noted that a unique grasp and release offsets is defined based on the move being performed and is unique to every individual fibre, for example the park release offset and the unpark release offset. The aim of testing the fibre movement process is to define a set of offsets for each individual fibre that can be applied based on the movement being performed. However, this is a chicken-and-egg calibration problem, as the offset values are not mutually exclusive and there is no absolute reference which may be used.

The grasp offset is defined as the change in position of the fibre once the gripper

jaws have closed around the button vane. Specifically, this is the shift in the centroids position, relative to the button vane. Manufacturing defects can effect this value, due to the varying shape along the vane's length, the relative alignment of the prism to the front end of the button, or the z height the robot picks the fibre up from. The carry offset is defined as the change in the fibre's position as it is carried by the gripper jaws whilst the robot traverses across the positioner's focal plane.

The park offset is the measured shift in position once the fibre has been released from the gripper jaws over its park position. This value is fed back into the software to adjust the positioning of the robot over the porch when it is about to perform an *Unpark* move. The same principle applies for the unpark offset, except that it is the shift in the measured position after moving the fibre from its park position to the field plate via the *Unpark* movement sequence.

The plate release offset, applied during the *Move* sequence, is the change in the position of the fibre when it has been carried from one position on the field plate to another target within its patrol region. These offsets must be measured and stored for each individual fibre, as the relative tilt angle of the plate to the robots, and each retractor porch to the robot is different. The magnitude of these angles is very small, however they have a significant impact on the shift in the fibre's coordinates.

Following the positioner's reassembly and prior to its mounting to the WRS, remote testing of the fibre movement sequence began in June 2021 due to travel restrictions as a result of the COVID-19 pandemic. Four web cameras were attached to the positioner on each side of its main structure to provide a live-feed of the robots kinematics and placement of the fibres from multiple points of view. To observe the state of the positioner, a floodlight system was arranged in an exterior frame which could be switched on and off remotely. This must be switched off when fibres are configured, as the additional light interferes with the centroiding software.

These calibrations were all completed in the zenith position. In addition to the

series of offset values needed per fibre, it is essential to identify any dependence in the offsets direction and magnitude as a function of position on the field plate. In principle, this should be accounted for through the map of the field plates surface profile, which has already been applied as a height offset in the fibre's placement.

Each fibre was initially placed radially on the field plate, before being moved to a sequence of 50 random targets within its patrol region. This varies by fibre tier due to the limits applied to their withdrawal radius. At the time of this exercise, the offset measured was the relative rotation of the fibre's position once the gripper jaws had been opened. The definition used to apply an appropriate offset to each fibre per move is outlined later in this section. This set of fibre exercises demonstrated that the offset's magnitude and direction are not correlated with the fibre's location across the field plate but are dominated by the relative tilt between the robot gantries and the field plate. For plate A, the magnitude of the offsets remains constant across the field plate. However, a consistent angular offset is observed for each vector, with a $\sim 180^\circ$ flip between Nona and Morta due to the mirrored setup of the robots. Plate B exhibited more complex behavior. When measuring the offsets using Nona, this tilt is relatively small as the direction of the offsets appears to be random. For Morta, this tilt is significantly larger as the offsets exhibit a distinct and consistent direction across the entire field plate. However, as mentioned in section 3.2.2 there appears to be a rotational offset which varies with Nona's position along the field plate. Further testing is required to understand this affect. This can be performed as a daytime activity during the ongoing commissioning process, where fibres at opposite sides of the field plate are positioned radially in a sequence of 500 moves using both Morta and Nona.

Unfortunately, the testing which began in June 2021 was eventually stopped due to numerous retractors stalling during the offset measurement sequence. This led to the intensive inspection and maintenance of entire optical fibre subsystem described

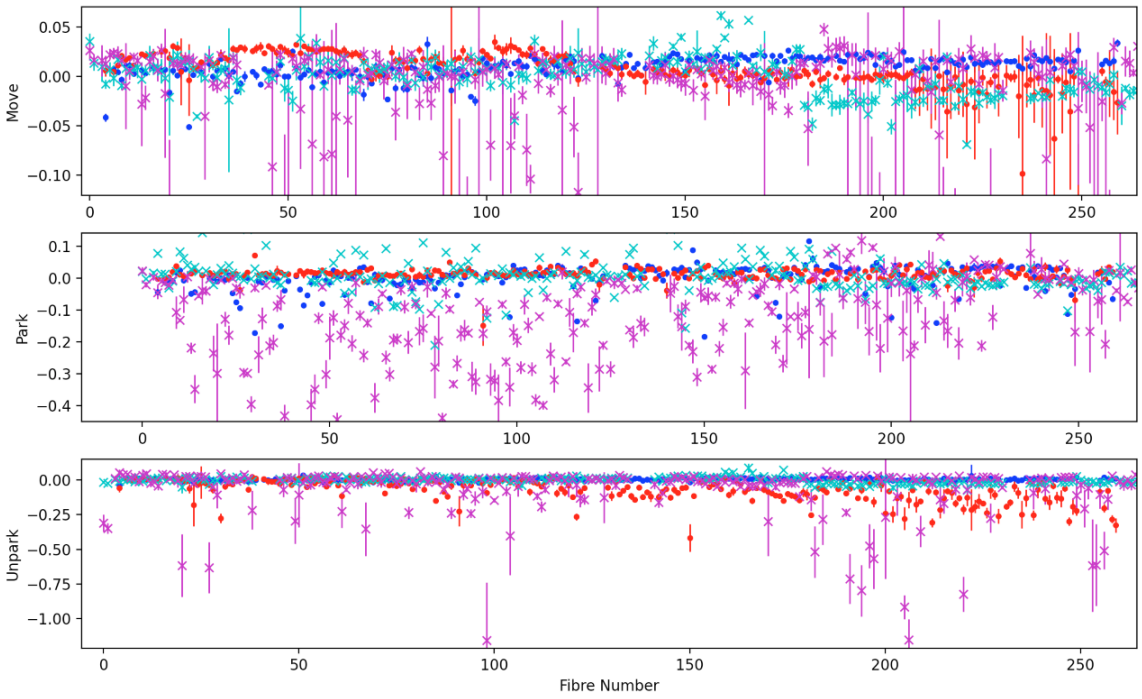


Figure 3.25: The distributions of the *Move* (*top*), *Park* (*Middle*), and *Unpark* (*bottom*) relative rotational release offsets for 250 fibres on plate A. Here, the cyan and magenta markers represent the x and y direction offset when a gripper force of 30 N was used. The blue and red markers represent the x and y direction offsets after this force was increased to 50 N.

in section 2.3.2. The reason for these failure modes is unclear, as there was no indication that they were linked to any systematic failures or isolated to regions around the field plates. They appeared to occur randomly. The most likely cause is a combination of quality control issues and the impact of the tumbler reassembly on the individual retractor components. The offset measurement procedure was paused until the positioner was successfully mounted to the telescope in July 2022, following a period of extensive maintenance.

Direct observations of the fibre movement sequence showed that several fibres were not securely held by the gripper jaws, resulting in a drooping of the button's back-end. The force applied by the gripper jaws was increased from 30 N to 50 N[2] to

[2]The maximum limit is 60 N, but 50 N was chosen to reduce the strain on the system.

improve the stability of the fibre as it is carried, and figure 3.25 displays the impact this had on the relative rotational offset distributions for 250 fibres for plate A. The cyan and magenta points represent the x and y offsets with a gripper force of 30 N, and the red and blue points represent the x and y offsets after the gripper force was increased to 50 N. A greater gripper force appears to reduce the dispersion in the y direction offset for all three movement types. This is expected, as the y-direction in gripper coordinates follows the length of the button vane, therefore by preventing the back-end from drooping the magnitude of the release offset is significantly reduced.

It eventually became that clear through these movement tests that a robust and efficient method of storing the offset values was required to conduct further analysis and apply the appropriate corrections to each fibre. I developed the positioner's fibre calibration database, which searches through the positioners logfiles to store the relevant information from the movement sequences of every fibre. The database stores the past 500 entries per fibre for each type of move, which includes information on the timings and the orientation of the telescope at the time of the move.

The relative rotational offsets calculated from the previous movement exercises had proven ineffective at reducing the number of iterations attempts the positioner takes to place a fibre, therefore a new definition is required. To position each fibre within tolerance on the first try, an offset value was calculated from the database information and applied to each movement type independently. The applied offset is defined as

$$\text{Move offset} = \text{Expected offset} - \text{Fibre position}, \tag{3.7}$$

where the "Fibre position" is the position of the fibre in mm relative to the centre of the gripper, in the grippers coordinate frame, once the jaws have been opened and the fibre has been placed. The "Expected offset" is the refined change to the fibre's position based on the current release offset and the new measured position of the fibre. The average offset calculated from the database is used to adjust the robots

positioning during the movement sequence. The entries stored in the database are continually updated as additional movement testing takes place and the fibres are configured for commissioning observations. It is not expected that these offsets will remain constant over the fibre's lifetime, however extreme values must be filtered out to avoid skewing the average and increasing the positioning error of the fibre.

The left-hand panel of figure 3.26 shows the x and y offset distributions for all three movement types for fibre 192 for plate A using Nona. This fibre has 500 entries in the database as a result of a repeated radial movement testing. This test aimed to observe the dispersion of the offset distribution for a single fibre. The red line represents the average offset value (which is currently being implemented) and the blue line represents a rolling average along the offset distribution, using the last 50 offset values. Figure 3.26 explores the use of sigma clipping as a method of excluding extreme offset values without removing the underlying trends.

Iterative sigma clipping is a data reduction technique designed to remove statistical outliers from a dataset in a robust and repetitive manner, using the assumption that outliers lie some multiple of standard deviations from the mean of the distribution. The aim of iterative sigma clipping is to remove these outliers whilst preserving the underlying distribution of the data. In this case, the outliers are predominantly due to instrumental effects, such as changes in temperature or manual interventions. A threshold is set in multiples of standard deviation, and data points that lie further from the mean than the threshold value are discarded. An updated mean and standard deviation is then calculated from the remaining data, and the process repeats until the dataset converges with no additional points being discarded.

The middle panel of figure 3.26 displays the offset distribution after values which lie $< 3\sigma$ from the mean have been excluded. This process was set to iterate up to five times or until the offset values had converged and no further data points were excluded. This method appears to be effective at removing the extreme values

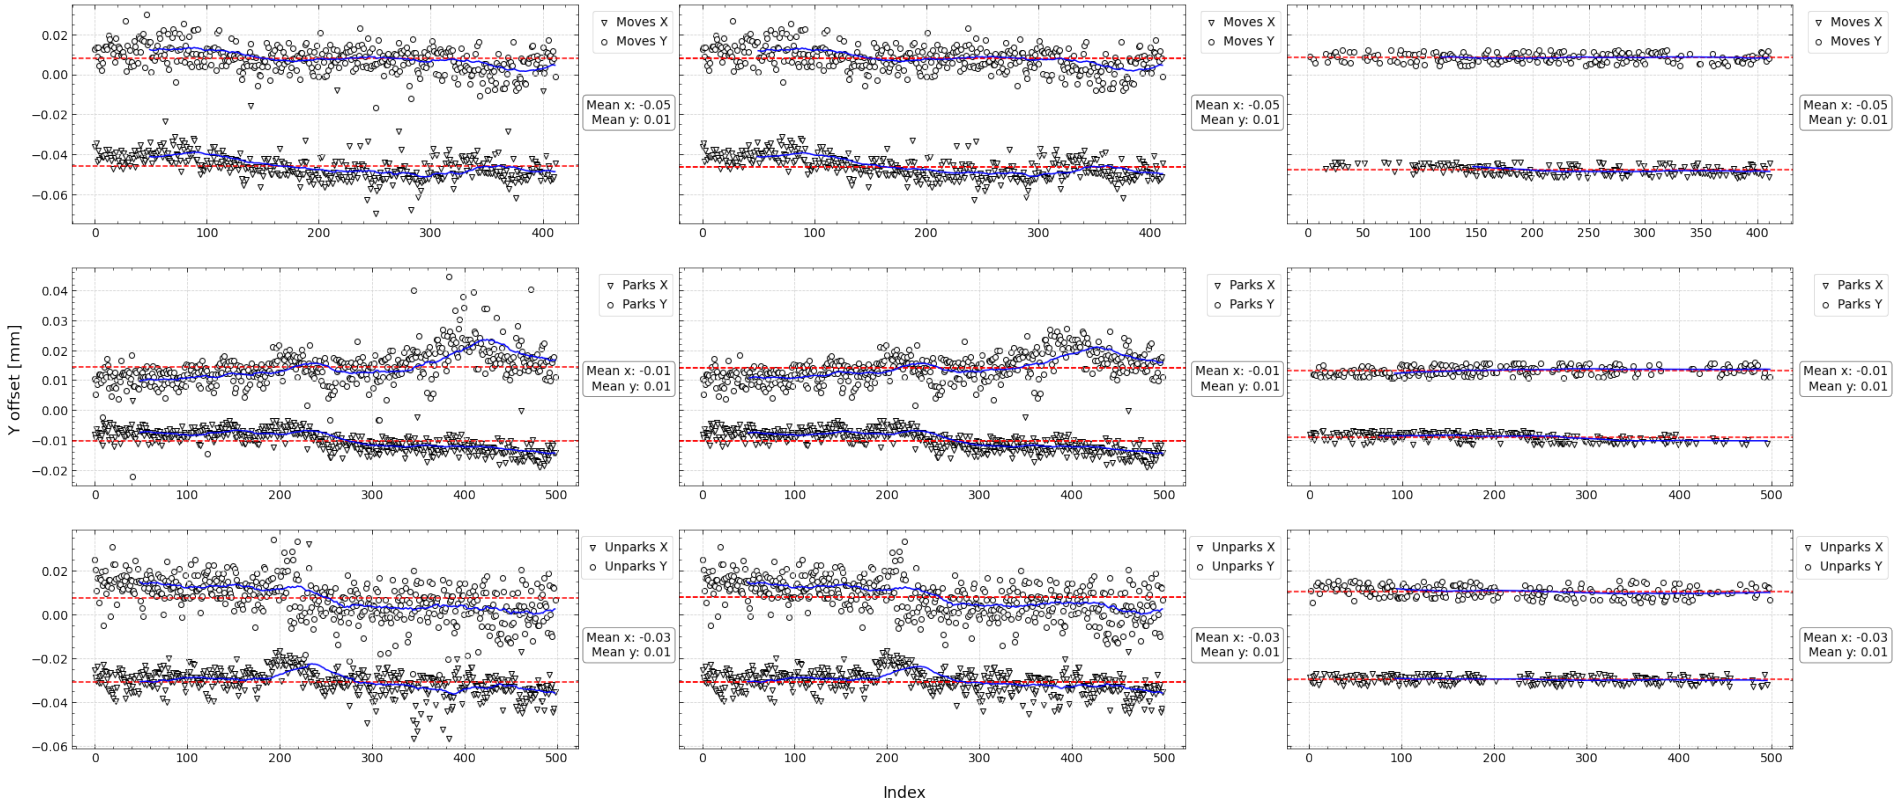


Figure 3.26: (*Left*) The *Move* (*top*), *Park* (*middle*), and *Unpark* (*bottom*) offset values for fibre 192 that have been recorded in the calibration database for plate A using Nona. (*Centre*) The offset values for fibre 192 with sigma clipping applied, where values 3σ away from the mean have been excluded. This has been conducted over 5 iterations or until the offset values have converged. (*Right*) The offset values for fibre 192 with an additional round of sigma clipping. With the extreme values excluded, sigma clipping has been applied by excluding values outside 1.5σ

however, the distribution still displays small peaks. These peaks are likely the result of variations in temperature of the gripper jaws, as these offsets were taken over several days and included unfavourable winter weather conditions.

Initially, the database was designed to store the last 10 moves a fibre had made, however this was increased in reduce the risk of outliers skewing the average offset. However, from the distributions shown in figure 3.25 it is clear that additional filtering is required. This is achieved through iterative applications of sigma clipping.

The right-hand panel of figure 3.26 shows the offset distribution after a second round of sigma clipping has been applied. With the extreme values excluded, a constrained sigma threshold of 1.5σ is used to remove the instrumental effects[3], limited to two iterations to prevent large quantities of data from being excluded. It appears to be very effective at constraining the dispersion of the distribution however, a significant number of points have been removed by this process. Across all three panels movement types in figure 3.26, the sigma clipping has had little impact on the average offset value applied in the movement sequence. However, fibre 192 is fully populated within the database. Due to commissioning timing constraints, it is not possible to populate the database with 500 entries for every fibre.

Figure 3.27 displays the equivalent processing steps on fibre 39, which has 20 entries recorded across all three movement types. The left-hand side of the figure displays the unprocessed (*top*) *Moves*, (*middle*) *Parks*, and (*bottom*) *Unparks* entries, before two rounds of sigma clipping have been applied, as shown by the right-hand side. This method appears to be very effective at removing extreme points at the cost of reducing the number of entries. As expected, the sigma clipping has a significant impact on the mean offset calculated for fibres with fewer entries, compared to fibres that have been exercised extensively. Based on these distributions, the option to include sigma clipping has been implemented as part of the offset calculation, where a

[3]For example, the small bump in the offset distribution has been traced to a drop in the gripper temperature.

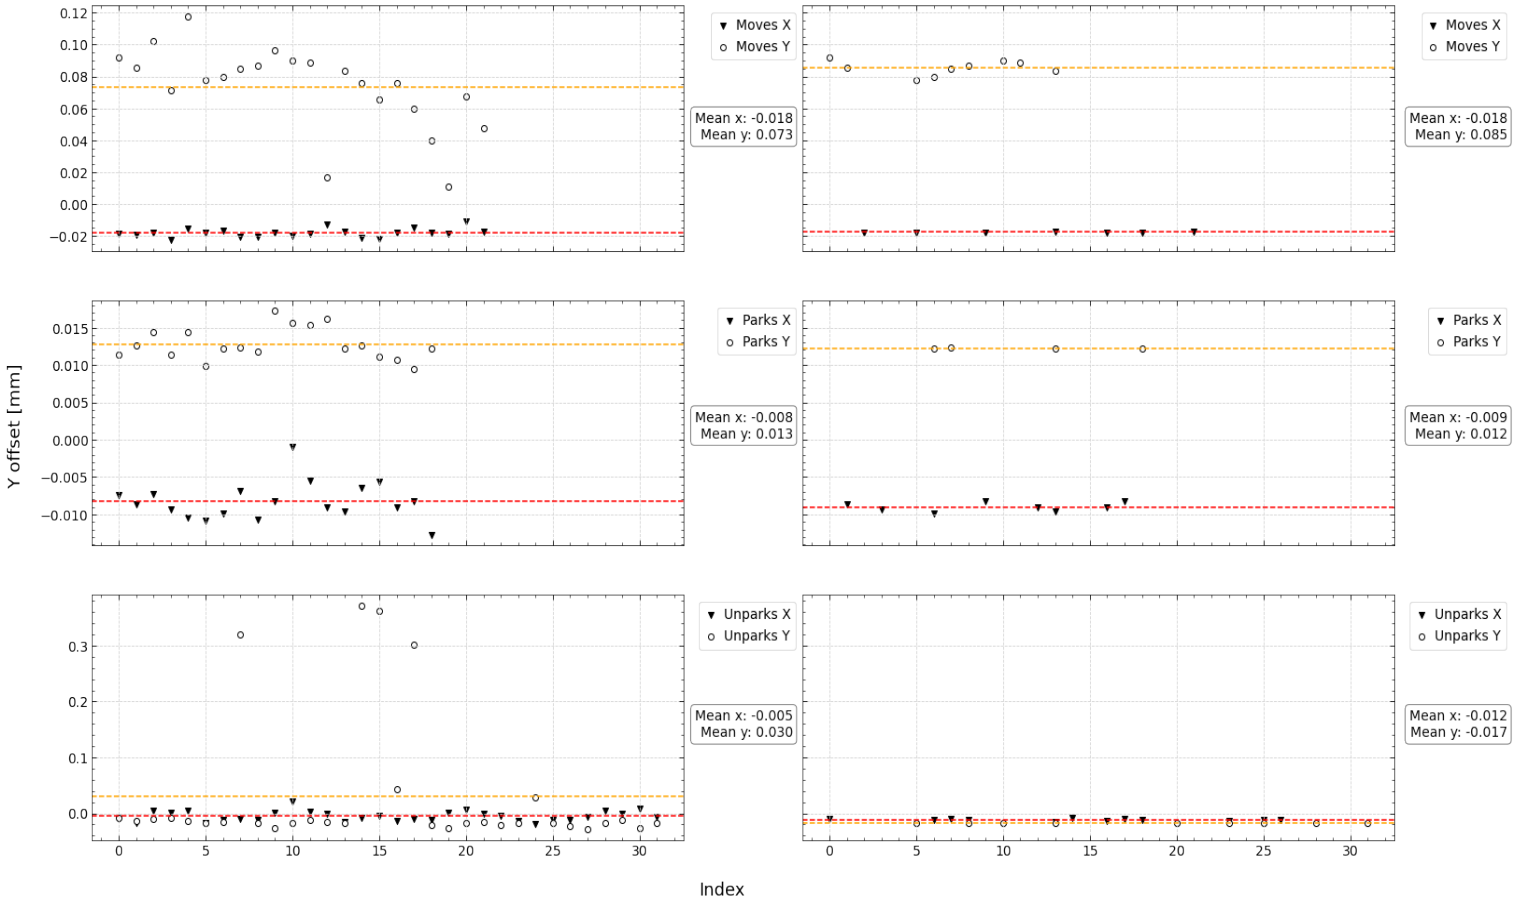


Figure 3.27: (*Left*) The *Move* (*top*), *Park* (*middle*), and *Unpark* (*bottom*) offset values for fibre 39 that have been recorded in the calibration database for plate A using Nona. (*Right*) The offset values with a two rounds of sigma clipping, where the red and orange lines represent the mean x and y offset values calculated.

rolling mean is used for fibres with more than 50 entries, however the exact level may be varied subject to further fibre movement testing. The effectiveness of this method will ultimately be determined by its ability to place a fibre within its positioning tolerance on the first attempt, which will be evaluated as the MOS observation rehearsals take place. Whilst the commissioning period continues, the database will be used to store all of the relevant calibration information and the logfile processing will become a part of the scheduled activities.

The offset calculations have primarily been carried out at ZD 50°. Further testing is needed to analyse how these offsets evolve with the telescope elevation angle. These measurements are important as they provide information on the effectiveness of the predicted park heights and plate maps. In principle, if the model predictions are correct then the distribution of the fibre offsets should remain relatively unchanged.

At the time of writing, it takes 47 minutes to park 900 fibres in the ring configuration on MOS A, using information recorded from the positioner logfiles. This gives an approximate configuration time estimate of 1 hour and 30 minutes. Given that the robots are running at a reduced speed, approximately 60% its capabilities, this is an excellent result. By reducing the number of attempts it takes to complete a single fibre placement and adjusting the kinematic acceleration limits for each axis, it seems highly likely that WEAVE will meet its 1 hour configuration time requirement once full survey operations commence.

Chapter 4

Field configuration optimisation

If you can't teach it to a duck then you don't know it

E. Schallig

As astronomical instruments expand both their capability and complexity, improvements in modern computing algorithms are needed to ensure they operate smoothly. Increasing the scale and number of components an instrument has introduces a greater number of combinations for the state of the system. This creates challenging combinatorial optimisation problems, which must be solved to operate the instrument efficiently.

These types of problems are referred to as the Travelling Salesman Problems (TSPs): given a set of n cities and the distances between them, a salesman must find the shortest possible route whilst visiting each city once before returning to the start. It is an example of a nondeterministic polynomial time problem (NP-hard) which is an active area of research in computer science and mathematics (Schrijver et al., 2003).

WEAVE contains two combinatorial optimisation problems; One when choosing which fibres to allocate to the assortment of targets in the field of view, and another when deciding which order each fibre should be placed and by which robot. The

former is not strictly a TSP and the later is classed as a TSP. This chapter describes the work completed in (Hughes et al., 2022b) to discuss how the MOS fibres are allocated to targets in the field of view using a program called *Configure* (Terrett et al., 2014), the solutions it provides, and how its assignment procedures can be modified to maximise the positioners scientific output.

4.1 Fibre Allocation

To allocate the positioners MOS fibre multiplex to the targets in the field of view, astronomers in the WEAVE collaboration implement a program called *Configure* (Terrett et al., 2014). *Configure* uses an algorithm known as Simulated Annealing (SA) (Kirkpatrick et al., 1983) to maximise the fibres used per configuration and assign them to targets according to their listed priority (or weighting). This software follows on from the assignment program used for 2dF (Miszalski et al., 2006) (Colless et al., 2001).

Thus far, *Configure* has only been tested with the goal of maximising the number of fibres assigned to each field, as this is critical for optimising the efficiency of the positioner. *Configure* also has the ability to prioritise the assignment of targets, which is necessary for the WEAVE Science teams to specify which targets must be included in their surveys.

In this chapter, we explore several input cases using *Configure* to study its behaviour and find adaptions which will both maximise the fibres deployed per observation and the number of high-priority science targets observed. The goal is to answer the following questions:

1. Does *Configure* distribute the fibres evenly across the field of view or are there structures in the fibre assignment distribution?

2. What impact does changing the target priority have on the field assignment?

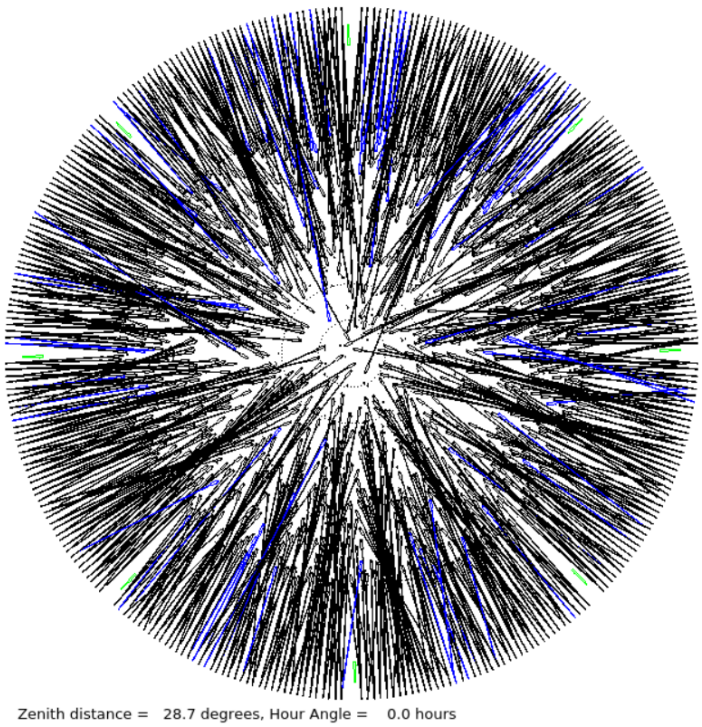


Figure 4.1: The allocation of WEAVE's $\sim$ 960 MOS fibre multiplex for a single observation field. This is produced by the program *Configure* for a field with 1400 targets which are randomly distributed over the field of view.

3. Is the nominal usage case the most effective method of using *Configure*? If not, then what is?

4. Are there any target selection effects that the survey teams need to account for?

To test *Configure*'s assignment procedure, 34 proto-fields were provided by the WEAVE-LOFAR (WL) team from their mid-tier survey. This tier will observe up to 1,000 deg^2 and has a selection flux of $\geq$ 1 mJy at 150 MHz. Further information on the WL survey can be found in (Smith et al., 2016).

To answer these questions, the following tests have been conducted and are described by this chapter:

- Visually inspecting the target completeness distribution across the FoV.

- Modifying the target priority scheme to identify potential improvements.

- Assigning the fibres to targets in a two-stage process based on the targets priority value.
- Assigning the fibres to targets simultaneously whilst using the multi-survey functionality, where targets are placed in pseudo-surveys based on their priority value.
- Calculating the fraction of fibres assigned in overlapping regions for the nominal and multi-survey usage case.
- Calculating the cumulative assignment distribution as a function of target flux, where several cuts have been applied to the target catalogue.

Since all survey teams in the WEAVE collaboration use this application, the results of this chapter will have a far reaching impact on the scientific observations completed with WEAVE.

