

Metrology, calibration, integration and testing of the WEAVE fibre positioner



Ellen Schallig
Magdalen College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2019

Voor Marten,

mede-avonturier en zekeraar

Acknowledgements

I would like to thank my supervisor Gavin Dalton for always taking the time to discuss research, help out with hardware, and read everything that I sent, even if it was on ridiculously short deadlines. You have always been very positive and motivational. Co-supervisor Ian Lewis, thank you for the support on the positioner and the calibration software. Both of you made me feel very welcome in the WEAVE group in Oxford.

Friends in Oxford, particularly the inhabitants of office 650 over the past few years (Becky, Francesca, Gus, Josie, Laura, Marisa, Matthew, Paul), other friends at the department, and friends met at the College exchange dinners (Alexandru, Alina), thank you for making these past years in Oxford so much fun.

Renée and Niels, fellow Groningen Dutchies, never compromise on being Dutch. Jamie, our time in Oxford didn't overlap that much, but your experiences really helped me, especially at the start and end of this journey. Shoko, our late-evening chats at different WEAVE meetings were invaluable. I would also like to thank my hockey friends, lots of random Wikipedia articles, and comics¹ for some much-needed distraction from doing research, and stroopwafels for being the best snack.

Thomas, Annemieke, thank you for reading through most, if not all of my thesis. The last few months of writing in the Netherlands were definitely more bearable because of your proximity and support.

I am very thankful for the trailblazing women before me, in particular Groningen feminist Aletta Jacobs, the first Dutch woman to complete a university degree. Proposition XI in her PhD thesis should be read by more people, even today ([Jacobs, 1879](#)).

Mam, pap, and Manon: your love, encouragement, and belief in me have always been so important. And lastly Marten, your support has meant the world to me. Thank you for following me to a different country and supporting me in all my endeavours. Cycling to work together in Oxford will be something to always remember.

¹In particular PHDComics ([Cham, 1997-2019](#)) and xkcd ([Munroe, 2006-2019](#)).

Abstract

This thesis is about a pick-and-place fibre positioner for WEAVE, a new fibre-fed multi-object spectrograph for the 4.2 m William Herschel Telescope with a multiplex of 960. First light is expected sometime in 2020. The research presented in this thesis has been a major step in ensuring that the fibre positioner will be able to perform to the design specifications, and meet WEAVE's scientific requirements.

The first piece of work presented in this thesis concerns the full positioner build, and the errors that are propagated by stacking components of the positioner. The guide to building a positioner shows that these errors can be measured and tracked, and ultimately compensated for.

The tumbler assembly can be built to a high precision, but as long as the robot positioners do not move accurately, it is impossible to position fibres in the correct places on the field plate. One reason why this movement can be inaccurate, is that the coordinate systems of the targets, the robots, and the field plates are not expressed in each other correctly. Another is that the positioning structure itself changes under external influences such as temperature and gravity. A novel algorithm ensures that the positioning structure can be recalibrated periodically, removing these kinds of effects.

Lastly, if the robot gantries cannot make repeatable accurate movements, this is a show-stopping problem too. Test procedures developed during this DPhil have shown that the X - and Y -gantries perform as expected, but the Z -gantry introduces unwanted movements in the x, y -plane when moving in the vertical direction. Because of the test procedures this design flaw was discovered and rectified by a new Z -gantry with a different internal structure. With these changes, the positioner robots are found to be repeatable to within the design specifications.

Contents

| | | |
|----------|---|-----------|
| 1 | Spectroscopy in astronomy | 1 |
| 1.1 | The science of spectroscopy | 1 |
| 1.2 | The spectrograph | 5 |
| 1.3 | Fibre spectroscopy | 7 |
| 1.4 | Multi-object spectroscopy | 7 |
| 1.4.1 | Integral field units | 8 |
| 1.4.2 | Automated positioning | 9 |
| 1.5 | Pick-and-place technology for the Autofib-2 and 2dF instruments . . | 10 |
| 1.6 | This research | 13 |
| 1.7 | Summary | 14 |
| 2 | Introduction to WEAVE | 15 |
| 2.1 | WEAVE subsystems | 16 |
| 2.1.1 | The prime focus corrector | 17 |
| 2.1.2 | The spectrograph | 18 |
| 2.1.3 | The fibre positioner | 20 |
| 2.2 | Science goals | 21 |
| 2.2.1 | Galactic archaeology | 21 |
| 2.2.2 | SCIP | 23 |
| 2.2.3 | White dwarfs | 23 |
| 2.2.4 | WEAVE-Apertif | 24 |
| 2.2.5 | WEAVE-LOFAR | 25 |
| 2.2.6 | Galaxy clusters | 26 |
| 2.2.7 | StePS | 27 |
| 2.2.8 | WEAVE-QSO | 27 |
| 2.2.9 | WEAVE Survey footprint | 28 |
| 2.3 | The WEAVE fibre positioner in detail | 29 |
| 2.3.1 | Positioner components | 29 |

| | | |
|----------|---|-----------|
| 2.3.2 | Positioner error budget | 32 |
| 2.4 | Summary | 35 |
| 3 | Positioner metrology | 38 |
| 3.1 | Methodology | 39 |
| 3.2 | Interface plate and handling trolley | 41 |
| 3.2.1 | Hardware status | 41 |
| 3.2.2 | Plate thickness measurements | 44 |
| 3.2.3 | Plate location measurements | 47 |
| 3.3 | Focal plane imager | 50 |
| 3.3.1 | Depth of focus and seeing | 50 |
| 3.3.2 | FPI plane and ledge | 53 |
| 3.4 | Tumbler | 57 |
| 3.4.1 | Measurement setup | 59 |
| 3.4.2 | Tumbler tilt adjustments | 61 |
| 3.4.3 | Height of RSD | 63 |
| 3.4.4 | Motor repeatability | 66 |
| 3.5 | Distance from field plate to interface | 68 |
| 3.5.1 | Interface plate thickness (C) | 68 |
| 3.5.2 | Top of interface plate to FPI ledge (B) | 68 |
| 3.5.3 | FPI ledge to RSD (A + field plate) | 70 |
| 3.5.4 | All measurements together | 71 |
| 3.6 | Summary | 74 |
| 4 | Robot calibration software | 75 |
| 4.1 | Coordinate systems | 75 |
| 4.2 | Robot limitations | 79 |
| 4.3 | Autofib-2 legacy | 81 |
| 4.4 | Differences between calibration approaches of Autofib-2 and WEAVE | 83 |
| 4.4.1 | Solid body rotation versus full 6-parameter fit | 85 |
| 4.4.2 | Unavailable grid points | 89 |
| 4.5 | Grid simulations | 92 |
| 4.5.1 | Algorithm | 93 |
| 4.5.2 | Separate contributions | 96 |
| 4.5.3 | WEAVE simulation | 100 |
| 4.5.4 | Simulation with only part of the grid available | 105 |
| 4.6 | Grid measurements | 108 |

| | | |
|----------|--|------------|
| 4.6.1 | Offset issues | 110 |
| 4.6.2 | Centroiding issues | 110 |
| 4.6.3 | Final result | 113 |
| 4.6.4 | Result as input for simulation | 116 |
| 4.6.5 | Forward look | 117 |
| 4.7 | Summary | 118 |
| 5 | Positioner performance | 119 |
| 5.1 | Fibre retractor performance | 120 |
| 5.1.1 | Rationale and retractor design | 120 |
| 5.1.2 | Assembly | 123 |
| 5.1.3 | Testing | 125 |
| 5.2 | Gripper optics | 127 |
| 5.2.1 | Plate scale | 128 |
| 5.2.2 | Rotation centre | 131 |
| 5.3 | Delta software | 132 |
| 5.4 | Robot repeatability performance | 135 |
| 5.4.1 | Preliminary results | 135 |
| 5.4.2 | Later understanding | 136 |
| 5.4.3 | Longer measurements | 138 |
| 5.4.3.1 | Servo only | 138 |
| 5.4.3.2 | Moving in x and y | 140 |
| 5.4.4 | Z-motor issues | 144 |
| 5.4.4.1 | Moving in the z-direction: upper focus | 145 |
| 5.4.4.2 | Moving in the z-direction: lower focus | 146 |
| 5.4.4.3 | Motor temperatures | 149 |
| 5.4.5 | New Z-motor | 151 |
| 5.4.6 | Conclusions and caveats | 152 |
| 5.5 | Summary | 153 |
| 6 | Discussion | 155 |
| 6.1 | Discussion of this work | 155 |
| 6.2 | Further work on the positioner | 158 |
| | Bibliography | 160 |

Glossary

| | |
|-----------|--|
| 2dF | Two-degree field, a fibre-fed multi-object spectrograph for the AAT with a pick-and-place fibre positioner. iv, 10–13, 29, 37, 77, 83, 84, 100, 120, 121 |
| AAT | Anglo-Australian Telescope. 10, 29, 77 |
| ADC | Atmospheric dispersion corrector. 17 |
| AGN | Active galactic nuclei. 25 |
| Apertif | A focal-plane array system for the Westerbork Synthesis Radio Telescope. 15, 21, 24 |
| APM | Plate measuring machine at Cambridge at the time of Autofib-2 testing. 85 |
| Autofib-2 | Automatic fibre positioner for the WHT. iv, v, 10–13, 29, 77, 80, 82, 83, 85, 86, 88–90, 93, 95, 100, 102, 115 |
| CAD | Computer-aided design. 58, 144 |
| CNC | Computer numerical control. 108, 117 |
| COSMOS | Plate measuring machine at Edinburgh at the time of Autofib-2 testing. 85 |
| FARO | 3D-measurement systems company. 39, 41, 42, 44, 47, 49, 53–55, 57, 58, 62, 65 |
| FEA | Finite element analysis. 36, 100 |
| Fortran | The Fortran programming language. A general-purpose compiled imperative programming language. 81, 93 |
| FPI | Focal plane imager, a robot unit on the WEAVE fibre positioner. 41, 50, 53, 55, 56, 58, 70, 75–77, 84 |

| | |
|-------------|---|
| FTU | Focus translation unit. 17 |
| Gaia | Space observatory designed for high-precision astrometry, aiming to observe one billion astronomical objects, mainly in the Milky Way. 15, 21, 22, 33 |
| GHRIL | Ground-based high resolution imaging laboratory, enclosure on one of the Nasmyth platforms on the WHT. 19, 20 |
| HETDEX | Hobby-Eberly Telescope dark energy experiment, also the area on the sky that this experiment targets. 27 |
| IFU | Integral field unit. 8, 16, 20, 24, 26 |
| IGM | Intergalactic medium. 27 |
| LED | Light-emitting diode. 77, 83 |
| LIFU | Large integral field unit, a WEAVE observing mode. 16, 20, 24–26, 66, 158 |
| LOFAR | Low-Frequency Array, a large radio telescope network. 15, 21, 25–27 |
| miniIFU | mini-IFU, 20 movable WEAVE units that can be positioned on the field plate by the WEAVE robot positioners. 20, 31, 32, 124 |
| Morta | One of the robotic positioners on WEAVE. 75–77, 135, 138, 153, 154, 157 |
| MOS | Multi-object spectroscopy. 7, 16, 20, 26, 27, 31, 32, 119, 120, 123, 124, 153, 157, 158 |
| Nona | One of the robotic positioners on WEAVE. 75–77, 109, 129, 138, 140, 153, 157 |
| PEEK | Polyether ether ketone, robust material used for tubing in the WEAVE context. 122, 123 |
| PFC | The WEAVE prime focus corrector subsystem. 17, 34, 35, 72 |
| PFR | Prime focus field rotator. 17 |
| POS | The WEAVE fibre positioner subsystem. 17, 36, 40 |
| prime focus | Focal surface of a telescope’s primary mirror. 7, 10 |

| | |
|-------|--|
| PSF | Point spread function. 34, 52 |
| QSO | Quasar (from ‘quasi-stellar object’), as WEAVE-QSO one of the eight WEAVE surveys. 27 |
| RMS | Root mean square, the error in each measurement of the same thing. 36, 58, 93, 96, 98, 99, 101–105, 107, 108, 113–115, 117, 118, 135, 136, 139–141, 146, 147, 151–154, 157 |
| RSD | Retractor support disc, a component in the WEAVE fibre positioner. 58, 59, 63–65, 67–69, 71–73 |
| RSS | Root sum of squares, the way to combine uncorrelated contributors to an overall uncertainty. 33, 35, 36, 61, 70, 117 |
| SCIP | Stellar, Circumstellar, and Interstellar Physics survey, one of the eight WEAVE surveys. 23 |
| SDSS | Sloan Digital Sky Survey. 1, 9 |
| SKA | Square Kilometer Array, future radio telescope network. 24, 25 |
| SNR | Signal-to-noise ratio. 32, 33, 35, 37 |
| WEAVE | WHT Enhanced Area Velocity Explorer, new fibre-fed multi-object spectrograph for the WHT, with two pick-and-place fibre positioners working in tandem. ii, iii, 13, 15–18, 20–29, 33–35, 37, 77, 83, 84, 86, 88, 89, 93, 94, 100, 102, 112, 120, 121, 130, 149, 155, 159 |
| WHT | William Herschel Telescope. 10, 15, 16, 19, 29, 33–35, 50–52, 77, 158, 159 |

Chapter 1

Spectroscopy in astronomy


Spectroscopy is an essential technique used in astronomy today. The notion that colour is an intrinsic property of light rays (Newton, 1671/72), and does not arise from the effect of light ‘pulses’ on the surroundings¹ is seen as the first successful foray into this. Spectroscopic analysis has completely transformed the field of astronomy from a mainly precision positional science, to also include what is now called ‘astrophysics’: understanding how the stars and galaxies around us form, evolve, and die.

1.1 The science of spectroscopy

Spectroscopy is used to learn more about the chemistry of celestial objects, their velocities, and, if they are at cosmological distances, their redshifts. This is done by examining an object’s spectrum: a plot of its light intensity as a function of wavelength. As an object such as a star emits light from its core, elements in its exterior absorb certain wavelengths. The remaining light passes through and some of it eventually hits a telescope. An example of such a spectrum can be found in figure 1.1². The

¹René Descartes and Robert Hooke, among others, explained light and colour in terms of pulses that are transmitted through the aether. The pulses produce different colours by making the aether particles rotate faster or slower (Descartes) or the pulses themselves change (Hooke) (Fara, 2015).

²The spectra in this chapter are random targets from SDSS. With the plate and fibre numbers specified in the plots the spectra can be requested from <https://dr12.sdss.org/basicSpectra>.



specific wavelengths of these absorbed features are compared to laboratory spectra of elements and molecules, and from the specific form of the absorption feature (depth, width and shape) we can find how much of this element is present.

In this way any light emitter (e.g. a star or a galaxy) or object that is between a light emitter and the Earth can be analysed. One can even examine planetary atmospheres, by comparing the starlight when the planet is in front of the star with when the planet is not there.

These abundance measurements require very high-resolution spectra, as they depend on the precise shape of very narrow and shallow features in the spectrum. But a wealth of information can be learned from lower-resolution spectra too.

Taking an image of an object shows us what the object looks like. Taking several images of the same object over a span of time can also show how much the object has moved over the sky, and therefore say something about the speed with which it is moving. However, images do not give any information about the third dimension, the radial direction.

When comparing the features in the spectrum ('spectral lines') to laboratory values³, very often one will find that the lines are not exactly at the same wavelength as the laboratory lines. On top of that, all features are shifted the same amount. This is due to the Doppler effect: for a source going towards Earth, all the light waves are compressed, and for a source travelling away from Earth, the waves are stretched. This compression or stretch manifests as a shift of the whole spectrum.

From this shift, the radial velocity can be calculated:

$$v_{\text{rad}} = \frac{\Delta\lambda}{\lambda_0} \cdot c, \quad (1.1)$$

where v_{rad} is the calculated radial velocity, λ_0 is the wavelength of the spectral line as measured in the laboratory, $\Delta\lambda$ is the shift measured from the spectrum, and c is

³Laboratory values are stationary measurements, that is, stationary with respect to Earth.



the speed of light⁴. A positive radial velocity means that the object is moving away from Earth.

The fraction $\Delta\lambda/\lambda_0$ is also called redshift (or blueshift) because of the shifting of waves towards red (or blue). This can happen because of the objects moving away from or towards the Earth as mentioned above, but also because of the expansion of space. On very large scales, we can measure the effect of galaxies moving away from each other as the space between the galaxies expands. Redshift therefore becomes an indirect measurement of distance, if we assume that the rate of expansion is a constant (the Hubble constant) and make some assumptions about the exact geometry of the Universe.

An example of a spectrum of a galaxy at a considerable distance can be found in the bottom half of figure 1.1. With a back-of-the-envelope approach we can approximate the redshift for this galaxy. First assume that the star of the top spectrum in the figure is close to Earth, and its spectrum can be taken as the laboratory values. Then compare any set of lines with the same label. In this case, take the H_β -line, which is close to 4850 Å in the star's spectrum, and around 5600 Å in the galaxy's spectrum⁵. Then apply the formula for the redshift z :

$$\begin{aligned} z &= \frac{\Delta\lambda}{\lambda_0} = \frac{\lambda_{\text{emitted}} - \lambda_0}{\lambda_0} \\ &= \frac{5600 - 4850}{4850} \\ &\approx 0.15. \end{aligned} \tag{1.2}$$

This is similar to the very carefully calculated redshift quoted in the figure.

The ability to measure a shift depends on an instrument's resolution. The resolution is the smallest difference between two wavelengths that can still be distinguished. In astronomy the resolving power R is usually quoted, which is related to resolution

⁴This particular equation holds true for objects that are travelling with velocities much smaller than the speed of light ($v_{\text{rad}} \ll c$), which is only in the near Universe.

⁵Figuring out which dip is which element line can be one of the hardest things in spectral analysis.



as follows:

$$R = \frac{\lambda}{\Delta\lambda}, \quad (1.3)$$

with $\Delta\lambda$ the resolution and λ the wavelength at which this smallest difference is evaluated. To make matters more confusing, this resolving power R is usually called ‘resolution’ too (even in this thesis).

The velocity of objects in the Milky Way is generally a few hundred km s^{-1} . Performing calculation 1.1 the other way round, to be able to see a radial velocity of 300 km s^{-1} at wavelengths around 500 nm , one needs to be able to see shifts of about 0.5 nm , which is equivalent to $R = 1000$. For smaller radial velocities, which occur often as most objects in the Milky Way do not travel radially outwards from the Earth, the shift would be smaller. A radial velocity of 30 km s^{-1} would mean to be able to measure a shift of 0.05 nm , or $R = 10000$.⁶

Spectra are essential for modern astrophysics. From abundances and velocities of groups of stars or galaxies, we can infer many different things about where and how they formed, and under what conditions they must have evolved into the objects that we currently see. We can even infer some properties of the Universe itself, from how the structures at the largest scales move. As such, spectra have opened the doors to studying the history of the Universe.

1.2 The spectrograph

To measure an object’s spectrum, the light from that object has to be dispersed and recorded. The simplest astronomical spectrographs have a dispersion element and a camera at the telescope. To ensure that only the object’s spectrum is measured, and not of the surrounding area, a mask can be placed in the telescope focus, with a hole

⁶When more lines are used for analysis (which is perfectly doable for instruments with long wavelength ranges), the effective resolution can be increased with the same hardware, and thus the minimum radial velocity that can be observed is decreased.



or a slit at the position where the light of the object falls on the focus.

There are different ways to select the objects in the focal plane. The simplest is a small hole in an otherwise fully blocked focus. After the starlight has gone through the dispersion element, it spreads into a spectrum on the detector. One dimension of the detector is chosen in such a way that the whole spectral area of interest fits on it. The other dimension can be as small as the hole. A better way to fill a detector is therefore by using a slit: a thin slice of the sky fills one direction of the detector (getting spatial resolution information), while the other direction is taken up by the spectra.

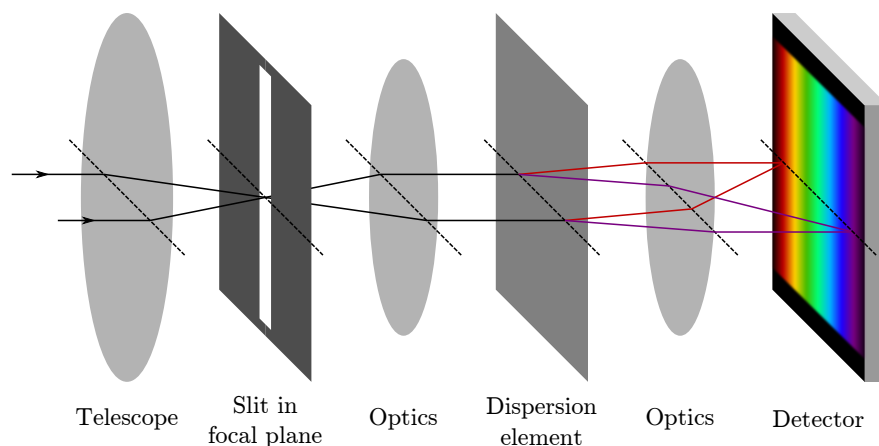


Figure 1.2: A simple spectrograph. Incoming light from the telescope is dispersed into a spectrum and falls onto a detector. Several optical elements ensure that the light is either collimated or focussed. The specific dispersion element, such as a diffraction grating, determines the resolution of the spectrograph.

Slit spectroscopy has been used for centuries. It is directly limited by the size of the detector, and only the objects that fall in the narrow slit are observed, plus possibly a large amount of emptiness around these objects. It can take a long time to observe any sizeable area of the sky in this way, but it is a very straightforward technique to use for high-precision observations of single objects. It also directly obtains spatial resolution in the direction of the slit. This is a good reason to keep the technique in use today.



1.3 Fibre spectroscopy


From a hardware point of view, another downside to the slit spectrograph is that the spectrograph has to be on the telescope, aligned to whichever focus of the telescope is used. Usually this means that the spectrograph is attached at the back of the primary mirror, in the telescope's Cassegrain focus. This puts restrictions on the size and weight of the instrument.

It would be convenient if the spectrograph could be placed off the telescope, allowing larger sizes and weights for larger detectors and optics. This became possible with the advent of optical fibres, where the light is trapped in flexible glass fibres. The efficiency through these fibres is very high: 90–95% transmission for wide-band fibres, and close to 100% for fibres optimised for narrow wavelength bands, for several tens of meters ([Sayède et al., 2015](#)). They can be used to transport the light from the telescope prime focus to a site off the telescope.

1.4 Multi-object spectroscopy

Another advantage of using fibres is that instead of being restricted to a slit shape to observe the sky with, now the fibres can be placed anywhere in the telescope focal plane, as long as the other end is still arranged compactly in a slit shape. The detector area is now used very efficiently, as all the fibres can be placed on interesting objects, and not too much of the detector area is lost on the emptiness of space.

Fibre-fed multi-object spectroscopy (MOS) has proved very successful. The flexibility of spectrograph placement has meant that the trade-off between performance and weight is now balanced towards better performance. The number of objects observed simultaneously is called the ‘multiplex’ of the instrument, and modern-day MOS instruments can have multiplexes in the thousands. This is useful, as astronomers want to see more of the Universe, and with better instruments, all the time.



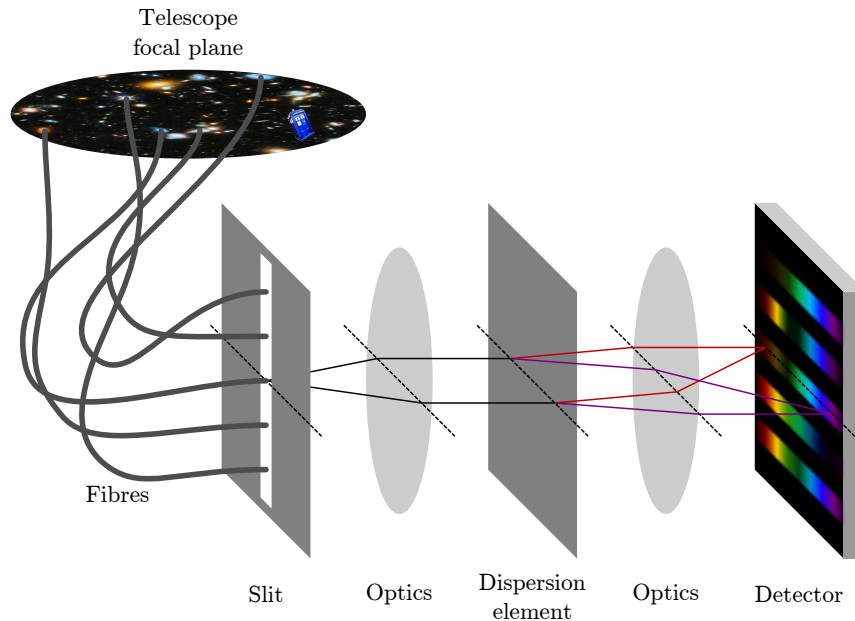


Figure 1.3: Fibres give the option of a larger distance between the telescope and the spectrograph, enabling different designs for spectrographs. They also make much more efficient use of the detector.

1.4.1 Integral field units

To be able to easily observe extended objects as a function of position or highly clustered objects in a small field, integral field units (IFUs) were developed. One of the methods is to bundle fibres together in a very closely packed array, thereby simultaneously obtaining positional information and spectra. This allows for a much more detailed study of extended objects, being able to distinguish between different areas. It also allows for a better sampling of the highly clustered fields, whereas single fibres would not be able to be placed so close together.

Some very successful surveys with fibre-fed IFU instruments are the SAMI survey (Croom et al., 2012) and the MaNGA survey (Bundy et al., 2015). The fibre bundles range from very few fibres to tens of thousands (e.g. VIRUS, Tuttle et al. (2016)).



1.4.2 Automated positioning

One of the best-known spectral surveys is the spectroscopic mode of the SDSS (Sloan Digital Sky Survey) (York et al., 2000; Owen et al., 1994), which started in 2000, and is currently in phase IV. The spectrograph started with 640 fibres, and nowadays has a multiplex of 1000. Up to 2014, the 640 fibres were manually plugged into plates (‘plug plates’) with pre-drilled holes at the exact locations of the objects to be observed. There are several sets of fibres, therefore making it possible to prepare several plates in advance. At the end of an observation, the next prepared plate can be easily switched into the focus and the fibres can be plugged into a new plug plate.

Of course this is very labour-intensive, and not very flexible. The holes for an observation are drilled for a planned time, and if the observation is not completed at that time, the holes do not line up with the targets anymore due to changing atmospheric refraction and the plate becomes useless. The labour involved in drilling and plugging the plates is massive, and requires many people at an easily-accessible site. Several ways have been devised to provide alternatives to this. The first robotic instrument was the MX spectrometer (Hill and Lesser, 1986), which had an arrangement where 32 stepper motors could position 32 fibres simultaneously. A reconfiguration took only 90 seconds, but the multiplex was very low.

Around the same time, ‘pick-and-place’ robotic instruments started being constructed, with Autofib (Parry et al., 1986; Parry and Sharples, 1988) commissioned on the Anglo-Australian Telescope in 1987 as alternative to the MX spectrometer. Its 64 fibres were positioned sequentially, because there was only one robot to move the fibres. However, these fibres could now be put anywhere in the focal plane of the telescope.

Other techniques also became available. Two important ones are tilting spine positioners and radial arm positioners. Fibres with tilting ends (‘tilting spines’ or ‘Echidna’, Kimura et al. (2010); Akiyama et al. (2008)) can all be actuated at the same



time, and the fibres can be re-angled very quickly for the next observation. They are usually in densely-packed configurations, and each fibre is limited in positioning area by the length of its spine. Large tilts can result in defocus. Radial arm positioning is a similar technique where all the fibres can be re-positioned at the same time, but instead of tilting the fibres, each fibre has two small motors which rotate the fibre end and a fibre ‘elbow’ in place. The fibre density is reasonably high, but the fibres remain relatively isolated because of the rotation motors.

This thesis is about a pick-and-place fibre positioner, so that is where the focus will remain.

1.5 Pick-and-place technology for the Autofib-2 and 2dF instruments

Two instruments should be highlighted in this thesis about a pick-and-place fibre positioner. The first is Autofib-2, commissioned on the William Herschel Telescope (WHT) in the mid-nineties (Parry et al., 1994) with 150 fibres, the second is 2dF, built for the Anglo-Australian Telescope (AAT) in the early 2000s (Lewis et al., 2002) with a multiplex of 400. The instrument about to be described in the next chapters can be seen as a direct descendant of these two instruments.

The WHT has a 1° field of view, and the AAT a 2° field of view. Both telescopes are therefore wide field, and are highly appropriate for survey instruments to obtain hundreds of thousands of spectra across the whole sky. Both instruments are located at prime focus.

Both instruments feature a positioner robot that positions fibres across a flat steel plate with an accuracy of a few tens of microns. Orthogonal carriages move the robot in x , y and z . This system allows the robot to reach any position above the field plate. The robot has a gripper with two jaws. With this gripper, the robot can pick



up the metal housing around the fibre ends (called ‘buttons’) to reposition a fibre. The gripper jaws can rotate freely around a θ -axis to align them to each button. Figure 1.4 shows the button designs for both instruments. It is clear that the design has remained very similar.





Figure 1.4: Fibre button designs for Autofib-2 (left) and 2dF (right). Figures from Parry et al. (1994) and Lewis et al. (2002).

The fibres are placed on the steel field plate, but their orientation is along the plate. A tiny prism glued to the end face of the fibre couples the light into the fibre. A permanent magnet in the base of the button keeps this fibre-prism assembly in place on the field plate. Each button is a few mm long and has a handle protruding up with which the gripper picks it up.

The gripper features an integrated camera system to measure a fibre position accurately after it is placed. Each fibre is back-illuminated so the core is visible in the image.

The fibres are placed at the edge of the field plate when not in use. In Autofib-2, the fibres are protected by stainless steel tubes. When the fibres are at the edge of the plate, these steel tubes protrude, which increases the overall size of the instrument. This is not a problem for Autofib-2 as it is quite small. However 2dF is a lot larger, and using the same steel tubes to store the fibres would mean that too much of the mirror is obscured by the instrument. Therefore on 2dF the fibre is stored securely in a ‘retractor’, a device that holds the fibre under a slight tension with a spring-pulley






Figure 1.5: The tumbler system in 2dF: general arrangement of tumbler unit, field plates and fibre retractors. Not to scale. Figure from [Lewis et al. \(2002\)](#).

system. This ensures the fibre moves back from the plate in a straight line when it is moved to the edge of the plate and does not tangle with other fibres in the field. Another advantage is that fibres can now cross over each other. This opens up more possibilities for making configurations.

The biggest difference between these two instruments is their multiplex. Where Autofib-2 has 150 fibres to reconfigure after every observation, 2dF has almost three times as many fibres to move. Whilst the time it takes to do these moves is acceptable for Autofib-2, it takes so long for 2dF that it would severely decrease the instrument's efficiency. Therefore 2dF incorporates a completely new element: a tumbler. Instead of having only one field plate and one set of fibres, there are two field plates and two sets of fibres which are placed on opposite ends of a rotating structure.

As one field plate is observing a field (on the underside in figure 1.5), the other



plate is reconfigured for the next observation. The only steps that remain after an observation is finished, are to rotate the two field plates 180°, and exchange the slits so the new set of fibres is in the spectrograph.

Both robots are calibrated for offsets, rotations and non-perpendicularity between the x - and y -axis, and temperature variations. The setup in place allows for recalibration at certain moments. The gantries on Autofib-2 can only be recalibrated every so often as it involves removing the field plate from the instrument, but this is fine as it does not need constant recalibration because of its small size. 2dF's gantries have to be recalibrated before every observation because of the tumbler. It does not always reset to the same zero point in x and y . Therefore the procedure is changed to make recalibration measurements before every observation possible and streamlined. However, 2dF cannot be recalibrated for every linear parameter that Autofib-2 can.

Pick-and-place fibre positioners are very versatile. They can easily deal with sparse fields, but objects close together are no problem either. Because of the specific shape of the button, the fibres can be placed quite close to each other. The closest approach is when two prisms approaching from opposite ends of the plate, which can get even smaller as the positioners become more accurate.

1.6 This research

This work is concerned with the metrology, calibration, integration and testing of the WEAVE fibre positioner. The main goals are to minimise the overall positioner error and deliver an instrument that performs to the specifications. Chapter 2 introduces the WEAVE instrument and its fibre positioner, and looks at what areas of science will be served by this instrument. Chapter 3 concerns the building of the positioner, and the ways in which this process is monitored to ensure as small errors as possible. Chapter 4 is about software to reduce the residual errors in the robot calibration



process, and shows the software working with models and with preliminary real-world measurements. Chapter 5 presents details on several positioner components, and shows the performance of each. A large part of this chapter is concerned with the positioner robots, particularly the *Z*-gantry. Lastly, chapter 6 summarises this work.

1.7 Summary

It is clear that the field of astronomy and astrophysics is an exciting and ever-changing one. With more data available every day, more can be learned about the lives of stars, galaxies and the whole Universe. However, the Universe is enormous, and even if there are many instruments already available to do observations with, it is abundantly clear that there are uncountable more objects that are or could be interesting to analyse. Much effort is spent on ensuring that each new generation of spectrographs is more efficient than the previous, by providing more spectra of increasing resolutions. Pick-and-place fibre positioners are versatile in that their fibres can be placed anywhere in the telescope focus, and this makes them very useful for a wide range of use cases.




Chapter 2

Introduction to WEAVE

In the late 2000s a European working group was set up to plan the future of Europe’s 2–4 meter class telescopes around the world ([ETSRC, 2010](#); [ASTRONET WFS Working Group, 2011](#)). From this a need for large multiplex wide-field spectrographs was identified, to do ground-based follow up on large-area surveys with space-based instruments (Gaia) and ground-based instruments (LOFAR, and later Apertif). To provide a large enough coverage, these spectrographs would have to be mounted on 4-meter telescopes with a ≥ 1 degree field of view, with a multiplex of at least 500, but preferably more. In the northern hemisphere the William Herschel Telescope (WHT) on La Palma emerged as the best choice, because of its large primary mirror of 4.2 m and its availability for the next decade.

With these considerations the WEAVE project started in 2010 ([Balcells et al., 2010](#); [Dalton et al., 2012, 2014, 2016](#); [Dalton et al., 2018](#)). WEAVE (the WHT Enhanced Area Velocity Explorer) is a ground-based fibre-fed multi-object spectrograph, specifically built to follow up on stellar targets from the Gaia satellite for Galactic Archaeology, and to follow up on galactic targets from the LOFAR and Apertif cosmological and local radio surveys.

Over a span of at least five years, WEAVE will measure the spectra in the visible



range for at least 10 million astronomical objects. There are three ways in which the instrument can be used: MOS mode with configurable single fibres; mini-IFU mode where up to 20 mini-IFUs can be positioned in the focal plane; and LIFU mode, which is one large IFU in the centre of the field. First light is expected in early 2020.

2.1 WEAVE subsystems

WEAVE consists of three distinct subsystems: a new prime focus corrector which also includes an atmospheric dispersion compensator, a new spectrograph, and a fibre positioner. The planned format with its three distinct subsystems is shown in figure 2.1. The positioning system will be at prime focus and a high-resolution spectrograph will reside at one of the telescope's Nasmyth platforms.




Figure 2.1: Overview of WEAVE on the WHT. The fibre positioner will be at prime focus with the prime focus corrector, and a high-resolution spectrograph at one of the telescope's Nasmyth platforms. The fibre cables that are routed from the positioner to the spectrograph are also shown.



