

New developments in
robotic fibre positioning for
astronomical multi-object spectroscopy



James Gilbert
Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2016

To Alice,

for all the pep talks.

Acknowledgements

I'd like to formally thank my supervisor, Gavin Dalton, and co-supervisor, Ian Lewis, for their support throughout the last 2.5 years. I'd specifically like to thank Gavin for always being so positive and accommodating.

I wish to extend my gratitude to everyone in and around Oxford who has assisted in making this outcome possible, particularly: Matthew Brock, Garret Cotter, Vanessa Ferraro-Wood, Peter Hastings, Ashling Morris, Mathew Newport, Jeff Niblett, Paul Pattinson, Ellen Schallig, Andy Stokes, Mike Tacon, Angela Taylor, and David Terrett.

I'd like to show my sincere appreciation for the cooperation of those at the Australian Astronomical Observatory throughout the latter stages of my research: Norrie Bennie, David Brown, Rebecca Brown, Jurek Brzeski, Tony Farrell, Garry Kitley, Urs Klauser, Slavko Mali, Rolf Muller, Vijay Nichani, Naveen Pai, Andy Sheinis, Scott Smedley, Nick Staszak, and Lew Waller. Special thanks go to Jon Lawrence, for his perpetual willingness, and to Will Saunders, for the chats and injections of enthusiasm.

My time at Oxford has been funded by the Science and Technology Facilities Council (STFC), with additional funds awarded by the Institution of Engineering and Technology (IET) in the form of the 2014 Leslie H Paddle Scholarship. I would therefore like to thank the STFC, the IET, and Mr Leslie H Paddle for this assistance.

Last but by no means least, I would like to thank my family. My heartfelt thanks go to my wife Alice, whose significance throughout this journey is too great for only a few words, but without whom I could not have succeeded. Also to my mum and dad, whose love, support and unshakable faith in me have always been so valuable.

Abstract

This thesis presents two bodies of work, both related to automated optical fibre positioning in highly-multiplexed astronomical instrumentation.

Part one concerns the development and prototyping of a pick-and-place fibre positioner for WEAVE, a new 1000-fibre spectrograph for the 4.2 m William Herschel Telescope. This research has helped shape the positioner's final design and has included development of three novel techniques to enhance predicted performance: i) a scheme for automatically calibrating the fibre measurement system; ii) an algorithm for finding the centre of an imaged fibre with superior accuracy and speed; and iii) a robot motion control scheme to greatly reduce field reconfiguration time. Following these developments, and with rigorous testing of the proposed industrial robotic components, it is found that the performance of the positioner is likely to meet WEAVE's scientific requirements.

Part two of this thesis presents a new design of piezoelectric motor for 'tilting spine' fibre positioners. Tilting spine technology already allows simultaneous positioning of fibres in high-density fields, but has two significant drawbacks: i) fibres are tilted, resulting in optical losses; and ii) actuators require signal amplitudes of >100 V, placing performance-degrading limitations on the control system. The new motor has a more efficient mechanism and a drive voltage of just ± 9 V, allowing new control approaches that vastly improve the technology's capabilities. Prototyping has shown a closed-loop positioning accuracy of <2.8 μm , lower than any other published technique. A 50% reduction in optical loss is shown to be feasible by trading some accuracy for a spine length increase. It follows that any survey using this motor would have a higher signal-to-noise ratio and/or would complete faster than with the existing technology.

Contents

Introduction	3
1 Spectroscopy with optical telescopes	3
1.1 Present-day optical telescopes	3
1.2 The spectrograph	5
1.3 Multi-object spectroscopy using optical fibres	5
1.4 Section summary	8
2 Fibre positioning systems	9
2.1 The need for fibre positioning	9
2.2 The need for automation	9
2.3 Robotic high-multiplex fibre positioning technologies	11
2.4 Section summary	19
I Prototyping of a new robotic fibre positioning system for WEAVE on the William Herschel Telescope	21
1 Introduction to WEAVE	23
1.1 WEAVE and the William Herschel Telescope	24
1.2 WEAVE science goals	24
1.3 The anatomy of WEAVE	26
1.4 WEAVE fibre positioner design	27
1.5 Goals of this research	32
1.6 Section summary	35
2 Commissioning of prototype positioner robot	37
2.1 Hardware assembly	37
2.2 Development of a PLC-based control interface	38

2.3	Section summary	41
3	WEAVE fibre placement metrology	43
3.1	Fibre position measurement principles	43
3.2	The measurement cameras	44
3.3	Design of a new measurement camera	46
3.4	Design of an in-situ camera calibration scheme	50
3.5	Verification of fibre back-illumination scheme	59
3.6	Analysis of techniques for finding a fibre's centre	66
3.7	Proof of concept for passive field plate fiducials	82
3.8	Measuring axis position uncertainty	85
3.9	Estimating total measurement error	86
3.10	Section summary	89
4	Verification of commercial motor performance	91
4.1	Repeatability testing	91
4.2	Improvements to positioner timings	93
4.3	Prototype robot limitations	98
4.4	Section summary	98
5	Final analysis of fibre positioning performance	99
5.1	Field reconfiguration speed	99
5.2	Measurement system accuracy	100
5.3	Robot repeatability and error tolerance	100
5.4	Section summary	101
6	Conclusions and prospects	103
 II Development of a new low-voltage piezoelectric motor for tilting spine fibre positioners		105
1	Introduction to tilting spine technology	107
1.1	Nomenclature	107
1.2	The anatomy of a tilting spine	109
1.3	Actuation principle	112
1.4	Spine control	114
1.5	Desirable improvements to tilting spine technology	117

1.6	The 4MOST instrument	120
1.7	Competing technologies	123
1.8	Goals of this research	124
1.9	The ideal tilting spine positioner	126
1.10	Section summary	126
2	Achieving low-voltage drive of tilting spines	129
2.1	The use of piezo actuators in micro-positioning	129
2.2	Design approaches	131
2.3	Section summary	136
3	Design and construction of a laboratory test rig for tilting spines	137
3.1	Test rig requirements	138
3.2	Test rig design	139
3.3	Section summary	144
4	New motor designs	147
4.1	The three-finger motor	147
4.2	The shear plate motor	148
4.3	The multilayer plate motor	150
4.4	Section summary	159
5	Development of control hardware and software	161
5.1	Exploring alternative control architectures	161
5.2	‘Blended’ moves for motion in any direction	164
5.3	Updating the control electronics design	166
5.4	Distributed control and ‘stand-alone’ spines	169
5.5	Power dissipation estimates	173
5.6	Software control principle	175
5.7	Adaptive calibration of spine performance	181
5.8	Choosing optimal waveform parameters	188
5.9	Section summary	189
6	The multilayer plate motor: Refinement, testing and analysis	191
6.1	Prototype motors and spines	191
6.2	Wear resistance	193
6.3	Closed-loop positioning performance	195
6.4	Gravity vector response	202

6.5	Further characterisation	203
6.6	Open-loop repeatability	206
6.7	Increasing spine length	210
6.8	Removal of ground connection	211
6.9	Longevity	212
6.10	Section summary	214
7	Conclusions and prospects	217
	Appendix	223
A	Piezo actuator background	223
A.1	Traditional materials and manufacturing processes	223
A.2	Typical voltages and displacements	224
A.3	Actuator stacks	224
A.4	Piezo actuation styles	224
	References	227

List of Acronyms

2dF	2-degree Field (instrument)
2dFGRS	2dF Galaxy Redshift Survey
4MOST	4-metre Multi-Object Spectroscopic Telescope
6dF	6-degree Field (instrument)
6dFGS	6dF Galaxy Survey
AAO	Australian Astronomical Observatory
AAT	Anglo-Australian Telescope
AESOP	Australian ESO Positioner
APERTIF	Aperture Tile in Focus (instrument)
API	application programming interface
ASCII	American Standard Code for Information Interchange (character encoding)
BOSS	Baryon Oscillation Spectroscopic Survey
CAN	Controller Area Network
CMM	coordinate measuring machine
CNC	computer numerical control
COTS	commercial off-the-shelf

CPU	central processing unit
DAC	digital-to-analogue converter
DC	direct current
DESI	Dark Energy Spectroscopic Instrument
E-ELT	European Extremely Large Telescope
ELT	extremely large telescope
ESO	European Southern Observatory
FIFO	first in, first out
FMOS	Fibre Multi-Object Spectrograph (instrument)
FoV	field of view
FPI	Focal Plane Imager (WEAVE/2dF subsystem)
FRD	focal ratio degradation
GMT	Giant Magellan Telescope
HV	high voltage
IC	integrated circuit
IFU	integral field unit
LCRS	Las Campanas Redshift Survey
LED	light-emitting diode
LOFAR	Low-Frequency Array for Radio Astronomy (telescope)
LV	low voltage

MLP	multilayer plate
MOS	multi-object spectroscopy
NOVA	Netherlands Research School for Astronomy
OCS	Observatory Control System
OLED	organic light-emitting diode
PCB	printed circuit board
PFC	prime focus corrector
PLC	programmable logic controller
PMN–PT	lead magnesium niobate–lead titanate ($\text{Pb}[\text{Mg}_{\frac{1}{3}}\text{Nb}_{\frac{2}{3}}]\text{O}_3\text{--PbTiO}_3$)
PSF	point spread function
PWM	pulse-width modulation
PZT	lead zirconate titanate ($\text{Pb}[\text{Zr}_x\text{Ti}_{1-x}]\text{O}_3$)
RAM	random-access memory
RMS	root mean square
RoHS	Restriction of Hazardous Substances (directive)
RSS	root sum of squares
RVS	Radial Velocity Spectrometer (Gaia instrument)
SDSS	Sloan Digital Sky Survey
SNR	signal-to-noise ratio
TCP/IP	Transmission Control Protocol/Internet Protocol
TMT	Thirty Meter Telescope

TRL	Technology Readiness Level
UKST	UK Schmidt Telescope
USB	Universal Serial Bus
VISTA	Visible and Infrared Survey Telescope for Astronomy
VLT	Very Large Telescope
VSSOP	Very Thin Shrink Small-Outline Package
WEAVE	WHT Enhanced Area Velocity Explorer
WHT	William Herschel Telescope

Introduction

1 Spectroscopy with optical telescopes

Spectroscopy is arguably the most useful technique in the astronomer’s toolbox. Its power lies in our ability to infer some of the most fundamental properties of light-emitting and light-absorbing objects from their spectra alone, such as their chemical composition (from spectral lines) and radial velocity (from spectral Doppler shifts).

Spectroscopy plays an important role in practically every facet of astronomical research, in revealing the nature and motion of stars and galaxies, mapping the structure of the universe, and even detecting the presence of extrasolar planets.

This section of the introduction aims to familiarise the reader with modern day telescope facilities and describe the emergence and evolution of multi-object fibre-fed spectroscopy as a major observational method.

1.1 Present-day optical telescopes

By and large, professional optical telescopes are built as general-purpose facilities. They are designed to be as versatile and flexible as possible, their basic function being to deliver photons to a diverse range of exchangeable scientific instruments.

The anatomy of a typical reflecting telescope is shown in Figure 1.1. Such telescopes are centred around their large primary mirror. The telescope’s on-sky field of view is set by the (often interchangeable) optics after the primary mirror, so that somewhere this field is focussed onto a virtual surface. This surface—sometimes flat, sometimes curved, but in any case misleadingly called the telescope ‘focal plane’¹—presents an image of the sky with which an instrument can do its specific job.

Reflecting telescopes tend to be classed by the diameter of their primary mirror, with larger mirrors collecting more light than smaller ones. Table 1.1 shows a gross generalisation of the size of the most advanced telescopes through recent decades, forming the widely-accepted notion of telescope ‘classes’.

¹‘Focal surface’ will be used throughout this thesis, in the interest of not being misleading.

INTRODUCTION




Figure 1.1: Large reflecting telescopes usually have a range of focus options to match different instrument requirements. Common instrument sites are at prime focus (blue path), Cassegrain focus (red path), and Nasmyth focus (yellow path). The telescope’s field of view is set by the optics between the primary mirror and the instrument, and can be as wide as a few whole degrees on the sky (several moon-widths). *Star field image by ESO.*

Table 1.1: Recent decades have seen a rapid increase in the size of the largest available telescopes, progressing from the ‘4 m class’ facilities of the 1970s to the ‘30 m class’ facilities (ELTs) currently under construction.

Designation	Mirror diameter	Decade(s)	Light-collecting area
4 m class	~4 m	1970s, 1980s	~10 m ²
8 m class	~8 m	1990s, 2000s	~50 m ²
ELTs	~30 m	2020s	500–1000 m ²

The next generation of 30 m class telescopes, more often referred to as the Extremely Large Telescopes (ELTs), will offer astronomers unprecedented light-collecting power and optical resolution. Three such telescopes are currently at a mature stage of development: the 25 m Giant Magellan Telescope (GMT), the 30 m Thirty Meter Telescope (TMT), and the 39 m European Extremely Large Telescope (E-ELT).

While larger telescopes will allow us to probe deeper into the faint universe than ever before, current- and past-generation facilities will not become irrelevant. Indeed, many 4 m class telescopes continue to be the workhorses of astronomical research, hosting some of today’s newest and most important instruments. A substantial part of this thesis is devoted to one project with this very purpose.

1.2 The spectrograph

In its simplest form, an astronomical spectrograph comprises a dispersive element and a camera (Figure 1.2) and sits behind the telescope's focal surface. Light from an astronomical object is dispersed into its component wavelengths and this information is extracted from a digitised image of the spectrum on an image detector.

A mask can be used at the focal surface to sample only a thin slice of an object. This 'slit' discards almost all of the image information in the direction of the spectral dispersion, which provides a cleaner spectrum for a particular section of the object.




Figure 1.2: A dispersive element such as a prism or a diffraction grating is at the heart of an astronomical spectrograph. Incoming starlight is spread into a spectrum and then captured by a camera. Slit masks are often placed at the focal surface to remove unwanted beams.

1.3 Multi-object spectroscopy using optical fibres

Multi-object spectroscopy (MOS) enables simultaneous spectroscopic observations of many objects, greatly increasing the data output of a single spectrograph. Long camera exposure times (often >1 h) make this all the more important, especially where faint objects can only be observed on the darkest nights of the lunar cycle. Moreover, large sky surveys are only feasible if spectra can be delivered at a constantly high rate. Parallelisation has therefore become the norm for many major spectroscopic instruments.

INTRODUCTION

While several methods of obtaining multiple spectra from a single observation have been used throughout history², no instruments have been as successful at multiplexing visible-wavelength spectroscopic observations as those fed with optical fibres.

Fibre-fed multi-object spectroscopy uses optical fibres to sample the light from individual objects across a telescope's focal surface and guide it to one or more spectrographs, as illustrated in Figure 1.3. The number of objects that can be observed simultaneously is commonly called the 'multiplex' of the instrument, and modern-day fibre MOS can support multiplexes in the thousands.

The fibre core diameter (commonly of the order 100 μm) can be chosen to match the optical format of the telescope, the type of target, and the quality of seeing at the observatory. The major advantages of fibre-fed multi-object spectroscopy are outlined below.




Figure 1.3: Optical fibres can be used to 'pick off' the light from individual objects within a telescope's field of view. These fibres can be arranged compactly at the spectrograph entrance, making full use of the available detector area. Fibres can be run for tens of metres without serious light loss, meaning that spectrographs do not have to be mounted on the telescope itself.

²'Slitless', 'long-slit' and 'multi-slit' spectrographs use physical masks to select targets across the field of view. The number of observable objects is limited by the size of the image detector and by spectral overlap. Contemporary multi-slit spectrographs are specialist cryogenic instruments for observing near infrared wavelengths, and are limited to around 50 objects [1][2].

INTRODUCTION

- **Very high multiplexes:** Optical fibres are very fine, meaning that the practical multiplex limit is not the number of objects you can sample but simply the size of your spectrograph(s). Modern day fibre-fed instruments have multiplexes in the thousands.
- **Efficient use of image detector real estate:** Fibres can be efficiently aligned in the spectrograph so that there is minimal wastage across the area of the image detector.
- **The option of an off-telescope spectrograph:** Spectrographs can be housed in purpose-built rooms with a fibre cable running from the telescope focus, eliminating the mass and size restrictions of an on-telescope instrument.

A more recent evolution of fibre-fed astronomy is to bundle several single fibres together in a closely-packed array while still obtaining separate fibre spectra. This allows *spatial* as well as spectral sampling of galaxies and other extended sources, allowing us to study the properties of different areas of the object. These fibre bundles are called Integral Field Units (IFUs) and can comprise as few as seven fibres, or many hundreds. Multi-object IFU instruments are currently gaining considerable momentum in the community [3][4][5].

1.3.1 The evolution of fibre-fed spectroscopy

The first multi-object observations with optical fibres were made in 1979 at the Steward Observatory, soon resulting in the 37-fibre MEDUSA instrument of 1980 [6]. The use of fibres in this way was quickly adopted across the world, with publications reviewing the history of the technique appearing as early as 1988 [7].

Throughout the late 1990s, fibre-fed MOS rapidly progressed in sophistication to become one of the major observational strategies of our time. Multiplexes increased, fields of view became wider, and the first massive ‘sky surveys’ began. Table 1.2 shows a cherry-picked selection of galaxy surveys made possible by highly-multiplexed fibre spectrographs. These large datasets have enabled some of the most important cosmological discoveries about the large-scale structure of the universe.

But there is still much more to see. Current and future surveys are routinely targeting many millions of stars and galaxies in an attempt to understand the birth, life and death of our own galaxy, our neighbouring galaxies, and of the entire universe.

To answer the biggest questions we must use our best technology.

INTRODUCTION

Table 1.2: A hand-picked selection of influential galaxy redshift surveys that relied on high-multiplex fibre spectroscopy, showing the number of spectra recorded, the approximate duration, the telescope size and field of view (FoV), and the instrument multiplex (mux.).

Survey name	Spectra	Years	Telescope	FoV	Instrument	Mux.
LCRS [8]	26 000	1991–1996	du Pont 2.5 m	2°	—	50
2dFGRS [9]	230 000	1997–2002	AAT 3.9 m	2°	2dF	400
6dFGS [10]	125 000	2001–2009	UKST 1.2 m	6°	6dF	150
SDSS [11]	900 000	2000–2008	Sloan 2.5 m	3°	(Dedicated)	640

1.4 Section summary

In this section we have reviewed the anatomy of modern optical research telescopes and their increase in light-collecting power from the 4 m class facilities of the 1970s to the 30 m class ‘ELTs’ currently under construction. Despite this advance, the basic design of reflecting telescopes seldom departs from a classic format, often with selectable foci to accommodate a range of interchangeable instruments.

The concept of multi-object spectroscopy using optical fibres has been introduced, with examples of historical surveys of up to a million objects that were only possible with this technique. Having made the connection between technological advancements and scientific output, it should now be clear that the future of astronomical research will be shaped by our ability to produce cutting edge instrumentation.

2 Fibre positioning systems

This section of the introduction focusses on the technical aspects of multi-object spectroscopy and the engineering challenges of reconfiguring fibre fields.

2.1 The need for fibre positioning

One of the major technical hurdles presented by multi-object fibre instruments is how their fibres are physically placed on targets before an observation begins. Fibres also have to be reconfigured afterwards, ready for the following observation.

High quality seeing conditions at most observatories mean that there is rarely much room for positioning error on the telescope focal surface. Accuracy requirements for fibre placement on point sources (stars) are commonly of the order 10 μm , otherwise a significant fraction of the light will miss the fibre core and be wasted.

Modern-day fibre numbers in the hundreds, or even thousands, make the system used for fibre positioning absolutely critical to the success of the instrument. These systems have two main functions: i) they must provide a way to place fibres at pre-defined locations¹ across the field of the telescope, with minimal positional error; and ii) they must minimise the time taken to switch from one field to the next.

2.2 The need for automation

Within five years of the first fibre-fed spectrographs going online, automated fibre positioning systems were being developed with a view to reducing the overhead associated with manual methods. But both approaches have their place, even today.

¹Target selection depends on the motives of the observing programme and can be automatic or specified by an astronomer. Translation of object sky coordinates into physical locations across the telescope field depends on the time and place of the observation and the optics of the telescope; these calculations are usually done by software.

2.2.1 Manual ‘plug plates’

Fibre ‘aperture plates’—later becoming ‘plug plates’—have been used with success since the first fibre-fed multi-object spectrograph, MEDUSA, commissioned in 1980 [6]. Plug plates are flat, interchangeable discs that are mounted at the telescope’s focal plane. Fibres are mounted in holes that have been pre-drilled to a high accuracy. A number of plates can be drilled ahead of time and queued up for the entire night, the only remaining task being to unplug and re-plug the fibres between each observation.

It’s tempting to berate the labour-intensive nature of manual plug plates for large astronomical surveys of a million or more objects, but this approach has been used with undeniable success in various sky surveys at the Sloan telescope since 2000. The ongoing SDSS Baryon Oscillation Spectroscopic Survey (BOSS) [12] has a fibre multiplex of 1000, relying on human effort to fully populate each plate. The system maintains a high duty cycle by buffering the plates, meaning that one set of fibres is used for observing while the other is being prepared. Progressing to the next field is simply a matter of swapping the plates.

An inherent disadvantage with a plugged fibre system is its lack of flexibility; observations must be executed at the planned time else the holes in the plate will be at incorrect locations due to changing atmospheric refraction. The method also requires on-site machinery to drill the plates, and a trained technical staff to keep things running. For these reasons, plug plates tend to be favoured by organisations with an existing trained workforce, or simply for programmes of a short duration or low multiplex.

2.2.2 Automated systems

Automated positioning systems are able to reconfigure fibre fields with minimal human input. The first robotic fibre instrument, the MX Spectrometer [13], was commissioned at the Steward Observatory in 1985 by the same group who pioneered multi-object observations using manual aperture plates.

The major advantages offered by robotic positioning systems over manual plug plates are listed below. It’s also important to remember that we may not want our instrument to constantly march forth in the pursuit of huge datasets, but that small-scale surveys or ad hoc programmes may need to be accommodated too.

INTRODUCTION

- **Enhanced versatility:** There is no commitment to an observation schedule prepared in advance. The observer can take advantage of the current conditions to optimise the schedule, which is particularly important at sites with imperfect weather.
- **Reduced labour overheads:** The dependence on a trained technical workforce is removed, as is the need for a workshop to produce drilled mounting plates every day.
- **Operation during observations:** Exposure time limits can be increased by fibres tracking objects throughout a spectrograph exposure, if the chosen positioning technology allows this.

2.3 Robotic high-multiplex fibre positioning technologies

In the pursuit of efficient automated fibre positioning, a number of distinct techniques have emerged for high-multiplex instruments. These techniques each have their own strengths and weaknesses, having generally been optimised for a certain type of use.

2.3.1 The variation of technical requirements

Before reviewing fibre positioning technologies, we should first consider how the requirements of an instrument can vary. The factors that affect the top-level technical requirements of a fibre positioner have two main origins: i) the scientific goals of the instrument; and ii) the optical design of the telescope that will host the instrument.

The main evaluation criteria for fibre positioning are as follows, with positioning accuracy and reconfiguration speed being the two metrics most relevant to this thesis.

- **Positioning accuracy:** Fibre positioning technologies live or die by their ability to meet an instrument’s positioning accuracy requirements, with errors of a few tens of microns often meaning that light from a star misses a fibre entirely. An instrument’s tolerance of positioning errors will depend on the width of the intensity profile of an object at the focal surface; this is set by a combination of the quality of seeing at the observatory, the design of the telescope optics, and the telescope’s plate scale². Accuracy requirements can be just a few microns.

²The plate scale of a telescope gives the relationship between angular distances on the sky and real-world linear distances on the focal surface. Larger telescopes generally have physically larger focal surfaces for the same field of view, meaning that on-sky object separations translate to larger physical separations on-telescope.

INTRODUCTION

- **Positioning speed:** All positioning technologies aim to minimise field reconfiguration times as much as possible, but some observing strategies are more tolerant to this than others. For example, a faint object survey that routinely requires 2 h exposures will be fairly insensitive to a reconfiguration overhead of 5 min. Naturally, there is always a trade-off between the reconfiguration speed and the achieved positioning accuracy.
- **Target density:** Some positioners naturally lend themselves to higher or lower average fibre spacings. The overall density of candidate targets across the telescope's field of view will vary depending on the instrument's science case and the telescope's plate scale. A higher target density will benefit from the choice of a positioning technology that supports a high fibre density. Conversely, high fibre densities will be wasted on sparse fields.
- **Closest approach:** Some positioners are able to place neighbouring fibres closer together than others. The science case and telescope plate scale will set the minimum likely physical separation between targets, so choosing a technology with a closer approach will benefit observations with smaller separations.
- **Clustering ability:** Many positioners have a limited range of motion for each fibre, so bringing many fibres together can be problematic. Field homogeneity has a strong dependence on the phenomena we wish to understand. A programme to study the kinematics of star clusters, for instance, would be hurt by a positioner that is inherently restricted to highly uniform distributions.
- **Focal surface curvature:** Telescope optics may or may not provide a flat plane of focus, meaning that some fibre fields must accommodate an amount of focal surface curvature. Some technologies are not well suited to this.
- **Angular position control:** Some fibre positioners are able to provide angular position control of fibre units. The recent emergence of multi-object spectroscopy using IFU fibre bundles means that the angular orientation of IFUs on targets may be important, especially if asymmetric bundles are used.
- **Object tracking:** Some technologies are able to move fibres without disturbing an observation. Multi-object instruments with wide fields of view often have their maximum exposure time limited by the differential motion of objects throughout the exposure due to atmospheric refraction. Some observation programmes would therefore benefit from being able to track this object motion.

INTRODUCTION

- **Optical losses:** Some positioning technologies inherently cause optical losses due to misalignment or defocus of the fibre as it is moved. It's hard to think of a scenario where this is a desirable trait, although the other advantages associated with the technology may make the trade-off worthwhile.

We can now appreciate how the problem of multi-object fibre positioning can take many forms. And naturally, different problems call for different solutions.

2.3.2 Description of major positioning technologies

The first automated multi-object spectrograph, the MX Spectrometer, used a novel arrangement of actuated arms that surrounded the telescope's focal plane [13]. These arms could swing and extend to reach different points across the field, bearing a resemblance to fishermen standing around a small pond. But the multiplex of these early instruments was rather low by today's standards, and this pioneering application of robotics to fibre positioning has not scaled well as multiplexes have increased.

If we now consider only the technologies that remain compatible with the modern day and so-called 'massively-multiplexed' spectroscopic instruments (fibre numbers in the hundreds or thousands), then we see four main types: i) *pick-and-place* positioners; ii) *tilting spine* positioners; iii) *radial arm* positioners; and iv) *roaming* positioners. These are discussed in more detail below, along with their associated strengths and weaknesses.

Conspicuously absent from the following comparison of technologies is any notion of cost. While this is obviously a key driver in projects, true costs are often distorted by contractual arrangements between consortium members, by fluctuations in the currency market, or by the availability of student labour.

2.3.2.1 Pick-and-place positioners

Pick-and-place technology was the first to be used on an automated fibre-fed spectrograph with a multiplex exceeding 100. It was pioneered on the AUTOFIB instrument (1987; AAT; 64 fibre units) [14] and later grew in size and sophistication to become a fibre positioning standard. Notable examples include the Norris spectrograph (1990; Hale; 200 fibre units) [15], AUTOFIB-2 (1994; WHT; 150 fibre units) [16][17], 2dF (1997; AAT; 400 fibre units) [18], 6dF (2001; UKST; 150 fibre units) [19][20], and OzPoz (2003; VLT; 170 fibre units) [21].

INTRODUCTION

Pick-and-place positioners generally comprise a set of mechanical fingers attached to a robotic gantry above a flat magnetic ‘field plate’. The robot sequentially picks up fibres and places them back down at the required position. Each fibre is terminated with a small magnetic assembly that the robot can safely grip, and that holds the fibre in place on the field plate. A small right angle prism is glued to the end of the fibre to reflect the incoming starlight into and along it. The fibre runs to the edge of the field plate where it is kept under slight tension by a spring-loaded system of pulleys.

The 2dF robot (Figure 2.1) on the Anglo-Australian Telescope (AAT) remains a shining example of pick-and-place positioning technology. Over the last two decades it has stalwartly caught the light from hundreds of thousands of stars and galaxies, and is still going. It has 400 fibres and a wide 2° field of view focussed on a field plate the size of a large pizza. Fibres are placed with an accuracy of a few tens of microns.

The sequential nature of pick-and-place positioning means that field reconfigurations are rather slow. This can be mitigated by using a buffered system of two field plates, each with its own set of fibres, so that one plate can be configured while the other is observing. 2dF has a motorised field plate exchange mechanism for doing just this.




Photo by Ángel R. López-Sánchez

Figure 2.1: The 2dF positioner robot sequentially positions 400 fibres that originate from spring-loaded retractors around the perimeter of the field plate.

INTRODUCTION

Pick-and-place positioners place relatively few restrictions on where fibres are allowed to go, and therefore are a good solution for a wide range of use cases. They are particularly suitable for observing highly-clustered fields, and provide a close approach for nearby object pairs.

Part I of this thesis will focus on a brand new 1000-fibre pick-and-place positioner with a range of new features and enhancements over existing facilities.

2.3.2.2 Tilting spine positioners

‘Tilting spine’ technology, also known as ‘Echidna’ technology, was developed in the early 2000s by the Australian Astronomical Observatory (AAO) for the Fibre Multi-Object Spectrograph (FMOS) instrument on the Subaru telescope [22][23]. Unlike pick-and-place positioners, Echidna technology allows multiple fibres to be moved simultaneously.

Figure 2.2 shows the FMOS–Echidna positioner, which has 400 fibres spread across a square field about the size of a football³. Each fibre is held in a rigid tube called a ‘spine’, which is 140 mm long and can freely tilt about a ball joint at one of its ends. This ball is mounted on a small piezoelectric actuator that can move the spine in discrete angular steps, therefore displacing the spine tip. The actuators are arranged in a regular grid with a spacing (pitch) of 7.2 mm. Every point across the field can be reached by at least three fibres, as illustrated in Figure 2.3.




Photo by Masayuki Akiyama

Figure 2.2: FMOS–Echidna was the first instrument to use AAO tilting spine technology. Each fibre is held in a rigid tube (spine) that can be tilted by a few degrees in all directions. The positioner hosts 400 fibres on a pitch of 7.2 mm.

³Soccer ball.

INTRODUCTION

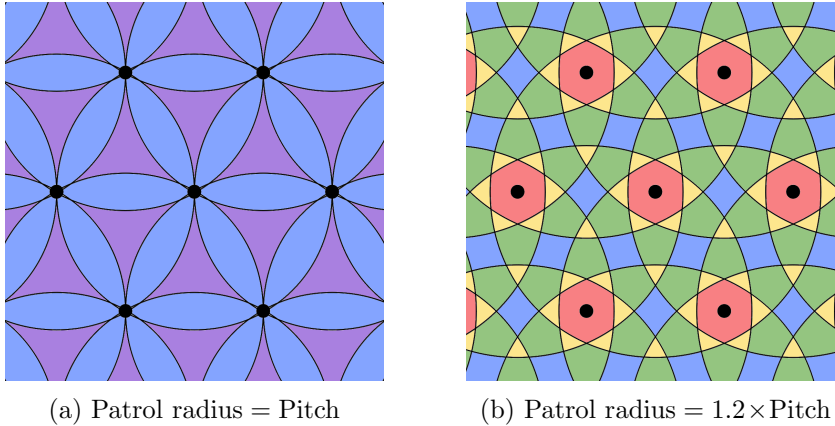


Figure 2.3: Tilting spine positioners have a large ‘patrol radius’. These diagrams show seven spines (black spots) viewed from their tips; any point in the field can be accessed by 3 (purple), 4 (blue), 5 (green), 6 (yellow) or 7 (red) different fibres. The FMOS–Echidna positioner has a spine patrol radius equal to the pitch (a), whereas more recent designs set the patrol radius slightly larger (b).

The Echidna concept has some unique strengths: it is able to support very high fibre densities while tolerating moderate target clustering, thanks to its generous overlap of the accessible ‘patrol areas’ of neighbouring spines; spines can also collide without damage, simplifying software routing algorithms. In addition to these, there are the natural advantages offered by a many-actuator positioner: all fibres can move simultaneously, making full field reconfigurations very fast, regardless of fibre count; and spines can also be moved during observations, albeit with limited accuracy.

An unavoidable weakness of tilting spines is the optical loss suffered due to non-telecetricity and defocus offsets when the spines are tilted, although for some instruments this may be compensated by increased target allocation yields [24][25].

The AAO is currently designing a much larger tilting spine positioner for the 4MOST instrument (VISTA) [26][24]. It will host ~ 2500 spines on a pitch of ~ 10 mm. All fibres are expected to be configurable to an RMS accuracy of $\sim 10 \mu\text{m}$ in ~ 2 min. Separate experiments by the AAO have recently shown that spine pitches smaller than 7 mm are plausible [27], making this technology the most compact available, by a considerable margin.

Part II of this thesis pursues further improvements to the current-generation tilting spine technology to meet the demands of future instruments, including the reduction of optical losses and an overall increase in accuracy and speed.

INTRODUCTION

2.3.2.3 Radial arm positioners

Another parallel positioning technique uses small robotic posts in a similar gridded arrangement to Echidna spines. Unlike Echidna spines, this breed of positioner does not significantly tilt or defocus the fibre. The principle was first put into practise on a large scale in the 4000-fibre LAMOST instrument, commissioned in the late 2000s [28]. Each independent LAMOST positioner has two miniature rotary motors in an eccentric-axis arrangement (Figure 2.4a). LAMOST’s fibre pitch is 25.6 mm and the fibre patrol radius is 16.5 mm, so every point across the field can be reached by at least one fibre (Figure 2.4b).

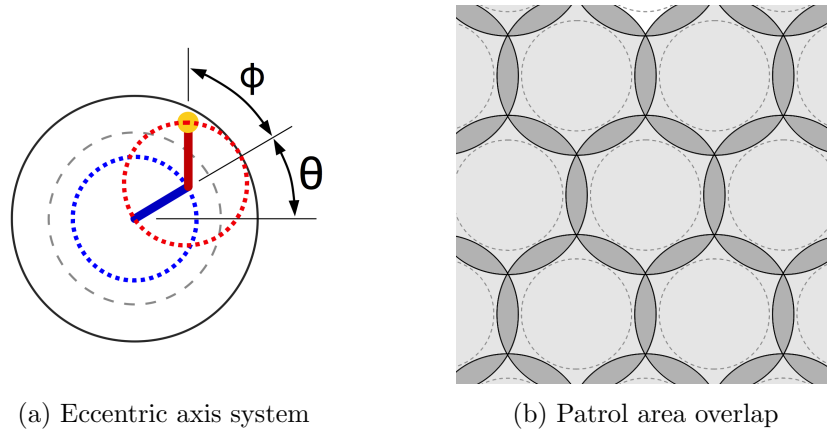


Figure 2.4: The most common format of radial arm positioner has two axes in a ‘theta-phi’ arrangement (a); the fibre (yellow spot) can access any point in a circular patrol area via a certain combination of ‘theta arm’ (blue) and ‘phi arm’ (red) angles. Points across the field (b) can be accessed by either one fibre (light grey areas) or two fibres (dark grey areas). Patrol area limits are drawn as solid black circles and the positioners’ footprints are drawn as dashed grey circles.