4.1.1 Configure

In comparison to brute force combination testing, *Configure* is significantly faster. It uses some dynamic programming elements of the original 2dF Oxford algorithm alongside SA to successfully assign the WEAVE fibre multiplex in several minutes using conventional computing power.

The SA algorithm is an established stochastic optimisation algorithm[1] used to solve complex combinatorial optimisation problems. It is based on the annealing process of materials, where crystalline structures form as the system is gradually cooled. SA can be thought of as a series of random perturbations introduced via the Metropolis algorithm, which simulates thermal motion (Metropolis et al., 1953). First, a random solution is generated. For the positioner, this is a random allocation

[1] it allows less optimal decisions to be made during the process to improve the chances of finding the global optimal solution.

of fibres to targets in the field of view. The thermal motion is applied as a series of random swaps between the fibres and the targets.

The temperature parameter controls the probability that these perturbations will be accepted. After each perturbation takes place the resulting change of energy ΔE is calculated. If ΔE is negative then the perturbation is accepted and used as the starting point for the next step; otherwise the acceptance is determined probabilistically using

$$P(\Delta E) = \exp(-\Delta E / k_B T), \tag{4.1}$$

where $P(\Delta E)$ is the probability of acceptance, k_B is the Boltzmann constant, and T is the effective temperature for optimisation. In practice, the temperature parameter has the same units as the cost or merit function (Bertsimas and Tsitsiklis, 1993), which is the error estimate for the current solution. If we set out to maximise a merit function M for a configuration and if δM is the change going from the current configuration to a new configuration, then the probability of accepting the change is equal to 1 if $M \geq 0$ and $\exp(M/T)$ if $M < 0$, where T is in units of the merit function. The temperature of the system is gradually decreased, which reduces the probability of perturbations being accepted and fixes a growing number of fibres in place until a locally minimal solution is found.

SA does not guarantee that the global optimal solution will be found in a finite amount of time, yet it is able to find an approximate locally minimal solution relatively quickly. Another advantage of this process is that it prevents extreme deviations from thermal equilibrium, reducing the likelihood of structures forming which are only optimal locally. SA is still a popular algorithm choice for TSPs however, new developments in Machine Learning (ML) methods, such as genetic algorithms (Grefenstette et al., 1985), and Reinforcement Learning (Gambardella and Dorigo, 1995) are likely to dominate this field of research in the future.

Before the SA algorithm can be implemented, a series of refinement steps are made by *Configure* to significantly improve its efficiency. A single observation of WEAVE has approximately 1400 targets and a maximum of 960 fibres available for assignment. This leads to $\sim 6.9 \times 10^{376}$ combinations, which is a computationally infeasible procedure. However, not all of these combinations are physically possible, due to the limited patrol radius applied to each fibre. This reduces the possible combinations by a factor of 60, which remains an infeasible number of combinations but rules out unphysical solutions. To incorporate this simplification process, *Configure* computes a collision matrix that records whether there is a conflict for every possible pair of fibre target allocations. These conflicts arise for multiple reasons. For example, if the fibre cannot physically reach that region of the field, or if the outline of the fibre will block another nearby target.

There are four stages involved in calculating this collision matrix and fibre allocations. They are summarised here and further details can be found in (Terrett et al., 2014). Firstly, the position of the targets in the focal plane are calculated, using parameters such as the atmospheric pressure, temperature, hour angle, and humidity. The targets location is used to isolate a list of fibres which may reach it, along with the area the fibre would cover, referred to as the fibre's footprint. *Configure* parallelizes this process for multiple targets simultaneously.

The second phase consists of a series of checks that are independent of the fibre allocation. For example, if two targets are positioned in such close proximity that a collision is certain, then there is no need for additional checks to be performed for that particular pair. The checks are then extended to to search for overlapping footprints between fibre pairs, and footprints which overlap nearby targets. Lastly, it identifies any fibres which may overlap with button outlines placed on a nearby target. The reliability of the collision matrix is a crucial factor, as fibres may be broken if collisions are not flagged. Therefore, this matrix has been tested extensively in software by

David Terrett and the initial configurations were conducted under carefully controlled conditions (see chapter 3).

Configure then combines the previous two stages by examining fibres that can be placed on the first target in a nearby pair. If there was a fibre-button overlap found for the first target in the second stage, it is also possible there could be an overlap for the second target. To ensure that all possible fibres reaching the first target are checked for this condition, a thorough examination is conducted. If this process results in a conflict, then the checks are initiated for the second target in the pair. By reducing the number of possible fibre target combinations early on in the process, these conflict assessments are significantly faster; when a clash is found, no further tests for conflicts are required.

It is important to note that for the WEAVE-MOS observations, a number of fibres are reserved for sky observations. These fibres are allocated at the end of the process and placed in available spaces across the field of view. This is typically set to ~ 100 fibres and caps the maximum fibres available for targets to ~ 850 per observation.

4.1.2 Completeness

Understanding the distribution of target assignments across the field of view is critical for the success of WEAVE science surveys. Certain regions of the field may be prioritized over others, which can result in an imprint similar to that caused by the Oxford 2dF algorithm, shown in figure 4.2. It is important to note that this field configuration algorithm (FCA) was designed for overlapping 2dF fields with 800 fibres and was not designed to be smooth (Colless et al., 2001). These imprints are not solely due to the Simulated Annealing algorithm, but can also arise from limitations in the movement of fibres based on their tier. Therefore, it is important to carefully consider the target assignment in each region to avoid biases, by using a representative sample of the population under study.

In astronomy, completeness refers to the fraction of objects in a given sample that have been detected by observations, compared with the total number that actually exist in the population. It is a measure of an instruments sensitivity and efficiency. Any regions of low assignment across the field of view caused by *Configure* must be identified, as the completeness would vary as a function of position, which could impact measurements of the luminosity distance or space density.

The definition of completeness in the context of fibre assignment, C, used in this chapter follows from Miszalski et al. (2006) where

$$C = \frac{n_a}{n_a + n_u}. \tag{4.2}$$

Here n_a represents the number of targets that have allocated fibres and n_u represents the number of targets that have not been allocated fibres. The additional scaling factor in (Miszalski et al., 2006) has been excluded, as the WEAVE survey teams have been advised to use approximately 1400 targets per survey field. This recommendation is given by (Terrett et al., 2014) to make use of the full fibre multiplex available. However, further study may be needed as open time for WEAVE observations becomes available, in the event that proposals with a varying number of targets are put forward.

To analyse the completeness distribution as a function of position, the activate area of the field plate was divided into regions measuring 5×5 mm (approximately 89" $\times$ 89"). This specific size was chosen to ensure that each region consistently contained a small number of fibres and targets; reducing this scale would result in regions that had only one or no targets at all.

To begin the analysis, the 34 proto-fields provided by the WEAVE-LOFAR team were processed using the nominal method with *Configure*. This involves supplying *Configure* with a field with no modifications to the annealing parameters and all

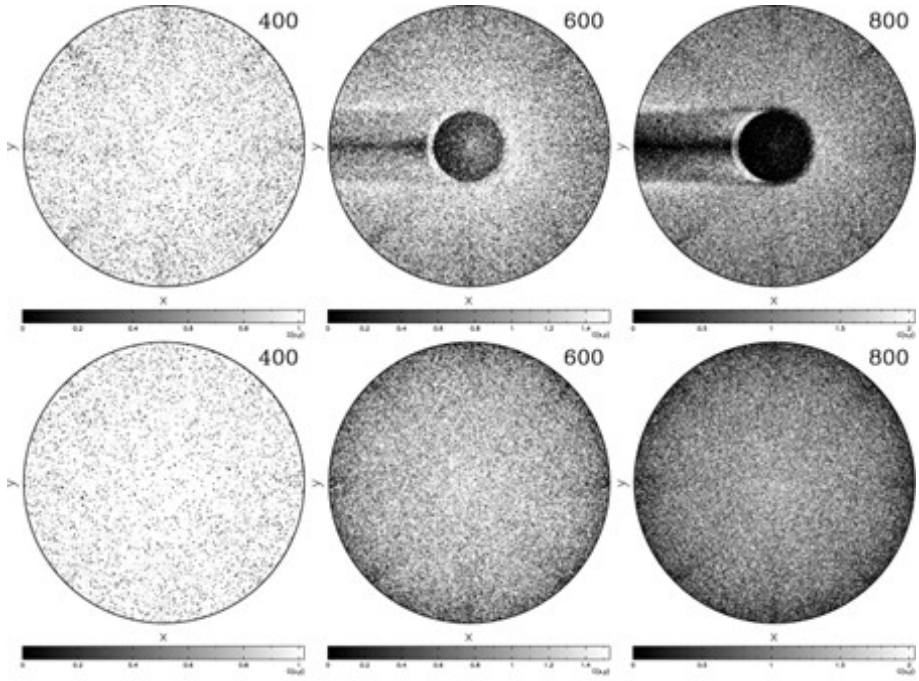


Figure 4.2: The completeness distribution presented in (Miszalski et al., 2006). $C(x, y)$ for the 'Oxford' (*top*) and SA (*bottom*) field configuration algorithm (FCAs) for fields including 400, 600 and 800 uniformly distributed programme targets and an additional nominal 30 fiducial targets. The 'Oxford' FCA imprints a circular deficiency and 'curtain-like' feature that becomes more prominent with increasing target density. Also visible is a 'flower-like' imprint by the eight guide fibres. The high degree of randomization used by the SA FCA eliminates much of this structure, save for a slight gradient in the radial direction, however it cannot remove the 'shadowing' caused by the guide fibres.

targets, regardless of their priority, are assigned simultaneously.

It is important to note that the WEAVE survey teams have agreed to use a shared target priority (or weighting) scheme. This prevents biased target assignment when survey teams share fields of view during observations. The priority scheme is based on integer numbers ranging from 1 to 10, where 1 indicates targets with the lowest priority and 10 indicates the highest priority. Each survey team has complete freedom to determine how they implement this priority range for their targets. However, *Configure* is not bound by this limit and the affect of expanding this range is discussed in section 4.1.3.

For the WL observations, a system where all the high priority science targets

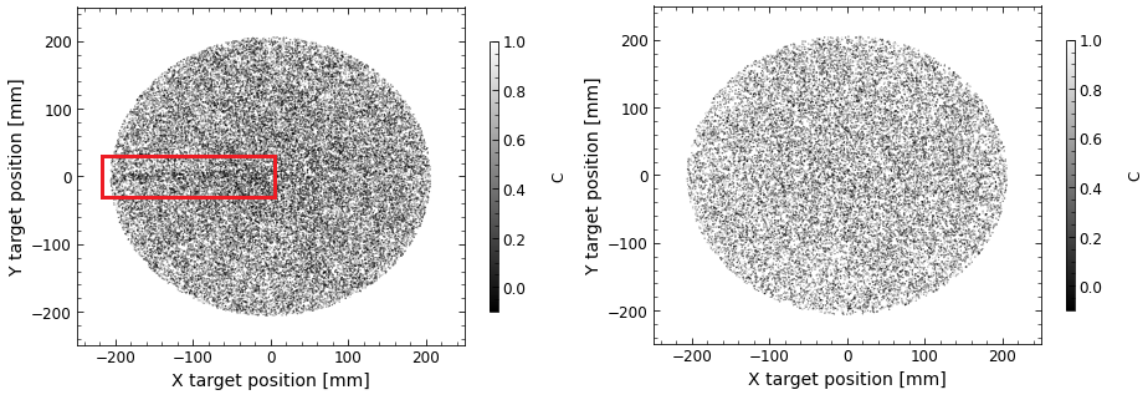


Figure 4.3: (*Right*) The distribution of the targets for the 34 mid-tier WEAVE-LOFAR fields, stacked together, where each field contains 1400 targets. Each point is shaded according to the completeness of the region it is in. Darker areas indicate regions of the field where the assignment is less complete, and the field is binned into regions of 5×5 $[mm^2]$. (*Left*) A stack of the initial mid-tier WL fields, prior to a 1400 target limit being imposed. The red box highlights a radial stripe, which is caused by low completeness as the same guide fibre has been assigned to the centre of each field.

were assigned a value of 10, and all other targets were assigned a value of 5 was initially implemented. Figure 4.3 shows the stacked distribution of all targets in the 34 proto-fields, coloured-coded according to the completeness value for each region. This completeness value is calculated on a field-by-field basis, with darker areas indicating consistently low areas of fibre allocation.

Upon plotting figure 4.3, a single dark radial segment was initially observed over the field of view. After reviewing *Configure*'s output files, it was discovered that a single guide fibre was being consistently allocated to the calibration star at the centre of the field for each observation, thereby blocking off the a section of fibres during the allocation process. *Configure* always assigns guide fibres to calibration targets prior before assigning the MOS fibres to ensure that telescope can accurately track targets during observations.

After randomising the guide fibre allocation process, the distribution in figure 4.3 displays no other visibly discernable features. This is a positive result, as there is no obvious imprint due to the fibres tiered arrangement.

4.1.3 Assignment cases

The nominal assignment case of *Configure* assigns all targets simultaneously. When introducing perturbations during the annealing process, the acceptance function is weighted according to the priority of the target being allocated in each fibre swap. In the 2dF algorithm, this is given in terms of the energy E, where

$$E = \sum_{i}^{N_{\rm piv}} \left[\beta^{P_i} + \delta \sum_{j}^{N_{\rm Assoc}} \beta^{P_j} \right] \left(\frac{\alpha_i - \alpha_{\rm max}}{\alpha_{\rm max}} \right)^{\gamma}, \tag{4.3}$$

and $\alpha_{\rm max}$ is the maximum fibre deviation, α_i represents the current angle deviation for this pivot, $N_{\rm piv}$ is the number of targets in the field, and P_i is the priority of the allocation to the target (Miszalski et al., 2006). The parameter $N_{\rm Assoc}$ is defined as the number of close pairs, where 'close' refers to pairs that have been excluded by the collision matrix. The additional parameters δ and γ are used to control the preference for close pair assignment and favour fibres which have radial trajectories with small angular deviations. The more positive δ is, the less likely an object with close pairs is to be accepted. Equally with γ, the more positive its value, the less likely it is that large angular deviations are accepted. The β term is used as a weighting base, described in (Campbell et al., 2004).

Configure uses a simplified version of equation 4.3, where the energy E of each individual fibre with a target assigned to it is given by

$$E = \frac{1 + s}{\text{target priority}} \tag{4.4}$$

(Terrett et al., 2014) where s represents the straightness of the fibre, which is much less than 1. The straightness is defined as $s = d \sin(\theta_{\rm gate})^2$, where d is the distance to the target in units of radius and $\theta_{\rm gate}$ is the gate angle of the fibre as it reaches the target. The target priority can be any positive integer value, where

zero excludes a target from the assignment process. Whilst the WEAVE survey teams use a maximum value of 10, there is no upper limit imposed by the software. The implications of changing this upper limit are discussed later in this section. The *Configure* algorithm searches for the configuration which gives the lowest energy value. The end temperature that *Configure* uses as a cut-off to stop the assignment process, is an adjustable parameter. By default, it is set to 50, and the nominal start temperature is 200 which reduces in steps of 0.01. The units are in terms of the energy, and are on approximate the same scale as the energy differences between randomly generated field configurations. If a fibre swap occurs in the perturbation process that lowers the total energy of the system, it will automatically be accepted. However, if the result increases the energy value, then the perturbation has a probability P of being accepted, where

$$P = \exp\left(1000\left(\frac{E_2 - E_1}{T}\right)\right). \tag{4.5}$$

The values E_1 and E_2 correspond to the energy before and after the fibre swap takes place, and T is the temperature of the system at the time. Since the temperature of the system gradually decreases as it "anneals", the probability of a fibre swap being accepted also decreases as the system cools.

To quantify the variation in fibre assignment associated with the nominal usage of *Configure*, I conducted 450 Monte Carlo (MC) iterations using a single WL proto-field. This number of iterations was used to balance the need for accurate uncertainty estimation with manageable computational time requirements, and will remain consistent across all subsequent MC distributions for the assignment test cases.

The distributions for the total number of targets assigned and the number of higher priority targets assigned for the MC realisations are shown in figure 4.4. The uncertainty boundaries were taken as the 16^{th} and 84^{th} percentiles of the MC distributions, which corresponds to an uncertainty of ± 2 in the total number of fibres assigned and ± 9 for the high priority targets.

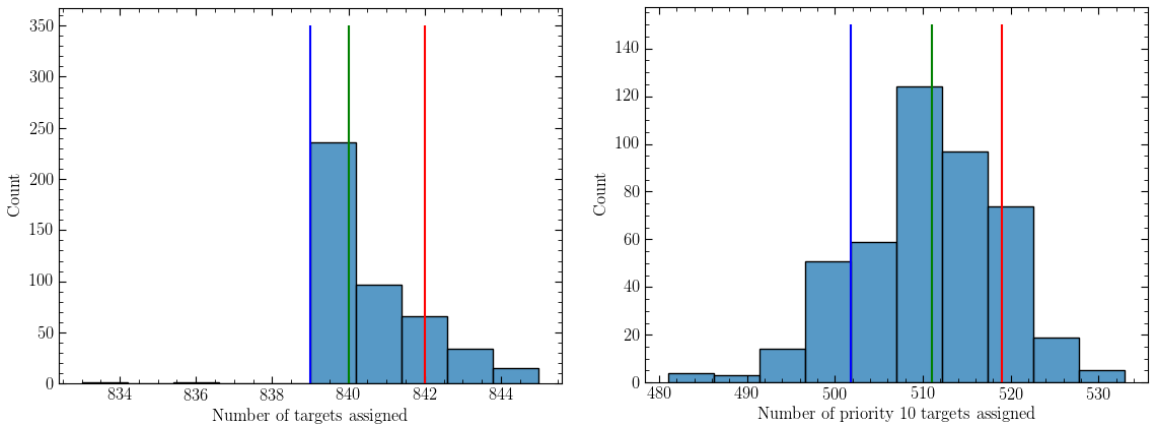


Figure 4.4: Monte Carlo assignment realisations for the nominal *Configure* usage case, processed 450 times with a single WL proto-field. (*left*) The total number of targets assigned for each iteration. (*right*) The number of priority 10 targets assigned per observation, in each realisation. The blue, green, and red lines represent the 16th, 50th, and 84th percentiles respectively.

Prior to the work completed in this chapter, the WL team had only used priority values of 5 or 10 in their field assignment process. For the mid-tier survey, a target with a flux density above 1 mJy is considered a high priority target and given a value of 10. All other targets are assigned a value of 5. The remaining WEAVE survey teams have implemented independent strategies for the target assignment of their fields, however these strategies are outside the scope of this thesis.

Using the 34 proto-fields from the mid-tier WL survey, I investigated how the divergence between these priorities impacts the number of high priority targets assigned per field. The first step was to maximise the value difference between the high and low priority targets, whilst remaining within the boundaries set by the survey team agreement. This meant reducing the low priority target value to 1 rather than 5, whilst the high priority target values remained at 10. Figure 4.5 shows the number of higher priority values assigned in each proto-field against the total number included, for both WL's nominal priority values and the maximised difference.

It can be seen in figure 4.5 that maximising the difference between the two priority values consistently increases the number of high priority targets allocated in each field,

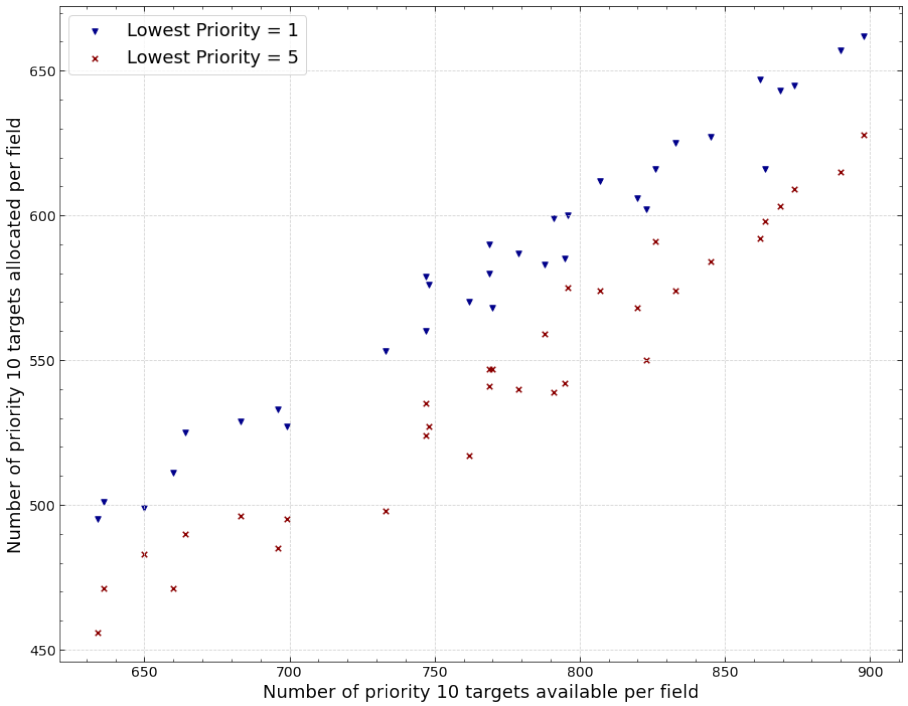


Figure 4.5: The number of priority 10 targets included in each field against the total number of priority 10 targets allocated by Configure, for each of the 34 proto-fields in the mid-tier WL survey. The red points represent the nominal target priority labelling applied by WL, and the blue points show the case where the difference between the high and low priority values has been maximised, whilst remaining within the 1 to 10 boundary given by the WEAVE survey teams.

when using the nominal annealing schedule. This has resulted in the WL modifying their priority value assignment scheme. To ensure the statistical integrity of these findings, further tests should be conducted with a larger subset of fields across all three of the WL survey tiers and across a range of annealing schedules.

To assess the wider impact of changing the target priority on the field assignment, I repeated this test beyond the boundary applied by the survey teams. Figure 4.6 shows the distribution for the 34 proto-fields given in figure 4.5, with the higher priority targets labelled with an extreme value of 10,000. This figure demonstrates that there is no significant impact after increasing the maximum priority value beyond the 1 to 10 range for the set of mid-tier WL fields. This should be conducted with a larger subset of fields across all three of the WL survey tiers to verify the result.

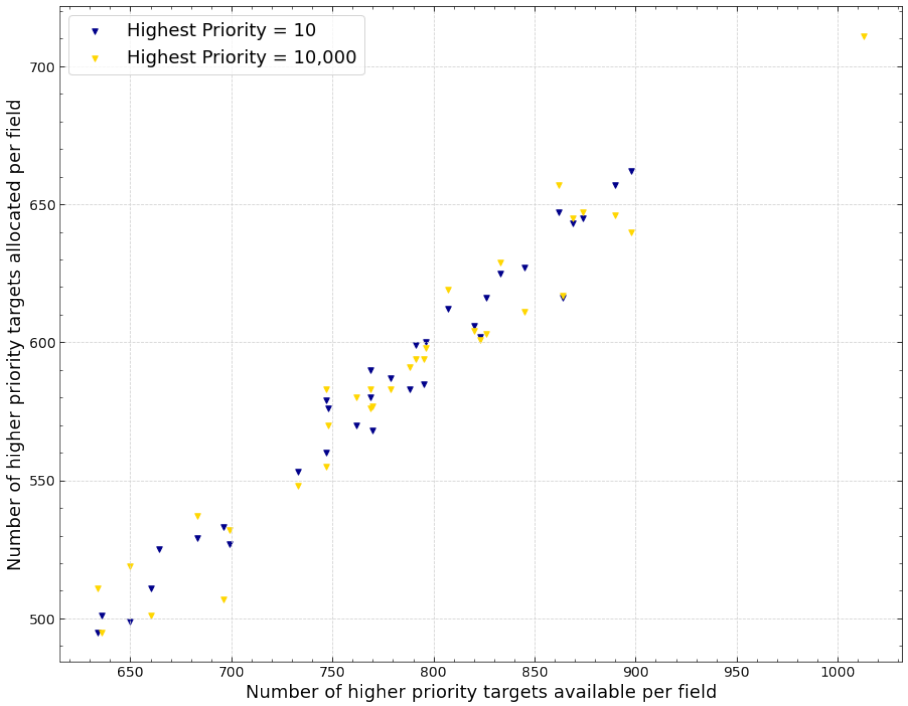


Figure 4.6: The number of higher priority targets assigned per field, against the number of higher priority targets included in each field. The higher priority values have been set to either 10 or 10,000 across the 34 WL proto-fields, represented by the blue and yellow points respectively.

Many of the WEAVE science teams share fields of view over the northern hemisphere. Over the 5-year survey, there are $1,858,420$ fibre hours allocated to shared fields of view, which corresponds to 71% of the total planned across the full 7-year survey (Trager, 2019). *Configure* has been designed with this need in mind, and each target can be assigned a survey within the input file prior to fibre allocation process. The input file allows the user to place limits on the number of fibres that each survey may use within the field of view, which prevents the preferential assignment of fibres across shared survey fields. However, there are no restrictions on the naming conventions that must be used. The number of fibres available between shared survey fields has been determined by the individual science teams and is based on the total number of fibre hours available for each survey within the project. The 1 to 10 priority scheme is designed to prevent preferential assignment cases for the shared fields,

although each survey has the freedom to decide on their own target priority policy.

To study whether the nominal usage case of *Configure* is the most effective method, I conducted a series of test cases which utilised the shared field assignment feature and *Configure*'s ability to exclude targets from the fibre allocation process. This method should, in principle, give more flexibility compared to the two stage method as the fibres assigned to higher priority targets are not fixed in place. In addition, this method allows us to place limits on the number of lower priority targets assigned.

These test cases were performed by modifying the input file for each of the 34 WL proto-fields, without adjusting the annealing parameters. Initially, I conducted short tests with small shifts in *Configure*'s ΔT parameter, which dictates the rate of the annealing process, to understand how the rate of the perturbations impacted the assignment statistics. This led to a significant increase in the computational time required to allocate a single field. In order to ensure that the field processing for each survey team was not slowed down, I decided that the default parameters would be used for the test assignment cases. However, further study is needed to explore this topic with a wider selection of fields. As the WEAVE survey is large in scale, even small yet consistent improvements to the fibre allocation process would be highly beneficial. Thus, a more comprehensive analysis of this issue could lead to significant gains for the survey as a whole.

The first assignment test case was to implement a two-stage annealing process, which had already been experimented with by science team members outside of the WL group[2].

As the name implies, the input files are processed through *Configure* twice. In the first stage, only the higher priority targets may be allocated fibres. The listings of the lower priority targets had their priority values set to zero in order to exclude them from the process without losing their information. The output file from this stage

[2]At the time of these tests, I was told this method was being implemented by some members of the WEAVE-Clusters group, however I do not have direct confirmation of this.

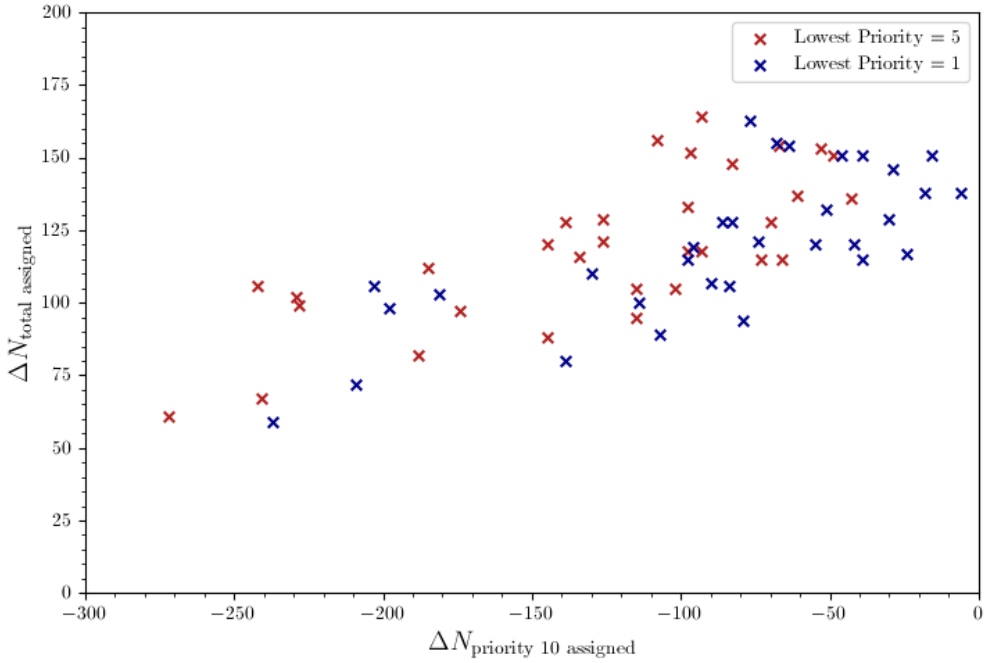


Figure 4.7: The difference between the nominal assignment and the two-stage processing method when comparing the total number of targets allocated per field against the number of higher priority 10 targets assigned, described as $\Delta N_{\text{total assigned}}$ and $\Delta N_{\text{priority 10 assigned}}$ respectively. Each case has been processed using a lower WL priority value of 5, or 1 for the maximum difference scheme.

is then fed through *Configure* a second time with the original lower priority target values restored. The higher priority targets now have a pre-allocated fibre number in their listing. The final output file contains the complete set of targets with an appropriate fibre allocation dictated by *Configure*.