2.1.1 The prime focus corrector

A new prime focus corrector (PFC) provides a field of view of 2 degrees, across a flat focal surface 410 mm in diameter (Agócs et al., 2014; Abrams et al., 2014). It also contains an atmospheric dispersion compensator (ADC) and a field rotator. The ADC can only compensate for atmospheric dispersion effects¹, which are the same over the whole field. It cannot compensate for differential refraction², which varies over the field in magnitude and direction and is an issue for telescopes with large fields of view. For WEAVE this is corrected for as much as possible in the fibre placement.




Figure 2.2: The WEAVE top end assembly in more detail. The prime focus corrector (PFC), prime focus field rotator (PFR), focus translation unit (FTU, 4 in total which provide focus and tilt correction of the whole system) and the fibre positioner (POS) are labelled. The positioner is bolted onto the prime focus rotator, which itself is attached to the prime focus corrector.

The PFC consists of six lenses, the largest of which is 1.1 m in diameter. The focus translation system (Canchado et al., 2016) is designed such that the four units

¹The different wavelengths are refracted by different amounts, and therefore do not fall onto the focal plane at the same location.

²The change in air column density as a function of zenith distance, which implies a varying trapezoidal distortion across the whole field of view.



| Parameter | Designed value |
|------------------------|-----------------------------|
| Field of view | $2^\circ \varnothing$ |
| Focal surface | 410 mm \varnothing , flat |
| Telescope focal length | 11.55 m |
| F-number | f/2.75 |

Table 2.1: Prime focus corrector design.

together allow a focus adjustment of ± 3 mm in the optical axis, and a $\pm 0.015^\circ$ tilt correction in one axis.

The positioner will live behind the prime focus corrector. The mechanical interface is where these subsystems are bolted together, between each interface plate, and the optical interface is at the aforementioned focal surface: the positioner field plate will be situated such that the positioned fibres are in the focal plane.

The prime focus corrector is outside the scope of this work, and will not feature further in this thesis. Updates on its building status can be found in e.g. [Lhomé et al. \(2016\)](#); [Tomás et al. \(2018\)](#).

2.1.2 The spectrograph

The spectrograph has been designed according to the WEAVE science goals, with a dual-beam configuration (a ‘red’ and ‘blue’ arm). This allows a broad wavelength coverage whilst permitting the anti-reflection coatings of the camera lenses to be optimised for less than a full octave of free spectral range. There are two configurations, a low resolution mode ($R \approx 5000$) over the full wavelength range, and a high-resolution mode ($R \approx 20000$) over several narrower ranges ([Rogers et al., 2014](#)). A drawing of the design is shown in figure 2.3, more details are in table 2.2.

The dispersing elements are a set of volume-phase holographic (VPH) gratings, two for the low-resolution mode and three for the high-resolution mode ([Bianco et al., 2018](#)). Each camera consists of eight lenses. The cameras are identical, apart from the spacing between the lenses and the anti-reflection coatings, which are optimised



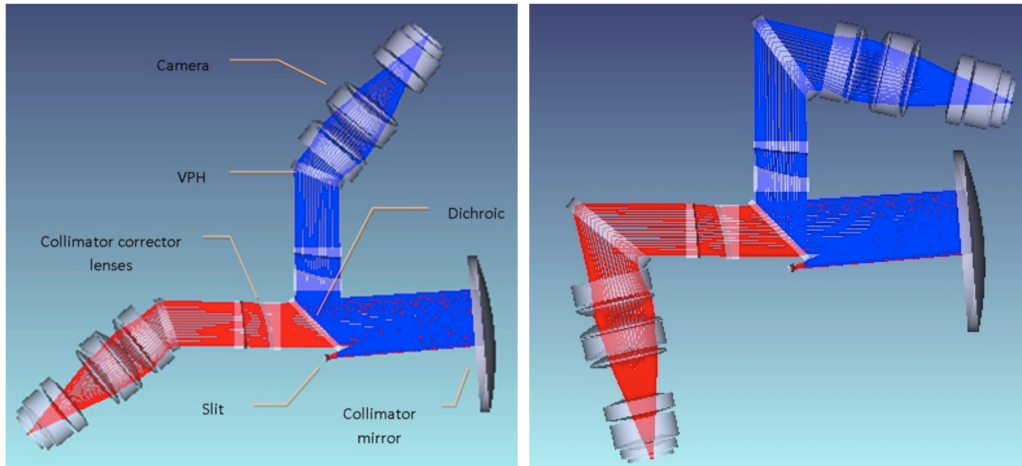



Figure 2.3: Spectrograph layout. *Top*: mechanical layout. The spectrograph is installed in the GHRIL room at the WHT. *Bottom*: optical layout. On the left the low-resolution mode, and on the right the high-resolution mode. The cameras move over the blue arcs in the top figure to switch modes.

| Parameter | Designed value | |
|-------------------------|------------------|-------------------------------------|
| | Low resolution | High resolution |
| Wavelength coverage (Å) | 3660–9590 | 4040–4650 4730–5450 5950–6850 |
| Resolution | 5750 (3000–7500) | 20000 (13000–25000) |

Table 2.2: Spectrograph design values. The values between parentheses are the resolution ranges, as resolution is dependent on wavelength.



for each waveband. The largest lens is 320 mm in diameter (Izazaga et al., 2018).

The interface between the spectrograph and the positioner is achieved by a set of four fibre cables, each 33 m in length, which route through the prime focus rotator, down the telescope structure and over the telescope elevation axis into the GHRIL enclosure.

2.1.3 The fibre positioner

The fibre positioner will take care of the accurate placement of the fibres on the field plate across the focal plane. There will be three working modes: a MOS mode with up to 960 movable fibres per field, an IFU mode with 20 movable small IFUs, and a single large IFU mode. The IFUs will be hex packed with 37 fibres in the small IFUs and 547 in the LIFU. A summary is given in table 2.3.

| Observing mode | Fibre size | Number of fibres in unit | Number of units | Unit field of view |
|----------------|------------|--------------------------|------------------|--------------------|
| MOS | 1.3" | 1 | 960 ^a | 1.3" |
| mIFU | 1.3" | 37 | 20 | 11" × 12" |
| LIFU | 2.6" | 547 | 1 | 1.3' × 1.5' |

Table 2.3: The three different observing modes for WEAVE, with the fibre configurations and sizes. ^aOne plate has 960 configurable fibres, the other plate has 940 fibres due to space constraints.

The fibre positioner is built around a central tumbler, which tumbles a configured plate into the focal plane, and tumbles the now observed plate back to the positioner robots to be reconfigured. The work in this thesis revolves³ around this positioner structure. A more in-depth introduction is given in section 2.3.

³Ha ha.



2.2 Science goals

Over a period of at least 5 years, several surveys will be carried out by WEAVE, producing more than 30 million spectra of about 10 million objects. The next sections contain very short summaries of the corresponding sections in [Trager and the Science Teams \(2015\)](#)⁴ and [Jin et al. \(in preparation\)](#). There are three main scientific motivations for WEAVE. The first one is to follow up on measurements made by the Gaia satellite for galactic archaeology. The second and third are the follow-up of galaxy measurements made by the LOFAR and Apertif radio telescope facilities to better understand galaxy evolution. Most of the available time will be allocated to these surveys, but more surveys have been devised as the plans for WEAVE became more concrete. All the surveys currently planned in the WEAVE Survey are listed in the next sections.

2.2.1 Galactic archaeology

The space-based Gaia instrument, launched in December 2013 ([Gaia Collaboration, 2016](#)), measures the positions and proper motions for a billion stars, with radial velocities for the brightest stars ($G \lesssim 16.5$). Gaia's goal is to acquire statistics on the orbits, ages, and chemical compositions of the stars in the Milky Way, enabling a complete transformation of the field with respect to the history of its formation and subsequent evolution. However, Gaia is limited to measuring radial velocities of the brightest stars only, and can only measure abundances of a few elements for a subset of all the stars, because of limitations in resolution to its on-board spectrograph.

To fulfil Gaia's goal, the measurements have to be complemented with ground-based spectroscopic follow-up surveys over the whole sky. WEAVE will provide these measurements for the northern hemisphere. It will provide radial velocities for stars fainter than about $G > 16$, with an expected precision of $\pm 1 - 2 \text{ km s}^{-1}$, together with

⁴Incidentally a great band name.



effective temperature, surface gravity and metallicity (expected precision ± 0.2 dex). Furthermore it will provide detailed elemental abundance measurements for the majority of stars with magnitudes $12 < G < 16$.

WEAVE's low-resolution $R = 5000$ mode will be used for the radial velocity measurements, and the high-resolution $R = 20000$ mode for the chemical composition measurements. Because of its location in the northern hemisphere, WEAVE will concentrate its efforts on the formation and evolution of the outer halo and disc. For the halo, observations of at least 2.5×10^6 stars (and 6×10^5 at high resolution) are required over at least 8500 deg^2 of the sky, and the disc survey requires observations of at least 1.5×10^6 stars over 2000 deg^2 of the sky near the Galactic Plane.

Because this is such a large task, the WEAVE Galactic Archaeology survey is divided into four sub-surveys. These are:

- The high-latitude low-resolution survey, a wide-area survey of 8500 deg^2 at high Galactic latitude of stars with magnitudes $G > 16$. This is to complement Gaia's precise astrometry and obtain full phase-space information for these stars. This information will constrain the halo's total mass and lumpiness to much larger distances than would be possible if only spectra of Gaia's onboard spectrograph were available.
- The disc-dynamics low-resolution survey, in which about 500 lines of sight through the Galactic disc are surveyed at very low Galactic latitudes, to probe fundamental Galactic disc dynamics processes.
- The high-resolution chemo-dynamical survey, a survey of 5600 deg^2 at intermediate and high Galactic latitudes. In this survey, stars with magnitudes $12 < G < 16$ are observed to complement Gaia's phase-space information with full chemical information.
- The open clusters survey, in which about 100 open clusters are observed, mostly



in high resolution. This survey will look at the chemical evolution of the Milky Way using these open clusters as probes, and provide insights on star formation and early stellar evolution.

2.2.2 SCIP


The natural focus of galactic archaeology is on the more common older stars. The Stellar, Circumstellar, and Interstellar Physics (SCIP) survey of the Galactic Plane on the other hand focuses on the short-lived phases of stellar evolution: young and old extremes, higher-mass stars, interacting binaries, and so on. Also important is the interstellar medium, both ionised and seen in absorption against background starlight, and its interplay with the stars.

Spectroscopy at $R \sim 5000$ over the entire optical wavelength range provides many types of information that cannot be derived from photometry alone. These include systemic/orbital radial velocities, kinematic evidence of binarity, interstellar absorption features, mass transfer/loss signatures, and signs of magnetic activity.

SCIP's low-resolution survey will cover almost 1100 deg^2 , and will all be in WEAVE's $R \sim 5000$ mode. Over half of the planned targets are luminous OBA stars, which creates the opportunity to map the Galactic disc structure out to heliocentric distances of about 10 kpc. The high-resolution survey also focuses mostly on OBA stars, but specifically in the Cygnus and Galactic Anticentre regions. In Cygnus, this is to build a full picture of kinematics and elemental abundances across the region, whereas in the Anticentre fields the focus is on the stars themselves and on how the effect of Galactic disc kinematics is imprinted on the population.

2.2.3 White dwarfs

Most stars (and all stars with masses $\lesssim 8 - 10 M_{\odot}$) will become white dwarfs in the end. Because many of these stars, those with masses $> 1.5 M_{\odot}$, have short main-



sequence lifetimes, this means that many stars that formed throughout the history of the Galaxy have already evolved into white dwarfs. Therefore, to constrain stellar evolution theories, a large sample of white dwarfs with accurate physical parameters is essential.


The WEAVE white dwarf survey will obtain low-resolution spectroscopy of about 70,000 white dwarfs with $G \lesssim 20$, which roughly triples the number of white dwarfs with high-quality spectroscopy in the northern hemisphere. This will also establish the first large and homogeneous spectroscopic white-dwarf sample that is not subject to complex selection effects.

Each WEAVE pointing will dedicate 5–15 fibres to white dwarfs, which can also be used for flux calibration purposes. This will provide effective temperature, surface gravity, atmospheric composition, magnetic field strength, and multiplicity. The mass, cooling age and progenitor mass can also be derived from these spectra.

2.2.4 WEAVE-Apertif

Recent local IFU surveys of galaxies suffer from an optical selection bias that makes it difficult to address the effects of environment in the evolution of galaxies. One possible approach to study these effects is to characterise the morphologies of galaxies seen in the atomic gas. To this goal a wide-area survey of resolved neutral atomic hydrogen is needed.

Apertif ([Verheijen et al., 2008](#)) is a new focal-plane array system on the Westerbork Synthesis Radio Telescope and an SKA pathfinder. Apertif will be used for a medium deep survey (out to $z \sim 0.2$), to obtain resolved kinematics from the neutral atomic hydrogen line for about 100,000 galaxies. For these galaxies, WEAVE will be able to provide the kinematics in the visible band, and also deliver spatially resolved stellar age distributions and star-formation rates. Using the LIFU in low-resolution mode an Apertif sample of about 400 galaxies will be observed in the WEAVE-Apertif survey.




A subsample of galaxies will be selected from the low-resolution dataset for follow-up spectroscopy with the high-resolution mode of the LIFU. These observations will provide information on the impact of the neutral atomic hydrogen accretion history on the chemo-dynamical properties of galaxies. This will be done by comparing the stellar kinematics and stellar properties of a sample of galaxies with perturbed neutral atomic hydrogen morphologies, to a sample of galaxies with no signs of those.

2.2.5 WEAVE-LOFAR

LOFAR ([van Haarlem et al., 2013](#)) is another SKA pathfinder at low radio frequencies, and is currently carrying out a set of three nested surveys, described in detail in [Röttgering et al. \(2011\)](#). An important driver for LOFAR is to advance the understanding of the formation and evolution of galaxies, galaxy clusters, and active galactic nuclei (AGN). LOFAR's widest survey covers the entire northern hemisphere, and its narrow, deep survey is deeper than any current radio imaging. This will yield samples of 10^7 sources, primarily a mix of distant star-forming galaxies and radio-quiet AGN. But LOFAR only provides continuum fluxes and so the redshifts for these objects cannot be measured.

WEAVE will perform the spectroscopic observations to provide not only these redshifts, but also source classifications. Only with spectra the distinction can be made between sources that are dominated by star formation (star-forming galaxies) and those that are dominated by accretion onto an AGN. The accretion mechanism of the AGN can also be determined; whether they are dominated by efficient accretion of cold gas, or inefficient accretion of hot gas. WEAVE will also provide information on velocity dispersions and metallicities, and therefore virial black hole mass estimates can be derived.

Because LOFAR sources are selected at radio frequencies, which are not affected by dust obscuration, WEAVE-LOFAR will be far less biased against obscured sources



than samples which are selected at optical wavelengths. This provides a much more representative view of populations of galaxies. The survey roughly mirrors the LOFAR strategy of three tiers, with a wide, mid and deep tier.


The WEAVE-LOFAR survey will obtain more than a million spectra of 150 MHz-selected sources in the MOS mode, and any of these can serve as follow-up targets for WEAVE's IFU modes. [Smith et al. \(2016\)](#) provides more information on the design of this survey.

2.2.6 Galaxy clusters

Galaxy clusters are the largest and most massive gravitationally bound structures known in the Universe. They are also excellent cosmological probes, and can provide constraints on cosmological parameters. According to current cosmological theories, galaxy clusters have grown hierarchically by the merging of smaller virialised haloes. These theories are very successful in describing most of structure formation, but do have a few problems. One of the largest problems is called the ‘missing satellite’ problem: the models predict more dwarf galaxies than we observe. Another issue is the star-formation efficiency in these dwarfs, which is highly debated.

Cluster research with WEAVE will focus on three main areas: the evolution of dwarf galaxies in nearby clusters, the effect of environment on galaxies and groups falling into clusters for the first time, and to constrain standard cosmological parameters.

In a single pointing with its (low-resolution) MOS mode, WEAVE will be able to obtain galaxy spectra from the cores to the suburbs of nearby galaxy clusters. In total, WEAVE will survey more than 200,000 cluster galaxies with its MOS mode, 1000 galaxies divided over 47 clusters with the small IFUs, and up to 100 cluster cores with the LIFU.



2.2.7 StePS

The aim for StePS (Stellar Populations at intermediate redshifts Survey) is to study galaxy evolution out to $z = 0.7$ in detail. It will look at galaxy evolution in the last 7 Gyr, the physical mechanisms driving the star-formation history of galaxies and quenching, and the interplay between environment and intrinsic galaxy properties.

StePS will carry out a MOS survey of about 40,000 galaxies at $0.3 < z < 0.7$ over four areas on the sky with a total area of about 25 deg^2 . By observing for 7 hours with WEAVE's low-resolution mode, the spectra will have high enough signal-to-noise per resolution element to undertake a detailed analysis of these faint galaxies.

StePS will precisely derive galaxy stellar ages, star-formation timescales, stellar and gas metallicities and dust attenuation, together with gas kinematics and stellar velocity dispersions. The large sample to be observed will provide a coherent picture of galaxy evolution over the last 7 Gyr.

2.2.8 WEAVE-QSO

The WEAVE-QSO survey (Pieri et al., 2016), which will observe around 400,000 high-redshift quasars, is designed to optimise quasar absorption science through the measurement of $\text{Ly}\alpha$ absorption and other intergalactic medium (IGM) absorbers. All quasars will have $z > 2.1$, which ensures that the $\text{Ly}\alpha$ transition is redshifted fully into the optical window. This will approximately double the number of observed $\text{Ly}\alpha$ -forest quasars.

The WEAVE-QSO fields are at high Galactic latitudes, so the fields will be shared with the Galactic Archaeology survey in both low- and high-resolution modes. Fields will also be shared with the WEAVE-LOFAR survey in the HETDEX region.



2.2.9 WEAVE Survey footprint

The eight surveys all have very different goals, with the science goals spanning from the nearby stars to some of the furthest astronomical objects known so far. Nonetheless, all these surveys use the same instrument on the same telescope. There are several regions on the sky in which some of the survey teams actively work together to obtain the most science out of every telescope pointing. An early figure of the survey areas (figure 2.4) show exactly this: wherever there are overlapping survey regions, the teams are discussing where and how to share fibres, and which configuration of the spectrograph to use. This has streamlined the process of deciding on the WEAVE pointings, and ensures that each WEAVE field has as many interesting targets as possible, improving the efficiency of the instrument.

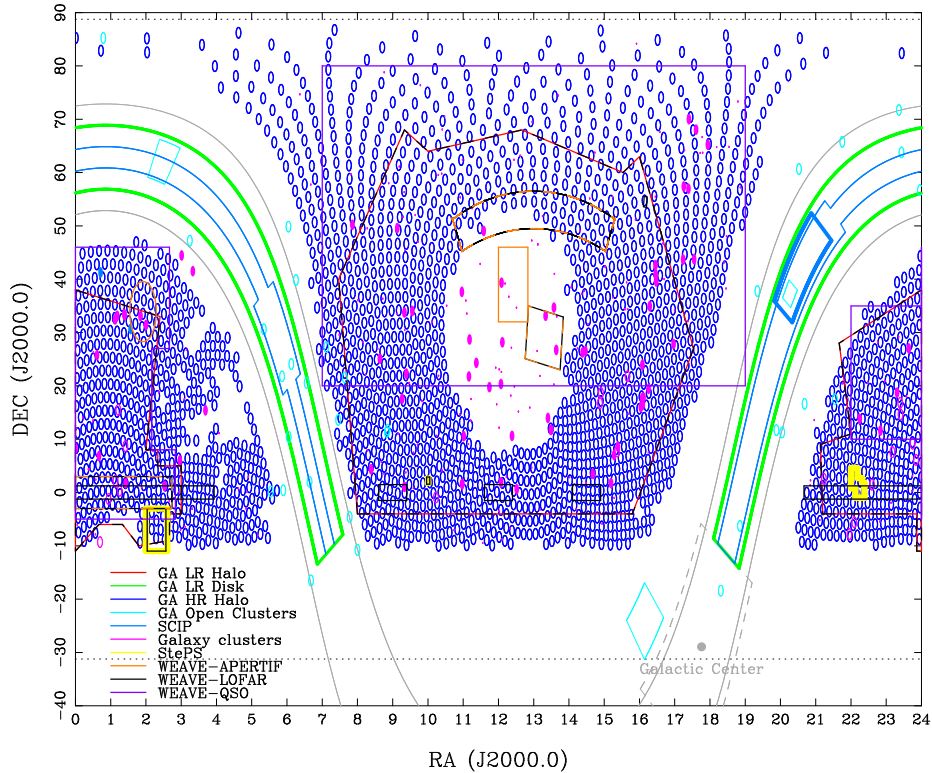



Figure 2.4: An overview of the different survey areas, correct in 2016. Since then, a lot of details have changed, but it still gives a reasonable visual impression. Each ‘circle’ is a single WEAVE pointing, each block is a region that will be fully observed, with a tiling strategy decided by that particular survey team.



2.3 The WEAVE fibre positioner in detail

The fibre positioner design is based on the ‘pick-and-place’ technology which was also used in Autofib-2 on the WHT (Lewis et al., 1997) and in 2dF on the AAT (Lewis et al., 2002). Robotic fibre positioners are able to place each fibre in its correct location on the field plate. The robot positions the fibres sequentially, which adds up to an unacceptable overhead when there are several hundred fibres that have to be reconfigured before a next observation can take place.


2.3.1 Positioner components

To reduce this positioning overhead the WEAVE positioner design, with figure 2.5 showing the status at the preliminary design stage (Lewis et al., 2014), includes a tumbler (figure 2.6), like the 2dF design. There are two field plates, one on either end of the tumbler, and the fibre positioner is located on the opposite end from the observing side. There are also two sets of fibres, one set for each field plate. This means that whilst one plate is in observation mode, the other can be reconfigured to the next field in the same time⁵. After a full observation, the tumbler only has to rotate the two plates 180° and the next observation can begin. Because there are more than twice as many fibres to be reconfigured in WEAVE compared to 2dF, the WEAVE design included two pick-and-place robots from very early on.

One of the key philosophies for this instrument is that the positioner is built with commercially available components wherever possible, to reduce cost and integration timescales. The tumbler drive and the positioner robots and gantries are all commercially available.

The two robots on (x, y, z, θ) -gantries are equipped with grippers to position the fibres anywhere on the field plate. Each of the fibres terminates in a structure called

⁵Timing checks were done by Gilbert (2016): the maximum average move time of 3.3 s/fibre can be achieved if the robot movements in the different axes overlap in a clever way.






Figure 2.5: General layout of the fibre positioner system. There are two field plates on either end of a tumbler, so one plate can be configured while the other is in use. Configuration is done by two positioner robots working together. Figure adapted from [Lewis et al. \(2014\)](#).





Figure 2.6: The key components of the tumbler. The motor is on one side, and rotates the whole structure around the red axis. Two retractors are shown; in reality each field plate will be surrounded by a ring of 168 retractors.



a button (figure 2.7), which has a vertical handle for the robots to grip the button and move the fibre around. A small magnet on the bottom ensures the button sticks to the field plate, and a prism glued to the end guides the light into the fibre. The prisms on the MOS fibres have a concave surface to transform the incoming beam from $f/2.75$ to $f/3.2$, to reduce fibre losses and simplify the spectrograph collimator design. The mIFU units have a separate lens which performs the same role. Furthermore, the fibres are secured in the button at a small tilt of 2° , to reduce the effect of slight non-telecentricity from the corrector.




Figure 2.7: Schematic of MOS and mIFU fibre magnetic buttons. In each drawing part 1 is the button housing made from magnetic stainless steel, 4 is a permanent magnet, 5 is a prism, 6 is the fibre ferrule. Part number 8 is the separate field lens for the mIFU button.

There is a separate button design for the guide fibre bundles (eight per field plate), which hold 3” coherent image bundles containing many thousands of fibre cores. These have larger (2 mm) prisms with no focal ratio conversion lens. The overall design is still very similar to the MOS buttons, to ensure that the robots can handle the different fibres in the same way.

The MOS fibres are kept in 168 retractors around the perimeter of each field plate. Each retractor houses up to six fibres and provides a three-tiered park porch to park the fibre button when not in use. A constant-force spring-pulley system keeps the fibre under tension, for safe storage without twists or kinks. The fibres are anchored



in the retractor by pressing the outer protective sleeve into a closely fitting square groove. A generous length of fibre is kept in a spare-length box, for the inevitable repairs that will have to be done on the fibre ends during the instrument lifetime of operations. Figure 2.8 shows a fully-built retractor, and some retractors in close up. The guide fibres are housed in the same retractors as the MOS fibres, but the mIFU bundles need a different design, as the fibre bundle is so much thicker. The MOS fibre retractors are discussed in more detail in section 5.1. The mIFU retractors were not part of the work for this thesis.




Figure 2.8: *On the left:* a fully-built MOS retractor. Spring-pulley systems keep the fibres under tension. *On the right:* a close up of the park porches of several populated retractors next to each other around the field plate.

2.3.2 Positioner error budget

The positioner error budget follows from the accuracy with which the fibres have to be positioned. A mismatch between the optimal and the actual fibre position results in less light entering the fibre. This positioning accuracy has an effect on measurements in two distinct ways. For a given magnitude limit, the signal-to-noise ratio (SNR) can be increased by more accurate positioning; and for a given observation time, the observed magnitude can be increased.



The exposure time is limited by differential refraction which depends on air mass⁶, and therefore the elevation. This manifests itself in two ways. First the objects at lower elevation in the field are refracted more than at higher elevation. This could be calibrated for if it were the only effect, but the second effect makes that impossible: as the telescope tracks the field during an exposure, the elevation (and thus air mass) changes continuously, and the objects at the lower edge of the field move relative to the ones higher up. After about 60 minutes the light loss in the edge fibres is too large to continue observing. At the edge of the field, it corresponds to a 0.5'' worst-case shift from the mid-point of the exposure. As the fibres cannot be moved during the exposure, the nominal exposure time for WEAVE is set at 60 minutes.

Early on in the project, the positioning accuracy requirement was set at 8.0 μm RSS. This was based on a total offset of 0.2'' divided equally between the prime focus corrector and fibre positioner subsystems. The description below follows the calculations and arguments made by the WEAVE team in 2014 (Dalton, 2014). Because WEAVE was proposed as a follow-up instrument for Gaia targets, it makes sense to use Gaia's magnitude limit as a starting point for the WEAVE calculations.

First the fibre aperture has to be chosen. It is ultimately a trade-off between the overall system cost, the complexity of the spectrograph design and the end-to-end system performance. Analysis showed that a fibre core diameter of 1'' gives a too-low SNR per resolution element, and a 1.6'' core would make the spectrograph too large and expensive. The design settled on a fibre core diameter of 1.3''.

Then taking $V = 20$ (Gaia's nominal magnitude limit), the expected performance of the prime focus corrector, the WHT's seeing conditions of 0.8''⁷, the fibre core diameter of 1.3'', and some anticipated values for the optical system throughput and detector readout noise, the SNR per resolution element can be calculated. This is

⁶The amount of air between the observed object and the observer. The air mass is at its minimum at zenith and increases towards the horizon.

⁷This is the reported 75th percentile from the WHT's site statistics. They are accessible at <http://www.ing.iac.es/astrometry/observing/conditions/>.






Figure made by Gavin Dalton, with additions from James Gilbert. It also appeared as figure 1.6 in [Gilbert \(2016\)](#). Figure and caption reused with gratitude.

Figure 2.9: Simulated signal-to-noise ratio per 3-pixel spectral resolution element for the WEAVE low-resolution mode and a nominal faint magnitude limit of $V = 20$, a nominal exposure time of 60 min, an airmass of 1.2, a detector readout noise of $2.5 e^-$, and an on-sky fibre core diameter of $1.3''$. Each sub-figure represents a different system throughput within the anticipated range; each line represents the post-PFC PSF width in different seeing conditions. The solid lines are for best-case (Gaussian) PSF shapes, and the dashed lines for worst-case (Moffat, $\beta = 3$) ones. The bold (red) lines correspond to a reported 75th percentile seeing of $0.8''$ for the WHT. The dashed vertical line shows the combined maximum offset requirement of $0.2''$ for the PFC and positioner subsystems combined.



shown in figure 2.9.

In this figure we see the SNR tapering off with higher fibre position offsets, as expected. Each panel shows the SNR for different anticipated system throughputs, and each line represents different seeing conditions. The red lines correspond to a 0.8'' seeing convolved with the expected performance of the PFC, and is adopted as the standard WEAVE value for evaluating the instrument specifications.

The vertical dashed line in each plot represents the 0.2'' accuracy limit for the PFC and positioner subsystems combined. With the WHT's plate scale⁸ of 56 $\mu\text{m}/\text{arcsec}$ this translates to 11.2 μm in distance on the focal plane. This cannot be taken up by the positioner wholly, as the PFC has its own errors to take into account, such as residual errors in the focal plane distortion map, changes in focal plane calibration with temperature, field rotator tracking issues, etc.

Splitting this total error budget of 11.2 μm evenly (RSS) over the positioner and PFC subsystems leaves 8.0 μm for each subsystem.

The 8.0 μm positioner budget is divided into several sub-budgets, shown in table 2.4. It shows the situation prior to the work in this thesis, with references to relevant sections where progress is made to update some of these values. As time and research has progressed, it is clear that the positioner is currently over budget. Looking again at figure 2.9, an increase of a few μm in offset does not make a large difference in SNR; it drops by a few percent only. The best-case scenario would of course still be to stay under 8.0 μm offset, but a slightly larger error will not pose a significant problem.

2.4 Summary


This chapter provides an introduction to the WEAVE instrument and its subsystems: the prime focus corrector, the spectrograph, and the fibre positioner. The design of

⁸The plate scale of a telescope gives the relationship between angular distances on the sky and real-world linear distances on the focal plane.



| Description | RMS value (μm) | Comments | Section |
|--|-----------------------------|--|--------------------------------|
| Fibre position measurement error | 1.0 | Includes fibre centre-finding error, camera noise, θ -axis reference error, pixel scale error, axis position uncertainty, and camera alignment error. | Gilbert (2016) |
| Gripper release repeatability | 3.6 | RSS of 2 μm placement and 3 μm release. | - |
| Positioner robot repeatability | 7.1 | Encoder errors, 5 μm per axis. | 5.4 |
| Positioner robot calibration | 4.0 | Residual after correction for non-squareness, thermal expansion, straightness of gantry axes. | 4.1–4.6 |
| Gantry flexure uncertainty | 2.0 | Flexure component which varies with robot position, value from FEA analysis. | - |
| Total misalignment contribution from POS | 9.2 | RSS of individual contributions | |

Table 2.4: Bottom-up error budget for the POS system. Initial values from [Middleton \(2018\)](#), with additions from [Gilbert \(2016, section 3.9, p. 86–88\)](#). The section numbers indicate where in this thesis the analysis is described to update some of these numbers from simulations to actual values.



the WEAVE fibre positioner follows the 2dF design very closely, with two field plates and two sets of fibres in a tumbler configuration, fibre retractors with spring-pulley systems and buttons with permanent magnets.

The science planned with WEAVE spans from stars in the Milky Way to some of the farthest known objects: quasars. The eight survey teams all have very clear science goals and want to utilise the WEAVE spectrograph in all the different modes that are available, in both low resolution and high resolution. The science cases require certain resolutions and signal-to-noise ratios (SNR), and it is clear that it is essential to build an instrument that will meet these science requirements.

The conversion from science requirements to technical requirements is shown by presenting the trade-off arguments between signal-to-noise ratio, fibre size, complexity of the spectrograph, and overall cost. Lastly, the positioning accuracy requirement is discussed, and it is shown that a small non-compliance with this requirement is not disastrous for the signal-to-noise ratio.



Chapter 3

Positioner metrology

Over the course of the past three years, the positioner has been fully assembled twice. The first time, called the initial build from here on, was to check we could make the positioner work. Did we have all the components, were they the correct dimensions, were the bolt holes in the correct place, etc. With a few modifications to some of the components, this initial build resulted in a fully functional positioner, albeit one where not all the counterweights were attached to the different gantries.

We did not take much care in ensuring that the build was within tolerance, preferring rather to focus on progressing the build. It was interesting to see how close we were able to get to the specifications by just building the positioner such that all the components worked. Another reason for not building to within the specifications the first time, was that the whole positioner had to be dismantled again for anodising (to prevent degradation of the surfaces and diffuse scattering). Any precision calibration would be lost.

This meant that for the second build, or the ‘final build’, we had a building plan, but not a fully tested metrology sequence. The measurements showed that there definitely is a need for a precision calibration, because compounded small offsets meant that some crucial components were out by up to a millimeter.



3.1 Methodology

Two elements essential to get right in the positioner are the accuracy in the positioner robots (x , y , and θ in local coordinates) and the overall focus (z in the same coordinate system). Whereas the robots' positioning capabilities can be calibrated after building, the overall focus is locked in during the build. There is no mechanism to piston the positioner up or down with respect to the prime focus corrector, and no mechanism to change the tilt between the two subsystems. Once the positioner is built and bolted onto the corrector, no changes can be made anymore.

The prime focus system is designed in such a way that the resulting focal plane is flat and a distance 271.6 mm away from the interface to the positioner. With a button height of 2.4 mm, this means that each field plate has to be 274 mm from the interface, with a tolerance of ± 0.1 mm in height, and 12" in overall tilt, when the plate is in the observing position. Figures 3.1 and 3.2 show an overview of the interface. The former emphasises the physical structure, whereas the latter focuses on all the different considerations that went into designing the interface.

The starting point is the interface between the positioner and the prime focus system. This interface is flat, and is where the positioner's interface plate is attached to the interface plate of the prime focus system. With the available measurement tools in the lab it is impossible to measure the location of this plane when the positioner is fully built, so it has to be done before the tumbler is put in. This means that I cannot take a direct distance measurement from the interface to the field plate, and I have to break up the metrology into separate measurements.

The measurement tools in the lab include a FARO Edge¹ with an accuracy of about 40 μm over 1.2 m, a digital dial gauge with a precision of 1 μm , and a micrometer, also with a precision of 1 μm . With these available tools, the metrology plan took

¹A measurement arm with a ceramic touch-tip (6 mm diameter) that measures points on surfaces, and a very temperamental piece of hardware. I spent several months ensuring that the lab was set up in such a way that reliable repeat measurements with the FARO Edge were possible.






Figure 3.1: Schematic of the interface between POS on the top and the prime focus corrector on the bottom. The distance from the bottom of the interface plate to the field plate (the thick line in red) should be 274 mm. The 2.4 mm height of the buttons above the field plate (not shown) make up the difference to the focal plane distance of 271.6 mm. The tolerance on this field plate is ± 0.1 mm and $12''$. The design on the side of the corrector has changed since this drawing, but the interface itself has not.




Figure 3.2: All the requirements needing to be considered for the interface between POS and the prime focus corrector. Figure from [Middleton and Dee \(2018\)](#).



the form of the following list. The letters in parentheses correspond to the letters in figure 3.3.

1. Measure the thickness of the interface plate and the location of this plate in free space. Use the dial gauge to measure the thickness through the bolt holes, and the FARO Edge to characterise the plate itself (C).
2. Set up the FPI gantry parallel to the interface plate, using the FARO Edge, and measure its distance to the interface plate (D and E). Use this to calculate B .
3. Set up the field plate in the observing position parallel to the FPI gantry, using the dial gauge attached to the gantry (A).
4. Tumble the field plates, set up other field plate in exactly the same way as the first field plate (also A).
5. Set up the positioner robots parallel to the field plate on the top for smooth fibre positioning (not described in this thesis).


In the next sections I will discuss each part of this plan, the parts of the positioner involved, and the measurements themselves.

3.2 Interface plate and handling trolley

3.2.1 Hardware status

The positioner is built on the ‘interface plate’: an aluminium plate that connects the positioner to the rotator and prime focus corrector. For ease of transport, the whole positioner is built on a blue handling trolley.