More recent radial arm designs include the Cobra positioner (SuMIRe/Subaru; 2394 fibres; 9.5 mm pitch) [29], the MOONS positioner (VLT; 1024 fibres; 50 mm pitch) [30], and numerous proposals for the DESI instrument (Mayall; 5000 fibres; 12 mm pitch) [31][32].

Radial arm positioners provide fast, accurate, and simultaneous fibre positioning. The fibre density is reasonably high although clustering ability is limited, making this technique best suited to consistently homogeneous fields.

INTRODUCTION

2.3.2.4 Roaming positioners

Roaming positioners are free to move about a field plate without significant restriction. In the context of fibre positioning, the only example of this style of device is the Starbug, conceived by the Australian Astronomical Observatory (AAO) in the early 2000s and later developed extensively in the early 2010s [33].

Starbugs (Figure 2.5) are miniature self-motile piezoelectric robots that are able to walk across a telescope’s field plate while carrying a variety of fibre-based payloads. They can have very wide ranges of motion, making them ideal for the large focal surfaces of next-generation telescopes. Indeed, Starbugs have been developed specifically for use on the Giant Magellan Telescope (GMT), which will have a field diameter exceeding 1 m [34]. This fibre positioning facility, called MANIFEST, is expected to be able to fully reconfigure the GMT field to an accuracy of $\sim 20\ \mu\text{m}$ in $\sim 2\ \text{min}$.



Image courtesy of the Australian Astronomical Observatory

Figure 2.5: Starbug fibre positioning robots can freely roam about a telescope’s field plate, carrying fibre payloads in their central cylindrical aperture. Current-generation Starbugs use a vacuum system to attach themselves to a transparent glass field plate, through which the starlight enters. The Starbugs pictured are $\sim 8\ \text{mm}$ in diameter.

Starbugs turn the traditional concept of fibre positioning on its head, attaching the positioner to the fibre rather than the fibre to the positioner. Their core philosophy is one of flexibility and modularity, enabling a range of tricks such as multi-instrument focal surfaces, object tracking, and angular positioning for IFU fibre bundles.

A generous patrol area means that Starbugs cope well with highly-clustered fields, although a rather large closest approach ($\sim 10\ \text{mm}$) limits their usefulness in telescopes with fine plate scales.

2.3.3 Summary of major positioning technologies

The strengths, weaknesses and idiosyncrasies of the four major fibre positioning techniques identified above are complex, interwoven, and practically impossible to compare in absolute terms. Table 2.1 is a sporting attempt at broadly summarising them, for no other reason than to show that their characteristics are unique.

Table 2.1: Broad-brush comparison of high-multiplex fibre positioning technologies against common selection criteria.

	Pick-and-place	Tilting spine	Radial arm	Roaming
Type	Serial	Parallel	Parallel	Parallel
Positioning accuracy	$\sim 20\ \mu\text{m}$	$\sim 10\ \mu\text{m}$	$\sim 5\ \mu\text{m}$	$\sim 10\ \mu\text{m}$
Positioning speed	Slowest ^a	Faster	Faster	Faster
Target density	Higher	Highest	Higher	Lowest
Closest approach	Better	Best	Worse	Worst
Clustering ability	Best	Better	Worst	Better
Focal surface curvature	Worse	Better	Better	Better
Angular control (IFUs)	Not possible	Not possible	Not possible	Possible
Object tracking	Not possible	Possible	Possible	Possible
Optical losses	Lower	Higher	Lower	Lower

^aThe sequential nature of pick-and-place positioning means that a buffered field plate exchange system and a second set of fibres are required to keep overheads low.

2.4 Section summary

In this section we have reviewed the motivation and history behind the major fibre positioning technologies for so-called ‘massively-multiplexed’ fibre-fed spectroscopic instruments, which can record spectra for hundreds or even thousands of objects in a single observation. A broad comparison of the four major types of positioner has shown their natural compatibility with different science goals and telescope formats.

Assuming that the scientific requirements of an instrument are satisfied, the suitability of a positioning system boils down to accuracy and speed: Fibre placement must be accurate so that photons aren’t lost, but fast so as to not waste valuable telescope time. As one might imagine, there is a natural trade-off between the two.

INTRODUCTION

This thesis will report on two separate areas of research in automated fibre positioning. The first body of work (reported in Part I) will focus on a new pick-and-place positioner with a fibre density never before seen on this type of system. The second body of work (reported in Part II) will focus on the development of a novel piezoelectric motor for tilting spine positioners, with characteristics that substantially improve on the technology's present limitations.

Part I

**Prototyping of a new robotic fibre
positioning system for WEAVE on
the William Herschel Telescope**

1 Introduction to WEAVE

The WHT Enhanced Area Velocity Explorer (WEAVE) is a next-generation wide-field fibre-fed multi-object spectroscopic facility for the prime focus of the 4.2 m William Herschel Telescope (WHT) [35][36]. Its planned format is shown in Figure 1.1. The WHT operates at Roque de los Muchachos Observatory on the island of La Palma in the Canary Islands.

This section describes the scientific motivations for the WEAVE project, the associated technical requirements for the instrument, and its overall design. The preliminary design of the fibre positioning system is described in detail, highlighting the need for extensive prototyping to retire the major technical risks relating to fibre positioning accuracy and speed. This was the primary aim of the research reported in this thesis chapter.




Image courtesy of Kevin Dee / Isaac Newton Group of Telescopes

Figure 1.1: WEAVE comprises a 960-fibre positioning system at the WHT prime focus and a high-resolution spectrograph at one of the WHT's Nasmyth platforms.

1.1 WEAVE and the William Herschel Telescope

WEAVE will turn the William Herschel Telescope (WHT) into a workhorse spectroscopic survey facility, capturing the visible-range spectra of more than a million astronomical objects in its first five years of service.

The instrument will feature 960 reconfigurable optical fibres, allowing 960 objects to be observed in a single 60 min exposure. The work reported in this thesis will focus on the automated reconfiguration of these optical fibres.

WEAVE will also host 20 deployable Integral Field Units (IFUs) for spatially-resolved spectroscopy. Table 1.1 shows a summary of the fibre types and multiplex for the entire instrument. In addition to these, a single large IFU will be available for observations of large diffuse objects such as nearby galaxies.

The WEAVE project has progressed through its final design stages and will shortly begin manufacture and integration. First light is expected by the end of 2017.

Table 1.1: WEAVE will have a variety of fibre types for different observing modes.

Fibre unit type	No. of fibres in unit	Unit field of view	No. of units
Single fibre (<i>MOS</i>)	1	1.3''	960
Mini IFU (<i>MIFU</i>)	37 (hex packed)	11'' × 12''	20
Large IFU (<i>LIFU</i>)	547 (hex packed)	90'' × 78''	1

1.2 WEAVE science goals

The main scientific motivation for WEAVE can be divided into two broad areas: i) follow-up of stellar targets from the Gaia satellite for galactic archaeology [37]; and ii) follow-up of galactic targets from the LOFAR and APERTIF cosmological radio surveys. These are explained below.

1.2.1 Gaia follow-up

The European Space Agency's Gaia satellite, launched in December 2013, is busy measuring the positions and proper motions of 10^9 stars in the Milky Way via precision astrometry [38][39]. In addition, Gaia's on-board Radial-Velocity Spectrometer (RVS) will measure radial velocities for those objects of magnitude $V < 17$.

PART I: THE WEAVE FIBRE POSITIONER

The limitations of Gaia’s instrumentation are such that more comprehensive spectroscopic follow-up by ground-based facilities is necessary. This is particularly true for elemental abundance measurements, which require broadband spectral coverage at good resolution. Furthermore, the accuracy of Gaia’s RVS instrument decreases at fainter magnitudes. An accuracy of 30 km s^{-1} is reported at $V \approx 17$, with performance decreasing sharply thereafter.

WEAVE aims to conduct a northern hemisphere spectroscopic census of 10^6 Gaia targets of magnitude $17 < V < 20$ at a resolution of $R \approx 5000$ across the visible wavelength range. In addition to chemical abundances, WEAVE will trace radial velocities with accuracies better than 3 km s^{-1} out to near the photometric limit of Gaia. The survey will target a uniform subset of Gaia objects, providing probability distribution information that can be applied to the whole catalogue.

WEAVE’s high-resolution mode ($R \approx 20\,000$) will allow accurate sampling of key elemental spectra, aiming to reveal chemical ‘fingerprints’ that link groups of stars with common formation histories. This will aid our understanding of how and when the Milky Way assembled its stellar mass.

1.2.2 Radio cosmology follow-up

WEAVE’s single fibre and IFU capability will allow spectroscopic follow-up of the Dutch-led LOFAR and APERTIF radio surveys. LOFAR’s wide-field continuum survey will detect 10^7 star-forming galaxies, but will not be able to infer their redshifts [40], meaning that this must be deferred to other instruments. WEAVE aims to determine the redshifts of LOFAR sources by detection of $\text{H}\alpha$ and [OII] emission lines for galaxies at $z < 1.3$ and detection of the $\text{Ly}\alpha$ line for galaxies at $z > 2$.

APERTIF, a radio focal plane array, will provide spatially-resolved kinematics from observations of the 21 cm HI line for 10^4 galaxies out to $z \approx 0.4$ [41]. With its deployable miniature IFUs, WEAVE will provide complementary kinematic information from the visible wavelength range, probing the relationship between star formation and neutral gas.

A single large IFU will be available on WEAVE to study the disks of nearby low surface brightness galaxies at intermediate spectral resolution, and determine the associated distributions of dark and luminous matter. Although WEAVE’s large IFU will be physically present on the fibre positioner structure, this thesis will focus only on the pick-and-place aspects of the positioner for multi-object observations.

1.3 The anatomy of WEAVE

The WEAVE instrument has three broad subsystems: i) the prime focus corrector; ii) the spectrograph; and iii) the fibre positioner. These were highlighted in Figure 1.1 (p. 23). The subsystems are summarised below, although the corrector and spectrograph are beyond the scope of this thesis.

1.3.1 The prime focus corrector

A new prime focus corrector (PFC) will be built for the WHT [42][43]. The PFC will provide a 2° field of view across a flat and near-telecentric focal surface 410 mm in diameter. Other optical specifications for the PFC are summarised in Table 1.2.

Table 1.2: Summary of the prime focus corrector design for WEAVE.

Parameter	Designed value
On-sky field of view	2.0°
Focal surface diameter	410 mm
Focal surface curvature	Flat
Telescope focal length	11.6 m
Speed	$f/2.75$

1.3.2 The spectrograph

A new spectrograph will be built for WEAVE [44], and has been designed according to the WEAVE science goals (see Section 1.2). It has a dual-arm configuration, together providing spectral coverage across the full visible range at $R \approx 5000$, or across narrower ranges at $R \approx 20\,000$. More details are shown in Table 1.3.

Table 1.3: Summary of the spectrograph design for WEAVE.

Parameter	Designed value	
	Low-resolution mode	High-resolution mode
Resolution	5000	20 000
Wavelength coverage (blue arm)	366–606 nm	404–465 nm 473–545 nm
Wavelength coverage (red arm)	579–959 nm	595–685 nm

1.3.3 The fibre positioner

The fibre positioner subsystem is responsible for the accurate automated placement of optical fibres (both single fibres and mini IFUs) on objects across the focal surface of the prime focus corrector. The positioner system also provides a focal plane imaging facility for field mapping and calibration of optical system distortions.

The primary requirements of the automated fibre placement system concern the accuracy and speed of fibre placement when reconfiguring a field. The WEAVE-related work reported in this thesis will focus almost entirely on these aspects of the anticipated fibre positioner performance.

Early on in the project, an accuracy requirement of $8\ \mu\text{m}$ RMS was set for the placement of fibres on WEAVE focal surface. A maximum field reconfiguration time of 60 min was also specified, although this is linked to the choice of fibre positioning technology (explained below).

1.4 WEAVE fibre positioner design

This section describes the motivation behind the early selection of ‘pick-and-place’ positioning technology for WEAVE, and how the WEAVE design inherits many features of the Anglo-Australian Telescope’s ‘2-degree Field’ (2dF) fibre positioning facility. Given this strong link, 2dF will be described in some detail. The WEAVE positioner design will then be presented with consideration of its differences to 2dF and how these unique challenges are addressed by the design.

1.4.1 Pick-and-place technology and the influence of 2dF

‘Pick-and-place’ positioning technology was chosen in the early stages of the WEAVE project due to its relatively low cost and technical risk. Of particular interest was the prospect of adopting commercial off-the-shelf (COTS) components, including industrial robotics products already used in high-precision applications.

Pick-and-place robotic positioning systems have a successful track record in the world of fibre-fed multi-object astronomy, examples being AUTOFIB (1987; AAT; 64 fibre units) [14], AUTOFIB-2 (1994; WHT; 150 fibre units) [16][17], 2dF (1997; AAT; 400 fibre units) [18], 6dF (2001; UKST; 150 fibre units) [19][20], and OzPoz (2003; VLT; 170 fibre units) [21]. While all of these are evolutions of the same principle, the AAT’s 2dF positioner is most akin to WEAVE given its prime focus location and similar telescope and field size.

PART I: THE WEAVE FIBRE POSITIONER




Image courtesy of Australian Astronomical Observatory

Figure 1.2: 2dF fibres are terminated with a small magnetic assembly called a ‘button’, allowing fibres to be picked up by the robot’s gripper. A right angle prism reflects incident starlight down the fibre core.

The 2dF robot, shown previously in Figure 2.1 (p. 14), sequentially positions fibres across a flat steel field plate using linear motors arranged in an orthogonal ‘H’ format. This (x,y) gantry structure allows an electro-mechanical gripper unit to access any point above the field plate. A third linear axis (z) provides up-down motion, while a rotary axis (θ) ensures the gripper unit is always aligned with the fibre being placed.

The positioner’s gripper unit has an integrated video camera system for accurately measuring the position of fibres before and after placing them. The fibres are illuminated during field reconfigurations so that their cores are detectable in the camera image for the purpose of image processing.

2dF’s 400 optical fibres originate from fibre ‘retractors’ mounted around the perimeter of the field plate. These retractors use a spring-loaded pulley system to keep every fibre under tension and therefore make their paths predictable.

Each fibre is terminated with a small magnetic assembly called a ‘button’ (Figure 1.2) measuring ~ 5 mm on its longest side. Each button incorporates a vertical handle or ‘vane’ with which the positioner robot picks up the fibre to move it. Light entering normal to the field plate is coupled into the fibre cores via small right angle prisms glued to their end faces.

The sequential nature of pick-and-place positioning means that field reconfigurations between successive observations would be unacceptably slow. To get around this, 2dF uses a pair of opposing field plates each with its own set of fibres. This means that when one field plate is in use, the other field plate is already being re-

PART I: THE WEAVE FIBRE POSITIONER




Image courtesy of the WEAVE team / University of Oxford

Figure 1.3: The WEAVE fibre positioner design adopts a dual-headed pick-and-place robot configuration, buffered dual field plates on a motorised ‘tumbler’, and is based heavily on COTS robotic components.

configured for the next observation. The field plates are quickly exchanged at the end of an observation by rotating a barrel-like ‘tumbler’, and the process repeats. It follows that 2dF is able to maintain a very high duty cycle, providing fields can be reconfigured within the duration of an average spectrograph exposure.

A further advantage of the 2dF fibre system is the ability to closely pack fibre buttons in order to observe dense or clustered fields. Fibre separations of <2.0 mm are achievable if buttons approach at large enough angles (prism-to-prism). This corresponds to an object separation of $\sim 30''$ on the WEAVE field plate.

1.4.2 WEAVE positioner preliminary design

As we have already seen, the WEAVE fibre positioner system will reside at the prime focus of the WHT, providing access to a 2° field of view via a robotic pick-and-place fibre positioning system with buffered exchangeable field plates.

Figure 1.3 shows a 3D model of the WEAVE positioner at its preliminary design stage [45]. While its similarity to the AAT’s 2dF positioner is clear, there are some important differences.

PART I: THE WEAVE FIBRE POSITIONER

The nature of pick-and-place technology is such that fibres cannot be moved once an observation commences, and therefore cannot track the movement of target objects caused by changing atmospheric refraction. The length of a useful exposure is consequently limited to ~ 60 min (for a nominal fibre diameter of $1.3''$) before targets at the edge of the field suffer unacceptable light loss. Hence this exposure time sets a goal for the field reconfiguration time, if we are to move to the next field with a minimal delay.

WEAVE's high multiplex of almost 1000 fibres means that it would be unrealistic for a single positioner robot to fulfil a 60 min reconfiguration time requirement. As a result, the WEAVE positioner will employ two pick-and-place robots, working simultaneously to reconfigure a single field plate.

In another significant departure from existing designs, all robotic axes and controllers in WEAVE will be commercially-available products from industry¹. The performance of modern day industrial equipment is such that its use in demanding applications such as ours is, for the first time, feasible. This is an important move in an area that has traditionally relied on building systems from scratch.

The high multiplex of WEAVE also means that the perimeter of the field plate is not large enough to accommodate every fibre. WEAVE's retractor units have therefore been designed in a three-tier format (Figure 1.4), allowing buttons to be parked on stepped platforms that resemble a sort of astronomical amphitheatre.




Image courtesy of the WEAVE team / University of Oxford

Figure 1.4: The WEAVE fibre retractor design features tiered platforms so that all fibres buttons can be parked around the perimeter of the field plate.

¹Schunk GmbH was chosen as the partner company for supplying WEAVE's robotic axes.

PART I: THE WEAVE FIBRE POSITIONER




Image courtesy of the WEAVE team / University of Oxford

Figure 1.5: WEAVE will have three types of fibre button: Type A is for use with guide bundles; Type B has a concave spherical upper surface to convert the beam speed for single science fibres; and Type C houses the miniature IFU fibre bundles.

WEAVE’s single-fibre buttons, while based on the 2dF button design, differ in two ways: i) their prisms have a concave face to convert the incoming beam from $f/2.75$ to $f/3.2$ for more efficient fibre injection; and ii) they incorporate a small tilt of the fibre to reduce the effect of slight non-telecentricity from the corrector.

The WEAVE fibre system will also feature 20 deployable miniature IFUs and 16 guide fibre bundles (eight per field plate), which will be positioned in the same way as single fibres. Figure 1.5 shows the design of all three WEAVE button types, although this is largely irrelevant to the operation and performance of the positioning system reported in this thesis.

A final part of the WEAVE positioner that should be described is the Focal Plane Imager (FPI), which is a twin camera system used to inspect the active field plate at the same time as imaging the sky. The primary purpose of the FPI is to determine the astrometric calibration of the focal plane, with a secondary use being a backup field acquisition tool. The FPI is mounted on a motorised gantry of the same type as the positioner robots. Again, this part of the positioner is of no direct relevance to the work reported in the following chapters.

1.5 Goals of this research

The remainder of this thesis chapter reports on work that was carried out during the final design stages of the WEAVE fibre positioner system, following its preliminary design review² and leading up to its final design review in January 2015.

The broad aim of this research was to retire the major technical risks associated with the design of the COTS-based pick-and-place system, better estimating its performance through laboratory prototyping and testing. More specifically, the work sought answers to two main questions:

1. Are the proposed off-the-shelf robotic axes capable of positioning WEAVE's optical fibres accurately enough to meet the scientific needs of the instrument?
2. If the accuracy is sufficient, then can this accuracy be achieved quickly enough to meet the required field reconfiguration time?

1.5.1 Fibre positioning accuracy requirement

The requirement for WEAVE's fibre positioning accuracy was set early in the project as part of the instrument development specification for the entire prime focus optical system. The requirement was set at $8.0\ \mu\text{m}$ (RMS), based on a total offset of $0.2''$ divided equally between the prime focus corrector and fibre positioner subsystems.

In the interest of completeness and diligence, the steps taken to arrive at the values above will be reproduced below and reviewed.

Since errors in positioning simply result in less light entering the fibre, there are two ways to look at the importance of fibre positioning accuracy: One way is to take the fact that exposure time is limited (see Section 1.5.2, below), and therefore see more accurate positioning as a way to extend the faintest magnitude limit of the instrument; or alternatively, we can take a nominal faint magnitude limit and see more accurate positioning as a way to increase the signal-to-noise ratio (SNR) of the data. Given that WEAVE's scientific motivation is so strongly interwoven with the Gaia satellite, the latter approach seems appropriate, with Gaia's faint magnitude limit (nominally $V=20$) used in the SNR calculations.

Taking the reported seeing conditions of the WHT site³, along with the expected performance of the prime focus corrector, WEAVE's known fibre core diameter, and

²The WEAVE fibre positioner PDR took place in March 2013, although the work reported here commenced in October.

³Site statistics currently accessible at <http://www.ing.iac.es/astronomy/observing/conditions/>

PART I: THE WEAVE FIBRE POSITIONER




Figure 1.6: Simulated^a resultant 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, solid lines being for a best-case (Gaussian) and worst-case (Moffat, $\beta=3$) PSF shapes. 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.

^aThese plots were produced using code originally written by Gavin Dalton, reused with gratitude.

PART I: THE WEAVE FIBRE POSITIONER

anticipated values for the optical system throughput and detector readout noise, we can simulate the number of photons entering the fibre as a function of the fibre’s radial distance from the centre of a point source with the relevant point spread function (PSF). This can be converted to SNR per resolution element in the final spectrum.

Figure 1.6 shows the simulated resultant SNR as a function of fibre offset for a $V = 20$ target with a 60 min exposure in WEAVE’s low-resolution mode. Each plot is for a different system throughput within the anticipated range. Each line on the plot represents different seeing conditions. The on-sky angular offset of the fibre has been converted to real distance units at the focal surface, by using the known focal length and speed of the telescope and PFC to calculate the plate scale ($56.2 \mu\text{m}/''$). The dashed vertical line shows the acceptable accuracy limit of $0.2''$ set by the project team, corresponding to a real world radial distance of $11.2 \mu\text{m}$.

The specified fibre positioning accuracy requirement of $8.0 \mu\text{m}$ was calculated by dividing the total allowable offset equally between the PFC and positioner subsystems and assuming that these would add in quadrature. The PFC portion of the budget includes target offsets caused by residual errors in the focal plane distortion map, changes in focal plane calibration with temperature, distortion map movement caused by various opto-mechanical systems, and field rotator tracking errors. Together these are expected to take up the entire $8.0 \mu\text{m}$ error budget allocated to the PFC.

Looking again at Figure 1.6 we can see that, while the $0.2''$ total error limit is a good target, a larger fibre offset would be far from catastrophic. Table 1.4 shows that the estimated SNR reduction with larger fibre positioning errors is actually quite modest; doubling the tolerated error weakens the estimated signal by only $\sim 6.5\%$ under the expected WHT site conditions. This will be important to remember once the achievable positioning accuracy of the positioner has been estimated later on in this thesis.

Table 1.4: Generously increasing the RMS fibre positioning error tolerance results in a relatively small decrease of the instrument’s estimated signal-to-noise ratio (SNR). These values are based on the data shown in Figure 1.6 and are for nominal observing conditions.

Tolerance change	RMS error	SNR change
Set requirement	$8.0 \mu\text{m}$	—
+50 %	$12.0 \mu\text{m}$	-3.9%
+100 %	$16.0 \mu\text{m}$	-6.5%

1.5.2 Fibre positioning speed requirement

The maximum allowable time for a full field reconfiguration is set by WEAVE's nominal exposure time, in order to make full use of the buffered field plate exchange system in the positioner.

WEAVE's exposure time is limited by target motion caused by residual atmospheric refraction. This effect is most pronounced at the edge of the field, where fibres suffer unacceptable light loss if an exposure is longer than ~ 60 min. This is based on previous calculations by the WEAVE team.

Because fibres cannot be moved during an exposure, the nominal exposure time for WEAVE has been set at 60 min, and therefore the maximum allowable time for a full field reconfiguration is also 60 min.

1.6 Section summary

In this section we have been introduced to the science goals of the WEAVE project and the instrument that will be built to achieve them. WEAVE will be a Gaia follow-up machine, providing accurate radial velocity measurements that complete the statistical picture for Gaia's census of 10^9 stars, and chemically labelling stellar groups in order to trace the mass assembly history of the Milky Way. WEAVE will also add galaxy redshift information to objects detected by the LOFAR and APERTIF radio surveys.

A detailed overview of the design of the WEAVE fibre positioner has been presented, and the choice of pick-and-place technology has been justified by WEAVE's similarities to the Australian 2dF facility, along with good prospects for using low-cost COTS components.

The flow-down of the fibre positioner's main technical requirement to achieve $8.0 \mu\text{m}$ RMS positioning accuracy has been analysed with respect to the optical design of the WHT and the anticipated observing conditions. The effect of non-compliance with this requirement has been shown to be non-catastrophic in the grand scheme of things, with a 100% increase in positioning error reducing the overall signal strength by only $\sim 6.5\%$.

The goals of this research have been defined. Prototyping of the proposed COTS robot axes was necessary in order to verify their performance against the requirements of WEAVE, thus retiring a major technical risk as the project progressed through its final design phase.

PART I: THE WEAVE FIBRE POSITIONER

2 Commissioning of prototype positioner robot

This section describes work done in commissioning a laboratory-based pick-and-place robot for the purpose of testing the WEAVE fibre positioner concept. This prototype system was instrumental in verifying the expected performance of the positioner in the final instrument, as reported in later sections of this chapter.

2.1 Hardware assembly

The available prototype hardware was a pick-and-place gantry comprising motors of the same or similar specification to the final WEAVE positioner design. Much of this was pre-assembled, but without a suitable gripper assembly. Figure 2.1 shows the assembled prototype hardware.




Figure 2.1: The prototype WEAVE positioner robot was built to allow pick-and-place positioning along a single axis for the purpose of laboratory testing.

The main linear axis of the prototype represents the y -axis of the WEAVE design, allowing positioning along a single metal strip representing the field plate. The z - and θ -axis are also present, along with the chosen gripper motor and the prototype measurement camera (described later).

Control of the existing hardware was only possible via a proprietary engineering interface on a desktop computer running a popular but non-ideal operating system. This software, which was nothing more than a commissioning tool, was unsuitable for controlling the axes in any meaningful way, not least because only one motor could be moved at a time. Pick-and-place actions were not possible.

In order to properly commission the prototype robot, a control system had to be built with the ability to move multiple axes at specified velocities and accelerations.

2.2 Development of a PLC-based control interface

A core philosophy of the WEAVE positioner design is to utilise COTS components and existing industrial standards and procedures wherever possible. To this end, a programmable logic controller (PLC) was to serve as the main interface to the robot axes. Moreover, a Bosch Rexroth PLC was chosen on the recommendation of the motor axes supplier, Schunk GmbH.

Substantial effort was put into the deployment and programming of a PLC in the lab, not only to enable full control of the prototype positioner from a computer workstation, but also to understand how a PLC-based architecture would integrate with the final instrument control system.

2.2.1 Positioner control architecture

There are various ways in which a PLC could be used in the WEAVE positioner. This is mostly a matter of how much intelligence one wishes to place in the PLC versus the external instrument control software.

Following discussions with the WEAVE positioner software lead, it was decided that the PLC's role in WEAVE should be a low-level one. The PLC was to serve as the communications interface between the instrument control software, which runs on a Linux workstation, and the axis motion controllers. Any operational or housekeeping tasks such as initialisation routines and interlocking were also to be handled by the PLC. Broadly speaking, the PLC would serve the purpose of exposing relevant motor functions and commands to the high-level control software in a friendly way.

PART I: THE WEAVE FIBRE POSITIONER

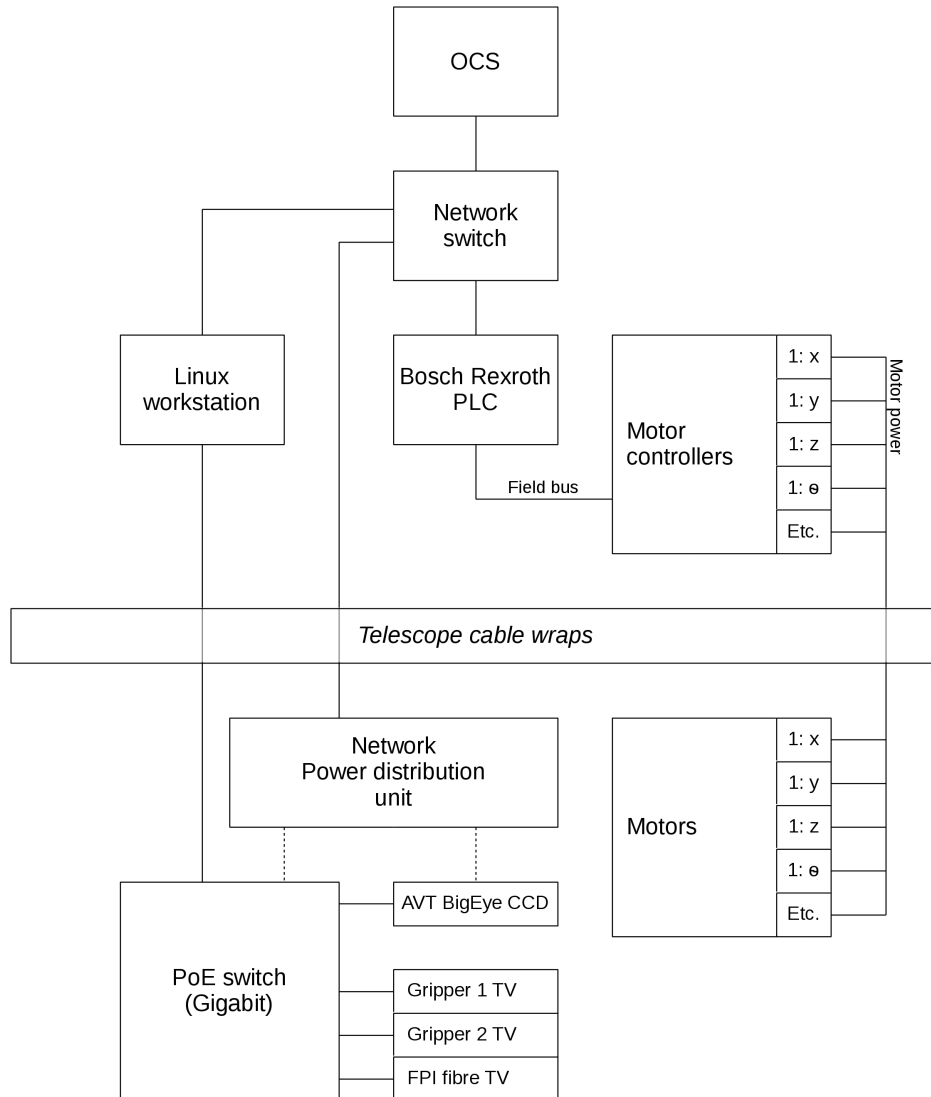


Figure 2.2: The control architecture for the final WEAVE positioner is its own subsystem below the Observatory Control System (OCS) and is centred around a Linux workstation running the positioner control software. The workstation handles all high-level positioning commands and processing of fibre images for the purpose of position feedback. The positioner PLC interface is connected to a local area network, using the TCP/IP protocol to communicate with the control software, and the Fieldbus protocol to communicate with the robot motors; all other connections are gigabit Ethernet unless specified.

PART I: THE WEAVE FIBRE POSITIONER

Figure 2.2 shows the control architecture agreed upon for the final positioner system. The positioner control software, running on a networked Linux workstation, handles nearly all the complex calculations involved in the reconfiguration of a field. These include calibration routines, coordinate conversions, image processing tasks, and so on. This makes good use of the workstation's computational resources while leaving the PLC to play to its main strength: the logical and deterministic execution of simple commands and procedures.

Gigabit Ethernet networks are used for communication between devices, including that between the workstation and the PLC. The decision was made to use the TCP/IP protocol here. The PLC and motion control hardware is kept local to the workstation in order to minimise mass at the telescope top end, with a long cable run to the motors on the positioner.

For the prototype positioner system, the architecture of the final instrument was mimicked in order to start meaningful development for the final design. Figure 2.3 shows the resultant control cabinet layout.




Figure 2.3: A prototype robot control system was built to mimic that of the full positioner (but for a single robot and no x -axis).

2.2.2 Definition of a custom command set

The positioner control architecture was designed with a simple custom command set in mind. These commands allow the positioner robots to be controlled and their status queried in a way that makes sense in terms of the positioner's physical design and function.

The commands comprise a simple sequence of 8-bit ASCII characters, in the spirit of simplicity and robustness. For example, the message

AP 50.123 500.456 90.789

is a command to move the x -, y - and θ -axis of robot A to the motor positions of 50.123 mm, 500.456 mm and 90.789° respectively. It is the task of the PLC to interpret this instruction and calculate the appropriate motion controller commands to move all three axes from their current position to their new position simultaneously. Information on the state of the robot system can be requested by the positioner control software at any time.

All coordinate conversions and calibrations are done within the instrument control software, removing the need for the PLC to know anything other than the robot's motor encoder values.

2.3 Section summary

This section has described one of the first tasks in this body of work, which was to assemble, commission and program a prototype fibre positioner robot for WEAVE.

The prototype robot has since enabled the laboratory testing and analysis reported in the remainder of this thesis chapter.

PART I: THE WEAVE FIBRE POSITIONER

3 WEAVE fibre placement metrology

The accurate determination of the location of optical fibres on the WEAVE field plate is a crucial aspect of the instrument, since being able to position fibres to a given accuracy relies on being able to measure to that accuracy. This section describes the requirements, principles, design and testing of the WEAVE measurement system.

Detailed analysis of likely error sources places the final estimated measurement uncertainty at $3.3\ \mu\text{m}$ RMS, which is well within the required fibre placement accuracy of $8.0\ \mu\text{m}$ RMS. Hence, the planned measurement system for WEAVE is shown to be suitable as feedback in the closed-loop pick-and-place fibre positioning process.

3.1 Fibre position measurement principles

The purpose of the fibre position measurement system is to facilitate the measurement of the position of optical fibres on the field plate beneath the positioner robots. The overall principle can be summarised in three statements:

1. We know where the robot axes are positioned in terms of axis coordinates from their built-in position encoders.
2. We can link fixed markers on the field plate to the axis coordinate system using cameras on the robots.
3. We can use the same cameras to measure science fibre positions with respect to the axes, and therefore with respect to the field plate.

This section introduces the core principles of the measurement system and how these principles tie together to allow accurate fibre position measurements.

3.2 The measurement cameras

A plan view of the WEAVE positioner is shown in Figure 3.1. Each of the two positioner robots has five motorised axes (x_0 , x_1 , y , z , and θ), provided by Schunk GmbH, with motor position feedback via built-in absolute encoders. These encoders give the position of the axes in real world distance units from a fixed but arbitrary origin on each axis. There is no intrinsic link to the coordinate system of the field plate below.

Attached to the θ -axis of each robot is the gripper assembly. The gripper assembly comprises motorised gripper jaws and a measurement camera that can image a fibre within the gripper jaws or, via a secondary focus, fibres or fiducial reference markers on the field plate below. The secondary focus is at a greater distance from the gripper assembly than the primary focus, so the robot can move around freely above the field plate and not collide with fibres that have already been placed.

The measurement camera is the all-important link between the robot and field plate coordinate systems. With the correct processing, images from the camera can be used to estimate the location of anything within its field of view, in terms of robot axis coordinates.




Figure 3.1: Each of the two fibre positioner robots has five motors, plus a gripper assembly with a measurement camera. Position information for all of these is combined to get a final measurement of the location of a fibre on the field plate.

3.2.1 Fibre measurement via back-illumination

A fibre back-illumination system is used to provide good contrast between an image of a fibre core and its surroundings. The system works by evenly illuminating the fibre slit block at the spectrograph, providing a relatively bright illumination of the fibre core at the field plate that can be imaged by a camera in order to obtain the fibre’s position. This is common among fibre-fed multi-object instruments, including the AAT’s 2dF positioner [18], from which the WEAVE positioner inherits many features.