Figure 4.7 shows the difference between the nominal and two-stage assignment process when comparing the distribution of the total number of targets assigned in each proto-field against the number higher priority targets allocated, represented as $\Delta N_{\text{total assigned}}$ and $\Delta N_{\text{priority 10 assigned}}$ respectively. This includes both the original WL target priority assignment and the maximum difference scheme. From this figure, it appears that using the two-stage method does not produce a significant difference between the nominal WL priority labelling and the maximum difference

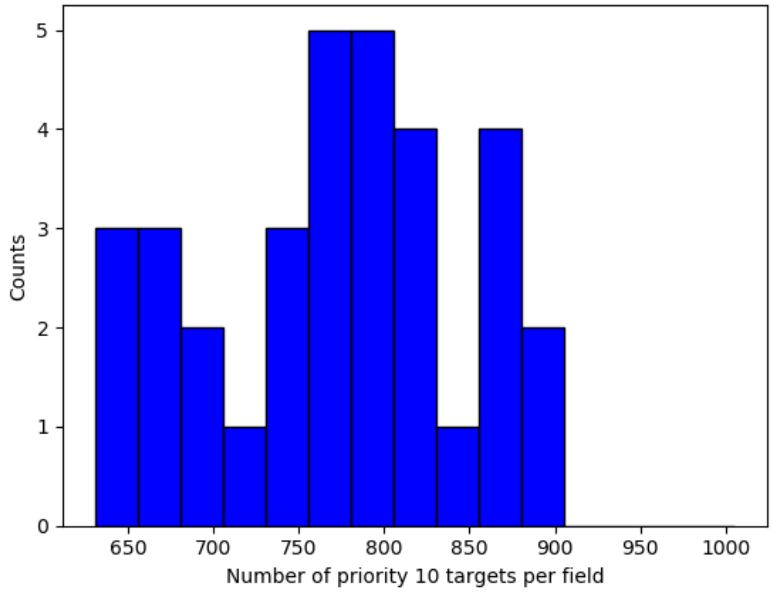


Figure 4.8: The distribution of high-priority targets across all 34 WL proto-fields. This distribution was used to divide the proto-fields into three sub-groups, with boundaries at 700 and 825 targets, corresponding to the bins with the lowest count.

scheme. However, it can be seen that as the difference in the total number of targets assigned between both methods increases, the gain from the two-stage method in terms of the number of higher priority targets assigned decreases. By comparing the number fibres assigned for each proto-field, the two-stage method consistently results in a decrease of the total number of fibres allocated in each field. The best case gives a 5% decrease and 20% at worst. The impact of this decrease on the survey science output is dependent on the importance of the high and low priority targets and would be unique to each survey.

These test cases have been dissected further to look at the impact of the two-stage process on the radial completeness across the field of view. The set of 34 WL proto-fields have first been divided into three sub-groups, based on the total number of priority 10 targets in the field. This is shown in figure 4.8, where the boundaries between groups have been set at 700 and 825 priority 10 targets, which correspond

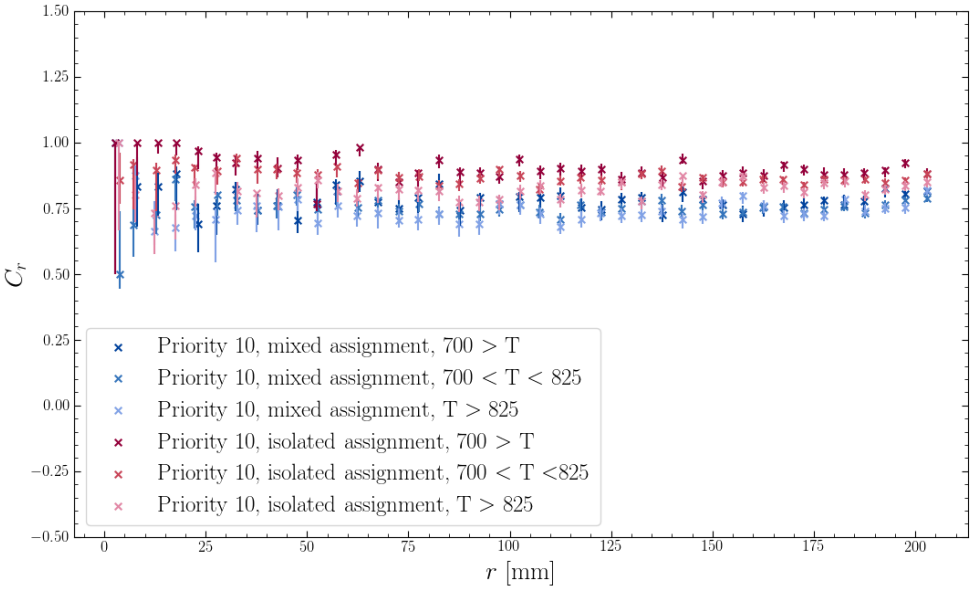


Figure 4.9: The radial completeness distribution C_r for both the nominal assignment and the two-stage assignment process. Each case has been divided into three groups based on the number of priority 10 values included in each field, with boundaries at 700 and 825 targets. This distribution shows the average completeness across all fields in each group as a function of radius.

to the lowest bin values of the distribution.

The radial completeness C_r builds upon the method used in figure 4.3, where the field is divided into annuli with a width of 5 mm. The completeness value is then calculated as before using equation 4.2, and the final value for each annulus is taken as the average across the sub-group. The radial completeness distribution for the two-stage and nominal processing for the maximised priority scheme is shown in figure 4.9.

This distribution demonstrates that the two-stage method does not introduce a radial dependency in the completeness function for all three sub-groups, which was possible due to the comparable number of higher priority targets to the total number of fibres available for each field. However, due to the reduced fibre usage per proto-field it should not be considered a viable processing method for the WL survey.

Based on the available evidence, it appears that this affect can be attributed to a lack of flexibility in the fibre assignment. The first stage "freezes" a large proportion of the available fibres into place. This means that fibres assigned to a high priority target cannot be moved to another, even if such a swap would allow several lower priority targets to be assigned fibres. Therefore, we take the first stage of this method as an estimate for the maximum number of higher priority targets it is possible to assign in a single field.

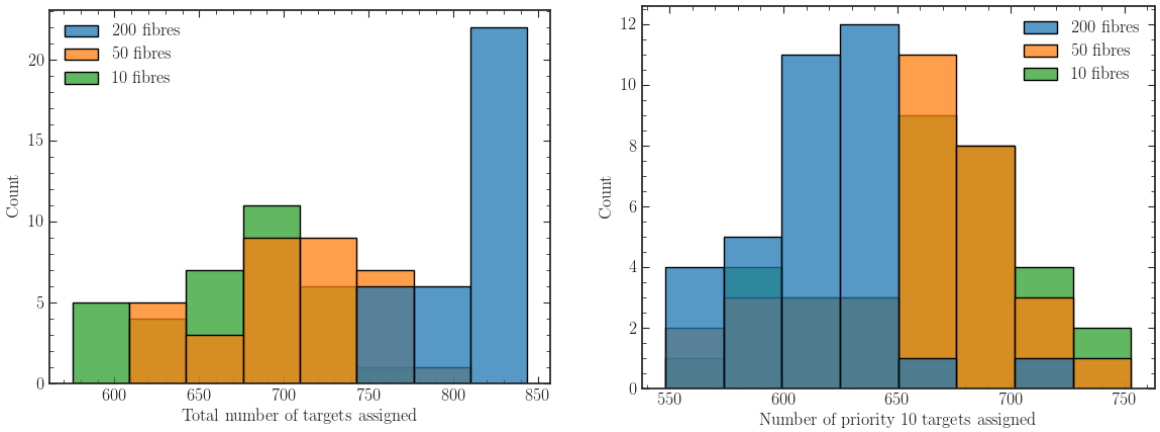


Figure 4.10: The distribution of assignment for the 34 mid-tier WL fields when configured under the multi-survey priority case. (*left*) The total number of targets assigned in each field, when either 10, 50, or 200 fibres are allowed to be used by the lower priority survey. (*right*) The number of priority 10 targets assigned, when either 10, 50, or 200 fibres are allowed to be used by the lower priority survey.

For the nominal usage case, the 50^{th} percentile of MC iterations yields 810 fibres, which is close to the available 860 fibre multiplex when accounting for the 100 fibres that are reserved for on-sky calibrations in each WL proto-field. Therefore, the aim is to find an alternative assignment method which brings us as close as possible to maximum number of higher priority targets that can be allocated, without compromising on the number of fibres used per observation.

To achieve this goal, I used *Configure*'s ability to assign shared fields of view to place limits on the fibres availability. Each target is given a pseudo-survey name based on its priority, with a limit that is defined in the input file. As this method

assigns the targets simultaneously, this method should prove to be more flexible than the two-stage assignment process.

The high priority pseudo-survey has a limit of 960, which is the maximum number of fibres available per field. The 100 sky fibres are not affected by imposing this limit. In contrast, the lower priority survey has a variable limit on the maximum number of fibres available, which I limited to range between 10 to 200 fibres. This range is based on the information shown in figure 4.7, which indicates that the maximum number of higher priority targets allocated in a single field is approximately 750. As a result, there are roughly 200 fibres left that can be allocated to lower priority targets in the ideal case.

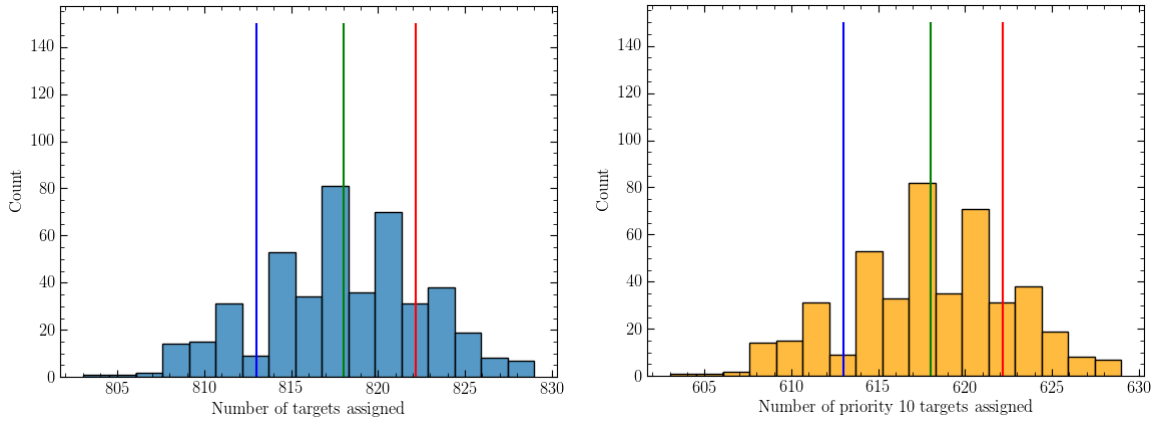


Figure 4.11: Monte Carlo assignment realisations for the multi-survey *Configure* usage case, repeated 450 times with a single field where 200 fibres are allowed for the lower priority survery (*left*) The total number of targets assigned for each iteration. (*right*) The number of priority 10 targets assigned for each iteration. The blue, green, and red lines represent the 16^{th}, 50^{th}, and 84^{th} percentiles respectively.

Figure 4.10 shows the assignment distribution for the 34 WL proto-fields using the multi-survey method, where the lower priority survey limit has been set to 10, 50 and 200 fibres respectively. It can be seen that the limit of 200 fibres consistently uses the majority of the available fibres for a large proportion of the proto-fields, whilst increasing the number of high priority targets allocated by $\sim$ 100 targets in comparison to the 50^{th} percentile of the nominal MC iterations shown in figure 4.4.

MC realisations were then conducted for a single proto-field using the multi-survey processing method, with a lower priority survey limit of 200 fibres. This is shown in figure 4.11. This method gives uncertainties of ±3 for both the total number of targets and the higher priority targets allocated.

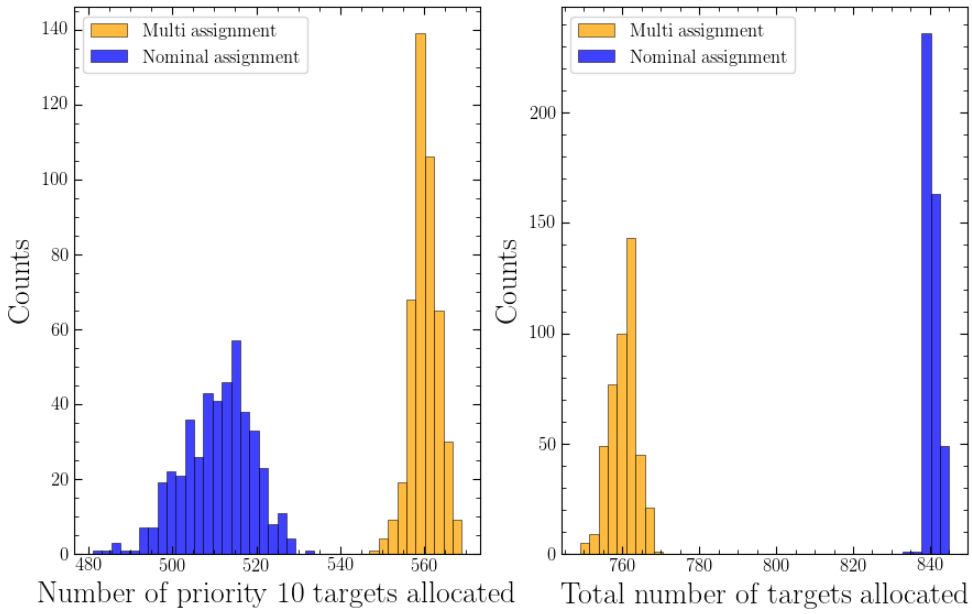


Figure 4.12: MC realisations of a single WL-protofield, processed by Configure 450 times for both the nominal case and using the multi-survey method. (*left*) Shows the distribution of priority 10 targets allocated per realisation. (*right*) Displays the total number of targets allocated in each realisation.

Figure 4.12 displays the MC realisations for both the nominal and the multi-survey processing for a direct visual comparison[3]. The peaks of these distributions suggest that the multi-survey method allocates approximately 50 additional high priority targets, at the cost of 80 fibres being excluded from the assignment process. This is a 9.5% increase in the number of high priority targets in the worst case scenario, and 12.3% for the best case. Clearly, this not an ideal trade-off, however this method of

[3]These MC realisations were conducted on different WL proto-fields. I do not expect this to significantly impact the results however, this test should be repeated as each field has a distinct setup and number of higher priority targets.

processing fields through *Configure* gives the survey teams an extra degree of freedom to decide whether they wish to prioritise key science targets or obtain a larger quantity of spectroscopic data. This method should be considered as an alternative starting point, which appears promising compared to the nominal assignment case.

4.1.4 Selection effects

It is important to understand *Configure*'s assignment process in the context of the target flux density distribution, as this will highlight potential selection affects introduced by the algorithm. These selection affects will impact the scientific analysis conducted with WEAVE data in areas of research such as stellar and galactic luminosity functions. This distribution has not been previously analysed by the WEAVE survey teams or in *Configure*'s initial testing.

This analysis will study the completeness of the assignment process as a fraction of the targets available in the field catalogue to determine whether it is flux dependent. In this case, the completeness is redefined as

$$N\left(F_i \geq f\right)/N_{\text{catalogue}}, \tag{4.6}$$

where $N\left(F_i \geq f\right)$ is the reverse cumulative sum of the assigned targets as a function of their flux density across all 34 proto-fields. For this analysis, a normalisation factor of 850 has been applied to represent the fibres available for each field which improves the scaling, and $N_{\text{catalogue}}$ is the number of targets in the WL mid-tier catalogue.

This analysis continued to use the subset of mid-tier WL fields, where targets with a flux density > 1 mJy are given a priority value of 10, and targets with a flux density < 1 mJy are given a value of 1. The results presented in this section are dependent on the survey fields used, however the overall behaviour is relevant for all MOS observations across all WEAVE surveys, when the number density of potential

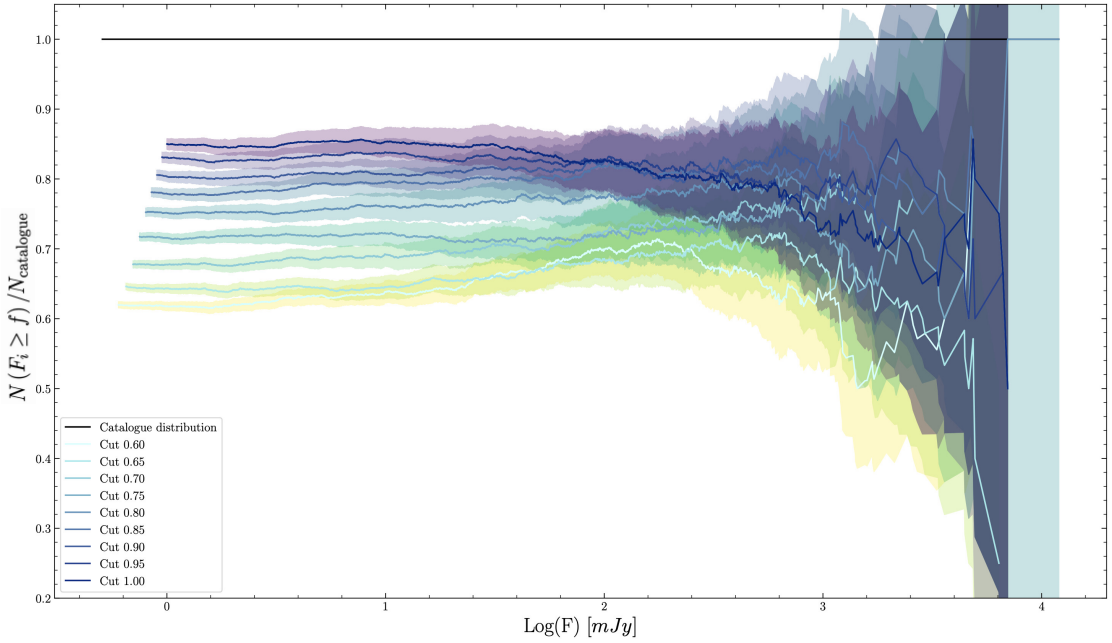


Figure 4.13: The cumulative fraction of targets assigned above the flux threshold value, compared to the flux value for each target. This is normalised by the total number of targets available in the field, and by the complete catalogue distribution of targets (represented by the black line). Each distribution has a selection cut ranging from 1.0 mJy to 0.6 mJy

targets correlates with decreasing flux or magnitude.

I applied a series of selection cuts to each of the 34 proto-fields based on the flux density, ranging from 1.0 to 0.6 mJy. Targets with a flux density greater than or equal to this threshold are included in the assignment process and their priority values are left unaltered. Targets with a flux density below the threshold are excluded from the assignment process by setting their priority value to zero. After applying this cut to each field, it is then processed once by configure using the nominal method.

The completeness distribution can be seen in figure 4.13, where a logarithmic scale has been used along the flux density axis to display the full catalogue range. Decreasing the selection cut threshold results in a larger combination of both high and low priority targets included in the field assignment process. At the high flux

density range of this distribution, the fluctuations are dominated by poisson errors with sharp variations. This is due to a reduction in the number of targets at this end of the scale and *Configure* finding the local minimum in the annealing procedure.

From figure 4.13, it can be seen that the 0.6 mJy cut displays a curved upward trend as a function of flux density. This behaviour is counter-intuitive, as it would be expected that the increase in lower priority targets, which have a lower flux density, would reduce the proportion of fibres assigned to targets with a higher flux density. This implies that there is indeed a flux dependent selection effect that varies with the configuration arrangement, however this should be confirmed using a greater subset of fields due to the effect of the Poisson errors at the higher end of the flux scale. Once verified, this selection affect must be accounted for in future observational analysis and studied across a range of WEAVE science surveys to confirm the behaviour of the configure algorithm.

4.1.5 Overlapping fields

To efficiently tile the whole northern sky, WL will implement a pointing strategy based on the mathematical description given by (Saff and Kuijlaars, 1997). Further details on this tiling can be found in (Jin et al., 2022). Based on this strategy, 94% of WL targets will fall within one field of view on average. The remaining targets will either be visited twice or not at all. This section only considers the mid-tier survey and these proportions are summarised in table 4.1. In this case, an observation is defined as the field of view a target falls within, based on the tiling pattern for the WL mid-tier survey.

The term "included" refers to a target that is allowed to be allocated during the processing of the field in *Configure*, while "assigned" refers to a target that is part of the field allocation process and has successfully had a fibre placed over its location in the field. Based on these definitions, it's possible for a target to be within an

Total number of targets in this subset	47600	%
Targets included in a single field	44891	94.31
Targets included in multiple fields	2709	5.69

Table 4.1: The number of targets included in the allocation process, across the 34 mid-tier WL proto-fields, and the proportions that appear in one or more WEAVE fields of view.

Targets included in one field	44891	%
Unassigned	19881	44.29
Assigned once	25010	55.71
Targets included in two fields	2709	%
Unassigned	329	12.14
Assigned once	1088	40.16
Assigned twice	1292	47.70

Table 4.2: The number of targets included in single or multiple fields for the nominal assignment case, and their respective allocation values.

observation but not included in the fibre allocation process, or for a target to be included in an observation but not allocated a fibre.

To assess the impact of *Configure* on target coverage across the WL mid-tier survey, the assignment (rather than just the visit statistics) must be simulated for the overlapping regions between fields of view. This includes quantifying the proportion of targets assignments in one or more fields. For *Configure*'s nominal usage case the assignment proportions are stated in table 4.2 and visualised in figure 4.14.

In this section the nominal usage case implements the maximised difference priority scheme, which has been adopted by the WL team members. As expected, the majority of the targets included in a single field are assigned (53.57%) however, a significant percentage remain unallocated (46.43%), with each field assigning 840 fibres on average.

When comparing the targets included in two fields, a smaller proportion of fibres are left unassigned. Nevertheless, the difference between the number of targets assigned twice versus once is only 13.7%. This outcome is not surprising given that the overlapping regions are situated near the edges of the field of view. If a target is

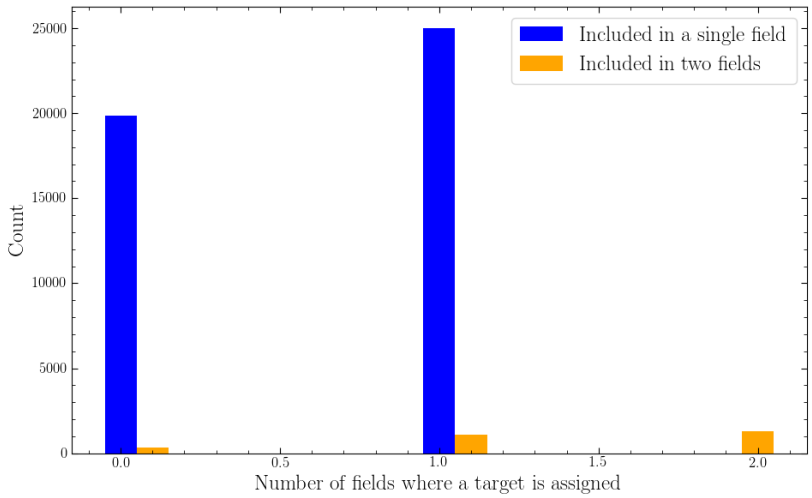


Figure 4.14: The proportion of the mid-tier WL catalogue targets that are included but not assigned a fibre, or are assigned a fibre in one or more fields with the nominal *Configure* processing. This uses the maximised priority labelling scheme.

easily accessible to the fibres in one field, it is also likely to be accessible in the other. However, this depends on the distribution of targets towards the center of the field.

To evaluate the variation in the fibre assignment for the higher priority targets across overlapping areas, I conducted 450 MC simulations on two WL proto-fields that have partially overlapping pointings. Figure 4.15 illustrates the resulting variation in each field, isolated to the overlapping region. The distribution of the fibre assignment is not consistent between the two fields of view, which indicates that targets located towards the centre of the field notably impacts the fibre assignment towards the edges.

Targets included in one field	44891	%
Unassigned	20845	46.43
Assigned once	24046	53.57
Targets included in two fields	2709	%
Unassigned	450	16.61
Assigned once	944	34.85
Assigned twice	1315	48.54

Table 4.3: The number of targets included in one or multiple fields for the multi-survey case, and their respective allocation statistics.

This work was expanded to assess the impact of the multi-survey processing on

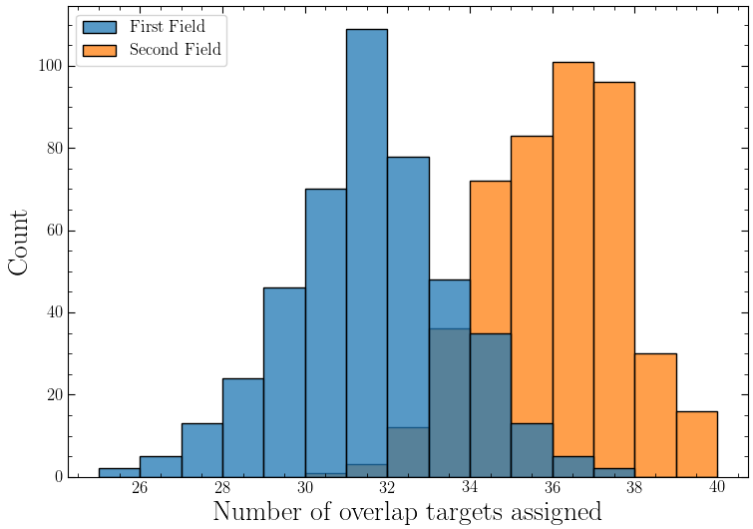


Figure 4.15: MC realisations for the number of high priority targets assigned for two proto-fields that share an overlapping region of the field. Each proto-field has been processed by *Configure* 450 times using the nominal method with the maximised difference in the priority weighting scheme.

these statistics, which are summarised in table 4.3 and visualised in figure 4.16.

For targets included in a single field of view, the total assignment only varies by $\sim 1\%$ compared to the nominal usage case with the maximised priority scheme. When considering targets included in two fields, there is a 4% increase in the number of targets which are unassigned using the multi-survey method, and a 1% increase in the number of targets assigned twice. This comparison is illustrated in figure 4.16 for the high and low target priorities. Figure 4.16 indicates that the multi-survey method assigns a greater number of higher priority targets twice in comparison to the nominal method, which is reflected in the lower priority assignment across both processing cases. These results are in agreement with section 4.1.3; that the multi-survey method allows for additional flexibility compared to the nominal case, at the expense of the total fibre assignment.