The tolerance on the flatness of the handling trolley was not specified, because it is only a handling trolley, and therefore it was assembled and painted blue quite quickly.






Figure 3.3: The setup for the distance measurements from the interface to the field plate, and the tilt between them. A , B , and C are designed values. A and C can be measured directly, D and E are measured to reconstruct B . The structure on the left shows the FARO Edge tip in its calibration cone and the distances measured (D and E) in that setup, on the right is the dial gauge which measures A .




Photo courtesy of Ian Lewis.

Figure 3.4: Blue handling trolley with the old interface plate on it.



In several places paint drippings can be seen, and the edge of the plate around the hole in the middle is much higher than the rest of the trolley, in several places more than 0.6 mm. There are also quite a few bumps.

The design for the interface plate changed due to developments in the design of the instrument rotator. This happened after the initial build of the positioner was complete, and so the original plate was used in that configuration. After dismantling the positioner for anodising, it turned out that this plate had bent whilst on the handling trolley with the positioner's weight resting on it. Because the redesigned interface plate is thinned to a light-weighted structure in many places, it can easily bend, and there would be a danger of distorting the physical interface to the prime focus corrector if this plate were placed directly onto the handling trolley. To prevent this from happening, we decided to leave the original interface plate on the trolley and start building from there.

Because of the light-weighting, one corner of the interface plate has already visibly bent upwards. It can therefore be assumed that this plate is not flat either in its current state.



Figure 3.5: Side view of trolley in blue, old interface plate in grey, and shims in orange. Note that the plate is flat but not necessarily horizontal. This does not matter for the measurements as everything is measured with respect to the interface plate. This situation assumes that the interface plate bends back under the positioner's weight.

The original interface plate is shimmed in 12 places, to raise the plate to about 0.5 mm above the trolley's edge. It is (loosely) bolted to the handling trolley in those places, and the new interface plate is pressed to the old plate with 16 clamps around the inside rim. The rest of the positioner is then built on top of this. With this setup,



we expect that the original interface plate relaxes to (mostly) its original, undeformed shape under the weight of the positioner in the time it takes to build and ship the structure. This procedure means that the new interface plate may also be able to bend back to its original flat shape, which makes putting together the positioner and the rotator that much easier.

The positioner on the handling frame will leave the lab through double doors which open directly to outside. The structure is designed to fit through these doors comfortably, both in height and width. This was tested with a wooden stand-in slightly larger than the design (figure 3.6), and it fit through the doors easily. The total height of the positioner was reduced by 107 mm after the Final Design Review by shrinking the outer covers. Therefore building on two interface plates instead of one, which adds 12 mm to the height, should not pose a problem for transport.

3.2.2 Plate thickness measurements

It is clear from figure 3.4 that we cannot access the actual location of the interface, i.e. the underside of the interface plate, as it rests on the trolley. To get an accurate measurement for the location of this interface, I therefore measure the location of the topside, and the thickness of the plate itself.

Both interface plates are machined with a precision that is lower than the tolerance on the position of the field plates. The old interface plate is also warped from the initial build, and test measurements with the FARO Edge on this plate showed the plate's flatness varies locally by up to 0.3 mm. It is therefore crucial to measure the new interface plate in such a way that these local differences do not impact the overall outcome.

The essential points on the interface plate are in the locations of the posts which form the main structure of the positioner. The rest of the interface plate is important for structural stability, but the actual thickness there is less critical. Each upright is






Photo courtesy of Ian Lewis.

Figure 3.6: April 2015: a slightly oversized wooden stand-in for the positioner, and PI Gavin Dalton holding a small invar plate with fibre buttons where one of the field plates will be. This structure fits easily through the doors off to the side on the right of the image.




held in place by two bolts, and the holes for these bolts give the opportunity to do a thickness measurement.

For the thickness measurement we put the new interface plate on top of the old one, slightly out of place such that the bolt holes did not line up. Two large weights on either side of the measurement site kept the two plates as close together as possible. I then measured the thickness by lowering the dial gauge into the hole until it touched the plate below, and noting the difference between the readings in the hole and on the topside of the new interface plate. Figure 3.7 shows an example reading for the topside and through the plate. Each hole had 12 readings on the topside, in a ring around the hole. Of those I discarded the highest and lowest reading to remain with 10, and there were 3 or 4 readings inside the hole in different places. The averages of these readings and the resulting thicknesses are given in table 3.1.



Figure 3.7: Two example readings of the interface plate. On the left a reading on top of the plate, on the right a reading through the hole, with the gauge touching the plate below. Not shown are the weights on either side of the measurement site.

Each pair of holes is connected through an upright, so the differences average out. These averages are also given in table 3.1. It is clear that the thickness is close to the specified 12 mm, but not the same in every corner. The differences are very small



| Location | Topside (mm) | Through hole (mm) | Thickness (mm) | Averaged thickness (mm) |
|----------|-----------------|----------------------|-------------------|-------------------------------|
| 1a | 14.349 | 2.328 | 12.021 | 12.017 |
| 1b | 14.346 | 2.333 | 12.013 | |
| 2c | 14.352 | 2.334 | 12.018 | 12.020 |
| 2d | 14.352 | 2.330 | 12.022 | |
| 3e | 14.354 | 2.349 | 12.005 | 12.002 |
| 3f | 14.351 | 2.352 | 11.999 | |
| 4g | 14.353 | 2.352 | 12.001 | 11.994 |
| 4h | 14.344 | 2.357 | 11.987 | |

Table 3.1: Measurements of the interface plate’s thickness in 8 key locations. The locations are drawn in figure 3.8. The last column gives the averaged thickness for each corner.

though. It is very likely that these differences are averaged out by compounded small offsets in other parts of the positioner. There is a very slight tilt from corners 1 and 2 to corners 3 and 4, but this tilt is 25 μm over 1.2 m, which translates to a tilt angle of about 3". This is in the direction of the tumbler rotation, and is even smaller than this mechanism can correct for.

3.2.3 Plate location measurements

The underside of the interface plate defines the interface. The assumption is that any bends and other non-flat regions in the interface plate will straighten out once it is bolted onto the instrument rotator flange. So how then do we factor in this change, when at the moment of measuring the plate itself is not necessarily flat? One way to approach this problem, is by fitting a plane through the measurements.

The measurements are taken with the FARO Edge, with seven readings around each upright. I have defined an external coordinate system with $(0,0,0)$ at the top of the lower left post in figure 3.9, with x and y pointing in the same directions as in the figure. Fitting a plane through these readings and then doing a coordinate






Figure 3.8: Locations of the bolt holes for the uprights in the interface plate, and the numbered corners. Most of the bolt holes are obscured by the notes, one set is visible in the oval in corner 3. Any warping in the image is due to the photographer, not because of the plate itself.

transformation² to make the fit coincide with the x, y -plane, gives the left plot in figure 3.10.

It is clear that the overall shape of the interface plate is more like a saddle than flat, and around the uprights it is not flat either. The overall saddle shape can be decreased by placing shims underneath the plate. Because of the assumption that the plate will become flat after bolting onto the prime focus corrector, these external and temporary shims are a valid way of making the plate flatter.

After multiple rounds of shimming and checking the interface plate, most of the saddle is taken out (right plot in figure 3.10). The spread around the uprights still remains, but this cannot be adjusted for with shims. This process was dependent on the availability of shim stock; each actual shim was within 10 μm of the calculated shim. The total shim thickness in each corner, under each upright, is given in table 3.2.

²I leaned quite heavily on Appendix 9.2 in [Cole \(2015\)](#) to do these transformations.






Figure 3.9: FARO measuring setup, looking down. The red dot on the lower left post is the location of $(0, 0, 0)$. *On the left:* the FARO Edge is fastened to an optical table that is secured to the trolley. From there the probe can reach each corner of the structure, and every possible location in between. *On the right:* there are seven measurement locations around each post.




Figure 3.10: Plate flatness. The measurements show a lot of structure and the plate is not flat at all. *On the left:* an overall saddle shape is clearly visible: two corners are lower than the fit, and the other two are higher than the fit. *On the right:* the saddle is corrected for. All the point clouds are now on both sides of the fit, and the individual readings have changed slightly.



As it is not possible to shim negatively, every upright was raised a base amount, and then a (smaller) negative shim was taken off of this.

| Location | Shim (μm) |
|----------|------------------------|
| Corner 1 | 357 |
| Corner 2 | 76 |
| Corner 3 | 250 |
| Corner 4 | 102 |

Table 3.2: Shims under the interface plate.

3.3 Focal plane imager

The focal plane imager (FPI) consists of two cameras on movable Y - and X -stages. One camera is pointed at the sky, the other is pointed at the interface plate. Together they relate the sky coordinates to the interface plate coordinates. See section 4.1 for more information on the different coordinate systems that exist in the positioner.

3.3.1 Depth of focus and seeing

The sky viewing camera is a Bigeye G-123B Cool thermo-electrically cooled camera³. Its (square) pixel size is $6.45 \mu\text{m}$. A point source cannot be imaged smaller than one pixel, so this directly influences not only the image quality, but also the tolerance on the focus. How close to the focus does the camera have to be placed, considering there is no autofocus option in the system? From geometric considerations of the particular setup on the WHT, we can find how much we can move the camera out of focus before the image quality starts to degrade. Let's call this the camera's depth of focus. Figure 3.11 shows the geometry involved.

³Produced by Allied Vision Technologies GmbH, the camera is cooled to $-20 \text{ }^\circ\text{C}$ to suppress the thermal noise background.






Figure 3.11: *On the left:* schematic of the WHT's optical system. *On the right:* how to compute the depth of focus for a given pixel size. It is a zoomed-in view of the schematic in the red dashed box on the left.

We compute the focal length f of our system with the following equation:

$$f = N \cdot D \quad (3.1)$$

where D is the effective aperture of our system, in this case the 4.2 m large primary mirror, and N the f-number of our system. The design of the prime focus corrector makes $N = f/2.75$. In the equations below we can see that the f-number pops up again.

Filling in equation 3.1, the focal length then becomes:

$$f = 2.75 \cdot 4.2 = 11.55 \text{ m.} \quad (3.2)$$

With this information, we can compute the angle θ , which is the missing link between pixel size and depth of focus, as visible in figure 3.11 on the right. This gives us:



$$\begin{aligned}
\tan(\theta) &= \frac{R}{f} = \frac{D}{2f} \\
&= \frac{2.1}{11.55} \\
&\rightarrow \theta = 10.3^\circ \\
&\& \\
\tan(\theta) &= \frac{\text{pixel size}}{2 \cdot \text{depth of focus}} \\
&\rightarrow \\
\text{depth of focus} &= \frac{\text{pixel size} \cdot f}{D} \\
&= \text{pixel size} \cdot N = 6.45 \cdot 2.75 = 18 \text{ } \mu\text{m}.
\end{aligned} \tag{3.3}$$

This means that the camera assembly should not be misaligned vertically by more than $\pm 18 \text{ } \mu\text{m}$, so something smaller than a pixel is not artificially made larger. However, atmospheric turbulence (‘seeing’) smears a point source like a star, and depending on the seeing, this spreading can be larger than the spreading due to pixel size, or smaller. Taking the WHT’s seeing conditions⁴ into account, the median seeing in the best month was about $0.7''$. With a plate scale of

$$\frac{11.55 \cdot 10^6}{206265} \approx 56 \text{ } \mu\text{m}/\text{arcsec} \tag{3.4}$$

each $6.45 \text{ } \mu\text{m}$ pixel is about $0.115''$ on the sky. A median seeing-smearred PSF is therefore sampled by at least 6 pixels. The best seeing was about $0.4''$, which still gives a sampling of about 3.5 pixels. For best-case seeing this means that the pixels are the correct size to capture the information, but for median-case seeing, there are about twice as many pixels as necessary. This means that there is some leeway in the tolerance of the depth of focus. The plate viewing camera has a much larger depth

⁴Site statistics for 2012-2015 at <https://www.ing.iac.es/astrophysics/observing/conditions/>.



of focus due to extra lenses and is not the limiting factor here.

3.3.2 FPI plane and ledge

Due to gravity, the Y -stage and X -stages supporting the FPI will flex. If all else is flat and the four uprights are exactly the same height, the best focus, measured with the FPI on a regular x, y -grid, will form a symmetrical bowl shape. If any of the four posts are slightly higher or lower, this will change the shape of the bowl, and increase the deviations from a flat plane even more. It is of utmost importance to characterise the shape of the deviations from the plane patrolled by the FPI, to ensure that the sky viewing camera can see every part in focus, or as much in focus as possible. For this, the same measurement setup as for the interface plate measurements can be used, but now enhanced with a measurement cone attached to the Y -stage, figure 3.12.

In this figure a small ledge is visible in black. This ledge is the location on which the FPI camera assembly is mounted. It is also important in the further building and metrology of the positioner. Due to space constraints, when the tumbler is put into the structure, the FARO Edge cannot be used anymore. The ledge is a good handover surface from one measurement method to another. It is very flat, with a FARO peak-to-valley measurement of $7\ \mu\text{m}$, and easily accessible.

It is unfortunately not possible to mount the measurement cone onto this ledge, but it does fit on the Y -motor. The cone is then located $113.32 \pm 0.012\ \text{mm}$ above the upper side of the reference ledge, measured with the FARO Edge.

Figure 3.13 shows an early measurement of the deflection in z when the Y -motor is moved along its stage, with the Y -stage at different x -positions. It is clear that some twisting and sloping is happening. This twisting gave us the idea to look at any twisting in the interface plate itself—this measurement is from before all the shimming as described in section 3.2.3 took place. The first (black) row is very out of place and not behaving in the same way as the other rows. This turned out to be






Figure 3.12: FARO measuring setup, same as figure 3.9. The red dot on the lower left post is the location of $(0, 0, 0)$. The moving X - and Y -gantries are explicitly visible. *On the left:* the FARO probe can reach the cone on top of the Y -motor. *On the right:* the grey circle in the middle is the extent of the focal plane. The dots show the approximate locations of the measurements taken by putting the FARO probe in the cone on top of the Y -motor and moving that around with the Y - and X -stages. The grid is slightly larger than the area of the focal plane on purpose.

over-flexing of the FARO joints, and by taking more care with the FARO Edge in that far corner, the discrepancies disappear. The measurements are not yet well-balanced around the centre of the focal plane, as there is an offset involved that I had not accounted for yet. Lastly, I measured each position three times, but I abandoned this method later as the readings were not truly independent, and thus gave a false sense of accuracy.

The uprights which support the stages were not all exactly the same dimensions, and there was some play in the bolts as well. This meant that height differences of about $100 \mu\text{m}$ had to be overcome for some of the uprights. Several iterations of shims between the uprights and the stages removed most of the sloping in both the x - and y -directions. I also narrowed the area of interest to extend only slightly past



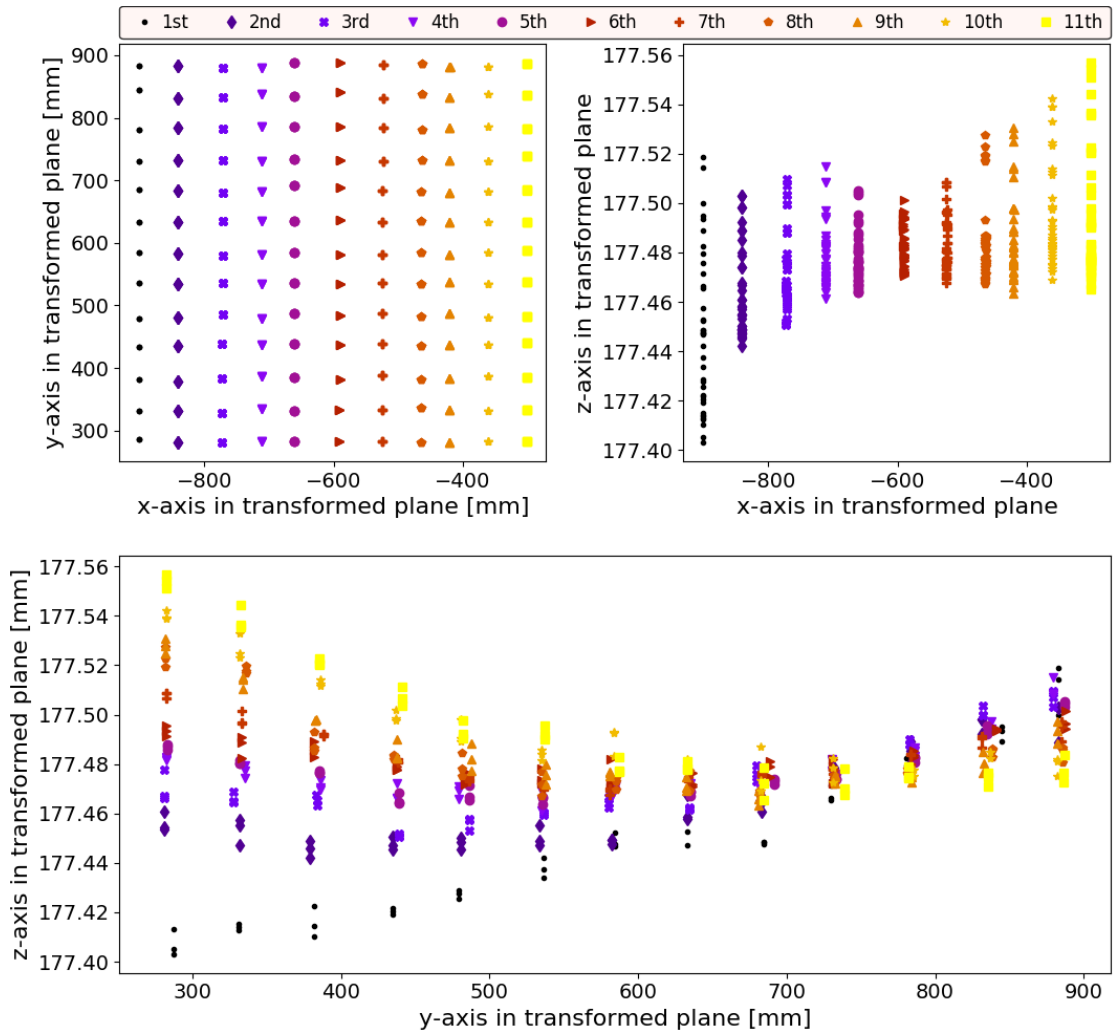


Figure 3.13: Initial result for the FPI measurements. Each colour is a measurement row. Every position is measured 3 times. The black measurements are off due to measurement protocol.

the size of the focal plane. The final result is shown in figure 3.14. The mean of all the readings, and therefore the distance of the measurement cone to the top of the interface plate, is 177.485 mm, and the deflection in the focal plane is only ± 0.02 mm. This is good news, as this is of the order of the depth of focus of the sky-viewing FPI camera.

It is important to note that even though the single-point accuracy of the FARO Edge is ± 30 μm , all the measurements are closer together than that. The measurements are all taken without taking the measurement probe out of the cone, and so



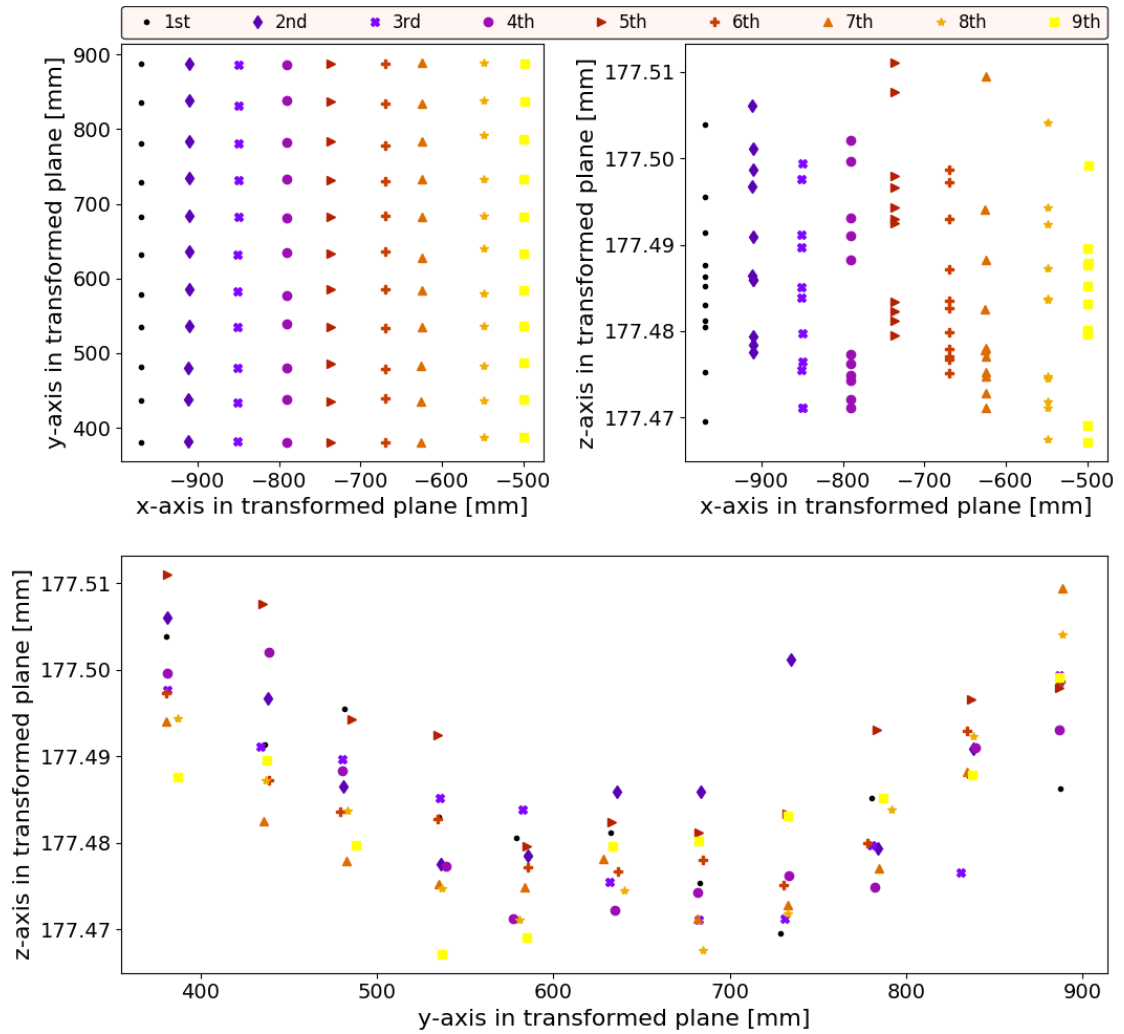


Figure 3.14: Final FPI measurements, centred around the centre of, and extending only slightly around, the focal plane. Each colour is a measurement row. Mean z of all the points is 177.485 mm, and the spread of most points is within $\pm 20 \mu\text{m}$.

they are all related. Any intrinsic error in the system is propagated through all the readings, and is therefore taken out when we talk about differences. So even whilst the results are closer than the single-point accuracy, there are strong indications that the trends seen in each measurement row are real.

Another important note is that there is an offset between the Y -motor and the place where the light enters the camera assembly. I have chosen to align the motor with the centre of the focal plane, as together with the counterweight it adds the most weight to that location, and so has the most influence on the deflection. The



camera is offset by about -4 cm in y , which means that its most extreme points are just positive of the smallest y plotted in the lower half of figure 3.14. In x the offset is about $+11$ cm, bringing it to the right of the largest x plotted in the top right, falling off the plot. However, as that plot shows a mostly flat relation between z and x , this should not pose a problem.

3.4 Tumbler

With the tumbler (overview in figure 3.15) in position, it is not possible to take measurements with the FARO Edge. The ledge described in section 3.3.2 can now be used to complete the measurements on the tumbler itself. The dial gauge is fastened to this ledge with a specially made bracket, as seen in figure 3.16, which also includes a detachable calibration bracket. The calibration bracket simulates the distance of the field plates to the reference ledge.




Figure 3.15: Tumbler assembly. In the build described in this section everything up to both of the retractor support discs is present.






Figure 3.16: The dial gauge attached to the FPI ledge. *On the left:* CAD drawing of the ledge, the dial gauge mounted on the bracket, and the L-shaped calibration bracket at the top. *On the right:* the dial gauge attached to the ledge on the Y-motor. The calibration bracket is taken off to be able to take measurements on the tumbler.

The bracket is designed in such a way that the dial gauge measuring one of the field plates will extend to about half of its travel range of 25 mm. With the calibration bracket attached, the dial gauge reads 12.312 mm (± 0.001 mm). One last measurement with the FARO Edge gives the distance from the upper side of the reference ledge to the calibration bracket: 197.04 ± 0.01 mm.⁵

As the field plates were not yet finished at the time of measuring, the tumbler measurements are taken directly against the retractor support discs (RSDs) underneath. Both field plates are designed to be 9 mm thick, and taking the difference from the reading from the dial gauge to the nominal value of 12.312 mm should give the ‘virtual thickness’ of the plate that is not actually there.

Because the field plates are not in hand yet, the results presented below are preliminary, and give a recipe for the way the final measurements should be done.

⁵The uncertainties on these measurements are much smaller than the single-point accuracy of the FARO. This is because every surface and point (in case of the cone) has been measured several times, and a best fit is taken through these point clouds. The FARO software shows how much the worst points deviate, which was 12 μ m or smaller for these objects. Adding these in RMS fashion gives the quoted uncertainties.



3.4.1 Measurement setup

The RSDs are not as flat as the field plates, as they are not supposed to be used for accurate positioning, and only made to support the field plate and retractors. Therefore it is important to keep in mind that 1) the results below are only preliminary, and 2) we cannot do any in-depth analysis of the surface of these RSDs. I have therefore taken two points on the plate in each direction, to maximise length l in figure 3.17. The difference in dial gauge readings in these positions gives height h , from which we can calculate tilt angle θ . This is true for both the x - and y -direction, and it does not matter whether the distances from the centre are the same or not.




Figure 3.17: Measurement locations on the retractor support disc (RSD).

In figure 3.18 the setup for reaching each location is shown. Along the x -axis, one of the counterweights is aligned to a mark on a ruler temporarily attached to the structure. Because the counterweight counters the mass of the stage itself, the movement is inverted, such that a positive movement of the counterweight is a negative movement of the stage. The image on the left therefore shows the dial gauge in the x_{low} location. The central mark is approximately the centre of the plate; it is just offset to avoid a small hole in the centre of the RSD which the dial gauge can get stuck in. With the field plates, this mark should of course indicate the centre of the plate, as there is no hole there. The counterweight should be in this position when measurements are done in the y -direction.






Figure 3.18: Two rulers. *On the left:* the x -axis. The counterweight is used to position the dial gauge in the correct locations. $l = 387$ mm. *On the right:* the y -axis. The Y -motor itself can be used to position the dial gauge in this direction. $l = 383$ mm. The ruler is taken away for the measurements on the y -axis; the marks are on the stage itself.

Along the y -axis, a more direct placement can be done. The marks on the stage show where the left side of the Y -motor should be aligned to, for the dial gauge to point at the corresponding y -locations. There is no inversion here. Again the central mark is slightly offset to streamline the measurements; the Y -motor is put in this location when the x -direction is measured.

This way of measuring the setup is not extremely precise in x or y ; the error on each placement is about ± 2 mm. How much does this affect the height measurements? The length l is made up of two placements, one on each end. The maximum permissible tilt of the field plate is $12''$, so the expected values for these measurements will be of that order. To take a conservative view of the effect of these inaccuracies in x and y on our measurement of z , we consider double this number, hence $24''$ or 1.164×10^{-4} rad. The small angle approximation is certainly valid for this angle, $h \approx \theta l$, therefore a percentage error in l will give the same percentage error in h (for the same θ). Adding the errors for l in quadrature gives $\pm 2.8 \approx 3$ mm, but even in the worst-case scenario of ± 4 mm, this will not have a measurable effect:



$$\begin{aligned}
h &= \theta l = 1.164 \cdot 10^{-4} \cdot 383 = 44.6 \text{ } \mu\text{m} \\
\text{percentage error in } l &\rightarrow \frac{4}{383} \approx 0.01 \\
\text{error in } h &\rightarrow 0.01 \cdot 44.6 \approx 0.4 \text{ } \mu\text{m}
\end{aligned}
\tag{3.5}$$

An error of under 1 μm is not measurable with the dial gauge, so these are definitely acceptable. Working backwards, a read-off uncertainty of 1 μm in h for an angle of 24" over a length of 383 mm gives about 9 mm of slack on that length. All the calculations with one of the lengths given in figure 3.17 are therefore valid for the other length as well. For clarity only $l = 383$ mm will be used in all the calculations below. This also means that the readings should be offset symmetrically, so from a glance at the readings themselves it should be easy to see whether the tilt is going one way or the other.


A reminder that this height is the difference between two readings. As each reading will have at least a 1 μm read uncertainty, the total read uncertainty will be (RSS) 1.4 μm .

The tumbler motor⁶, which can be used to change the tilt in the x -direction, has a minimum step of 0.001°, or 3.6". Over 383 mm that translates to a height h of 7 μm . Anything smaller than that cannot be compensated for by the motor.

3.4.2 Tumbler tilt adjustments

The tumbler measurement setup explained, it is time to do the measurements. Each plate has three readings in both the x - and y -directions, composed of $[x, y]_{\text{high}}$ — $[x, y]_{\text{low}}$. These are completely independent, as the measurement sequence is low—

⁶The motor is a Newport RV350PE rotation stage, with a minimum incremental motion of 0.001°, a guaranteed uni-directional repeatability of 0.002°, bi-directional repeatability of $\pm 0.002^\circ$, and absolute accuracy of $\pm 0.005^\circ$. The typical specifications are about twice better than these values, except for the incremental motion. (RV datasheet, 10/2013)



high—low—high—low—high. These three readings are averaged to give one value.

I found the correct motor settings by measuring the two locations in x , finding the matching angle, and changing the motor angle by this amount. After a few iterations the measurements converged to $+60.095^\circ$ for plate A, and -119.916° for plate B. Together they are just over 180° , which is reasonable, because the motor now corrects for the tilt in x of each plate.

Any tilt in x is thus easily corrected, but this is not the case for the tilt in y . There is no mechanism to correct for tilt in this direction after the building phase. Fortunately the tumbler axis is attached to the plates with large M8 screws, which means that there is a possibility for shifting the tumbler up or down slightly in the holes for these screws. Figure 3.19 shows the bearing plate with the bearing ring (and the rest of the tumbler) installed.




Figure 3.19: The tumbler is installed. The zoom in shows a clearer view of the screws that provide a vertical adjustment opportunity. The orange circles are around the visible screws, on the bearing end (this side) and the motor end (opposite side).

The bearing ring can be attached to the plate with the screws in the lowest possible positions in the holes, or the screws can be pushed up in the holes as much as possible. I had measured the height difference between these states previously, with the FARO Edge, and found a shift of 0.59 ± 0.04 mm. Even though the whole tumbler axis is



much longer than just the width of the focal plane, this difference is still a lot and can be used to fix the tilt and potentially to fix any overall height issues too. This height requirement is less important though, as this can also be corrected for (by up to a few mm) by re-focusing the telescope.

To change the y -tilt, do the following. First, with the dial gauge, measure the height of the side that will remain fixed. With the lab crane supporting the weight of the tumbler on the other side, that side can be detached from the plate. By lifting or lowering the tumbler end with the crane, the previous result can be matched with the dial gauge now on the detached end. When the screws are tightened again, and the slings are removed, the recently attached end will lower a bit more. So aim for about 80 μm too high, and the tumbler will lower into place after all the weight is taken off the slings⁷.

In this case, before lifting one side, the stationary side (y_{high}) was 2.987 mm and the other side (y_{low}) was 3.251 mm. The pivoting point and y_{high} are still quite far apart, so the value for y_{high} changes too. After a few attempts of following the above method, y_{high} was 2.789 mm, and y_{low} came to 2.786 mm. This is an angle of 1.6", which is well within the tolerance.

Of course this only fixes the tilt in one plate. The other plate has to be shimmed to remove the tilt. This can be done by putting a shim between the RSD and the field plate. We have not done this, as the field plates are not finished yet, but the procedure should be relatively simple.

3.4.3 Height of RSD

With the tilt in both x and y measured and corrected for, at least for one of the plates, it is time to determine the overall height of each RSD below the zero point of the dial gauge. Each point on plate A has 3 readings, which are averaged. All the

⁷This is not the whole story, see section 3.4.3 where I argue that another 24 μm should be added to get to the correct height.



readings are within $8\ \mu\text{m}$ of each other, so the errors are at most $4\ \mu\text{m}$.

It is important to note that a dial gauge works by pushing in its tip. A fully-extended tip is $0.000\ \text{mm}$, and any lowering of the tip produces a larger read-off value. Therefore larger read-off values mean that that section of the RSD is closer to the reference ledge, and where the RSD is further away, smaller values are read off.




Figure 3.20: Baseline measurements for plate A and plate B. The values are the averaged dial gauge values for each location. These values are taken as the zero values for figures 3.21 and 3.22.

Figure 3.20 shows all the locations for the height measurements, with the averaged dial gauge value quoted, both for plate A and plate B. The z -values for x_{low} and x_{high} for both plates are almost the same, as the differences are calibrated out by the motor angle adjustments. Any remaining difference is smaller than one step in the motor.

The y_{low} value of $z = 2.810\ \text{mm}$ for plate A is unexpected. In the previous section the tilt in the y -direction was corrected for, and the measured values clustered around $z = 2.786\ \text{mm}$. The measurements in figure 3.20 were done several days later, and therein could lie the answer to this difference. It is not unreasonable to assume that the full relaxation of the structure after reattaching the tumbler beam took more time than was given between reattachment and measurement. In that case the beam should be lifted an extra $24\ \mu\text{m}$ above the final dial gauge reading. Section 3.4.4



shows that the value for y_{low} did not change significantly over a day of tumbling and measurements.

Another puzzling thing is the big difference between the z -values for the x - and y -locations, for both plates. Plate B is tilted in the y -direction, but even with this tilt taken out, the z -values will not match those for the x -direction. In plate A there is at least 60 μm difference, and in plate B it is more of the order of 180 μm . The RSD is not as stiff as the field plate, but it is still a large difference. With 168 retractors with 6 fibres each attached to each RSD and top plate in the near future, the RSDs should stiffen up as well. A measurement done previously with the micrometer found that the thickness of both RSD's inner rims vary around 4.960 ± 0.015 mm. Assuming that the rest of the RSD shows the same behaviour, this could account for part of the variation between the z -values for x and y in both plates. These changes are conceivably averaged out when the field plates are placed on the RSDs and their surfaces are measured.

Lastly, the height for plate A clusters between $z \sim 2.7 - 2.8$ mm, and for plate B $z \sim 2.4$ mm (from figure 3.20), which is an overall height difference of 300–400 μm . A separate set of measurements I had done previously was to check the height of each drum, the tall and wide tube that the RSD is bolted on. These measurements, also with the FARO Edge, showed that one of the drums is 307.96 mm, and the other 308.01 mm, for a nominal length of 308 mm. After anodisation it was impossible to tell which drum was which, but it stands to reason that part of the overall height discrepancy between plate A and plate B is explained by this difference. Another element could be the tumbler beam. This part was never measured, but it could have been attached slightly off-axis to the tumbler motor, which then results in an increase in height on one side, and a decrease on the other side.



3.4.4 Motor repeatability

All that remains is to check whether the measurements remain stable after multiple tumbles. To check this, we simulated tumbles for a whole night, measuring first plate A in the observing mode. After that, tumble to plate B, and repeat this until each plate has been measured seven times. 14 plates is more than a whole night, but 14 tumbles does allow for movements to the LIFU and back during one night. The setup explained previously in section 3.4.1 results in 2×3 readings each for x and y . The readings are averaged and the results are presented in figures 3.21 and 3.22.