The WEAVE back-illumination system is part of the spectrograph subsystem, but is an important interface to the positioner subsystem. A prototype unit is evaluated for suitability in Section 3.5.

3.2.2 Field plate fiducial markers

Placing fibres at definite locations on the WEAVE field plate comes with the problem of accurately mapping field plate coordinates to robot axis coordinates in an environment with changing temperature and varying mechanical flexure due to the telescope’s zenith angle.

A regular grid of reference markers on each of the WEAVE field plates will allow plate coordinates to be mapped to robot axis coordinates. The specific pattern of the markers on the each field plate will be measured to a high degree of accuracy at the point of manufacture, so this pattern can be fitted to the marker locations observed by the positioner’s measurement camera. This on-telescope measurement of the field plate is called a field plate ‘survey’. The survey is quickly executed before every field reconfiguration, providing compensation for axis encoder scale changes due to temperature, and mechanical flexure due to the current gravity vector. Only a small subset of markers needs to be measured in order to provide sufficient information to update the fitting of the grid.

Some proof of concept experiments for the proposed fiducial marker design will be presented in Section 3.7.

3.2.3 Fibre position measurement procedure overview

The procedure for measuring the position of a fibre currently on the field plate, using the principles outlined above, is summarised below. The procedure is the same for both positioner robots.

PART I: THE WEAVE FIBRE POSITIONER

1. A field plate survey is carried out for the current telescope position, to map the field plate to the robot axis coordinate system.
2. The measurement camera on the robot gripper captures an image of the back-illuminated fibre on the field plate.
3. Image processing routines find the fibre in the image and calculate the pixel coordinates of its centre with respect to the θ -axis.
4. The fibre's image coordinates are converted to real-world distances using the known pixel scale of the camera.
5. The fibre location is added to the motor encoder values to get the fibre's position relative to the axis coordinate system.
6. Finally, the position of the fibre on the field plate is found by applying the coordinate mapping from the field plate survey.

3.3 Design of a new measurement camera

Each of the WEAVE positioner robots will carry a camera for the purpose of measuring the position of fibres relative to the position of the robot axes. The camera is housed alongside the gripper motor to form the 'grripper assembly', which is fixed to the robot's rotating θ -axis.

This section describes the design and prototyping of a new measurement camera as part of the WEAVE final design phase. A pre-existing camera had a number of issues that made this redesign necessary. The revised design is shown to meet all necessary requirements.

3.3.1 Camera requirements

The camera had to satisfy the optical requirements of the measuring system as well as the mechanical constraints of the gripper assembly itself. Table 3.1 shows the requirements that were set for the system, including the need for two focal distances. This is because the camera must be able to focus on a fibre button within the gripper jaws as well as on stationary buttons and fiducial field plate markers from a greater height.

PART I: THE WEAVE FIBRE POSITIONER

Table 3.1: The design of the WEAVE measurement cameras began by defining suitable requirements for the system.

Parameter	Requirement	Rationale
Primary focus	Fibre face of gripped button	For measuring fibre position after gripping and releasing
Secondary focus	>7.0 mm below gripper assembly	Minimum clearance to avoid collisions with buttons
Pri/sec magnification	Equal	Allows common image processing routines for both foci
Field of view (diag.)	≥ 4.0 mm	Allows identification of individual button housings
Pixel scale	≤ 4.25 $\mu\text{m}/\text{px}$	Equates to a resolvable circle of ~ 20 px diameter for single fibres
Pixel aspect	1:1 (square)	Simplifies image processing by guaranteeing circular fibre images
Detector scan	Progressive scan, global shutter	Most robust against illumination changes and motion
Comms protocol	GigE with Power over Ethernet (PoE)	Easy integration and fast, robust protocol
Detector unit brand	Basler	Pre-existing software and experience with Linux API




Figure 3.2: The pre-existing measurement camera prototype was unsuitable for various technical reasons, including the off-axis jaw motor and a lack of dual focus.

3.3.2 Problems with pre-existing gripper assembly

A pre-existing prototype camera and gripper assembly, shown in Figure 3.2, had a number of issues that necessitated a redesign.

The gripper assembly in this case kept the camera *on-axis* with the robot's θ -axis motor while placing the gripper motor *off-axis*. This also meant that the gripper motor was off-axis from the fibre button being gripped, resulting in an uneven gripping force at the tip of the jaws. In real terms, the fibre buttons were seen to freely pivot about the single point of contact when gripped, which is not desirable for accurate and controlled pick-and-place manoeuvres.

The camera detector chosen for the pre-existing design was also unsuitable; it had a small format, oblong pixels and an interlaced readout, which added needless layers of complexity to fibre position measurement routines. A final problem was the lack of a secondary focus distance in the camera's optical design.

3.3.3 New design and prototyping

Given the available space envelope, the magnification and focal length of the lens system was limited. This, and the requirement of matching magnifications for each focus, pointed to a design with a fixed lens assembly and two incident optical paths combined with a beamsplitter. This is preferable to a single optical path, as it ensures identical magnification at two foci and removes the need for electro-mechanical systems to switch or refocus the optics. Correct placement of the two focal surfaces is then just a matter of steering the beams appropriately.


A working combination of an off-the-shelf lens doublet and a detector unit from the Basler Ace range of industrial machine vision cameras was found¹, which is shown in Figure 3.3a and detailed in Table 3.2.

Contrary to the pre-existing camera design, the revised design keeps the axis of the redesigned gripper jaws² aligned with the axis of the θ -axis motor and places the camera system off-axis instead. The optics provide a 1:1 magnification and the chosen detector has 1280×960 $3.75 \mu\text{m}$ square pixels, for a theoretical pixel scale of $3.75 \mu\text{m}/\text{px}$ and diagonal field of view of 6.0 mm. A prototype of the gripper assembly is shown in Figure 3.3b. The prototype successfully meets all technical requirements.


¹Acknowledgement goes to Ian Lewis for his contributions to the optical design.

²Acknowledgement goes to Matthew Brock for his assistance with the gripper jaw design.

PART I: THE WEAVE FIBRE POSITIONER



(a) Optical design



(b) Working prototype

Figure 3.3: The new measurement camera has two foci, one of which is steered with a beamsplitter and a mirror. The ‘near focus’ is for fibres that are gripped in the robot jaws, and the ‘far focus’ is for viewing field plate reference markers and fibres from a height.

Table 3.2: List of optical components for the new measurement camera design.

Component name	Part number
Detector unit	Basler Ace acA1300-30gm GigE mono
Lens (L1)	Techspeg 30mm achromatic doublet lens
Beamsplitter (BS1)	Edmund Optics 50R/50T plate beamsplitter
Mirror (M1)	Edmund Optics 4-6 λ first surface alu mirror

The planned assembly and integration procedure for the camera is to align and focus the system mostly via mechanical tolerances. An initial ‘set and forget’ fine focus operation will be carried out by moving the detector unit relative to the optics, after which the gripper assembly is fixed in place. Long-term focus drift is expected to be minimal and therefore can be compensated as part of the positioner operations procedure, by adjusting the pick-up height of the fibre button (z -axis position) as necessary for the near (button) focus. The corresponding change in the far (field plate) focus is similarly compensated by a small z -axis offset.

3.4 Design of an in-situ camera calibration scheme

This section describes the design and prototyping of an independent optical reference system that can be permanently in place on the WEAVE positioner and allows fast and accurate calibration of important camera parameters. The concept is shown to be very effective at focussing, aligning, and obtaining the pixel scale of a WEAVE measurement camera. In addition, the determination of the centre of rotation of the θ -axis using the reference system is shown to have an uncertainty of only $0.32\ \mu\text{m}$.

3.4.1 Camera calibration parameters

The correct operation of the overall measurement system relies on the proper setup and calibration of the measurement cameras. Four things must be known for each camera: i) the focal distance of the camera; ii) the alignment of the gripper unit's θ -axis; iii) the centre of rotation of the θ -axis; and iv) the pixel scale of the camera.

- **Camera focus:** The metrology cameras on the positioner's gripper units have a fixed optical design (described in Section 3.3) with a magnification close to unity at two different focus distances. These focus distances must be known so that the gripper units are moved to the correct height above the fibre buttons or the field plate, depending on which focus is being used.
- **θ -axis alignment:** The θ -axis of each gripper unit must be properly referenced so that the orientation of the gripper is known with respect to the field plate coordinate system.
- **Camera pixel scale:** The pixel scale of the camera system, measured in microns per pixel, ties image detector coordinates to real world distances. A good measurement of this value will ensure that the reported positioning errors are accurate.
- **θ -axis centre of rotation:** There is a pixel coordinate on each camera detector that corresponds to the centre of rotation of the θ -axis. Defining this as the origin, and therefore the place that fibres should be placed, removes any complex offsets as a function of the θ -axis angle. Since all measurements will be referenced to this point, an accurate calculation of its location is important.

3.4.2 The case for in-situ camera calibration

Finding camera calibration parameters could be treated as an engineering operation that only occurs when commissioning the instrument or at infrequent maintenance intervals thereafter. Given the importance of accurate metrology on WEAVE, however, it would be prudent to design an in-situ calibration scheme. Such a scheme would: a) quickly account for focus and pixel scale drift due to changes in temperature; and b) allow easy ‘sanity checking’ of the cameras in the event of a problem.

One can imagine a simple case where science fibres are used for camera calibration, in tandem with the absolute distance encoders already present on the axes of the robot gantries: Focus could be checked by measuring the intensity profile of the back-illuminated fibre; pixel scale and θ -axis alignment could be found by moving the robot’s x - or y -axis a certain distance and measuring the apparent fibre shift.

A more robust approach to measurement camera calibration would be to provide an independent optical reference that does not rely on the position encoders within the robot axes at all. An ideal solution would allow fast and accurate recalibration of all four parameters (focus, alignment, pixel scale, rotation axis) so that a routine could be built in to the normal operation mode of the instrument. This would easily remedy long-term environmental variations, and even short-term variations such as mechanical flexure caused by telescope pointing. A built-in reliable reference also makes life much easier during instrument commissioning and troubleshooting.

3.4.3 Automatic determination of focus, alignment and pixel scale

A system for automatically aligning, focussing, and obtaining the pixel scale of the measurement cameras has been designed, based on an optical Ronchi ruling. Ronchi rulings are optical targets comprising constant-interval black and white, or opaque and transparent, bars with high-contrast edges. They are very accurate, with off-the-shelf examples having typical line width tolerances of $0.3\ \mu\text{m}^3$.

The system uses a back-illuminated Ronchi ruling that is either parfocal with the nominal focal surface of the instrument, or at a known offset from it. The target is placed away from the field plate, perhaps under the location at which the gripper assemblies are parked when not active. This allows camera calibration to take place away from the fibres, although there may be some advantages to placing several targets around the perimeter of the plate instead.

³Specification obtained from Edmund Optics Ltd.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.4: A back-illuminated Ronchi ruling provides an accurate real-world optical reference for calibrating the positioner’s measurement cameras.

The lines on the Ronchi ruling must also be physically aligned with the x - and y -axis of the positioner, or at a known angular offset. The ruling must be fixed in place securely. The back-illumination of the target must be even and diffuse; electroluminescent or organic light-emitting diode (OLED) panels would be suitable choices. Figure 3.4 shows a mock-up of the system.

The principle behind the calibration system begins with the assumption that the measurement camera is already roughly in focus (Ronchi lines discernible) and roughly aligned (Ronchi lines within a few degrees of vertical), either from a previous calibration or from manual setup by a human. From this point, analysis of an image of the Ronchi ruling can reveal a relative measurement of how well the camera is focussed and aligned. The procedure is as follows.

1. The image is collapsed vertically, by taking the mean pixel value of each column, to produce a one-dimensional array representing the entire width (x -axis) of the image.
2. An approximation of the derivative of the 1D pixel array is obtained by taking the finite difference between each pixel value; the difference will be highest at the boundary between black and white lines.
3. A peak-detection algorithm is used to find the location of the contrast peaks, and therefore the location of the line boundaries on the Ronchi ruling. This is illustrated in Figure 3.5.
4. The mean peak value is calculated to provide an overall ‘score’ for the image; this value will be higher when the camera is better aligned or better focussed, because the changes in pixel brightness will be more abrupt.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.5: The sub-pixel location of a line transition is inferred by calculating the centroid of the finite difference across a vertically-collapsed image of the Ronchi ruling.

The above procedure can be used as feedback in a closed-loop iterative routine adjusting the θ -axis (alignment) and z -axis (focus) of each robot until optimal values are found, according to a predefined tolerance.

An added feature is that, once the camera has been aligned and focussed, the locations of the contrast peaks provide an excellent measure of the camera pixel scale by dividing the real world line width of the Ronchi ruling by the interval between the peaks. This interval can be accurately inferred by fitting a straight line to the locations of each successive peak. This is demonstrated in the following pages.

3.4.4 System prototyping and testing

In prototyping the image processing routines required for this system, care was taken to make the peak detection algorithm simple yet accurate. The algorithm also has to be robust when the image is out of focus or misaligned. The procedure is as follows.

1. Take the absolute value of all array elements, so that all peaks become positive.
2. Set a threshold value, above which a peak is defined. This is set at one tenth of the maximum element value, a value which was chosen empirically and sits above expected noise levels.
3. For broad peaks consisting of more than one element (e.g. defocussed line transitions), take the centroid of the values around the peak as the location of the peak.
4. Finally, estimate the value at the peak location. This is done using simple linear interpolation.

PART I: THE WEAVE FIBRE POSITIONER

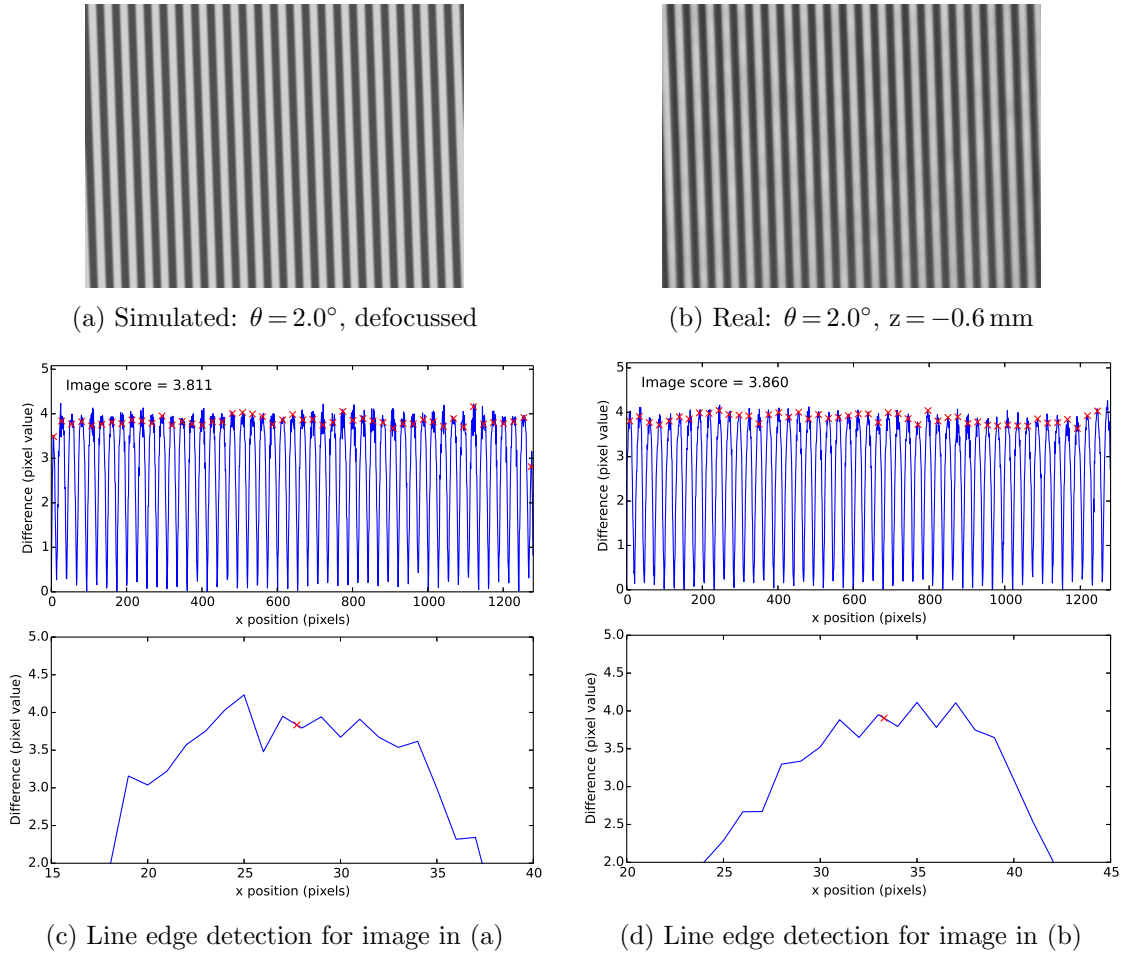


Figure 3.6: Simulated and real images of an optical Ronchi ruling agree well for a misaligned and defocussed case, with successful line edge detection and peak-finding.

Simulated images of Ronchi rulings were generated in order to prove the camera calibration concept. These were based on a readily available line width of $100\ \mu\text{m}$, which provides good sampling across the camera’s field of view. The simulated images attempted to mimic the optical properties of the WEAVE measurement cameras, including the addition of Poisson noise and Gaussian blurring. Optical distortion was not added, but was considered negligible given the camera’s narrow field of view. Indeed, a useful feature of this method is a high tolerance of optical distortion.

The results of processing the simulated images are shown alongside real images of a $100\ \mu\text{m}$ Ronchi ruling in Figures 3.6 and 3.7. The simulated images agree well with reality, and the peak-finding algorithms were able to infer the line transitions in the case of a defocussed image with a 2.0° offset, as well as the aligned and focussed case. Figures 3.7c and 3.7d illustrate the fitting of a straight line to the locations of the peaks, the gradient of which gives the line spacing on the detector.

PART I: THE WEAVE FIBRE POSITIONER

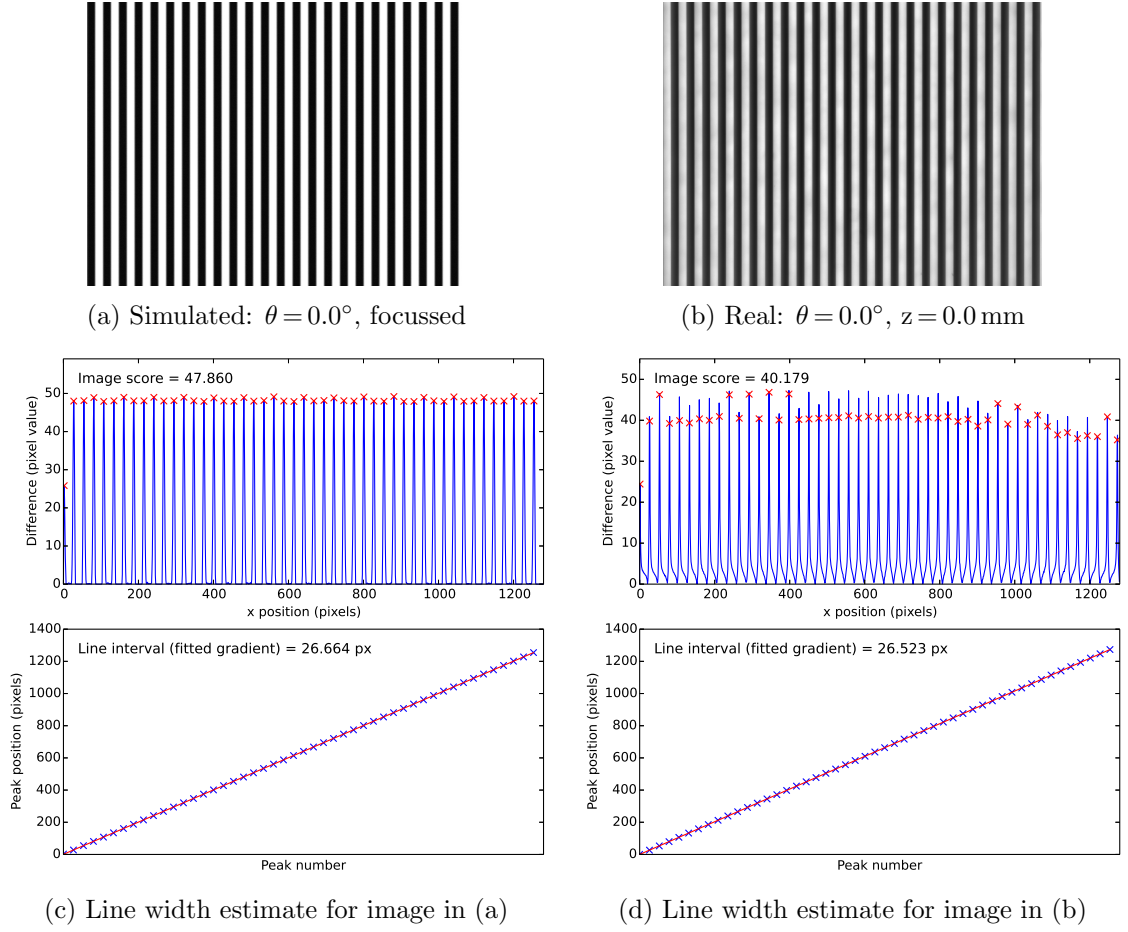


Figure 3.7: Simulated and real images of an optical Ronchi ruling agree well for the aligned and focussed case, with line-fitting reliably revealing the camera pixel scale.

With a spacing of 26.523 px, the pixel scale of the prototype camera is given by

$$\begin{aligned}
 \text{Pixel scale} &= \frac{\text{Line width}}{\text{Measured spacing}} \\
 &= \frac{100.0}{26.523} \\
 &= 3.770 \mu\text{m}/\text{px}
 \end{aligned} \tag{3.1}$$

Since we know that the pixel size on the detector is $3.75 \mu\text{m}$, we can verify the magnification of the camera and compare it to the designed value of unity.

$$\begin{aligned}
 \text{Magnification} &= \frac{\text{Pixel size}}{\text{Pixel scale}} \\
 &= \frac{3.750}{3.770} \\
 &= 0.995
 \end{aligned} \tag{3.2}$$

PART I: THE WEAVE FIBRE POSITIONER

Everything seems to be in order. Returning to the calibration routine, it's useful to know by how much the θ - and z -axis can be initially offset while still reporting a meaningful image score. Figure 3.8 gives us an idea of this, showing the score as a function of axis offset, with data obtained with the prototype positioner robot and prototype measurement camera.

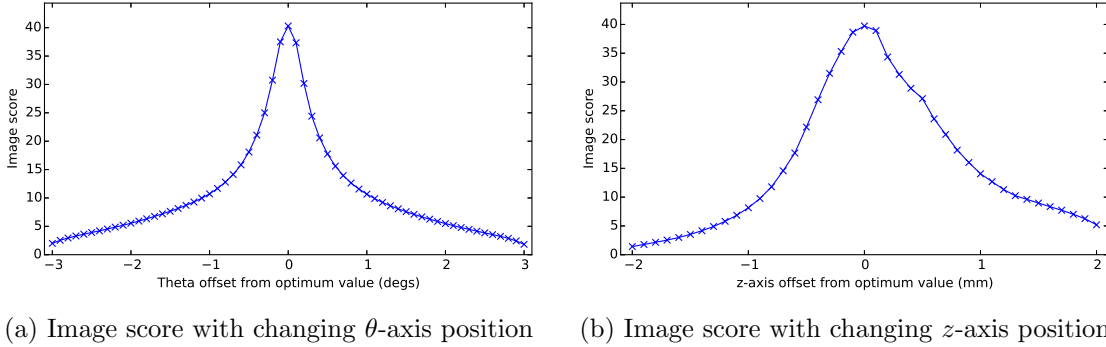



Figure 3.8: The auto-calibration procedure aims to find the highest ‘image score’ by changing the camera alignment (θ -axis position) and focus (z -axis position). These plots show the change in image score for real calibration images on the prototype system.



(a) Camera view before automatic alignment and focus ($\theta = 3.0^\circ$, $z = -2.0$ mm) (b) Camera view after automatic alignment and focus ($\theta = 0.0^\circ$, $z = 0.0$ mm)

Figure 3.9: The automatic calibration routine was able to focus, align, and obtain the pixel scale for this example scenario using the positioner’s prototype measurement camera and an off-the-shelf Ronchi ruling.

The final stage of prototyping the calibration procedure was to make it truly automatic. There are many ways to approach this, but the following method managed to successfully align and focus the camera from an initial offset of 3.0° and -2.0 mm for the θ - and z -axis respectively (see Figure 3.9). The initial stages of the procedure are based on how a human might begin the same task, followed by a simple binary search to home in on the optimal axis positions. The procedure is fairly unsophisticated and assumes only a global optimum, but is only intended to prove the feasibility of an auto-calibration feature on the final instrument.

PART I: THE WEAVE FIBRE POSITIONER

1. The image score is calculated for the initial conditions.
2. The θ -axis is rotated in a negative direction by a pre-defined step size (e.g. 0.5°), an image is taken and its score compared to the previous score.
3. The above stage is repeated until the score begins to decrease. This position is the lower limit.
4. The θ -axis is stepped in a positive direction, again until the score begins to decrease. This position is the upper limit, and we assume that the maximum score (optimal position) lies between the two limits.
5. A binary search is performed, whereby the score at the position half way between the two limits is calculated, then compared to the limit scores. The central position, plus the highest-scoring limit position become the new limits.
6. When the position is within a pre-defined tolerance (e.g. 0.02°), the search is complete and the axis is considered to be aligned.
7. The same procedure is carried out for the z -axis, until the image is considered focussed.
8. The θ -axis procedure is repeated, in case severe defocus affected the precision of the result.
9. A final focussing procedure is carried out with the z -axis.
10. Pixel scale is inferred from the spacing of the lines, as described earlier.

We now have a way to obtain three of the four key parameters listed in Section 3.4.1 (focus, alignment, pixel scale) using only a Ronchi ruling. This leaves the θ -axis centre of rotation, which can be incorporated into the auto-calibration procedure by adding another optical target, described below.

3.4.5 Automatic determination of θ -axis centre of rotation

An extension of the optical reference system described above is to allow the fourth and final camera calibration parameter to be automatically determined: the pixel coordinates of the centre of rotation of the θ -axis. This is the point to which all image-domain measurements should be referenced, because it is this point that relates to the known position of the robot motors and is independent of the θ -axis position.

The procedure for finding the θ -axis uses a single circular pinhole aperture, which is back-illuminated by a panel (as with the Ronchi ruling) so that the camera sees a bright circle on a dark background, similar to fibres themselves. The pinhole aperture

PART I: THE WEAVE FIBRE POSITIONER

diameter matches the apparent diameter of the science fibre cores, so that the same image processing techniques can be used to accurately estimate the centre of this circle (these techniques will be discussed in Section 3.6). A real fibre could be used for this purpose, but it is assumed that an independent system placed away from the field plate and science fibres is preferable.

Given a suitable illuminated target, the pixel coordinates of the θ -axis can be estimated from a pair of measurement camera exposures taken at θ -axis motor positions that are 180° apart. The θ -axis location is simply taken as the mid-point between the measured fibre position for each of the images. This is shown in Figure 3.10.

The above technique was tested with the prototype WEAVE positioner robot and a back-illuminated pinhole aperture of $80\ \mu\text{m}$ diameter, an image of which is shown in Figure 3.11 for some context.

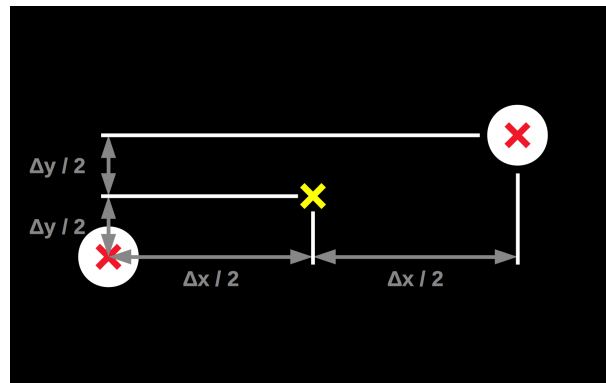


Figure 3.10: The point on the measurement camera detector corresponding to the centre of rotation of the θ -axis (yellow cross) is mid-way between the centroids (red crosses) of two fibre images taken 180° apart.




Figure 3.11: A back-illuminated off-the-shelf $80\ \mu\text{m}$ pinhole can be used to infer the centre of rotation of the θ -axis. This is close to the diameter of the science fibres.

3.4.6 Calibration uncertainty

Errors in the camera pixel scale, camera alignment, and θ -axis reference will ultimately cause fibre positioning errors, so these should be estimated for the sake of our error budget.

Pixel scale error The pixel scale error can be estimated from the manufacturer's $0.3\ \mu\text{m}$ tolerance for the line width on the target Ronchi ruling. This corresponds to an error of $0.3\ \%$ for the chosen line width of $100\ \mu\text{m}$. Since the calibration routine uses 46 lines, we can assume the line width tolerance is random and therefore decreases by a factor of \sqrt{N} :

$$\begin{aligned} \text{Pixel scale error} &= \frac{0.3}{\sqrt{46}} \\ &= 0.00044\ \% \end{aligned} \tag{3.3}$$

That is, a measurement over a distance of $1.0\ \text{mm}$ could have an error of up to $0.44\ \mu\text{m}$.

Camera alignment error An upper limit for camera alignment error can be taken as the alignment tolerance specified in the auto-calibration routine (see item 2 in above list). 0.02° was found to work well, but could be tightened at a cost of increased calibration time. A camera misalignment of 0.02° would introduce positioning errors of $0.35\ \mu\text{m}$ for every $1.0\ \text{mm}$ from the θ -axis centre of rotation.

θ -axis reference error By analysing the scatter of the inferred θ -axis coordinates over multiple test runs, we can estimate the magnitude of the error in this calculation. Shown in Figure 3.12 is a scatter of $\sigma = 0.32\ \mu\text{m}$ for this, which seems reasonable.

3.5 Verification of fibre back-illumination scheme

This section describes the testing of a prototype fibre back-illumination unit. The design of the unit is shown to provide good quality illumination despite its design compromises, with the only significant source of uneven illumination being the lateral and angular offset of some fibres in the slit block. The measurement error caused by this is shown to be acceptably small, at $0.20\ \mu\text{m}$ for the worst case.

PART I: THE WEAVE FIBRE POSITIONER

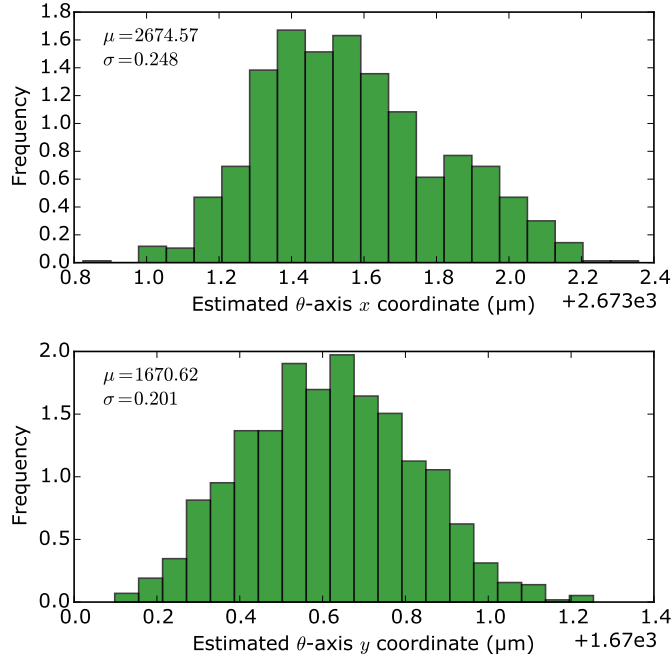


Figure 3.12: After 1000 calculations of the θ -axis reference coordinate on the measurement camera detector, the total uncertainty of this calculation is seen to be $\sim 0.32 \mu\text{m}$.

3.5.1 Back-illumination on WEAVE

Back-illumination is a technique commonly used to aid the measurement of the position of optical fibres at the field plate of a fibre-fed instrument. Each fibre is illuminated from its nominal exit, so that the entrance face can be imaged by a camera and the location of the fibre core measured accurately.

The light used for back-illumination should be as uniform and diffuse as possible, as this guarantees symmetric illumination across the core of the fibre and simplifies the measurement process.

The WEAVE back-illumination system will use an actuated slit exchange unit at the spectrograph to switch between fibre illumination while positioning, and spectrograph use while observing. This makes the design of the back-illumination unit part of the spectrograph rather than the positioner [44].


The back-illumination system is a crucial interface between the spectrograph and positioner, because positioning performance ultimately depends on the illumination being done well. For this reason, an important piece of work was to test a prototype WEAVE back-illumination unit supplied by the Netherlands Research School for Astronomy (NOVA). NOVA's design features two banks of off-the-shelf LEDs in a curved housing. The curvature of the housing roughly matches that of the fibre slit

PART I: THE WEAVE FIBRE POSITIONER


block, and it was hoped that a simple arrangement of diffuse white reflectors would provide enough scattering to illuminate all fibres equally and evenly. Figure 3.13 shows a photograph of the prototype unit and a cross-section of its optical design.

Testing of the prototype illumination unit using the prototype WEAVE fibre measurement camera aimed to answer four important questions:

1. **Is the illumination bright enough?** The supplied prototype illumination unit had two banks of 39 white LEDs, operating at their nominal rated current of 20 mA. Are the required exposure times for the measurement camera acceptable at this brightness?
2. **Does the LED diffuser work well enough?** The use of standard LEDs, which are encapsulated in a lensed body and produce a relatively narrow beam, means that the performance of the diffuse reflectors is especially important. Is there a significant variation in illumination as a fibre is moved across the front of the back-illumination unit?
3. **Would a single bank of LEDs be sufficient?** NOVA would prefer it if one of the rows of LEDs could be removed, meaning that the diffuser would be illuminated from one side only. How would this affect the quality of illumination across the fibre?
4. **Will offset fibres be illuminated sufficiently?** The shape of the fibre slit block means that some fibres would have a lateral and/or angular offset from the incident illumination. Do the worst-offset fibres suffer a reduction in illumination quality that causes measurement errors?



(a) Prototype unit



(b) Optical cross-section

Figure 3.13: The WEAVE back-illumination unit is curved, with banks of LEDs along its sides and a white diffuse reflector.

3.5.2 Test setup

Without a fully-populated fibre slit block, an alternative setup was devised in order to test the prototype back-illumination unit. The setup needed to simulate fibres as they would be mounted at different points along the block, including those offset from the illumination unit due to the block's curvature. To this end, a pivoting arm was designed that allowed a single prototype science fibre to be swept across the face of the illumination unit while maintaining the correct illumination distance (see Figure 3.14). The test fibre could be mounted on-axis with the unit, or at a lateral and angular offset that corresponds to the worst possible case for all positions in the slit block design (see Table 3.3). This provided the conditions necessary to answer the above questions through a series of tests⁴.




Figure 3.14: A test rig with a pivoting arm simulates fibre placement at different points along the would-be fibre slit block. Note that the back-illumination unit will be mounted vertically in the spectrograph, not horizontally as shown here.

Table 3.3: General specifications of the WEAVE fibre slit block.