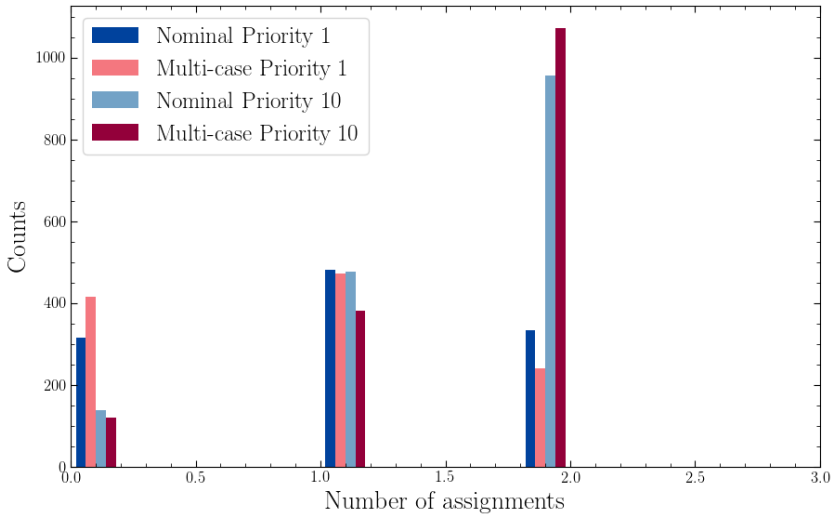


Figure 4.16: The target assignment distribution across the 34 mid-tier WL survey fields. This is separated by priority for both the nominal assignment, and the multi-survey assignment cases. Only targets that are included in at least one observational field are shown.

4.2 Conclusion

In light of the results obtained, it can be concluded that the *Configure* algorithm represents an effective approach for assigning the WEAVE fibre multiplex to targets in the field of view. However, the algorithm presents certain challenges with respect to its detailed control, making it difficult to achieve specific outcomes in the fibre assignment.

This work has explored a subset of processing methods which does not modify the *Configure* algorithm and can be easily implemented by the WEAVE survey teams. These methods return consistent results that provide the survey teams with an extra degree of freedom to prioritise aspects of the survey outcome. The allocation of each field, for shared regions in particular, is highly subjective to the configuration setup for each scientific survey.

In the context of WEAVE-LOFAR, the main conclusion of this work is that I have identified a potential underlying flaw in the adopted weighting scheme, which would lead to a systematic undersampling of fainter sources within the survey. This requires

further investigation.

As a result, this chapter represents a starting point in the detailed study of *Configure*, which has the capacity for further developments to optimise the scientific output from WEAVE. In particular, this work should be expanded after the science verification period for WEAVE is completed to highlight potential selection effects for all planned surveys. These findings provide important insights into the factors that influence fibre assignment accuracy and highlight the need for further research to optimise the design of fibre assignment algorithms for astronomical surveys.

Chapter 5

Commissioning and First-light

All hail the spectrograph

Remko

In May 2022, the combined WEAVE Rotator System (WRS), fibre positioner, and optical fibres were successfully mounted to the William Herschel Telescope (WHT) and integrated with the spectrograph system. The new top-end structure is shown in figure 5.1[1]. This achievement was shortly followed by the commissioning period, where the combined performance of WEAVE's sub-systems is tested and evaluated. Chapter 5 describes this phase of WEAVE, which officially began in July 2022 and is currently ongoing. This includes evaluating the image quality, guiding system, and astrometry development. Here, I present the first-light data taken for the LIFU observing mode.

5.1 Astrometry

Astrometry is precise measurement of the positions of stars, and so in the case of instrument commissioning, astrometry is the determination of the mapping of the

[1]The final covers are still pending

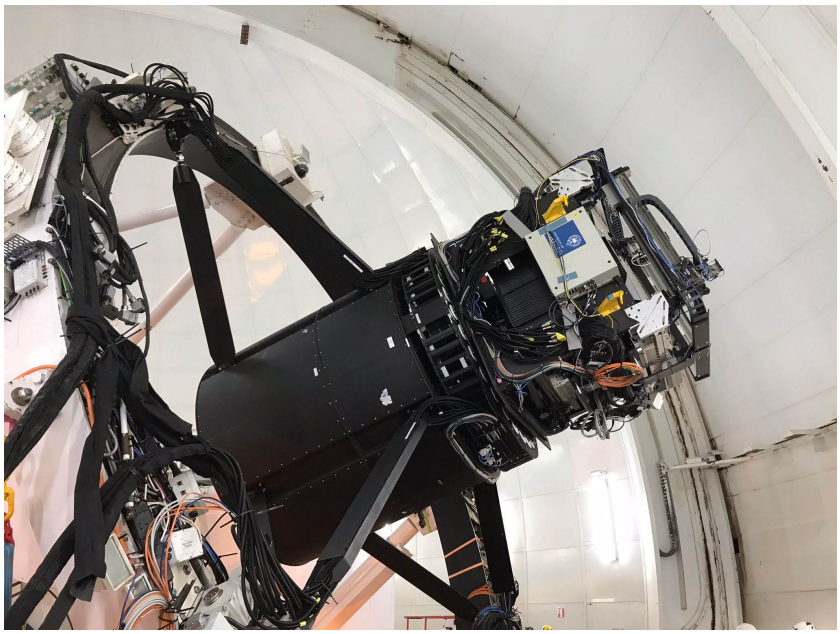


Figure 5.1: The assembled top end of WEAVE on the William Herschel Telescope in La Palma, Spain. This photo was taken on the 25th of May 2022.

coordinate reference frames of the instrument to the sky. When commissioning new astronomical instruments, it is crucial to establish that the astrometry of the system is accurate. These measurements are used to align and model the pointing of the telescope, calibrate the optical alignment, and track and guide the telescope during observations. By carefully calibrating the astrometry, astronomers can obtain high-quality and reliable data that can be used to study the properties and dynamics of the universe.

5.1.1 Image quality

In order to obtain the best astrometric calibrations, it is essential that the system is delivering the best image quality. To ensure that the corrector did indeed deliver the as-designed image quality, an imager was used behind the WEAVE prime-focus corrector (PFC) after its integration in early 2021 at the WHT. The on-axis images displayed round point sources with a Full-Width Half Maximum (FWHM) consistent with an external Differential Image Motion Monitor (DIMM).

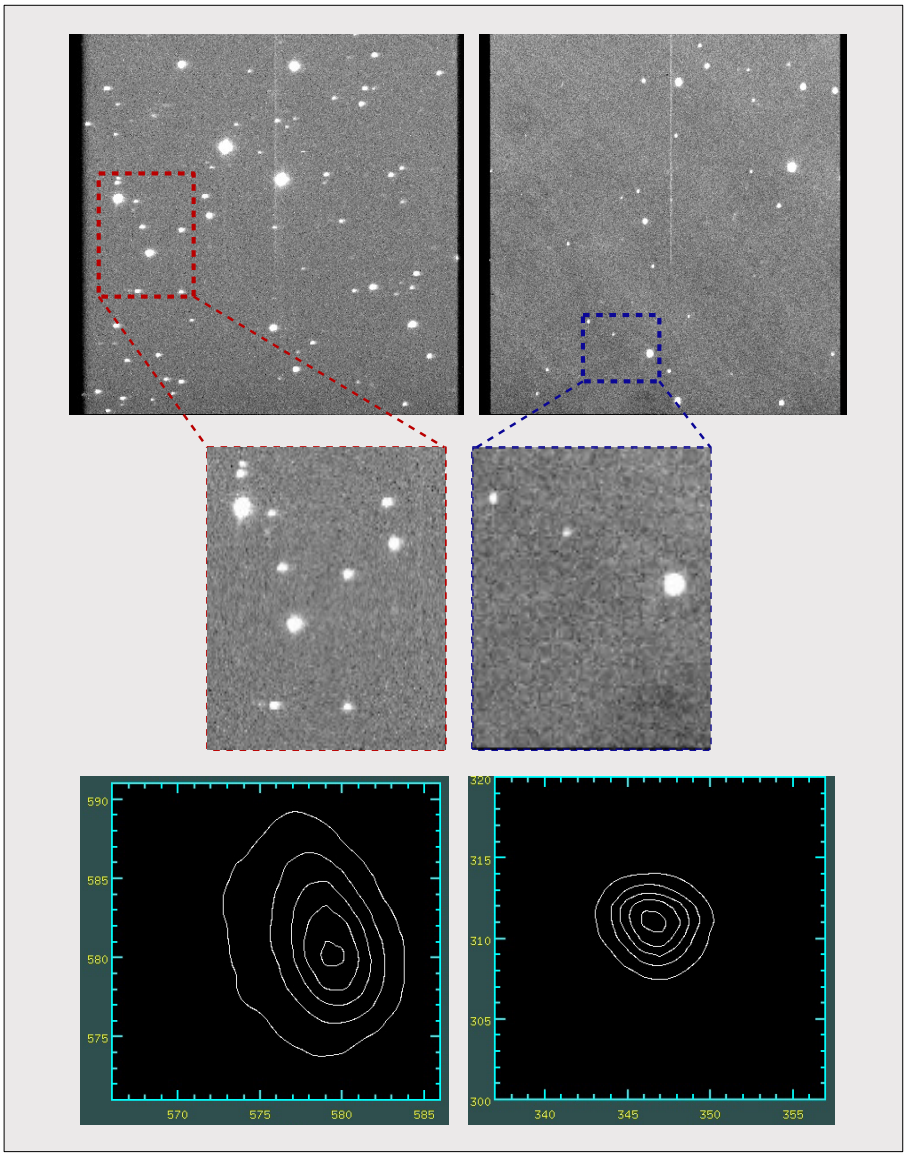


Figure 5.2: (*Left, top*) The elongated images taken by the WHT following the new top-end installation on the 12/09/2022. (Right, top) Images taken by the WHT on the 12/09/2022 after a 65° rotation was applied to one of the ADC prism lenses. It can be seen that this rotation removes the elongation in the point sources caused by chromatic aberration. (*Bottom*) Contour plots of a carbon star before and after the ADC lens was rotated to remove the elongation.

In July 2022, following the mounting of the completed top-end structure, repeated observations indicated that these images were now elongated. The result for off-axis point sources, such as stars, is that they appear to have a comet-like tail. The

elongated images initially produced by the new-top-end can be seen in figure 5.2. It is important to note that the elongated images taken by the FPI would not have been identified if the suggestion of incorporating a filter had gone ahead.

Initially, it was thought that this elongation was caused by Coma (Lynden-Bell, 2002). In optics, this is defined as a variation in magnification over the entrance pupil and is caused by imperfections in the optical system, which implied a misalignment between the top-end relative to the optical axis. However, this explanation could not reconcile the out of focus images with the image quality seen prior to the mounting of the positioner-WRS systems.

The solution was identified by observing extremely red carbon stars, which resulted in images with round point sources, as the effect of atmospheric dispersion increases for blue wavelength ranges. Eventually, it was discovered that manually changing the rotation of the Atmospheric Dispersion Corrector's (ADC's) secondary prism reduced the magnitude of the elongation. This meant that the elongated images were a result of the atmospheric dispersion.

In the period the between the first imaging tests and the top-end being mounted, some maintenance procedures were completed on the ADC. This resulted in the prisms becoming un-initialised and it was found that the secondary prism was rotated by an additional 65° from its home position. Testing the initialisation of the ADC has now been incorporated as part of the WHT's routine maintenance procedures, to prevent this affect from occurring during WEAVE's survey operations.

5.1.2 Telescope guiding and tracking

To obtain astronomical images with the best signal-to-noise ratio possible, astronomers take images over long exposure times. However, as celestial objects appear to move across the night sky due to Earth's rotation, telescopes must track changes in their position to keep them fixed within the field of view. To achieve this, telescopes use

a combination of tracking the telescope axes and a continuous adjustment known as autoguiding. Tracking refers to the telescope's pointing and movement with the target, it should correct for errors such as those caused by the periodic drive gears. Guiding is a technique used to improve the accuracy of the tracking by removing the telescope drift and image motion. This section will describe the MOS autoguiding system and the calibration work completed on it during the commissioning period.

The MOS guiding system consists of two sets of image bundles, one per plate, each with 8 coherent guide fibres, which was also used for AUTOFIB-2 (AF2) on the WHT (Parry et al., 1994). It is possible to guide using only two fibres in the field of view, however extra fibres were included for contingency and to improve the accuracy of the autoguiding adjustments. With a field of view of each guide fibre measuring approximately 4.5", the auto-guiding system is designed to mitigate the bulk effect of differential diffraction by tracing the stars' location across the imaging bundle. The MOS auto-guider has two image bundles of 8 guide fibres for MOS A and B, that rest on a mechanical slide. Whilst one field is observing, the other bundle is back-illuminated to enable the guide fibres to be positioned for the following observation. The LIFU has its own dedicated autoguider, as there is no way to deploy guide fibres at the LIFU's location. The mIFUs use the MOS B guide fibres for their observations, and the mIFU/MOS/guide fibres are designed to be co-focal.

The autoguider software works by tracking the position of the point source within each guide fibre relative to the centre of the field of view. As the telescope pointing changes to track the centre of the field, the positions of the point sources will move relative to the fibre. These movements are traced and used to predict the next adjustment, as differential refraction produces a shear of the field with more distortion at the lower elevation side of the plate. This can be corrected to first order by a rotation. Therefore, it is vital that they are co-focal and the orientation of each guide fibre is determined to ensure that an accurate pointing correction can be applied.

It was not possible to adequately determine the arrangement and focus of the MOS autoguider until the WEAVE commissioning period began. The focus of the bundle was set by making small, manual adjustments to the bundle's position relative to the guide slide, following a series of exposures taken whilst the guide fibres were deployed. However, setting the focus does not provide any information on the orientation of each fibre, which must be measured to refine the telescope tracking. It is important to note that each guide fibre is rotated within its image bundle; the bundle may then be rotated with respect to the autoguider CCD; and each guide fibre will be rotated according to its position on the field plate.

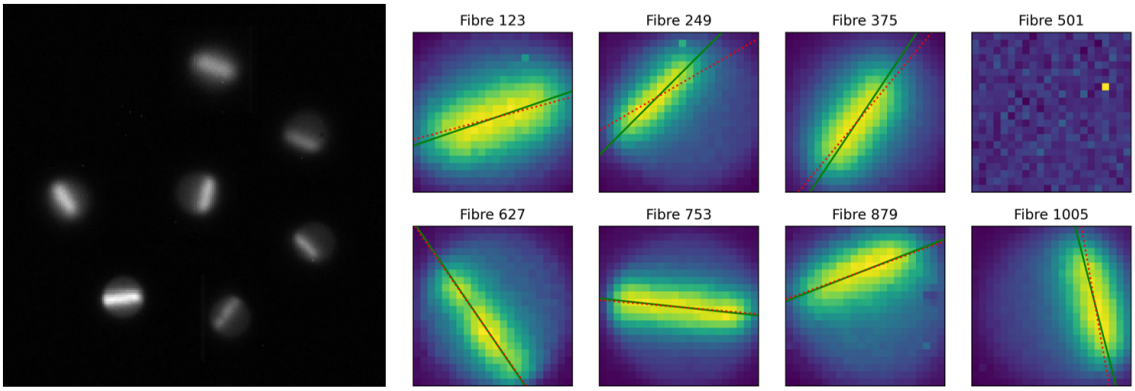


Figure 5.3: Example images from the autoguiding system, where exposures were taken as a group of stars were slowly moved across the autoguider's field of view (*right*). This image used $V \sim 15$ magnitude stars with an exposure time of 25 seconds, whilst 5° away from the full moon. This exposure time produced long lines across each fibre, allowing the orientation of each guide fibre within the bundle to be calculated (*left*). In these images, seven out of eight guide stars were acquired, with the exception being fibre 501 which is disabled (Grafton et al in prep) [image credit to Jure Skvarc and Emmanuel Grafton; used with permission].

Figure 5.3 presents the observations used to determine the position and orientation of each guide fibre within its image bundle. These images were produced using relay-optics, which consisted of two back-to-back Canon lenses, which are needed to focus the light from the fibres onto the autoguider CCD. In reasonable seeing conditions (< 2") configurations that use only the guide fibres observed isolated stars with a

magnitude of appropriately $V \sim 15$. With a single star appearing inside each guide fibre, the telescope pointing is offset in right ascension by 3" as a 30 second exposure is taken, so that stars move slowly across the guide fibres. This image is displayed on the left-hand side of figure 5.3. The aim is to produce streaks of light across the CCD, which are then used to calculate the fibres orientation, shown in the right-hand panel. This method was combined with additional on-sky observations, where a series of shifts were applied to the guide star's position to observe its relative movement within each bundle. This was then used to determine the exact rotation angle of each fibre.

The orientations of each fibre have been incorporated into the observing software by Jure Skvarc and Emmanuel Gafton (Grafton et al in prep). The software is designed to adjust the telescope pointing during observations by tracking the guide fibre targets, to maintain the MOS science target positions within the fibres field of view[2].

5.1.3 Distortion maps

To ensure that the astrometry of WEAVE is correct, the positions of objects in the night-sky must be mapped to the focal plane. This is not a simple one-to-one mapping, as all astronomical objects lie on the inside of the celestial sphere, which must be the projected onto the 2D focal plane. This is known as a tangent plane projection, and can be thought of as the inverse of a globe to a 2D map. A diagram of this projection can be seen in figure 5.4. The result is a shift in the objects position in the focal plane, which this section will refer to as the distortion map, that is exclusive to the optical system used in observations.

This process begins by calculating the position of each star, for a given observation, on the positioner's field plate using the default focal plane map. This is referred to

[2]The autoguiding software development will likely be written up for publication at some point in the near future.

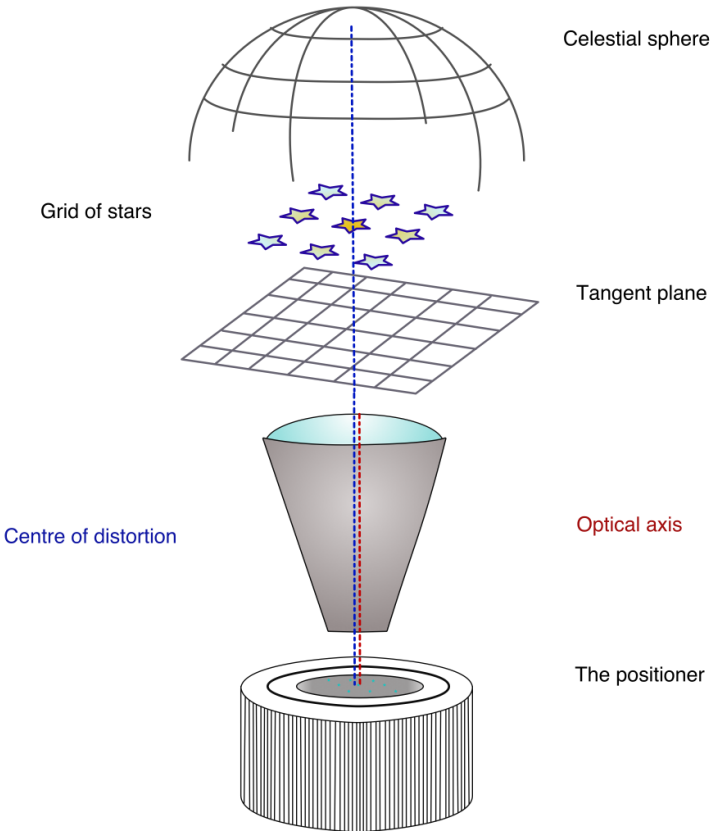


Figure 5.4: A diagram of the optical distortion that affects celestial objects as they are observed across a 2D focal plane.

as the distortion map. A field with 320 bright stars was generated using stellar objects taken from the *Gaia* catalogue. Isolated stars with comparable magnitudes across WEAVE's field of view were chosen, as they produce consistent images over a constant exposure time. The majority of these stars are from the outer regions of the Galactic plane.

The position of the centre of the field in the focal plane and the centre of rotation, referred to as x_0, y_0, and r_0 respectively, are not related to the distortion map and are determined independently. As the distortion map is radial, to first order these values do not impact the coefficients determined by the astrometry runs. The distortion map for WEAVE is modelled as an odd 7th order radial polynomial $f(r)$, where

$$f(r) = ar + br^3 + cr^5 + dr^7. \tag{5.1}$$

At the time of this thesis, the coefficients a, b, c, and d have been set to default values of 202.96, 3.019, and -0.01 for plate A and 202.99, 3.036, and -0.088 for plate B to two significant figures. This model was produced by David Terrett, based on the output from the ray-tracing software *Zemax*. Once on-sky, these coefficients were found to be adequate for the MOS field acquisition process with an RMS of 0.8", however further investigation is needed to determine the best fit of the optical distortion across the focal plane.

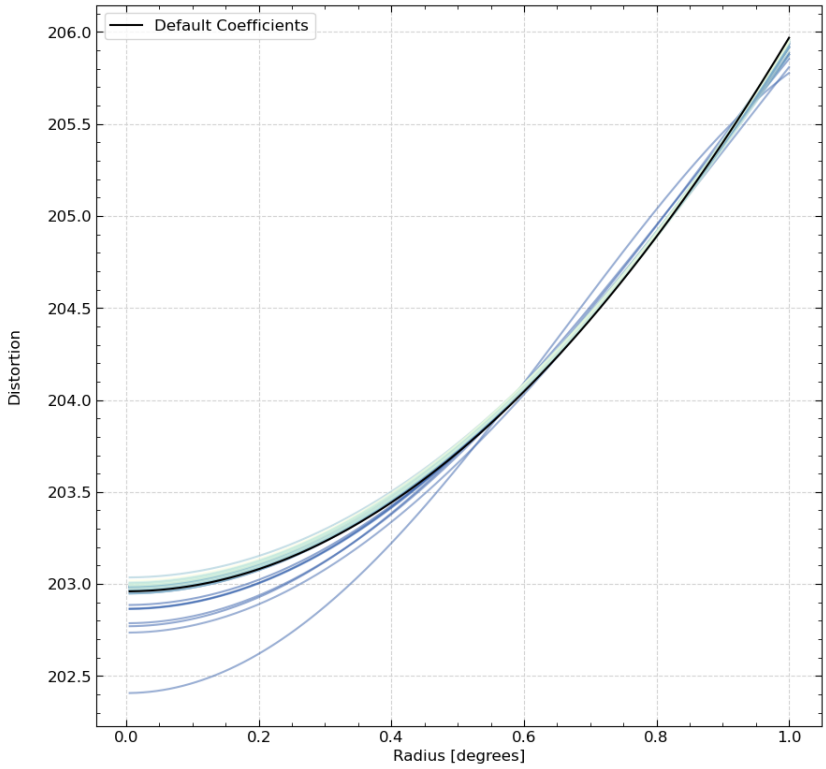


Figure 5.5: The distortion polynomial coefficients measured for a series of astrometry runs in November 2022. The black line displays the default coefficients that the distortion map uses. Each blue line represents the polynomial coefficients for a single astrometry run, which takes approximately 45 minutes to complete.

5.1.4 Astrometric measurements

The next phase requires the Focal Plane Imager (FPI) to image each star in the field, and the difference between its expected and actual position in the focal plane

is calculated using the default distortion map. This field is observed by the WHT in favourable weather conditions, with good seeing ($\sim$ 0.7") and low humidity. When taking these images, it is only possible for the telescope to track, not guide, as the FPI moves around the focal plane. The sequence of astrometry measurements for every star in this field is referred to as one "run". It should be noted that each run takes approximately 45 minutes to complete with only 160 stars imaged, which causes the telescope elevation and rotation to progress by approximately 10°. Given this length of time, the positions of the stars are recalculated for the actual time and field position at which the FPI observes them. Figure 5.5 displays the variation in the coefficients measured for a series of astrometry runs taken over the course of a few nights.

A fit is then performed using an affine transformation to determine the coefficients of the distortion map on-sky. The divergence between the anticipated and observed positions of the stars in each field are represented as vector plots, as demonstrated in figure 5.6. The top-left displays the raw differences prior to any additional processing, where the expected positions have been calculated using the default distortion map. This stage is used to remove stars from the observation that introduce additional noise to the astrometry. Typically, this is due to the presence of additional stars in the FPI image, which affects the centroiding software. The positions of the stars are measured using a similar centroiding program used for the grid spots; a box is defined around the stars expected position, it then shifts its position around the images highest pixel intensity and centres a smaller box around the new position. The presence of brighter nearby stars can easily shift this measurement, therefore the field is filtered to exclude these objects from the measurement process. Typically, this is only necessary for the first astrometry run for a newly generated observation field.

The astrometry run is then repeated following the fields refinement and these measurements are processed further, represented by the additional panels in figure

5.6. The raw measurements following the exclusion of problematic images is shown in the top-right panel. Any shift in the centre of the field (x_0, y_0) is calculated and applied as a correction to each measurement. This is shown in the top-left panel. Coefficients for a new distortion map are then calculated using the Standards of Fundamental Astronomy (SOFA) software (IAU SOFA Board, 2021). This software first calculates the standard coordinates of each star, which is their position relative to the field centre as an angle. This is converted into a radial distance in mm across the focal plane before a fit of the 7th order odd polynomial is performed.

The new distortion map is then applied to the measurements with the shifted field centre to provide a final root-mean squared error value (RMS), displayed in the bottom-right panel of figure 5.6. The WEAVE specifications state that this must be < 0.1" to achieve the desired scientific output. The initial astrometry runs completed in October 2022 gave an RMS on the order of 0.8". This is clearly not acceptable value. To reduce the magnitude of the RMS, further investigation is required.

It is important to state that for a configuration with only six guide fibres deployed, this "tweak" method resulted in stars appearing within the centre of each guide fibre on the first acquisition attempt without any further adjustments. The exception to this was a slight offset seen by the guide fibre placed at the centre of the field, however the acquisition of the guide fibres towards the edge of the field is a notable achievement.

It was suspected that the main contribution to the RMS was a result of a drift in the telescope tracking, which requires an offset to correct. The shift in position of the centre of the field was studied by making a repeated measurement of the central star throughout the astrometry run. It was discovered that the telescope was making small, sudden tracking shifts that appeared to be slightly periodic. To counteract this affect, a "tweak" method was introduced to the astrometry run which significantly improved the measured RMS values. The FPI moved to measure three stars,

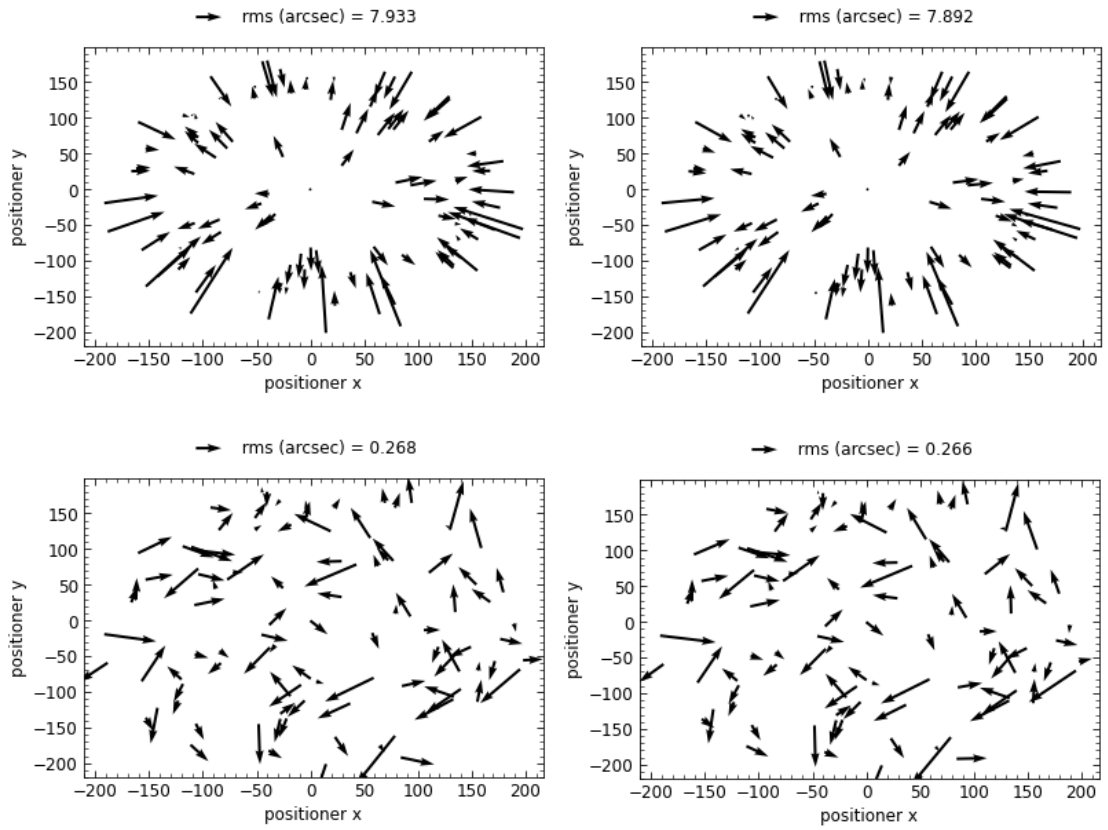


Figure 5.6: The residuals between the expected and measured positions of each star in the astrometry calibration field, measured using the Focal Plane Imager (FPI) in xy coordinates. (*Top-left*) The unprocessed differences between each stars measured and expected position on-sky. (*Top-right*) The unprocessed differences between each stars measured and expected position following the exclusion of problematic targets and images from the field. (*Bottom-left*) The difference between the measured and expected positions after the field plate position x_0, y_0 has been calculated and applied as a shift to the measured coordinates. (*Bottom-right*) The difference between the measured and expected positions after applying the field plate shift and fitting a new distortion map to the data.

ordered randomly across the field, before moving back to the centre of the field and re-measuring the central stars position. The measured shift was then applied to the telescope pointing before the measurement process continued. However, incorporating this "tweak" removed our ability to trace the shift in the telescope's pointing. The RMS of the astrometry runs reduced to approximately 0.3" in February 2023 using this method, shown in figure 5.6. Whilst it has improved, additional effort is required

to improve the precision of the astrometry. As a collection of astrometry runs has now been completed in favourable weather conditions, there are ongoing discussions as whether stacking the measurements collected would improve measurements of the distortion map coefficients.