Figure 3.21: Tumbler repeatability tests for plate A. The small angle differences ($< 2.7''$) are explained by measurement error, the large shifts in x are probably the tumbler over- or understepping by 1 step.

Figure 3.21 shows the result for plate A. All the measurements are compared to the first measurement, displayed in figure 3.20, so only the difference is shown. On the left the average height difference measurements for each x_{low} , y_{low} , x_{high} , y_{high} are plotted. On the right the results are translated into angles. Again, these are not the absolute angles, just the difference to the first measurement.

The vertical dotted line just after the sixth measurement shows where we rehomed the tumbler to check whether that would have any influence on the positioning. Re-







Figure 3.22: Tumbler repeatability tests for plate B. The y -limits are the same as for figure 3.21. Again the small angle differences are explained by measurement error, and the last jump in x is due to the tumbler motor.

homing involved switching off the motor, then switching back on, initialising and homing. This means that the flow of the measurements was interrupted, and the motor encoders were able to do an extra check on the position.

The largest contributor to the error on the measurements is the averaging of three readings to get one point; some of the readings for x_{high} can differ up to $8\ \mu\text{m}$. The error on the average is then about $4\ \mu\text{m}$. With a read-off error of $1\ \mu\text{m}$ added in quadrature, the error remains $4\ \mu\text{m}$. The errors for x_{low} are slightly lower overall, $2 - 3\ \mu\text{m}$, and the errors for y_{low} and y_{high} are all $< 2\ \mu\text{m}$. The different readings are likely caused by the local variation in RSD thickness due to machining.

The results show a small variation in angle. The smallest change the tumbler itself can make is $3.6''$ and this ‘jump’ is visible in the x -direction at A1, A2, A4, and A6. In y no such jump is expected, as it is perpendicular to the tumbling direction, and the overall spread of the angle differences is even smaller than in the tumbler direction. The smaller variations are due to the measurement errors: with the errors for the low and high point added in quadrature, the measurement error in x is up to $5\ \mu\text{m}$, which gives a maximum angle error of $2.7''$.



Similar results can be seen in figure 3.22 for plate B. The results show that the repeatability in the tumbler motor is more than sufficient to keep the field plates in the correct orientation. The tumbler motor will sometimes correct by going 1 step further or back than previously, but this is allowed in the tilt budget of 12". The y -results show that tumbling has an almost negligible effect on changes in angles perpendicular to the tumbling direction.

3.5 Distance from field plate to interface

Figure 3.23 shows the different lengths involved to find the distance between the interface of the positioner and the field plate. The designed height of each field plate above the interface is $A + B + C$. However, as stated above, the field plates are not in hand yet, so all the measurements are up to the RSDs. Let's look at all the separate components first, before adding them up to the total height.

3.5.1 Interface plate thickness (C)

The measured thickness of the interface plate, 12.00 ± 0.02 mm (section 3.2.2), agrees well with the designed thickness of 12 mm. The local variations under the positioner supports are expected to be averaged out by those supports. Any variations in other places do not matter, because they do not add to the distance of the field plate to the interface.

3.5.2 Top of interface plate to FPI ledge (B)

This distance is not directly defined in the drawings, but from the lengths of other parts in the drawings we can derive the nominal distance of 65 mm. The measurement consists of two components. The first component is the distance from the top of the interface plate to the top of the Y -motor (D), and the second component is from the






Figure 3.23: Figure 3.3 again, with some additions: the calibration L-bracket and the RSD are now located. The L-bracket is made to emulate the location of the field plate, but is not necessarily exactly at the same height (it may also be higher than the field plate). A , B , and C are designed values, D and E are measured values to reconstruct B , and F is a measured value to reconstruct A . The dial gauge read 12.312 mm when touching the L-bracket. Table 3.3 has the designed and measured values.



Y-motor back down to the ledge (E).

Both of these measurement sets are described in section 3.3.2. E is found to be 113.32 ± 0.012 mm and D , from figure 3.14, is 177.485 ± 0.02 mm. These measurements are independent, so the errors can be added in quadrature (RSS) for the distance B :

$$\begin{aligned} B &= D - E \\ &= 177.485 \pm 0.02 - 113.32 \pm 0.012 \\ &= 64.17 \pm 0.023 \text{ mm.} \end{aligned} \tag{3.6}$$

This value is much smaller than the nominal distance, and could pose a problem if all the other heights are as designed (or shorter) too. Distances that are too small (i.e. the field plate surface is too close to the interface flange) cannot be fixed within the tumbler; the whole structure would need to be shimmed with respect to the interface. It could also be that the FPI ledge itself is not in exactly the correct place, which would affect the height of the sky- and field-viewing cameras, and therefore whether the sky and field plate are in focus or not.

If this is the case, the exact positions of the cameras can be adjusted to ensure that their focus is again on the field plate and the sky. There is enough leeway for this in the FPI box that houses both cameras.

3.5.3 FPI ledge to RSD ($A +$ field plate)

The distance from the ledge to the field plate is also not directly defined in the drawings, but can again be derived from the parts involved in this part of the build. It is supposed to be 197 mm. Each field plate is designed to be 9 mm thick, for a combined height of 208 mm.

This distance is found by measuring the distance from the FPI ledge to the L-bracket (A), noting the value on the dial gauge when the gauge touches this bracket,



and then measuring the height of the RSDs in dial gauge coordinates to find the difference between these two values (A). The L-bracket is made to be the same distance from the ledge as the field plate.

The measured distance A is 197.04 ± 0.01 mm, so it is quite close to the designed value. With the dial gauge in the L-bracket, the display reads 12.312 mm. This value is important for the calculation of distance F , which is the difference between this reading and the readings on the two RSDs.

The difference in dial gauge readings between the L-bracket and the RSD (F) can be thought of as the ‘virtual’ or ‘derived’ field plate height. There are two values, one for each RSD (section 3.4.3). In case of plate A, the average reading on the RSD is 2.75 ± 0.05 mm, whereas for plate B, it is 2.4 ± 0.1 mm. RSD B is more uncertain, because the plate was still at an angle when measured.

The derived field plate heights then become:


$$\begin{aligned} \text{plate A: } & 12.312 \pm 0.001 - 2.75 \pm 0.05 = 9.56 \pm 0.05 \text{ mm;} \\ \text{plate B: } & 12.312 \pm 0.001 - 2.4 \pm 0.1 = 9.91 \pm 0.11 \text{ mm.} \end{aligned} \tag{3.7}$$

These values are much larger than the nominal field plate values, which suggests that the RSDs are too far away from the ledge. Before we can conclude anything from this however, we have to take all the measurements into account.

3.5.4 All measurements together

The designed and measured values are collated in table 3.3.

The nominal height of each RSD above the interface, going from the bottom to



| Distance | Designed (mm) | Actual (mm) | Derived (mm) |
|----------|------------------|--|-----------------|
| <i>A</i> | 197 | 197.04 ± 0.01 | - |
| <i>B</i> | 65 | - | 64.17 ± 0.023 |
| <i>C</i> | 12 | 12.00 ± 0.02 | - |
| <i>D</i> | - | 177.48 ± 0.02 | - |
| <i>E</i> | - | 113.32 ± 0.012 | - |
| <i>F</i> | 9 | to RSD A 9.56 ± 0.05 to RSD B 9.91 ± 0.11 | - |

Table 3.3: Designed, actual, and derived values for all the distances defined in figure 3.23, the last two with their measurement errors. Length *B* is calculated by subtracting *E* from *D*. There are no tolerances on the design values. Many different parts together make up the final assembly and some of the differences between the designed and actual values for parts can cancel each other out. There is no specific tolerance on sub-assemblies, as only the total counts. The parts are typically machined to within about 5–10 μm from the specification. The total tolerance is ± 0.1 mm for the distance between the field plate and the interface to the PFC, here given by $A + B + C$, on assembly.

the top in figure 3.23, is given by:

$$\begin{aligned}
C + B + A + F &= \\
&= 197 + 65 + 12 + 9 \\
&= 283 \text{ mm.}
\end{aligned} \tag{3.8}$$

The measured height is, again going from the bottom to the top in figure 3.23:

$$C + (D - E) + A + F. \tag{3.9}$$

Filling in this equation for side A gives:

$$\begin{aligned}
C + (D - E) + A + F &= \\
&= 12.00 + 64.17 + 197.04 + 9.56 \\
&= 282.77 \text{ mm}
\end{aligned} \tag{3.10}$$

$$\text{with error} = \sqrt{0.02^2 + 0.023^2 + 0.01^2 + 0.05^2} = 0.06 \text{ mm.}$$



So the actual distance between the RSD and the interface is slightly too short, which means that the field plate needs a negative shim. As this is impossible, one solution would be to shim the whole positioner with respect to the rotator. The deficit is 0.23 mm, or 0.13 mm to just be within the tolerance amount, which could also be made up if one of the field plates is too thin by that amount. This is not likely to happen, as it is a rather large amount to make up with one part.

An option to fix this would be to raise the tumbler axis as much as possible with the procedure described in section 3.4.2, and then readjusting the y -axis tilt for one of the plates. The previous time the tilt was adjusted, it involved a fair amount of lowering of both ends, which means that there is room to raise both sides somewhat. As the field plates have to be installed and the tilt has to be rechecked anyway, this is a good solution. Most likely both field plates then need a positive shim, but this is not a problem, and certainly preferred compared to shimming the whole positioner with respect to the rotator.

On side B we get:

$$\begin{aligned}
 C + (D - E) + A + F &= \\
 &= 12.00 + 64.17 + 197.04 + 9.91 \\
 &= 283.12 \text{ mm}
 \end{aligned}
 \tag{3.11}$$

$$\text{with error} = \sqrt{0.02^2 + 0.023^2 + 0.01^2 + 0.11^2} = 0.11 \text{ mm.}$$

The distance for RSD B to the interface is 283.12 mm, and so also not within the specification. This however, is easily remedied by adding a shim on the tumbler. It is encouraging that only one of the sides is slightly too long, as it makes it easier to fix this.



3.6 Summary

The measurements show that it in principle the field plates can be set up in such a way that they occupy the same space and tilt (to within 0.1 mm and 12") when they are in the observing position. Some extra shimming may be needed. When the field plates are in hand, the measurements from section 3.4 onwards have to be redone to find the actual distances to the field plates, instead of the above derived distances.

Any deficit in this distance can most likely be compensated for by lifting the whole tumbler axis on its M8 screws. As the tilt of the y -axis may have to be changed anyway when the field plates are installed, this may be the most straightforward solution.



Chapter 4

Robot calibration software

The positioner not only has to be built to a very high standard (previously described in chapter 3), various subsystems also need very careful calibration. It is not enough to have the field plate in focus when it is in the observing position, the fibres have to be placed in the correct (x, y) -position on that field plate too. As the combined tolerance on this ‘correct’ position is only 8 μm (see also table 2.4), this means that there is not much leeway.

Each of the different steps involved in the accurate placement of these fibres add a certain error. The goal is to minimise each of these errors. In this chapter I will detail how a certain position on the sky is translated to fibre positioner coordinates, and how this translation can be reverified and recalibrated whenever is necessary. The grid marks on each field plate are essential for this, as is the software that I have written, which I will introduce at the end of section 4.3.

4.1 Coordinate systems

The positioner works with and exists in several coordinate systems. They are the sky coordinates in (RA, Dec), the robot coordinates in (x, y, z, θ) for Morta and Nona on the positioning side, the robot coordinates (x, y) for the FPI on the sky imaging side,



and the plate coordinates (x, y) for each field plate. To be able to place a fibre button in the correct position, all these coordinate systems have to be expressed relative to each other: for a target's location on the sky and given the desired telescope pointing centre its corresponding location on either field plate has to be calculated, and the robots need to know that location to actually put the button down correctly.




Photo kindly provided by Jonathan Brodie of Angel Sharp Media.

Figure 4.1: The fibre positioner robots Morta (on the left) and Nona (on the right). They are on their own Y, Z, Θ -gantries, and share the two X -gantries. Situation shown in May 2018, before anodisation of the positioner took place, and with the stand-in aluminium field plate.

A good way to link these coordinate systems is to make a grid of reference marks on each field plate that can be surveyed by all three of the robots' cameras. This correlates the field plate coordinates to each set of robot coordinates. The field plate coordinates are then linked to the sky coordinates with the FPI that surveys the grid, and the Bigeye camera that looks at the sky. The grid marks are expressed in plate coordinates.

An additional benefit of surveying a grid, is that it can be used to determine parameters of the measurement system itself. A grid with known reference marks, for example by measuring it with a precise measuring tool like a coordinate-measuring



machine, can then be measured by the positioning robots Morta and Nona, and the FPI robot. Any change in the grid mark locations is caused by the measurement system. These can be linear changes such as an offset, a scale change, a rotation, or a shear (non-perpendicularity of the robot axes), or non-linear changes due to a non-straightness in some or all of the robot axes, and flexure of the axes due to gravity. As the robots are used for precise positioning and location finding, any error in their actual location as opposed to where they encoders think they are, has to be corrected for as much as possible.

This technique of measuring an existing grid was used on past instruments, in particular Autofib-2 on the WHT (Lewis et al., 1997) and 2dF on the AAT (Lewis et al., 2002). Autofib-2 had a 10 mm spaced grid of dots on a 14-inch square glass plate and calibration was performed off the telescope. This calibration did not involve the field plate, only the robot axes, as the field plate itself had to be removed to install the calibration plate. 2dF used 21 reference marks on the field plate for offset and rotation information, and because of the tumbler had to do these calibration measurements after each tumble (and before reconfiguring any fibres). The grid was not fine enough to correct for other linear parameters or non-straightness of the axes. This had to be done separately by surveying fields of known bright stars at different (RA, Dec), and were called *poschecks* in 2dF terminology.

On WEAVE those two approaches are combined: each field plate has a reference grid with dots every 20 mm machined into it. Each hole of 0.5 mm diameter is filled with black paint¹ to make it stand out against the plate itself. The plate is illuminated by a ring of LEDs around the robot camera's plate-viewing focus to increase the contrast. Figure 4.2 shows part of the prototype field plate, with a few grid marks.

With this combined approach the FPI-Bigeye system and the fibre positioner robots can now be recalibrated after every tumble. This is needed as due to the tumbling

¹Section 3.7.2 in Gilbert (2016) has more information on the prototyping.






Figure 4.2: A close-up of the grid machined into the prototype field plate, in this case made of aluminium. Each point is 0.5 mm wide, and the ones visible here are filled with black paint and lightly sanded.

the field plates will not be in exactly the same place as before. If not corrected for, this would make both the conversion from sky to field plate coordinates and the fibre positioning itself less precise. It is also necessary to recalibrate after each time the telescope is slewed to a different field on the sky, as the angle with respect to the gravity vector has now changed. It therefore makes sense to do a grid calibration run every time just before a field plate is reconfigured².

Another new component in this approach is that it is now possible to simulate the behaviour of the positioner, and confirm how each external influence (e.g. temperature or zenith distance) affects the proper imaging of this grid. Correcting for the linear and non-linear components together is complex, as different methods to correct for the non-linear components affect the correction of the linear components. This will be explained in detail in section 4.5.

²More information on how often science calibrations will have to be done is detailed in [Benn et al. \(2017\)](#). The details will be finalised during commissioning, depending on the instrument's performance.



4.2 Robot limitations

Surveying the grid and the placement of the fibre buttons will not be 100% accurate because of limitations in the positioner robots. Positioning of the buttons is not only a combination of movement in the x -, y -, z -, and θ -directions of the robot, but also of the amount of twist or wave (non-straightness) in the x - and y -axes, and the actual angles between the axes. On top of that, the verification of a position depends on the precision with which the camera systems can find the locations of the grid marks and the fibres themselves. The biscuit cutting method described in [Gilbert \(2016, p. 76-77\)](#) is precise up to N^{th} of a pixel, with $N = 10$ or about $0.5 \mu\text{m}$ for the large (0.5 mm) grid marks, and $N = 50$ or about $0.09 \mu\text{m}$ for the fibres.

Two distinct mechanisms make the physical stages not perfectly straight. The first is due to the installation, as each stage is only clamped down to a flat surface. This causes a fixed wave pattern in the stage, perpendicular to the axis. The second is flexure due to gravity and the movement of the robots on the axes. Together these two mechanisms result in what is called ‘non-straightness’ of the axes throughout this thesis.

Not taking the z - and θ -axis of each robot into account here, two axes (x and y) work together to move each robot to the correct position above the field plate. This introduces four possible sources of error that can be described by six parameters. These are:

1. An offset in both x and y , between the centre point of the grid and the corresponding encoder values in the robots;
2. A scale factor in both x and y , for example because the robot gantries do not move not completely parallel to the field plate in z . If a gantry is at an angle in z , a move measured with the encoders will give a smaller real move in that axis. Temperature differences will also cause scale changes, because of expansion or







Figure 4.3: Schematic view of XY -straightness of one of the robots. *On the left:* non-straightness at zenith (gravity perpendicular to robot plane) is only due to installation. *On the right:* non-straightness for zenith distance $\gg 0$, where both the fixed installation effects and a sag due to gravity make up the total.

contraction of the axes;

3. A rotation, for example because the field plate is not fully lined up with the robot axes;
4. A non-perpendicularity, because the y -gantry and the x -gantries are not installed completely perpendicular to each other. This is also called shear.

Prior to installation the grid will first be measured with an independent measuring tool, which has its own precision. The uncertainties are all uncorrelated, and are therefore random errors for each grid mark. The next measurements are done with the robots, which currently have a combined positioning uncertainty of $\sim 5.5 \mu\text{m}$ for the X - and Y -gantries together³. All these positioning errors are random in nature too.

³See section 5.4.3.2 for the measurements. Investigations are ongoing to improve these values.



4.3 Autofib-2 legacy

The Autofib-2 instrument (Parry et al., 1994) had a separate photographic glass plate with a square grid on it (10 mm spacing) that would be put in place of the field plate whenever recalibration had to be done. This meant that for every recalibration, the instrument had to be taken off the telescope, the field plate removed, and the calibration plate put in its place. Every time this was done, the calibration plate would have a slightly different rotation, as this could not be very accurately controlled⁴. These changing rotation values proved to be important for the development of the calibration software.

The calibration software developed by Ian Lewis and Ray Sharples (Lewis et al., 1997) used a Fortran library called `slalib` (Wallace, 1994), software made specifically for positional astronomy. The routines in this library use linear transformations to remove offset, rotation, scaling and non-perpendicularity of the axes from a set of (x, y) -coordinates when compared to a reference set of coordinates. After these linear components to the errors in the grid were taken out, the non-linear components (the non-straightness in the axes) could be corrected for by averaging the error over a row or column of points in the grid.

Because the rotation was such a large value, the decision was made⁵ to first do a partial linear fit with a solid body rotation, then to correct for the non-linear components, and lastly to do a full linear fit where all the parameters were used independently to find the best solution. The differences between these two fits are explained in section 4.4.1.

The recipe to do the recalibration of the robot axes was this:

1. Measure the grid with the positioning robot;

⁴A rotation did not matter for the field plate, as that was circular without any preference for alignment.

⁵Communication with Ian Lewis.



2. Feed those (x, y) -values into the recalibration software, together with the expected (x, y) -values acquired with an independent measurement tool;
3. Use the `slalib` routines with a solid body rotation to take out the offset, some scaling effects, and rotation between the two sets. The non-straightness of the axes is left over, together with shear and some other scaling effects;
4. Calculate the non-straightness of the axes by averaging the remaining offset in each row and column in the grid. Each point in the grid will have two correction values: one for the mean offset of the row it is in, and the other for the mean offset of the column it is in;
5. Take the two non-straightness correction arrays and correct for the straightness in a copy of the *original* measured grid, i.e. the one in item 1;
6. Then use the `slalib` routines again on this straightness-corrected set, now without any limitations on the fitting parameters to also correct for any leftover scaling and shear effects that were not caught in the non-straightness arrays.

As detailed in section 4.2, the (non)-straightness of an axis is defined as the amount of ‘wobble’ perpendicular to that axis. Every measurement with a fixed value for x (but different y values) has the same amount of y -offset because of physical distortions at that particular location, and the other way round.

These straightness corrections for Autofib-2 were calculated in the following way. First the data is sorted into columns. For each column, so for each group of points with a similar x -value, the mean offset in y is calculated. This is the X -straightness. The same is done for each row, where the mean offset in x is calculated for each group of points in y , and is called the Y -straightness. Figure 4.4 shows an example of what a non-straightness in a stage could look like (visualised using the modern software explained below).






Figure 4.4: Example of straightness arrays. The offset in y for each point in x is the X -straightness, and the offset in x for each point in y is the Y -straightness. The stage in the x -direction is simulated with 3 sine waves with amplitude $20\ \mu\text{m}$, and the stage in the y -direction has 3.5 waves, shifted by 90° , with an amplitude of $14\ \mu\text{m}$.

This recipe for Autofib-2 is the basis for the software that I have written for WEAVE, which I have called `GridKit`⁶, to do these linear and non-linear corrections in this manner, and not only on ‘real’ measured grids, but also in simulations. The plots in the next sections are made with this software, and a comprehensive description can be found in section 4.5.

4.4 Differences between calibration approaches of Autofib-2 and WEAVE

Both Autofib-2 and 2dF showed that having a grid to calibrate the robot axes is extremely useful. But both of the approaches had their problems too. Autofib-2 had a photographic plate with a grid on it that had to be attached to the instrument to do the calibration, so every time the plate was used, there would be extra offset and rotation errors that had nothing to do with the robot axes. The photographic plate also had to be illuminated to see the grid dots, and this changed the temperature of the glass and the robots. This meant that an extra temperature compensation had to be done to find the grid at operational temperature.

⁶`GridKit` is written in Python ([Python Software Foundation](https://www.python.org/)) version 2.7.13 with NumPy ([van der Walt et al., 2011](https://numpy.org/)) version 1.11. This version of Python is now unfortunately deprecated, but `GridKit` can still be run on machines with Python 2.7 installed.



The solution to the problem of the changing offset, rotation, and temperature errors was to include a grid into the field plate. This was done for 2dF, where 21 reference marks were embedded in a regular grid pattern. These reference marks consisted of polished optical fibres inserted into holes in the field plate, with the other ends of the fibres illuminated by LEDs (Lewis et al., 2002). Unfortunately, this approach introduced new problems. The number of fibres was too small to find information on the non-straightness of the axes. On top of that, even illumination of the fibres proved to be difficult because of the LEDs and the alignment involved. This led to the recalibration with the grid being done at the start of every observation (after every tumble), but the `poschecks` only after each re-installation of the instrument on the telescope.

Including a reference grid in the field plates themselves proved to be the correct approach, as, like 2dF, WEAVE now also includes a tumbler that tumbles the field plates from the configuring side to the observing side. It becomes critical to check what happens to the field plate's position from one side of the instrument to the other, as the coordinates for the robots on the configuring side are not directly linked to the sky coordinates anymore.

With a fine grid (20 mm spacing) machined directly into the field plates, it is again possible to do a full recalibration for both the linear and the non-linear parameters (no separate `poschecks` necessary). Another bonus is that the recalibration can be done at any time. For every reconfiguration, a quick recalibration scan of the grid can be performed, on both sides of the tumbler. On the configuring side this immediately takes into account the flexure at the zenith distance for which the next set of observations is reconfigured. On the observing side any shifts in the field plate due to tumbling can be monitored.

The caveat is that not all of the points in the grid are available for surveying, as up to 65% of them (Lewis et al., 2017) are estimated to be obscured by the previous



configuration of fibres. On the configuring side the robots have to know the new situation before they can accurately remove the fibres from the old configuration into the new one, and the best time to do this is before any fibre is removed. On the observing side, the fibres obviously stay as they are, and the FPI has to work off what is still visible between the fibres.

4.4.1 Solid body rotation versus full 6-parameter fit

One of the `slalib` routines used for the calibration is called `SLA_FITXY`⁷. This routine fits a linear model to two sets of (x, y) coordinates, here to an ‘expected’ grid and a ‘measured’ grid. The ‘expected’ grid is the reference grid that is measured with an independent measuring system, in the case of Autofib-2 with APM and COSMOS.

`SLA_FITXY` can fit these sets of (x, y) coordinates with two models: one where all six coefficients are independent, and one where only four are independent. When all six coefficients are independent, offset (x and y), scale (x and y), rotation, and non-perpendicularity are modelled. The other option is the solid body rotation, where the offset and rotation are modelled unchanged, but the scale coefficients now have to be the same in absolute value, and there is no non-perpendicularity.

This is important because to take out all the components, the Autofib-2 algorithm will first fit as much as possible with linear transformations. This means that some of the non-linear effects are also somewhat compensated for by the linear components. This is unavoidable. By first fitting with a solid body rotation, no fictitious shear can be induced by the non-linear components. However, the flip side is that any real non-perpendicularity will then be caught in the non-straightness. The reverse is true for rotation. If the first fit is a solid body rotation, any non-straightness can only be compensated for by an offset and a rotation. There will be more rotation than is truly present, which would happen less if the full 6-parameter fit is available. In that

⁷See <http://starlink.rl.ac.uk/star/docs/sun67.htx/sun67ss92.html> for information on this routine.



case both rotation and non-perpendicularity share the non-straightness effects they are trying to compensate for.

It depends on what the unadjusted grid looks like, whether it is better to choose the 4-parameter fit or the 6-parameter fit. If there is a non-negligible shear, it is better to go for a full 6-parameter fit. If there is no shear, or its effects are buried under the effects of the other parameters, it can be better to use the 4-parameter fit, the solid body rotation.

For Autofib-2 the rotation value that changed for each calibration drowned out all the other contributions, so it was more useful to first correct for the rotation, before doing any other corrections. This is not necessarily true in WEAVE.

Figures 4.5 and 4.6 show how simulations with only non-linear errors are corrected for with a 4-parameter fit, or a 6-parameter fit, before going to the non-linear calculations.

First, two cosine waves are simulated into the x -gantry. A cosine is perfectly symmetrical around both the x - and y -axis, and therefore `GridKit` cannot try to compensate for this with any linear routines. All of the non-linearity is left for the non-linear routines to do something with. The top straightness array shows that the cosines are recovered perfectly. The amplitude of the waves is such that it is easy to see what is happening in the grid plot, it has no relation to real-world effects.

The next row in figure 4.5 shows what happens to two sine waves. They are not symmetrical around the y -axis, and that is why the software introduces a rotation. This is a non-real rotation, and is therefore corrected again by the non-linear part of `GridKit`. The last row shows four sine waves. Two sine waves on each side make the plot more symmetrical, which is why the introduced rotation is smaller.

If the same is done with the full 6-parameter fit first, the outcome is as in figure 4.6. The simulation returns the same value for the rotation as in the 4-parameter fit, but now the y -axis has not rotated with the x -axis. This is because the induced rotation in



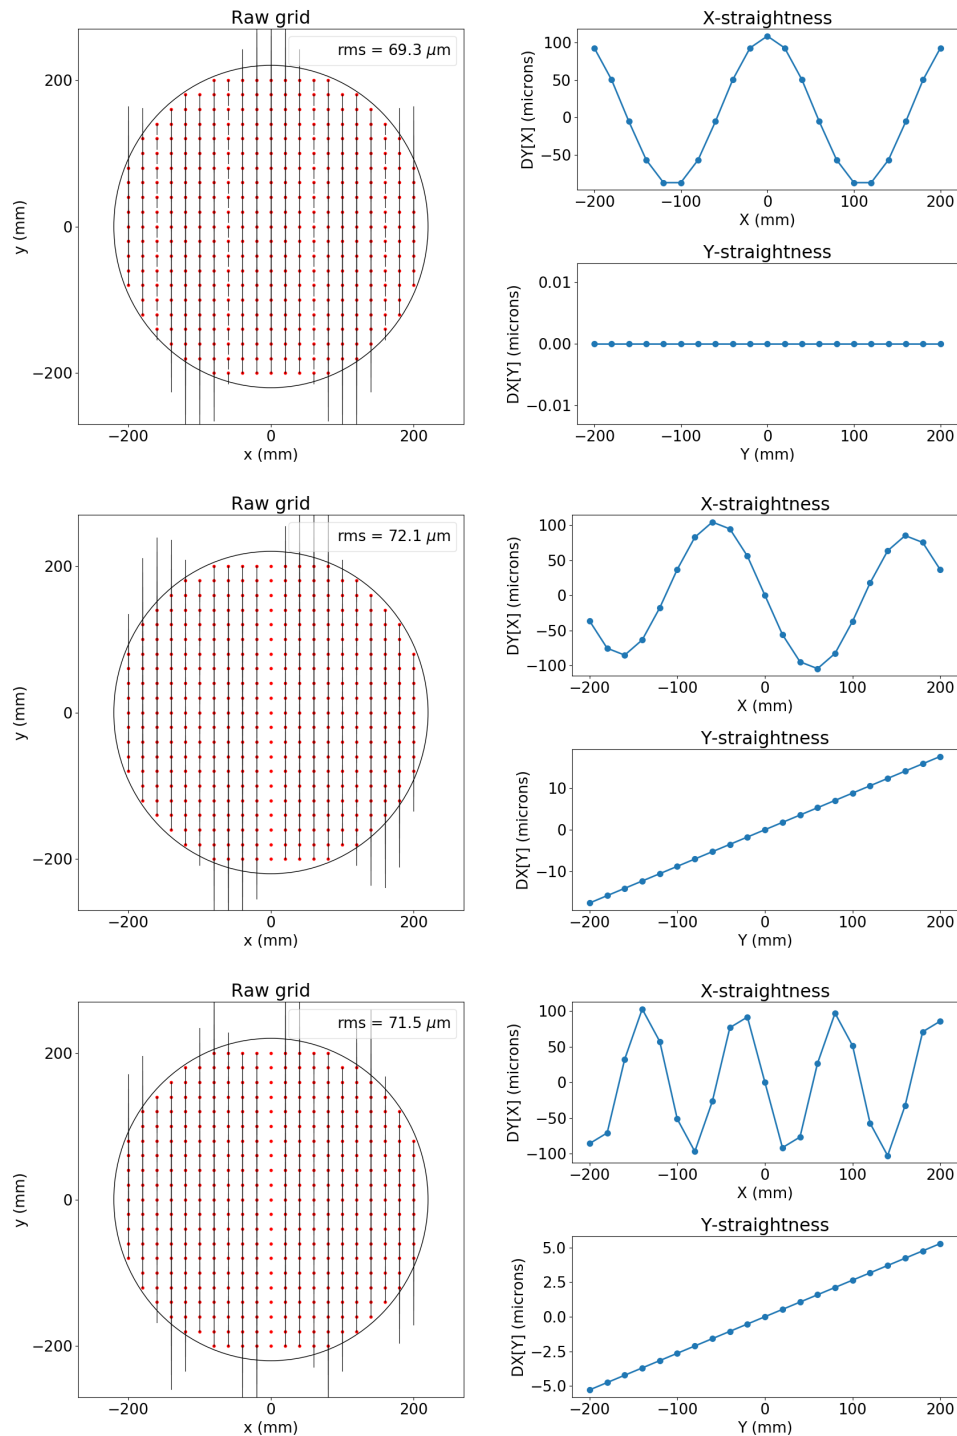


Figure 4.5: 4-parameter fits. The residual vectors are scaled such that the distance between two dots in the grid is 20 μm . *Top row:* a non-linearity in the X-gantry with an amplitude of 0.1 mm. The software recovers this perfectly. *Middle row:* the same non-linearity, but now phase-shifted 90°. There is an overcompensation in rotation. *Bottom row:* 4 sine waves with the same amplitude as before. The overcompensation is less obvious.



the y -axis is now compensated with a shear. The two sine waves introduce a rotation of -0.005° , which is the same as in figure 4.5, but the non-perpendicularity is 0.01° . Similarly, the rotation for four sine waves is -0.0015° , and the non-perpendicularity is 0.003° .




Figure 4.6: The same non-linearities as in the middle and bottom row of figure 4.5, but now with the 6-parameter fit. The linear component is taken out of the Y -straightness, but the effect on the X -straightness is slightly worse.

Of course mathematically it does not matter whether the shear and rotation, or the non-linearity, are compensated for by the linear routines, or the non-linear routines. It is interesting from a physical perspective, where each parameter means something for the positioner as it is built. If the parameter values align more closely to the ‘real world’ values, it may be possible to physically correct for some of the larger ones. For example, if a large rotation is present, the field plate could be rotated the other way to counteract this. This cannot be done if a parameter is overestimated. The most important thing though, is that the different sources of error are removed as much as possible. And if that is through a combination of overestimation of one parameter and underestimating another, that is fine too.



4.4.2 Unavailable grid points

With the calibration grid incorporated in the field plates, recalibration can be done whilst the instrument is in use. This presents some unique challenges for the WEAVE case that were not present in Autofib-2. With the fibres still configured for the previous field, up to 65% of the grid marks can be obscured. This means that any software has to be able to deal with changing numbers of grid points in a flexible way. This is most important in the case of the straightness calculation.

The algorithm described in section 4.3 (first a solid body rotation, then the straightness correction, and lastly a full 6-parameter fit) on a partial grid gives some surprising results. Figure 4.7 shows what happens when the Autofib-2 algorithm is directly applied to a grid with an x -scale of 1.0015, and y -scale 1.0, missing about 10% of its points. This is a clear problem. When all the grid points are available, the software neatly recovers the parameters without any residuals, and flat straightness arrays.

It looks as if the straightness calculation is off for some rows and columns, but it turns out that what happens here, is again a sort of overcompensation. The first linear correction can only fix the scale in specific ways, which here results in the scale in x being evenly distributed over x and y . This means that for each point in the grid, the errors now have components in x and y . Looking at the points in one column, the straightness is calculated by taking the mean offset. For a full column each point in the column is balanced by a point on the other side of zero, but if a point is obscured, the mirror point is not balanced. This gives rise to a non-zero non-straightness value for that column; the results are plotted in the lower left of figure 4.7.

So the mean non-straightness for each column and row is minimised by the non-linear calculations, as expected. The last step is then to correct for any global effects that were not included before. In this case, the scale. The scale is now corrected for globally, but each column that already had a correction, is now overcorrected.