Radius of curvature	312.734 mm
Maximum vertical fibre offset	2.5 mm
Maximum angular fibre offset	4.0°
Illumination unit distance	14.4 mm

⁴Thanks go to Ellen Schallig for repeating several sets of measurements in the following sections.

3.5.3 Illumination brightness

An initial test of an on-axis test fibre showed ample illumination from the two banks of LEDs, with a measurement camera exposure time of 1.8 ms being suitable. This is considered a sufficiently small fraction of the camera's 33 ms overall acquisition time for a single frame (i.e. at 30 frames per second).

An issue that arose from these initial trials was that successive images of the same fibre showed a noticeable variation in brightness. This was being caused by the driver circuit for the LEDs, which used pulse width modulation (PWM) for intensity control. The on-off switching of PWM drive circuits is unsuitable for imaging applications unless the switching period is much shorter than the exposure time, for the very reason that was observed here. The PWM driver was replaced with a constant current source (a single resistor) and the problem was eliminated. Consequently, PWM circuits will not feature in the final back-illumination system.

The use of white LEDs was questioned, if only for the reason that using broadband light sources without good justification isn't best practice in a sensitive spectroscopic instrument. After studying the data sheet of the measurement camera's detector and discussing it with the spectrograph designers, it was agreed that narrowband LEDs would be used, at a wavelength at or close to the peak sensitivity of the detector. This is likely to be an amber colour (590 nm). This will have the effect of decreasing the brightness of the illumination by a factor of ~ 2.5 , increasing the required exposure time to ~ 4.5 ms, but this is still acceptable.

3.5.4 Fibre-to-fibre variations

The effectiveness of the diffuse reflector in the back-illumination unit was evaluated by looking for changes in fibre illumination at many different locations across the unit. This issue is less to do with a change in overall brightness than the possible asymmetric illumination of a fibre, causing errors in fibre position measurement.

A simple centroid method was chosen for estimating the centre of the imaged fibre face, which finds the fibre's 'centre of brightness' by weighting image pixels according to their intensity. For an evenly-illuminated fibre, the centroid would coincide with the centre of the fibre core. Any asymmetry in illumination, however, would cause a shift in the centroid.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.15: The scatter in the fibre centroid for 78 equally-spaced fibre positions across the active section of the back-illumination unit ($0.042\ \mu\text{m}$ radial) is no greater than the known measurement noise for the test ($0.040\ \mu\text{m}$ radial), meaning that imperfect illumination across the back-illumination unit causes no measurable increase in fibre measurement uncertainty.




Figure 3.16: The mean centroid position of a back-illuminated fibre shifts by a negligible amount when only one LED bank is used (blue, left) compared to two opposing banks (orange, right). The total radial shift is $0.08\ \mu\text{m}$. These histograms are for 78 equally-spaced fibre positions across the active section of the back-illumination unit.

PART I: THE WEAVE FIBRE POSITIONER

Figure 3.15 shows centroid measurements for 78 equally-spaced fibre positions across the active section of the back-illumination unit. For an ideal system, the scatter of these measurements would simply be the measurement noise ($0.040\ \mu\text{m}$ for this test), and indeed this is the case (observed radial scatter $0.042\ \mu\text{m}$). Therefore we can conclude that the imperfect LED diffusers cause no measurable increase in fibre measurement uncertainty.

3.5.5 Single-sided illumination

Figure 3.16 shows centroid measurements for 78 equally-spaced fibre positions across the active section of the back-illumination unit for the single and double LED bank cases. The difference in the mean centroid values gives a radial shift of just $0.08\ \mu\text{m}$ when one bank is disabled. There can be little doubt that this is negligible compared to WEAVE's $8.0\ \mu\text{m}$ error budget, therefore NOVA's desire to remove one of the LED banks is no problem.

3.5.6 Fibre offsets

As with the single-sided illumination tests, centroid measurements were used to compare the illumination of an on-axis fibre with the worst-offset fibre that will exist in the slit block. This fibre would be in the middle of the block, offset from the centre of the illumination unit by $2.5\ \text{mm}$ and 4.0° to normal with the front of the unit.

Figure 3.17 shows centroid measurements for 26 equally-spaced fibre positions across the central third of the back-illumination unit for the on-axis and offset cases. The difference in the mean centroid values gives a radial shift of $0.20\ \mu\text{m}$ for the offset case. This is acceptable, especially when we consider that the average misalignment across all fibres is much smaller. Therefore the slight geometrical difference between NOVA's back-illumination unit and the fibre slit block is considered acceptable.

3.5.7 Measurement errors caused by non-ideal fibre images

It has been shown that the design of the WEAVE back-illumination unit provides good quality illumination for all fibres in the slit block, even when illumination is provided by LEDs along only one side. The most significant change in illumination evenness is caused by the lateral and angular offset of some fibres with respect to the front of the back-illumination unit. This has been shown to introduce a radial shift in the image centroid of $0.20\ \mu\text{m}$ for the worst fibre.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.17: The mean centroid position of a back-illuminated fibre shifts by an acceptably small amount when the fibre is mounted with a worst-case lateral and angular offset (orange, right) instead of being on-axis with the back-illumination unit (blue, left). The total radial shift is $0.20 \mu\text{m}$. These histograms are for 26 equally-spaced fibre positions across the central third of the back-illumination unit.

In considering how well the *true* centre of a fibre can be estimated from an image, imperfect back-illumination is not the only error source. Fibre termination quality, measurement camera aberrations, and even dirt on a fibre will cause further centroid inaccuracies that are difficult to predict. It follows that a simple centroid measurement is not a reliable way to estimate the location of the centre of the fibre at high magnifications. A more robust method of determining the pixel coordinates of a fibre’s centre warrants investigation.

3.6 Analysis of techniques for finding a fibre’s centre

This section presents and compares four candidate image processing techniques for estimating the sub-pixel coordinates of the centre of an imaged fibre. One of these techniques is a new approach developed specifically for this task. The techniques are found to produce significantly different results for the same set of fibre images.

3.6.1 The importance of accurate centre-finding

The overall purpose of the WEAVE fibre positioner is to place the centre of a science fibre at a specified point on the field plate⁵. This raises a fundamental and crucial question: How is the centre of a fibre found? This is a question of image processing.

There are many scenarios in astronomy that call for the inference of the centre of a spot in an image to sub-pixel accuracy. Stars are a common example, but they come with the prior knowledge that they are point sources. This means that a simple centre of brightness measurement (centroid) can produce a robust result without much additional processing.

In the case of metrology for fibre-fed instrumentation, centroiding may also be the technique of choice for low-magnification measurement cameras, where the shape of an illuminated fibre face is unresolvable and its image is dominated by the PSF of the camera; multi-object spectrographs such as 4MOST [26] and DESI [31] are contemporary examples that have adopted this approach. But for high-magnification measurement cameras such as WEAVE, where an 85 μm fibre has a diameter of more than 20 pixels in the image, a centroid is not robust. This is illustrated in Figure 3.18, which shows an extreme example of uneven core illumination and how it influences centroid calculation.

The risk of positioning errors caused by inaccurate centroid results is compounded by anecdotal evidence⁶ that the AAT's 2dF facility regularly miscalculates fibre locations in this way. It follows that a better method should be sought for WEAVE.




Figure 3.18: This fibre image (a) shows an extreme example of uneven core illumination, resulting in a centroid calculation not at the centre of the fibre; the non-central intensity peak can be seen clearly in the 3D profile (b).

⁵The placement of fibre bundles (guide fibres, IFUs) is taken to be an extension of the more basic task of measuring single fibres.

⁶Anecdotal evidence of 2dF centroid issues amounts to personal email exchange and conversations with AAO software engineers.

3.6.2 Evaluation criteria for centre-finding on WEAVE

Before seeking better techniques for estimating the centre of fibre images, we need to decide how to evaluate them. Some effort was put into a software simulator to produce fibre images with uneven illumination, but this presented the problem of deciding what the simulated intensity profile should look like, and how drastic the effect should be. Later, however, testing a prototype of WEAVE's back-illumination unit (described in Section 3.5) presented a reasonable solution: Instead of simulating fibre images, real images from these tests were a good representation of the actual conditions on the final instrument.

Three criteria for evaluating the performance of fibre centre-finding techniques were defined as follows. Absent from this list is a metric for the absolute accuracy of the measurement, which will be discussed later (see Section 3.6.9).

1. **Tolerance of fibre offsets in slit block:** Better techniques will show a smaller shift in their results for the fibres in the slit block that are known to have a lateral and angular offset with respect to the back-illumination unit.
2. **Tolerance of measurement camera noise:** Better techniques will be more resilient to detector pixel noise from the measurement camera.
3. **Time taken to calculate result:** While faster is generally preferable, the practical limit is defined by the positioner's requirement for ≤ 100 ms fibre measurements. This is also somewhat linked with item 2, above, as a poorer noise tolerance may be mitigated by stacking multiple frames (with a time penalty).

Since a modest commodity computer has already been selected for development of the WEAVE positioner software⁷, this will be used as the baseline for available computing resources when evaluating the speed of algorithms.

3.6.3 Pre-processing of measurement camera images

All centre-finding techniques share a common initial 'pre-processing' stage, designed to get image statistics and crop around the fibre to be measured.

A new fibre measurement begins with a digital frame from one of WEAVE's two fibre measurement cameras. These images are 1280×960 px in size, with a pixel scale of $3.77 \mu\text{m}/\text{px}$ and a diagonal field of view of 6.0 mm. This means that WEAVE's

⁷Workstation spec: Dell OptiPlex 755; dual core 3.0 GHz CPU; 4 GB RAM.

PART I: THE WEAVE FIBRE POSITIONER

85 μm single science fibres will appear in the image as 22.5 px diameter spots⁸, and that multiple spots may be present in cases where neighbouring fibre buttons are close by.

An image pre-processing stage has been written [47] and is described below. It is able to detect the presence of multiple fibres, isolating the fibre image closest to the centre of the gripper assembly if necessary. Images containing no fibres are also detected. The routine⁹ computes various useful statistics about the image, including an estimate of a sensible value for the background pixel value (or ‘black point’). It does this quickly, by coarsely sampling the input image.

1. The pixels at every n^{th} row and column are visited and their values added to a histogram for the whole image. $n=4$ works well for the WEAVE fibre diameter.
2. The histogram is used to determine the most common pixel value (the mode), which is a good estimate of the background signal level of the camera. The minimum and maximum pixel values are also found, ignoring zero-value and saturated pixels in order to protect against dead or hot pixels. An overexposure warning is triggered if there are many saturated pixels. The value representing the p^{th} percentile of the histogram is calculated, where p is the percentage of pixels expected to be part of the fibre. If this value is close to the mode, then there is not a bright enough fibre image and the frame is rejected. Otherwise, the value half way between the mode and the p^{th} percentile pixel value is used as the ‘threshold value’, below which pixels are considered not part of a fibre.
3. The binary centroid of the coarsely-sampled image is found, treating all pixels above the threshold value with equal weight.
4. All pixels above the threshold value have their pixel moment calculated with respect to the centroid position. The sum of the magnitude of these moments can be used to detect the presence of multiple fibres, as illustrated in Figure 3.19. If the sum is high, indicating more than one fibre, then the image is progressively cropped about its centre until only one fibre is present.
5. Finally, the original image is cropped around the fibre of interest, with an appropriate margin to ensure the whole fibre is present in the cropped image. The cropped image can now be passed to a function that precisely determines the centre of the fibre. These functions are explored in the following sections.

⁸This number may need correcting in the final positioner to account for the lenses in the fibre prisms. Their effect on magnification will depend on the final measurement camera design.

⁹Thanks go to Tony Farrell of the AAO for useful discussions on this procedure, and his advice on possible improvements to existing methods used by the AAT’s 2dF positioner.




Figure 3.19: The presence of more than one fibre (white spots) in an image will produce a higher sum of absolute pixel moments (grey arrows) with respect to the image centroid (red dot) than a single fibre image.

The time taken for the pre-processing operation (code compiled with Cython [48]) has been measured at 12.2 ms on the WEAVE positioner development workstation. This is well within the 65 ms available processing window.

3.6.4 Centre-finding technique A: The centroid

The shortcomings of a simple centroid measurement for WEAVE fibre images have already been discussed in the introduction to this section. They chiefly concern the centroid's inherent bias to brighter areas of an image, and hence its dependence on flawless fibre images. A further drawback is the requirement to choose a suitable threshold pixel value for the process (below which pixels are ignored). Different threshold values produce different centroid results. Automatic determination of the threshold value therefore becomes important, and the whole technique becomes less robust.

A centroid measurement may also be called a centre of mass, centre of gravity, or centre of brightness measurement. In the world of image processing, the centroid of an image is defined in terms of image moments. Moments of an image are given by

$$M_{pq} = \sum_x \sum_y x^p y^q I(x, y) \quad (3.4)$$

where x and y are the coordinates of a pixel, and I is the intensity of the pixel (pixel value). The 'order' of the moment is given by $p+q$. The components of the image centroid (\bar{x}, \bar{y}) are defined as

$$\begin{aligned} \bar{x} &= \frac{M_{10}}{M_{00}} \\ \bar{y} &= \frac{M_{01}}{M_{00}} \end{aligned} \quad (3.5)$$

PART I: THE WEAVE FIBRE POSITIONER

Implementing a centroid calculation on a fibre image is a fairly straightforward process. Indeed, it is a standard function in image processing libraries for many programming languages.

The results for a centroid measurement of sample fibre images are summarised in Table 3.4. These are the same results previously reported for the back-illumination tests in Section 3.5. The results will be compared to those of other techniques later in this section.

Table 3.4: Results of centre-finding tests using a simple centroid.

Test criterion	Metric	Result
Offset fibres	Mean radial shift in result	0.20 μm
Measurement noise	RMS scatter for static fibre	0.07 μm
Speed	Mean time to calculate result	0.14 ms

3.6.5 Centre-finding technique B: A parabolic fit

This technique fits a second-order polynomial to separate one-dimensional collapsed images of the fibre along the x - and y -axis. Figure 3.20 shows these profiles for a real fibre image. Only values within an appropriate percentile are fitted to, in order to exclude the wings of the profile.




Figure 3.20: The parabolic fit method fits two second-order polynomials to 1D versions of the fibre image collapsed along the pixel rows and columns.

PART I: THE WEAVE FIBRE POSITIONER

The centre of the fibre is given by the vertices of the two component parabolas. Each vertex is given by

$$x_{vertex} = -\frac{b}{2a} \quad (3.6)$$

for the parabola given by

$$y = ax^2 + bx + c \quad (3.7)$$

The results for a parabolic fit measurement of sample fibre images are summarised in Table 3.5. These results will be compared to those of other techniques later in this section.

Table 3.5: Results of centre-finding tests using a parabolic fit.

Test criterion	Metric	Result
Offset fibres	Mean radial shift in result	0.21 μm
Measurement noise	RMS scatter for static fibre	0.05 μm
Speed	Mean time to calculate result	2.39 ms

3.6.6 Centre-finding technique C: Cross-correlation with a circle

Making use of prior knowledge is often a good way to optimise a process. Given that we know the fibre images will be circular, and of a known diameter, we can remove most of the effect of unevenness across the face of the fibre by looking at the cross-correlation of a real fibre image with an idealised fibre image.

Because the cross-correlation will be dominated by sudden intensity changes at the edge of the circle, the relatively minor variations in intensity across the face of the fibre will have little effect on the optimum position. In optical terms, this essentially gives you centre coordinates for the point of highest encircled energy.

The prototype of this technique uses a modified algorithm for 2D Fourier space cross-correlation, originally written for the translational alignment of diffuse astronomical images¹⁰. The code is mostly written in pre-compiled C and is optimised to find the best sub-pixel result quickly.

The idealised fibre image was simply a white circle on a black background, pre-generated by a circle-drawing function at a diameter matching that of the fibre core, as shown in Figure 3.21.

¹⁰Acknowledgement goes to Adam Ginsburg, author of the Image Registration Python module.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.21: The cross-correlation centre-finding technique takes the centre of the fibre to be the point of highest cross-correlation between real and idealised fibre images.

The results for a cross-correlation measurement of sample fibre images are summarised in Table 3.6. These results will be compared to those of other techniques later in this section. An added advantage of this approach is that it doesn't involve any image thresholding, and therefore its result does not rely on a chosen threshold value. However, performance does depend on choosing the correct diameter for the idealised fibre image.

Table 3.6: Results of centre-finding tests using cross-correlation with an idealised image.

Test criterion	Metric	Result
Offset fibres	Mean radial shift in result	0.18 μm
Measurement noise	RMS scatter for static fibre	0.05 μm
Speed	Mean time to calculate result	98.7 ms

Figure 3.22 shows the response of the cross-correlation method to non-optimal reference circle diameters, for two different test fibres at the centre and edge of the camera field, respectively. The edge fibre shows a more dramatic response due to increased optical aberrations at the edge of the field. In both cases, the 'correct' fibre diameter is easy to spot, and matches the expected nominal value.

The plots show that it is far better to underestimate the fibre diameter by a few percent than to overestimate it. This makes sense, as cross-correlating an oversized circle will mean the overhanging area around its edge will be correlated with dark, noise-dominated pixels outside of the fibre core image.

WEAVE will use Molex/Polymicro 'FBP 085/102/120' fibre, which is specified at a core size of $(85 \pm 3) \mu\text{m}$ ¹¹. This places an upper limit of 3.5% on variation either side of the nominal core diameter for all fibres on the field plate.

¹¹Data taken from the WEAVE fibre system final design review documentation.

PART I: THE WEAVE FIBRE POSITIONER

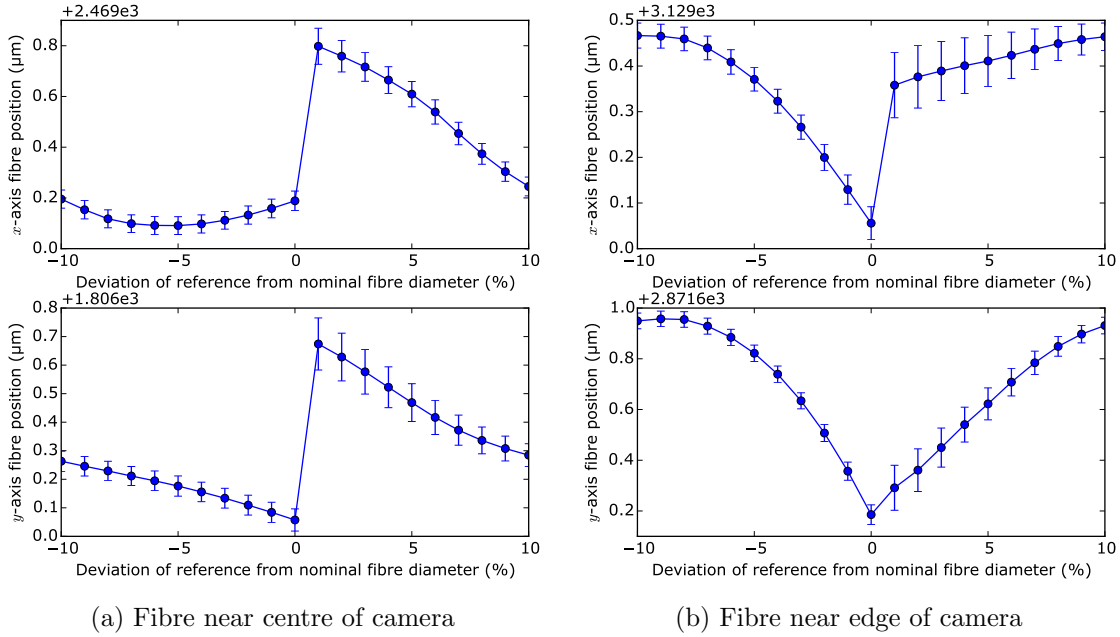


Figure 3.22: Two sets of cross-correlation results for the same set of 100 images of a static fibre, but varying the diameter of the reference circle. Underestimating the optimal diameter is better than overestimating it, as oversized circles become sensitive to asymmetric camera PSFs (and noise) surrounding the fibre. The consequences of a non-optimal diameter are less severe near the centre of the camera’s FoV.

Looking at Figure 3.22a, we can say that if the reference circle diameter is set around 3.5% smaller than the nominal core size then errors in fibre centre estimates should not exceed 0.25 μm as a result of core diameter variation. This comes with the proviso that measurements are done on fibres close to the centre of the camera’s field of view (as is planned), noting that errors could approach 1.0 μm for the worst-case scenario at the edge of the camera field.

3.6.7 Centre-finding technique D: The ‘biscuit cutter’ algorithm

A new algorithm has been developed with the goal of obtaining similar performance to the cross-correlation method above, but in a fraction of the time. Again, this approach makes use of prior knowledge of the fibre’s shape and size.

This technique is based on pixel masks, such as the one shown in Figure 3.23. These masks are stored arrays, the values in which reflect the area of the pixel bound by a perfect circle of the required diameter. These values are between 0 (fully outside the circle) and 1 (fully inside the circle). Much like cutting a circle from biscuit dough, multiplying this pixel mask with an image is equivalent to cutting a perfect circle of

PART I: THE WEAVE FIBRE POSITIONER

information from the image. The algorithm is referred to as the ‘biscuit cutter’ for this reason, with intentional avoidance of the more prevalent term ‘cookie cutter’ for the purpose of keeping things British.

Continuing the bakery metaphor, this technique can be thought of as trying to cut the heaviest biscuit from unevenly-rolled dough. Starting at a given point, one can cut the dough, weigh it, then put it back and cut the dough in a slightly different place before weighing it again. If the cut dough is heavier than before, then you continue cutting in the same direction until you reach a maximum, before changing direction and repeating the process to the point where you’ve found the optimum cutting position. The same is true for applying the pixel mask to different areas of a fibre image, except it is moved by a single pixel in x or y , and the masked pixel values are summed rather than weighed. This is a quick way to find the centre of the circular image to a precision of a whole pixel. To all intents and purposes this is the same as regular cross-correlation, so far.




Figure 3.23: A circular mask of an arbitrary diameter can be generated using the equation of a circle. The mask pixel values reflect the area bounded by the (green) theoretical circle, between a value of 0 and 1.

The procedure above has limitations. Apart from having the unacceptably poor precision of one pixel, summing many pixels over and over again is rather inefficient. The latter issue can be mitigated once we appreciate that, if the mask is moved by only one pixel, most of the pixels within the mask need not be re-summed; only the pixels at the edge of the mask are changing in value. The difference in value of a given pixel after masking is therefore

$$\begin{aligned}
 P_{new} &= P_{old} - IM_{old} + IM_{new} \\
 &= P_{old} + I(M_{new} - M_{old})
 \end{aligned}
 \tag{3.8}$$

where P is the masked pixel value, I is the image pixel value, and M is the mask pixel value.

PART I: THE WEAVE FIBRE POSITIONER

This means that if you already know the previous masked value of a pixel, then the difference between the mask in its old position and the mask in its new position is all you need to calculate the new masked value. Hence, the $(M_{new} - M_{old})$ term in equation 3.8 can be expressed as another mask entirely, with far fewer non-zero elements. For example, the ‘difference mask’ shown in Figure 3.24 is the result of subtracting a full circular mask (Figure 3.23) from a version of itself that has been shifted by one pixel in the negative y direction.

The advantage of using this mask is that only the first calculation needs to sum pixel values over the entire circle, then subsequent iterations only need to add the leading-edge pixels (shown in red) and subtract the trailing-edge pixels (shown in blue) from the previous sum total. A difference mask, like a full mask, can be pre-calculated and stored in memory. This speeds up operations considerably, particularly for large circles.




Figure 3.24: An initial application of the circular mask (top-left) to the raw fibre image (bottom-left) produces an initial value of the total ‘brightness’ of the masked image (second column); next, another mask can be used to calculate only the *difference* in brightness that results from a shift of the mask in a given direction (third column); summing the initial value and the difference value gives a new brightness for the shifted position, allowing comparison of which is ‘better’, i.e. which produces the highest total value.

In solving the problem of having only whole pixel precision, the principle of the difference mask presents a solution. Since a difference mask is the result of subtracting a full mask from a shifted version of itself, it is trivial to produce a difference mask for whatever shift you like, as mask values are calculated based on the equation of a circle about any given centre coordinate. Sub-pixel precision is therefore achieved by generating a catalogue of difference masks representing equal subdivisions of a

PART I: THE WEAVE FIBRE POSITIONER

pixel (e.g. $\frac{1}{10}$, $\frac{2}{10}$, $\frac{3}{10}$, etc.). The precision of the result is set by the number of pixel subdivisions, which is limited only by the memory required to store the masks. The number of masks required for an accuracy of an N^{th} of a pixel is proportional to N^2 , since we are subdividing a single pixel in two dimensions. The actual number of masks stored is

$$\text{No. of masks} = 4(2N - 1)^2 \quad (3.9)$$

for all four shift directions $(\pm x, \pm y)$. Or, if we are happy to exploit the symmetry in x and y and rotate the masks as necessary,

$$\text{No. of masks} = (2N - 1)^2 \quad (3.10)$$

The memory required to store the masks is proportional to the number of masks, the diameter of the circle we are trying to locate, and the number of bits required to store each mask value. A conservative scenario for the case of WEAVE might be a precision of 0.02 px (0.08 μm), a fibre diameter of 23 px, and 64 bit double precision floating point numbers. The memory required for mask storage is therefore

$$\begin{aligned} \text{Memory required} &= ((2 \times 50) - 1)^2 \times (23 + 1)^2 \times 64 \\ &= 45 \text{ MB} \end{aligned} \quad (3.11)$$

This is perfectly acceptable on a modern-day computer. Furthermore, resolutions much finer than 0.02 px are likely to be dominated by camera noise, so are of little benefit.

In summary, the procedure for taking a measurement is as follows.

1. Make an initial rough estimate of the fibre centre location using a fast method such as centroiding.
2. Apply a circular mask to the image at the initial coordinate, multiply the mask with the image, and add these values to get an initial ‘sum of brightness’.
3. Apply a single-pixel difference mask for every direction $(\pm x, \pm y)$ and update the sums for each.
4. If any of the new sums are larger than the old sum, update the current best estimate of the fibre centre coordinate accordingly.
5. Repeat steps 3–4 until the best single-pixel estimate has been found.
6. Switch to a sub-pixel difference mask, apply the masks which correspond to a shift in every direction from the current coordinate, and calculate new sums for each.

PART I: THE WEAVE FIBRE POSITIONER

7. If any of the new sums are larger than the old sum, update the current best estimate of the fibre centre coordinate.
8. Repeat steps 6–7 until the best sub-pixel estimate has been found. This is the final measurement.

The results for a biscuit cutter measurement of sample fibre images, using ‘fast’ C-type code compiled with Cython [47][48], are summarised in Table 3.7. As with the cross-correlation method, an advantage of this approach is that it doesn’t rely on threshold values being chosen. Its dependence on the correct mask size is also the same as for the cross-correlation method (refer to Figure 3.22, p. 74).

Table 3.7: Results of centre-finding tests using the biscuit cutter technique.

Test criterion	Metric	Result
Offset fibres	Mean radial shift in result	0.17 μm
Measurement noise	RMS scatter for static fibre	0.05 μm
Speed	Mean time to calculate result	0.18 ms

3.6.8 Comparison of centre-finding techniques

A comparative summary of the results from the previous sections is shown in Table 3.8. All four techniques perform similarly in terms of fibre offset and noise tolerance, with the values being too small to be of concern anyway. The biscuit cutter algorithm appears to be the best performer overall.

Table 3.8: Similar performance is seen for all four candidate centre-finding methods, with the biscuit cutter being the best overall performer, particularly in terms of speed.

	Centroid	Parabolic fit	Cross-correlation	Biscuit cutter
Offset fibres	0.20 μm	0.21 μm	0.18 μm	0.17 μm
Measurement noise	0.07 μm	0.05 μm	0.05 μm	0.05 μm
Speed	0.14 ms	2.39 ms	98.7 ms	0.18 ms

3.6.9 Absolute accuracy

We saw in Table 3.8 that speed is the only factor that differs significantly between the four centre-finding algorithms under evaluation. But the metrics in this table do not capture anything to do with the actual accuracy of the algorithms.

Recall that the principal purpose of the centre-finding operation is to estimate the true centre of the fibre in the image. If we plot the results of 100 measurements of a static fibre for all techniques then we see two things: i) their tolerance of camera noise, given by the scatter (already reported in the previous sections), and ii) the difference in the actual calculated centre position, given by the mean result.

Figure 3.25 shows the disagreement between techniques for two test fibres. We see differences of up to $0.8\ \mu\text{m}$ when the fibre is close to the centre of the measurement camera's field of view (e.g. Figure 3.26a), but this greatly increases when the same fibre is moved towards the edge of the field (e.g. Figure 3.26b), with differences of up to $3.1\ \mu\text{m}$. A larger error at the edge of the field points to optical aberrations in the camera as the cause, further evidenced by a radial bias in the fibre's intensity profile. Although fibres are expected to be close to the centre of the field in most cases, discrepancies of a few microns are more than enough to raise concern, especially when compared to WEAVE's total error budget of just $8.0\ \mu\text{m}$.

The plots in Figure 3.25 also reveal other things: The cross-correlation and biscuit cutter methods are in good agreement, which would be expected because the biscuit cutter is essentially a cross-correlation operation with shortcuts. The slight difference between these two is likely caused by the fixed resolution of the biscuit cutter method ($0.01\ \text{px}$ in this case) and/or floating point precision errors. The methods also appear in a different sequence for each plot, so do not correlate with each other.

Without an advanced measurement setup, we cannot say for sure which method is closest to the true centre of the fibre. The best we can do, therefore, is to consider the theoretical performance of the different approaches.

We have already established that a simple centroid cannot provide a robust estimate in the presence of non-uniform fibre intensity profiles because of its bias towards brighter areas of an image. Centroiding is also more sensitive to camera noise.

Fitting a parabolic function relies on the accurate computation of a threshold value to exclude the wings of the image profile (see Section 3.6.5). It is also a 1D operation, so does not make use of all the information in the image. As with the centroid, there is an inherent bias towards brighter parts of the profile, which is what we are trying to avoid. Indeed, the plots in Figure 3.25 show an agreement with the centroid rather than the other methods.

PART I: THE WEAVE FIBRE POSITIONER

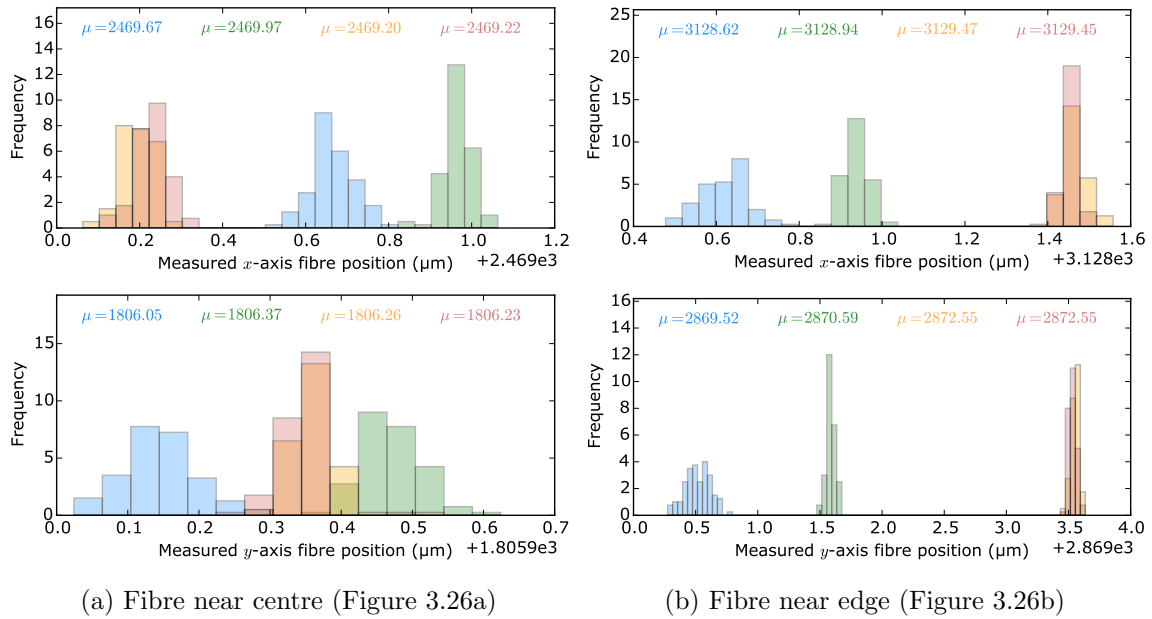


Figure 3.25: There is significant disagreement between centre-finding techniques when applied to the same 100 fibre images. The disagreement is larger at the edge of the camera’s field of view (3.1 μm) than near its centre (0.8 μm). Techniques are centroid (blue), parabolic fit (green), cross-correlation (orange), and biscuit cutter (pink).




Figure 3.26: Increased asymmetry is seen in the fibre image intensity profile when closer to the edge of the camera’s field of view. The intensity contours around the fibre image point to the centre of the field, indicating optical aberrations are the cause. Green markers show orientation of magnified profile [46].

We are left with the simple argument that finding the point of greatest encircled brightness should be the best estimate of the true fibre centre, in theory. For this reason, the biscuit cutter method should be the technique of choice due to its speed advantage over full 2D cross-correlation. Both of these techniques depend on prior knowledge of the fibre core diameter in the image, but it has been shown that they are robust to the worst possible mismatch between the expected and actual diameters.

3.6.10 Centre-finding summary

Four different image processing techniques for finding the centre of back-illuminated WEAVE fibres have been evaluated: i) a simple centroid; ii) a parabolic fit; iii) cross-correlation with an ideal circle; and iv) the new ‘biscuit cutter’ algorithm, which is a fast version of (iii).

All techniques perform similarly in terms of camera noise tolerance and their response to off-axis fibre illumination at the spectrograph. However, there were significant differences in the calculated fibre centre coordinates in the tests that were performed, especially when fibres were not near the centre of the camera image.

The new biscuit cutter algorithm is recommended as the better technique for WEAVE, given its theoretically-superior accuracy over the centroid and parabolic fit methods, both of which favour locally brighter areas of an imaged fibre core.

The biscuit cutter is extremely fast, taking on average only 0.18 ms to compute a result. The total fibre measurement time is given by

$$T = N(t_{expose} + t_{transfer}) + t_{pre-process} + t_{process} \quad (3.12)$$

where N is the number of camera frames to be stacked (co-added), t_{expose} is the exposure time per frame, $t_{transfer}$ is the time taken for the image to arrive in local memory from the camera, $t_{pre-process}$ is the time taken to pre-process the image and crop around the fibre being measured, and $t_{process}$ is the time taken to calculate the fibre’s centre coordinates. Filling in the known values so far,

$$T = N(9.0 + 33) + 12 + 0.18 \text{ ms} \quad (3.13)$$

which allows up to $N=2$ frames to be stacked within the 100 ms available measurement time, for a $\sqrt{2}$ reduction in camera noise. The total time is therefore

$$\begin{aligned} T &= 84 + 12 + 0.18 \\ &= 96 \text{ ms} \end{aligned} \quad (3.14)$$

PART I: THE WEAVE FIBRE POSITIONER

Further speed improvements may be possible if: a) the initial frame used for pre-processing was sub-sampled; and b) subsequent images used an area-of-interest readout around the fibre being measured. This would allow many more frames to be stacked for a further reduction in noise. This is not considered necessary, given the already low measurement noise.