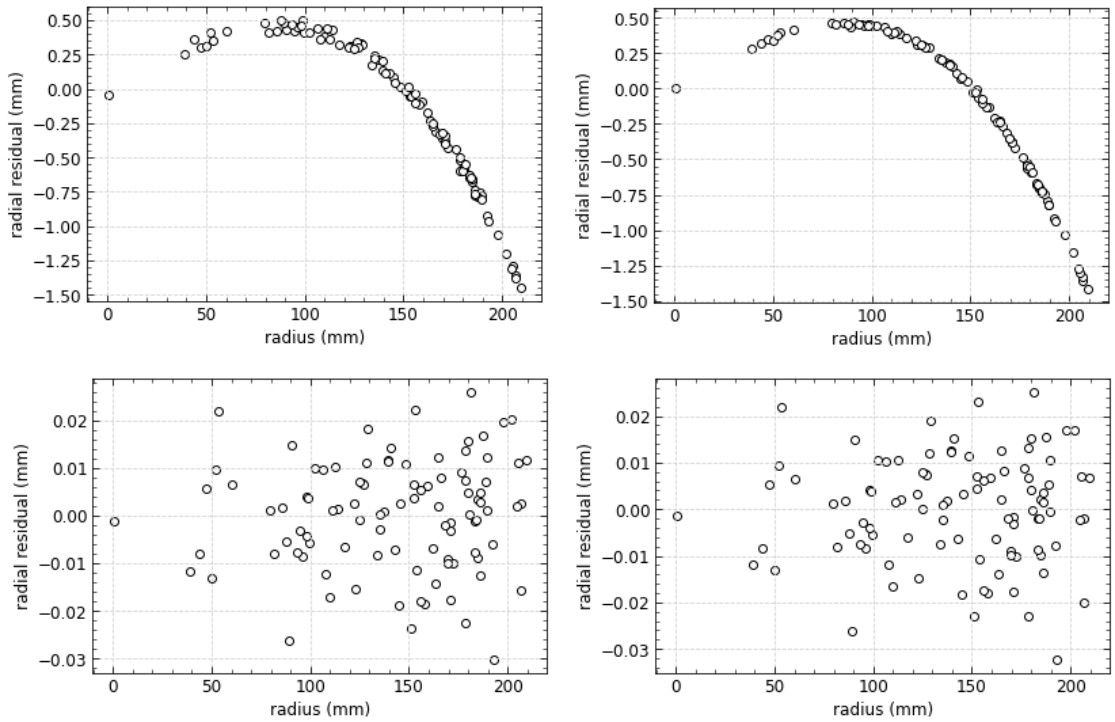


Figure 5.7: The residuals between the expected and measured astrometry targets, in radial coordinates, as a function of the targets position in the focal plane. (*Top-left*) The raw unprocessed residuals. (*Top-right*) The residuals after applying a pointing correction to x_0, y_0. (*Bottom-left*) The residuals after subtracting the polynomial fit from the default distortion map. (*Bottom-right*) The residuals after subtracting the polynomial fit from the newly measured distortion map.

Figure 5.7 presents the data from figure 5.6 in the radial coordinates. The top-right panel displays the unprocessed radial residuals, whilst the top-left panel displays the residuals following a correction applied from the measured x_0, y_0 position. Despite this correction, the plot demonstrates that a 50 mm change in radius corresponds to 1.25mm of distortion. Scaling this down demonstrates that a 0.5 mm error in the x_0, y_0 position corresponds to a 12 μm shift in the distortion, which exceeds the entire

allowed POS misalignment contribution stated in table 2.1. Therefore, to meet the misalignment budget, the shift in x_0, y_0 should not exceed 100 μm, which may not be possible by observing a field of stars due to the changing elevation and rotation position of the field itself.

A proposed solution is to predict the shift in the distortion map centre (x_0, y_0) as a function of the telescopes orientation and compare it to the shift in the centre of the field plate (x, y) as measured by the FPI grid scans. This should reduce the RMS of the astrometry runs but may not be the only contributing effect. To make this comparison, a series of grid tests and astrometry runs are being conducted at varying elevation and rotation angles. This enables the movement of the field centre to be characterised across the WHT's pointing range using the grid scans, which can be performed as a daytime activity. A model of (x_0, y_0) as a function of elevation and rotation angle will then incorporated into the positioner software and the *Configure* program to correct the positions of the targets in the field. The result is a configuration where the fibres are placed based on the expected location during the observation itself, rather than the telescopes position as the fibres are being placed in the configuring orientation.

Figure 5.8 displays the shift in the distortion maps x_0, y_0 and the field plate centre x, y for the for a series of telescope rotation and elevation angles. There rotation and elevation angles in this figure represent the mean between the start and end of each observation, with a 5 minute grid scan performed at the start of each run. There appears to be a clear correlation in the direction of their shift for both sets of measurements across the full orientation range, however there is a notable difference in magnitude. This is to be expected, as the x_0, y_0 corresponds to the position of the optical axis, which is subjected to a greater degree of flexure in comparison to the FPI relative to the focal plane. Further investigation is needed to produce a complete predictive map, although these initial results appear promising. The primary concern

at this stage is the repeatability of these measurements, which will determine whether it is possible to implement these corrections as a map within the positioner's software.

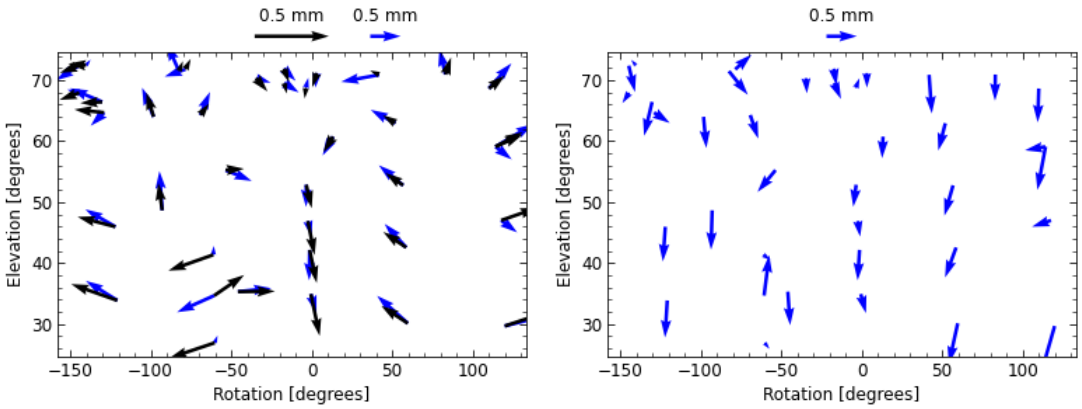


Figure 5.8: (*Left*) The shift in the centre of distortion x_0, y_0 (blue) and the field plate centre x, y (black) from the mean, measured by the Focal Plane Imager, as a function of elevation and rotation angle. (*Right*) The shift in x_0, y_0 position after being rotated into a common reference frame.

At this stage, the astrometry still remains dominated by corrections to the telescope pointing. The main contribution is from the shift in x_0, y_0 over the course of the 45 minutes the astrometry run takes to complete. This is exacerbated by the WRS, which has been found to shift in absolute position by approximately 1°. This magnitude is enough to prevent the observers from acquiring the MOS fields correctly by introducing changes to r_0. Recently, its calibration procedures have been modified to initialise the system each night prior to observations, which should minimise the change in its absolute position. This is an ongoing issue that the engineering team at the observatory are actively resolving and it is expected that the telescope pointing will improve over the duration of the commissioning period. At this stage, refining the astrometry measurements further will not make a significant impact until an accurate model of x_0, y_0, produced using the grid scans, has been implemented in the positioners software.

5.2 Fibre mapping

As spectra are being obtained through the commissioning observations, they must be mapped to the appropriate fibre that produced them. These mappings form a crucial component of both the Core Processing System (CPS) and the Science Processing and Archiving (SPA). Whilst the location of the MOS fibres within the positioner are well documented, this is not the case at the interface to the spectrograph. Only the ordering of the slitlets and the groups of fibres which belong to them are documented for the MOS and mIFUs. The LIFU is treated independently, as a fibre map was generated as part of its assembly process, which must be confirmed on-sky via image reconstruction.

5.2.1 Slitlet mapping

As mentioned in section 2.4.2, 24 configurations were created for MOS A and B to determine each fibres position along the spectrograph detectors. An example of the raw spectra for a single slitmap configuration with the red-arm of the spectrograph can be seen in figure 5.9. Each bright horizontal line represents the spectrum transmitted by a single optical fibre. As the ordering of the slitlets along the detector is documented, these images have been processed by the CASU team to identify which fibre ID the spectra belong to. The deployed fibres are illuminated using the WHT's calibration unit.

The complete set of slitmaps exposures have been taken for both MOS A and MOS B, and have been incorporated into the pipeline to define the fibre trace and their position along the detector CCDs. This may be subject to change, due to the improvement of the alignment of the spectrograph. Recent observations have revealed a defocus effect that varies with the fibres position along the slit. To identify the cause of the defocus, the position of the cameras will be jacked on a small scale to adjust the

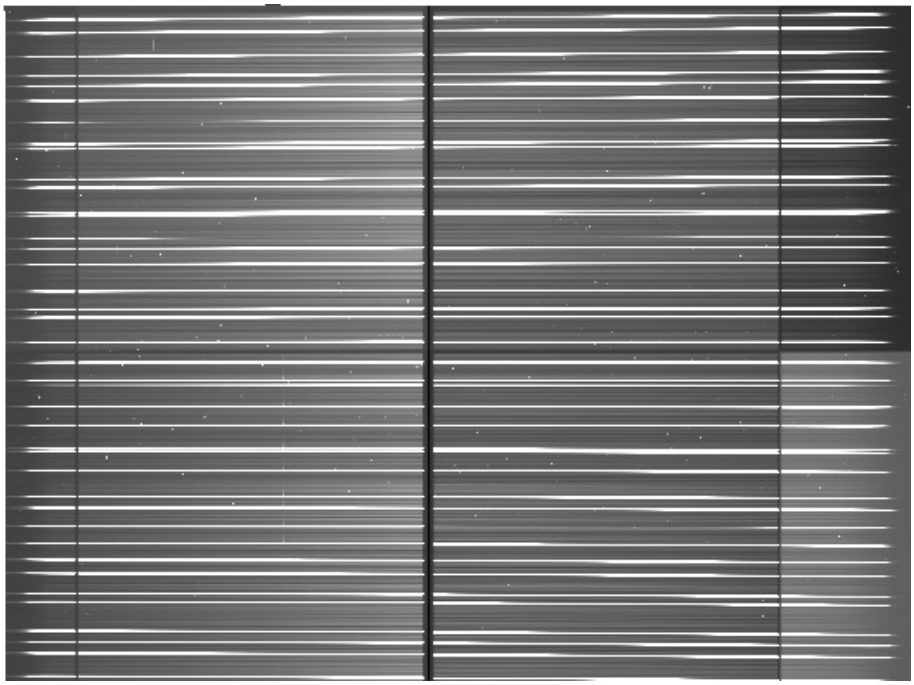


Figure 5.9: The raw spectral data taken for a single slitmap configuration on MOS A. This image was taken using the red-arm of the spectrograph with a one second exposure. The calibration unit illuminated the fibres using the quartz halogen lamp.

spatial focus of the fibres with respect to the detector and note any changes. This is the easiest form of adjustment which can currently be made, however further testing is required on the collimator alignment, which was unable to be completed prior to the spectrographs transport to the observatory.

5.2.2 IFU mapping

The commissioning period began with the LIFU observing mode, as its positioning requirements are significantly simpler than the MOS and mIFU modes, and it has been designed to conduct observations in poorer seeing conditions. Unlike the MOS modes, the fibre locations of the IFUs across the positioners field of view are somewhat unknown. For the mIFUs, only the fibre ID within each bundle is recorded, not its exact location. The Large Integral Field unit (LIFU) has a preliminary fibre map for its aperture which was generated during its assembly and must be tested. This is

displayed in figure 5.10.

The LIFU was assembled by Ellen Schallig and members of NOVA in the Netherlands. It is composed of 547 carefully stacked optical fibres that form the LIFU head assembly. These fibres were arranged in layers to prevent spectral cross-talk. When assembling the LIFU head, neighboring fibres along the slit were intentionally placed far apart from each other to avoid interference. The map of the LIFUs optical fibres can be seen in figure 5.10.

To conduct observations with the LIFU, its position within the focal plane and with respect to its autoguider must be measured on-sky. In principle, the centre of the LIFU autoguider should be -74 mm away in the y-direction. To begin the mapping process, a bright star is observed at the centre of the LIFU autoguider CCD chip before taking an image to record its position. The telescope is then offset and defocused slightly, to convert the point source into a "doughnut" like shape in the exposure. This is shown in figure 5.11. These doughnut images were used to confirm that the fibre mapping of the LIFU head in figure 5.10 was correct, by testing the image reconstruction of the processing software. Some fibres appear as 'outliers' in the footprint due to a shift in their position as the LIFU head was assembled by-hand.

To confirm the LIFU head fibre mapping across the entire aperture, observations of White Dwarf (WD) stars were taken as the telescope pointing was shifted to produce a trailing effect. This technique was repeated multiple times across the width and height of the LIFU footprint on-sky. To begin this process, the fibre ID of the stars location within the LIFU footprint was determined by visual inspection. A series of offsets were applied to the telescope pointing to shift the stars position to align with the central fibre of the LIFU head. The distance from fibre centre to centre in the LIFU aperture corresponds to 3.38" on-sky. The positioner tumbler was then rotated to the MOS observing mode, and a scan of the field was taken before the FPI was driven to (0, 20) mm in the focal plane. An image of a bright star was then taken with

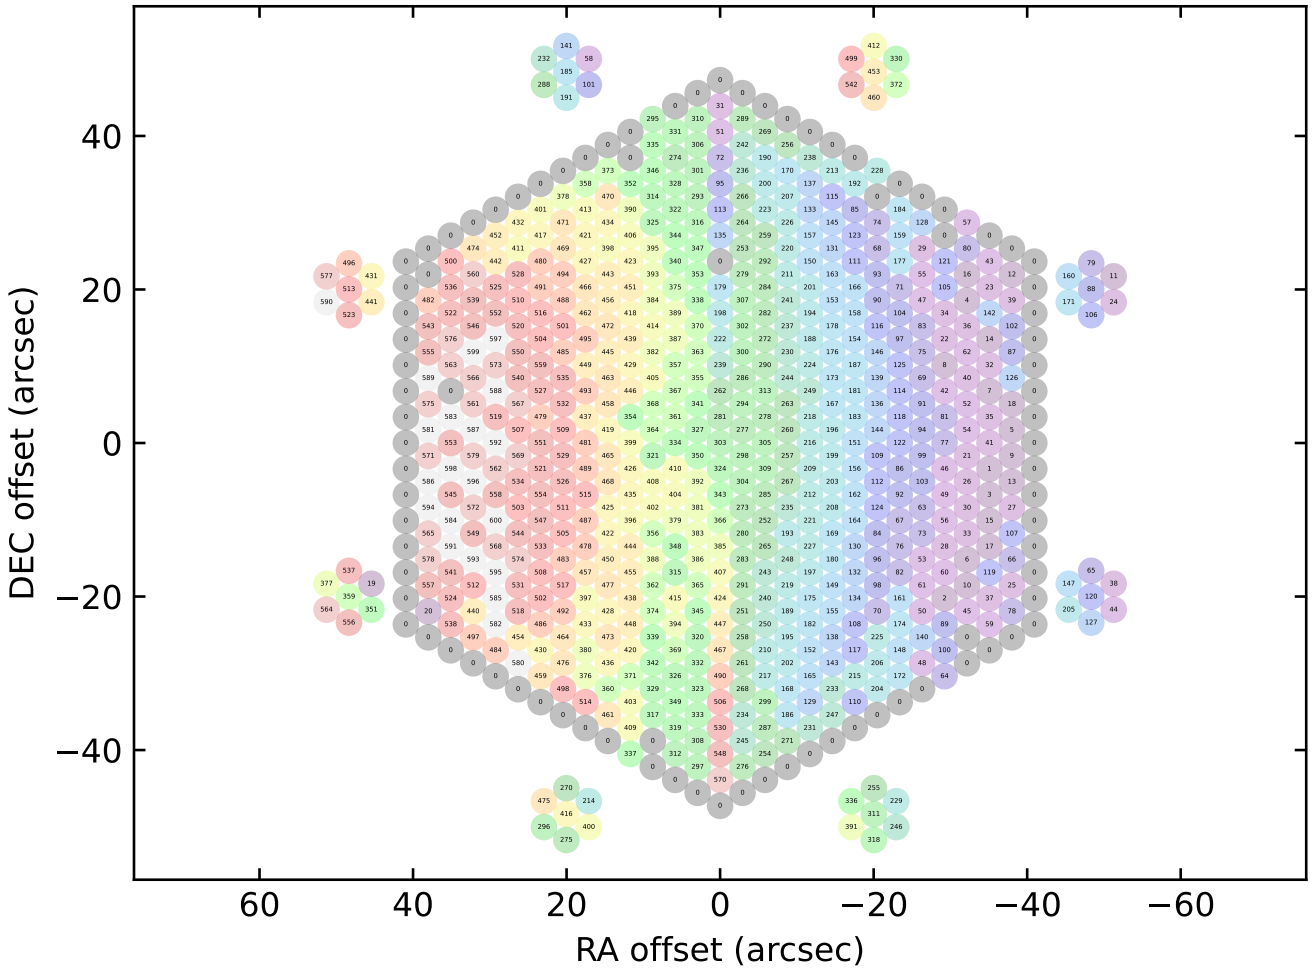


Figure 5.10: A diagram of the fibre mapping within the LIFU footprint. Each circle indicates one fibre, which is labelled with its ID number and coloured according to its position along the LIFU slit. This mapping was created to reduce spectral cross-talk along the detector [image credit Scott Trager; used with permission].

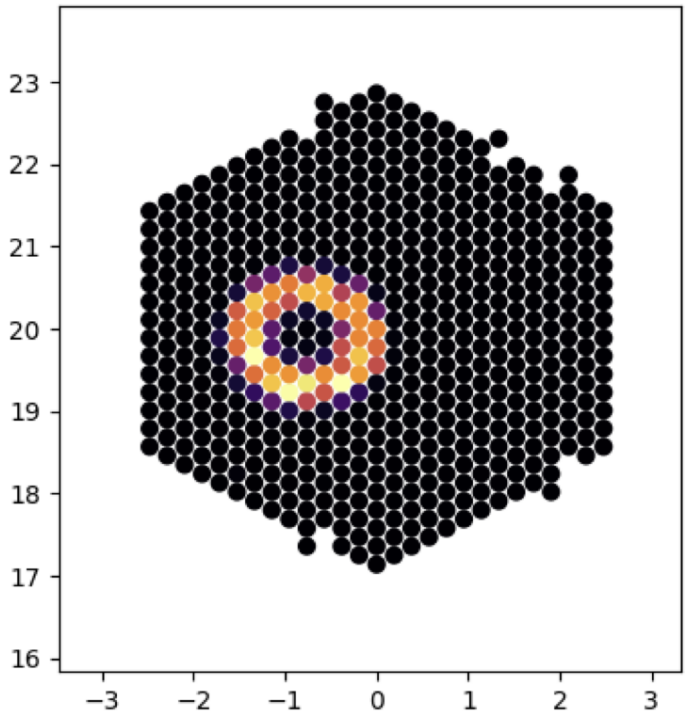


Figure 5.11: An image of a $V \sim 8$ magnitude star taken using the LIFU aperture. The telescope has been intentionally defocused to produce a "doughnut" like shape that can be used to test the fibre mapping of the LIFU head [image credit to Scott Trager; used with permission].

the FPI as a reference point, and the tumbler was rotated by 90° back to the LIFU observing mode. This image was taken to identify possible shifts in the pointing as a result of moving the tumbler between plate A or plate B from the LIFU position.

With the WD star aligned to the central fibre of the LIFU, the boundaries of the aperture on-sky were then defined. The telescope was first moved north by 94 mm in plate coordinates. This should place the star as close to the centre of the autoguider as possible. After determining its position in pixels on the autoguider chip, the star can then be moved back to the central fibre. The telescope is then moved iteratively to the furthest corners of the LIFU head before re-centering the star. These measurements can then be used to determine the location of the LIFU aperture within the focal plane and with respect to the autoguider.

To test the reconstruction of the LIFU, the first galaxy observed with the LIFU is

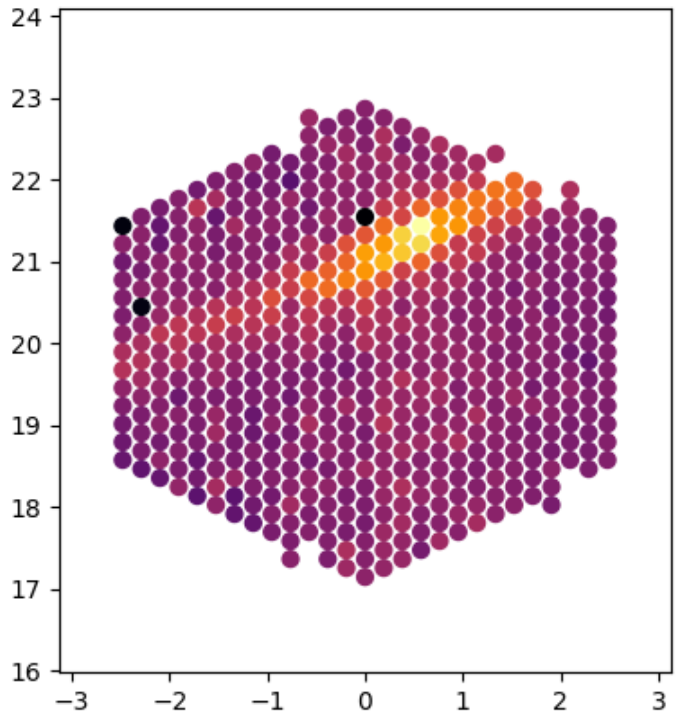


Figure 5.12: A reconstructed image of a LIFU observation of the super-thin galaxy UGC 711 [image credit to Scott Trager; used with permission].

shown in figure 5.12. The galaxy UGC711 was chosen for its long, thin outline, which was used as a visual aid when testing the LIFU reconstruction, as its length spans the width of the aperture. Once this was confirmed, further acquisitions of bright stars were observed at various dither patterns with predicted guide star positions to determine the precise x_0, y_0, r_0 values for the LIFU 'plate'. The majority of the LIFU commissioning tasks have now been completed, and it is expected that the LIFU survey operations will begin in the coming weeks. However, a small number of these commissioning tasks may need to be repeated as a result of adjustments to the spectrograph. At this stage, the remaining commissioning tests for the LIFU involve characterising the focus variation of the autoguider as a function of the telescopes on-sky position angle.

The mini Integral field Units (mIFUs) follow a similar procedure to the LIFU in terms of the fibre mapping through image reconstruction. It should be noted that

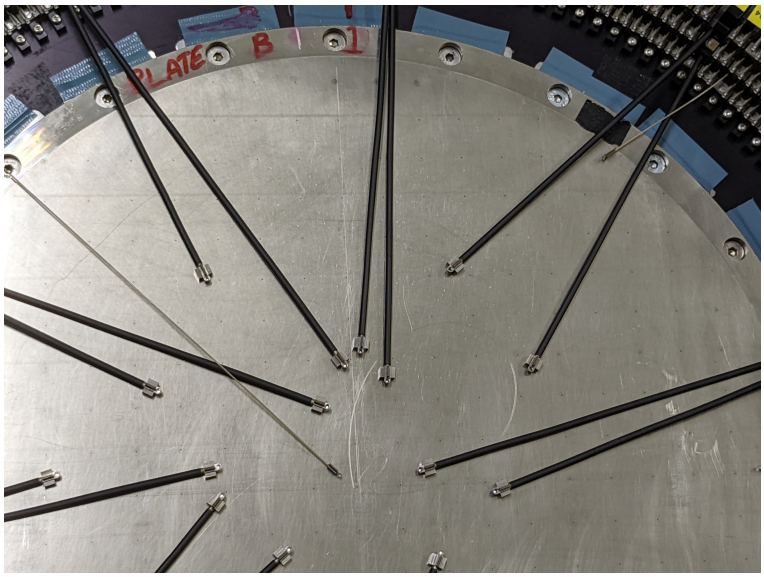


Figure 5.13: An image of the position with all 20 mIFUs deployed in a configuration by the positioner.

the mIFUs share the MOS B autoguider system therefore, the procedures required to calibrate them follow a combination of those required for the LIFU and MOS modes. The mIFUs were deployed in a configuration to centre an isolated star within each bundle, shown in figure 5.13. The fibre mapping was then determined by defocusing the telescope to produce a doughnut image inside each bundle, which was then used as part of the image reconstruction. The results of the fibre mapping are presented in figure 5.14. Further observations are needed to confirm that the mappings of both the LIFU and the mIFU are correct, which will take place during their respective commissioning and science verification phases.

The mIFUs are subject to the same positioning tolerances as the MOS mode. However, as only 20 mIFUs are deployed at most for a single observation they are not subject to the same timing constraints. Therefore, it is not crucial to place the mIFUs on the first attempt, and the iteration software has been proven using the MOS fibres to be effective. Measurements of the mIFU handling offsets is ongoing, however this should not inhibit its commissioning and science verification. The mIFU mode commissioning will take place after the MOS mode, which takes priority.