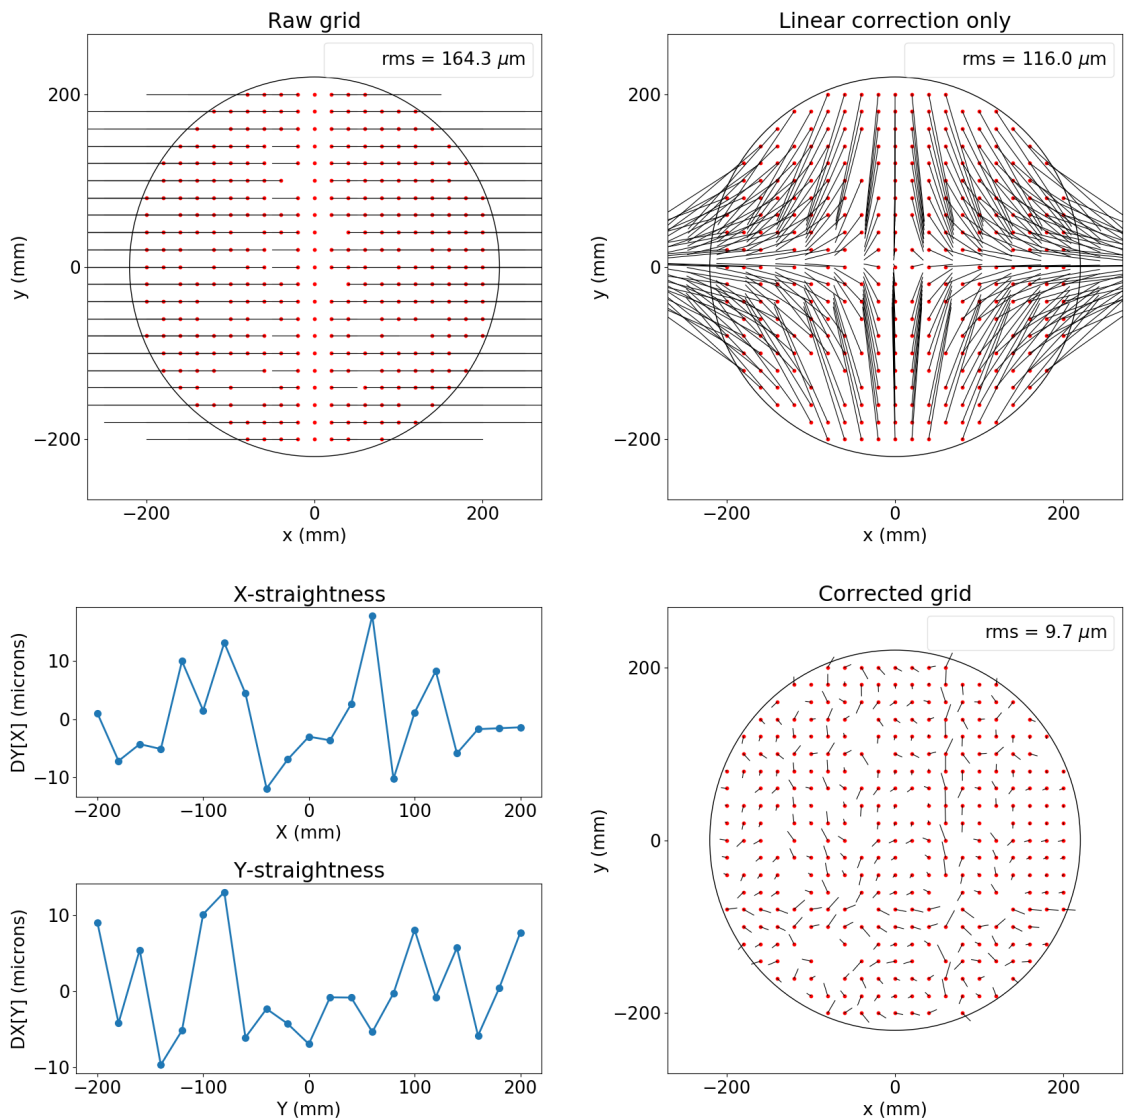


Figure 4.7: Autofib-2 algorithm applied to a grid where 10% of the points are obscured. *Top left*: simulated grid with a large scale change for x only, for illustrative purposes. *Top right*: the grid after going through the solid body rotation. *Bottom left*: the straightness arrays as calculated by the software. As there is only a scale change, we would expect flat lines on zero. This does not happen as the grid points do not cancel each other out anymore. *Bottom right*: the residuals after the full calculation. There is still a clear structure visible, in rows and columns around the ‘missing’ points, for example in the column on $x = 60$ mm.



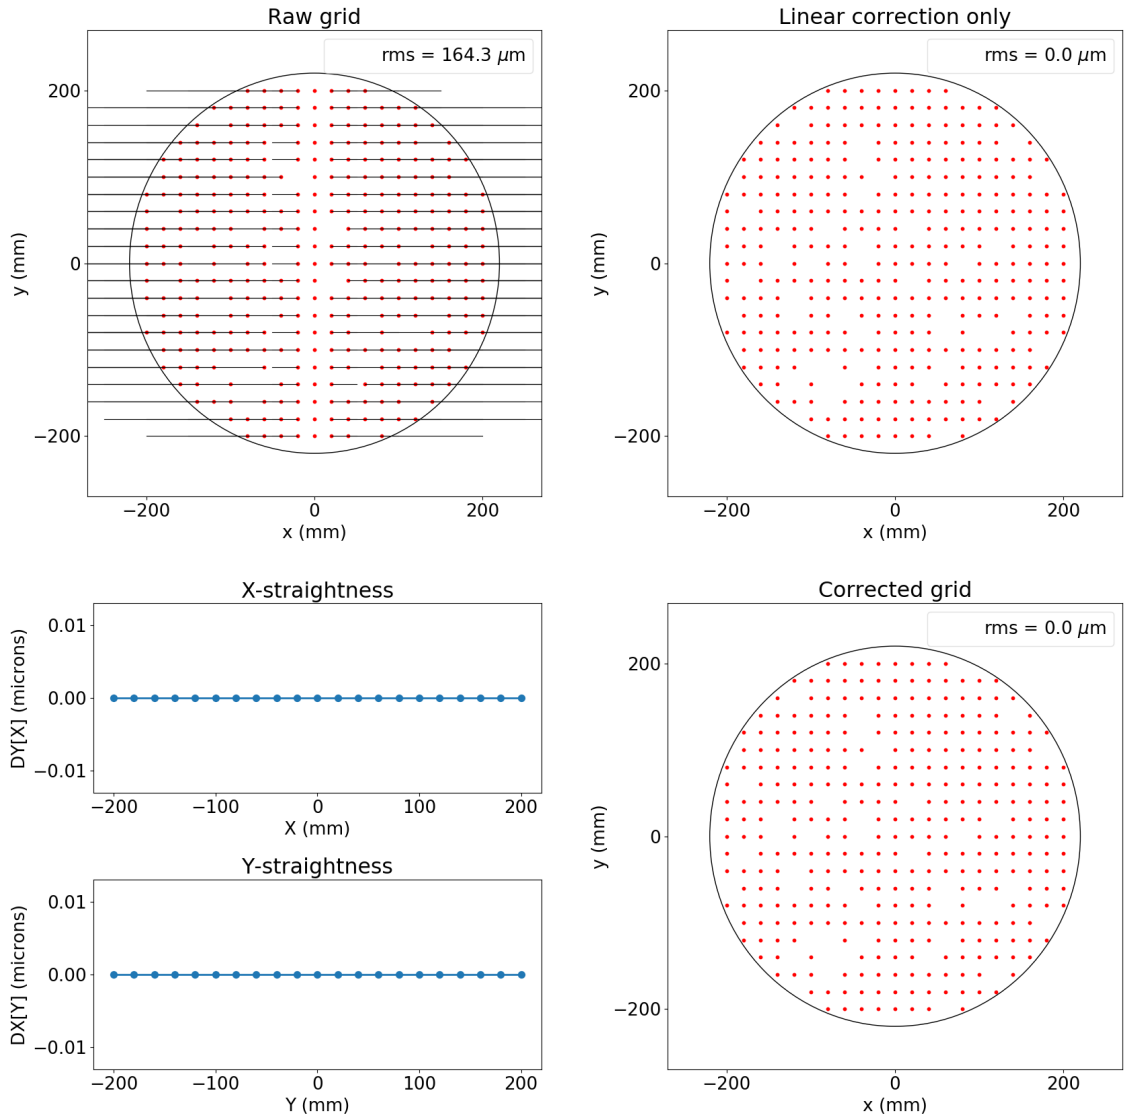


Figure 4.8: **GridKit** applied to a grid where 10% of the points are obscured. Compare this with figure 4.7, where first only a 4-parameter fit is done. *Top left*: simulated grid with a large scale change for x only, for illustrative purposes. *Top right*: the grid after going through the full 6-parameter fit, fitting for the scale change in the process. No residuals are left. *Bottom left*: the straightness arrays as calculated by the software. Because the scale change has been corrected for, the arrays are zero. *Bottom right*: the fully corrected grid has no residuals at all, as expected.



This cannot be changed at this point in the calculations, so the only conclusion from this is that the Autofib-2 algorithm cannot be used for WEAVE as-is, with first the 4-parameter fit before the linearity correction.

If the scale is not changed at all before the straightness calculation, the results for non-straightness are much the same, only just for one axis. This shows that the scale has to be corrected for completely before the straightness calculation, if any points are obscured. From this it follows that a 6-parameter fit has to be done before trying to fit for any non-linearity.

The results of the 6-parameter fit, as implemented in the `GridKit` software, are shown in figure 4.8. With the 6-parameter fit, the scale is taken out before the non-straightness of the axes is calculated. There is now no imbalance for the points that do not have their mirror points, so the straightness arrays are flat as expected. As the only simulated component was a scale change in x , the residuals are exactly zero.

4.5 Grid simulations

The influences of all the different errors described in section 4.2 can be investigated with a software model. Those sources of error can be simulated by introducing small Gaussian errors into a regular (perfect) grid of points and then adding offsets, scale effects, rotation, and non-perpendicularity. The waviness of the axes can be approximated by adding sine waves with certain amplitude and shift, and the flexure due to gravity is modelled as a centre-loaded simply-supported beam. I have simulated all this in `GridKit`, which I also used to make various plots in this chapter.

When simulating these sources of error, care has to be taken to first add the effects of the robots (the straightness effects), and only then of the plates. This matters, because correcting for these effects is non-commutative, and the linear effects are corrected for first. It also makes sense physically, because the gantries already have



wobbles independently of whether they measure a grid.

After introducing all these effects, the objective is to remove them as best as possible, and recover the original inputs to check for accuracy, with the caveat that rotation and non-perpendicularity will be over- or underestimated in the presence of non-linear effects. For offset, rotation, scaling and non-perpendicularity of the axes this is fairly easy, as these are linear transformations and every point in the grid is affected equally. For this the `slalib` library is still very useful. There are now Python wrappers available for these routines⁸, so `GridKit` is written in Python instead of in Fortran.

As the non-straightness of the axes and the small random errors are non-linear effects, the routines in the `slalib` library cannot remove those. More importantly, the routines will try to minimise the non-linear effects as well, as seen in section 4.4.1. Therefore the software has to try to decouple these effects, by first removing the linear influences, and then to eliminate the non-linear effects as best as possible with other routines.

4.5.1 Algorithm

With the discussion in the previous sections in mind, the most important step to change, is to change the 4-parameter fit to a 6-parameter fit as the first step. The solid body rotation introduces more problems than it would potentially solve. An obvious next step to simplify the algorithm is to apply the straightness fits to the partially corrected grid immediately. Indeed, as the arrays are calculated to produce the minimum RMS for that particular set of errors after the linear fit is already taken out, this combination of linear and non-linear fit is the best solution. This eliminates the need for an extra copy of the measured grid, and the mental juggling to know which grid is in the process of being modified.

⁸Thanks goes to Scott Ransom of the `pySLALIB` package ([Ransom, 2010](#)).



Moreover, the order in which the corrections are done, is now reversed. In both algorithms, the straightness correction is calculated after a rotation is done, and now the straightness correction is actually done on the linearly corrected grid, instead of on the uncorrected grid. The effect is tiny for the expected rotations for WEAVE, and was smaller than the random errors in Autofib-2, but it is now more physically correct.

If one wants to be very pedantic, one could do another 6-parameter fit after the straightness correction. For some cases of non-linearity the linear correction introduces a small scale factor to lower the overall error, which cannot be corrected for by the non-linear correction afterwards, and becomes an over-correction. The solution is a local minimum, but would be in a lower local minimum if the scale factor were taken out again by another linear fit. However, as the results are more than satisfactory (remaining error of about 10 nm) for over ten times the expected values of the non-linearity, I have decided to keep the simulation simple instead of making it slightly more accurate by adding another step, for something that will be lost in measurement error anyway.

Including these last deliberations, the steps for the algorithm for WEAVE are outlined below, and depicted in figure 4.9.

1. Start with a grid that has both linear and non-linear transformations applied to it, in the case of a simulation, or a grid that is obtained by measuring a physical grid. This is the measured grid.
2. Also have an expected grid: a grid without the errors introduced by the robot systems. This can be a simulated grid without the transformations, or the reference grid measured with an independent measuring system.
3. Use the `slalib` routines with a 6-parameter fit to calculate the offset, scaling, rotation, and non-perpendicularity between the two grids. Most of the non-



straightness of the axes is left over.

4. Correct the measured grid using these linear transformations.
5. Calculate the non-straightness of the axes by averaging the remaining offset in each row and column in the partially corrected grid from point 4, and save these as ‘non-straightness arrays’.
6. Remove these arrays from each point in the corresponding row or column of that partially corrected grid. This is now the fully corrected grid.

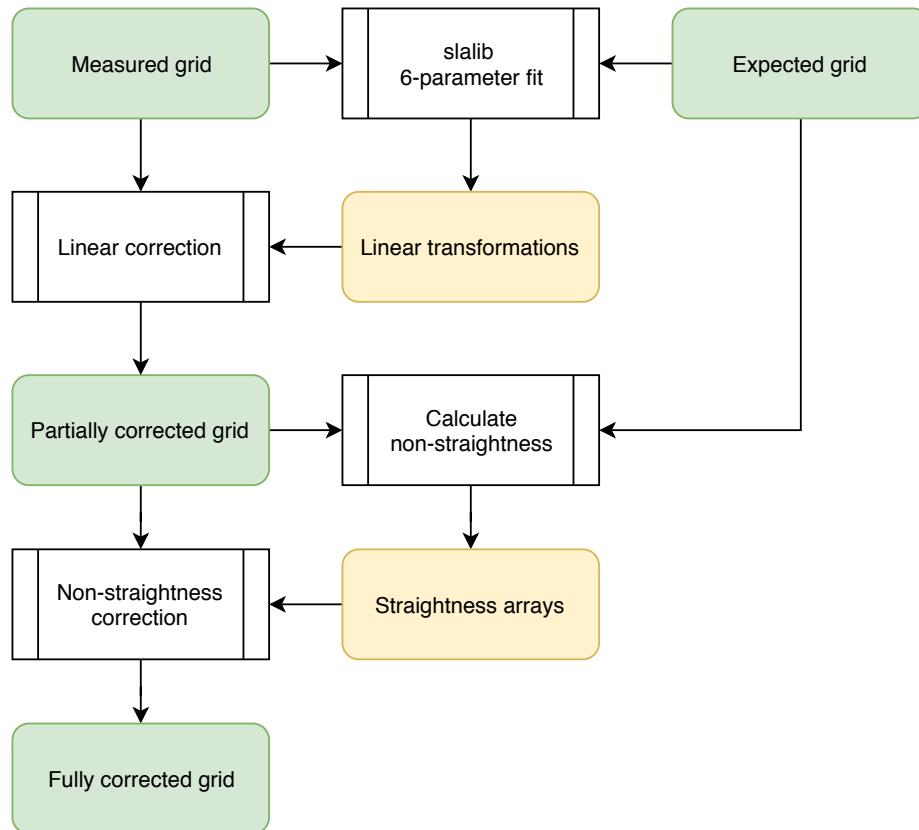


Figure 4.9: GridKit flowchart: from a measured and an expected grid to a fully corrected grid. The boxes with rounded corners contain data, the boxes with vertical lines contain calculations.

For completeness the Autofib-2 algorithm is also included in the software, to be able to compare the two algorithms, as shown in section 4.4.



The non-straightness arrays are important outside of this immediate calculation because they provide corrections to the grid that will not change from measurement to measurement, even with the instrument at different zenith angles. See also figure 4.3 again. If the primary measurements to find the non-straightness are done with the instrument fully vertical, the only component to these arrays is the intrinsic waviness. The sag because of gravity is perpendicular to the measurements in the (x, y) -plane, and does therefore not influence the measurement. This is useful, because it can then be corrected for before any other calculations are made. Furthermore, any change in the non-straightness after that has to come from gravity pulling on the axes.

After initial calibration off the telescope, the expected grid therefore includes corrections that will not change from measurement to measurement. The remaining residuals are smaller, and on-telescope calibration becomes easier, because the effects of the linear and non-linear errors are further decoupled.

In these simulations the effect of tumbling is not simulated. After reconfiguring a grid, the tumbler will tumble the plate into the observing position. This will probably give an overall offset, but this can be adjusted for by re-pointing the telescope, if necessary.

4.5.2 Separate contributions

Even though some of the effects of linear and non-linear transformations were already discussed in previous sections, it is useful to look at how each effect contributes to the distortion of a grid. Figure 4.10 shows some examples of linear distortions. The plotted radius of the field plate is 220 mm, which is larger than the nominal 200 mm or the designed 205 mm, but shows the effects at the edges well. Each plot shows the expected grid with red dots, and the difference to the measured, uncorrected grid with black lines. The magnification on these lines is 1000, so they are still visible. It also means that a length of for example 20 mm on the plot is translated to 20 μm for



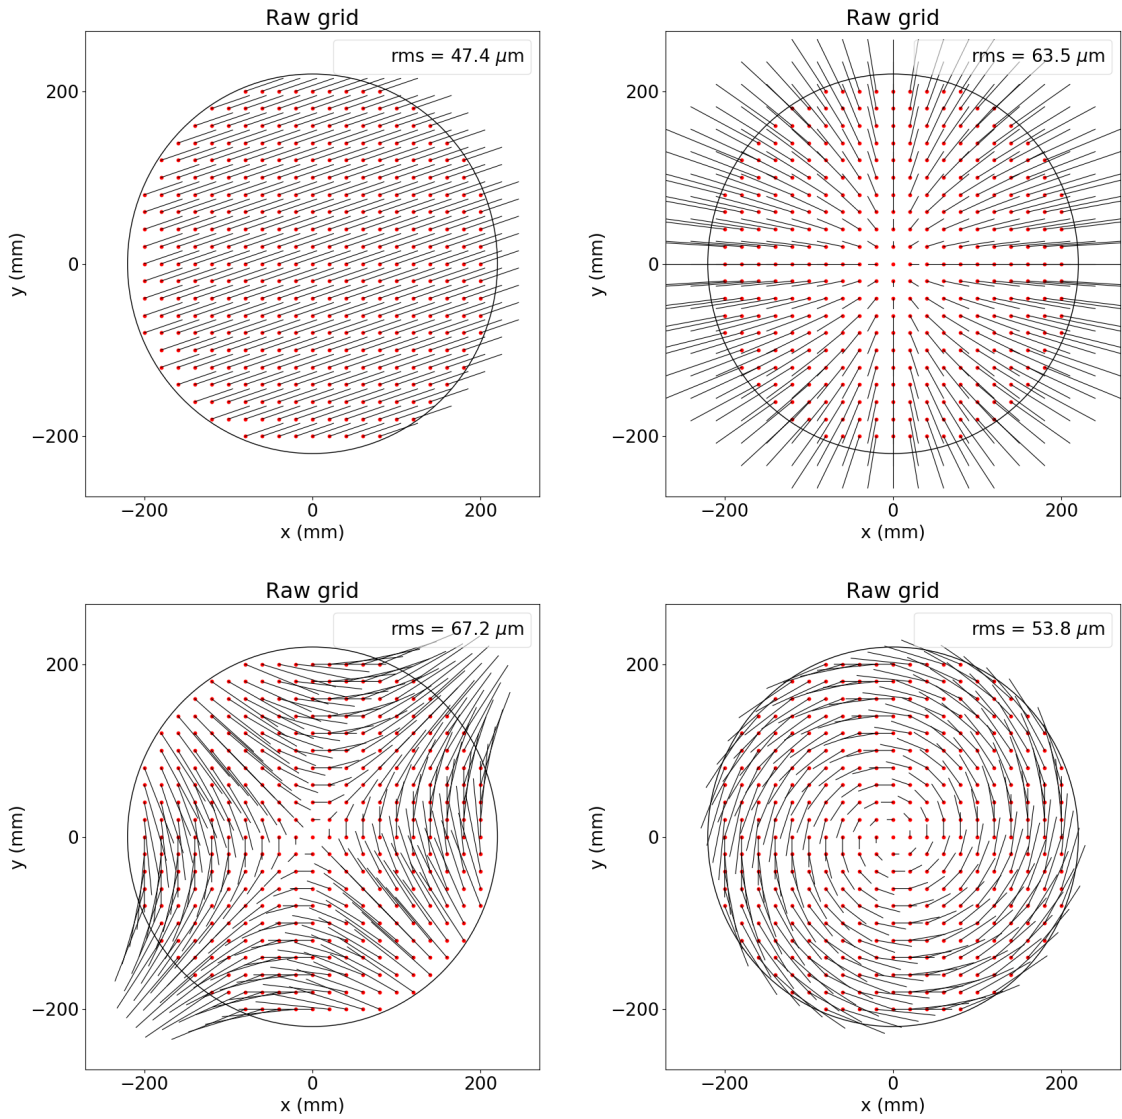


Figure 4.10: The linear transformations possible on a grid of points. *Top left:* an offset of 45 μm in x and 15 μm in y . *Top right:* a scale factor of 1.0005 for x and 1.0003 for y . *Bottom left:* a non-perpendicularity of the axes of 0.05° , which shears the grid. *Bottom right:* a rotation of -0.02° around the origin. In this software, a positive rotation is clockwise.



the errors. Each plot also displays the total RMS-value.

The plots show an offset, a scale change, a non-perpendicularity and a rotation. Each of these contributions is easily removed by the procedure outlined in section 4.5.1, and 0.0 μm total RMS is left over (to be precise, each calculation has less than 10^{-13} mm RMS left over). `GridKit` also recovers the parameters to counter these effects. For the offset, the recovered values are -45 and -15 μm , for the scale they are 0.9995 and 0.9997, the non-perpendicularity is -8.727×10^{-4} rad, which is -0.05° , and the rotation is 3.491×10^{-4} rad or 0.02° .

The values for the offset and the non-perpendicularity gives the additive inverse from what was put into the simulation, and the scale is the reciprocal, as expected. The rotation is also the opposite value, but the rotation direction is the other way round than expected. In `slalib`, a positive angle is defined as going clockwise, and a negative angle is anti-clockwise. This is important to remember if something on the positioner has to be physically rotated to lessen the rotation error.

Pure non-linearities will not be corrected for solely by the non-linear routines, because the linear routine will try to find a fit first. As we saw previously, this results in a small non-perpendicularity and rotation, and has to be re-corrected for in the straightness fits. For the case of the four sine waves in figure 4.6, repeated top left in figure 4.11, the rotation is -0.0015° , and the non-perpendicularity is 0.003° . The RMS left over after the all the corrections is 0.05 μm , so the software is not doing too badly there.

An order of magnitude decrease in the amplitude of the sine waves (more towards real-world values) results in an order of magnitude decrease in the rotation and non-perpendicularity as well. This can also be seen in figure 4.11, where the top row's straightness plots look exactly the same, just with a tenfold smaller y -scale on the right. From this follows that the rotation is ten times smaller, and thus the non-perpendicularity is ten times smaller too. The parameter values confirm this, as the



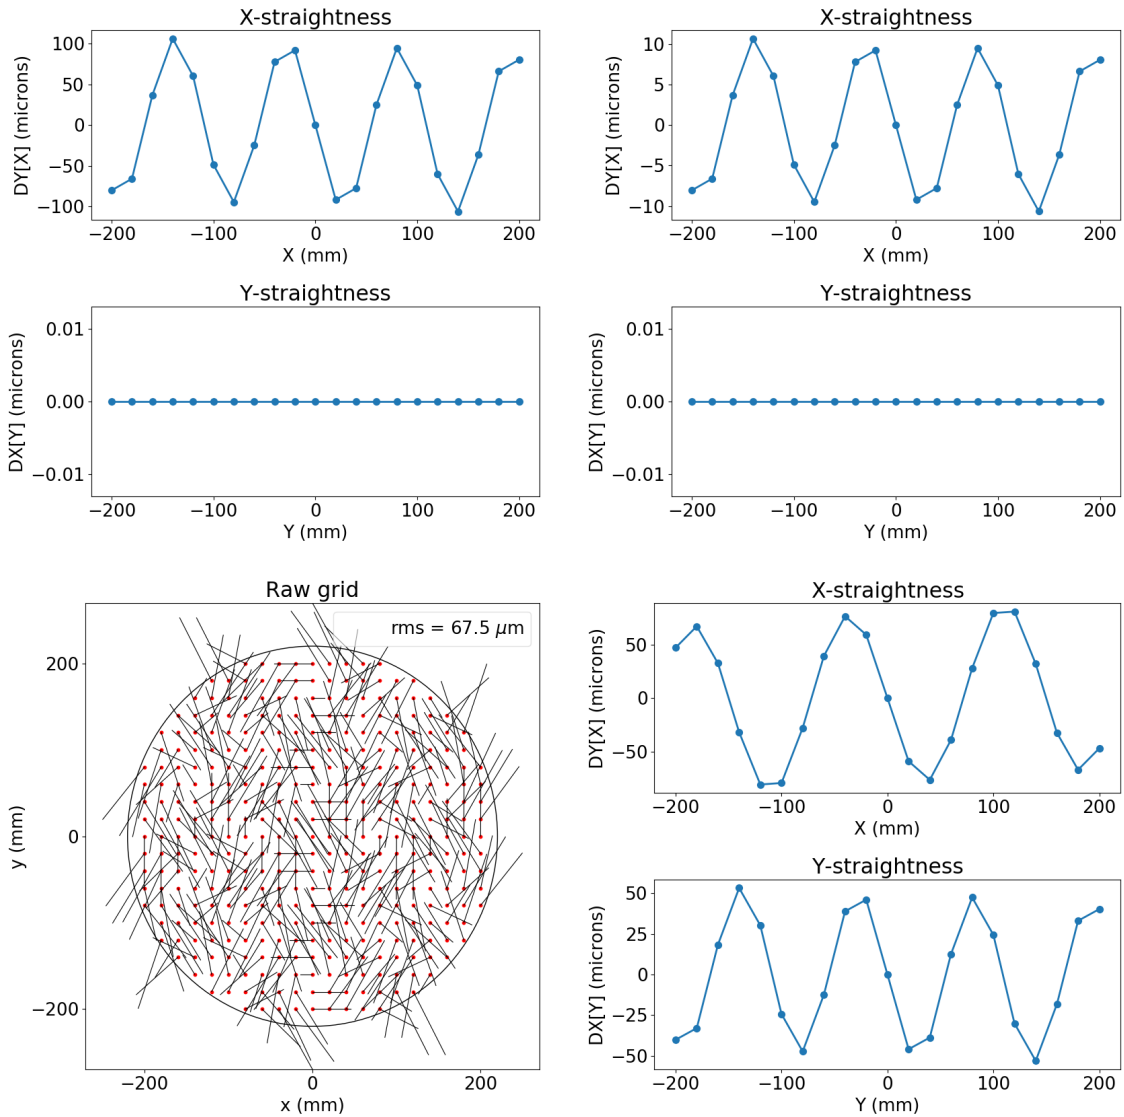



Figure 4.11: Non-linearity of the axes, focusing on the amplitude. *Top left*: four sine waves with an amplitude of $100\ \mu\text{m}$. *Top right*: four sine waves with an amplitude of $10\ \mu\text{m}$. The ‘overcompensation’ in non-perpendicularity and rotation is ten times smaller too. *Bottom left*: an exaggerated example of sine waves in both directions. *Bottom right*: the corresponding straightness arrays to the grid on the left, after -0.0023° non-perpendicularity and 0.0027° rotation have been removed. After the full correction $0.1\ \mu\text{m}$ RMS remains.



software now returns $-1.5 \times 10^{-4}^\circ$ and $3 \times 10^{-4}^\circ$. The RMS is now of the order of 10^{-4} μm .


For non-straightness in both directions, the procedure remains mostly the same. There is some rotation and non-perpendicularity (and a tiny bit of scale change), but the ratio between them changes. The bottom row in figure 4.11 shows a random example of this.

Lastly, it is important to show the flexure due to gravity. Each gantry is best approximated by simply-supported beams: the ends of the gantries are allowed to rotate, but not deflect. Therefore the flexure is modelled as a centre-loaded simply-supported beam. It transforms a flat line into all positive or all negative values compared to that line. Because of this, `GridKit` will always perform a shift first, so the overall error is minimised. This makes the straightness arrays look slightly weird, because the shift is then incorporated into the arrays again. This is illustrated in figure 4.12, where the middle grid shows the vertical shift to lower the overall error, and in the array on the right the curve extends past $y = 0$.

Because the gantries extend out far beyond the grid (each gantry is 1100 mm long), there is a considerable deflection at the outer grid points. Therefore most of the error is removed by performing a shift in the right direction. The two robots share the same X -gantries. This means that if one robot moves, both of them will experience a change in deflection. The centre loading is strictly only correct for when the robots are in the centre of their gantries, and not moving, but as by far most of the deflection is a shift, it can be taken as a good approximation for the overall deflection.

4.5.3 WEAVE simulation

To accurately predict the errors in positioning due to the aforementioned effects, we can do a simulation with some good estimations for the different influences. These



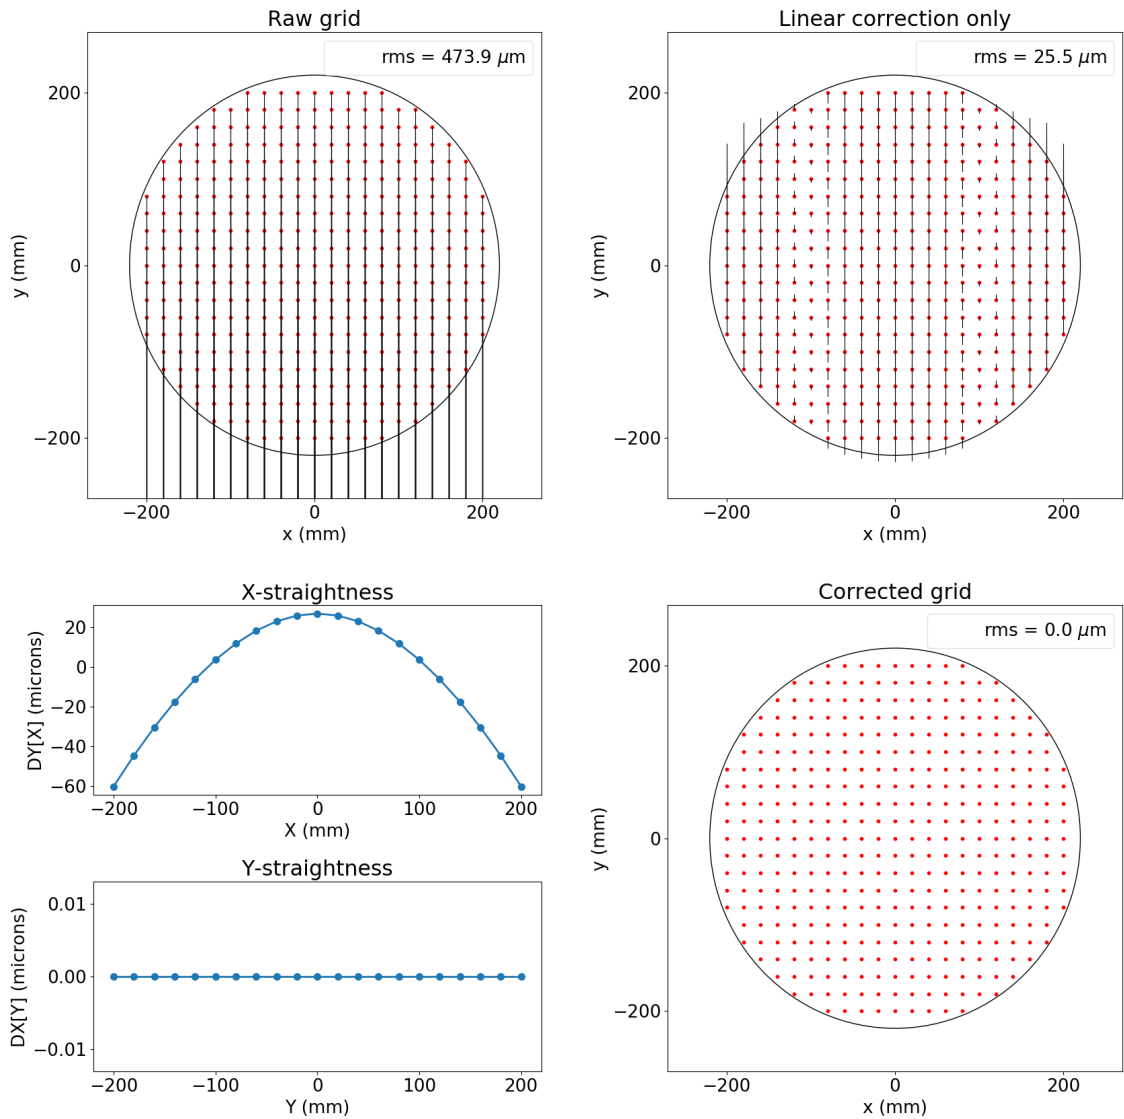


Figure 4.12: A simulated centre-loaded simply-supported beam, with a maximum deflection of -0.5 mm at $x = 0$, and zero deflection on the ends of the beam at $x = -550$ and $x = 550$ mm. This particular deflection is chosen to show the behaviour, it is not close to the expected value at all. *Top left:* the measured grid, without any corrections. *Top right:* after the linear corrections: a vertical shift upwards of 0.473 mm. *Bottom left:* the straightness arrays to minimise the errors that were left. *Bottom right:* the residuals are as good as gone. The remaining RMS is 0.006 μm .



come from known uncertainties in the robot gantries and the machining equipment, or from previous experience with the Autofib-2 and 2dF instruments. Typical expected values for WEAVE are given in table 4.1.

| Source of error | Value |
|---|--------------------------------------|
| Independent measurement error | 2 μm |
| Measurement error (RMS) ^a | 4 μm |
| Offset x | 100 μm |
| Offset y | 100 μm |
| Scale error | 1.0005 to 0.99995 |
| Non-perpendicularity ^b | 0.0005° |
| Rotation | 0.005° |
| Non-straightness X -gantry ^b | 5 μm amplitude, 3 periods |
| Non-straightness Y -gantry ^b | 5 μm amplitude, 3 periods |
| Maximum deflection X -gantry ^c | -15 μm |
| Maximum deflection Y -gantry | -5 μm |

Table 4.1: Expected values for the WEAVE robots, from [Lewis et al. \(2017\)](#). The ‘independent measurement error’ is the expected error due to measuring the field plates with a different measurement system than the robots, e.g. a coordinate-measuring machine or even the machine that produced the grid marks, if it is precise enough. ^aNoted in section 4.2. Included in this value is a noise term of 1 μm . ^bFrom Autofib-2 measurements. ^cAt zenith distance = 60°.

FEA modelling shows that the difference in displacement due to gravity for both of the positioning robots at the edge, versus both of them in the centre of travel, is about 3 μm . This is only for the displacement in y of the X -gantry, and the displacement in x of the Y -gantry. This is the maximum extra deflection that happens due to the robots moving around while repositioning the fibres. This cannot be corrected for, but the effect can be modelled, by changing the maximum deflection for a second grid, whilst leaving all the other effects (bar the random errors) exactly the same. This grid’s deflection value for the X -gantry then becomes 18 μm . Note that this approach models the maximum x - and y -deflection simultaneously, which will not occur in practice.