The final estimate of the centre-finding error using the biscuit cutter technique is $0.5\ \mu\text{m}$, as shown in Table 3.9. This includes two sources of error: i) inaccuracy of the approach in general, taken as half of the discrepancy seen earlier in Figure 3.25a, and ii) errors caused by core diameter variation between fibres, as analysed in Section 3.6.6. Both of these require that the fibre being measured is close to the centre of the measurement camera’s field of view, which is already the planned procedure.

The error in the biscuit cutter result introduced by camera noise for a stack of two frames is negligible, at $0.04\ \mu\text{m}$. Assuming a mask resolution of $\frac{1}{50}$ px, the quantisation error introduced by this is also negligible at $0.08\ \mu\text{m}$.

Table 3.9: Summary of estimated errors in the biscuit cutter centre-finding technique.

Error source	Error magnitude
General inaccuracy in result	$0.4\ \mu\text{m}$
Fibre core diameter variation	$0.25\ \mu\text{m}$
RSS total	$0.5\ \mu\text{m}$

3.7 Proof of concept for passive field plate fiducials

This section describes some preliminary work on a new type of fiducial marker for the WEAVE field plate. Two techniques for embedding non-illuminated fiducials into an Invar plate were trialled and evaluated in terms of image quality. An on-camera illumination system was designed, prototyped, and shown to produce good contrast images that are suitable for measuring in the same way as back-illuminated fibres.

3.7.1 Comparison with 2dF

Field plate fiducials are used on the AAT’s 2dF positioner [18], from which the WEAVE positioner inherits many features. 2dF has fiducials that are actually a grid of 21 illuminated optical fibres, recessed into its field plate.

PART I: THE WEAVE FIBRE POSITIONER

While actively-illuminated markers make imaging with a camera easy, the difficult process of embedding fibres into a field plate is best avoided if possible. To this end, and given that we wish to substantially increase the number of fiducials for WEAVE in order to reduce coordinate mapping errors, an alternative approach was investigated.

The proposed WEAVE field plate will have a regular grid of more than 1300 markers (on a 10.0 mm pitch) that are not illuminated themselves, but from above by a spotlight on the positioner robot.

3.7.2 Fiducial system prototyping

Some proof of concept testing was carried out on sample markers, illuminated by a ring of wide-beam LEDs on the measurement camera assembly. A prototype of this on-board spotlight is shown mounted to the prototype measurement camera assembly in Figure 3.27.

Two methods of creating dark markers on an Invar plate were investigated: i) laser burning; and ii) paint-filled pockets machined by a CNC mill. The results of these processes are shown in Figure 3.28. There is a clear difference in the edge definition, laser burning being inferior. Laser burning also provides a lower contrast image. These factors are likely to make accurate measurements from the image difficult.

Of all the shapes one could make the markers, filled circles are considered preferable on the grounds that the same image processing routines used for fibre measurement could be used for these as well, following an inversion of the image pixel values.




Figure 3.27: An LED spotlight for illuminating field plate fiducial markers was designed and prototyped, and provides good quality illumination without any reflections of the individual LEDs in the camera images.

PART I: THE WEAVE FIBRE POSITIONER




Figure 3.28: Laser burning markers into Invar (bottom) produces an inferior result in terms of edge definition and contrast when compared to milling a pocket with a CNC machine and filling it with black paint (top).

The value range of an 8-bit image of a sample fiducial marker was measured at 161, compared to 242 for a typical back-illuminated fibre. This reduced contrast, coupled with the potential issue of reflections or scratches on the field plate, could challenge the measurement system. But further investigation in this area requires a more representative sample of the final field plate, manufactured using appropriate processes, and this is beyond the scope of this work.

3.7.3 Field plate survey errors

The WEAVE field plates will each have a fiducial reference grid pitch of 10.0 mm. The AUTOFIB-2 fibre positioner [16], commissioned on the WHT in 1994, has proven that this density of points allows good calibration of robot gantry flexure and straightness. The specific location of points will be measured with a high-accuracy coordinate measuring machine (CMM) and this data stored for each field plate. The known grid, assumed to never change apart from a small amount of (correctable) radial temperature-induced expansion, can then be fitted to observed fiducial marker positions from the robots' measurement cameras.

Simulations by other WEAVE team members¹² have shown that the RMS of the residuals for fitting a reference grid, with $\sim 60\%$ of the points obscured by deployed fibres, is $3.1\ \mu\text{m}$ ¹³. This accounts for all expected uncertainties relating to the measurement and recovery of the fiducial grid pattern in order to calibrate robot axis straightness and flexure.

¹²Acknowledgement goes to Ellen Schallig and Ian Lewis for their work on this.

¹³Data obtained from WEAVE final design documentation.

In summary, the robot axis coordinate system can be calibrated with an RMS error of $3.1\ \mu\text{m}$ by quickly surveying all fiducial markers not obscured by fibres, for any zenith angle. This is the standard mode of operation before beginning any new field reconfiguration, and this estimated error is a significant contribution to the total measurement error for the positioner.

3.8 Measuring axis position uncertainty

This short section presents test results to estimate the uncertainty of the robot axis motor positions at any given moment. The uncertainty comes from position encoder noise, but also from communications lag between the encoders and the positioner control software, during which time the axes may have moved slightly. The combined uncertainty for the x - and y -axis is found to be $0.69\ \mu\text{m}$, which is almost entirely generated by noise in the encoders.

3.8.1 Sources of axis position uncertainty

All fibre position measurements from the measurement cameras must be referenced to the position of the robot at the time the fibre image was captured, so that the actual position of the fibre on the field plate can be found.

When the positioner control software requests the current position of the robot axes, there is a delay between requesting this information and a response being sent. During this time, the axes may have moved slightly due to their servoing about a particular point, which will introduce an error in the reading. Noise in the encoder system will also be present in this error.

3.8.2 Position uncertainty testing

As a way to estimate axis position uncertainty, the prototype positioner robot was set to move to a specific point and to stay there. A static back-illuminated fibre was placed beneath the measurement camera, and its apparent position in the camera image was compared to the latest encoder reading received by the control software from the linear axis drive. The RMS difference between these two values over many measurements gives an estimate of the position uncertainty for a single axis.

Results from the test are shown in Figure 3.29, revealing a scatter of $\sigma = 0.49\ \mu\text{m}$ for the y -axis. Assuming the same performance for the x -axis, the RSS total uncertainty becomes $0.69\ \mu\text{m}$.




Figure 3.29: Comparing the difference between the apparent position of a static fibre in a measurement camera image and the latest axis encoder values over 100 runs revealed an RMS uncertainty of $0.49\ \mu\text{m}$ per linear axis.

The stability of the motor servo loop becomes important here, as a badly tuned motor is likely to move more during encoder reading delays. Since the motor tuning parameters were provided by the motor manufacturers, it is assumed that the tuning is already close to optimal. A measurement of the axis position with the motor drive disabled and clamped in place showed that the total error above is dominated by the pure encoder noise ($\sim 0.40\ \mu\text{m}$), which further proves the stability of the control loop.

3.9 Estimating total measurement error

Having evaluated the main stages of measuring the position of a single fibre on the WEAVE field plate, the average error of doing so can be estimated. This is the overall ‘measurement error’ of the fibre positioning system, which will be a substantial contribution to the overall positioning error.

3.9.1 Review of measurement error sources

The many sources of error in measurement are perhaps best reviewed by walking through the measurement procedure, starting with an image of a fibre from a measurement camera, and following the analysis of this image all the way ‘up’ to an absolute measurement of position. The following paragraphs attempt to do this, each source of error being given a name for the purpose of referring to the source later on.

PART I: THE WEAVE FIBRE POSITIONER

Given an image from the measurement camera, the fibre within the image is found and its centre estimated using a suitable technique. The result of this measurement will not exactly be at the centre of the fibre (*fibre centre-finding error*). The noise from the measurement camera detector will also introduce a shift to this measurement (*measurement camera noise*). The severity of these effects will depend on the chosen centre-finding technique.

Next, the fibre centre coordinates are referenced to the point on the camera detector that the θ -axis passes through. This point will have a particular uncertainty (*θ -axis reference error*). This offset is converted to a real-world distance using the pixel scale of the camera, which was calibrated against a fixed optical reference. The quality of the optical reference and the technique used to infer the pixel scale will both introduce an error to this conversion (*pixel scale error*).

The encoders in the robot motors are interrogated to get a snapshot of their position. There is a delay between requesting this information and a response being sent, during which time the motors may have moved. Assuming that no axes are on a journey to some new destination, the difference between the reported position and the actual position (*axis position uncertainty*) is a combination of instability in the motor servo loop, the noise in the encoder systems, and the communications lag between the control computer and the motor controllers.

The real-world fibre position with respect to the measurement camera detector is now mapped to the robot axis coordinate system as a function of the current θ -axis angle. Any inaccuracy in the calibration of the camera alignment with respect to the θ -axis will introduce an error here (*camera alignment error*). The alignment calibration depends on the quality of the optical reference used for the alignment and the tolerance on the alignment algorithm.

At this point we have the fibre position, in real-world distance units, relative to the origin of the robot's coordinate system. This origin has no specific definition, which is why the field plate survey now becomes important. The field plate survey would have been carried out prior to the field reconfiguration in order to map the robot coordinate system to the field plate coordinate system, subtracting any flexure effects in the robot gantry structure. The survey is done by imaging a grid of fixed fiducial markers on the field plate and fitting the known grid pattern accordingly. Uncertainties in the original measurement of the grid pattern and errors in the fitting process (*field plate survey error*) will therefore add to the total uncertainty of any measurement by the camera system.

The position of the imaged fibre is now known with respect to the field plate.

3.9.2 Final measurement error calculation

Table 3.10 shows all the measurement error sources identified in this chapter. Together, these give an estimate of the total measurement uncertainty for the WEAVE positioner metrology system, at $3.3\ \mu\text{m}$.

Some errors get worse when fibres are further from the θ -axis reference point (which is nominally in the centre of the camera's field of view). For these, a distance of $1.0\ \text{mm}$ is assumed, which is conservative given that fibres within the gripper jaws should be very close to the reference point by design. Other reported errors are maximum values rather than RMS values, so are generally overestimated.

We see that the total measurement error is dominated by plate survey errors and reference grid fitting. This comes from a comprehensive analysis by other group members (Schallig and Lewis) and includes many of the more serious uncertainties such as robot axis straightness, non-orthogonality, and gravity-induced flexure.

We can conclude that the overall RMS measurement error of $3.3\ \mu\text{m}$ is good news for the positioner, being substantially smaller than the required placement accuracy of $8.0\ \mu\text{m}$. This will be combined with axis repeatability data in the following pages, to give a total estimate of the RMS positioning error for a single pick-and-place operation.

Table 3.10: Summary of measurement error sources from throughout this chapter and resultant estimate of total measurement error. Section numbers are included for reference.

Source	Value	Comments	Section
Fibre centre-finding error	$0.50\ \mu\text{m}$	Using biscuit cutter method	3.6.10
Measurement camera noise	$0.04\ \mu\text{m}$	Using biscuit cutter method and stacking two frames	3.6.10
θ -axis reference error	$0.32\ \mu\text{m}$	Scatter over multiple tests	3.4.6
Pixel scale error	$0.44\ \mu\text{m}$	For a fibre $1.0\ \text{mm}$ from θ -axis reference	3.4.6
Axis position uncertainty	$0.69\ \mu\text{m}$	Combined x - and y -axis uncertainty of latest encoder readings	3.8
Camera alignment error	$0.35\ \mu\text{m}$	For a fibre $1.0\ \text{mm}$ from θ -axis reference	3.4.6
Field plate survey error	$3.1\ \mu\text{m}$	Simulation results for various error sources	3.7.3
RSS total	$3.3\ \mu\text{m}$		

3.10 Section summary

This section has thoroughly described, tested and analysed WEAVE's entire fibre position measurement system, arriving at a final estimate of $3.3\ \mu\text{m}$ for the system's total RMS measurement error.

Arriving at the final error estimate involved a series of tasks: WEAVE's fibre measurement camera was redesigned, prototyped and tested; an automatic camera calibration scheme was devised, designed, prototyped, tested, and its performance analysed; the proposed WEAVE fibre back-illumination unit was evaluated, resulting in recommendations to improve its design but otherwise confirming its suitability; various image processing techniques for inferring a fibre's position were compared, resulting in the development of the 'biscuit cutter' algorithm, a new method with improved accuracy and speed over the best alternatives; the uncertainty of the reported positions of the robot axes was calculated through testing; and finally, the total contributions of all of these error sources were summed to produce a final estimate of measurement system error.

The measurement error was actually found to be dominated by uncertainties in the calibration of the WEAVE field plate, which have been modelled by other WEAVE team members. These are included in the final error sum.

Overall, we can be satisfied that the WEAVE metrology system should be accurate enough to achieve the required fibre positioning accuracy of $8.0\ \mu\text{m}$.

PART I: THE WEAVE FIBRE POSITIONER

4 Verification of commercial motor performance

This section describes performance testing of the prototype WEAVE positioner robot described in Section 2. The robot axes are shown to perform within manufacturer specifications, indicating that positioning at a suitable accuracy for WEAVE is achievable. Improvements to positioning speed are also presented, improving the prospects of fast and accurate positioning in the final instrument.

The work in this section will be published in the paper “Verification of commercial motor performance for WEAVE at the William Herschel Telescope”, as part of the proceedings of the Multi-Object Spectroscopy in the Next Decade conference in Santa Cruz de La Palma, Spain, 2015 [49].

4.1 Repeatability testing

General axis performance was measured with a simple test of repeatability. The purpose of this test was mainly to verify the manufacturer’s published values of positional accuracy for the chosen axes.

4.1.1 Linear axis repeatability

The prototype WEAVE robot was instructed to position 500 mm away from its origin and then return. This routine was repeated 500 times and the apparent location of a static reference fibre extracted for each cycle. The scatter in this measurement provides an upper limit for positioning accuracy of the axis. Each measurement was taken 10s after the axis controller reported the axis as being in position, meaning that mechanical vibrations from the move should not be present. It follows that the scatter in measured position, which was taken relative to the camera itself, should give a picture of the stability of the servo loop and position encoders on the axis.

PART I: THE WEAVE FIBRE POSITIONER




Figure 4.1: An upper limit of $2.55\ \mu\text{m}$ on the scatter of axis position over 500 moves away from and returning to a fixed point was measured (including $\sigma = 0.17\ \mu\text{m}$ measurement noise). This is better than the manufacturer’s published value of $5.0\ \mu\text{m}$.

The result (Figure 4.1) shows a scatter of just $2.55\ \mu\text{m}$, which is well within the published accuracy of $5.0\ \mu\text{m}$. An even distribution shows no obvious creep in error, meaning that encoder readings remain stable over 500 moves.

4.1.2 Gripper release repeatability

The gripper jaws were tested for repeatability when releasing fibres. A fibre was picked up, placed back down, and released. This was repeated 200 times, using two fibre buttons at two different locations along the y -axis. The difference in fibre location between gripped and released states was measured for every cycle.

A significant shift is expected whenever a fibre button is released. This is because the fibre button vane (the handle with which it is gripped) may not be perfectly normal to the base of the button. In theory, this lateral shift is constant for each fibre button. The intention is therefore to subtract this shift from the target placement location. This strategy relies on repeatable gripper motion (i.e. both jaws always working in unison), which is tested in this experiment.

The results (Figure 4.2) show a mean lateral shift of order $10\ \mu\text{m}$ when releasing a button. Crucially, however, the scatter of this offset is sub-micron. This indicates that, as hoped, subtraction of the mean offset for each fibre button during positioning is feasible.

PART I: THE WEAVE FIBRE POSITIONER




Figure 4.2: A sub-micron scatter in the measured fibre shift when releasing fibre buttons shows excellent gripper repeatability. This makes subtraction of the mean shift a feasible strategy to account for imperfections in the button geometry. The test was done for two different locations along the prototype robot field plate.

4.1.3 Axis performance vs. measurement error

It is important to distinguish the above reports of axis performance from the total measurement error presented in Section 3.9.

The total measurement error defines the minimum placement accuracy you can ever be sure to achieve. In this regard, it is the most critical number for the WEAVE positioning system, as accurate positioning is impossible without accurate feedback.

Axis performance, on the other hand, shows the ability of a motor or axis to be useful in our mission to position. Here, repeatability is the governing factor. What the results above have shown is that it is reasonable to expect consistent results when placing a fibre on a target. This is especially true for the gripper release repeatability, as any lateral shift in the fibre button when it is let go must be predicable ahead of time. In other words, good axis performance increases the probability that fibres will be placed correctly *on the first attempt*.

4.2 Improvements to positioner timings

The expected field reconfiguration time for WEAVE has been significantly improved by parallelising the pick-and-place operation. This section describes these improvements, without which the target field reconfiguration time for WEAVE was unfeasible.

4.2.1 Timing requirements

The WEAVE positioner has a top-level requirement of being able to reconfigure an entire field in 60 min or less. The maximum available time per average pick-and-place move is therefore calculated by dividing the total available time by the number of pick-and-place operations, split between the two positioner robots. Here we assume that the division of labour between the robots is 90 % efficient, and that every fibre must be re-parked at its retractor before being placed on a new target, and hence moved twice¹. For WEAVE's nominal 960 fibres and eight guide fibres,

$$\begin{aligned} \text{Max. average move time} &= \frac{2 \times \left(\frac{\text{Total time}}{\text{Number of fibres}} \right)}{2 \times 0.9} \\ &= \frac{2 \times \left(\frac{60}{968} \right)}{1.8} \\ &= 3.3 \text{ s/fibre} \end{aligned} \tag{4.1}$$

Therefore, an average of only 3.3 s is available for an entire pick-and-place operation. This includes axis motion and fibre position feedback measurements. The required RMS placement accuracy of 8.0 μm must also be achieved.

4.2.2 Simulating an average move

In lieu of a fully-assembled positioner, the prototype robot system was used to simulate pick-and-place moves for the purpose of timing tests. The lack of an x -axis on the prototype was assumed to be insignificant, since this would be moving simultaneously with the y -axis anyway.

Calculating the average pick-and-place move time requires the average move distance. If we assume an even distribution of fibres across the field plate area, then the average distance to any fibre button from the centre of the field plate is the radius of a circle encompassing half of the total science field area, given by

$$r = \sqrt{\frac{R_{fp}^2}{2}} \tag{4.2}$$

where r is the radius of the positioning distance from the centre of the field plate and R_{fp} is the radius of the WEAVE science field.

¹This approach is in line with the positioner's design documentation.

PART I: THE WEAVE FIBRE POSITIONER

Since fibres originate from the retractors around the perimeter of the plate at a radius R_{ret} , the average move distance is

$$\begin{aligned}d_{average} &= R_{ret} - \sqrt{\frac{R_{fp}^2}{2}} \\ &= 282.3 - \sqrt{\frac{205^2}{2}} \\ &= 137.3 \text{ mm}\end{aligned}\tag{4.3}$$

It follows that pick-and-place moves covering a distance of 137.3 mm, using the prototype WEAVE positioner robot, should give an accurate indication of the expected average move time.

4.2.3 Unsuitability of a sequential timing scheme

The most straightforward operating mode for the positioner would be to carry out each of the stages of a pick-and-place operation in sequence, moving to the next stage only when axes are within a specified distance from their commanded destinations. The pick-and-place stages are as follows, from an initial state where the z -axis is raised and the gripper jaws are open.

1. Move the x -, y - and θ -axis to the ‘pick’ position (call this position A), all three axes moving simultaneously to arrive at their target positions at the same time.
2. Check the position of the fibre below with the fibre measurement camera.
3. Lower the z -axis.
4. Close the gripper jaws.
5. Raise the z -axis.
6. Move the x -, y - and θ -axis to the ‘place’ position (call this position B), at the same time as ascertaining the exact fibre position in the closed jaws, refining B as necessary.
7. Lower the z -axis.
8. Open the gripper jaws.
9. Confirm correct placement with the fibre measurement camera.
10. Raise the z -axis.

PART I: THE WEAVE FIBRE POSITIONER

The mean time taken for the above sequence was measured over 100 cycles as 4.5s with the axis motors at their maximum practical velocity and acceleration². A breakdown of the total time taken in terms of each positioning stage is shown in the top portion of Figure 4.3. This duration exceeds the maximum allowable value of 3.3s and would result in an overall field reconfiguration time 20 min longer than a nominal spectrograph exposure. A better solution must be found.

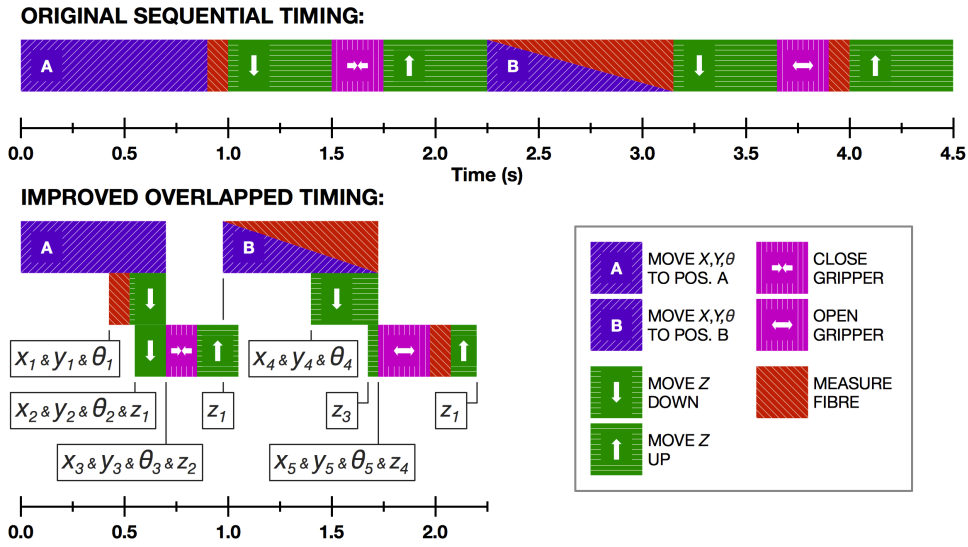


Figure 4.3: Safely overlapping parts of the pick-and-place sequence greatly reduces the average move time. These charts show the timings for a simple sequential scheme (top) and an overlapped scheme (bottom) with various position tolerances defined in Table 4.1. The durations shown are mean values measured over 100 pick-and-place cycles.

Table 4.1: Definitions and example values for the various tolerances shown in Figure 4.3.

Label	Description	Nominal value
x_1, y_1, θ_1	x -, y - and θ -axis within ‘sloppy pick’ tolerance	Pos. $A \pm 2.0$ mm, 1.0°
x_2, y_2, θ_2	x -, y - and θ -axis within ‘coarse pick’ tolerance	Pos. $A \pm 200$ μ m, 0.5°
x_3, y_3, θ_3	x -, y - and θ -axis within ‘fine pick’ tolerance	Pos. $A \pm 100$ μ m, 0.1°
x_4, y_4, θ_4	x -, y - and θ -axis within ‘sloppy place’ tolerance	Pos. $B \pm 2.0$ mm, 0.5°
x_5, y_5, θ_5	x -, y - and θ -axis within ‘fine place’ tolerance	Pos. $B \pm 5.0$ μ m, 0.01°
z_1	z -axis above ‘safe height’	>20 mm above f.p.
z_2	z -axis within ‘fine pick’ tolerance	\downarrow pos. ± 100 μ m
z_3	z -axis at ‘hover height’	≤ 0.5 mm above f.p.
z_4	z -axis within ‘fine place’ tolerance	\downarrow pos. ± 5 μ m

²The velocities and accelerations for the motors were determined empirically, by increasing them until the motion controllers reported a following error, then reducing slightly.

4.2.4 An improved timing scheme

The average pick-and-place time was reduced from 4.5 s to 2.2 s by carefully overlapping axis movements and implementing a range of axis position tolerances. Figure 4.3 shows a graphical representation of the improved scheme from laboratory timing tests, with the different tolerance levels explained in Table 4.1.

The improved scheme is based on the fact that the tightest positional tolerances (e.g. 5 μm) are only necessary the moment before a fibre is placed on the field plate. Furthermore, very loose tolerances (e.g. 2 mm) are acceptable whenever the gripper jaws are high enough above the field plate to clear all obstructions (fibres, buttons, etc.). Tolerances become gradually tighter as the z -axis descends, leading to a hierarchical tolerance system that significantly reduces the time spent waiting for fine servo motion. Other salient points of this approach are as follows.

- The z -axis can begin descending before the x -, y - and θ -axis have finished moving.
- Once the gripper jaws are below the minimum ‘safe height’, the z -axis can continue descending provided the x -, y - and θ -axis are close enough to their target positions not to cause a button or fibre collision.
- Fibre buttons can be gripped (picked) with a looser x -, y - and θ -axis tolerance than when releasing (placing).
- The x -, y - and θ -axis can begin to move to a new position as soon as the gripper jaws are above the minimum safe height.
- The only time that the axes need to be within the tightest of tolerances is when a fibre button is being held very close to the field plate, the moment before being placed down in their target position.

An average move time of 2.2 s means a total reconfiguration time of 39.4 min. There is a little more time to be gained by distinguishing between moves where fibres are being re-parked and moves where fibres are being placed on a science target. The former can be done with a looser tolerance. If we take this to be the same as the ‘fine pick’ tolerance in Table 4.1, then the measured timings indicate a reconfiguration time of 36.7 min.

The increased positioning speed achieved with the new timing scheme is a promising result. It has made meeting the positioner performance requirements feasible, where previously this was impossible.

4.3 Prototype robot limitations

The results presented in this section are for a prototype robot system, hence they are not representative of the final positioner assembly. Moreover, no facility for changing the gravity vector is available, so robot performance when the telescope is not at zenith cannot be simulated. Given the WEAVE project schedule, this situation has been unavoidable. Nevertheless, it is important to remember that significant risks remain as the project enters the integration phase.

Specifically, the measured pick-and-place times do not include any time spent waiting for the gantry structure to settle after a move. This cannot be tested until a gantry has been built to the final mechanical design. That said, it seems reasonable to remain optimistic given that the total estimated positioning time now occupies less than 65% of that available, leaving a generous margin in which to trade off axis acceleration with placement accuracy.

A further assumption is that fibres will always be placed correctly on the first attempt. The positioner's ability to perform well will depend on this, as re-placing many fibres would come at a significant time cost and result in failure to reconfigure the field during a nominal exposure period.

4.4 Section summary

This section has reported the results of laboratory testing of the COTS robot axes that will be used for the WEAVE fibre positioner. Analysis has found that the axes perform better than manufacturer specifications and are therefore suitable for achieving WEAVE's 8.0 μm fibre positioning requirement.

An important distinction has been made between the accuracy of the positioner robot and the accuracy of the fibre measurement system; poor robot axis repeatability does not directly affect the achievable accuracy of the fibre positioning system, but does increase the probability of fibres not being placed adequately on the first attempt, therefore incurring a time penalty.

The anticipated field reconfiguration time for WEAVE has been improved from ~ 80 min (non-compliant with the 60 min requirement) to < 40 min, by efficiently overlapping certain robot actions. However, field reconfiguration times are likely to increase significantly when mechanical vibrations are considered. This will likely only be known once the final positioner has been built.

5 Final analysis of fibre positioning performance

This section brings the previous two major sections together, reviewing our estimates of fibre measurement error and robot axis performance and considering what this means in terms of the bigger picture of fibre placement accuracy and speed in WEAVE.

What we can conclude is that fibre placement accuracy and field reconfiguration time are related, with better accuracy being achievable if particularly poor fibre placements are re-attempted by the robot. A scheme for managing this is proposed.

5.1 Field reconfiguration speed

Let us first revisit the estimated speed of the positioner reported in Section 4.2. It was shown that, with efficient motor timings, a field reconfiguration time of 36.7 min should be achievable. This occupies only 61 % of the available 60 min window.

The estimated field reconfiguration speed assumes that WEAVE's two positioner robots spend 10 % of their time waiting for each other to complete a move. Furthermore, the timing tests only considered the time taken for all robot axes to come to a stable rest *according to the motor controllers*. No additional time was allowed for the settling of the structure connecting the robot gantry to the field plate itself, because the structure used in the laboratory was not comparable to that of the final instrument.

Assuming the measured pick-and-place move times are accurate, the spare 39 % portion of the available field reconfiguration window will be available for spending in two ways: i) waiting for the positioner structure to settle, so that extra errors are not introduced; and ii) re-placing fibres that are deemed to have been placed poorly.

5.2 Measurement system accuracy

Now let us return to the accuracy of the measurement system, reported in Section 3. Measurement uncertainty is critical for the WEAVE positioning system, as accurate positioning is impossible without accurate feedback. The estimated accuracy of $3.3\ \mu\text{m}$ RMS is substantially smaller than the required fibre placement accuracy of $8.0\ \mu\text{m}$ RMS, so we can be quite confident that WEAVE will be capable of detecting significant fibre placement errors.

5.3 Robot repeatability and error tolerance

Finally, let us revisit the performance of the robot axes reported in Section 4. The linear axis and gripper repeatability are summarised below in Table 5.1.

Table 5.1: Summary of robot motor repeatability measurements.

Parameter	Value (σ)
Linear axis repeatability	$2.55\ \mu\text{m}$
Gripper release repeatability	$0.9\ \mu\text{m}$

We can estimate the accuracy of a single pick-and-place fibre positioning operation by summing the linear axis repeatability (assuming these are the same for the robot's x - and y -axis), the gripper release repeatability, and the measurement system uncertainty of $3.3\ \mu\text{m}$. Assuming these can be added in quadrature,

$$\begin{aligned}\sigma_{total} &= \sqrt{2.55^2 + 2.55^2 + 0.9^2 + 3.3^2} \\ &= 5.0\ \mu\text{m}\end{aligned}\tag{5.1}$$

From this we can conclude that, in an ideal world, every fibre should be placed well enough on the first attempt. Naturally we have made many idealised assumptions, not least treating all error sources as statistically random, hence the pure RMS positioning accuracy is likely to be worse on the final instrument.

Since we may have enough time to repeat some fibre placements, we can think about setting a maximum tolerance on placement error. The plot in Figure 5.1 is based on the $5.0\ \mu\text{m}$ RMS positioning uncertainty calculated above. What this plot shows is the probability that a fibre will be placed within a certain error tolerance *in a single attempt*.

PART I: THE WEAVE FIBRE POSITIONER




Figure 5.1: Data on the repeatability of the WEAVE robot axes has been used here to estimate the probability that a fibre will be placed with an error less than or equal to a specified tolerance, in a single attempt.

For example, setting an error tolerance of $10.0\ \mu\text{m}$ (equal to 2σ) would result in 95% of moves being acceptable after one pick-and-place attempt, i.e. 5% being re-attempted.

Optimising a tolerance such as this, either with respect to the average field reconfiguration time, or with an on-the-fly calculation *during* field reconfigurations, may be a good way to increase the average quality of WEAVE’s data.

5.4 Section summary

This section has combined what we know about the accuracy of the WEAVE measurement system and the prototype positioner robot, in the context of the real-life operation of the positioner.

The trade-off between positioning accuracy and field reconfiguration speed has been restated, highlighting that our fixed reconfiguration time of 60 min allows a strategy to be adopted to make best use of this time. A simple approach has been suggested, whereby the fibre positioning error tolerance is dynamically adjusted based on the predicted time remaining in the current configuration. This could guarantee an upper limit on positioning error, while working to make the overall positioning error lower than this whenever possible.

Optimally utilising WEAVE’s 60 min field reconfiguration window will be a real opportunity to increase the signal-to-noise ratio of the WEAVE data, and therefore improve its scientific impact.

PART I: THE WEAVE FIBRE POSITIONER

6 Conclusions and prospects

Part I of this thesis has reported on work that was carried out during the final design stages of the WEAVE fibre positioner system, following its preliminary design review and leading up to its final design review.

The broad aim of this research was to retire the major technical risks associated with the design of the COTS-based pick-and-place fibre positioner system, better estimating its performance in order to answer two main questions:

1. Are the proposed off-the-shelf robotic axes capable of positioning WEAVE's optical fibres accurately enough to meet the scientific needs of the instrument?
2. If the accuracy is sufficient, then can this accuracy be achieved quickly enough to meet the required field reconfiguration time?

These questions, which are somewhat linked, have been answered with a combination of engineering design, prototyping, testing and analysis.

The proposed robotic axes from Schunk GmbH have been shown to perform well enough to meet the $8.0\ \mu\text{m}$ RMS fibre positioning accuracy requirement of WEAVE, with an RMS repeatability estimate of just $5.0\ \mu\text{m}$ for all axes combined. Since the positioning will be a closed-loop process, the story is only complete if the position feedback system is also up to the task. Following rigorous development and testing, the RMS accuracy of the WEAVE fibre measurement system has been estimated at $3.3\ \mu\text{m}$. This is also sufficient to achieve our required $8.0\ \mu\text{m}$ positioning accuracy.

Pick-and-place reconfiguration of all of WEAVE's fibre units within the requisite 60 min window has been shown to be feasible, with $5.0\ \mu\text{m}$ RMS robot accuracy being achievable in as little as 37 min.

Importantly, the positioner's estimated performance does not account for vibration of the robot gantry structure or increases in errors caused by changes in the local gravity vector, neither of which could be tested with the prototype system.

A changing gravity vector caused by telescope pointing should be insignificant to the counterbalanced robot axes, but may alter the calibration of the measurement

PART I: THE WEAVE FIBRE POSITIONER

system. The in-situ camera calibration procedure proposed in this thesis would allow quick compensation of gravity effects should the need arise. This could be done for every new pointing at a modest time cost of <1 min.

The risk of increased positioning errors due to mechanical vibrations is somewhat lessened by the robot's measured accuracy being well within WEAVE's requirements. By returning to the plots of estimated spectral SNR in Figure 1.6 (p. 33) we are reminded that fibre positioning tolerance could be increased without catastrophically affecting signal strength. Indeed, doubling the fibre offset tolerance to $16.0\ \mu\text{m}$ would only reduce the estimated SNR by 6.5%.

A further relief is that the estimated 37 min field reconfiguration time would allow robot moves to be slowed down by $\sim 40\%$, therefore reducing the severity of mechanical vibrations.

All things considered, it seems reasonable to conclude that the WEAVE fibre positioner will be able to fulfil its function without significantly affecting the quality of the instrument's data. If the flexure issue is tightly controlled in the final assembly, as it may well be, then the WEAVE positioner has the capacity to far exceed performance expectations.

Three novel processes have been developed in order to meet the required accuracy and speed of the positioning system:

1. **An in-situ camera calibration scheme:** A system for automatically aligning, focussing, and obtaining the pixel scale of the measurement cameras has been designed. Proper camera calibration is key to achieving consistently accurate measurements.
2. **The 'biscuit cutter' algorithm:** A new image processing algorithm has been developed that can infer the location of the centre of a WEAVE fibre core with superior accuracy and speed. This is another critical aspect of the fibre position feedback process.
3. **Overlapped robot motion:** A robot motion control scheme has been created that safely overlaps pick-and-place motor commands via a hierarchical system of motor position tolerances. This scheme has practically halved the predicted field reconfiguration time; without it, WEAVE would incur a severe and unacceptable time penalty of ~ 20 min per field.

The outcomes of this research assisted in the WEAVE fibre positioner system passing its final design review in January 2015.

Part II

Development of a new low-voltage
piezoelectric motor for tilting spine
fibre positioners

1 Introduction to tilting spine technology

Tilting spine positioners enable fast and accurate field reconfigurations of hundreds or thousands of optical fibres simultaneously. Each fibre is held within a rigid ‘spine’ that can be tilted in any direction, thus moving the fibre tip. Each spine is independently actuated by a simple and compact motor mechanism, meaning that the technology can support very high fibre densities.