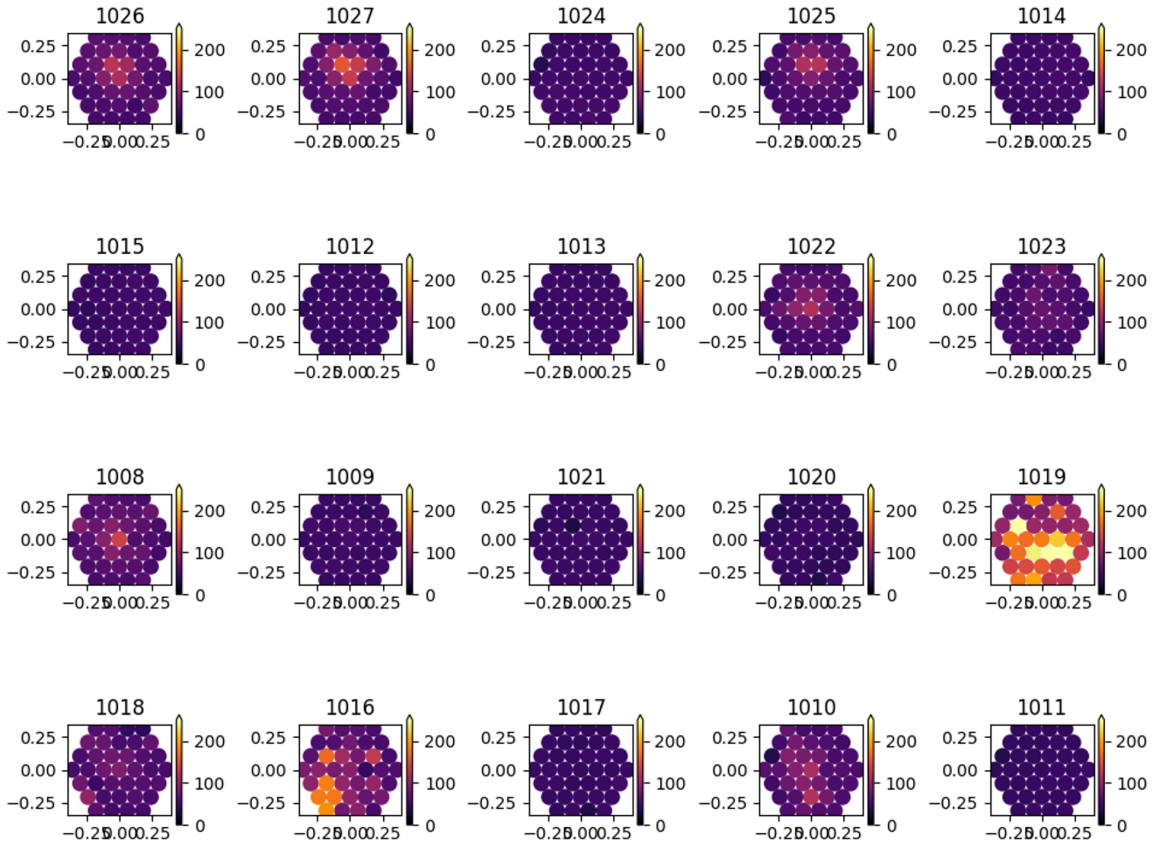


Figure 5.14: Fibre maps of the 20 mIFUs, which present the results of the image reconstruction software for each fibre bundle. Galaxies were used as targets for these mIFU observations [image credit to Scott Trager; used with permission].

5.2.3 MOS observations

At the time of writing this thesis, the commissioning of the MOS observing mode has now commenced and progress continues to be made at an increasing pace. Unfortunately, a 2-month delay has occurred from December to February 2023 as a result of poor weather conditions and unexpected events such as the dome shutter cable breaking. A significant period of time has also been dedicated to the LIFU astrometry, which was compounded by errors introduced by the WRS absolute positioning.

Alongside commissioning, investigation of retractor failures has continued. It is thought that some of the retractors which have failed may be the result of fibres being

left on the field plate after a field has been configured, although it is unknown how factors such as the temperature and humidity inside the dome affect the retractor components. As WEAVE was not yet in full-operation, MOS fields were being deployed well ahead of time and were initially being observed over multiple nights. It is possible that this has caused some "memory" to develop in the shape of the fibres protective tubing, known as the fibre peek. This bulge may cause the top pulley to derail. It was recently discovered that despite a top pulley derailment, it is possible, at least for tier 1 fibres, to continue being withdrawn and retracted, meaning that would appear to operate as normal. These fibres may continue to work for many cycles until they randomly fail, making it difficult to trace the exact cause. It is possible that these retractors will continue to work until the bulge in the peek eventually derails the bottom pulley if the fibre is extended close to the centre of the field. Due to the compact arrangement of the retractors, it is not possible to identify these fibres until the retractors have been manually extracted from the positioner.

The observatory protocols have been updated to ensure that the fibres are parked between operational test runs. Over time, it is expected that the instances of these events will reduce significantly as full operations begin, and with scheduled maintenance procedures the multiplex will be within the $\sim 5\%$ disabled tolerance. It should be noted no significant delays have occurred as a result of fibres being disabled and repaired. This will be performed as a daytime task over the course of a few days although the exact schedule will be determined by the ING team once survey operations begin.

Whilst the calibration process for the MOS observing mode is still ongoing, observations have begun using an increased buffer to protect the fibres (from 50 μm to 100 μm). For the MOS commissioning, fields provided by the Stellar Populations at intermediate redshifts Survey (StePS) have been used to characterise the quality of the spectra obtained by WEAVE. This is shown in figure 5.15. These are the raw

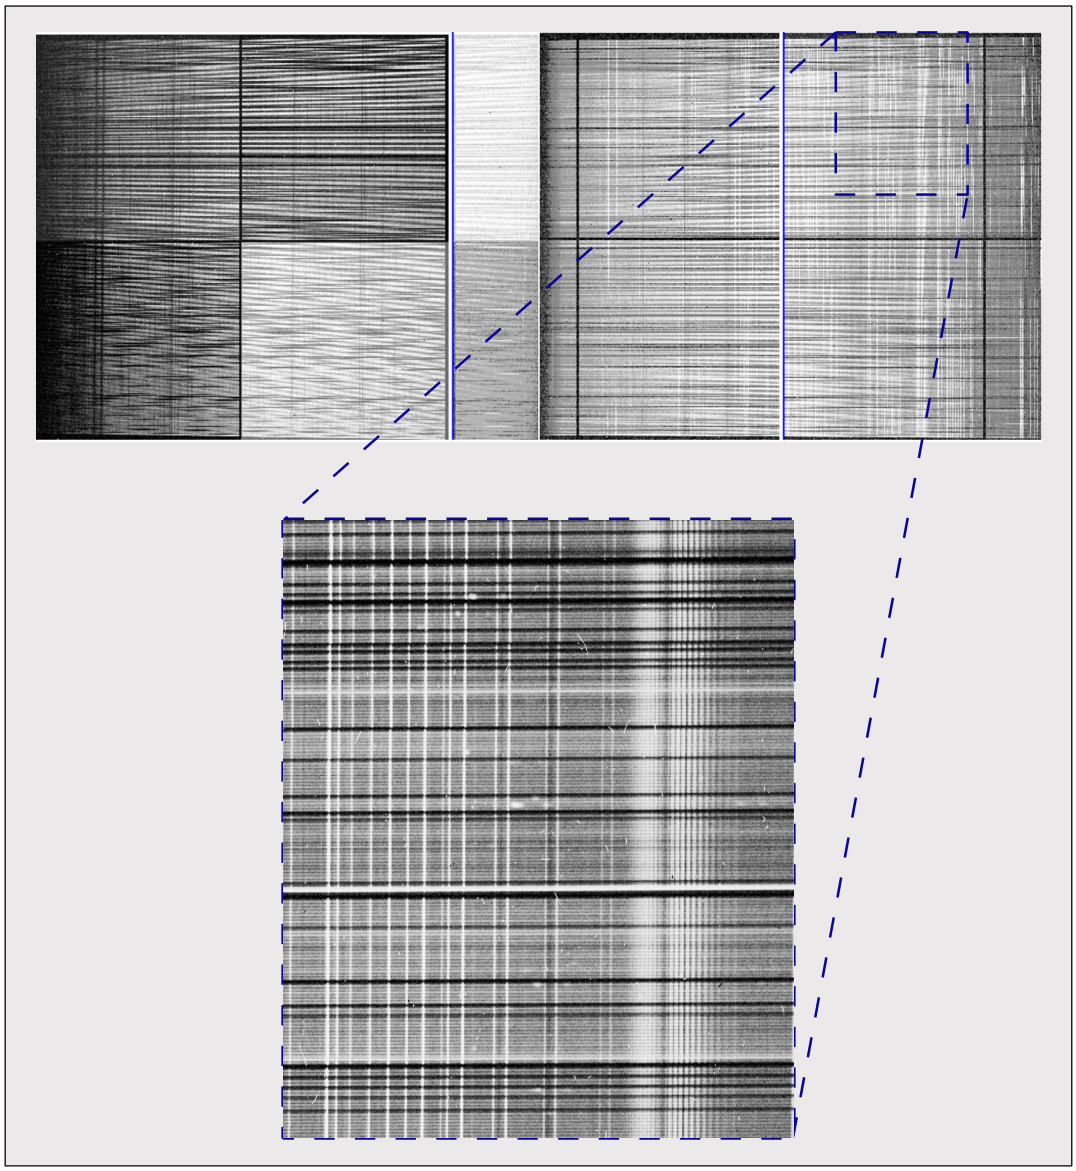


Figure 5.15: The MOS spectra taken for a 1 hour observation using a StePS science field, across both the red (*left*) and the blue (*right*). Whilst the spectra have not been processed, features such as Hα and [N II] are still visible and appear promising.

spectra which have not been processed by the pipeline yet, however features are still visible. The data obtained from this one hour observation is very promising and demonstrates the potential WEAVE has once the surveys operations begin. At this stage, the MOS pipeline has the information available to define the trace of each fibre along the detector, along with the slitlet maps, and the wavelength calibration data.

It is expected that fully-processed MOS spectra will be available in the coming weeks.

5.3 First-light data

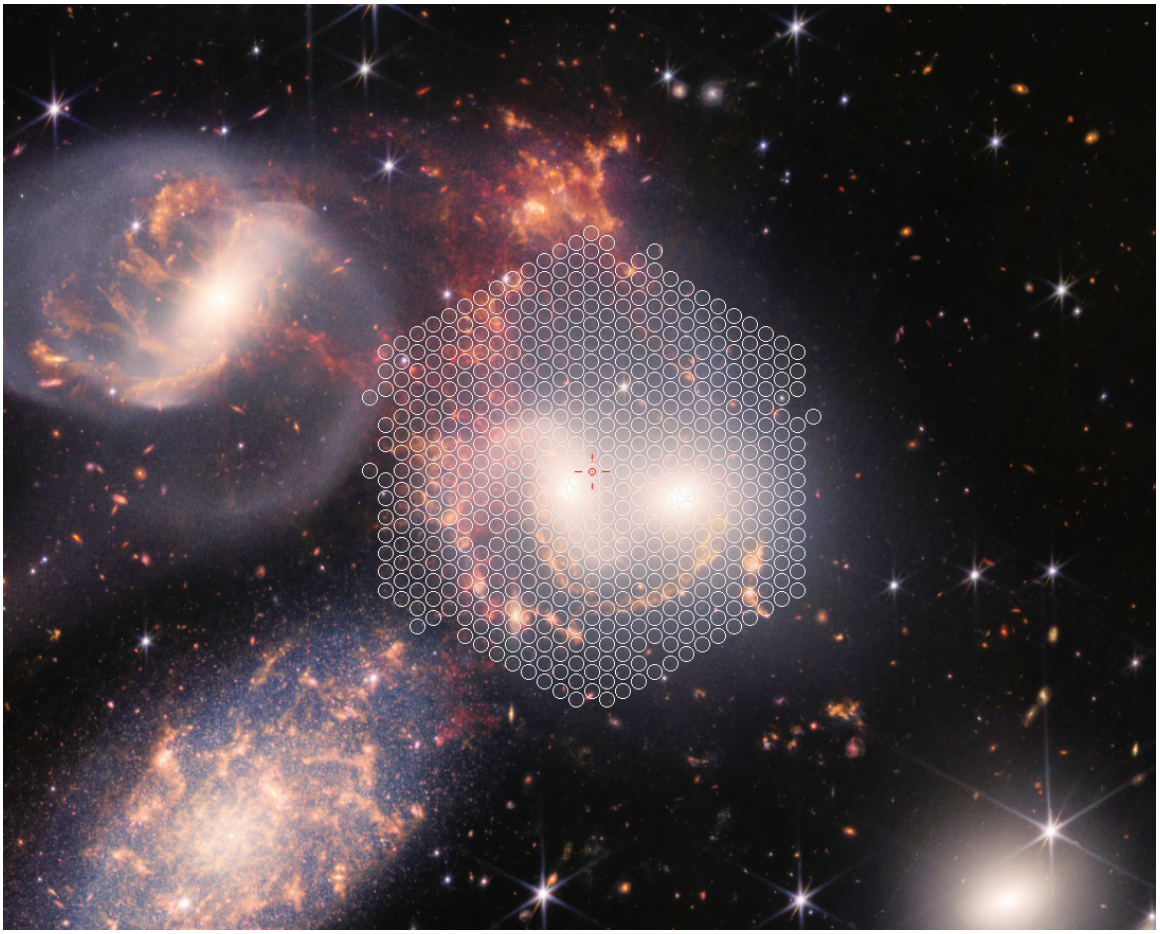


Figure 5.16: The JWST image with the WEAVE LIFU pointing at Stephan's Quintet for the first-light observation. The LIFU gathers light from 547 points on the sky for analysis by the WEAVE spectrograph (each circle indicates an optical fibre 2.6" in diameter). The observation provides physical information from each separate region of each galaxy as well as their immediate environment, covering 120,000 light years from top to bottom. Credits: NASA, ESA, CSA, STScI (background image); This composite has made use of "Aladin sky atlas" (Bonnarel et al., 2000) [image credit to Scott Trager with additions from ING; used with permission from the WEAVE first-light announcement *WEAVE First Light* (2023)].

On the 12th December 2022, the WEAVE first-light observations were announced. The LIFU was used to take a 2-hour exposure of two galaxies within Stephan's Quin-

tet, NGC 7318a and NGC 7318b. This was chosen as the first-light target due to the quintet's interesting dynamics and its recent observation by the JWST (Appleton et al., 2023). Figure 5.16 shows a composite image of the LIFU fibre footprint over an image of Stephan's quintet taken by the JWST, using the "Aladin sky atlas" developed at CDS, Strasbourg Observatory, France (Bonnarel et al., 2000).

The purpose of this section is to present and discuss the quality of the first-light observations. It does not include a scientific analysis, however the data obtained through this 2-hour exposure alone has led to the development of at least six scientific projects within the WEAVE consortium, with the aim of publishing their analysis. These projects are ongoing at the time of writing and the data available will significantly increase as the LIFU survey operations begin.

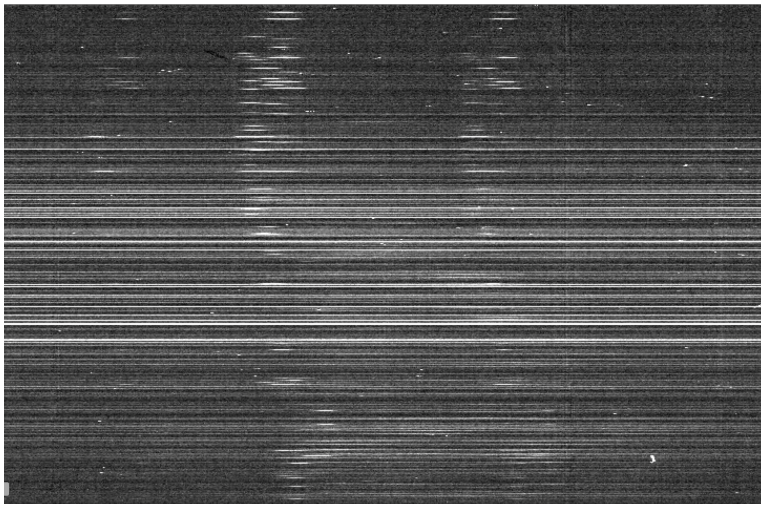


Figure 5.17: The unprocessed spectra for the 2 hour first-light observation of Stephan's Quintet, taken using the high-resolution mode of the LIFU. This image displays the Hα and [N II] lines within the NGC 7318a and NGC 7318b galaxies.

The raw spectral data produced by WEAVE during the LIFUs high-resolution (HR) observation can be seen in figure 5.17. The high-resolution mode has $R \sim 21,000$ and the presence of Hα lines can clearly be seen prior to being processed. The data was extracted by the Cambridge Astronomical Survey Unit (CASU) team and the images have been sent through the Science Processing and Archiving pipeline

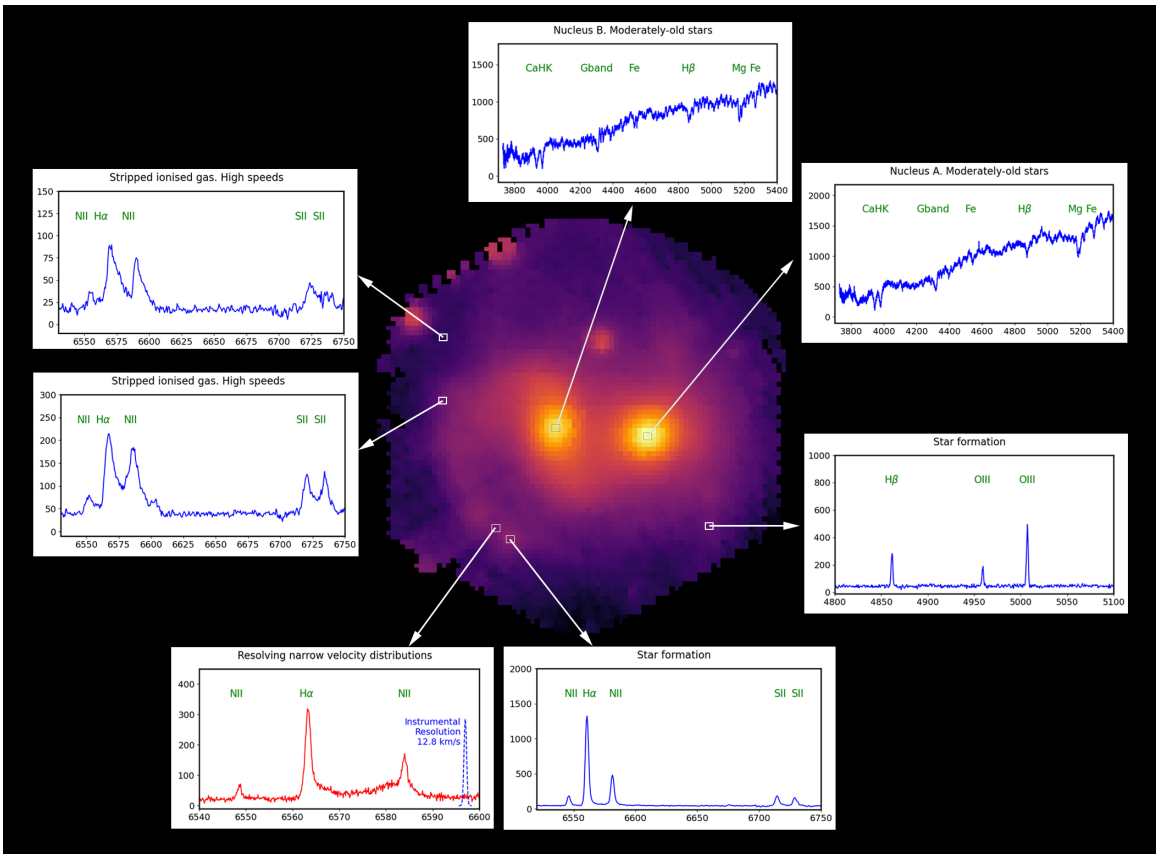


Figure 5.18: The first-light observation of Stephan's quintet by the LIFU in the LR-mode after being processed by the Core Processing System (CPS). With small shifts of telescope pointing, WEAVE/LIFU produced spectra at a total of 31500 locations in and around the galaxies, in just 2 hours. The intensity of light from the fibres builds the image of the galaxies shown in the centre. The individual spectra (intensity at each wavelength; seven examples shown) provide a wealth of information about the physical conditions at each location. At the two galaxy nuclei (top-right) the spectra indicate moderately-old stars (one billion years) and no on-going star formation. The narrow, peaked spectra in the lower-right are typical of gas (hydrogen, oxygen, nitrogen, sulfur) heated to over 10,000 degrees by very young stars, whereas the broad, asymmetric peaks in the spectra shown on the left indicate turbulent shocks between gas clouds. WEAVE is particularly accurate at measuring wavelengths, or velocities. In the bottom-left panel (in red) obtained in the high spectral resolving power mode, velocity distributions as narrow as 12.8 km/s can be measured [image credit to the WEAVE Team; used with permission from the WEAVE first-light announcement *WEAVE First Light* (2023)].

(SPA), which combined this exposure with calibration data such as biases, darks, and detector and fibre flats to extract and calibrate the spectra before combining it into a stacked spectral data cube. This worked as a proof of concept for the pipeline

and figure 5.18 displays the stacked LR spectral datacube, with several 1D spectra extracted to display interesting features of the galactic dynamics.

The HR datacube represents the highest resolution IFU observations that have been obtained for these galaxies within Stephan's quintet. Iglesias-Páramo et al. (2012) observed Stephan's quintet with the Potsdam Multi-Aperture Spectrophotometer (PMAS) over a $5305 - 6809$Å with R$\sim$ 600. This resulted in severe blending of the Hα and [N II]$\lambda 6548, 6583$Å. Figure 5.19 displays these spectral lines observed by WEAVE, which have been extracted from two regions of the shockfront for both the HR and LR datacubes, which can be seen on the left-hand side of figure 5.18.

It is important to note that the absolute flux scale of figure 5.19 is likely incorrect as the spectra were extracted from an early test of the pipeline, however the relative flux scale between the Hα and [N II]$\lambda 6548, 6583$Å lines is correct. The red and green lines represent the spectra taken from the LR and HR datacubes. The top plot was taken from the lower left of the shockfront, and highlights that in regions dominated by a single line component the spectra still display complexity. Tails can be seen on the right-hand side of each line, however the lines themselves are sharp and clearly distinguishable both resolution modes. The bottom plot displays spectra extracted from the north of the middle of the shock, which exhibits the most complex behaviour. For the low-resolution mode it is significantly harder to disentangle the velocity structures compared to the high-resolution mode due to blurring effects. This spectrum has also been smoothed using a 5-pixel wide box-car filter.

Both figure 5.18 and figure 5.19 highlight the spectral data produced by the first-light exposures and the potential of future observations with WEAVE. The first-light data has been fully-processed by the pipeline, however there are ongoing improvements being implemented in areas such as the cosmic ray extraction and the drizzling solutions, along with adjustments to the focus and fibre traces as a result of changes to the spectrograph alignment. This will require the data to be re-processed, however

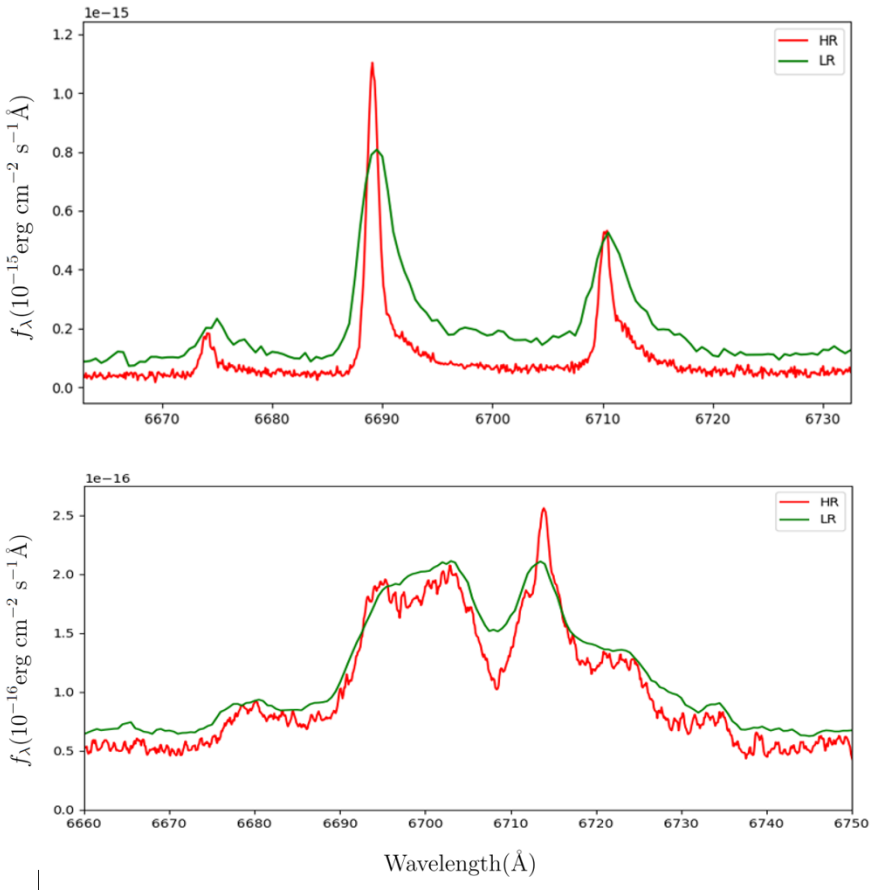


Figure 5.19: The flux taken during the 2-hour LIFU first-light observation of Stephan's Quintet in both the high and low resolution modes. It highlights the evolution of the velocity structure within the [N II] and Hα lines across the shockfront between NGC 7318a and NGC 7318b. (*Top*) A spectrum extracted from the lower left of the shockfront. (*Bottom*) A spectrum extracted from north of the middle of the shock, which exhibits the most complex behaviour. This noise of this spectrum has been smoothed using a 5-pixel wide box-car filter [image credit to Scott Trager; used with permission].

the quality of this data demonstrates that WEAVE is now ready to enter the science verification phase.

5.4 Specification comparison

In the original schedule for WEAVE, 50 nights were allocated in total for commissioning, with a minimum of 25 nights with good weather conditions needed. Un-

fortunately, the commissioning period for WEAVE has taken longer than expected, primarily due to the impact of winter weather conditions and unexpected circumstances, such as dome shutter cable malfunctions. Despite these setbacks, progress continues to be made and the members of the WEAVE collaboration are making full use of the available data to characterise and improve the six observing modes. Table 5.1 summarises the wavelength coverage of the spectrograph, which is now stable, across the low and high resolution modes (LR and HR).

The RMS obtained using the astrometry measurement are approximately 0.3". This does not currently meet the < 0.1" science requirement therefore, further refinement and investigation is needed. It is expected that the implementation of a predictive model for the field plate centre with telescope orientation should reduce the RMS to this tolerance once the drift introduced by the WRS has been resolved.

	Wavelength range Å
LR Blue	3660-5550; 5615-5955
LR Red	5800-7580; 7670-9470
HR1 Blue	4005-4555; 4535-4700
HR2 Blue	4705-5345; 5325-5515
HR1 Red	5905-6400; 6380-6815

Table 5.1: The results of the WEAVE spectrograph wavelength coverage, supplied by the Cambridge Astronomical Survey Unit (CASU) team. The two low resolution (LR) ranges represent the wavelength range on each side of the gap between the two detector chips. For the high resolution (HR) mode, the wavelength ranges also represent the coverage on either side of the gap between the two detector chips, however as the x position of a given wavelength is strongly curved, there is a 20 Å overlap between the two ranges.

It is crucial to assess the throughput of WEAVE over the designed exposure time for observations, as this defines the signal-to-noise performance of the instrument. The throughput is defined as the fraction of incoming photons which are detected by WEAVE and is calculated as

$$\text{Throughput} = \text{Atmosphere} \times \text{Reflectivity} \times \text{PFC} \times \text{Fibres} \times \text{Spectrograph} \times \text{QE}, \quad (5.2)$$

where Atmosphere is the atmospheric throughput, Reflectivity is the reflectivity of the WHT primary mirror, PFC is the throughput of the prime-focus corrector, Fibre is the throughput of the positioner's optical fibres, Spectrograph is the throughput of the spectrograph, and QE is the quantum efficiency of the CCD detectors.

Due to the ongoing alignment of the spectrograph, it is not possible to give an estimate of this number. Simulations of the WEAVE system indicate that the throughput exceeds the minimum survey requirements by at least 25%. However, this margin must be determined as the throughput value is the driver behind the survey operations. The Galactic Archeology (GA), WEAVE-LOFAR (WL), and Quasi-Stellar Objects (QSO) surveys would be notably impacted by any increases to the exposure time due to the throughput, as they are based on the area covered during observations rather than the signal-to-noise ratio (SNR). Following the spectrograph alignment, further modelling will be required to assess the impact of the WEAVE systems on the scientific survey designs.