Figure 4.13 shows the simulation with the values from table 4.1, an x -scale of 1.0005, and a y -scale of 0.9998. The measurement error of 4 μm together with the



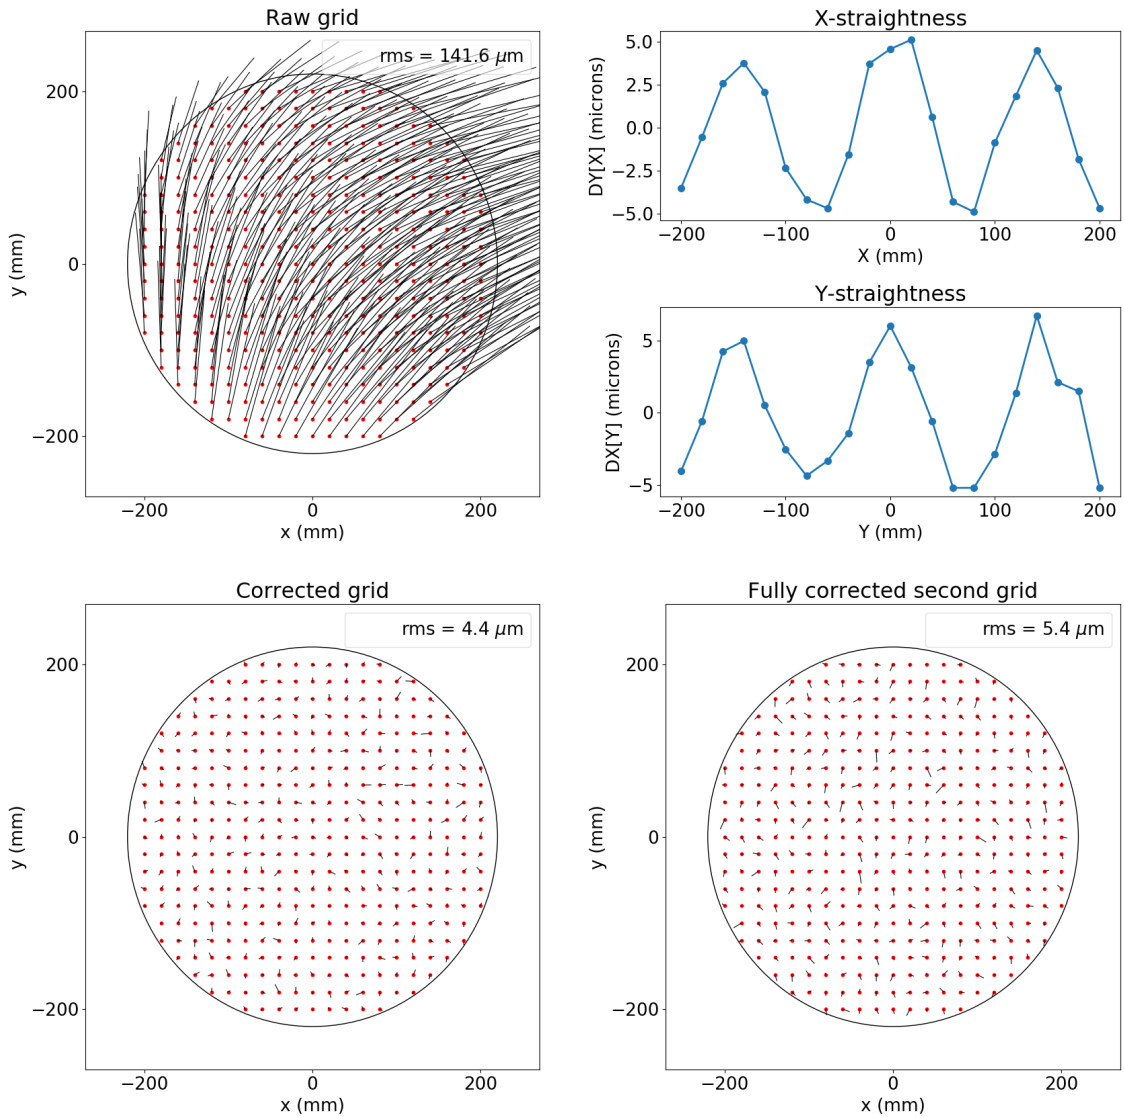


Figure 4.13: *Top left:* the uncorrected grid with all the simulated sources of error. *Top right:* the straightness arrays calculated after the 6-parameter linear fit is removed. *Bottom left:* the same grid now fully corrected. Randomness cannot be taken out. *Bottom right:* a second grid with a different x -flexure and randomness fitted with the exact same solution as the grid on the left. The RMS is larger, but still dominated by the measurement error.



independent measurement error of $2\ \mu\text{m}$ are the main contributing factors to the total error of $4.4\ \mu\text{m RMS}$ ($\sqrt{4^2 + 2^2} \approx 4.5\ \text{mm}$). With the extra (uncorrected for) flexure in the second plot, the overall error grows to $5.4\ \mu\text{m}$. This flexure is mostly a shift in y and could quite easily be corrected for if this were a static situation. This is not the case though, because the robots are moving around all the time and so the amount of flexure changes too.




Figure 4.14: Compared to figure 4.13, the measurement error is reduced from $5\ \mu\text{m}$ to $2\ \mu\text{m}$ (RMS). *On the left:* the fully corrected grid. Randomness cannot be taken out. *On the right:* a second grid with a different x -flexure and randomness fitted with the exact same solution as the grid on the left. The residuals now start to be dominated by the flexure difference.

If the robot repeatability improves to an RMS-value of $2\ \mu\text{m}$, the corrected grid and a second grid with the same solution (but a different value for the flexure) are the ones in figure 4.14. The measurement errors still dominate the residuals: $\sqrt{2^2 + 2^2} \approx 2.8\ \text{mm}$. For the second grid the error changes to $4.1\ \mu\text{m}$, and if the errors were not so small, a trend would start to be visible in the residuals. This is because the maximum flexure difference of $3\ \mu\text{m}$ is almost the same over the whole grid, and this shift is slightly larger than the random error.

For a set of measurements without any random errors, the software performs



| Parameter | Expected output | 4 μm error | 2 μm error | no error |
|-------------------------------------|-----------------|-----------------------|-----------------------|-----------|
| x -offset (mm) | -0.1 | -0.095239 | -0.095118 | -0.095100 |
| y -offset (mm) | -0.1 | -0.085508 | -0.085519 | -0.085556 |
| x -scale | 0.999500 | 0.999501 | 0.999500 | 0.999500 |
| y -scale | 1.000200 | 1.000200 | 1.000200 | 1.000200 |
| shear ($^\circ$) | -0.0005 | -0.000431 | -0.000482 | -0.000500 |
| rotation ($^\circ$) | -0.005 | -0.005033 | -0.005026 | -0.005000 |
| Residual (RMS in μm) | | 4.4 | 2.8 | 0.0006 |

Table 4.2: Overview of the linear parameters included in the simulations. The input values are the same as the ones in the list above. The output values are grouped by the simulated measurement error, with a fixed independent measurement error of 2 μm (and no error for the last column). The remaining residuals show that the simulation performs extremely well. The differences in offset are due to the flexure.

exceptionally well. `GridKit` is able to get back the input parameters to much better than can be easily shown here, apart from the offset. This has a good reason too: almost all of the flexure is first corrected for by a shift. A negative deflection causes a positive shift, so the overall shifts are smaller than expected. The $-15\ \mu\text{m}$ deflection in the X -gantry is a deflection in the $-y$ -direction, which is why the y -offset is just under $-0.086\ \text{mm}$ instead of $-0.1\ \text{mm}$. The same happens for x , due to the deflection in the Y -gantry. This is summarised in table 4.2.

4.5.4 Simulation with only part of the grid available

The software performs extremely well when all of the grid points are visible. As stated before, this will not usually be the case. Up to 65% of the points can be obscured by the arrangement of fibres on top of the field plate. How well does the software perform with only 35% available?

In the next simulation (figure 4.15) a random selection of grid points is made unavailable for measurement, and thus for the software. This does not cover all of the possible permutations of availability, and does not take into account that some points are more likely to be obscured than others, but it does give a clear indication



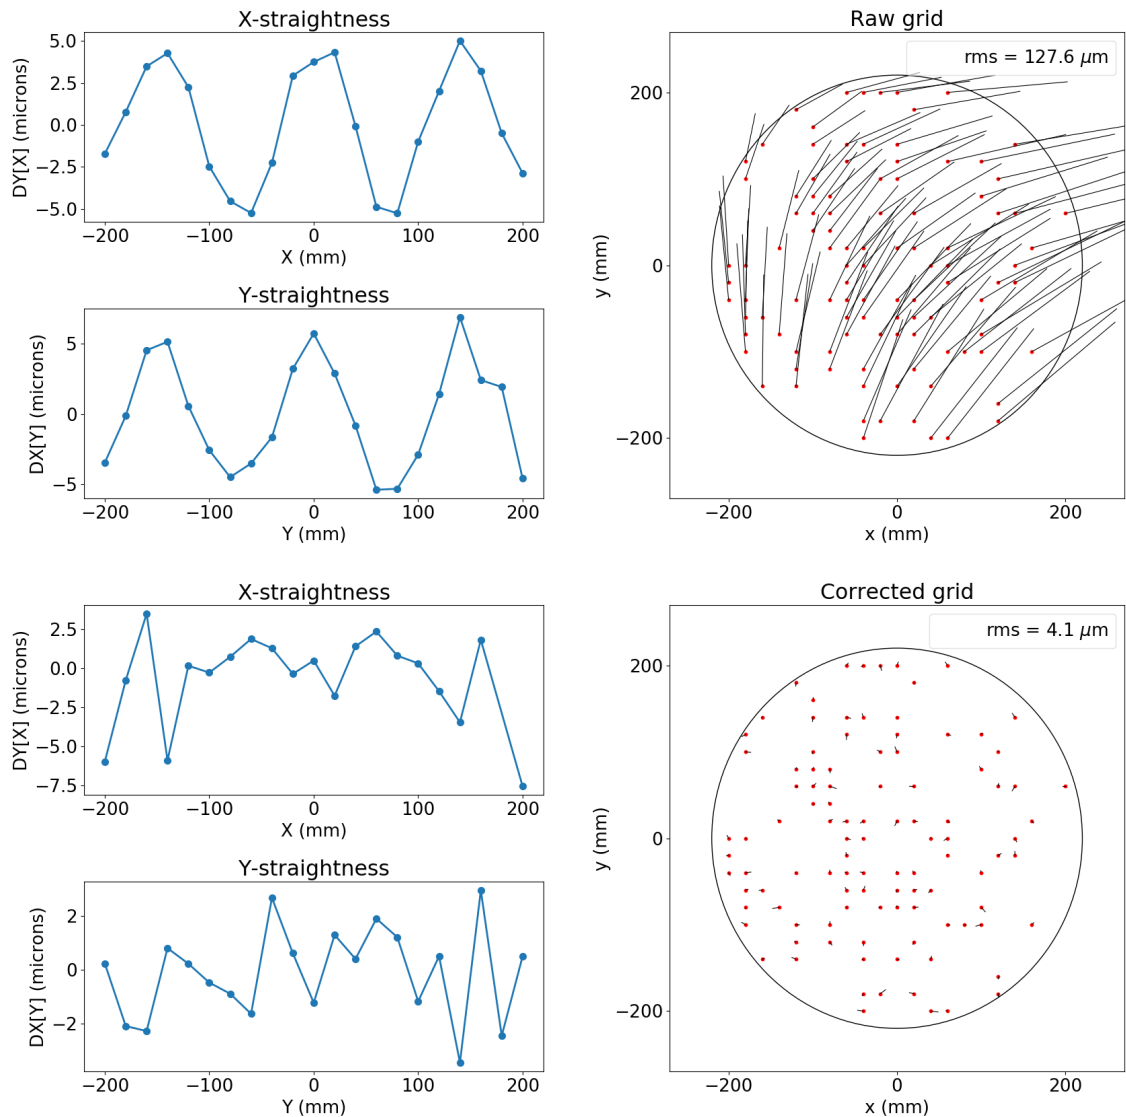


Figure 4.15: *Top left*: the straightness arrays for a simulated measurement with the positioner pointing at zenith; only the (static) installation effects are included. *Top right*: a grid with only 28% of the grid points available for measurement. In this plot the non-straightness due to installation effects are already removed. *Bottom left*: the gravity effects are somewhat visible in the remaining straightness plots, but particularly the Y -straightness is affected by the random errors. *Bottom right*: the grid after all the corrections. Only $4.1\ \mu\text{m}$ RMS remains.



of how well the software works. Only 28% is still available in this simulation, to test whether the software still performs below the minimum 35%, if for whatever reason even fewer points are available.


Compared to the previous section, `GridKit` takes one extra step in determining the input parameters and correcting for them. The non-straightness in the axes due to installation effects is assumed to remain the same over the positioner's lifetime (this can be checked with a measurement at zenith every month or so), so before any measurement goes through the analysis, this non-straightness can be subtracted from the measurement. After that, the procedure stays the same.

Because of this extra step, the only effects that remain in the current measurement's straightness arrays are the deflection due to gravity and possible structure due to the interplay between the random errors and the amount of available grid points in a row or column.

The results show exactly this. The linear input values are slightly different ($x = 0.1$ mm, $y = 0.09$ mm, x -scale = 1.0005, y -scale = 0.9997, shear = 0.0003°, rotation = 0.0055°) to the grid providing the overall non-straightness information (the input values from the previous section). This models the changes due to temperature effects, wear, etc. The measurement error is again 4 μm. The remaining non-straightness shows some deflection in the X -straightness, compounded with the random errors, and the Y -straightness is mostly noise.

As before, the simulation shows only one particular set of random errors and random available points, but this particular set has not been selected over any other possible set. Simulations with other random errors and grid points show similar results.

Table 4.3 summarises the parameters found by `GridKit` with the above simulation. It also summarises the situation where the straightness information has not been removed as a first step. The main difference is in the amount of shear that is applied



| Parameter | Expected output | 28% of points, installation effects removed | 28% of points, no extra step |
|-------------------------------------|-----------------|---|------------------------------|
| x -offset (mm) | -0.1 | -0.094614 | -0.094979 |
| y -offset (mm) | -0.09 | -0.075159 | -0.075130 |
| x -scale | 0.999500 | 0.999498 | 0.999497 |
| y -scale | 1.000300 | 1.000293 | 1.000292 |
| shear (°) | -0.0003 | -0.000496 | -0.000162 |
| rotation (°) | -0.0055 | -0.005645 | -0.005750 |
| Residual (RMS in μm) | | 4.1 | 4.1 |

Table 4.3: A comparison between a calculation with the non-straightness effects taken out before the other corrections, and one without this step. In this case the same random seed was used; the ‘randomness’ is exactly the same in both simulations, to only distinguish between the effects of having prior knowledge of the straightness or not. The main difference is in the amount of shear that is applied, and this is a known effect of the way the software works. The input values are slightly different from table 4.2, but as the input error is 4 μm , it is comparable to that column in table 4.2.

to correct for non-straightness, which is an effect that was seen before in the other sections. The software still recovers the parameters surprisingly well, even without the extra step of removing the non-straightness. The conclusion is therefore that as long as the grid mark measurements are highly repeatable⁹, the software will not only find a solution to minimise the overall error, but will also find the proper solution to the actual state of the instrument.

4.6 Grid measurements

In each field plate, a grid is machined with a dot diameter of 0.5 mm and 20 mm spacing. The CNC machine used for this has an absolute accuracy of 5 μm . To measure this grid, each robot system has a camera mounted above the jaws on the gripper, to

⁹The measured centroids of each grid point have to stay the same under different positioner orientations, possibly changing light conditions, a slightly different position of the robot above the grid mark, etc. This is subject to further study, and conclusive results can only be gotten from the final field plates, which are made from invar instead of aluminium.



be able to look at the field plate directly below. For the measurements described in this section, the robot is sent to each location where a grid marker is expected, and then takes an image. These expected locations are the nominally expected locations: $(0, 0)$ with increments of 20 mm in both x and y until the full grid is sampled. This does not take into account any prior knowledge of the real locations of the grid marks.




Figure 4.16: One of the grid marks imaged by Nona’s camera. *On the left:* the painted mark imaged as-is, displayed with the SAOImage `ds9` software. *On the right:* the image after inversion and applying a square-root scale, imaged with `matplotlib`. This image (light dot on a darker background) is now ready for the biscuit cutter software.

The images are saved as fits files, after which they can be processed further¹⁰. An example of an image is shown in figure 4.16. Centroiding software, in this case James Gilbert’s biscuit cutter software specifically made for backlit fibres (Gilbert, 2016, section 3.6)¹¹, but also useable for grid marks on a plate, finds the centre of the mark in pixel coordinates. Together with the plate scale (more on that in section 5.2.1) the physical location in (x, y) -coordinates is calculated. The robot coordinates for each image are taken as the expected location for each grid point. These grids of expected and measured/calculated locations are then fed into the grid software.

¹⁰Fits file handling is done with `astropy.io`, from <https://www.astropy.org>, a community-developed core Python package for Astronomy (Astropy Collaboration, 2013, 2018).

¹¹The biscuit cutter method lives in a software package called `fibmeasure` (Gilbert, 2016), which makes use of the package `image-registration` (Ginsburg, 2012).



4.6.1 Offset issues

The grids in the figures below are not centred around $(0,0)$ because the grid marks are imaged in the upper focus. This focus is offset by 30 mm from the lower focus. The motor encoders are set in such a way that when the θ -axis rotation centre (which is located in the lower focus) coincides with the centre of the field plate, or $(0,0)$ in field plate coordinates, the motor encoders are also set to $(0,0)$ in motor coordinates. With $\theta = 0^\circ$, the orientation in which the below data is taken, this offset happens to translate to the x -direction only. The robot is therefore sent to +30 mm in the x -direction, to be able to image each grid mark. The central mark now has coordinates $(30,0)$. This does not introduce any offset between the measured and expected grid, as the expected grid locations are taken to be the robot locations.

Another contribution does result in a few mm offset between the expected and measured grids. At the time of measuring the grid, the above method of setting the motor encoders to the robot's rotation axis had not happened in detail. There was still an offset of about 2.1 mm in x and 2.9 mm in y . Therefore the location of the grid mark is off by these values from the expected location. This offset is the same in each image, as long as the rotation of the robot does not change. After removing this offset by hand in between the centroiding stage and `GridKit`, none of this is visible in the figures below.

4.6.2 Centroiding issues

After preparing the data in the way described above, we can feed it into the biscuit cutter centroiding software and then into `GridKit`. Unfortunately, not every image is centroided as expected. This is clearly visible in figure 4.17. Not only are grid points missing, for which no centroid has been calculated at all, there are also values that are clearly wrong. The vectors are going in a completely different direction than the surrounding vectors are.



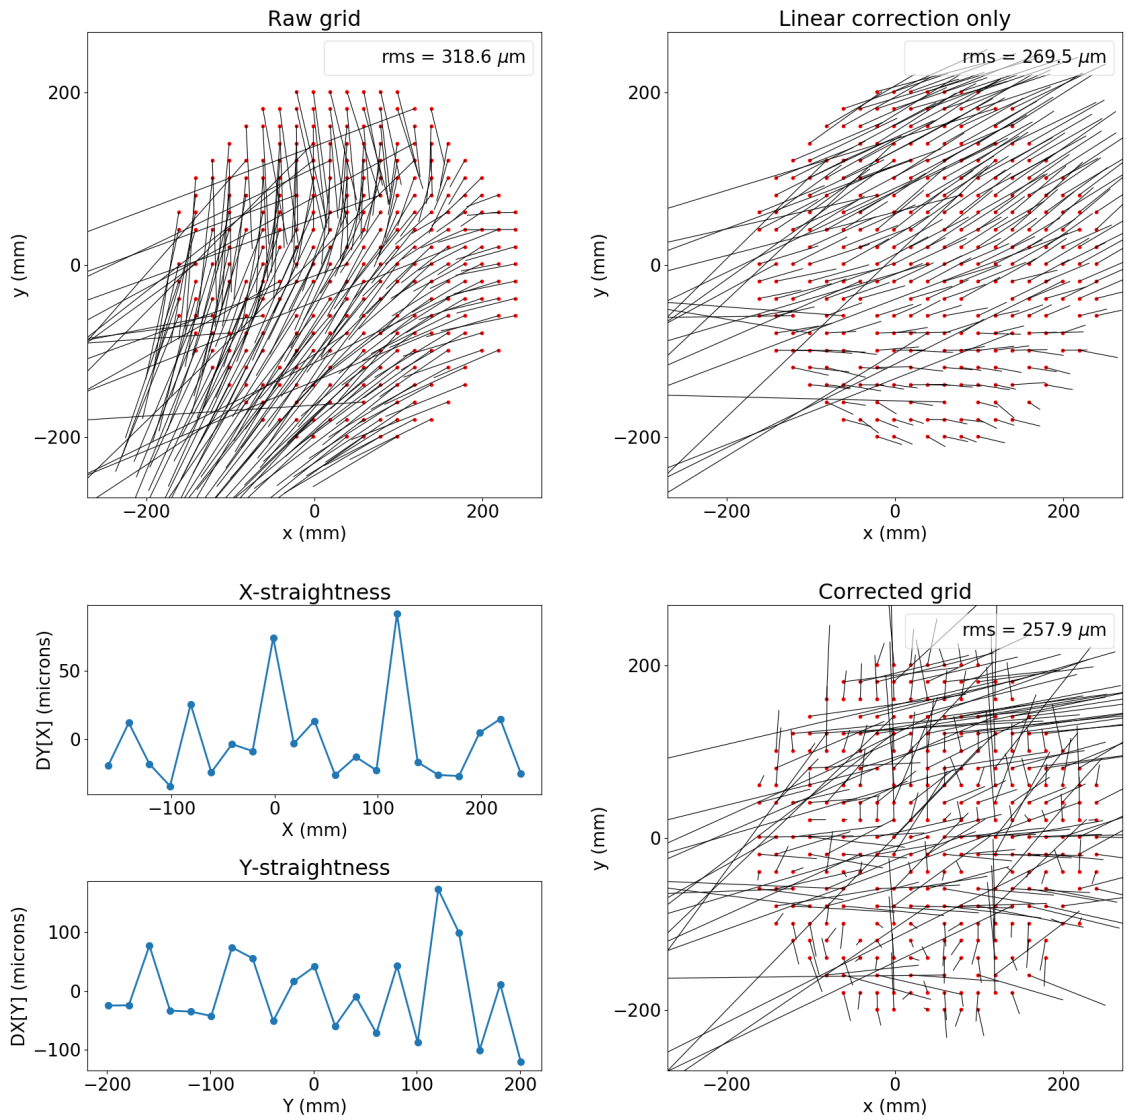


Figure 4.17: A measurement of the grid on the stand-in aluminium plate. All the grid marks have been painted and sanded. The grid is not in the centre due to the 30 mm offset between the upper and lower focus. It is also clear that the grid marks with unexpected centroiding results are influencing the results significantly, and therefore have to be removed.



This is only a problem because we have currently defined the expected grid as a regular grid with a set distance of 20 mm between the grid marks, which does not take into account the reality of finding the centre of these marks (see also footnote 9). For these grid marks, the found ‘centre’ does not actually have to be in the centre, as long as each subsequent measurement of a mark gives the same result. This information of where each grid mark’s expected ‘centre’ can be found, has to be incorporated into the expected grid in a next pass through the software.

For now, we can deal with the off-centre centroids in a different way. First we can ignore the points for which no centroids have been found, as this is essentially the same outcome as when a grid mark cannot be measured due to it being obscured by buttons and fibres on top of them. We saw in section 4.5.4 that missing out on about 14% of the points (48 of 349) should not pose any problems for `GridKit`.

There are two steps in the process where some images fail to get through to the next stage. The first step where the biscuit cutter software in `fibmeasure` sometimes fails is that it eternally gets stuck on trying to centroid an image. I have dealt with this by taking that image out of the list that is fed into `fibmeasure`. Another failure mode is where `fibmeasure` does manage to process the image, but returns `None` instead of the coordinates. In the measurements described here there were 10 files that completely locked up the computer, and 38 that returned no centroids. If an extra 14% of the 35% unobscured grid marks is unavailable (worst case scenario for grid mark unavailability), that leaves only 30% of the total number of marks. Looking at section 4.5.4, this will not pose a problem for `GridKit`, the next stage in the processing pipeline.

I have not found where exactly in `fibmeasure` the bug originates, because that would have meant trying to understand a completely unknown code base. For WEAVE on the telescope, the software provided by James Gilbert and me is being rewritten in C++, with more consideration given to edge cases. It is fair to assume that these



kinds of bugs will be solved during that rewriting process.


Which grid mark results have to be removed? The cut-off is always a bit artificial, but the selection can be done in two ways. The first one is to reject any centroid that falls outside a certain range of pixels in the image. This only works if the camera and grid are aligned properly (so there is no rotational misalignment between the two), but is very easy to implement. The second one is to calculate the distance between each expected and each measured coordinate and to reject any point above a certain distance. This only works if there is no major offset or other large-scale deformation of the grid.

4.6.3 Final result

To make figure 4.18, I have decided to implement the first selection method. Again, this only works because the grid on the field plate and the robot gantries align almost perfectly. If there were any rotational misalignment, the centroid would move around on the image and implementation of this method would become a lot harder. I have determined empirically that a box with dimensions of 60 px in x and 50 px in y around a common centre is a good start to eliminate the few worst points. The diameter of a grid mark is about 120 px, so this is a very generous box.

With this method, 20 extra grid marks are ignored. This leaves 281 out of a maximum of 349, or 81%. This is shown in figure 4.18, and now much more underlying structure can be calibrated out. The top left image shows that no vectors are going in the opposite way anymore.

Table 4.4 shows the parameters found by `GridKit` for the filtered measurements. The remaining RMS error is 9.9 μm . This looks like a large error, certainly if the maximum permissible overall error is only 8 μm , but is actually encouraging. I have already shown that the assumption that the expected grid is regular, is not correct. I expect that most of this 9.9 μm error is because the found centroids are not in fact



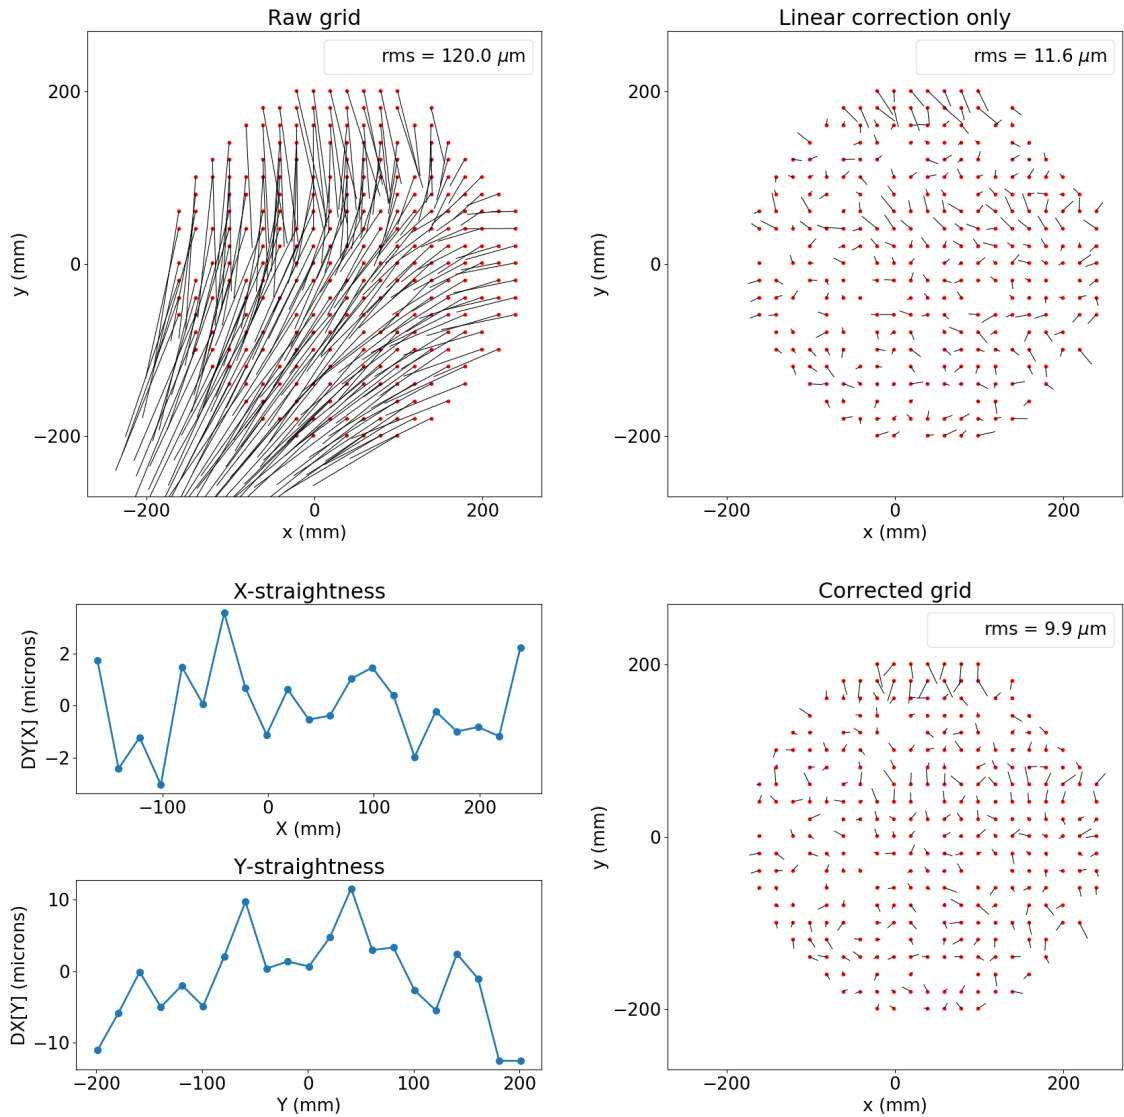



Figure 4.18: The same data as in figure 4.17, but with the 20 worst outliers ignored. The linear part of the routine gets rid of most of the error. The amplitude in the X-gantries is too small to exceed the noise levels, but there is a clear structure in the non-straightness in the Y-gantry, which is most likely due to the installation. These arrays are subsequently removed too, to remain with 9.9 μm RMS error.




| Parameter | Output |
|-------------------------------------|-----------|
| x -offset (mm) | 0.050148 |
| y -offset (mm) | 0.109889 |
| x -scale | 1.000036 |
| y -scale | 1.000029 |
| shear ($^{\circ}$) | -0.043368 |
| rotation ($^{\circ}$) | 0.002544 |
| Residual (RMS in μm) | 9.9 |

Table 4.4: The linear parameters that `GridKit` returns after going through the calculations. These parameters are fed into the simulation part of `GridKit` (together with random errors and non-straightness) to get the plots in figure 4.19. The largest sources of error are the shifts and the shear, or non-perpendicularity of the axes.

in the centre of each grid mark. Another $5\ \mu\text{m}$ is the accuracy of the machining, without any extra independent information this error will remain.

The parameters show that there are small remaining offsets after having removed the large offsets of several mm by hand. Because of that manual removal, not much can be learned from these numbers, other than that another shift has been done to make the expected and measured grids more similar. The scale parameters are very close to 1, on their own these parameters would give an error of only $5.0\ \mu\text{m}$ over the whole grid. Similarly, the rotation on its own gives an error of $6.8\ \mu\text{m}$. By far the largest contributors to the error in the raw grid are the shear and the offsets. The shear, or non-perpendicularity, is much larger than expected (0.04° instead of 0.0005°), but this does not matter, as it is easily corrected for.

After having corrected for the parameters in table 4.4, the non-straightness arrays show that there is a clear single deflection in the Y -gantry of about $15\ \mu\text{m}$. Any deflection in the X -gantries is so small that it is lost in the noise. This could be because any deflection in one of the two X -gantries cancels out the deflection in the other gantry. This shows that there is an overall non-straightness due to installation, but the sinusoid does not have three periods as `Autofib-2` had.



4.6.4 Result as input for simulation

When the parameters in table 4.4 are taken as inputs for the simulation software (with the minus signs accounted for), plus a sinusoid with an amplitude of $15\ \mu\text{m}$, a period of 0.5 and shift of -90° , a $5\ \mu\text{m}$ error on the expected grid, and a $9\ \mu\text{m}$ error on the measured grid, the simulation returns the plots in figure 4.19.

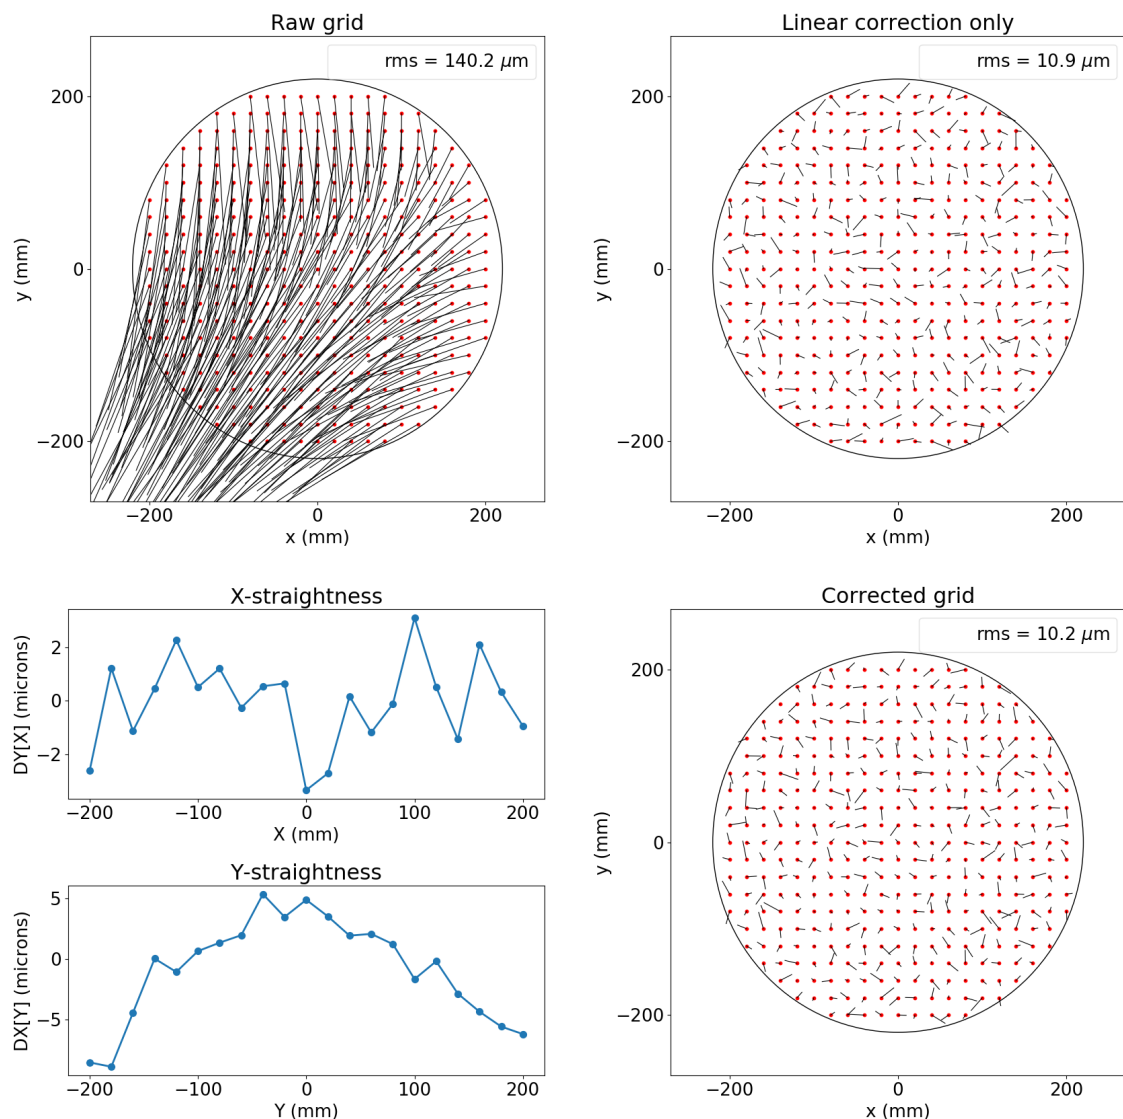


Figure 4.19: The results when the linear parameters together with random errors of $5\ \mu\text{m}$ and $9\ \mu\text{m}$ for the expected and measured grid respectively, are put into the simulation part of `GridKit`. Also included is a half-period sinusoid with an amplitude of $15\ \mu\text{m}$ and a shift of -90° . The results are remarkably similar to the real measurements in figure 4.18.