Tilting spine technology was created by the Australian Astronomical Observatory (AAO) in the early 2000s for the purpose of the 400-fibre FMOS–Echidna instrument on the Subaru telescope (Figure 1.1) [22][23]. Since then, Echidna technology has progressed through several stages of development [50][51][27][24] and will soon be incorporated into the 2436-fibre 4MOST spectrograph on VISTA [26]. Despite the many refinements made, however, the nature of spines and their actuators remains essentially unchanged.

The work reported in this thesis looks at more radical ways to improve the design of the Echidna motor mechanism so that its scientific capabilities are improved.

This introductory section aims to do four things: i) explain what current-generation tilting spine positioners look like and how they work; ii) highlight the technology’s main weaknesses; iii) describe the 4MOST instrument in more detail, to give context; and iv) define the specific goals of this research.

1.1 Nomenclature

Since this chapter will be describing new developments to an existing technology, there are several historical (and new) naming conventions that should be described here for later reference.

A number of different terms and phrases actually mean the same thing, and some very similar phrases actually mean different things. Table 1.1 aims to remedy this.

PART II: A NEW TILTING SPINE MOTOR

Table 1.1: Nomenclature and naming conventions for this chapter.

Term	Explanation
Tilting spine technology	The name given to the general positioner concept developed by the Australian Astronomical Observatory in which a fibre is held by a long tube that is tilted by actuators; refers to the existing technology in all its forms
Echidna technology	Used interchangeably with ‘tilting spine technology’
FMOS–Echidna	A 400-fibre tilting spine positioner, built by the Australian Astronomical Observatory for the Fiber Multi-Object Spectrograph (FMOS) instrument on the Subaru telescope and commissioned in 2008 [52]; the first instrument to use tilting spine technology
Drive waveforms	The underlying (sawtooth) waveforms used to drive the spine actuators and achieve the stick–slip motion of spines
Drive signals	Used interchangeably with ‘drive waveforms’
Waveform voltage	The peak voltage of a drive waveform, often bipolar (i.e. $100\text{ V} \Rightarrow \pm 100\text{ V}$)
Waveform amplitude	Used interchangeably with ‘waveform voltage’
Spine actuator(s)	Usually refers to the individual piezoelectric element(s) within a spine motor
Spine motor	The active assembly that holds and moves a spine, comprising piezo actuator(s) and a magnetic mount
Actuator assembly	Used interchangeably with ‘spine motor’
Spine mount	The non-active part of the spine motor that holds a spine, comprising a magnet within a metal ‘cup’; the spine ball sits in/on the cup
Magnetic mount	Used interchangeably with ‘spine mount’
Cup assembly	Used interchangeably with ‘spine mount’
Spine assembly	The assembly comprising a carbon fibre tube, ball and counterweight, with a fibre held inside; a passive assembly that is mounted within a spine motor
Spine	Used interchangeably with ‘spine assembly’




Figure 1.1: FMOS-Echidna (a), so called because of its resemblance to the spiny Australian mammal of the same name (b), was the first instrument to use AAO tilting spine technology. Each fibre is held in a rigid tube (spine) that can be tilted by a few degrees in all directions. FMOS-Echidna has 400 spines with a spacing of 7.2 mm. The image in (a) shows the tips of the spines, hence starlight would approach from the bottom of the image. The area shown is about the size of a football.

1.2 The anatomy of a tilting spine

The basic anatomy of a tilting spine is shown in Figure 1.2. Each fibre is held in a rigid carbon fibre tube that can freely tilt about a ball joint at one of its ends. The tube is balanced with respect to the pivoting ball and is itself a passive assembly.

The spine is inserted through the centre of a cylindrical piezoelectric actuator assembly (‘the motor’) and magnetically held in place, as shown in Figure 1.3. The motor can apply small displacements to the surface of the ball, resulting in discrete angular steps of the spine and therefore shifting the spine tip. The actuators are fixed to a solid mounting surface in a regular triangular grid of a certain spacing (pitch), as illustrated in Figure 1.4.

1.2.1 Spine length, pitch and patrol area

The length of the spines, the pitch of their actuators, and the ‘patrol area’ that each fibre can access are set to meet the needs of the instrument, and of course have some limitations.

Spine pitch is set by the physical size of the telescope’s focal surface and the number of fibres that the instrument requires; the number of fibres is most often limited by the size of the instrument’s spectrograph(s), likely rounded down to arrive at a geometrically-pleasing total for efficiently packing the telescope’s field of view.

PART II: A NEW TILTING SPINE MOTOR




Image courtesy of the Australian Astronomical Observatory

Figure 1.2: Spines are passive assemblies comprising a lightweight rigid tube cemented into a pivoting ball. They are balanced about this ball and therefore operate regardless of the local gravity vector. The fibre passes through the counterweight, through the tube, and is terminated in a ferrule at the tip of the spine. The specified length of a spine is taken from the centre of its ball to its tip.




Image courtesy of the Australian Astronomical Observatory

Figure 1.3: The spine is held within a cylindrical actuator assembly by means of a magnetic ‘cup’ mount that is cemented to a piezo tube. The piezo tube is held by a solid mounting structure behind the focal surface of the telescope. This structure has an oversized clearance hole for the spine, allowing it to tilt by a few degrees in all directions.




Image courtesy of the Australian Astronomical Observatory

Figure 1.4: Actuators are tightly packed across the field in a triangular grid. A printed circuit board (shown in green) provides access to the actuator electrodes from around the edge of the field. This view is from behind the spines and shows only the spine balls (gold) mounted in the actuator cups (blue).

PART II: A NEW TILTING SPINE MOTOR

The patrol area of a spine is set by a combination of its length and the maximum tilt permitted by the motor. A spine’s patrol area is naturally circular, and the radius of this circle relative to the fibre pitch will influence the target allocation efficiency of the final instrument.

Figure 1.5 shows the desirable case to have a patrol radius equal to or exceeding the pitch multiplied by a factor of $\frac{2}{\sqrt{3}}$, with which any point across the field can be accessed by at least four and as many as seven spines, providing excellent fibre allocation yields.




Figure 1.5: Tilting spine positioners have a large ‘patrol radius’. This diagram shows seven spines (black spots) viewed from their tips. By setting the patrol radius greater than $\sim 1.15 \times$ the pitch, any point in the field can be accessed by either four (blue), five (green), six (yellow), or seven (red) different fibres.

Spine length has to be sufficient to achieve the required fibre patrol radius, but will also affect the optical throughput of the entire system. An unavoidable weakness of tilting spine positioners is that their tilt introduces non-telecentricity and defocus losses (explained later), which have an approximately-quadratic relationship with the spine’s tilt angle [24]. It follows that, for a given pitch, longer spines can achieve the same patrol areas at smaller tilt angles and hence with smaller losses.

Lengthening spines may reduce optical losses, but it also has the effect of increasing the minimum step size of the positioner, thus reducing the achievable positioning accuracy. A rather complex trade-off exercise is therefore required to find the optimum spine length for a given instrument. FMOS–Echidna spines were set at a length of 140 mm. More recent designs (such as for 4MOST, described later in Section 1.6) have been able to increase this to 250 mm while retaining acceptable positioning performance.

1.3 Actuation principle

The electro-mechanical design of the existing Echidna technology is able to move spines in discrete steps by means of a stick–slip mechanism.

1.3.1 Piezo tube actuators

An Echidna positioner spine is centred around a type of cylindrical piezoceramic actuator historically used for steering electron beams in scanning tunnelling microscopy. These piezo tubes have four electrodes covering their outer wall and a single electrode covering their inner wall, as illustrated in Figure 1.6. The tube can be made to bend by applying a potential across its walls.

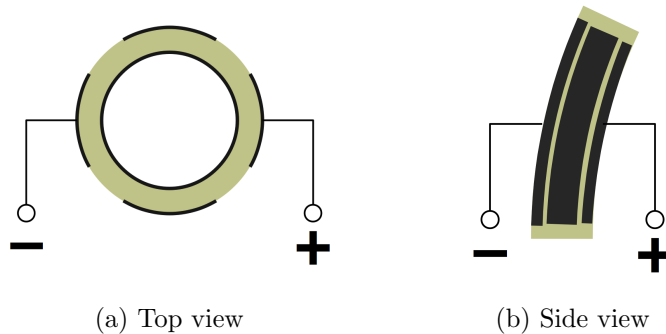


Figure 1.6: Piezo tubes have four equally-spaced electrodes on their outer wall and a single electrode on their inner wall. Applying voltages of opposite polarity to opposing electrodes will cause the tube to bend. Displacement is exaggerated.

1.3.2 The Echidna ‘stick–slip’ mechanism

Echidna positioners exploit the bending modes of piezo tube actuators in order to produce a dynamic stick–slip mechanism between the spine ball and the ball’s magnetic mount on the actuator.

The stick–slip mechanism exploits the fact that most material pairs have a lower frictional coefficient when rubbing together than when static. By varying the acceleration (i.e. force) of an actuator, we can selectively enter the regime of either static or kinetic friction. Cycling between these regimes while alternating the direction of the force results in a progressive ‘ratcheting’ of the materials as they become ever more displaced with respect to each other.

PART II: A NEW TILTING SPINE MOTOR

Figure 1.7 shows how Echidna’s stick–slip mechanism is achieved by driving a piezo tube with a simple sawtooth waveform: A steady ramp bends the tube, tilting the spine slightly as the ball ‘sticks’ to the magnetic mount on the actuator; the drive signal then quickly returns the tube to its starting position, the high acceleration of this motion causing the ball and mount to ‘slip’ with respect to each other and the inertia of the spine keeping it, more or less, in the displaced position.

The result of a single waveform cycle is a discrete angular step of the spine, and hence a discrete lateral step of the spine tip and the fibre. The actuator has returned to its neutral state at this point and therefore requires no further power.

A single step at the spine tip is of order $10\ \mu\text{m}$ in size, and many waveform cycles in quick succession (of order 100 Hz) can produce a smooth motion of the spine. The quadrant electrode format on the piezo tube means that this motion is in one of four selectable directions, as illustrated in Figure 1.8.




Image courtesy of the Australian Astronomical Observatory

Figure 1.7: The tilting spine ‘stick–slip’ mechanism is a simple way to achieve discrete displacements of the spine tip using a single piezo tube actuator and a straightforward sawtooth drive waveform: Starting from a neutral position (a), the ramp of the sawtooth steadily bends the actuator and tilts the spine within it (b); the actuator then suddenly returns to its neutral position, leaving the spine behind (c).

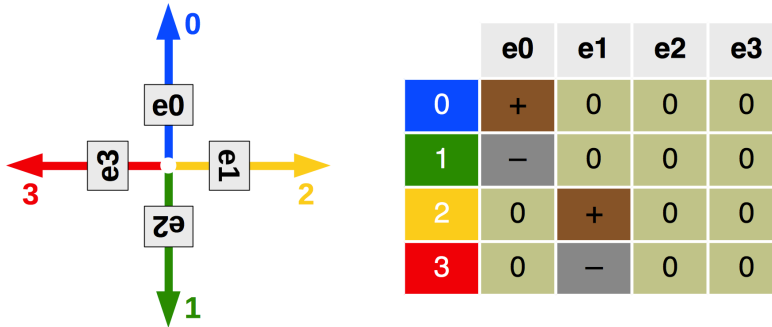


Figure 1.8: Four electrodes on the piezo tube actuator (e0, e1, e2, e3) and two waveform polarities (+ and -) yield spine movement in four orthogonal directions (0-3), depending on the waveform connections. The signal combinations shown in the table keep electrodes 2 and 3 constantly grounded in order to simplify the electrical architecture.

1.4 Spine control

When moving a spine, there are three electrical settings: i) which electrode is connected to which drive signal, thus selecting one of four orthogonal move directions; ii) the amplitude of the drive waveforms, which sets the size of the steps; and iii) the frequency of the waveforms, which sets the speed of the move. These settings are all handled by the positioner control electronics and software. The software also decides how best to resolve the intended move into the available component directions (including how many steps to take in each), and how to avoid collisions.

1.4.1 The control loop

Spine positioners operate within an iterative closed-loop control system. The differences between the desired and actual positions of the fibres are calculated, then a corrective move of every spine is executed and the position error is calculated again. The process repeats until all fibres are sufficiently close to their targets, or after a set maximum number of iterations.

Echidna spines have no built-in position feedback system, so fibre locations are measured via a camera looking at the back-illuminated fibre tips. One thing to note is that every spine must complete its move before a positioner-wide iteration is considered complete. The fibre view camera then images the entire field at once, and all fibre locations are extracted from the image. This takes a significant amount of time (say ~ 10 s for a few thousand fibres), so the number of required moves is important. This system also means that move times are set by the spine that has to travel the farthest.

PART II: A NEW TILTING SPINE MOTOR

Spine moves are calculated so that spine collisions are avoided, sometimes having to stop and wait for a neighbouring spine to get out of the way. An advantage of the Echidna technology is that collisions cause no damage to the spines or actuators, because the stick–slip mechanism simply slips until the obstruction is removed. This means very close fibre approaches can be attempted without risk.

1.4.2 Spine characteristics

Tilting spine actuators require drive voltages (i.e. sawtooth wave amplitudes) of ~ 150 V to achieve an appreciable displacement of the fibre tip [29].

The step size of a spine scales linearly within the operable voltage range: Drive amplitudes of ~ 150 V produce step sizes of ~ 50 μm , and drive amplitudes of ~ 75 V produce step sizes of ~ 10 μm ¹. This has led to a two-stage approach to positioning: a ‘coarse’ move stage at a higher voltage, for fast moves over long distances, and a ‘fine’ move stage at a lower voltage, for when close to the target fibre position.

1.4.3 Spine calibration

Every spine is calibrated for each of the coarse and fine modes, so that the control software knows how far a fibre will move in one step. This is done by driving the spines in each of their four movement directions for a fixed number of steps and then calculating the mean resultant step size (and exact angle of movement). These ‘calibration vectors’ are stored for each fibre and used by the control system to calculate the number of steps required to arrive at a given destination.

Coarse move waveform frequencies are ~ 100 Hz (i.e. 100 steps of ~ 50 μm per second), giving a fibre movement speed of a few millimetres per second and therefore allowing a nominal ~ 20 mm patrol diameter to be traversed in just ~ 4 s.

Fine move waveform frequencies may be reduced to < 10 Hz to avoid possible non-linearities caused by spine resonance.

1.4.4 Electronics system

The task of the electronics system is to provide waveforms of a particular amplitude and frequency to the correct electrodes on the spines’ piezo actuators. But in order for this to be feasible for high numbers of spines, some shortcuts are necessary.


The control circuit for a single spine is shown in Figure 1.9a. A series of high-voltage switches are used to connect two of the piezo electrodes to either the positive

¹Data obtained from the AAO, regarding their current tilting spine prototypes.

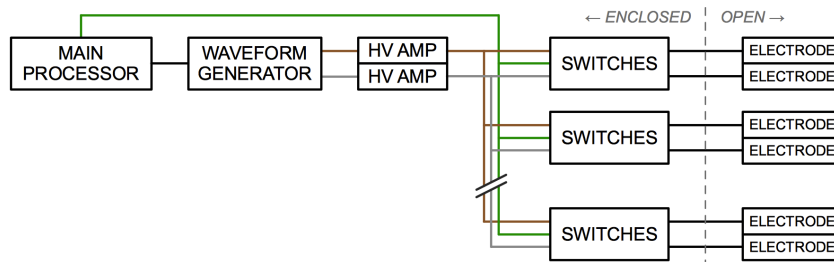
PART II: A NEW TILTING SPINE MOTOR

or negative drive waveform, or to ground. The remaining two electrodes are grounded permanently. This arrangement—having only one active electrode for each direction instead of driving opposite electrodes in a bipolar format—sacrifices some efficiency, but means that only two signal connections must be made to each actuator across the positioner. Without this design, it would not be feasible to produce circuit boards with enough tracks to reach every actuator across the field.

The high signal amplitudes demanded by the spine actuators call for specialist power amplifiers to drive them. These amplifiers tend to be expensive and bulky (Figure 1.10). They can, however, provide enough current to drive many actuators together. This is the reason that switches are used to re-route drive signals rather than changing the waveforms themselves. Figure 1.9b shows the resultant control architecture for the positioner, with one pair of drive waveforms (one positive, one negative) being shared by many (sometimes hundreds) of actuators. The consequences of this ‘shared waveforms’ architecture are serious and will be discussed shortly.



(a) Single spine control circuit



(b) Multiple spine control architecture

Figure 1.9: The existing tilting spine control architecture routes two ‘master’ waveforms to many spines using an array of switches. Waveforms must be shared in this way because the high-voltage amplifiers required to produce them are large and expensive. Hundreds of spines will often be driven by a single amplifier stage.




Figure 1.10: High voltage amplifiers tend to be large, their size dominated by the heat sinks they require. This heat sink is from an existing spine control system, where the pair of amplifiers attached to it can drive more than 600 spines. The coin shown is an Australian 5 cent piece (diameter ~ 20 mm).

1.5 Desirable improvements to tilting spine technology

In considering the finer points of tilting spine technology, we can identify some technical factors that would significantly improve it, either in terms of its weaknesses when compared to competing positioners, or in terms of the more challenging aspects of designing and building a full-size tilting spine system.

1.5.1 The curse of high voltages

The high-voltage nature² of the spine drive waveforms is perhaps the most hated ‘feature’ of the tilting spine technology from an engineering standpoint, and the cause of many an operational headache. It also severely affects spine performance. These issues are explained below.

1.5.1.1 Performance issues

The ‘shared waveforms’ control architecture described in Section 1.4.4 means that hundreds of actuators must be driven *at the same frequency and voltage*. This is a major design compromise, because performance variations between actuators cannot

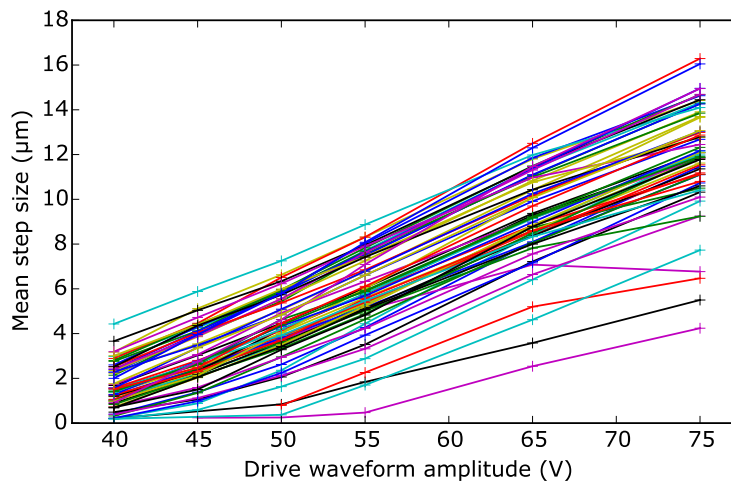
²Note that the term ‘high-voltage’ is used here in a relative sense, with respect to the order ~ 10 V potentials usually found in electronics systems. It does not refer to the various international standards that tend to consider ‘high voltages’ to be $\gg 1$ kV.

PART II: A NEW TILTING SPINE MOTOR

be accounted for. Instead, these waveform parameters must be chosen to ‘please all’, or more likely, to please the poorest-performing actuator in the group.

Sharing waveforms like this has a pronounced effect on the fine positioning mode, because setting the drive amplitude too low will cause the weaker spines to not move at all. Consequently, an amplitude must be chosen that is high enough to guarantee that all spines move at least a small amount. *Ipsa facto*, the more sensitive actuators produce larger steps and the average accuracy of the entire positioner is worsened.

The extent of this problem can be seen in Figure 1.11. The apparent chaos in this plot only emphasises the point: Having to pick a single drive amplitude results in a huge range of step sizes across multiple spines.



Courtesy of Scott Smedley / the Australian Astronomical Observatory

Figure 1.11: Spine step sizes can vary wildly for the same actuator drive amplitude. This data was obtained from the AAO’s 64-spine AESOP fibre positioner prototype (introduced later). Each of the 64 spines is represented by a line. Note how some spines reach their lower operational limit at 50–55 V, which means the global amplitude has to be set higher than this. This particular prototype system had its fine mode amplitude set at 75 V, which was sufficient for the relatively loose accuracy requirements of the project. Even so, this setting yields step sizes in the large range of 4–16 µm.

1.5.1.2 Operational issues

Once a high-voltage instrument is commissioned, staff must be trained in maintaining it safely. The various safety systems may also make fault-finding (either during or after commissioning) significantly harder.

1.5.1.3 Engineering issues

High voltages tend to greatly increase the size of an electrical system: Cables and connectors must be properly rated and are often bulky as a result, making cable runs impractical; all exposed conductors, including those on circuit boards, must be generously spaced (more so if being used at altitude); low-voltage control circuits must be isolated from their high-voltage load, requiring special opto-coupled switching devices; small quiescent currents in power supplies and amplifiers end up producing excessive heat and require large heat sinks; and systems must be interlocked and fail-safe, often calling for a separate safety systems controller.

It follows that removing hazardous voltages results in a smaller, leaner and cheaper control system that is easier to design.

1.5.2 Tilt-induced optical losses

Figure 1.12 illustrates how tilting fibres away from the telescope pupil introduces non-telecentricity and defocus. The non-telecentricity causes focal ratio degradation (FRD) at the exit of the fibre due to azimuthal scrambling of the light entering at an angle, which means that the spectrograph collimator is likely to be over-filled. The result of this is an overall throughput loss in the instrument. The defocus of the tilted fibre also causes a loss in throughput, although this is generally much less significant ($\sim 10^{-1}\times$) than the tilt-induced FRD [24].




Figure 1.12: The tilt angle of a spine, θ , introduces unavoidable optical losses due to two phenomena: i) focal ratio degradation caused by non-telecentricity between the fibre and the incoming beam (grey dashed line); and ii) defocus of the fibre tip, d . The yellow cones illustrate the acceptance angle (numerical aperture) of the fibres. Tilt and defocus are exaggerated.

PART II: A NEW TILTING SPINE MOTOR

A spine’s optical loss is a major drawback of the technology to the eyes of anyone deciding between competing fibre positioning technologies. Indeed, this was a major factor in the choice of a radial arm positioner over a spine positioner for the planned Dark Energy Spectroscopic Instrument (DESI) [31].

Being able to increase spine length, while retaining suitable positioning performance, is highly desirable.

1.5.3 Mechanical crosstalk

A subtle but potentially serious effect that has been observed in past AAO experiments³ concerns vibrations from the movement of spines inducing a sympathetic movement of *other*, non-moving spines across the field. It is believed that the motion of the actuators parallel to their mounting surface may be to blame here.

This is exacerbated by a control architecture that routes the same sawtooth waveforms to many actuators, because the vibrations are increased enormously when hundreds of piezo tubes are oscillating in lock-step.

Measures that reduce the risk of this phenomenon in future spine positioners would be valuable.

1.6 The 4MOST instrument

The 4-metre Multi-Object Spectroscopic Telescope (4MOST) is a new, wide-field, massively-multiplexed fibre-fed spectrograph facility being developed for the Cassegrain focus of the European Southern Observatory’s (ESO’s) VISTA telescope [26].

4MOST will use a tilting spine positioner to configure 2436 optical fibres across a 2.5° diameter field. The fibres will feed one high-resolution ($R \approx 20\,000$) and two low-resolution ($R \approx 5000$) spectrographs covering the visible wavelength range.

The 4MOST project is being introduced here because it is a contemporary example of an instrument perfectly suited to an Echidna-style positioner. The project has driven many of the AAO’s recent refinements of the technology [24], and the instrument is therefore a good yardstick against which we can consider more radical improvements that result from the work reported in this thesis.

4MOST is currently approaching the end of its preliminary design phase and is expected to begin operations by the end of 2020.

³This issue was communicated informally and has not been published.

1.6.1 4MOST science goals

In 4MOST’s most recent overview paper (de Jong et al., 2014) [26], its main science drivers are listed as galactic archaeology, high-energy physics, galaxy evolution and cosmology. Much like WEAVE in Part I of this thesis, it will also provide spectroscopic follow-up of large area survey instruments such as Gaia. In addition to this, 4MOST will operate a “flexible, general-purpose operations concept . . . surveys ranging from scarce targets all-sky to high target densities in specific areas”.

The clear desire for flexibility, especially in terms of varying field homogeneity, presents a strong case for the use of tilting spine technology on 4MOST. More critical, however, is 4MOST’s use of different spectrographs, with one third of its fibres being permanently assigned to a high-resolution observing mode. Tilting spine technology is the only parallel positioning solution that will ensure every point in the field is accessible by this subset of fibres, thanks to the spines’ large overlapping patrol radii.

1.6.2 AESOP: The 4MOST fibre positioner

A major part of 4MOST will be the Australian ESO Positioner (AESOP), a tilting spine fibre positioner built by the AAO. A 3D model of AESOP’s preliminary design is shown in Figure 1.13. Its specifications are summarised in Table 1.2.

Table 1.2: Summary of specifications and requirements for the AESOP fibre positioner.

	Low-resolution	High-resolution
Number of fibres	1624	812
Spine pitch	~9.6 mm	~16.6 mm
Spine patrol radius	~11.5 mm	
Spine length	250 mm	
RMS spine tilt [25]	~1.4°	
Closest approach	<1 mm	
Positioning accuracy	10 μ m (RMS)	

The fibre allocation yields for AESOP are examined at length in Saunders et al., 2014 [25], with results from mock catalogues indicating an overall target allocation completeness⁴ of 85.3% for low-resolution targets and 78.9% for high-resolution targets. As previously mentioned, these high target allocation yields come at the price of optical losses due to the tilt of the spines.

⁴Target allocation completeness is a measure of the achievable vs. desired object observations for a single pointing/exposure.

PART II: A NEW TILTING SPINE MOTOR





Image courtesy of the Australian Astronomical Observatory

Figure 1.13: The AESOP fibre positioner for the 4MOST instrument has 2436 tilting spines arranged in a hexagonal array ~ 0.5 m wide. Note the considerable space occupied by the (brown and beige) electronics enclosures on all four sides.



Courtesy of Will Saunders / the Australian Astronomical Observatory [25]

Figure 1.14: Untilted fibres produce a far-field output profile (black curve) within the speed of the spectrograph collimator (black vertical line, $f/3.7$). Maximally-tilted fibres (red curve) have the worst FRD and lose significant light outside the collimator speed limit. The overall loss for the whole positioner is given by the simulated RMS tilt (green curve), with a predicted collimator loss of 11.3%.

1.6.2.1 Tilt-induced optical losses in AESOP

The effect of different spine tilts on the far-field profile of light exiting the fibre in a spectrograph similar to 4MOST’s is shown in Figure 1.14. For the RMS tilt across all spines, the plot estimates a loss of $\sim 11\%$ of the fibres’ light at the collimator due to the induced FRD. This is perhaps not a show-stopping penalty, but it is significant. The approximately-quadratic relationship between tilt angle and throughput loss means that even a small increase to the spine length could have a substantial effect on a survey’s overall signal-to-noise ratio.

1.7 Competing technologies

Various fibre positioning technologies were reviewed as part of the main introduction to this thesis (p. 2.3.2) along with their associated strengths and weaknesses. Still, it is worth comparing aspects of their performance that might be matched or beaten by a better tilting spine design. Table 1.3 summarises the performance of other multi-actuator positioners (including tilting spine instruments) using data that has either been published or provided on request.


Table 1.3: Broad performance summary of existing parallel positioning technologies.

Technology	Type	RMS error	Moves	Reference
ES–CH θ – ϕ	Radial arm	$3.6\ \mu\text{m}$	6	Fahim et al., 2015 [32]
Cobra	Radial arm	$<5.0\ \mu\text{m}$	6	Fisher et al., 2014 [29]
AESOP	Tilting spine	$8.1\ \mu\text{m}^a$	6	(Informal contact)
FMOS–Echidna	Tilting spine	$\sim 6.3\ \mu\text{m}$	7	(Informal contact)

^aNote that the AESOP positioner is only designed to achieve $10\ \mu\text{m}$ RMS accuracy.

The best-performing positioner appears to be the ‘ES–CH’ (Spanish–Swiss) radial arm positioner led by the Universidad Autónoma de Madrid [32]. This technology was developed as a candidate positioner for the DESI project [31] and is at a prototyping stage. These positioners, which use miniature DC motors, are reported to achieve an RMS positioning accuracy of $3.6\ \mu\text{m}$ in six moves, or $3.9\ \mu\text{m}$ in just four moves.

Where details on the accuracy vs. moves relationship have been published, the available data points are displayed in Figure 1.15. These two plots are a good way to visualise the performance of any positioner, as they concern the two main metrics used in positioning requirements. These plot styles will become useful for the purpose of comparisons later.



(a) RMS positioning error


(b) Cumulative convergence (error $\leq 5.0 \mu\text{m}$)

Figure 1.15: These plots relate to two important performance metrics for fibre positioners in terms of the number of closed-loop iterations a technology needs to achieve either: (a) a general accuracy across the whole field; or (b) a fraction of fibres within a given accuracy. The data points shown are taken from the publications referenced in Table 1.3.

1.8 Goals of this research

The primary aim of this body of work, broadly speaking, was to make tilting spine technology more competitive against other positioner designs by improving its scientific capabilities. The secondary broad aim was to remedy some of the engineering challenges associated with building and operating a tilting spine positioner.

PART II: A NEW TILTING SPINE MOTOR

More specifically, and in light of the issues presented and discussed throughout this introduction, the following development goals were set.

- Reduce the drive voltage of spine motors, so that: a) spines do not suffer the performance penalty of sharing drive waveforms with others; and b) the control electronics system becomes smaller, leaner, safer, and easier to design. The specific aim was to reduce the voltage range from ± 150 V to ± 24 V (standard industrial DC voltage), with a strong desire to further reduce this to ± 15 V (standard op-amp supply voltage).
- Improve the closed-loop positioning accuracy of the spines to compete with current radial arm positioners, aiming to attain ≤ 5.0 μm RMS error in six moves, with an ultimate goal of ≤ 3.6 μm RMS error in six moves (based on the ES-CH radial arm positioner). This was most likely to come from a reduction of the minimum step size and therefore required a better motor mechanism. Improved spine control (through non-shared waveforms, as above) would also help here. A maximum tolerable error of 10 μm in six moves was also targeted.
- Increase the maximum spine length from the current 250 mm limit in order to reduce tilt-induced optical losses, while maintaining good positioning accuracy. The approximately-quadratic relationship between tilt and throughput loss means that only a small increase could be a significant improvement, but at the cost of increased minimum step size. No specific goal was set for this, as lengthening the spine will be limited by the achievable accuracy.
- Improve the overall open-loop repeatability of spine moves, with the intention of: a) producing faster convergence on a target, i.e. requiring fewer moves; and b) improving the prospect of tracking objects during observations. Object tracking requirements are highly dependent on instrument and telescope design, so no target was set other than to do better than the existing technology.
- Reduce the transfer of actuator forces along the plane of the actuator mounting surface, in order to reduce the risk of mechanical crosstalk between neighbouring spines. This hinted at moving away from the lateral bending motion of the existing piezo actuator design and towards something more efficient.

These research goals are all linked together in one way or another, thanks to the idiosyncrasies of the tilting spine concept. But the good news here is that one radical design improvement can yield benefits in all areas.

1.9 The ideal tilting spine positioner

Finally, it can be useful to consider what an idealised tilting spine system might look like, in order to get a feel for the trade-offs introduced by reality. Table 1.4 summarises an idealised spine in terms of its physical characteristics and performance. Not included in the table are the more generic practicalities such as cost, power consumption, maintainability, and so on.

1.10 Section summary

In this section we have seen how the ‘Echidna’ tilting spine technology works, and reviewed its strengths and weaknesses.

We have seen how, due to the high voltages that traditional Echidna actuators require, the performance of an entire tilting spine positioner is negatively impacted by many spine motors having to share the same drive waveforms.

The relationship between a spine’s tilt angle and the resultant optical throughput losses has been explained, highlighting the desirability of longer spines in order to reduce the average tilt angle.

The 4MOST project has been introduced as the next major instrument to use tilting spine technology. Its science objectives have been summarised, highlighting the need for a fibre positioner that can handle both homogeneous and clustered fields for a range of survey programmes.

The main goals of this research have been set, the salient points being to reduce the spine drive voltage, improve positioning accuracy, and maintain improved accuracy with even longer spines so that losses are reduced too. This called for radical changes to the design of the tilting spine motor.

PART II: A NEW TILTING SPINE MOTOR

Table 1.4: Characteristics of an ideal tilting spine positioner, and real life limitations.

Characteristic	Ideally	In reality
Minimum step size	Extremely small (e.g. sub-micron)	The imperfect stick–slip mechanism has a threshold, below which the actuator movements are too small to induce any spine motion at all, which limits the minimum achievable step size.
Step size repeatability	Zero scatter in spine response when driven with an identical signal, regardless of spine position	Piezo actuator hysteresis and tribological factors mean that moves will be different; also, imperfect spine balancing and ball irregularities can change performance based on the current spine position.
Gravity vector response	Spine performance unchanged by changing gravity vector	Changes in the local gravity vector (e.g. due to telescope pointing) have a small effect on performance due to imperfect spine balancing and a finite magnetic holding force in the actuator assembly.
Long-term performance stability	Stable despite the effects of time and temperature	Actuator age and environmental conditions can have a small but significant effect on spine behaviour.
Number of steps vs. distance moved	Directly proportional	Some non-linearity due to imperfect stick–slip mechanism and dynamic factors such as spine inertia.
Step size vs. drive voltage	Directly proportional	Non-linear for small signals due to actuator hysteresis, tribological factors, and non-infinite acceleration in mechanical ‘slip’ phase.
Frequency dependence	Spine performance independent of drive (stepping) frequency	Some frequencies yield better performance than others; resonant modes in the spine can impact repeatability at certain drive frequencies; higher frequencies (faster moves) are more prone to inertial effects.
Maximum spine length	Unlimited; longer spines reduce tilt-induced optical losses	Length is limited by mechanical factors such as spine resonance, inertia, and tolerable mass; longer spines also increase the minimum achievable step size.
Minimum spine pitch	Not limiting for specified science case	Spine motor assemblies have a minimum practical size that will limit the packing density of fibres.

PART II: A NEW TILTING SPINE MOTOR

2 Achieving low-voltage drive of tilting spines

This section discusses the possible strategies behind creating a low-voltage tilting spine motor. A key driver behind the commencement of this research was to make use of new and emerging technologies in the piezo industry to do this.

The advantages of using piezoelectric actuators are reviewed, along with their principal disadvantage of requiring high voltages. This shows that we need to look beyond classical piezo materials and manufacturing techniques in order to achieve our goals. Three core strategies will be defined in this regard.

2.1 The use of piezo actuators in micro-positioning

Piezoelectric actuators are unbeatable for many micro-positioning applications. They offer precise control of small displacements, fast response times, and high forces.

Appendix A (p. 223) provides some background on piezo actuators, reviewing piezoceramic material properties, traditional manufacturing methods, and piezoelectric motor actuation styles (including the stick-slip principle).

2.1.1 The inverse piezoelectric effect

Piezoelectric materials exhibit an accumulation of electrical charge when under mechanical stress. This phenomenon is called the direct piezoelectric effect. This effect also works in reverse, meaning that the presence of an electric field will induce stress in the material, resulting in a small displacement. This is called the inverse piezoelectric effect, and is the basis of all piezo actuators.