It is crucial that the spectral cross-talk between adjacent fibres and potential contamination of sources is evaluated both on-sky and after data-processing. Contamination can occur from many instrumental effects, such as scattered light and ghosting, and needs to be carefully measured. For WEAVE, it is essential that these contributions are $< 5\%$ prior to data processing and $< 0.05\%$ after the pipeline is applied. The aim of this requirement is to minimise the noise contribution of a bright object that is adjacent to a faint object. The signal-to-noise ratio of the fainter object should not be degraded by more than a factor of two, on the assumption that the magnitude difference between two objects in the same slit is up to 5.

Unfortunately, the spectrograph alignment work has delayed the evaluation of the cross-talk between neighbouring fibres, as their position along the detectors is likely to change. However, at the time of writing no issues have been identified which have caused concern and the LIFU on-sky commissioning observations have been

completed. Most aspects of the LIFU's performance have been tested by ING and it is expected that a decision will be made shortly regarding the end of the LIFU commissioning period and the beginning of its survey observations, subject to further tests required as a result of adjustments to the spectrograph alignment. The MOS commissioning period continues, with the astrometry runs taking priority alongside positioning tolerance and timing testing as a daytime activity.

Chapter 6

Conclusion

> Even the smallest mIFU can change the world.
>
> G. Dalton

The work contained within this thesis describes the final assembly of the WEAVE fibre positioner, the development and evolution of its calibration procedures, an analysis of its field assignment performance, and the first-light results. Whilst the commissioning period for WEAVE continues, the results thus far appear promising and it is expected that WEAVE will meet its design specifications by the time survey operations begin. An overview of the calibration processes described in this thesis and how they interlink is shown in figure 6.2.

Following the completed assembly, metrology measurements of the tumbler's position relative to the interface plate were retaken due to the additional weight applied by the MOS fibre retractors. Prior to the retractors installation, these measurements were completed using a DTI gauge and a calibration bracket setup by Ellen Schallig as part of the original metrology measurements as the positioner's structure was built. From the new set of measurements, it was found that the rotation axis of the tumbler was off-axis and the structure as a whole had shifted downwards by approximately 1 mm towards the interface plate. Combined with the tilt of the field plates, the result

was a cone-like structure which required interventions to meet the ± 15 μm limit to the field plate height when in the observing position. Prior to shipping, the change in height along the y-axis was 15 μm and 10 μm for plate A and B respectively. Following these adjustments to the tumbler, calibration procedures were then developed by Gavin Dalton and me in the Oxford laboratory with the positioner in the zenith position. Measurements of the invar field plates surface profile were taken using gauge blocks to produce a smooth contour map that can be used to estimate the z-height above the plate the robot should use when positioning each fibre.

With the retractors successfully installed, work began to test the fibre movement procedures that had been previously developed using proto-type fibres. Before any of the fibres can be moved, their park positions must be correctly placed in position on the retractor porch. As the fibres were installed by hand their orientations were random, and the robot gripper unit was used to calculate and adjust their position whilst minimising the risk of damage. Nona was used to drive over the fibre's expected park position, with a z-height slightly above the button, to inspect the fibre's orientation and adjust the position of the gripper jaws relative to the button vane. The centroiding software then calculated the fibre's position and applied a correction. To ensure that each fibre is parked correctly, the heights of the retractor porches were measured manually using the gripper jaws and the z-axis encoders. The height of the retractors exhibited a periodic motion for both robots that corresponded to differences between the x-stages as the robots move around the edge of the retractor support disc (RSD). It is maximised at the furthest points across the field along the y-direction due to differences in the rest height of the y-stages across each x-axis.

Unfortunately, the positioner was damaged during its journey to the observatory, and several months were spent repairing the tumbler structure after several bolts had sheared off due to the freely moving counterweights that had not been detached prior to shipping. Consequently, the metrology measurements of the tumbler struc-

ture were repeated and a new robot counterweight system was installed to prevent further incidents. Following the tumbler reassembly, the positioner was mounted to the WEAVE rotator system (WRS) and aligned via the interface plate, before both components were mounted onto a testing rig which could manually change in elevation angle. This setup was used to simulate the positioner's orientation when mounted to the telescope, and a series of measurements were conducted using a laser tracker to determine the flexure of the tumbler structure with respect to the focal plane of the corrector. Initially measurements displayed a curved profile with a maximum displacement of 1000 μm and mechanical interventions were required. Inserting stiffening panels along the central beam, a clamp on the bearing side of the rotator axis, and a reinforcement strut to restrict the movement of the drums resulted in a significant reduction to the magnitude of this displacement. After repositioning the tumbler structure and the shims beneath the field plates, the maximum displacement of the tumbler is approximately 200 μm and the flatness across the field plates in the observing position is within ± 14 μm for the vast majority of the elevation and rotator angles within the $\pm 155^\circ$ rotation limit.

The field plate surface profile was also remeasured following the tumbler reassembly, using a DTI gauge due to restricted access to the positioner on the testing rig however, this method removed the uncertainties introduced by the gauge blocks. The DTI was attached to the robot axis in the position of the θ-axis counterweight, which restricted the range of coordinates that could be measured with a consistent θ angle. Additional measurements were taken at $\pm 90^\circ$ rotation angle to cover the full active area of the field plate, alongside measurements at (0,0) for each rotation angle to ensure the height measurements could be converted into a single reference frame. With a new base contour map defined at the basis used for the flexure measurements ZD 40°, measurements of the field plate surface profile were taken for a series of rotation and elevation angles across the observable range. The evolution of the field plates

surface profile was used as a basis to predict changes to the retractor porch z-heights. The heights of the guide fibre porches were measured to evaluate the predictions given by the surface profile evolution, and the largest residuals were approximately 40 μm, which can be compensated for by over-driving the z-axis when parking the fibres.

In June 2021, measurements of the fibre's movement procedure began. This process aimed to measure the required adjustment to each fibre when completing either a *move*, *park*, or *unpark* sequence. Initially, 50 moves were completed by each fibre within its allowed patrol radius, which demonstrated that the offsets were not dependent on the targets position along the field plate. I then created a database to store the offset information for the previous 500 moves for every fibre, so that a rolling average can be taken as the database is populated. Sigma clipping has been an effective tool for excluding extreme variations, which may be caused by errors in the movement sequence, without removing the underlying distribution of the offsets. Further experimentation and analysis is required as the database continues to be populated across a series of elevation and rotation angles. The fibre movement testing has been used as a basis to estimate the field configuration timing. At the time of writing, it takes $\sim$ 45 minutes to deploy every fibre from park, providing a rough timing estimate of 1 hour 35 minutes to complete a single configuration. Given that eventually the robot gantries will be running at the maximum speed possible, without exceeding the kinematic specification defined for each axis, it is expected that the positioner will reach its 1 hour timing limit by the time survey operations begin by implementing additional optimisation techniques.

Chapter 5 presents the first-light data obtained through a 2-hour observation of Stephan's Quintet, specifically of the two galaxies NGC 7318a and NGC 7318b. The low-resolution (LR) LIFU mode has been fully-processed by the WEAVE pipeline into a datacube which astronomers within the WEAVE collaboration can access. There are currently six projects being conducted by the collaboration using this observation

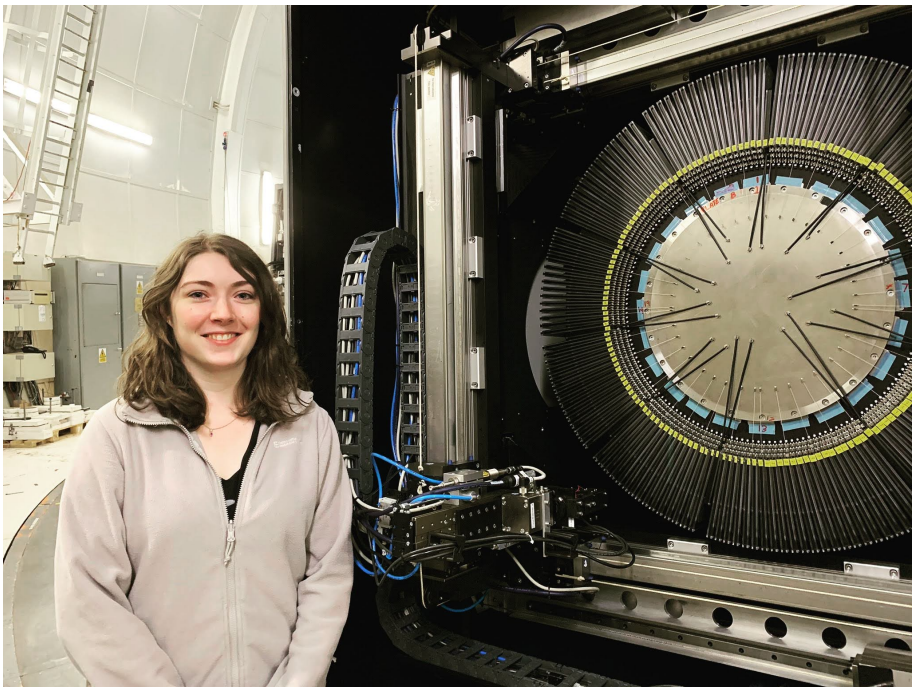


Figure 6.1: The author alongside the completed top-end of WEAVE.

alone, with the aim of publishing their analysis in the near future. This data presents the capabilities of WEAVE and is a proof of concept for the pipeline processing.

Due to the ongoing alignment testing of the spectrograph, it is not currently possible to make comparisons of the throughput and cross-talk with the WEAVE design specifications, as the fibre's position with respect to the spectrograph detectors is subject to change. Recent measurements of the positioner's distortion map have resulted in an RMS of 0.3" for the astrometry measurements, which exceeds the 0.1" survey science requirement. Estimations of the shift in the center of the distortion as a function of telescope orientation appear promising, and the implementation of this model will enable fibres to be placed according to where the field plate will be during observations, rather than the plates position at the time of configuring. It is expected that this refinement will bring the astrometry measurements within tolerance once the drift in the rotators absolute position has been resolved.

The LIFU observing mode has now reached the end of its commissioning period and a large amount of data is expected to be processed in the coming weeks as it

begins its science verification observations. It is expected that the MOS and mIFU modes will follow suit in the coming weeks.

6.1 Further work

At the time of writing this chapter, WEAVE is still in its commissioning and science verification period[1]. The LIFU observing mode has been fully commissioned, subject to a few remaining tests, and is about to begin its science verification phase. Whilst significant progress has been made for the MOS and mIFU modes, more work is required before these observations may begin.

The key tasks remaining for the MOS observation mode are the astrometry refinement and the configuration timings. As commissioning observations are taking place in the MOS mode, the fibres are being exercised and the calibration database is being gradually populated. The offsets calculated by the exercises are being fed back into the positioning sequence, and over time it is expected that it will take fewer attempts to position each fibre, with the ultimate goal of placing each fibre on its first attempt. This must be conducted across the elevation and rotation ranges used to generate the surface profiles and retractor heights, as the majority of the offsets have been measured at ZD 50° rotation 0°. The first stage is to complete a series of configurations at ZD 50° rotation 0° using a 50 μm fibre footprint buffer where the majority of the fibres are placed in a single attempt. Once it has been confirmed that these configurations are reproducible, the acceleration parameter specifications on the robot kinematics can be increased, to reduce the time taken to configure a single field, without compromising on the number of iterations taken to position a single fibre. It is expected that this will be an ongoing refinement process that should meet the 1 hour configuration timing requirement when survey operations begin.

The astrometry runs are now a high priority observation task to be completed in

[1]This has taken longer than expected, as is usual in instrumentation projects

good seeing conditions, and in recent weeks an increasing number of measurements have been conducted. Figure 5.8 demonstrates the impact of the shift in the field plate position on the centre of the distortion map. Scans of the field plate grid are now being conducted as a daytime activity to study the shift in the field plate center as a function of elevation and rotation angle. The intention is to produce a model which refines the position of the fibres to place them based on where the field plate will be at the time the observation takes place. After the ongoing drift in the rotators absolute position is resolved by the ING engineering team, it is expected that this modelling will reduce the astrometry RMS to be within the required 0.1".

The WEAVE pipeline continues to improve as astronomers within the collaboration explore the quality of the data available. Improvements in the cosmic ray extraction, wavelength solution, and dichroic model are ongoing and the first-light data will be reprocessed in the coming weeks.

6.2 The future of WEAVE

WEAVE is an incredibly flexible and versatile instrument, and discussions have already begun about upgrades that could be implemented in the future.

Ideally, these solutions would avoid removing the WEAVE from the top-end, as this would likely require significant re-calibration work before observations could restart. One possibility would be to introduce a small detachable secondary mirror to the positioners observing field plate. This would enable smaller instruments to be installed on the opposing Nasymth platform to the WEAVE spectrograph. It would allow the WHT to conduct a diverse range of observations, without the need to remove WEAVE from the top-end.

Any upgrades to the spectrograph itself would require significant intervention however, it is possible that it could be reconfigured to include an additional grating

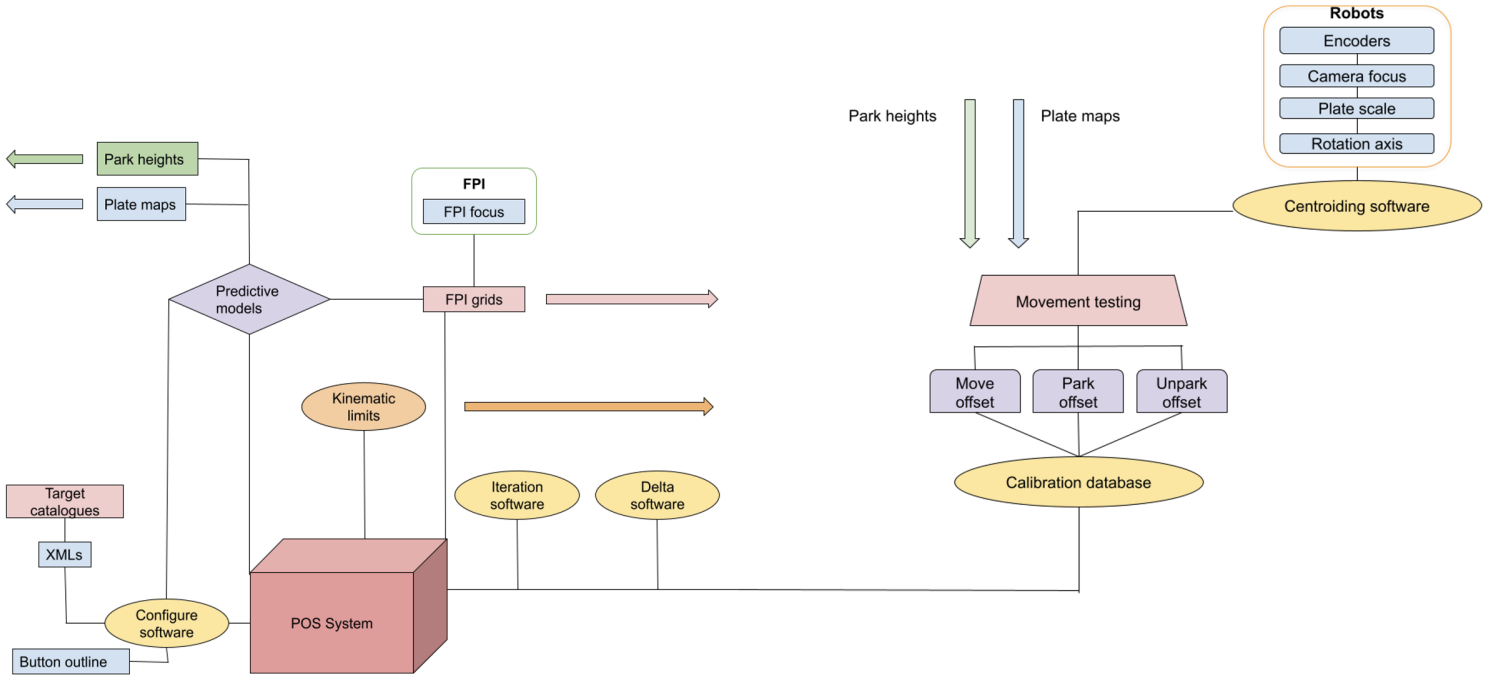


Figure 6.2: A diagram highlighting the calibration processes involved in the WEAVE fibre positioner and how they interlink. This is not an exhaustive list of the processes involved in the POS system. This diagram is designed to give an overview of the calibration stages described within this thesis.

and railing on the red arm. The grating would be used to increase the wavelength coverage to include the 7700Å to 8600Å window for the high-resolution mode, which is beneficial to astronomers. In particular, this would extend to the calcium triplet lines Ca II 8498, 8542, and 8662 Å which are strong, widely spaced emission lines useful for galactic kinematic studies (Dressler, 1984). Depending on the wavelength window used, it could be used to include oxygen lines which can be used to study the stellar chemical evolution produced directly by supernovae. At this stage, this wavelength is only available for the low-resolution mode. However, it is worth noting that it would be difficult to find a grating design configuration for these specifications which maintains a high throughput.

Another possibility would be to insert a new grating that would improve the blaze profile and throughput at the edges of the HR mode. Lee and Poutous (2022) have begun to fabricate plasma etched surface relief gratings, which introduces ripples on the surface that force the photons through the grating. They have the potential to be a high-throughput alternative to Volume Phase Holographic gratings, and replacing them would be a modest upgrade to WEAVE.

In conclusion, the WEAVE fibre positioner is on the brink of beginning its survey operations across all three modes, and the vast quantities of data it will provide are certain to make a significant impact on our understanding of the universe. Whilst WEAVE has a designated 5-year survey, it has a lifetime of $\sim$ 20 years. It is unknown what scientific questions WEAVE will uncover yet it is clear that the future of the WHT will continue to be shaped by it for many years to come.

The end of THE END is the best place to begin THE END, because if you read THE END from the beginning of the beginning of THE END to the end of the end of THE END, you will arrive at the end.

$\sim$ *Lemony Snicket* $\sim$

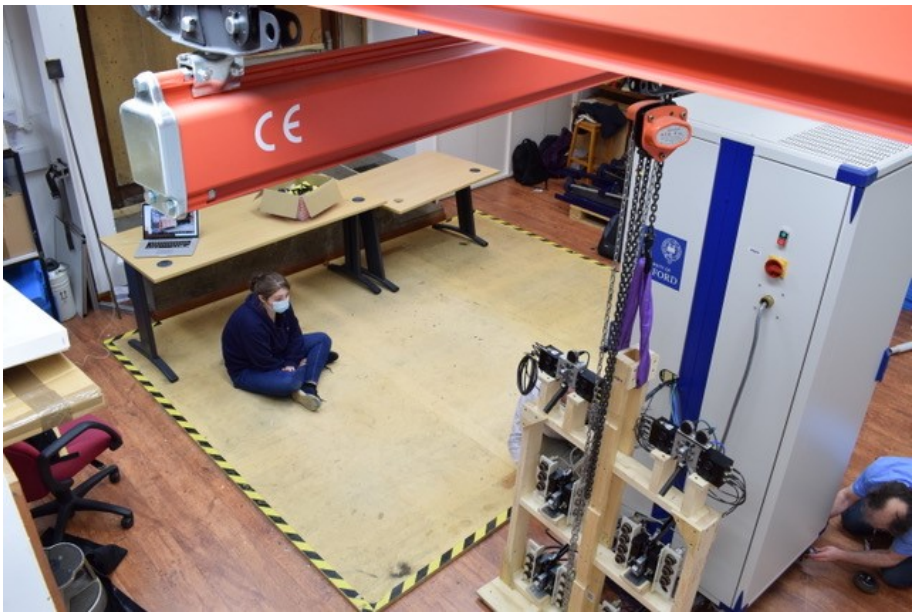


Figure 6.3: The forlorn graduate student, sat on the floor of the empty WEAVE Oxford laboratory.

Appendix A

The WEAVE science surveys in detail

A.1 Galactic Archaeology and Evolution

The largest survey planned for WEAVE is the galactic archaeology survey. Its goal is to obtain comprehensive data of positions, ages and chemical abundances of all major Milky Way stellar structures in the northern hemisphere. It can be summarised into four sub-surveys, each with their own goals and observational requirements.

By studying the galactic halo surrounding our galaxy, the formation history can be traced due to its unique sub-structure and stellar chemistry. Recent surveys have focused on the inner halo out to 40 kpc, which demonstrated structural properties formed by debris from satellite galaxies (Bell et al., 2008). The WEAVE survey plans to extend this study to the outer halo regions to further constrain formation models and also study RGB stars. These have only recently had considerable catalogues produced (Xue et al., 2014). The HR survey will measure 1.5 million stars over an area of 8500 deg^2, covering a magnitude range of $15 < r_{SDSS} < 17$ to compliment limitations of Gaia's RV measurements.

The Milky way's thin and thick disc structure is the main stellar component of our galaxy, yet many questions remain about its formation, dark matter distribution and kinematics. The low resolution (LR) WEAVE survey is designed to measure the spatial structure and mean velocity distribution of the disc populations, as well as the distribution of their metallicities and α elements. Constraints can then be placed on first order non-axisymmetric perturbation models (Monari et al., 2016), as well as other dynamics such as radial migration (Sellwood and Binney, 2002). Approximately 500 lines of sight through the galactic disc will be observed with $> 10^6$ stars, chosen to be mainly red-clump giants with a magnitude range of $15 < r < 19$.

The kinematics of many stellar components of our galaxy overlap. An efficient way to untangle and understand their interactions is to obtain detailed spectra for chemical abundance analysis. The HR chemo-dynamical survey will complement GAIAs phase-space distribution to provide full chemical abundances for stars with a magnitude between $12 < G < 16$ over an area of 5600 deg^2. In regions where GAIAs parallax accuracy is better than 10%, stellar parameters such as effective temperature and metallicity provided by WEAVE will allow for stellar ages to be constrained to a sphere with a radius of $\sim$ 2 kpc. These stellar ages will be used to constrain disc assembly models (Schonrich and McMillan, 2017), as well as identify predicted sub-structures (Helmi and White, 1999).

The study of Open clusters (OCs) provides vital information on the star formation (SF) of our galaxy, and allows evolutionary changes to be traced across all ages and masses of stars. The open clusters survey aims to sample $\sim$ 300 old and new clusters, that cover a star forming region that complements GAIAs newly discovered OCs up to 5 kpc from the sun (Castro-Ginard et al., 2019). Most of these observations will be made in the HR mode to provide full chemical abundances.

A.1.1 Stellar, Circumstellar and Interstellar physics

The stellar, circumstellar and interstellar physics (SCIP) survey plans to improve ongoing research of stellar and interstellar processes through local studies, which can then be applied to the younger universe. Like many other areas of astronomy, larger statistical samples are needed to remove observational bias towards more clustered environments. For this reason, observations are focused on the northern galactic plane due to its vast range of targets at different evolution stages, making use of both the LR and HR modes of WEAVE.

The LR survey is divided into four main areas; massive stars, ISM sight-lines, young stellar objects, and specific target groups such as Cephid variables and compact objects. The majority of fibre allocations target massive stars such as OBA stars and RSGs ($\sim 60\%$), with the aim of mapping their true unbiased spatial distribution and the ejection frequency from clustered environments. Current observations of massive OB stars have been limited to a magnitude of $V \sim 11$ and a distance of 2.5 kpc (Apellániz et al., 2016) (Garmany et al., 1982). WEAVE will expand this magnitude range up to $B \sim 20$ and conduct spectroscopic characterisation across all the OBA stars evolutionary stages to improve models of high-mass stellar-evolution theory. Additionally, the spectra will be used to trace the 3D distribution of extinction in the ISM, as current models cannot agree on extinction law parameters (Nataf et al., 2016).

The HR survey is directed towards the Cygnus region, due to its close proximity at 1.5 kpc and high star formation rate, building on data from GAIA DR2. The primary goals are to measure the distributions of several parameters, such as; rotational velocities, the binary fraction, and stellar multiplicity. Chemical abundances for O3-B9 stars with a magnitude range of $11 < B < 16.5$ will be used to study internal-mixing and self-enrichment processes, and kinematic data will be used to map both the stellar population and the ISM clouds for the Cygnus region.

A.1.2 White Dwarf Stars

Every star that has a mass ≤ 8 $M_{\odot}$ will become a White Dwarf (WD) at the end of its life. Due to the short lifetime of stars with $M \geq 1.5$ $M_{\odot}$, the vast majority of A/F-type stars that formed in our galaxy will already be WD stars. In the WEAVE White Dwarf Survey, approximately 100,000 WD stars with $G \leq 20$ will be sampled with LR spectroscopy, which is triple the current sample in the northern hemisphere (Gentile et al., 2015).

The target stars have been chosen using the data from GAIA DR2, as the spectra from GAIA is too faint to determine any key features. Every WEAVE pointing will have 15 fibres allocated to WD observations ($\sim 1.5\%$ of all fibre allocations), as they were originally intended to be used as a calibration tool before being expanded into their own survey.

The large wavelength coverage of WEAVE in the blue range (down to 3700 Å) will cover the higher Balmer lines and show any pollution due to accretion of planetary debris. It is suggested that approximately 1% of all WD stars accrete planetary debris, giving a sample size of ~ 1000 for this survey. Furthermore, it is thought that approximately 1 in 1000 WD stars are a rare species, e.g. near the core-collapse boundary (Gaensicke et al., 2010) or formed from unknown binary interactions, so it is likely a significant sample will be discovered.

By increasing the sample size, further constraints can be made on a number of models, such as the mass distribution of WDs formed by binary mergers (Wegg and Phinney, 2012), the mass-radius relation (Tremblay et al., 2017), and the age-velocity dispersion relation. The RVs provided by WEAVE will be used in conjunction with GAIA data to provide full 6D kinematics, allowing the thin disc, thick disc, and halo populations to be distinguished (Anguiano et al., 2017). The study of WD stars are also important for other areas of such as particle and high-density plasma physics, due to their extreme environments that cannot be recreated in a laboratory.

A.1.3 Apertif

Apertif is a new focal-plane receiver array that began operations in 2019 on the Westerbork Synthesis Radio Telescope (WRSST), and is one of telescopes involved in the Square Kilometer Array (SKA) project (Verheijen et al., 2009). Apertif is designed to look at the 21cm HI line; a crucial tool for mapping the distribution of neutral hydrogen in galaxies (Rohlfs and Wilson, 2004). Observations of this line are sparse but show extended neutral discs in many galaxies, with evidence for physical processes including minor mergers, stellar feedback, and the fuelling of AGN. Apertif will resolve the kinematics of approximately 100,000 galaxies.

The WEAVE-Apertif survey aims to understand the evolution of galaxies by studying the local group. It is the only survey planned that makes use of both the $1.5' \times 1.3'$ LIFU and the HR mode of WEAVE, with 400 galaxies outlined as LIFU targets. This survey will combine the spatially resolved kinematics of neutral gas with distributions of stellar populations and chemical abundances to study the following three areas; galaxy bimodality, the accretion history of spiral discs, and disc distributions of baryonic and dark matter.

The LIFU component of the survey will cover a wavelength range of $3660-9840$ Å at a resolution R$\sim$ 2500. Targets have been selected based on parent surveys that have resolved HI distributions from Apertif's wide-area and medium-depth surveys. These galaxies include members of both the red sequence and blue cloud for bimodality, as well as a broad range of morphologies and environments.

The HR mode will target 100 galaxies from the subset observed by the LIFU survey. These have been chosen specifically to study the influence of HI on the evolution of bars and spiral arms, as well as the distribution of baryonic matter for low and high surface-brightness galaxies.

A.1.4 Galaxy clusters

To study unresolved areas of galaxy formation, observations of galaxy clusters are essential. They contain a diverse range of environments at a variety of evolutionary stages. The WEAVE galaxy cluster survey is divided into three sub-sections to effectively meet its science goals; Nearby clusters, wide-field surveying, and the cosmological survey.

There are many questions when tracing the formation history of low-mass galaxies for a range of environments. One example of this is the missing satellite problem. Topics that are highly debated include the efficiency of star formation and the 'downsizing' of galaxies. The nearby survey will advance the study of these areas by observing the largest sample of Dwarf galaxy cluster members to date, using a combination of the mIFU and LR mode for each pointing. The sample will consist of 47 clusters within $0.01 < z < 0.04$ and obtain reliable velocities down to $M_r \sim -16$. A sample of this size is large enough to observe the effects of cluster properties on formation, such as the mass, galaxy density, and the cluster velocity dispersion.