The error in the uncorrected for simulated grid is a bit larger (140 μm instead of 120 μm) than that for the real measurements. This could be because there may be some systematics in the grid marks and therefore in where the centroiding software `fibmeasure` finds the centre of each mark. This could end up as a systematic error per row or column, something that cannot be simulated in the current version of `GridKit`. This would explain some of the spikes in the straightness arrays in the real measurements in figure 4.18.

All the other behaviour, from the amplitude and shape of the straightness arrays to the remaining RMS error after the corrections (10.2 μm for the simulation and 9.9 μm for the real measurements) is captured very well by the parameters found by `GridKit`.

4.6.5 Forward look

The measurements described in this section have been done on the prototype aluminium field plate. As described in section 4.6.2, some more verification steps will have to be completed to ensure the repeatability of the grid mark measurements. These can only be done once the final invar field plates are finished. They are the following steps:

- First ensure that each robot can image all the grid marks reliably under different conditions (lighting, slightly different motor positions, etc.). The centroiding software still has to find the same centroids under those different conditions. These results do not have to be the same for both robots, as long as each robot on its own has good repeatability (preferably under 1 px or 4.4 μm).
- Then record these centroid measurements for each grid mark, and calculate the average centroid location including its RMS error. These should include the above varying conditions.



- These positions can then be used as the expected grid, with the RMS errors added in RSS fashion to get the overall measurement error on the expected grid (added to the error of the CNC machine that made the grid marks in the first place).

It is also important to realise that in this way we do not have a completely independent measurement of the grid, because the cameras are located on the robot gantries which are being verified. This is not ideal but it is the only way to perform these measurements. An independent measuring tool cannot be instructed to find the same centroid for these large grid marks as the biscuit cutter software does.

4.7 Summary

In this chapter I have shown that the software I have written to recalibrate the positioner robots for each reconfiguration (`GridKit`) works as expected. Certainly in simulations, `GridKit` recovers the input parameters to a very high (sub- μm) precision. It performs very well in the never-before seen situation where only part of the grid is available due to obscuration by fibres.

At the moment, the real measurements provide a recalibration with 10 μm RMS error. This is above the permissible total error of 8 μm and the budgeted for calibration error of 4 μm , but is because the grid marks themselves are not properly calibrated yet. The calibration process itself is sound. To calibrate the grid marks, a few steps still have to be done. These are outlined in section 4.6.5.



Chapter 5

Positioner performance

With the positioner fully built (during either the initial build or the final build) it is possible to test several aspects of it. In particular anything that moves has to be tested for accuracy, repeatability, and lifetime. Each section in this chapter is about a different part of the structure.

One such part is the set of 326 MOS retractors which had to be lifetime tested to ensure that no potential wear would have any negative effects on the performance of the retractors, or on the fibres themselves. This is described in section 5.1.

With the robots fully operational, it is relatively easy to check the robot cameras' plate scale, field of view, and location of the rotation centre on the sensor. Section 5.2 deals with these checks.

A fully operational positioner also enables us to check the Delta software that is written to control both of the robots during the repositioning of all the fibres from one configuration to the next. More on this software is in section 5.3.

Lastly, I have checked the repeatability of the robots. This is of utmost importance for the accuracy of each fibre's location on the field plate, as when the robots do not perform as expected, the fibres will definitely be out of position. Section 5.4 goes into detail about the repeatability of each of the axes, and the issues that arose.




5.1 Fibre retractor performance


5.1.1 Rationale and retractor design

Each of the 1900 MOS fibres are kept under constant tension, whether they are positioned on the field plate, put back into their park position, or during observations. This is to ensure that each fibre remains straight on the field plate, and therefore obscures other targets as little as possible, but more importantly, it also ensures that the fibres themselves will not twist or kink, particularly as they are moved out of the field. This would be disastrous, resulting in anything from fibres tangling up and becoming unusable, to fibres breaking.

The fibres are kept under tension in a device called a fibre retractor. The principle of the WEAVE retractors is based on the well-tested 2dF retractors (Lewis et al., 2002), with a few modifications. For 2dF each retractor held up to ten fibres, with all the unused fibres parked on the edge of the field plate. The left side of figure 5.1 shows most of the fibres parked around the edge of the field plate, with the top parts of the retractors visible around the edge. Each fibre was kept under tension with a constant-force spring and two pulleys: one moving, and one stationary. The ten fibres per retractor shared the same axle for the stationary pulley. The 2dF design showed the problems of fibres sharing space in a retractor: if one fibre comes loose for whatever reason, it could interfere with all the others in the same retractor. Furthermore, by having the pulleys share axles, any friction in one system can easily be transferred to another system and impede that fibre's proper positioning. The right side of figure 5.1 shows a retractor without any fibres; the stationary pulleys sharing the same axle are at the top of the image.

For WEAVE, with over twice as many fibres per field plate (up to 968 instead of 400), and a similar sized field plate, there is simply not enough room to accommodate all the retractor mechanisms next to each other. They have to be stacked somehow






Photos courtesy of Ian Lewis.

Figure 5.1: The 2dF retractors. *On the left:* twenty fibres are placed on the field plate, with the other fibres parked around the edge. The fibre retractors are visible around the plate. *On the right:* close up of an empty retractor. The clear box is for spare fibre length. All the pulleys are directly next to each other in the black casing and the fibres have to share this space.

away from the field plate. The redesign was also a good moment to address the problems the 2dF retractors experienced. The WEAVE design settled on six fibres per retractor with each fibre in its own space, and because of the field plate restriction, the design includes three tiers. Each retractor is now a lot thinner, but also much deeper. This has direct consequences for the size of the positioner, as the retractors tumble with the plates, and have to clear the surrounding structure.

Figure 5.2 shows an assembled retractor unit on its side, with the fibre buttons parked on the tiered park porches at A. There are three fibre slots on each side, for a total of six fibres per retractor. Each slot's depth is determined by the bend radius of the fibre. Two tiers would not give enough fibres overall, and with four tiers the positioner as a whole would become too large for operation.

The tensioning mechanism is delivered by a constant-force spring-and-pulley system. Each fibre loops around a stationary pulley and a moving pulley attached to a spring wound around a stationary hub. This system keeps the fibre under a constant







Figure 5.2: Assembled fibre retractor unit: (A) fibre buttons on park porches, (B) spare-length box, (C) fibres exiting retractor, (D) central plate with fibre secured in square groove, (E) fibre under tension, (F) fixed pulley, (G) tensioned moving pulley.

tension of 13 N. The moving pulley at G in figure 5.2 moves to the left when the fibre is pulled out by one of the positioner robots. This travel also defines the maximum length the fibre can be pulled out, and is designed such that every part of the field plate can be reached by multiple fibres.

The retractor holds the final 2 m of each fibre, of which at least 1 m goes into the spare-length box at C. This extra length is for the inevitable repairs that will have to be done on the fibre ends during the instrument lifetime of operations. Each fibre has a 120 μm outer diameter, loosely sleeved in flexible PEEK tubing with 500 μm outer diameter, and 350 μm inner diameter. The fibres are anchored in the retractor by pressing them in a closely fitting square groove between the spare-length box and the moving pulleys. The PEEK tubing ensures that the fibres themselves are not damaged by this procedure, and that any stretching is not transferred to the fibres.

The central plate is made from aluminium, because the machining on aluminium ended up much more precise than on the polycarbonate prototypes. The design






Photo courtesy of Ian Lewis.

Figure 5.3: 0.5 mm PEEK tubing (outer diameter) pressed into a square groove, and holding a 120 μm fibre. The fibre has room to move around, whilst the tubing is held securely by the accurately machined groove. This prevents the fibre from being stressed or stretched.

had to be severely light-weighted to keep the overall mass the same, resulting in the minimalistic see-through product, and again this was much easier to do in aluminium. The plate is anodised black to give it a clean, smooth finish. This enables someone to easily spot any debris when populating the retractor. It also prevents any oxidation on the surface. The covers are made of clear polycarbonate. Machining does not need to be as precise for these components (but still precise enough to match up the hole locations on the covers and the central plate) and it is useful to be able to see the fibre routing during assembly and testing.

5.1.2 Assembly

As each retractor houses six MOS fibres, we need 160 retractors for the 960 fibres for plate A. There are eight guide fibres per plate, and on plate A they will live in their own retractor. This gives a total of 168 retractors for plate A. Plate B has the same



number of retractors, but ten of them are modified to fit the mIFUs on that plate. This leaves 158 retractors for the 940 MOS and eight guide fibres, which fits as long as the guide fibres share their retractors with five MOS fibres.




Figure 5.4: A single retractor unit component breakdown, a total of 106 individual parts: (a)-(b) covers, (c) centre plate, (d)-(g) spare-length-box covers, (h) fixed pulley assemblies, (j) moving pulleys with constant-force spring assemblies: pulley, yoke, spring, and hub, (k) fasteners, (m) steel park porches.

The parts for 360 fibre retractors (to include spares) have all been manufactured. Each retractor consists of 106 individual parts, as seen in figure 5.4, which makes checking each part and assembling all retractors a very time-consuming task. I timed our assembly procedure with a test fibre bundle of 24 fibres (Sayède et al., 2016), and assembling one retractor with six fibres takes about 45 minutes. This excludes the time to check each part, does not take into account the time it takes to thread the fibres through the positioner (before they can be put into the retractor), and assumes the constant-force spring assemblies (at *j* in figure 5.4) have already been put together.

The decision was taken to check all the parts and pre-assemble the spring assemblies before the fibres arrived, as it is a completely separate process. This consists of



fitting an axle in each wheel, fitting yokes to half of the wheels, winding springs onto their hubs, and attaching the spring to the yokes. Given that most if not all of this will be done before the fibres arrive, the expectation is that two people together can assemble up to one fibre bundle of 24 fibres, or four retractors, per day. Integration of the fibres into the positioner will therefore take at least two months per plate.

Figure 5.5 shows three assembled and populated retractor units together with some empty retractor units. Twenty retractors (12% of the total around one plate) were installed next to each other around the field plate to check the build up of the thickness tolerances. They fit well, which gives us confidence that the other retractors will fit too.





Figure 5.5: Three populated retractors alongside some empty retractors at the edge of the field plate. They all fit around the field plate without any problems.

5.1.3 Testing

I conducted lifetime testing on one retractor (which I also reported on in [Schallig et al. \(2016\)](#)). This comprised testing the spring lifetime, assessing wear on moving components, and noting possible failures, such as slack in the fibre causing it to jump off the pulley. The test consisted of pulling out and retracting a fibre for a number



of cycles representative of a lifetime's operation. A lifetime is defined as 6 fields per plate per night over 70% of the total telescope time for five years, which amounts to about 8000 operations per fibre. An operation is defined as moving the fibre with the gripper back and forth over a distance of 175 mm, coming to a complete stop at each end. This distance was chosen as 175 mm is well over the average move distance of about 137 mm (Gilbert, 2016, p. 95).

After 8000 operations I opened the retractor to check the status of the different components. We found a tiny amount of debris on one of the prototype pulleys, which was due to a polishing effect from the turning of the prototype wheel. This did not inhibit its correct operation. The aluminium centre plate showed no signs of wear at all. Another 10000 operations were completed with different accelerations. The fastest tested acceleration was 2 ms^{-2} , with the prototype positioner robot. The force on the fibre due to the spring was strong enough that there was no slack as it was going back into the retractor, even at the highest accelerations. There was no visible wear on the centre plate and other components, and no extra wear on the wheels.

The tests did show that care has to be taken to make sure that the components are all completely clean. A speck of dried glue that had somehow landed on one of the axles prevented a pulley from moving correctly, with it slipping rather than turning. This was easily corrected, but it has to be noticed before something can be done¹.

The pulleys, spring drums and attachment yokes in the final version are made by injection moulding ABS plastic, and are all carefully checked. A number of pulley wheels had to be rejected because not enough material had been injected into the mould. Moreover, each wheel must be sanded by hand to ensure that all the injection moulding debris comes off, and that the pulley can spin freely in the yoke. This is a lot of extra work, and with over 4000 wheels and the knowledge of hindsight, the question arises whether it may have been better to go the machining route.

¹No glue is necessary in the assembly. As long as the assembly environment is clean, this should not happen again.



5.2 Gripper optics

Two robots on (x, y, z, θ) -gantries equipped with grippers are able to position the fibres anywhere on the field plate. The robots share the x -gantries, which consist of two rails. Each robot moves independently and software ensures that they do not run into each other. Each y -axis supports a z -gantry on which a rotational θ -axis is mounted, which in turn holds the gripper assembly. The final gripper design is based on a commercial two-fingered gripper unit from Schunk GmbH with custom fingers to grip the handle of the fibre buttons. See figure 5.6 for an overview. Each robot has an optical system that is mounted next to the gripper under the θ -axis, and together they form the gripper optical system.




Figure 5.6: Each robot is mounted on its own (y, z, θ) -axes, and shares the two x -gantries with the other robot. The θ -axis is aligned with the z -axis. The gripper consists of the gripper unit, custom-made jaws, and a camera.

The optical system has two foci with the same image scale, but in different places. One focus is located close to the rotation centre of the gripper and images a fibre held in the gripper jaws. The second focus images the field plate with the gripper






Figure 5.7: The gripper and camera unit. *On the left:* a schematic display. *On the right:* the actual unit in the lab. In both images the green lines denote the two paths the light can travel into the camera. The straight-through line is the plate-viewing focus, the longer path is the fibre-viewing focus.

raised, as a tool for calibrating fibre locations. Figure 5.7 shows the setup both as a schematic and the working unit in the lab. The green lines show the two foci.

5.2.1 Plate scale

This optical system needs to be calibrated for plate scale and rotation centre. A glass distortion target² can be imaged to measure the field of view, plate scale, and orientation. The pattern on the target has 125 μm spacing between the dot centres and 62.5 μm dot diameter. The field of view follows from the number of pixels in the height and width multiplied by the plate scale.

The biscuit cutter centroiding method (Gilbert, 2016) can be used on the distortion target too. Because there are multiple dots in the single image, an approximate centre or reference point (in pixels) has to be given to the software for each of the dots that we are interested in. It is quite time-consuming work to select these reference points. No fixed distance can be assumed between the dots, as this is exactly what

²Edmund Optics glass distortion target, 25x25 mm. Overall accuracy ± 0.001 mm. <https://www.edmundoptics.eu/p/25-x-25mm-0.125mm-spacing-glass-distortion-target/15362/>



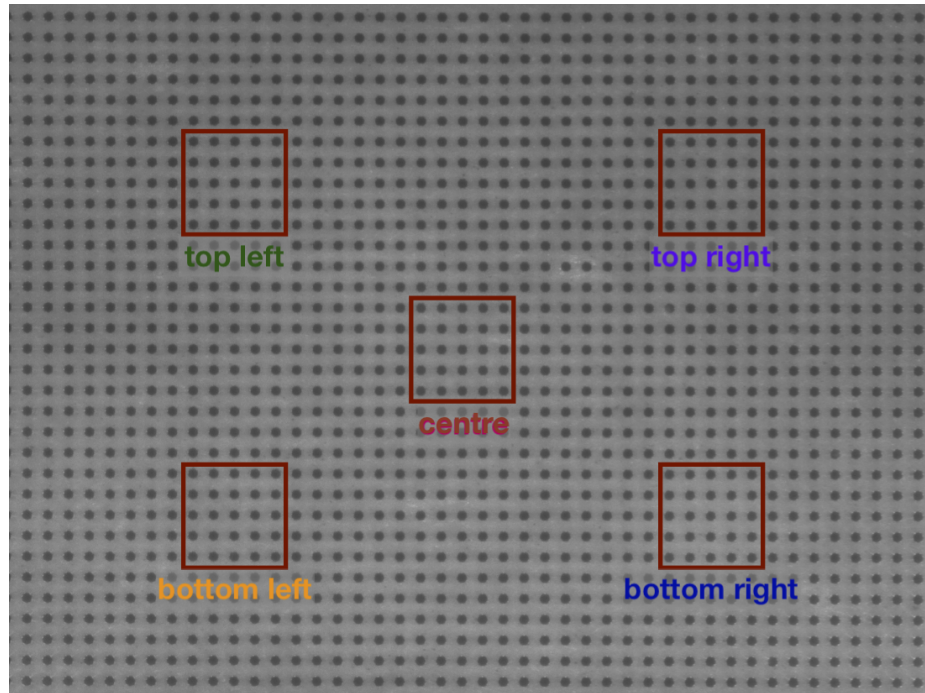


Figure 5.8: Image of the glass distortion target with $125\ \mu\text{m}$ spacing and $62.5\ \mu\text{m}$ dot diameter, imaged with Nona's camera in the plate-viewing focus. The five squares show the locations where the plate scale analysis is done. Counting the dots reveals that the field of view is about $5.63\ \text{mm}$ wide and about $4.13\ \text{mm}$ high. $(0,0)$ is at the top left corner in this image.

this exercise means to find out.

I selected five regions to calculate the local plate scale for, with the possibility of selecting more regions if the local plate scale varies substantially. The field of view of the plate-viewing focus is so large, that if a target is imaged close to the edge, either it or the camera is completely out of position and a reposition of the target or camera has to be done to view it closer to the centre. Therefore we can ignore the outer edge of the distortion target. The five regions are in the centre and at the corners of the relevant area.

Each dot in a region is centroided with the biscuit cutter method, and each top left dot is taken as the comparison point. The distance of the other dots to that top left dot is calculated in pixels and in physical length (in μm). Dividing these numbers gives a plate scale value in $\mu\text{m}/\text{px}$ for the 24 distances.



The differences can be visualised in a histogram, but plotting histograms over each other is not very insightful. Therefore I have plotted the lines connecting the top of each bar in the histogram. Figure 5.9 shows the results. The plate scales are all very similar, with a very small shift from left to right in the image, with the peaks from 4.35 to 4.37 $\mu\text{m}/\text{px}$. This shift is so small, that it is insignificant over the typical distances that are important for WEAVE. Over a distance of 1.0 mm this plate scale shift amounts to 1.0 μm .




Figure 5.9: Plate scale results. *On the left:* the histogram of all the regions together. The line segments connecting the top of each bar are plotted over the histogram, to show how the plot on the right was produced. *On the right:* the results split into each separate region. There is a very small shift in plate scale from right to left.

The bin widths in the histogram of all the data in figure 5.9 are smaller than necessary for our purposes. By binning the data with bin widths of 0.1 $\mu\text{m}/\text{px}$, shown in figure 5.10, we get a clear result for the plate scale: 4.36 $\mu\text{m}/\text{px}$. One could argue that this is still unnecessarily precise for the distances that we will typically see in WEAVE applications, in which case the plate scale can be quoted as 4.4 $\mu\text{m}/\text{px}$.

We assume that the beamsplitter and fold mirror do not distort or magnify the image through the longer path, and so the results for the fibre-viewing focus should be the same as for the plate-viewing focus.






Figure 5.10: Binning the plate scale data with bin widths of $0.1 \mu\text{m}/\text{px}$ shows that most of the computed plate scales are around $4.36 \mu\text{m}/\text{px}$. Extrapolating from this histogram, we can also see that almost all of the points are within the $4.35\text{--}4.40$ range.

The field of view of the far focus then becomes:

$$\begin{aligned}
 4.36 \mu\text{m}/\text{px} \cdot 1280 &= 5581 \mu\text{m} \\
 &= 5.58 \text{ mm}; \\
 4.36 \mu\text{m}/\text{px} \cdot 960 &= 4186 \mu\text{m} \\
 &= 4.19 \text{ mm}.
 \end{aligned}
 \tag{5.1}$$

5.2.2 Rotation centre

The rotation centre of the optical system is the reference axis for the fibre positions. All fibre centres in the images are measured relative to this. It can be found by imaging a stationary backlit fibre with different gripper orientations. For this we use the fibre-viewing focus. The other focus cannot be used as the distance from that focus to the rotation centre is too large to fit enough orientations onto the sensor. The different images stacked on top of each other describe a circle, the centre of which is the mechanical rotation axis.






Figure 5.11: *On the left:* a backlit fibre button as seen by the plate-viewing (straight through) focus. In normal operation only the spot in the prism is visible. *In the middle:* the result of rotating the gripper and camera unit about its axis and imaging the fibre with the fibre-viewing (longer path) focus. Each image represents a 30° rotation. The lower left corner cannot be filled because the button vane is in the way. *On the right:* a close up of the ring with the stacked images of the backlit fibre.

We can fit a circle through the fibre centres (found with the biscuit cutter method) and from there find the pixel that is in the centre of that circle. The distance to the (0,0)-value of the image is the distance to the physical location of the motors as described by the motor encoders. Using the plate scale, this pixel distance can be converted into a physical distance.

The rotation centre changes every time something is changed on the optical system. At the time of writing, there are still some final adjustments to be made. Therefore the above should be taken as a recipe to find the rotation centre, and there are no final results to report on.

5.3 Delta software

Where Configure is the program that searches for an optimal fibre layout for a given set of target objects, the actual order in which the fibres are moved to go from one configuration to another is handled by the software component called Delta, written by David Terrett ([Terrett et al., 2014](#)).

To minimise the time needed to move all the fibres from one configuration to the



next, the amount of moves has to be minimised. Delta looks at how the current configuration is built up, and attempts to find the sequence of moves that minimises the total number to go to the next configuration. The optimal way is to move as many fibres as possible directly to their new position. There are a few important rules to take into account to prevent any mishaps. Fibres that are blocked by other fibres crossing over them, cannot be moved. If the new position for the fibre is not empty, the fibre cannot be moved there directly, but it can be parked on its porch. If the porch is blocked by a fibre from a higher porch, the fibre cannot be parked, but it can potentially be placed in its new position.




Figure 5.12: Frame of a video showing a configuration from start to end. I combined a video made with the camera positioned above the field plate looking straight down with the software mimic of the same exact configuration. The blue and black lines are the placed and parked fibres in the mimic, the grey rectangles are the robots.

In principle, the order of the moves can be determined in advance, after which a set of instructions can be passed to the positioner to move the robots in the correct order. In practice however, the time taken to move a fibre can vary (if the fibre is not set down in the correct position on its first try, the robot has to try again) and



is therefore not predictable. One robot could end up blocking the way for the other one, making it wait until it has moved out of the way. In that time, the other robot could have started to position another fibre. The decision about which fibre to move next therefore has to be taken in real time, and is informed by the current robot trajectories.

Testing Delta on the real positioner (instead of in simulations) made for some nerve-racking moments where the two robots came quite close to each other. This was not a problem with Delta, but with the exclusion rules around each robot and between the robots from which Delta takes its information. After adjusting these values, the positioner configured several field plates' worth of configurations without any problems. Figure 5.12 shows a frame from a video produced by combining a real video of the moving robots shot from above, with the software mimic that shows the same movements. In this run, all the fibres are parked in the initial state.

The tests were done without any fibres and with adjusted z -values for the field plate and the park porches (high above the actual plate), to prevent the grippers from accidentally running into anything. Matching the locations of the retractor porches in software to their actual locations around the field plate will be essential to get the correct positioning behaviour.

It took about 40 minutes to configure a field from all parked fibres. The fibres were all placed in the correct spot on the first try as there are no real fibres to check the position of. Re-placing will add more time. The robots were set at lower speeds than when in operation on the telescope, but still there is a need for improvement. The video taken from above for this configuration run showed where some more efficiency improvements could be made, particularly at the start and the end of the sequence where only one robot is working. David Terrett is working on these improvements.



5.4 Robot repeatability performance

We can test the repeatability of the robot gantries by measuring a reference point many times with the gripper optical system, moving away from the point in between measurements. This reference point is a stationary, backlit fibre on the field plate, exactly as in figure 5.11. Its position is the centroid of the fibre within the image, in pixel coordinates, when the robot has returned to the same position. Repeating this measurement between gantry movements should yield the same result.

To make sure that the move has completely finished, each measurement is taken at least 1 second after the robots have reported the end of the move (the dwell time). Because of the way the software was written, an image is taken each time a move is completed. This means that for every cycle of moving away and back again, two images are taken each time, and the total wait time is twice the dwell time for one image. The total time between measurements is therefore twice the dwell time plus the time it takes to do the two moves.

The repeatability is always measured in the x, y -plane, as that is what is important for fibre positioning and starlight acquisition. Also important to note: the image x - and y -axes are not necessarily aligned with the robot X - and Y -gantries, as the gripper's orientation (θ) is not always the same.

5.4.1 Preliminary results

During the initial build phase, we measured the repeatability in the X -, Y -, and Z -gantries with Morta, and compared this to the repeatability for stationary measurements. This was originally reported in [Schallig et al. \(2018\)](#). The stationary measurements showed a spread of about a pixel in x , and a very small total RMS: < 0.2 px or < 0.8 μm for both sets of measurements depicted in figure 5.13. The results are clustered into rows and columns, and this is because of the chosen resolution



of 0.02 px of the biscuit cutter used to find the fibre location in the image.

Figure 5.14 shows the repeatability results for when the gantries have actually moved away and come back to the original location. The left plot shows the influence of movement in the x, y -direction, which gives an RMS of 0.8 px or $3.5 \mu\text{m}$. On the right the results are less good, where moving in the z -direction produces a strong correlated movement in x and y . The central bulge looks to be about the same size as the total spread in the graph on the left. From this, we started to investigate with the manufacturers whether we could adjust the mechanical preload to reduce these pitch and yaw errors.




Figure originally published in *Schallig et al. (2018)*.

Figure 5.13: The gripper optical system measuring a stationary fibre, 60 data points each. The colours denote a time series to check for temporal effects, with blue the early measurements and red the later ones. (This turned out to be wrong, see section 5.4.2.) Each measurement is taken 2 seconds after the previous measurement. *On the left:* the robot motors are completely switched off. *On the right:* the robot motors are switched on, but there is no movement away from the fibre in between measurements.

5.4.2 Later understanding

Unfortunately, there turned out to be a mistake in the code that produced the previously described plots³. The measurements were not plotted according to time, but in

³It turns out that Python's `glob.glob` routine does not read in the files alphabetically, but in the fastest way. This of course then changes the order of the measurements.






Figure originally published in [Schallig et al. \(2018\)](#).

Figure 5.14: Effect of moving the gripper optical system away from and back to a stationary reference point, 100 data points each. The colours denote a time series to check for temporal effects, with blue the early measurements and red the later ones. (This turned out to be wrong, see section 5.4.2.) The dwell time is 2 seconds, so each measurement is taken about 5 seconds after the previous measurement. *On the left:* movement in the x - and y -direction, 100 mm in each. *On the right:* movement in the z -direction, about 22 mm, with the images taken in the upper focus.

a ‘random’⁴ way. So any conclusion or gut feeling drawn from the temporal scatter was wrong. We only found out about this after the final build, about a year later, shortly before representatives from the manufacturers visited.

Fixing this mistake resulted in some interesting new insights. The updated plots are in figure 5.15, and immediately more structure can be seen. It is clear that there is a lot of temporal structure, especially in the first two plots where the motors are off and stationary, and the last plot where the z -motor moves up and down.

The structure in the first two plots proves that something else is going on with the camera-gripper system. As the robots cannot move at all in the first plot, it must be something in the camera system itself. The spread is much too structured for it to be vibrations in the lab or other such outside effects, as originally suggested in the paper. In the last two plots the data points still drift in the last measurements (the yellow dots), so it is clear that the next measurement series have to be much longer.

⁴Not really random, but definitely not easily predictable. (It depends on where in the tree structure of the file system a particular file ended up, and by which particular algorithm that tree structure is read.)






Figure 5.15: Figures 5.13 and 5.14 replotted, now with the correct temporal data and a better colour scheme. The figure ratio and the limits are exactly the same as in the previous plots. The temporal scale now runs from blue through green to yellow.

5.4.3 Longer measurements

Before taking the measurements described below, representatives from Schunk adjusted the mechanical preload (more on this in section 5.4.4) in the Z -motor. The original settings were out of spec, so we expected the adjustment to give us better results in the z -direction.

In this section, I discuss the longer measurements done with the robot stationary, and in the xy -direction, and any long measurements to do with the z -direction are located in section 5.4.4.

5.4.3.1 Servo only

We decided to redo most of the measurements taken earlier, but now with many more data points and on much longer timescales. This was now done with Nona, whereas the earlier measurements were done with Morta. The preload was left as before on Morta, for comparison purposes if needed. Figure 5.16 shows the results for when the system is on and servoing. There are almost 10,000 points in the dataset, and the total time for this measurement is 1 hour and 23 minutes. There was a 20 minute delay between powering up the positioner and starting the measurements.

There is a trend towards lower image y -values, and then going back up again slightly. This has to be due to parts of the positioner warming up. Because the motors have been servoing for 20 minutes already, it is not entirely clear how much






Figure 5.16: Measurements with no deliberate movements in the gantries, only servo movements. Each measurement is taken 0.5 seconds after the previous one, and $\theta = 0^\circ$. *On the left*: the spread of the measurements compared to the x - and y -means in pixel units. *On the right*: the same information as on the left, now shown as two separate time series. The RMS-value is that of the total distance to the mean.

of this is due to the motors, and how much is because of the camera itself warming up⁵. It is evident however, that the robot has to be exercised before a steady state is reached.

This behaviour can also be seen in the preliminary results in figure 5.15, where the first plot is with no power to the robots at all. The downwards shift there has to be caused by the camera itself warming up. The data for the second plot is taken immediately after the robot motors have been switched on, and even in the short time frame of these plots (120 seconds), the vertical shifts are evident. These plots can be directly compared to figure 5.16, because the gripper's orientation is the same in all three measurements.

The errors are small; the largest offset is less than a pixel ($< 4.4 \mu\text{m}$). The spread in y stays within 0.2 px after about 1500 measurements, which takes about 13 minutes. To completely eliminate the creep, the camera should be taking pictures often (or be in continuous shot mode) for about a quarter of an hour before the start of operations. Because of the 20-minute delay it is not clear whether the robots take

⁵The camera works with Power over Ethernet (PoE) and is always on, but does not use much energy when not actively taking images.



the same amount of time warming up, or longer. Powering up the positioner half an hour before the start of operations should give the system enough time to mostly reach a steady state. In any case, the errors are still small, and this is only a small contribution to the overall error budget.

5.4.3.2 Moving in x and y

Now we look at the results when the movement is in the xy -plane. Before taking each data point in figures 5.17 and 5.18, Nona moves away 100 mm in both the x - and y -direction, for a total distance of 141 mm. The dwell time after moving and before taking the image is 1 second in both plots. Because of the way the software works, after each move an image is taken. Therefore the time between each data point is twice the dwell time plus twice the move time. This is about 3 seconds in total. The fibre is imaged in the upper focus. The only difference is that there is a 90° rotation between the two sets of measurements.

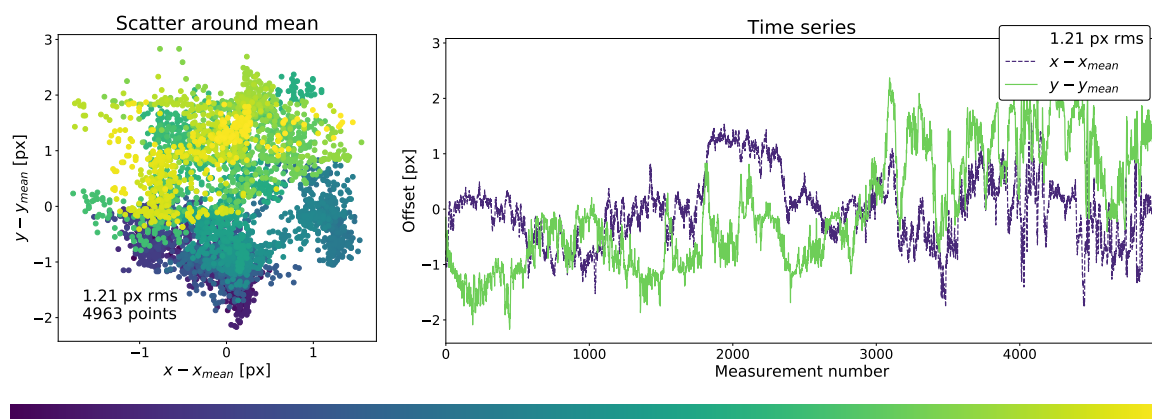


Figure 5.17: The repeatability after moving the robot in both x and y and back, 100 mm in each. The dwell time is 1 second, and $\theta = 90^\circ$. *On the left:* the spread of the measurements compared to the x - and y -means in pixel units. *On the right:* the same information as on the left, now shown as two separate time series. The RMS-value is that of the total distance to the mean.

The results of both plots are comparable, even with the 90° rotation difference, if the first 30-odd points are not taken into account in figure 5.18. In both plots it looks






Figure 5.18: The repeatability after moving the robot in both x and y and back, 100 mm in each. The dwell time is 1 second, $\theta = 0^\circ$. Apart from the first 30-odd points, the results are comparable to figure 5.17.

like a ‘spot’ of roughly a pixel diameter slowly drifts over the plotting area. There is no preference in direction in figure 5.17, whereas in figure 5.18 x and y are slightly more correlated. This is more easily seen in the time series parts of the plots. We do not know what could cause this, and why it would show up only sporadically. The fact that rotating one of the plots in the direction of the other plot does not make the plots look more alike, does point towards some random effect.

Both sets of measurements took about 4 hours in total, which is much longer than the time it takes to reconfigure a whole plate of fibres. However, only the longer timeframe shows the accuracy of the system in x and y , as opposed to the repeatability. The positioner error budget (table 2.4) includes an entry for the robot gantry accuracy, which is $7.1 \mu\text{m}$ ($5 \mu\text{m}$ for each axis). The accuracy measured here is the RMS-value multiplied by the plate scale, or:

$$1.25 \text{ px RMS} \cdot 4.4 \mu\text{m/px} = 5.5 \mu\text{m RMS}. \quad (5.2)$$

For these measurements, we also read out the encoder values for the gantries. The positioner provides the encoder values every tenth of a second or so, and the last provided encoder value at the time of the request is recorded. So whilst this is not




always the encoder value at the time of the measurement (or even at the time of the encoder value request, these actions are not happening at the same time), it is very close to the true value.

In figure 5.19 we see that the feedback loop for the X -, Y -, and θ -motors is very good. The Z -motor reported two very different values from the mean value, which is puzzling, as the robot did not get any commands to move in the z -direction. This is slightly worrying, because if this is a true offset, this will have an effect on image quality in the upper focus. Even worse in the lower focus, this could mean that the robot could run into buttons already placed on the field plate, or not reach a button if it stays too high. The offsets reported here are of the order of $270\ \mu\text{m}$, but we will see later (in figure 5.23) that the offsets can be even larger.

Even if this is not a true offset, but a glitch in the reporting, it would still have an effect on the positioner. To ensure that the robots cannot run into the side of the (parked) buttons, the software Delta has exclusion zones built into the volume above and extending past the field plate. If the reported height dips into these exclusion zones, the positioner will go into an error state, and will not be able to position autonomously until the error has been solved.

With the reported X - and Y -gantry encoder values so close to the actual values the robot is sent to, most of the uncertainty in the real measured position depends on the way the whole gripper-camera system is built and attached to the positioner. This uncertainty, at least for moves in the xy -plane, is smaller than the estimated error quoted in the error budget (table 2.4). This is good, because this is expected to be the largest contribution to the overall error.