2.1.2 Piezo polarisation and modes

The vast majority of piezo actuators are made from synthetic ceramics. These ceramics are manufactured using a sintering (baking) process, which produces a bulk ceramic with randomly-oriented crystalline grains. These grains must be aligned so that the inverse piezoelectric effect occurs in a common direction. The alignment is achieved using very large electric fields, in a process called polarisation.

Polarisation of piezoceramics causes a permanent change¹ in their structure and, by extension, defines their mechanical response to the future application of electric fields in terms of the inverse piezoelectric effect. In other words, a piece of piezo material will move differently depending on the direction of an applied electric field relative to the direction of polarisation. The different combinations of these with respect to the geometry of the actuator are called piezo *modes*.

Piezo modes are defined with reference to a numbered axis system (Figure 2.1) representing three orthogonal directions in space and their associated axes of rotation. The direction of polarisation is always in direction 3.




Figure 2.1: Piezo materials use a numbered axis system to describe the direction of electric fields and displacements. It is equivalent to the Cartesian coordinate system. The direction of polarisation is always along axis 3.

¹The polarisation of a material is permanent as long as it is not further exposed to very high electric fields, in which case re-polarisation of the material can occur.

PART II: A NEW TILTING SPINE MOTOR

Two of the most common piezo modes are described below, along with a special case for cylindrical actuators. Refer to Table 2.1 for illustrations.

- **Longitudinal (d_{33}) mode:** This is perhaps the simplest mode. Plated electrodes on opposite faces of the actuator create an electric field in the same direction as the polarisation, resulting in a longitudinal extension of the actuator along axis 3.
- **Shear (d_{15}) mode:** This mode produces a shearing of the electrode faces, because the electric field is being applied orthogonal to the polarisation direction.
- **Cylindrical tubes:** This is not a piezo mode per se, but the geometry of a tube can produce a useful motion when electric fields are applied across its walls. The tube can be made to bend if the outer electrode is sub-divided, and this produces a lateral displacement of the tip that is dependent on the tube's length. This is the mechanism used in existing AAO Echidna positioners.

2.1.3 Typical voltages and displacements

The actuator displacement that results from a given electric field strength depends on the properties of the piezo material, but in any case is rather small for traditional piezoceramics. Table 2.1 gives example displacements for the common actuators.

The practical limitations of such actuators are made clear when we consider driving them at the voltages usually found on electronic circuit boards, such as the standard op-amp supply rails of ± 15 V. Here the largest theoretical displacement is for cylindrical actuators, and is only of the order ± 100 nm.

Traditional Echidna motors already use piezo tube actuators, and still require voltages of ± 150 V. This highlights the difficulty in designing a new spine motor mechanism that can match the coarse step sizes of the current system *and* run at voltages an order of magnitude lower.

2.2 Design approaches

Three main strategies were identified at an early stage as ways to realise a low-voltage tilting spine motor: i) exploit new piezo materials; ii) exploit new industrial processes; and iii) improve the efficiency of the motor mechanism. These were the cornerstones of this research, and are described in more detail in the following pages.

Table 2.1: The shape of a piezoceramic actuator has a substantial effect on the displacement it can produce. Longitudinal mode actuators are the least sensitive in terms of mechanical displacement, whereas a similar material used in shear mode will move much farther. A cylindrical actuator is a special case producing even larger displacements when used in its bending mode (as in the existing FMOS-Echidna motors). Piezoceramic charge constants, d , are highly dependent on the material compound, so all values are approximate and based on specifications from various datasheets.

	Longitudinal (d_{33}) mode	Shear (d_{15}) mode	Cylindrical tube
Diagram			
Displacement formula ^a	$\Delta x_3 = d_{33} V$	$\Delta x_5 = d_{15} V$	$\Delta x_{lat} = \frac{2\sqrt{2} d_{31} V L^2}{\pi (ID + t) t}$ [53]
Example displacement at ± 200 V	± 80 nm ^b	± 460 nm ^c	± 1870 nm ^d
Displacement at ± 15 V	± 6 nm	± 35 nm	± 140 nm

^aThe piezoelectric charge constant, d , gives the sensitivity of a piezo material when used in a particular mode, in metres per volt. It is described with a pair of subscript numbers: The first of these numbers gives the direction of the applied electric field, and the second number gives the direction in which the displacement is being utilised.

^bData obtained from Physik Instrumente stock materials datasheet for 'PIC255/PIC252' PZT.

^cData obtained from Noliac A/S stock actuator datasheets.

^dTube parameters taken from the existing FMOS-Echidna fibre positioner.

2.2.1 Exploiting new piezo materials

New piezoelectric materials are appearing on the market that offer superior sensitivities over traditional piezoceramics, meaning lower voltages are required for the same mechanical displacement.

Synthetic piezoelectric materials have traditionally been ceramic compounds that are mixed as a powder and formed with a sintering process. This creates an imperfect, polycrystalline structure in the material that affects its electromechanical efficiency.

Recent years have seen the introduction of ‘single crystal’ piezo materials [54] to the commercial market. These monocrystalline materials have a near flawless lattice structure, and hence they do not exhibit the inefficiencies seen in polycrystalline materials. Single crystal material can be cut and polarised after the growth stage to produce piezo actuators with very high piezoelectric charge coefficients.

Of particular interest is lead magnesium niobate-lead titanate, or PMN–PT. This new-generation single crystal material can have a piezoelectric charge coefficient more than five times that of traditional ceramics (PZT)². This could allow a tilting spine motor to be built with operating voltages in the range we are aiming for.

The relative immaturity of single crystal actuator manufacture means that accessing the technology is currently expensive. The majority of the cost comes from the amount of material used, so a new motor design using PMN–PT material would have to use small actuators if costs are to be kept low.

2.2.2 Exploiting new industrial processes

There have been significant advances in piezo manufacturing technology in recent years, specifically in the area of layered ceramics and electrode deposition for small devices. Suppliers now offer multilayer plate actuators, which comprise many thin piezoceramic layers with integral interstitial electrodes (Figure 2.2). This greatly increases the electric field strength throughout the ceramic, thus requiring lower voltages for the same displacements.

Multilayer plates are generally only available in simple block shapes, meaning that their use in a tilting spine motor would necessitate moving away from the current cylindrical tube design.

²Charge coefficient comparison based on PMN–PT d_{33} value specified by APC International Ltd.

PART II: A NEW TILTING SPINE MOTOR




Figure 2.2: Multilayer piezo plate actuators have many tens of thin ceramic layers separated by interstitial electrodes. The close spacing of the electrodes results in a much higher electric field strength, which means that much lower voltages are required.

2.2.3 Improving the mechanism efficiency

This approach concerns the efficient translation of force and motion from an actuator to the spine itself. If we look at the existing design, we see that it may not be an efficient mechanism at all: The lateral ‘snap’ of a piezo tube with a conical ball contact has component forces that drive *into* the ball rather than at a tangent to it. It’s conceivable that this would even momentarily launch the ball into the air.

An ever-present challenge associated with developing this technology is knowing for sure what mechanisms are present when a spine is in motion. Movements are microscopic and fast, requiring high-bandwidth nano-displacement sensors or high-speed video cameras at very high magnifications. This equipment is beyond the scope of many laboratories, including those at Oxford Physics at this time. Nevertheless, it’s still possible to arrive at educated conclusions based on classical mechanics.

Figure 2.3 shows the argument that the existing Echidna drive mechanism is inherently inefficient. We know that driving opposing piezo tube electrodes at opposite polarities produces a ‘banana bend’, causing the spine mounting cup to lean sideways during the slow ramp of the drive waveform’s ‘stick’ phase. If we now think about the moment that the actuator begins its sudden move back to a neutral position in the ‘slip’ phase, it seems reasonable, perhaps for the first time ever, to compare astronomical fibre positioning to golf. The proposition here is that the high-acceleration lateral ping of the wedged mounting face into the side of the ball is approximately equivalent to hitting a golf ball with a 9-iron.

PART II: A NEW TILTING SPINE MOTOR




Figure 2.3: The lateral motion of a current-generation tilting spine mount does not result in an efficient rotation of the ball. In this exaggerated diagram, the force exerted by the actuator (F) on the ball at the moment it ‘snaps’ back to a neutral position (left) is shown as being analogous to hitting a golf ball with a steep club (right). The resultant acceleration (a , not drawn accurately) could well launch the ball into the air.

We can attempt to quantify the efficiency of the Echidna motor design by comparing the lateral displacement of the piezo tube actuator with the resultant rotation of the ball. For an Echidna-spec actuator driven at³ 140 V, the theoretical displacement at the tip of the tube is $1.3\ \mu\text{m}$ (refer to Table 2.1 for equation). Data obtained from the AAO shows a mean fibre step size of $\sim 40\ \mu\text{m}$ at this voltage for 250 mm-long spines with 7.94 mm diameter balls. Figure 2.4 shows how this translates to a mechanical advantage of $\frac{250}{3.97} = 63$ between the surface of the ball and the fibre tip. Therefore, if the entire actuator travel was being transferred to the rotation of the ball, the fibre step size would actually be $83\ \mu\text{m}$. This reveals an ‘efficiency’ of less than 50 % with the Echidna drive.




Figure 2.4: The mechanical advantage (MA) between the surface of a spine’s ball and its tip is equal to the ratio of the ball radius to the spine length.

Even if the mechanical coupling between the Echidna ball and cup was perfect, it would not be enough to lower drive voltages by the amount we desire. A more efficient mechanism could, however, be combined with other modifications to yield enough of an improvement. Furthermore, any refinements to the stick–slip scheme that make spine movements less like chipping a golf ball would probably benefit open-loop repeatability.

³This is the drive voltage for which a resultant step size was readily available.

2.2.3.1 Actuator amplification and resonance

An extension of creating a more efficient mechanism to achieve lower drive voltages is to employ one of two common tricks used with piezo actuators to achieve higher displacements and/or lower voltages: i) mechanical amplification of the piezo's displacement; or ii) exploitation of actuator resonance.

Actuator amplification involves the incorporation of a mechanical advantage between the piezo element and the surface being moved. The drawbacks of this approach are added complexity, cost and bulk.

Resonant piezo drives use the resonant frequency of a piezo actuator as a natural means of increasing its efficiency. By driving the actuator close to this frequency (typically in the range 10–100 kHz, depending on design), much less energy is required to produce dynamic displacements. The nature of resonance means that this approach is unsuitable for mechanisms that rely on actuator motions with varying acceleration (e.g. stick–slip designs requiring sawtooth waveforms). In other words, actuators driven at resonance will exhibit approximately sinusoidal motion. Driving Echidna spines in this way is a possibility, but would require a fundamental shift in their principle of operation.

While actuator amplification and resonance were considered in the development of a new tilting spine motor, other approaches proved more successful and so neither were pursued in detail.

2.3 Section summary

This section has introduced the three avenues that were explored in producing an improved tilting spine motor mechanism: i) the use of new, highly-sensitive single crystal piezo material instead of classical ceramics; ii) the use of new multilayer actuator types to boost the performance of traditional materials; and iii) changing the way that the displacement and force from an actuator is transferred to a spine's ball so that the mechanism is more efficient than the current design, which has been shown to 'waste' more than 50% of its available actuator travel.

3 Design and construction of a laboratory test rig for tilting spines

With no existing infrastructure for testing tilting spine positioners at Oxford, an initial hurdle was to build a laboratory-based ‘test rig’ (Figure 3.1) that would allow meaningful evaluation of improvements to the technology.

This section describes the requirements, design and assembly of a complete system for testing the performance of a single piezo-based tilting spine motor. This is presented for the sake of completeness, but later sections are not contingent on knowing the details of the test system hardware or software.




Figure 3.1: A laboratory ‘test rig’ was designed and built, including software, drive electronics and a metrology system for accurate fibre position measurement. The measurement camera and spine mount are within the central pivoting cage.

3.1 Test rig requirements

This section describes the technical and operational requirements for the tilting spine test rig. These requirements are based on the existing Echidna control architecture and the typical performance metrics used to evaluate fibre positioning systems.

3.1.1 Control system requirements

The control system architecture was to allow control, via direct user input or a software script, of positioning for a single tilting spine.

Position feedback was to be provided, to allow iterable closed-loop control. This was to be via a measurement camera imaging the tip of the spine, with an overall measurement accuracy of $\sigma \leq 0.5 \mu\text{m}$ and a circular field of view $\geq 30 \text{ mm}$ diameter, to cover the largest anticipated spine patrol area.

The control electronics were to provide arbitrary, user-defined waveform generation over four channels, with selectable amplitudes in the range $\pm 200 \text{ V}$ (noise $\sigma \leq 50 \text{ mV}$) and selectable frequencies in the range 1–150 Hz. These values encompass the current operating conditions of Echidna actuators, but would also allow flexible conditions for lower-voltage designs.

3.1.2 Testing requirements

The test system would allow tests to be run that can evaluate tilting spine technology against the key metrics of independently-actuated fibre positioners. The associated requirements are as follows.

- **Closed-loop positioning performance:** A facility for closed-loop positioning to a specified target in the measurement camera’s field of view, and provision for recording the positioning error after each iteration.
- **Open-loop positioning performance:** Appropriate test routines to measure the repeatability/scatter of a spine’s step size and direction, and provision for recording the error of any single move together with the distance of that move.
- **Performance across patrol area:** A way to record open- and closed-loop performance with respect to the location of the spine in its patrol area.
- **Gravity vector response:** Provision for changing the angle of the spine mount with respect to gravity through a zenith angle range of 0–180°, to simulate telescope pointing with both prime focus and Cassegrain instruments.

PART II: A NEW TILTING SPINE MOTOR

- **Calibration stability:** Appropriate test routines to measure a spine’s changing calibration over repeated calibration cycles.
- **Long-term performance:** Provision for recording open- and closed-loop performance throughout longevity tests that simulate typical instrument lifetimes.
- **Safety:** Sufficient mechanical and electrical safety measures to guard against accidental human contact with hazardous voltages.

Performance testing outside of normal laboratory conditions (temperature, pressure, humidity) was considered beyond the scope of this work.

3.2 Test rig design

This section describes the approach taken in meeting the requirements of the tilting spine test rig. The overall aim was to allow flexible control options that are completely configurable by software. In the interest of simplicity and cost-effectiveness, all optical, electronic and mechanical systems used commercial off-the-shelf (COTS) products wherever possible.

3.2.1 Control architecture design

The general architecture of the control system (Figure 3.2) was based on that of the existing AAO tilting spine control architecture, but with added flexibility. It is centred around a commodity Linux workstation. This computer handles high-level control commands and communicates with the control electronics over a USB serial link.

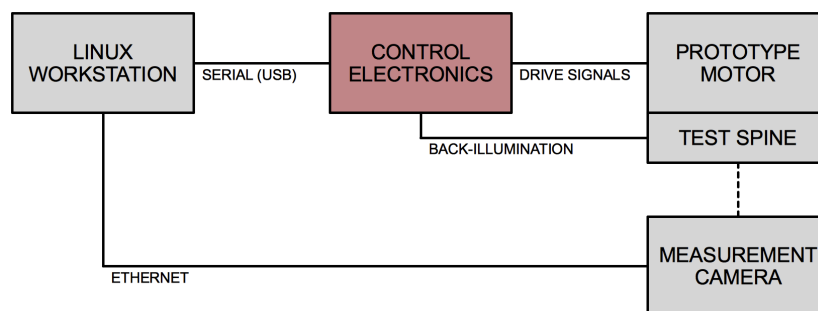


Figure 3.2: The test rig control architecture is a small-scale implementation of a full tilting spine control system for closed-loop positioning. It comprises a control computer, control electronics, and a visual feedback system via a digital camera.

PART II: A NEW TILTING SPINE MOTOR

Spine actuators are connected to the control electronics (described below) and their physical connections are specified in a simple software configuration file. A machine vision camera (also described below) provides positional feedback to the main control computer for full closed-loop operation.

3.2.2 Electronics design

The electronics system (Figure 3.3) is managed by a microcontroller running purpose-built firmware. Fortunately, the Arduino range of hobbyist microcontroller boards offers a model that was suitable for this system at a very low cost. The microcontroller firmware interprets a custom command set and generates output waveforms through a quad channel digital-to-analogue converter (DAC). The output waveforms then pass through a custom-built preamp stage¹ and on to a set of off-the-shelf high-voltage power amplifiers. The four amplified signals can be used to drive actuators directly, or a switched output can be used. The three switched outputs can be toggled between two waveforms, which mimics the design of the existing Echidna electronics.

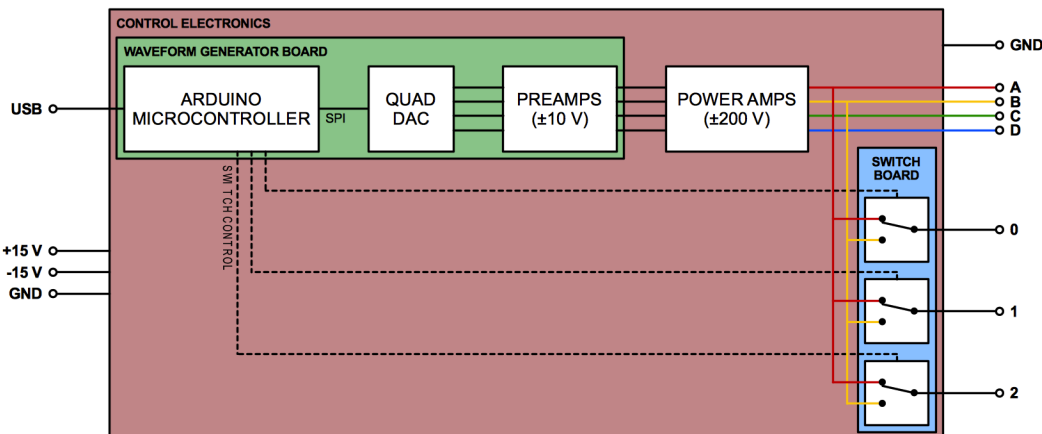


Figure 3.3: The test rig control electronics system provides flexible and reconfigurable waveform generation for driving tilting spine actuators. Prototype positioners can be driven directly through outputs A–D, or connected through a switched (classical Echidna) architecture at outputs 0–2. Arbitrary waveform generation is accomplished with a hobbyist grade Arduino microcontroller, at a sample rate of ~ 15 kHz across four channels.

In summary, this electronics design was created for its ability to emulate the high-voltage regime of existing tilting spine technology, at the same time as allowing lower voltage operation with modified signal paths for the new prototype devices that were to be developed.

¹Acknowledgement goes to Peter Hastings of the Oxford Physics Central Electronics Group for his assistance manufacturing this board.

3.2.3 Measurement system design

Positional feedback is provided by a machine vision camera system that images a back-illuminated optical fibre held within a prototype spine. The image of the fibre is processed by the control computer in order to extract its precise location.

The technical requirements for the camera design imposed significant restrictions, given the desire to be able to rotate the entire spine and camera assembly to simulate telescope pointing. This called for a large and high-resolution image detector and a suitable lens for imaging the fibre at close quarters.

The final measurement system layout is shown in Figure 3.4a. The chosen lens assembly is an off-the-shelf unit², providing a magnification of $\sim 0.24\times$ and giving a (calibrated) pixel scale of $24.9\ \mu\text{m}/\text{px}$ on a one-inch square format detector with $5.5\ \mu\text{m}$ pixels³.

The camera’s field of view is approximately circular, with a diameter of $\sim 46\ \text{mm}$. This size was chosen to accommodate the accessible area of a cluster of seven hexagonally-packed spines, should this be required in the future.

The focal distance of the camera is 110 mm from the front lens element. This, together with the substantial physical length of the camera ($\sim 220\ \text{mm}$ to the rear of the detector unit, plus cabling) meant that the optical path had to be folded through 180° (Figure 3.4b), so that the camera system and spine could be accommodated in a compact rotating ‘cage’. The optical design is also telecentric at the fibre image plane in order to reduce tilt-induced dimming of the fibre image.

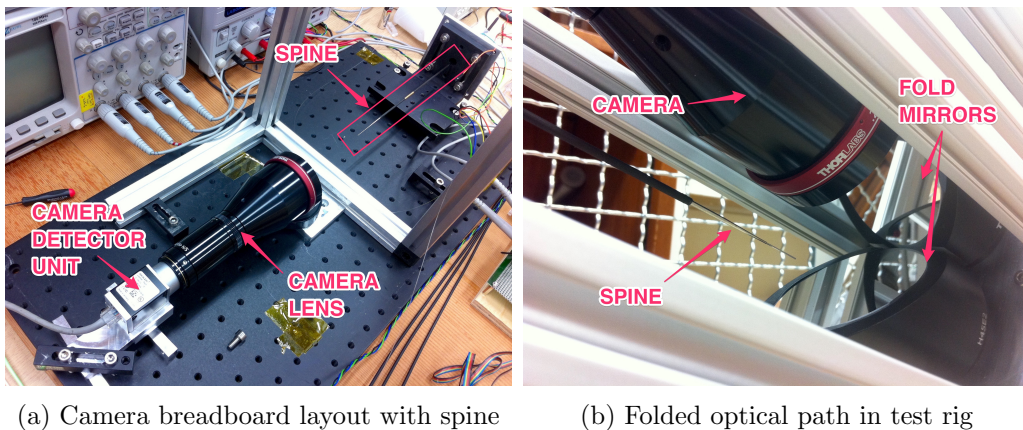


Figure 3.4: The length of the measurement camera system (a) necessitated a fold mirror (b) in order to keep the test rig’s spine and camera assembly compact.

²The lens used is a Thorlabs MVTC23024 bi-telecentric machine vision lens.

³The detector unit used is a Basler Ace acA2040-25gm, with a Gigabit Ethernet connection.

PART II: A NEW TILTING SPINE MOTOR

Figure 3.5 presents sample images from the measurement camera, showing the full camera field of view and a sample feedback image with a typical spine patrol area overlaid. Figure 3.6 shows a plot of centroid measurements for a stationary spine, revealing an RMS measurement system noise of $\sim 0.22 \mu\text{m}$.




Figure 3.5: The tilting spine test rig measurement camera has a $\sim 46 \text{ mm}$ diameter circular field of view (a) and a pixel scale of $24.9 \mu\text{m}/\text{px}$. These are long exposure sample images from the assembled system. A typical spine patrol radius (yellow) is overlaid in (b).




Figure 3.6: The measurement system noise and stability was found by imaging a stationary spine multiple times over a period of $\sim 10 \text{ min}$. Only a small amount of (thermally-induced) drift, manifest as elongated wings, was apparent. The magnitude of the scatter of the centroid measurements comes out at $\sim 0.22 \mu\text{m}$, which is within the goal of $\leq 0.5 \mu\text{m}$. This is for sets of nine stacked frames.

3.2.4 Mechanical design

A modular system of aluminium profiles provided a mounting structure ‘cage’ for the measurement camera optics and prototype spines under test. The same system of profiles was used to form a pivoting mount for this cage, allowing the simulation of different telescope pointing angles.

The entire test rig is surrounded by a larger enclosure with a mixture of removable plastic and mesh panels. This provides ventilation for the electronics, while ensuring hazardous voltages aren’t a risk to human safety.

Figure 3.7 shows the enclosed system, including the camera, electronics, and a prototype spine assembly.




Figure 3.7: The test rig enclosure provides ventilation for the electronics while keeping hazardous voltages contained.

3.2.5 Software design

The control software is responsible for the general operation of the control system (commanding the electronics, processing feedback images) as well as running test routines. These functions were implemented using a custom-built ‘suite’ of software modules written in Python [55].

Figure 3.8 shows the software architecture. It comprises a set of higher-level test scripts and lower-level driver-like files for both the electronics and the position feedback system. New test routines can be written as needed, each with their own data visualisation functions for producing appropriate plots.

The hardware is controlled by calling ‘friendly’ functions defined at the top layer of the driver modules. These functions are specific to the design of the spine motor, but include things such as pre-defined calibration routines and absolute movement commands to specified coordinates.

The layers of the driver files gradually decrease in their level of abstraction, the lowest level communicating with the main electronics microcontroller over a USB serial link. A camera driver module interfaces with the measurement camera over an Ethernet connection via the manufacturer’s software API. A fibre measurement module presents friendly measurement functions that return the current spine position.

This software suite provides a very flexible testing and control environment for the tilting spine test rig hardware. It allows long-term automated testing with only a few lines of Python code. Similarly, different hardware configurations can be accommodated with simple changes to configuration files.

3.3 Section summary

This section has described the laboratory test rig that had to be built before any spine motor development could be done. The details of its design are not important to remember, but as a testing facility it was an instrumental part of this research. All of the prototype testing results reported in the remainder of this thesis were obtained using this test rig system.

PART II: A NEW TILTING SPINE MOTOR

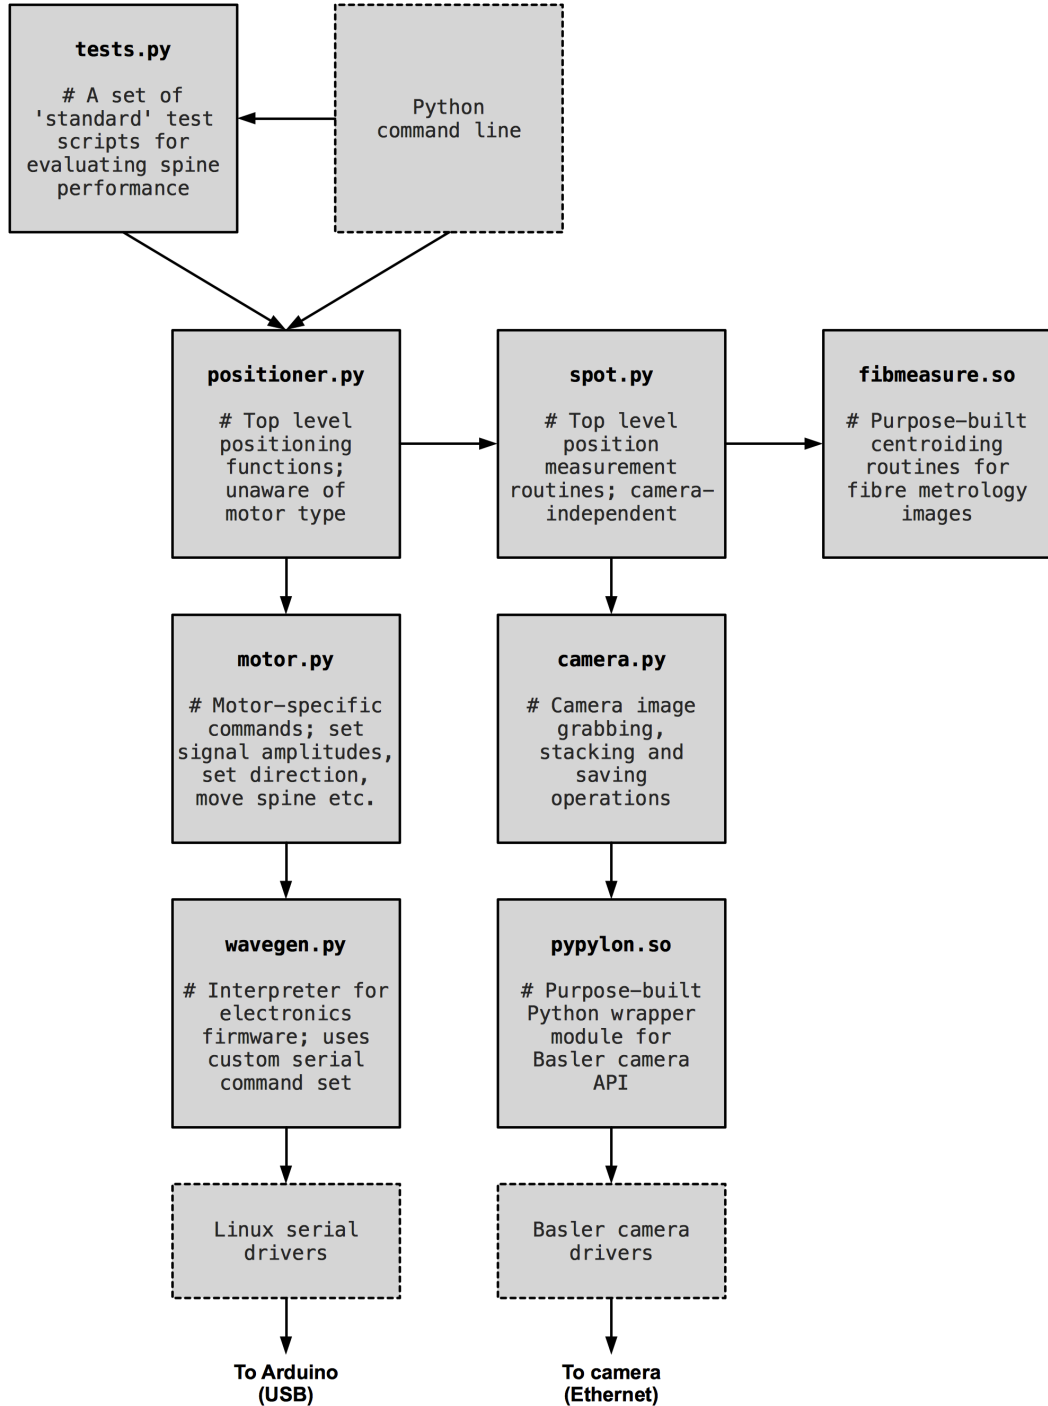


Figure 3.8: The test rig control software comprises a suite of Python class definitions of varying levels of abstraction.

PART II: A NEW TILTING SPINE MOTOR

4 New motor designs

This section describes three novel piezoelectric stick–slip motor designs that were created in the pursuit of a low-voltage, high-performance tilting spine motor. The first two designs—the ‘three-finger motor’ and the ‘shear plate motor’—were discarded for technical reasons, so will be only briefly summarised. The third design—the ‘multilayer plate motor’—was a success, and is explained in detail.

4.1 The three-finger motor

The three-finger motor was based on a stick–slip mechanism and aimed to use single crystal longitudinal mode piezo actuators to reduce drive voltage requirements.

A 3D model of the motor mechanism is shown in Figure 4.1a. It featured three vertical ‘fingers’ that tilted the spine’s ball. They were to be driven in the same way as the existing Echidna motor, producing a similar stick–slip motion but with the actuators moving ‘up-down’ rather than ‘side-to-side’.




Figure 4.1: The three-finger motor concept used three actuators to roll the spine ball (a) and hence tilt the spine. It was later decided that the magnetic force between the motor assembly and the ball would put undue outward stress on the actuators (b).

PART II: A NEW TILTING SPINE MOTOR

The three-finger motor design was abandoned at the concept phase over concerns for its mechanical robustness and cost.

The primary technical concern was that the attractive force between the magnet and the ball would cause an undue amount of stress on the actuators in the outward (radial) direction, as illustrated in Figure 4.1b. The brittle nature of piezo actuators meant that such stress was likely to result in mechanical failure.

In terms of cost, an estimated actuator pricing of USD 90 per spine was likely to be prohibitively expensive when compared to Echidna’s piezo tube pricing of around AUD 20 per spine.

4.2 The shear plate motor

The shear plate motor was based on a stick–slip mechanism and aimed to use single crystal shear mode piezo actuators.

A 3D model of the motor mechanism is shown in Figure 4.2a. Rather different to the current Echidna design is that the actuators are fixed directly to the top of the magnetic mount (cup). The idea here was to achieve an ideal tangential force with respect to the ball (Figure 4.2b), with the hope of improving performance.

The motor’s stick–slip cycle was similar to the existing Echidna motor: By slowly shifting one of the actuators radially outwards while shifting the other two inwards, or vice versa, the ball ‘sticks’ to the actuators and rolls about its centre. Quickly reversing this move will cause a ‘slip’ of the ball, leaving it displaced. Six combinations of this action are possible, hence spines can move in six directions with this motor.




Figure 4.2: The shear plate motor has three piezo actuators fixed directly to the magnetic cup assembly (a); these actuators are slightly proud of the cup, providing a moveable kinematic mount for the spine ball. This provides a near-perfect ‘rolling’ of the ball. The section diagram (b) shows two opposing actuators, which is inaccurate but better illustrates the tangential forces.

4.2.1 Prototyping

A working prototype of the shear plate motor, shown in Figure 4.3, was built in order to prove the design concept. This prototype used COTS PZT actuators in place of the intended single crystal devices. To all intents and purposes, the only difference was that these versions require a higher operating voltage to produce the same result. The plan was to replace these with identically-sized PMN–PT actuators if the design worked as expected.




Figure 4.3: The assembly procedure for the prototype shear plate motor involved bonding the bare actuators to the cup with an electrically-conductive epoxy (a), bonding insulating ceramic plates to the tops of the actuators, and then using a custom-made lapping tool to thin the plates to a height only slightly above the rim of the cup (b). Fine wires were connected to the actuators’ top electrodes; a ground connection was provided via the cup assembly itself, which bridged all the actuators’ bottom electrodes.

4.2.2 Poor performance and manufacturability issues

The first tests with a spine produced motion as expected. A test spine was kindly supplied by the AAO, with a length of 250 mm and a ball diameter of 7.94 mm.

The repeatability of spine motion was found to be poor, with a strong dependence on the orientation of the motor (zenith angle). Attempted closed-loop positioning failed to reach even 50 μm RMS accuracy in ten moves. All of these issues may have been fixed by further design iterations, possibly by improving the actuator–ball interface material, or by increasing the magnetic holding force.

It then materialised that the required single crystal actuators could not be manufactured, despite initial indications to the contrary from two suppliers. This only came to light after several manufacturing runs had been attempted and had failed.

Given that the design presented no real benefit without the prospect of low-voltage actuators, the shear plate motor was abandoned and a third design was pursued.

4.3 The multilayer plate motor

This section presents a successful low-voltage tilting spine motor design that will be the focus of the remainder of this thesis. Its design and principle of operation are described, along with initial prototyping and testing. Early results showed promising performance at waveform amplitudes of $<10\text{ V}$.

4.3.1 Design and principle of operation

Following two unsuccessful concepts that focussed on exploiting new *materials* in the piezo actuator industry, a new approach was taken that aimed to exploit new *processes* instead.

A 3D model of the ‘multilayer plate’ (MLP) motor concept is shown in Figure 4.4. This design is more akin to the original Echidna concept than the previous two designs, but makes use of the modern-day multilayer manufacturing techniques described previously in Section 2.2.2.




Figure 4.4: The multilayer plate motor has three multilayer piezo actuators fixed underneath the magnetic cup assembly; the actuators act as ‘pillars’ that are able to tilt the cup slightly. An embedded magnet holds the ball in the cup, as in the current Echidna design.

The magnetic cup assembly is similar to the existing Echidna design, with a smooth chamfer forming a direct line contact with the spine’s ball. The difference is that the traditional piezo tube is replaced by three individual actuators that can tilt the cup with a change of their height.

4.3.2 An improved stick–slip mechanism

Unlike Echidna, which uses a bending tube to shift the cup sideways under the ball, the MLP motor tilts the cup more or less in its place by lengthening or shortening three actuators arranged like pillars. This is shown in Figure 4.5. Already this goes a long way to achieving our goal of reducing the transfer of forces orthogonal to the spines in order to reduce mechanical crosstalk.




Figure 4.5: The multilayer plate motor uses three small actuators (blue blocks) to apply vertical forces (red arrows) to the spine’s magnetic cup assembly, tilting it slightly. The section view (left) shows two opposing actuators, which is inaccurate but illustrates the principle well; actuators are actually equidistant as shown on the plan view (right). Spine tube not shown; drawings not to scale; displacements exaggerated.

The steady lengthening and shortening of opposing piezo actuators under the cup produces a tilt of the cup, ball and spine together (‘stick’ phase), and then the sudden reversal of the motion tilts the cup in the opposite direction while the spine stays in its tilted position (‘slip’ phase).