The wide-field survey is designed to investigate the infall of galaxies towards the centre of a cluster and to map the filaments which feed the cluster itself. The virialised regions of galaxy clusters have been studied in great detail, but little time has been spent on the outer regions. Therefore, the outskirts of these clusters is the ideal area to study the accretion of matter from the cosmic web. A sample of 16 clusters have been chosen up to five times the virial radius (r_{200}) at $z = 0.056$, which will be observed using the LR mode of WEAVE.

The cosmological survey will observe the largest sample of clusters in comparison to the other sub-surveys, with ~ 70 clusters chosen from the Planck Sunyaev-Zel'dovich redshift survey (PSZ). All targets have $z < 0.5$ and will be observing with both the LIFU and LR mode of WEAVE. The objective is to study the stellar populations of bright galaxies within the central region of the cluster. This is because

the Planck SZ survey suffers from observational biases when estimating cosmological parameters, which is not improved by increasing the sample size. As this is the main source of error for these estimators, an independent measure is needed of the SZ scaling relation, which WEAVE will provide. This allows for further constrains to be made on cosmological parameters such as $\Omega_{\rm m}$ and σ_8.

A.1.5 STePS

The objective of the Stellar Populations at intermediate redshifts Survey (StePS) is to understand the processes that cause deviations between dark matter halos and luminous structures. There has been recent evidence that star formation episodes occur in passive galaxies with a redshift $z < 0.8$, but further observations are needed to constrain theoretical models and simulations.

The survey has been constructed to combine two different analysis techniques; archaeological reconstruction and the look-back approach. Archaeological reconstruction is a difficult type of analysis, as most local galaxies are too old to accurately resolve their earlier star formation history (Gallazzi and Bell, 2009). The look-back approach compares images of galaxies at different cosmic epochs, to trace the evolution of galaxies as a population (Perez-Martinez et al., 2020). The difficulty with this approach is in linking the galaxies with their high-redshift progenitors. A similar survey has already been conducted with SDSS data, and STePS aims to extend this back further to ~ 7 Gyrs.

WEAVE will obtain spectra over a redshift range of $0.3 < z < 0.7$ for ~ 1000 galaxies to measure the age of stellar components, the star formation timescale, metal abundances, and AGN activity in each galaxy. The wide wavelength range of WEAVE includes the NUV and U rest bands, so that the youngest star components in galaxies can be constrained.

A.1.6 LOFAR

LOFAR was built with the purpose of exploring poorly studied regions of MHz frequencies to understand galaxy evolution and areas such as the epoch of reionisation and AGN activity. WEAVE has been chosen as the main source of spectra for LOFAR, as typical photometric redshifts are inadequate. This is because a higher S/N is needed for broad band detections and to cover well defined continuum breaks. As a result, Photometric redshifts would be insufficient for low-mass galaxies or those with dust extinction, which is crucial given the evidence that $\sim 85\%$ of star formation is dust obscured between $1 < z < 3$. Recent studies also show that the star formation rate is proportional to the radio luminosity (Bell et al., 2008).

Aside from redshift measurements, the spectral data will be used to distinguish between star forming galaxies and AGN, as well as the AGN accretion modes. To study the merger events which trigger AGN activity, velocity dispersions are required in addition to distinguishing between cluster members and merging galaxies. Additionally, the spectral data can characterise radiatively efficient and inefficient accretion.

The targets have been selected using LOFAR frequencies, to eliminate observational biases caused by dust extinction. The most important region of this survey is at $z \sim 2$, where proto-clusters are predicted to be found. WEAVE will observe 3300 galaxies between $1 < z < 3$ in the LR mode. A resolution of at least 1,000 is required to convert a 2D projection into 3D, but 3,000 is desirable to peer through the OH forest at the red end of the spectrum. The LR of WEAVE is higher compared to the resolution of both MUSE and VIRUS, meaning that further constraints can be placed on galactic models.

The WEAVE-LOFAR survey will be the primary source of samples over the Northern Euclid deep field area. The LIFU has a similar resolution to LOFAR, so comparisons can be easily made between maps to quantify the distribution of 10^6 K plasma which gives off cool gas. In comparison to the other magnitude limited surveys,

this study has a vast discovery space and it is likely that unexpected objects will be encountered.

A.1.7 Quasi-Stellar Objects

The WEAVE Quasi-Stellar Object survey (QSO) was established to understand why the expansion of our universe is accelerating, and how stars form in intergalactic gas. WEAVE will observe $\sim 400,000$ high redshift quasars over a $10,000$ deg^2 area to make tomographic measurements of the intergalactic medium (IGM) and large scale structures, which requires an extensive study of the Lyman α forest. As a result, all quasars chosen have a redshift $z > 2.1$.

The parameters that this survey will measure include the angular diameter distance using Baryonic acoustic oscillations (BAO's), 3D structures such as topology and void abundance, and the small 1D power spectrum including neutrino mass and the IGM temperature.

The main 6000 deg^2 target area follows the same footprint as J-PAS, which contains the region covered by HETDEX. These fields will be shared with the Galactic Archeology survey and will include WEAVE-LOFAR targets found to be quasars. Outside of the J-PAS footprint another region 6000 deg^2 will be observed, with $\sim 20,000$ target quasars. There will be a high completeness for bright quasars, as they will be supplemented with data from GAIA, WISE, and BOSS.

Appendix B

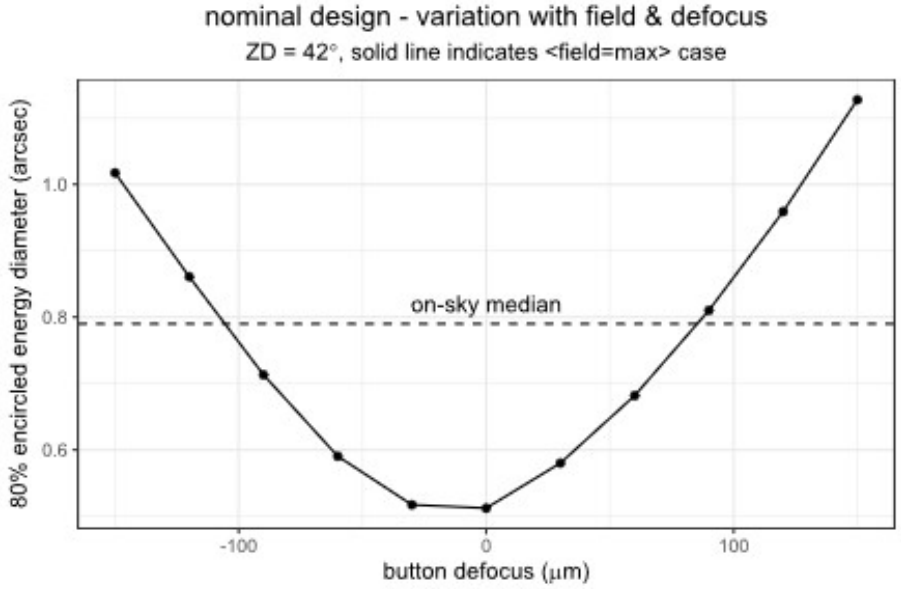
Flexure defocus model

Initial measurements made using both a laser tracker and DTI gauges showed that neither plate A or B were positioned correctly with respect to the WRS interface, and far exceeded the allowed tolerances for the telescope focusing capabilities. A clamp, an aluminium flat with its corresponding reinforcement bracket, and two RSD struts were installed into the tumbler's structure to reduce the magnitude of the displacement. The final results are shown in figure 3.15, however there is still a maximum displacement of 200 μm, which could potentially restrict the range of positions that can be used during observations. This will be analysed during on-sky commissioning.

Despite the preventative measures in place, some flexure behaviour is still observed that potentially restricts the range of orientations that can be used when observing to remain within the focus tolerance. Calculations have been made by Kevin Dee as to how the final flexure results will affect the point spread function (PSF) observed with WEAVE. Previous estimations do not account for any variation in co-planarity of the field plates or any tilt, meaning that they are underestimating the residual sum of squares (RSS) of defocus. Here, the updated RSS, $|\Delta x_{\mathrm{Tot}}|$, is defined as

$$|\Delta x_{\mathrm{Tot}}| = \sqrt{\Delta x_{FP}^2 + \Delta x_0^2}, \tag{B.1}$$

where Δx_{FP}^2 is the measured field defocus and Δx_0^2 is the RSS of defocus without the affect of the field plates.



Produced by Kevin Dee, figure and caption reused with gratitude.

Figure B.1: The encircled energy diameter of the PSF as a function of button defocus, having accounted for the plate position and tilt in the new focus budget. The worst case scenario with nominal seeing conditions is 1.2" at +150 μm defocus.

The previous calculation assumes a symmetrical 2D gaussian profile for the PSF. This assumption underestimates the energy in the wings of the profile, and does not allow the full-width half-maximum (FWHM) to be well defined. As a result, a Moffat function has been used for these calculations, which is defined as

$$f(x, y, \alpha, \beta) = 2\left(\frac{\beta - 1}{\pi\alpha^2}\right)\left[1 + \frac{x^2 + y^2}{\alpha^2}\right]^{-\beta}. \tag{B.2}$$

where β and α are seeing dependent parameters. The result of modelling the PSF as a Moffat function with a FWHM of 0.8" and a β value of 3, then convolving it with typical seeing conditions at the observatory is shown in figure B.1. The worst case scenario is a FWHM of 1.2" at +150 μm defocus. The allowed angular travel range of the WRS system when observing is $\pm 155°$ rotation, and these results do not

imply that any orientations need to be avoided. However, the effect of the flexure behaviour will need careful consideration, particularly as results are received from commissioning and first light.

Bibliography

Aghamousa, A. et al. (2016). "The DESI experiment part I: science, targeting, and survey design". In: *arXiv preprint arXiv:1611.00036*.

Allington-Smith, J. et al. (1990). "Low dispersion survey spectrograph for the William Herschel Telescope". In: *Instrumentation in Astronomy VII*. Vol. 1235. SPIE, pp. 691–701.

Allington-Smith, J. (2006). "Basic principles of integral field spectroscopy". In: *New Astronomy Reviews* 50.4-5, pp. 244–251.

Allington-Smith, J. and R. Content (1998). "Sampling and Background Subtraction in Fiber-Lenslet Integral Field Spectrographs". In: *Publications of the Astronomical Society of the Pacific* 110.752, p. 1216.

Andrade, E. d. C. (1959). "Doppler and the Doppler effect". In: *Endeavour*.

Anguiano, B. et al. (Aug. 2017). "The kinematics of the white dwarf population from the SDSS DR12". English. In: *Monthly Notices of the Royal Astronomical Society* 469.2, pp. 2102–2120.

Apellániz, J. M. et al. (May 2016). "The galactic O-star spectroscopic survey (GOSSS). III. 142 Additional O-type systems". English. In: *The Astrophysical Journal Supplement Series* 224.1, p. 4.

Appenzeller, I. (2013). *Introduction to astronomical spectroscopy*. Cambridge University Press.

Appleton, P. N. et al. (Jan. 7, 2023). “Multi-Phase Gas Interactions on Subarcsec Scales in the Shocked IGM of Stephan’s Quintet with JWST and ALMA”. In: arXiv:2301.02928. preprint.

Bell, E. F. et al. (June 2008). “The Accretion Origin of the Milky Way’s Stellar Halo”. In: *The Astrophysical Journal* 680.1, pp. 295–311.

Bertsimas, D. and J. Tsitsiklis (1993). “Simulated annealing”. In: *Statistical science* 8.1, pp. 10–15.

Bonnarel, F. et al. (Apr. 2000). “The ALADIN Interactive Sky Atlas: A Reference Tool for Identification of Astronomical Sources”. In: *Astronomy and Astrophysics Supplement Series* 143.1, pp. 33–40.

Bryant, J. J. et al. (Mar. 2015). “The SAMI Galaxy Survey: instrument specification and target selection”. English. In: *Monthly Notices of the Royal Astronomical Society* 447.3, pp. 2857–2879.

Bryant, J. J. et al. (2016). “Hector: a new massively multiplexed IFU instrument for the Anglo-Australian Telescope”. In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Vol. 9908. SPIE, pp. 375–393.

Bundy, K. et al. (2015). “Overview of the SDSS-IV MaNGA Survey:mapping nearby galaxies at Apache Point Observatory”. English. In: *The Astrophysical Journal* 798, Bundy, K, Bershady, M A, Law, D R, Yan, R, Drory, N, MacDonald, N, Wake, D A, Cherinka, B, Sánchez–637X/798/1/7.

Campbell, L. et al. (June 2004). “The Tiling Algorithm for the 6dF Galaxy Survey”. In: *Monthly Notices of the Royal Astronomical Society* 350.4, pp. 1467–1476.

Castro-Ginard, A. et al. (July 2019). “Hunting for open clusters in Gaia DR2: the Galactic anticentre”. In: *Astronomy and Astrophysics* 627, A35, A35.

Colless, M. et al. (Dec. 21, 2001). “The 2dF Galaxy Redshift Survey: Spectra and Redshifts”. In: *Monthly Notices of the Royal Astronomical Society* 328.4, pp. 1039–1063.

Dalmases, F. et al. (2018). "The WEAVE prime focus correction:from design to integration". In: 10706.

Dalton, G. et al. (2014). "Project Overview and Update on WEAVE: The next Generation Wide-Field Spectroscopy Facility for the William Herschel Telescope". In: *Proceedings Volume 9147, Ground-based and Airborne Instrumentation for Astronomy V*. Vol. 9147. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. SPIE, p. 91470L.

Dalton, G. et al. (2020). "Integration and Early Testing of WEAVE: The next-Generation Spectroscopy Facility for the William Herschel Telescope". In: *Ground-Based and Airborne Instrumentation for Astronomy VIII*. Ed. by C. J. Evans et al. Vol. 11447. SPIE, pp. 191–200.

Dressler, A. (1984). "Studying the internal kinematics of galaxies using the calcium infrared triplet". In: *The Astrophysical Journal* 286, pp. 97–105.

Drew, J. et al. (2009). *Report by the ETSRC on Europe's 2-4m OIR telescopes over the next decade*.

Gaensicke, B. T. et al. (2010). "Two White Dwarfs with Oxygen-Rich Atmospheres". In: *Science* 327.5962, pp. 188–190.

Gaia Collaboration et al. (Nov. 2016). "The Gaia mission". In: *A&A* 595, A1, A1.

Gallazzi, A. and E. F. Bell (Dec. 2009). "Stellar Mass-to-Light Ratios from Galaxy Spectra: How Accurate Can They Be?" English. In: *The Astrophysical Journal Supplement Series* 185.2, pp. 253–272.

Gambardella, L. M. and M. Dorigo (Jan. 1, 1995). "Ant-Q: A Reinforcement Learning Approach to the Traveling Salesman Problem". In: *Machine Learning Proceedings 1995*. Ed. by A. Prieditis and S. Russell. San Francisco (CA): Morgan Kaufmann, pp. 252–260.

Gansicke, B. T. et al. (Sept. 2007). "SDSS J104341.53+085558.2: a second white dwarf with a gaseous debris disc". In: *Monthly Notices of the Royal Astronomical Society: Letters* 380.1, pp. L35–L39.

Garmany, C. et al. (Dec. 1982). "The initial mass function for massive stars." In: *The Astrophysical Journal* 263, pp. 777–790.

Gentile, F. et al. (Apr. 2015). "A photometric selection of white dwarf candidates in Sloan Digital Sky Survey Data Release 10". English. In: *Monthly Notices of the Royal Astronomical Society* 448.3, pp. 2260–2274.

Gilbert, J. et al. (2012). "Starbugs: all-singing, all-dancing fibre positioning robots". In: *Modern Technologies in Space-and Ground-based Telescopes and Instrumentation II*. Vol. 8450. SPIE, pp. 446–459.

Gilbert, J. M. (2016). *New developments in robotic fibre positioning for astronomical multi-object spectroscopy*. eng.

Gillingham, P. R. et al. (2000). "Echidna: a multifiber positioner for the Subaru prime focus". In: *Optical and IR Telescope Instrumentation and Detectors*. Vol. 4008. SPIE, pp. 1395–1403.

Gray, P. M. and K. Taylor (1990). "Fiber optic instrumentation for the wide-field AAT prime focus: the 2dF". In: *Instrumentation in Astronomy VII*. Ed. by D. L. Crawford. Vol. 1235. International Society for Optics and Photonics. SPIE, pp. 709–720.

Grefenstette, J. et al. (1985). "Genetic Algorithms for the Traveling Salesman Problem". In: *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*. Psychology Press.

Hearnshaw, J. (2009). *Astronomical spectrographs and their history*. Cambridge University Press.

Hearnshaw, J. B. (1990). *The analysis of starlight: one hundred and fifty years of astronomical spectroscopy*. CUP Archive.

Hecht, E. (1998). *Optics. Hecht*. Addison-Wesley.

Helmi, A. and S. D. M. White (Aug. 1999). "Building up the stellar halo of the Galaxy". English. In: *Monthly Notices of the Royal Astronomical Society* 307.3, pp. 495–517.

Hill, G. et al. (2008). "The Hobby-Eberly telescope dark energy experiment (HETDEX): description and early pilot survey results". In: *arXiv preprint arXiv:0806.0183*.

Hill, G. J. et al. (2020). "VIRUS2: a next generation replicated integral field spectrograph with wide field and broad wavelength coverage". In: *Ground-based and Airborne Instrumentation for Astronomy VIII*. Vol. 11447. SPIE, pp. 208–221.

Hill, J. (1988). "The history of multiobject fiber spectroscopy". In: *Fiber Optics in Astronomy*. Vol. 3, p. 77.

Hill, J. and M. Lesser (1986). "Deployment Of The Mx Spectrometer". In: *Instrumentation in Astronomy VI*. Ed. by D. L. Crawford. Vol. 0627. International Society for Optics and Photonics. SPIE, pp. 303–320.

Hu, H. et al. (2003). "New type of optical fiber positioning unit device for LAMOST". In: *Large Ground-based Telescopes*. Vol. 4837. SPIE, pp. 548–555.

Hughes, S. et al. (2020). "Final assembly, metrology, and testing of the WEAVE fibre positioner". In: *Ground-based and Airborne Instrumentation for Astronomy VIII*. Vol. 11447. SPIE, pp. 1678–1692.

Hughes, S. et al. (2022a). "Calibration at elevation of the WEAVE fibre positioner". In: *Ground-based and Airborne Instrumentation for Astronomy IX*. Vol. 12184. SPIE, pp. 2067–2082.

Hughes, S. et al. (2022b). "Optimisation of the WEAVE target assignment algorithm". In: *Ground-based and Airborne Instrumentation for Astronomy IX*. Vol. 12184. SPIE, pp. 2048–2066.

IAU SOFA Board (2021). *IAU SOFA Software Collection*.

Iglesias-Páramo, J. et al. (Mar. 2012). "New Insights on Stephan's Quintet: Exploring the Shock in Three Dimensions". In: *Astronomy & Astrophysics* 539, A127.

Izazaga, R. et al. (July 2018a). "WEAVE Spectrograph Cameras: The Polishing of the Spherical Lenses". In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III*. Vol. 10706. SPIE.

Izazaga, R. et al. (July 2018b). "The Polishing of WEAVE Spectrograph Collimator Mirror". In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III*. Vol. 10706. SPIE.

Jin, S. et al. (2022). "The wide-field, multiplexed, spectroscopic facility WEAVE: Survey design, overview, and simulated implementation". In: *arXiv preprint arXiv:2212.03981*.

Kirkpatrick, S. et al. (May 13, 1983). "Optimization by Simulated Annealing". In: *Science* 220.4598, pp. 671–680.

Krabbe, A. et al. (1997). "3D: a new generation imaging spectrometer". In: *Optical Telescopes of Today and Tomorrow*. Vol. 2871. SPIE, pp. 1179–1186.

Lagrange, A.-M. and J. Andersen (Aug. 2006). "ASTRONET: strategic planning for European astronomy". English. In: *Proceedings of the International Astronomical Union* 2.14, pp. 537–538.

Lee, H. and M. K. Poutous (Oct. 2022). "Reactive Ion Plasma Etched Surface Relief Gratings for Low/Medium/High Resolution Spectroscopy in Astronomy". In: *Journal of Astronomical Telescopes, Instruments, and Systems* 8.4, p. 045002.

Lewis, I. J. et al. (June 2002). "The Anglo-Australian Observatory 2dF facility". In: *Monthly Notices of the Royal Astronomical Society* 333.2, pp. 279–298.

Lindegren, L. et al. (2007). "The Gaia mission: science, organization and present status". In: *Proceedings of the International Astronomical Union* 3.S248, pp. 217–223.

Lynden-Bell, D. (Aug. 21, 2002). "Exact Optics: A Unification of Optical Telescope Design". In: *Monthly Notices of the Royal Astronomical Society* 334.4, pp. 787–796.

Massey, P. and M. M. Hanson (2013). "Astronomical spectroscopy". In: *Planets, Stars and Stellar Systems* 2, pp. 35–98.

Metropolis, N. et al. (1953). "Equation of State Calculations by Fast Computing Machines". In.

Middleton, K. (2018). *WEAVE-SYS-017: Positioning Error Budget at Prime Focus*.

Mignot, S. et al. (2020a). "Fiber Links for the WEAVE Instrument: The Making Of". In: *Modeling, Systems Engineering, and Project Management for Astronomy IX*. Ed. by G. Z. Angeli and P. Dierickx. Vol. 11450. SPIE, 114502F.

— (2020b). "Fiber links for the WEAVE instrument: the making of". In: *Modeling, Systems Engineering, and Project Management for Astronomy IX*. Ed. by G. Z. Angeli and P. Dierickx. Vol. 11450. International Society for Optics and Photonics. SPIE, 114502F.

Miszalski, B. et al. (Aug. 2006). "Multi-object spectroscopy field configuration by simulated annealing". In: *Monthly Notices of the Royal Astronomical Society* 371.4, pp. 1537–1549.

Monari, G. et al. (Apr. 2016). "Modelling the Galactic disc: perturbed distribution functions in the presence of spiral arms". English. In: *Monthly Notices of the Royal Astronomical Society* 457.3, pp. 2569–2582.

Moore, A. M. et al. (2003). "Spine development for the Echidna fiber positioner". In: *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*. Vol. 4841. SPIE, pp. 1429–1439.

Nataf, D. M. et al. (Mar. 2016). "Interstellar extinction curve variations towards the inner Milky Way: a challenge to observational cosmology". English. In: *Monthly Notices of the Royal Astronomical Society* 456.3, pp. 2692–2706.

Parry, I. and R. Sharples (1988). "AUTO FIB Current Status". In: *Fiber Optics in Astronomy*. Vol. 3, p. 93.

Parry, I. (2006). "Optical fibres for integral field spectroscopy". In: *New Astronomy Reviews* 50.4-5, pp. 301–304.

Parry, I. R. et al. (1994). "Autofib–2: an automated fiber positioner for the prime focus of the William Herschel Telescope". In: *Instrumentation in Astronomy VIII*. Vol. 2198. SPIE, pp. 125–133.

Perez-Martinez, J. M. et al. (Feb. 2020). "Galaxy kinematics across different environments in the RXJ1347-1145 cluster complex". English. In: *Astronomy and Astrophysics*.

Pillow, J. (2018). "Lecture 3B Notes: Least Squares Regression". In.

Rogers, K. et al. (2014). "The design of the WEAVE spectrograph". In: *Ground-based and Airborne Instrumentation for Astronomy V*. Ed. by S. K. Ramsay et al. Vol. 9147. International Society for Optics and Photonics. SPIE, pp. 2043–2057.

Rohlfs, K. and T. L. Wilson (2004). *Tools of radio astronomy*. Springer.

Saff, E. B. and A. B. J. Kuijlaars (Dec. 1997). "Distributing Many Points on a Sphere". In: *The Mathematical Intelligencer* 19.1, pp. 5–11.

Sayède, F. et al. (2016). "First results of tests on the WEAVE fibres". In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II*. Vol. 9912. SPIE, pp. 674–686.

Schallig, E. (2019). *Metrology, calibration, integration and testing of the WEAVE fibre positioner*.

Schonrich, R. and P. J. McMillan (Jan. 2017). "Understanding inverse metallicity gradients in galactic discs as a consequence of inside-out formation". English. In: *Monthly Notices of the Royal Astronomical Society* 467.1, pp. 1154–1174.

Schrijver, A. et al. (2003). *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer.

Sellwood, J. A. and J. J. Binney (Nov. 2002). "Radial mixing in galactic discs". English. In: *Monthly Notices of the Royal Astronomical Society* 336.3, pp. 785–796.

Smith, D. J. B. et al. (Nov. 8, 2016). *The WEAVE-LOFAR Survey.*

Smith, G. A. et al. (2004). "AAOmega: a multipurpose fiber-fed spectrograph for the AAT". In: *Ground-based Instrumentation for Astronomy.* Ed. by A. F. M. Moorwood and M. Iye. Vol. 5492. International Society for Optics and Photonics. SPIE, pp. 410–420.

Stuik, R. et al. (2018). "Integration and testing of the WEAVE Spectrograph". In: *Proc. SPIE.* Vol. 10702, pp. 10702–275.

Terrett, D. L. et al. (2014). "Fibre positioning algorithms for the WEAVE spectrograph". In: *Software and Cyberinfrastructure for Astronomy III.* Ed. by G. Chiozzi and N. M. Radziwill. Vol. 9152. International Society for Optics and Photonics. SPIE, pp. 216–224.

Tomás, A. et al. (July 2018). "The WEAVE prime focus correction: from design to integration". In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III.* Vol. 10706. International Society for Optics and Photonics. SPIE.

Trager, S. et al. (2013). *WEAVE-SCI-001: WEAVE science requirements.*

Trager, S. C. (Nov. 18, 2019). *The WEAVE Survey Plan.*

Tremblay, P.-E. et al. (Mar. 2017). "The Gaia DR1 mass–radius relation for white dwarfs". English. In: *Monthly Notices of the Royal Astronomical Society.*

Van Cappellen, W. et al. (2022). "Apertif: Phased array feeds for the Westerbork Synthesis Radio Telescope-System overview and performance characteristics". In: *Astronomy & Astrophysics* 658, A146.

van Haarlem, M. P. et al. (2013). "LOFAR: The LOw-Frequency ARray". In: *A&A* 556, A2.

Verheijen, M. A. W. et al. (2009). "Apertif, a focal plane array for the WSRT". English. In: *Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array*. AMER INST PHYSICS, pp. 265–271.

Verheijen, M. et al. (2008a). "Apertif, a focal plane array for the WSRT". In: *AIP Conference Proceedings*. Vol. 1035. 1. American Institute of Physics, pp. 265–271.

Verheijen, M. et al. (2008b). "The evolution of galaxies through the neutral hydrogen window". In: *AIP Conference Proceedings*. Vol. 1035, pp. 265–271.

Vicente, A. S. et al. (2020). "New prime focus rotator system for the WHT". In: *Ground-based and Airborne Telescopes VIII*. Ed. by H. K. Marshall et al. Vol. 11445. International Society for Optics and Photonics. SPIE, 114454S.

Wang, X. et al. (Sept. 1, 2013). "Effect of Aging Temperature on Dimensional Instability of Invar Alloys". In: *Optoelectronics and Advanced Materials, Rapid Communications* 7, pp. 754–759.

WEAVE First Light (2023). URL: `https://www.ing.iac.es/PR/press/weave_LIFU_first_light.html` (visited on 03/02/2023).

Wegg, C. and E. S. Phinney (Oct. 2012). "White dwarf kinematics versus mass". English. In: *Monthly Notices of the Royal Astronomical Society* 426.1, pp. 427–439.

Wynne, C. (Nov. 1, 1982). "A More Efficient Faint-object Spectrograph". In: *Optica Acta: International Journal of Optics* 29.11, pp. 1557–1571.

Xue, X.-X. et al. (Mar. 2014). "The SEGUE K giant survey II. A catalog of distance determinations for the SEGUE K giants in the galactic halo". English. In: *The Astrophysical Journal* 784.2, p. 170.