Of course this value only holds if movements in the z -direction do not add any more error in the xy -plane, but as we saw in the preliminary results, this is not true. We took more measurements in the z -direction to look at this issue in more detail, which are described in section 5.4.4.






Figure 5.19: Encoder values for the measurement points in figure 5.18. Top to bottom is X, Y, θ , and Z. The repeatability is mostly very good, but the two high points in Z (270 μm off) are slightly worrying.



5.4.4 Z-motor issues

From the preliminary results it was already clear that there is a significant problem with the Z -gantry. To understand why moving in the z -direction results in such strongly correlated values for x and y , we look at the Z -gantry in more detail. Figure 5.20 shows close ups of the gripper system on the Z -gantry.




Figure 5.20: A close up of the Z -gantry. The motor is stationary, and the gantry (with the gripper system attached) moves through the motor. The red line shows the edge of the motor in all three images, the yellow line is the edge of the gantry. *On the left:* a CAD model of the Z -gantry. The exact details of how the gripper system is attached to the Z -gantry are not correct here, but the model does show clearly the location of the Z -motor with respect to the Y -gantry. *In the middle:* the gripper in the lowest position, ready to pick up a fibre off the field plate. *On the right:* the gripper in an intermediate position, travelling to the parked fibres on the edge of the field.

The model on the left shows that the Z -motor is directly attached to the Y -motor. The motor itself is stationary, and the gantry moves through the motor. This is different from the other gantries in the positioner, where the gantry is stationary, and the motor moves along it. The gripper system is bolted to the side of another rail (at the height of the θ -motor, with the horizontal Schunk label in figure 5.20) that is attached to the top and bottom ends of the gantry. This is different from the (out-of-date) CAD model, which shows a horizontal bracket connecting the gripper



system to the gantry at the top only. Obviously the system as built is much better than the model shown here, but it still has major problems.

The gantry is not clamped to anything, its only restraint is the Z -motor. This motor is very precise in the vertical direction, but has a bit of play in the other two directions in order to be able to move relative to the gantry. This play is constrained by three pairs of crossed bearings, with one fixed shoe and an adjustable shoe, shaped like v-grooves. On this adjustable shoe the preload can be set. Increasing it pushes the internal rail more into the v-groove, at the cost of increasing the wear in the system and reducing the lifetime.

The two v-grooves point along the y -axis, so any scatter we expect in the y -direction. From the initial measurements we saw however that the scatter was not only in the y -direction, but in a diagonal in the x, y -plots. The cantilever of the gripper and θ -axis load is along the x -axis, and this together with the v-grooves in the y -direction could account for this diagonal behaviour.

The measurements in the next sections shed some more light on the exact behaviour of the z -axis.

5.4.4.1 Moving in the z -direction: upper focus

To check the behaviour of the z -axis, the gripper unit moved down about 12 mm and back up again, after which an image was taken in the upper focus. Doing this for about 2.5 hours results in figure 5.21. The data suggests that if there is an x, y -correlation, it is not very strong, but there does appear to be a systematic drift.

This result is quite different from the initial result in figure 5.14, where there is a clear correlation between x and y . Is this only due to measuring over longer timescales, or is there more to be investigated? The other important location to measure in is the lower focus, where most of the measurements on the telescope will be done. In that location it is even more important that the robot behaves as designed.






Figure 5.21: Z -measurements in the upper focus, $\theta = 90^\circ$. The dwell time is 0.5 s; the overall duration is about 2.5 hours. *On the left*: the spread of the measurements compared to the x - and y -means in pixel units. *On the right*: the same information as on the left, now shown as two separate time series. The RMS-value is that of the total distance to the mean. There is no strong overall correlation, although there does appear to be a systematic drift.

5.4.4.2 Moving in the z -direction: lower focus

So far, all the measurements have been done in the upper focus. This is for two reasons. It is much safer: there is no fibre button to bump into by accident with the gripper jaws. Also, the lower focus is designed to fall just in front of the jaws, on a line that goes between the two jaws, so a fibre can be imaged whilst it is being carried. At the moment though, this is not properly aligned yet, so the jaws are also physically in the way of the focus. Therefore, in order to take measurements in the lower focus, we needed to temporarily (carefully!) remove the gripper jaws.

Figure 5.22 shows the results of the measurements in the lower focus, with $\theta = 0^\circ$. There is a very strong anti-correlation between x and y . This is more similar to the initial results, apart from the fact that the correlation has now turned into an anti-correlation. This is due to a 90° rotation difference between the two measurement sets, so the image x - and y -axes now align with the motor y - and x -axes respectively, instead of x and x aligning.

The extent of the correlated scatter is ± 4 px. With a plate scale of $4.4 \mu\text{m}/\text{px}$,






Figure 5.22: Z -measurements in the lower focus, $\theta = 0^\circ$. The dwell time is now 1 s. *On the left:* the spread of the measurements compared to the x - and y -means in pixel units. *On the right:* the same information as on the left, now shown as two separate time series. The RMS-value is that of the total distance to the mean. A strong anti-correlation is visible between x and y .

this is equivalent to $\pm 18 \mu\text{m}$, which is horrible. Clearly something needs to change in the setup, and changing the preload even more will not reduce this amount of scatter enough. Furthermore, the cantilever of the gripper and θ -axis load is not affected by this, and therefore an offset will remain in the robot's x -direction.

The encoder values for this set of measurements are shown in figure 5.19. As before, the x -, y -, and θ -values are all very similar, as expected. The z -values are mostly very similar too, apart from one data point. This one is 1.56 mm away from the rest of the values, which is an enormous distance on these scales. The z -direction is higher numbers closer to the field plate, so at least (if this reported value is real) it would not mean that the robot would crash into the field plate, but it is still a concern.

A check of several measurements close in time to the odd one suggest that the odd measurement is shifted towards a higher y -pixel value by about 2 px, and this is reflected in the centroid values. However, the y -encoder values are all the same, and only the z -value differs. On the other hand, the image does not look different from the other images; it is a very similar focus. 1.65 mm height difference would be






Figure 5.23: Encoder values for the measurement points in figure 5.22. Again the repeatability in x , y , and θ is very good, but the low point in z is extremely worrying, as it is 1.56 mm away from the other points. All the other points are within 3 μm from each other.



clearly visible as a defocus. This suggests that the encoder value sent through the system is not the real value the motor is at. This could be due to the fact that the encoder values are not requested at the same time as the move is finished, but this immediately raises another question: why does this happen only once or twice every few thousand measurements?⁶

5.4.4.3 Motor temperatures

The motors in the positioner are optimised for a certain internal temperature, which is for the usual conditions (such as factories) these motors work in. WEAVE's fibre positioner is not an average environment. The motors do not have to work much at all compared to a production environment, do not get very warm, and therefore may not work optimally. On top of that, telescope domes are cold compared to factory floors. To see whether a different temperature has an effect on the positioning, we can exercise the motors by moving them around at fast speeds. However, this is quite cumbersome, as the Z -motor has to move up and down a few times before taking each image, and this takes time.

Both gripper-theta-camera systems have a 'magspring' or constant-force magnetic spring. It is attached to passively compensate for the weight of this system that the Z -motor has to hold in a steady state. Another way to increase the temperature in the Z -motor is to disconnect this magspring and have the motor deal with the extra weight. In this way the Z -motor is exercised more not with extra movements, but with extra weight.

Figures 5.24 and 5.25 show the results with and without the magspring connected. These are long moves in the vertical direction, almost the whole available z -travel (35 mm), with the images taken in the lower focus. After several hours of continuous motion with the magspring on, the temperature in the Z -motor only peaked at 32 °C.

⁶Postscript: this turned out to be a bug in the robot control software and has since been fixed. See also the end of section 6.1.






Figure 5.24: Magspring on, $\theta = 0^\circ$. Over almost 10 hours of measurements the internal temperature does not rise significantly, staying within 1°C for each motor. The one drop in Z -temperature is due to a break between subsets of measurements, during which the Z -motor had time to cool down slightly. The scatter on the left is strongly anti-correlated.

For almost 5000 measurements each of the motor temperatures stayed comfortably within 1°C .

With the magspring disconnected, the temperature rose to about 38°C within 2000 measurements, and to 40°C after all the measurements, with the whole lab warming up with it too. It is unclear how many measurements there were exactly in this test, as a large chunk of them somewhere in the middle of the set was lost just after the test was completed⁷. The internal temperature still stayed well below the maximum permissible temperature of the Z -motor, so there was no danger of overheating.

With the magspring removed, the RMS also increases, but the scatter increases only along the already established correlation. Because it is unclear how long each set of measurements took—the test without the magspring could well have taken an hour or so longer—we cannot be certain that this spread is entirely due to an

⁷The two sets of measurements (with and without magspring) actually consisted of several sets of measurements each. The robot would sometimes time out on a move (this bug has since been fixed in the robot movement software), and only when checking back into the lab this would be seen (a stopped robot) and a new run would be started. It was therefore not clear how many points were in each run. Some of these runs were lost with the initial analysis.






Figure 5.25: Magspring off, $\theta = 0^\circ$. Several measurement runs after another. The sudden jumps around 2000 and 2100 are because chunks of data were lost. The temperatures of the other motors rise because the lab is getting warm from the effort of the Z -motor. In the left plot, the scatter also increases, but only along the correlation.

increased temperature. It could also be due to the effect that the ultimate extent of the scatter had not been reached yet even after several hours, and by taking more data, the scatter therefore increases. The measurements do make very clear though that the problem of the highly correlated large scatter in the lower focus persists over several days and several measurement runs in different circumstances.

It is also unclear why the later measurements on average favour to go to lower y -pixel values when the magspring is on, and to higher y -pixel values when the magspring is off.

5.4.5 New Z -motor

With all this information in hand, a replacement Z -rail of a different design was requested from Schunk. This new gantry has many more bearings on each rail, a longer engagement of the motor on the rail, and adjustable preloads on both sides. This should diminish the scatter in the y -axis. It also allows us to reduce the perpendicular distance from the gripper and θ -axis to the z -axis, therefore reducing the scatter in the x -axis.

The rail was delivered during the last month of writing up, and therefore the results






Figure 5.26: The data with the new Z -motor. *On the left:* the scatter is much smaller than for the z -measurements with the previous motor, but some structure still persists. *On the right:* the problem with some z -values being different still persists too. Here, one data point is 0.21 mm lower than the others.

shown here are only exploratory. Figure 5.26 shows the new situation. The overall scatter is much smaller, although there is still a correlation, and the shift persists too. However, this could probably still be fine-tuned by adjusting the preload on either side, or changing the exact configuration of the gripper system on the Z -motor.

The repeatability of this axis is:

$$0.85 \text{ px RMS} \cdot 4.4 \text{ } \mu\text{m/px} = 3.7 \text{ } \mu\text{m RMS.} \quad (5.3)$$

The potential problem with the encoder z -values persists as well. This ‘wrong’ data point is now the same one as the yellow data point on $(-3, 0.8)$ in the left plot, which could indicate a correlation between these two offsets.

5.4.6 Conclusions and caveats

From the stationary measurements the advice arises to power up the motors and exercise the camera at least half an hour before any accurate positioning can be done.

The X - and Y -motors behave as expected, and even go below the tolerance limit stated in the positioner error budget in table 2.4. The Z -motor was much more



problematic, and the only way to fix this was to order a different Z -motor from Schunk. This new motor performs better, but with its current settings it is still not doing as well as it should do. More adjustments and measurements are needed to bring down the scatter to a manageable level. However, even with this issue still open, the calculated repeatability for all the motors together becomes:

$$\sqrt{5.5^2 + 3.7^2} \approx 6.7 \text{ } \mu\text{m RMS}, \quad (5.4)$$

which is already smaller than the allocated budget.

These longer measurements are done with one robot only (Nona), but the earlier measurements show that Morta has the same issues. To check that the issues are solved completely with the new Z -motor setup, the measurements should be repeated with Morta, once it has a new Z -motor.

When the Z -motor performs more as expected, a set of measurements should be taken where the robot moves in x , y , z , and θ in one move, to better simulate the moves that each robot will have to make to pick up or position a fibre.

The issue of the outlier measurements in the encoder z -values remains open.

5.5 Summary

This chapter deals with a few distinct parts that are essential to the positioner. First, I showed that the retractors work well and will work well until far past their designed lifespan. They are compliant with the size tolerances, and filling them with the MOS fibres will take several months of work.

I have also calculated the plate scale and field of view of Nona's camera, in the far or plate-viewing focus. This focus and its properties is necessary for the grid measurements described in chapter 4. Similar calculations still need to be done for Morta's camera system.



The Delta control software written by David Terrett works, and my video has highlighted some areas for improvement.

The largest part of this chapter was devoted to the robot repeatability measurements. The x - and y -axes are accurate and repeatable to within the tolerances set by the error budget, but the z -axis posed large problems, and the measurements revealed a design flaw in this part of the robot. The only way to improve this, was to order a new z -axis with different properties. The initial measurements on this new axis show that there is a big improvement on the repeatability. The calculated repeatability for the X -, Y -, and Z -motors together is now $6.7 \mu\text{m}$ RMS, which is smaller than the allocated budget, but there is still room for improvement on the Z -motor, which would lower this value even more.



Chapter 6


Discussion

6.1 Discussion of this work

This work is concerned with the metrology, calibration, integration and testing of the WEAVE fibre positioner. In the previous chapters I have shown that we are well underway to delivering an instrument that performs to the required specifications.

I have devised a way to accurately build up the positioner and check every step of this process. The challenge is that the two field plates have to occupy the same space and tilt (within 0.1 mm and 12") when in the observing position. This is crucial for ensuring that all the fibres are in focus, and thus that the amount of light coupled into the fibres is strong enough for high-quality observations. The two field plates are on opposite ends of the tumbler structure, so there are many possible elements that can prevent the challenge to be met. The measurements to check the building process show two important points:

1. It is actually possible to take all the different measurements with the precisions required to check every step of the process.
2. The building process delivers a positioner that is partially compliant, and I have made recommendations to improve this.



One of the remaining problems is that distance of one of the field plates to the interface of the instrument is too short, that is, the field plate is too close to the corrector optics. To solve this problem, we cannot just put in any shims, as this would only exacerbate the problem. Therefore I have suggested to lift the whole tumbler axis on its M8 screws, giving up to half a mm extra space between the field plates and the corrector optics. The plate that is now too far away from the ideal position can be shimmed to move it closer to the optimal position.

Close to finishing this thesis, the invar field plates were finished and their flatness and thickness measured. The flatness is acceptable¹, within 20 μm peak-to-valley on most of the area of each plate. They are however much thicker than expected, about 9.4 and 9.8 mm instead of 9.0 mm. In the meantime the tumbler has been lifted on its screws as high as it can go already, so a new solution has to be found to gain almost a mm of space. Shimming the whole positioner with respect to the tumbler is an option, but it is a last resort, as it is very cumbersome to check (the positioner-corrector assembly has to be put in prime focus to check and taken off for any adjustments) and can introduce unwanted and unknown side effects.

Another possible solution is to remove up to 1 mm of material from each tumbler drum, to shorten the tumbler assembly and regain that distance to the corrector optics. About 0.5 mm would be removed from either side of each drum, to ensure that the flanges remain thick enough to provide stability for the rest of the assembly. This also has consequences for the screws that go through these flanges; shorter screws would be needed to attach all the parts without screws protruding into solid material.

The second piece of work concerns calibration software for the fibre positioner robot gantries. The software has been shown to work extremely well as a modelling tool for linear and non-linear errors that appear when a grid is surveyed with gantries that are affected by these errors. The linear errors and non-straightness of the axes

¹It was very difficult to get the plates within the flatness requirements, as the plates deformed several times over time after lapping.



are calibrated out exceptionally accurately, and only random errors remain. This is very useful to know when the software is used on real data from the positioner: remaining large errors are not a feature of the software, but are there because of assumptions made in the measuring process.

I made a start on addressing these assumptions, but work remains to prepare the grid measurements for these calibration purposes. In particular, the centroiding procedure for the grid marks should be evaluated. Even with these assumptions, the remaining error after recalibration on a preliminary dataset is $10\ \mu\text{m}$ (RMS), which is close to the desired range of $4\text{--}6\ \mu\text{m}$.

A significant part of this thesis concerns the performance of several distinct components in the positioner. The fibre retractors perform to the specifications, and a test retractor has undergone several lifetime tests without showing any degradation. The retractors are size compliant, and will all fit around the perimeter of each field plate. The next step is to fill the retractors with the MOS fibres, when they arrive.

I have outlined the procedure for calculating the plate scale and field of view of the robot cameras, and have done so for Nona's camera. This plate scale is $4.36\ \mu\text{m}/\text{px}$ and this value is critical for every measurement in which the camera is utilised. The same still needs to happen for Morta's camera, and the fields of view of the near foci have to be determined.

Accurate fibre positioning cannot happen if the robotic positioners themselves are not accurate or repeatable. It is therefore of utmost importance to check the behaviour of each axis separately compared to a stationary target. The repeatability of the X - and Y -axes is sufficient to within the tolerances set by the error budget, but the original Z -gantry introduced unexpected and unwanted deflections in the x, y -plane. These measurements revealed a design flaw in the original robot design, and a new Z -gantry with different internal properties was acquired. Initial measurements on the new gantry show that a lot of the problems have been resolved, but there is



still room for improvement.

Confirmation measurements done just before this thesis was finished show that this improvement can be made by calibrating out a temperature effect. As the Z -motor heats up, the gripper jaws move away from the Y -axis in a predictable way that can be calibrated. Furthermore, the measurements confirmed that the odd z -values reported on throughout chapter 5 are a result of a bug in the robot control software. Due to a threading issue two sets of commands were executed in the wrong order, which meant that the robot was already in another position when the Z -encoder reading was queried. This has now been resolved.

6.2 Further work on the positioner

Several larger and smaller jobs still have to be done before the fibre positioner is ready to work on the WHT. The most time-consuming task by far is to integrate the 1900 MOS fibres into 336 fibre retractors. This will take at least several months. Before that can happen though, all the fibre bundles have to be threaded through the tumbler. This has to happen in such a way that no prisms are knocked off their fibres, and cables do not get stuck behind each other. As this has not been done before in this capacity, there is no way of knowing how long this part of the process will take, but it is clear that it is more of the order of weeks than of days.

The LIFU has to be finished at NOVA in the Netherlands, after which it can be integrated in the positioner. This too has to be threaded through the tumbler, and this will most likely be hampered by the MOS fibres already in place. On the other hand, the large LIFU fibre cable will hinder threading the MOS fibres. As the whole instrument is behind schedule, preference will probably be given to whichever fibre cables arrive first, to not slow down the process even more.

The positioner still has to be extensively tested on fibre handling. There are many



little parameters that affect the positioning process, from jaw gripping strength to the location of the retractor park porches in the control software. Finding a back-illuminated fibre on the plate and aligning the jaws to the button vane to be able to grip it involves several steps that have been tested separately, but not together. These kinds of tests should be done with the test fibre cable, which minimises the risk to the science fibres. Ideally this should be finished before any science cables arrive to be integrated.

The Delta positioning software has to be fine-tuned to reach the desired re-configuration time of < 60 min from field to field.

After these integration and verification steps have been successfully completed, the positioner will be shipped to the WHT. Care has to be taken to ensure that no vibrations or sudden shocks loosen components such as screws or prisms, or deform crucial parts of the positioner. Commissioning off and on the telescope then involves a whole suite of tests to ensure not only the correct behaviour of the fibre positioner, but also of the prime focus corrector, fibre positioner and spectrograph as one unit. One of the positioner-specific tests that I recommend be done is to perform flexure measurements under different angles, as this was impossible in the Oxford lab.

In conclusion, the WEAVE instrument and the fibre positioner in particular are almost finished and getting close to start surveying the night sky. Many crucial steps still have to be done before observations can begin, but certainly the positioner system has progressed very nicely in the past few years.



References

- Don Carlos Abrams, Kevin Dee, Tibor Agócs, Emilie Lhome, José Peñate, Attila Jaskó, Evelin Bányai, José A. Burgal, Gavin Dalton, Kevin Middleton, Piercarlo Bonifacio, J. Alfonso L. Aguerri, S. C. Trager, and Marc Balcells. The mechanical design for the WEAVE prime focus corrector system. In *Proc. SPIE*, volume 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91472K, Aug 2014. doi: 10.1117/12.2055894.
- Tibor Agócs, Don Carlos Abrams, Emilie Lhomé, Kevin Dee, Gavin Dalton, Kevin Middleton, Piercarlo Bonifacio, J. Alfonso L. Aguerri, and S. C. Trager. Final optical design for the WEAVE two-degree prime focus corrector. In *Proc. SPIE*, volume 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 914773, Jul 2014. doi: 10.1117/12.2055306.
- Masayuki Akiyama, Scott Smedley, Peter Gillingham, Jurek Brzeski, Tony Farrell, Masahiko Kimura, Rolf Muller, Naoyuki Tamura, and Naruhisa Takato. Performance of Echidna fiber positioner for FMOS on Subaru. In *Proc. Spie*, volume 7018 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 70182V, Jul 2008. doi: 10.1117/12.788968.
- ASTRONET WFS Working Group. Report on a wide field, highly multiplexed spectrograph, 2011. <http://www.astronet-eu.org/IMG/pdf/D26-Vsdef-2.pdf> [Accessed: 25-09-2016].
- Astropy Collaboration. Astropy: A community Python package for astronomy. *Astronomy & Astrophysics*, 558:A33, October 2013. doi: 10.1051/0004-6361/201322068.
- Astropy Collaboration. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *The Astronomical Journal*, 156:123, September 2018. doi: 10.3847/1538-3881/aabc4f.



- Marc Balcells, Chris R. Benn, David Carter, and 21 other authors. Design drivers for a wide-field multi-object spectrograph for the William Herschel Telescope. *Proc. SPIE*, 7735:77357G–77357G–15, 2010. doi: 10.1117/12.856947. <https://dx.doi.org/10.1117/12.856947>.
- Chris Benn, Lilian Domínguez, Ceclia Farina, Chris Evans, Sergio Pico, Marie Hrudkova, Raine Karjalainen, Juerg Rey, and David Terrett. WEAVE Concept of Operations. *Internal Document WEAVE-SCI-004 version 1.97*, 2017. Date 26-01-2017.
- Andrea Bianco, Giorgio Pariani, Matteo Aliverti, Alessio Zanutta, James Arns, and 11 other authors. VPHGs for WEAVE: design, manufacturing and characterization. In *Proc. SPIE*, volume 10706 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107064X, Jul 2018. doi: 10.1117/12.2311693.
- Kevin Bundy, Matthew A. Bershad, David R. Law, Renbin Yan, Niv Drory, Nicholas MacDonald, David A. Wake, and 61 other authors. Overview of the SDSS-IV MaNGA Survey: Mapping nearby Galaxies at Apache Point Observatory. *The Astrophysical Journal*, 798(1):7, Jan 2015. doi: 10.1088/0004-637X/798/1/7.
- Manuel Canchado, Antonio Romero, Óscar Maroto, Albert Tomas, and 13 other authors. The WEAVE focus translation system: from design to construction. In *Proc. SPIE*, volume 9912 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99126D, Jul 2016. doi: 10.1117/12.2232985.
- Jorge Cham. PHD Comics, 1997-2019. <http://phdcomics.com/>.
- I.R Cole. Modelling CPV. *Doctoral Thesis*, Loughborough University, 2015. https://repository.lboro.ac.uk/articles/Modelling_CPV/9523520.
- Scott M. Croom, Jon S. Lawrence, and 23 other authors. The Sydney-AAO Multi-object Integral field spectrograph. *Monthly Notices of the Royal Astronomical Society*, 421(1):872–893, Mar 2012. doi: 10.1111/j.1365-2966.2011.20365.x.
- Gavin Dalton. WEAVE Trade-Off Summary. *Internal Document WEAVE-SYS-018 version 1.0*, 2014. Date 13-03-2014.
- Gavin Dalton, Scott Trager, Don Carlos Abrams, and 40 other authors. WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope. *Proc. SPIE*, 8446:84460P–84460P–12, 2012. doi: 10.1117/12.925950. <https://dx.doi.org/10.1117/12.925950>.



Gavin Dalton, Scott Trager, Don Carlos Abrams, and 50 other authors. Project overview and update on WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope. *Proc. SPIE*, 9147:91470L–91470L–11, 2014. doi: 10.1117/12.2055132. <https://dx.doi.org/10.1117/12.2055132>.

Gavin Dalton, Scott Trager, Don Carlos Abrams, and 69 other authors. Final design and progress of WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope. *Proc. SPIE*, 9908:99081G–99081G–10, 2016. doi: 10.1117/12.2231078. <https://dx.doi.org/10.1117/12.2231078>.

Gavin Dalton, Scott Trager, Don Carlos Abrams, and 71 other authors. Construction progress of WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope. In *Proc. SPIE*, volume 10702 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107021B, Jul 2018. doi: 10.1117/12.2312031.

ETSRC. Report by European Telescope Strategic Review Committee on Europe’s 2-4m telescopes over the decade to 2020, 2010. http://www.astronet-eu.org/IMG/pdf/PlaqueT2_4m-final.pdf [Accessed: 25-09-2016].

Patricia Fara. Newton shows the light: a commentary on newton (1672) ‘a letter ... containing his new theory about light and colours...’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373, 2015. <https://doi.org/10.1098/rsta.2014.0213>.

Gaia Collaboration. The Gaia mission. *A&A*, 595:A1, Nov 2016. doi: 10.1051/0004-6361/201629272. <https://ui.adsabs.harvard.edu/abs/2016A&A...595A..1G>.

James Gilbert. New developments in robotic fibre positioning for astronomical multi-object spectroscopy. *Ph.D. Dissertation*, University of Oxford, 2016. <https://ora.ox.ac.uk/objects/uuid:ca88e3e2-cd47-422e-96a7-d644a88176d0>.

James Gilbert. fibmeasure: Python/Cython module to find the center of back-illuminated optical fibres in metrology images, Mar 2016. <http://ascl.net/1603.014>, <https://github.com/labjg/fibmeasure/tree/a86e3f4a1ed97d08d4eae6b7564ef2d7d4a94673>.

Adam Ginsburg. Image Registration Methods for Astronomy,



2012. https://github.com/keflavich/image_registration/tree/43e750ba98df5599ad6509e9a78d4269accf9daf.
- J. M. Hill and M. P. Lesser. Deployment of the MX spectrometer. In David L. Crawford, editor, *Proc. SPIE*, volume 627 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 303–320, Jan 1986. doi: 10.1117/12.968104.
- Rafael Izazaga, Esperanza Carrasco, Andrea Hidalgo, and 10 other authors. WEAVE spectrograph cameras: the polishing of the spherical lenses. In *Proc. SPIE*, volume 10706 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107063J, Jul 2018. doi: 10.1117/12.2314124.
- Aletta Henriette Jacobs. *Over localisatie van physiologische en pathologische verschijnselen in de groote hersenen*. PhD thesis, Rijksuniversiteit Groningen, 3 1879. <https://hdl.handle.net/11370/dea668cc-b108-4d28-b6db-20468380b8f1>.
- Shoko Jin et al. The WEAVE Survey: Design, Overview and Simulated Implementation. in preparation.
- Masahiko Kimura, Toshinori Maihara, Fumihide Iwamuro, and 40 other authors. Fibre Multi-Object Spectrograph (FMOS) for the Subaru Telescope. *Publications of the Astronomical Society of Japan*, 62:1135–1147, Oct 2010. doi: 10.1093/pasj/62.5.1135.
- Ian Lewis, Matthew Brock, James Gilbert, Andy Ridings, and Johan Pragt. Fibre Positioner System FDR Document. *Internal Document WEAVE-POS-009 version 1.5*, 2017. Date 24-04-2017.
- Ian J. Lewis, R. M. Sharples, I. R. Parry, and 4 other authors. Autofib-2: commissioning results of a robotic multiobject fiber system for the William Herschel Telescope. In *Optical Telescopes of Today and Tomorrow*, volume 2871 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 1318–1324, March 1997.
- Ian J. Lewis, R. D. Cannon, K. Taylor, and 20 other authors. The Anglo-Australian Observatory 2dF facility. *MNRAS*, 333:279–299, June 2002. doi: 10.1046/j.1365-8711.2002.05333.x. <https://dx.doi.org/10.1046/j.1365-8711.2002.05333.x>.



- Ian J. Lewis, Gavin B. Dalton, Matthew Brock, James Gilbert, and 5 other authors. Fibre positioning concept for the WEAVE spectrograph at the WHT. In *Proc. SPIE*, volume 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 914734, Jul 2014. doi: 10.1117/12.2055883.
- Emilie Lhomé, Tibor Agócs, Don Carlos Abrams, Kevin M. Dee, and 16 other authors. Manufacturing process for the WEAVE prime focus corrector optics for the 4.2m William Herschel Telescope. In *Proc. SPIE*, volume 9912 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 991245, Jul 2016. doi: 10.1117/12.2230648.
- Kevin Middleton. Positioning Error Budget at Prime Focus. *Internal Document WEAVE-SYS-017 version 2.7*, 2018. Date 28-02-2018.
- Kevin Middleton and Kevin Dee. Two Degree Prime Focus / Fibre System / Fibre Positioner Interface Control Document. *Internal Document WEAVE-ICD-009 version 2.5*, 2018. Date 01-03-2018.
- Randall Munroe. xkcd, 2006-2019. <https://xkcd.com>.
- Isaac Newton. A letter of Mr. Isaac Newton, Professor of the Mathematicks in the University of Cambridge; containing his new theory about light and colors: sent by the author to the publisher from Cambridge. *Philosophical Transactions, R. Soc.*, 6(80):3075–3087, 1671/72. <https://doi.org/10.1098/rstl.1671.0072>.
- Russell E. Owen, Walter A. Siegmund, Siriluk Limmongkol, and Charles L. Hull. Fiber feed for the SDSS spectrograph. In David L. Crawford and Eric R. Craine, editors, *Proc. SPIE*, volume 2198 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 110–114, Jun 1994. doi: 10.1117/12.176689.
- I. Parry and R. M. Sharples. AUTO FIB Current Status. In Samuel C. Barden, editor, *Fiber Optics in Astronomy*, volume 3 of *Astronomical Society of the Pacific Conference Series*, page 93, Jan 1988.
- I. Parry, R. Ellis, and P. Gray. Astronomy goes automatic. *New Scientist*, 110:36–39, Apr 1986.
- I. R. Parry, I. J. Lewis, R. M. Sharples, G. N. Dodsworth, J. Webster, D. W. Gellatly, L. R. Jones, and F. G. Watson. Autofib–2: an automated fiber positioner for the prime focus of the William Herschel Telescope. In *Instrumentation in Astronomy*



VIII, volume 2198 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 125–133, June 1994. doi: 10.1117/12.176708. <https://dx.doi.org/10.1117/12.176708>.

M. M. Pieri, S. Bonoli, J. Chaves-Montero, I. Pâris, M. Fumagalli, J. S. Bolton, M. Viel, P. Noterdaeme, J. Miralda-Escudé, N. G. Busca, H. Rahmani, C. Peroux, A. Font-Ribera, and S. C. Trager. WEAVE-QSO: A Massive Intergalactic Medium Survey for the William Herschel Telescope. In *SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, pages 259–266, Dec 2016.

Python Software Foundation. Python Language Reference, version 2.7.13. Available at <http://www.python.org>, however version 2.7 is now deprecated.

Scott M. Ransom. The pySLALIB Package, 2010. <https://github.com/scottransom/pyslalib/tree/fcb0650a140a8002cc6c0e8918c3e4c6fe3f8e01>.

Kevin Rogers, Remko Stuik, Iain A. Steele, and 20 other authors. The design of the WEAVE spectrograph. In *Proc. SPIE*, volume 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91476H, Jul 2014. doi: 10.1117/12.2055860.

Huub Röttgering et al. LOFAR and APERTIF Surveys of the Radio Sky: Probing Shocks and Magnetic Fields in Galaxy Clusters. *Journal of Astrophysics and Astronomy*, 32:557–566, Dec 2011. doi: 10.1007/s12036-011-9129-x.

F. Sayède, Y. Younes, G. Fasola, S. Dorent, and 15 other authors. First results of tests on the WEAVE fibres. In *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II*, volume 9912, page 991220, July 2016. doi: 10.1117/12.2231170. <https://dx.doi.org/10.1117/12.2231170>.

Frédéric Sayède, Scott Trager, Youssef Younes, and Isabelle Guinouard. Fibre System Design Report. *Internal Document WEAVE-FIB-005 version 0.2*, 2015. Date 25-09-2015.

Ellen Schallig, Ian J. Lewis, James Gilbert, and 9 other authors. Developments in fibre-positioning technology for the WEAVE instrument at the William Herschel Telescope. *Proc. SPIE*, 9908:99087U–99087U–11, 2016. doi: 10.1117/12.2231626. <https://dx.doi.org/10.1117/12.2231626>.



- Ellen Schallig, Ian J. Lewis, Gavin Dalton, and 10 other authors. First lab results of the WEAVE fibre positioner system. *Proc. SPIE*, 10702:290, 07 2018. doi: 10.1117/12.2312703. <https://dx.doi.org/10.1117/12.2312703>.
- D. J. B. Smith, P. N. Best, K. J. Duncan, N. A. Hatch, M. J. Jarvis, H. J. A. Röttgering, C. J. Simpson, J. P. Stott, and 22 other authors. The WEAVE-LOFAR Survey. In *SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, pages 271–280, Dec 2016.
- D. L. Terrett, I. J. Lewis, G. Dalton, D. C. Abrams, J. A. L. Aguerri, P. Bonifacio, K. Middleton, and S. C. Trager. Fibre positioning algorithms for the WEAVE spectrograph. In *Software and Cyberinfrastructure for Astronomy III*, volume 9152 of *Proc. SPIE*, page 91520P, July 2014. doi: 10.1117/12.2055844.
- Albert Tomás, Manuel Canchado, Joan Manel Casalta, Francesc Dalmases, Oscar Maroto, Carlos Martín-Nuño, Antonio Romero, and 13 other authors. The WEAVE prime focus correction: from design to integration. In *Proc. SPIE*, volume 10706 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070606, Jul 2018. doi: 10.1117/12.2313718.
- Scott Trager and the Science Teams. The WEAVE Science Case. *Internal Document WEAVE-SCI-002 version 2.0*, 2015. Date 30-09-2015.
- Sarah E. Tuttle, Gary J. Hill, Brian L. Vattiat, and 26 other authors. VIRUS early installation and commissioning. In *Proc. SPIE*, volume 9908 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99081I, Aug 2016. doi: 10.1117/12.2231253.
- Stéfan van der Walt, S. Chris Colbert, and Gaël Varoquaux. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science and Engineering*, 13(2):22–30, Mar 2011. doi: 10.1109/MCSE.2011.37.
- M. P. van Haarlem et al. LOFAR: The LOW-Frequency ARray. *A&A*, 556:A2, Aug 2013. doi: 10.1051/0004-6361/201220873.
- M. A. W. Verheijen, T. A. Oosterloo, W. A. van Cappellen, L. Bakker, M. V. Ivashina, and J. M. van der Hulst. Apertif, a focal plane array for the WSRT. In Robert Minchin and Emmanuel Momjian, editors, *The Evolution of Galaxies Through the Neutral Hydrogen Window*, volume 1035 of *American Institute of Physics Conference Series*, pages 265–271, Aug 2008. doi: 10.1063/1.2973599.



P. T. Wallace. The SLALIB Library. In *Astronomical Data Analysis Software and Systems III*, volume 61 of *Astronomical Society of the Pacific Conference Series*, page 481, 1994.

Donald G. York et al. The Sloan Digital Sky Survey: Technical Summary. *The Astronomical Journal*, 120(3):1579–1587, Sep 2000. doi: 10.1086/301513.