This design assumed that glue compliance in the actuator and cup assembly would permit the (sub-micron) offsets in actuator height without causing significant stress or damage to the motor over time. Discussions with actuator vendors indicated that this was acceptable. Longevity tests have since agreed, and will be reported later.

The force applied to the ball during the tilting of the cup is not perfectly tangential, and its direction actually depends on the cone contact angle of the ball in the cup. The decision was made to keep the cone angle at 45° for the first prototypes, to match the existing Echidna design.

A beneficial feature of the three-actuator design is that a spine can be made to move along one of six vectors, rather than the four vectors available in the traditional design. The direction of movement depends on the combination of positive and negative waveforms applied to the actuators, as shown in Figure 4.6.

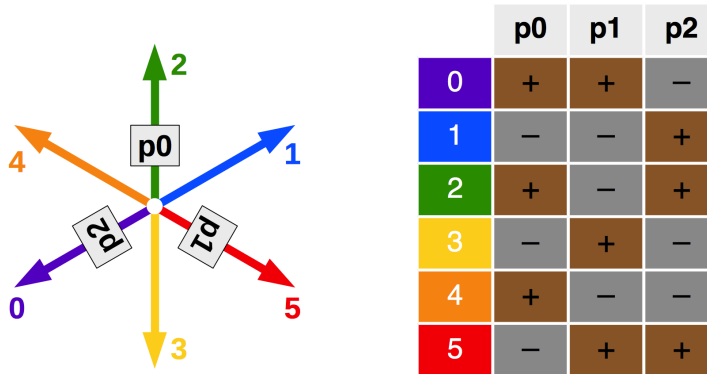


Figure 4.6: Three actuators (p0, p1, p2) and two waveform polarities (+ and -) yield spine movement in six possible directions (0–5), depending on the waveform connections. Note that a ‘negative’ waveform is simply the inverse of the standard sawtooth shape, i.e. a negative initial ramp instead of a positive one.

4.3.3 The use of multilayer actuators

The design is based around three MLP actuators operating in longitudinal mode. These actuators have 80 thin PZT layers with interstitial electrodes (refer to Figure 2.2, p. 134), are fully encapsulated in a passive ceramic shell, and are cubes measuring just 2.0 mm on each side. An added bonus is that they are off-the-shelf products.

The alternating electrode polarities of the actuator’s layers are bridged on two opposing sides of the actuator, making signal connections straightforward. Note that the reason for the actuators protruding beyond the cup is to ensure that the outside electrodes do not make electrical contact with the cup itself.

Having such thin layers of PZT means that the high electric field strengths required by these actuators can be achieved with much lower voltages. These actuators have a specified free stroke¹ of $\pm 3.0 \mu\text{m}$ at just $\pm 60 \text{ V}$.

As a first-order estimate, assuming a linear voltage–displacement relationship, the predicted actuator displacement is $\pm 0.75 \mu\text{m}$ at our ideal drive voltage range of $\pm 15 \text{ V}$. Now assuming that this displacement is entirely transferred to the ball, and taking a mechanical advantage of 63 for a current-generation 250 mm spine (referring to Section 2.2.3, p. 134), this predicted displacement becomes $2 \times 0.75 \times 63 = 95 \mu\text{m}$ at the spine tip. This is around double the typical coarse mode step size for the existing Echidna technology, so sufficient step sizes within the desired $\pm 15 \text{ V}$ range certainly looked plausible from the outset.

¹An actuator’s free stroke is its maximum travel in an unloaded state.

4.3.4 Assembly advantages

Recall that the existing Echidna technology provides signal connections to all the actuators across the field by means of a printed circuit board (PCB). This can be seen in the 3D model of the preliminary 4MOST/AESOP design for a spine ‘module’ shown in Figure 4.7.

The planned assembly routine for AESOP is to glue the cylindrical piezo actuators into a hole in the base and then apply solder to the actuator electrodes to connect them to the PCB. The PCB must also have a clearance hole (~ 6.5 mm diameter) for the actuators, which leaves little room for routing the signal tracks between them. This means that the PCB must have 17 layers, which calls for significant design and manufacturing effort.




Image courtesy of the Australian Astronomical Observatory

Figure 4.7: This 3D model of a single ‘module’ from the preliminary design of the AAO’s AESOP fibre positioner for 4MOST shows how the electrodes on the spine actuators are connected via a ‘module PCB’ running along the length of the ‘module base’ structure. This PCB has 17 layers.

PART II: A NEW TILTING SPINE MOTOR

The MLP motor design allows the motor assembly to be mounted on the positioner base structure in a similar way to the existing Echidna actuators, as demonstrated in Figure 4.8. It actually offers two advantages over the Echidna method: i) the diameter of the hole in the PCB and base structure reduces from ~ 6.5 mm to ~ 3.0 mm, meaning that the base and the PCB can be made thinner; and ii) the motor actuators can probably be bonded directly to the PCB rather than buried in the base like the piezo tubes have to be. The MLP motor's ground connection was later shown to be unnecessary (discussed later), allowing an even simpler arrangement (Figure 4.8c).

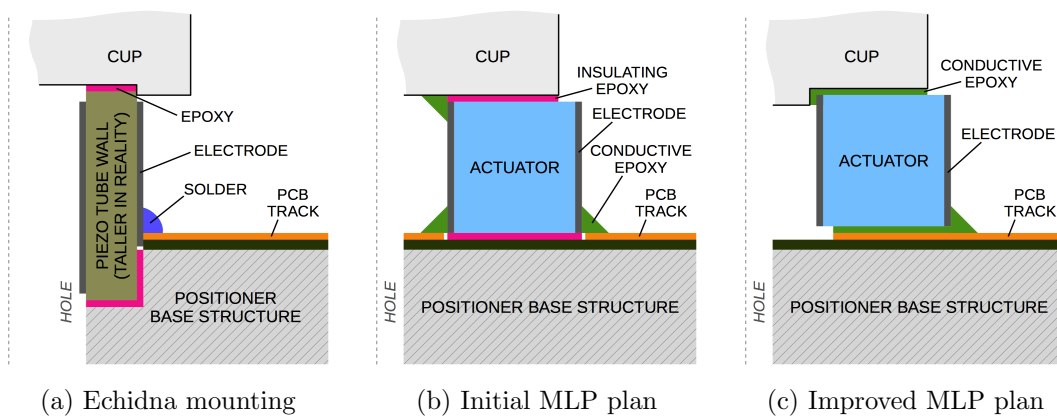


Figure 4.8: The existing Echidna tube actuators have to be bonded to the base structure of the positioner (a), which means that a large clearance hole is required; the Echidna actuators are also much taller (not shown). The MLP motor can be directly bonded to the PCB that runs across the base structure (b), requiring a smaller clearance hole. An even simpler assembly plan was later developed (c), which is explained later.

4.3.5 The lack of actuator preload

Manufacturers of layered piezo actuators often advise that tensile forces should be avoided wherever possible, recommending the use of a mechanical preload so that the bonds between the layers are not stressed. Such a preload might be a strong spring applying constant compression to the actuator, for example.

The MLP motor would become significantly more expensive, complex and bulky if a preload was required. This was discussed with the actuator manufacturer, who advised that the tensile forces in this motor would probably be low enough for a preload to be unnecessary. The same actuators are known to be used without preload in other stick-slip motors, routinely operating for tens of millions of cycles without degradation. The conclusion was to continue with prototyping and confirm that no damage will occur through long-term testing (reported later).

4.3.6 Cost

Budgetary estimates for the MLP motor actuators have placed them at around EUR 8 each in quantities of ~ 5000 , i.e. a total actuator cost of EUR 24 per motor. This equates to approximately AUD 35 at the time of writing. The PZT tube actuators used in the existing Echidna technology are a little cheaper, at around AUD 20. The current cup assembly, which is essentially unchanged in the new design, is estimated at AUD 30 in large quantities².

Costs of this scale are likely to be dominated by assembly labour, so full comparative costing is difficult at this stage. With this in mind, it seems fair to say that the MLP motor materials cost, while slightly higher, is not prohibitive.

4.3.7 Prototyping

Two first-generation MLP motor prototypes were built, the first of which was accidentally destroyed. The second prototype, identical in form, performed well. Further prototyping iterations followed, along with extensive testing, but these will be reported in the main testing and analysis section.

4.3.7.1 Prototype 1




Figure 4.9: The first MLP motor prototype was assembled and mounted on a simple break-out PCB for the purpose of testing.

Figure 4.9 shows the first prototype assembly of the MLP motor design. Assembly was reassuringly straightforward; the procedure was as follows.

²Costings are rough estimates from AAO engineers.

PART II: A NEW TILTING SPINE MOTOR

1. Cement the actuators to the base of the cup using a 3D-printed alignment jig and an electrically-insulating epoxy resin such as Araldite. Allow to cure.
2. Apply a small amount of electrically-conductive epoxy to the inside edges of the actuators, connecting the electrodes to the cup to form a common ground.
3. Cement the actuator and cup assembly onto the PCB using an insulating epoxy resin. Use an alignment pin to centre it over the hole. Allow to cure.
4. Apply conductive epoxy to the outside edges of the actuators, connecting the electrodes to the PCB tracks; also apply to one or more inside edges (ground).
5. Cement the ring magnet into the cup (will align and hold itself).
6. Allow the entire assembly to fully cure (room temperature or oven bake³ $\leq 100^\circ\text{C}$).
7. Verify connections with a resistance check ($\leq 1\ \Omega$) and actuator capacitance check (all within $\sim 10\%$).

In an unfortunate and perhaps ironic event, the actuators on the very first MLP motor prototype were destroyed by being driven at too high a voltage. In the spirit of good research, it was noted that this new motor design does not perform well with drive signal amplitudes of 190 V.

4.3.7.2 Prototype 2

The second MLP motor prototype was assembled using the same procedure as the ill-fated Prototype 1, and looked very much the same once complete. Amplifier gains were triple-checked before proceeding with testing.

Initial testing involved the completion of several ‘calibration routines’, which involved driving a test spine in all six motor directions for a certain number of steps and recording the mean step size. Waveform frequency and amplitude values were chosen empirically for this, after getting a hands-on feel for the motor’s behaviour.

Initial test results were excellent, with even calibration vectors (Figure 4.10) and good repeatability (Figure 4.11). We see acceptable coarse-mode step sizes ($\sim 30\ \mu\text{m}$) at waveform amplitudes of just $\pm 7\ \text{V}$, which is less than half of the goal operating range of $\pm 15\ \text{V}$. This is very exciting in terms of the primary aim of this research.

³Too high a temperature may de-pole the magnet or actuators; check their Curie points.

PART II: A NEW TILTING SPINE MOTOR




Figure 4.10: The prototype MLP motor showed very good calibration results, with relatively equal spine step sizes in each of its six component move directions. The direction of movement appeared to be stable at different voltages, and small step sizes were achievable. D0–D5 are component direction numbers.





Figure 4.11: A run of 80 successive calibration routines showed little scatter in the resultant calibration vectors. D0–D5 are component direction numbers.


A $\sim 10\ \mu\text{m}$ fibre step size was observed at a drive voltage of just 3.75 V. To first order, this corresponds⁴ to a positioning accuracy of $\frac{10}{2} = 5\ \mu\text{m}$, which is our target RMS error limit. Further analysis would reveal that step sizes were stable down to even smaller magnitudes, but this will be covered later.

⁴The general assumption is made that if a spine is more than a half but less than a whole step size away from its target, then taking a single step will overshoot the target, but by less than half of the step size.

PART II: A NEW TILTING SPINE MOTOR



(a) RMS positioning error



(b) Cumulative convergence

Figure 4.12: Closed-loop positioning results for a run of 500 random targets within a ~ 10 mm patrol radius at $Zd = 180^\circ$ (spine pointing down). The overall positioning accuracy leaves something to be desired, although nearly all spines get within $10 \mu\text{m}$ of their target in 10 moves. This suggests that a small number of outliers may have reduced the RMS error values.

PART II: A NEW TILTING SPINE MOTOR

Spine behaviour also appeared to be unaffected by the orientation of the motor, indicating that the magnetic force between the cup and the ball was sufficiently high. This was expected, as the cup assembly was based on a previous AAO design.

A snapshot of the prototype's closed-loop positioning performance is shown in Figure 4.12. The RMS positioning error over 500 targets approaches $10\ \mu\text{m}$ after ten moves, which leaves something to be desired with respect to our goal of attaining $\leq 5.0\ \mu\text{m}$ RMS error in six moves. Given the encouraging calibration results, however, it was decided that closed-loop performance may be improved by further development.

4.3.8 Further development and testing

At this stage it was clear that the MLP motor concept showed promise. Several new prototypes were to follow, along with improvements to the design and a great deal of testing and analysis. This will be reported in Section 6.

As we have already seen, a low-voltage motor presents many advantages in terms of positioner performance and engineering simplicity. Many of the highest-impact advantages concern changes to the hardware and software of the control system. These were investigated incrementally throughout the months that followed the first successful MLP motor prototype, but are best described in one place. This will be done in the next section.

4.4 Section summary

This section has presented three low-voltage tilting spine motor concepts, one of which was successful and two of which were not.

The three-finger motor was discarded before prototyping due to concerns over mechanical robustness and cost. The subsequent shear plate motor design was prototyped and showed some signs of promise, but the single crystal actuators it relied upon could not be manufactured after several attempts. Only a generalised description of these two designs has been given, despite this work having taken several months.

Finally, we have been introduced to the MLP motor design and the encouraging results seen after initial prototyping. The motor is able to produce spine step sizes comparable to the existing Echidna technology, but with waveform amplitudes of $< 10\ \text{V}$. This is well within the goal of $\pm 15\ \text{V}$ operation, which is excellent news.

This MLP motor was developed much further. These developments will be detailed in the remainder of this thesis.

PART II: A NEW TILTING SPINE MOTOR

5 Development of control hardware and software

This section describes how the new MLP motor is controlled. It focusses on a new approach to control, made possible by the motor’s modest electrical demands. A major change of the electronics architecture allows the spine to move in any direction, improving performance and simplifying collision avoidance when multiple spines move together. Control software developments are also presented, including new motor calibration algorithms and control strategies.

5.1 Exploring alternative control architectures

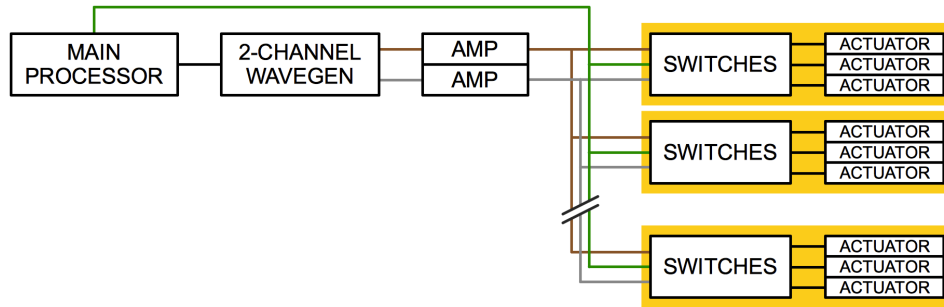
The successful design of a low-voltage tilting spine motor has opened up new possibilities in terms of electrical design. In Section 4.3.7 we saw promising initial test results using the existing Echidna-style control architecture built into the tilting spine test rig, but the low voltages required by the MLP motor mean that small-package COTS components can now be used to drive it instead. This gives us the option of moving away from the ‘shared waveforms’ architecture that was necessary with high-voltage drive signals and towards something far more flexible.

5.1.1 Limitations of the ‘shared waveforms’ architecture

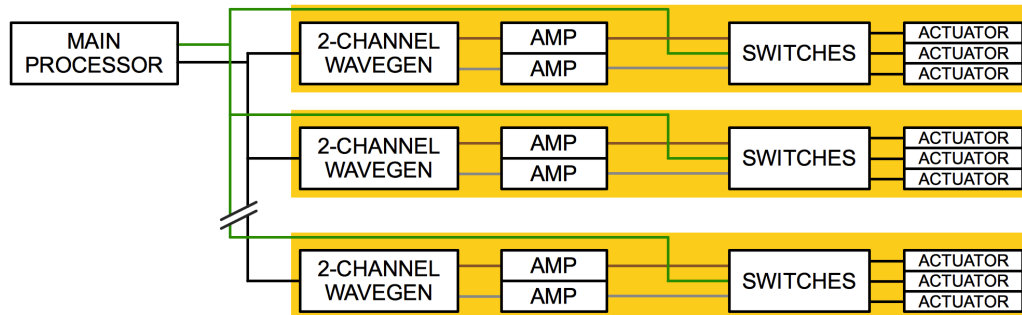
Recall that the size, inefficiency and cost of Echidna’s high-voltage power amplifiers meant that their outputs had to be shared among many spines. This ‘shared waveforms’ architecture is shown in Figure 5.1a.

For example, the AAO’s preliminary design for the AESOP fibre positioner has hundreds of spines sharing one pair of high voltage drive waveforms. The frequency and amplitude of these waveforms must be chosen to satisfy all spines, so the overall positioning performance is unavoidably limited by the ‘weakest spine in the pack’.

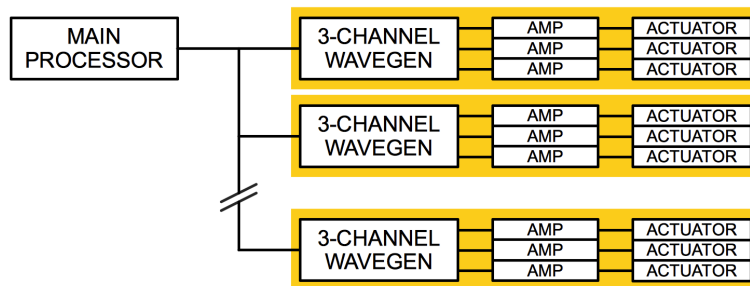
PART II: A NEW TILTING SPINE MOTOR



(a) 'Shared waveforms' architecture



(b) 'Waveforms-per-spine' architecture



(c) 'Waveform-per-actuator' architecture

Figure 5.1: The original tilting spine control architecture (a) shares drive waveforms among many (hundreds of) spine motors via a network of switches. With the new low-voltage motor, it becomes feasible to generate waveforms for every spine motor (b), or even every actuator within a motor (c). This has various performance advantages, but also improves modularity. The level of modularity is highlighted in yellow.

PART II: A NEW TILTING SPINE MOTOR

This is particularly significant when selecting the fine positioning amplitude, because this amplitude sets the minimum step size for each fibre.

In a 64-spine test system for 4MOST, the AAO has reported¹ that the finest step sizes range from 5.0 μm to 18.9 μm . In other words, some spines have less than a third of the positioning resolution of others, purely due to the fact that drive waveforms cannot be tuned to the specific needs of each spine actuator.

At the simplest level, the low-voltage MLP motor could retain the shared waveforms architecture, replacing the high-voltage amplifiers with smaller and cheaper low-voltage devices. Right off the bat, the MLP motor presents potential cost savings while keeping the architecture exactly the same. But the motor's real advantages are far more radical in nature, redefining the control principle entirely.

5.1.2 The ‘waveforms-per-spine’ architecture

Small and inexpensive amplifiers make it feasible to have dedicated drive waveforms for each spine in a low-voltage positioner. Using two dedicated amplifiers *per spine*, as illustrated in Figure 5.1b, presents a major advantage in that the waveform amplitude and frequency can be chosen for every spine individually. This has many benefits relating to the goals of this research:

1. **Increased positioning accuracy:** There is general increase in positioning accuracy by being able to ‘tune’ waveforms for each spine motor. This is by far the most significant performance gain.
2. **Reduced mechanical crosstalk:** The phase of drive signals across all spines can be staggered, reducing mechanical vibrations. This reduces the risk of inducing sympathetic movement of other spines.
3. **Reduced peak power demands:** Another result of staggering the phase of waveforms is a substantial reduction in the peak power demands of the system. Otherwise, many hundreds of spines are driven at high accelerations at the same moment in time, producing large current spikes.
4. **Improved modularity:** The architecture allows the option of a distributed control system where each spine is a self-contained unit with its own drive electronics. The merits of this approach will be discussed later.

¹Data obtained from internal design review documentation for the 4MOST fibre positioner.

5.1.3 The ‘waveform-per-actuator’ architecture

We can take the principle of dedicated drive signals even further, this time using a dedicated amplifier *per actuator* (i.e. three amplifiers per spine). This is illustrated in Figure 5.1c, where additional improvements to modularity are evident.

With the waveform-per-actuator architecture, the need for signal-routing switches—costing around AUD 12 per spine in the current system—is removed entirely, because each actuator’s dedicated waveform can be changed instead.

A hugely-beneficial effect of being able to choose arbitrary waveform parameters for each actuator in the motor is that spines are no longer limited to moving in six fixed directions. Rather, spines can move in *any* direction, which is a quantum leap in the capabilities of tilting spine positioners. These types of moves, which have been called ‘blended’ moves, will be explained in more detail shortly.

The waveform-per-actuator architecture offers the following benefits *in addition* to those listed for the waveforms-per-spine architecture above.

1. **Simplified software:** Because spines can move in any direction, it becomes much easier to avoid collisions between spines when reconfiguring a field, reducing software design complexity.
2. **Faster reconfigurations:** Spines can now move to a target position in a single move instead of two, which reduces positioning errors and therefore reduces the average number of closed-loop iterations to achieve a desired positioning accuracy. The improved collision avoidance mentioned above also saves time when positioning, because fewer spines will need to wait while their neighbours get out of the way.
3. **Further improved modularity:** Modularity now exists at an even more basic level. Each spine can have a self-contained control circuit that only requires rudimentary instructions from the main controller processor.

5.2 ‘Blended’ moves for motion in any direction

The waveform-per-actuator control architecture allows each of the three actuators on an MLP motor to be driven at an arbitrary voltage. What this means is that spines are no longer limited to moving in fixed directions, because this limitation was an effect of having to route fixed-amplitude waveforms in the old design.

5.2.1 The blended move principle

Spines are driven in any direction simply by adjusting the ratio of the drive signal amplitudes between each actuator. The principle is illustrated in Figure 5.2 and the mathematics of it are explained in Section 5.2.2, below.

These moves have been called *blended* moves², in contrast to the *component* moves used previously, where spines only ever travel in fixed directions. An alternative name might be *interpolated* moves, which is often used in robotics to describe cooperation of many axes to achieve a net result.

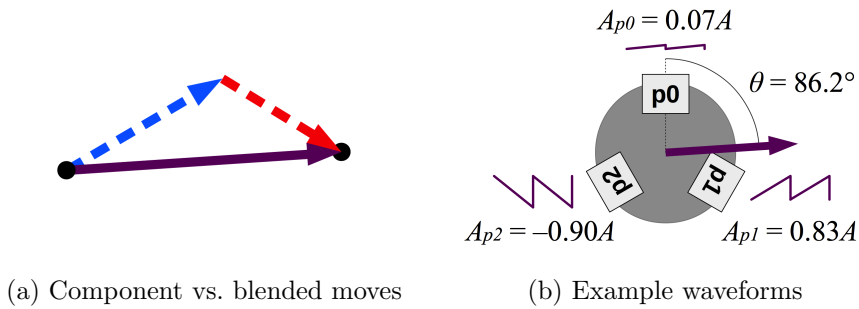


Figure 5.2: Driving the MLP motor actuators with a particular ratio of amplitudes can make a spine move in any desired direction, rather than being limited to only six component directions. These have been called ‘blended’ moves.

5.2.2 Calculating actuator amplitudes with a ‘drive angle’

In developing the control approach for blended moves so that the idea can be tested in the lab, the term ‘drive angle’ was adopted as a way to specify the relative amplitudes of the actuator waveforms and therefore the resultant direction of movement.

The drive angle is the clockwise angle from the actuator ‘p0’ in the motor assembly, as illustrated in Figure 5.2b above. Relative actuator amplitudes (or amplitude ‘gains’) in the range ± 1 are calculated from the drive angle using simple trigonometry:

$$\begin{aligned}
 A_{p0} &= A \cos(000.0 - \theta) \\
 A_{p1} &= A \cos(120.0 - \theta) \\
 A_{p2} &= A \cos(240.0 - \theta)
 \end{aligned}
 \tag{5.1}$$

where A is the overall requested amplitude for the move, or in other words, the maximum amplitude. This sets the overall step size.

²With acknowledgement to Will Saunders for this term, which seems to have stuck.

5.3 Updating the control electronics design

The decision was made to implement the waveform-per-actuator architecture in the lab in order to assess its feasibility for use on a large scale.

The first task was to find and test a commercially-available, small and inexpensive low-voltage amplifier that could drive an MLP motor. Having done this, a conceptual design of a replacement spine control circuit, including local waveform generation, was created and found to require less space and cost no more than the existing high-voltage Echidna control system.

5.3.1 Sourcing and testing low-voltage amplifiers

Several small, low-voltage COTS amplifiers were identified and one was chosen for lab testing. The chosen drive amplifier was the LM7322 op-amp from Texas Instruments, shown in Figure 5.3.

The LM7322 will operate on a $\pm 15\text{ V}$ supply and is specified to easily handle the capacitance ($\sim 150\text{ nF}$) and peak current demand ($\sim 40\text{ mA}$, shown in Figure 5.4) of an MLP actuator. It has two separate amplifiers in a single package and is available in a VSSOP surface-mount format measuring only $3.0 \times 5.0\text{ mm}$. It costs just USD 1.20³.




Figure 5.3: The Texas Instruments LM7322 dual-channel power op-amp is a small and low-cost device that can easily drive two MLP motor actuators. The LM7322 shown here is a large-package version soldered to a (blue) breakout PCB for the purpose of prototyping. The coin shown is an Australian 5 cent piece (diameter $\sim 20\text{ mm}$).

³Component cost based on current Texas Instruments budgetary estimates for quantities exceeding 1000 units, taken from the TI website.

PART II: A NEW TILTING SPINE MOTOR




Figure 5.4: This oscilloscope screenshot shows the measured current demands of the MLP motor actuators when driven with a nominal 8.0 V coarse-mode waveform (green channel 4). The displayed current (magenta channel 3) and power (red channel M) values are $30\times$ the actual values due to the gain of the current probe. Hence we see a peak current demand of ~ 40 mA per actuator during the high-acceleration ‘slip’ stage of the waveform cycle. These current spikes highlight the importance of staggering the phase of waveforms between spines in order to smooth positioner power demands.

The laboratory test rig control system was ‘upgraded’ by replacing its original high-voltage drive amplifiers with LM7322 devices. This not only allowed verification of the small amplifiers’ capabilities, but also enabled (with some software changes) testing of the blended moves principle.

The majority of tests reported in the coming testing and analysis section used the miniaturised test rig control electronics, with independent waveform generation on three channels via two LM7322 ICs. These devices have required no heat sinks or active cooling arrangements, and have been operating continually for more than 2000 h at the time of writing.

5.3.2 Conceptual design of a miniature drive circuit

Having found and tested a suitable small-format amplifier to drive the new motor, the remaining design hurdle concerned waveform generation. As we know, the shared-signal arrangement of the original Echidna technology meant that one pair of drive waveforms was distributed among many spine motors. Waveform generation could therefore be handled by a single microcontroller, the cost and footprint of which

PART II: A NEW TILTING SPINE MOTOR

was not a major issue. But the size and cost of an electronics design becomes very important if it is to be replicated for every single spine.

What we need is a simpler approach to generating sawtooth waves. One possible solution is to use small, crystal-less microcontrollers⁴ to generate the waveforms using high-frequency pulse-width modulation (PWM)⁵. This would not require a DAC stage, which keeps the circuit lean.

Figure 5.5 shows a conceptual design of an entire control circuit for a single spine, purely as an example of what an electronics engineer might come up with when building a control system with the waveform-per-actuator architecture.

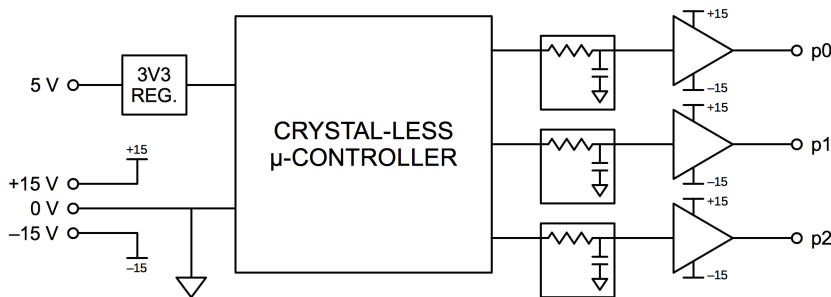


Figure 5.5: An MLP motor could have a completely independent control circuit based on PWM waveform generation: A small microcontroller would produce the waveforms, which would pass through a basic analogue filter and then through small amplifiers such as the LM7322s described earlier.

The estimated cost of this circuit is compared to that of the existing high-voltage shared waveforms circuit in Table 5.1. We see an approximate cost per spine that is slightly lower than the existing system. Of course, in reality costs are dominated by the labour involved. Recent discussions with AAO electrical engineers ended in an agreement that the real cost of the new architecture would be “about the same” as the old system. This is good news when we consider the advantages that it can offer.

Table 5.1 also includes an estimate of the circuit board area occupied by the new and old circuit designs. Here we can be quite sure that the new design is no worse in this regard. In other words, the size of the control cabinets around the positioner field would, if anything, shrink. Actually, a reduction of the overall size of the control hardware is very likely if we consider that we no longer require large high-voltage power supplies or amplifiers with tall heat sinks. A dedicated safety interlock system, included in the 4MOST/AESOP design, is also rendered obsolete.

⁴Crystal-less microcontrollers require no external oscillator for their clock, which saves space.

⁵PWM waveforms are generated by varying the duty cycle of a simple square wave at high frequencies and then smoothing this output with a basic analogue filter.

PART II: A NEW TILTING SPINE MOTOR

Table 5.1: We can be fairly sure that the proposed waveform-per-actuator architecture will be no more demanding in terms of material costs or circuit board area than the existing system. Costs are budgetary estimates for quantities of order 10^3 where applicable, and PCB area calculations do not account for signal routing.

Component	Quantity	Example device	Approx. total area	Approx. total cost
CONCEPT (LV)				
Microcontroller	1	STM32F042	25 mm ²	AUD 2
Amplifier ICs	2	LM7322	30 mm ²	AUD 4
Ancillary	—	—	50 mm ²	AUD 3
Totals			105 mm²	AUD 9
EXISTING (HV)				
Amplifier ICs	$\frac{2}{609}$	PA93	$\frac{7420}{609}$ mm ²	AUD $\frac{600}{609}$
HV switch ICs	3	AQW614	195 mm ²	AUD 12
Ancillary	—	—	Negligible	Negligible
Totals			207 mm²	AUD 13

5.4 Distributed control and ‘stand-alone’ spines

Previous tilting spine positioner designs have placed the control electronics system away from the telescope focal surface, around the edge of the fibre field. This has the advantage of allowing easy active cooling of the electronics cabinets, so that very little heat is dissipated into the open air.

This section discusses the feasibility of making each spine’s drive electronics part of the spine assembly itself, so that every spine is a stand-alone unit that only requires power and basic control instructions. This is the method adopted by the radial arm positioning technology being favoured by some of today’s fibre instruments, so the feasibility of doing this for a tilting spine positioner is worth exploring now that we have the luxury of low-voltage motors.

5.4.1 Conceptual design of a miniature control board

Using the component choices from the conceptual control circuit presented in the previous section, the approximate minimum size of an independent PCB containing the circuit can be investigated.

PART II: A NEW TILTING SPINE MOTOR




Figure 5.6: It seems plausible to fit an entire MLP motor control circuit onto a PCB measuring just 6×25 mm.

Figure 5.6 shows a simple 3D model of a 6×25 mm PCB with two copper layers and the main component packages on only one side. This particular concept requires four power connections (± 15 V, 5 V, 0 V) and two data connections for a communications bus (e.g. CAN bus). Boards can be daisy-chained via flexible PCBs.

A master control computer would load into the microcontroller the waveform amplitude and frequency parameters, the motor drive angle/direction, the number of steps required, and finally a ‘go’ command. The control board would return acknowledgements, ‘done’ messages when a move is complete, and optional status and telemetry information.


We can conclude from this short design exercise that an entire MLP motor control circuit, allowing arbitrary drive amplitudes and frequencies to be selected for each motor and permitting blended moves in any direction, would probably fit on a single 6×25 mm PCB if need be.

5.4.2 Conceptual design of a ‘stand-alone’ spine assembly

Taking the idea of fully modular spines to the finish line, a design concept was created that incorporated the control PCB design from the previous section into a complete stand-alone assembly. Models of the concept are shown in Figure 5.7.

This idea will be taken no further, but is presented as a demonstration of the new engineering possibilities that the MLP motor has created. The pros and cons of this design are reviewed below.


PART II: A NEW TILTING SPINE MOTOR



(a) The stand-alone spine concept



(b) Actuator close-up



(c) Daisy-chained spines

Figure 5.7: The MLP motor ushers in a new level of modularity for tilting spine systems. It could be possible to produce stand-alone spine assemblies that have a minimal dependence on ancillary electronics systems. This concept design has its control electronics incorporated via a vertical mounting vane, which is cemented to a 3D-printed base. The MLP motor could be assembled onto the base piece and connected to the local control board via a flex-PCB. There would even be the option to surround the entire assembly with a metal tube for enhanced magnetic flux and electric field containment. Many spines could populate the science field with daisy-chained flex-PCB connections (c). Although difficult to see, there is sufficient clearance around the spine's counterweight for it to tilt freely.

PART II: A NEW TILTING SPINE MOTOR

5.4.2.1 Advantages of stand-alone spines

- **Easy testing and verification:** Spines can be tested in the lab as single units, with minimal requirements for ancillary electronics.
- **Reduced in-house design work:** Having fewer custom circuit boards means that the entire control system is easier to design and build.
- **Small replicable units:** Manufacture of many identical spine control boards can be outsourced and would be cost-effective.
- **Reduced control system size and mass:** Most of the control electronics racks that previously housed the old architecture are no longer needed.
- **Efficient use of space behind the spine:** There may not be much room for electronics around the edge of the field.
- **Easier assembly and maintenance:** Population of the base structure with the spine assemblies may be easier than the current system. Replacement of faulty spine assemblies may also be simplified.
- **Best practice signal generation:** Generating waveforms that have high-frequency components close to the load is better for signal quality and electromagnetic compatibility than sending them down long copper tracks.

5.4.2.2 Disadvantages of stand-alone spines

- **Fiddly signal connections:** Daisy-chaining spines via thousands of flex-circuit connectors or small cables may end up being far too difficult.
- **Less room for advanced electronics features:** Advanced telemetry features, in-place microcontroller firmware upgrades etc. may not be feasible.
- **Possible increase in electromagnetic noise:** Contrary to the advantage above, it's possible that unshielded microcontrollers and communications buses may *increase* the overall level of electromagnetic interference.
- **Minimum pitch bottleneck:** Pursuing this type of design may end up limiting the minimum achievable spine pitch for very high fibre densities.
- **Heat dissipation:** Observatories have strict requirements on the amount of heat generated near the optical path of a telescope, and even careful electrical design may be unable to prevent this. A challenging active cooling arrangement would then be required behind the telescope focal surface.

5.5 Power dissipation estimates

Power dissipation is probably the greatest physical challenge when placing electronics in an open space near the focal surface of a telescope, as the dissipated heat can disrupt the incoming optical path. An observatory’s thermal restrictions can be a show-stopping limitation in this regard.

5.5.1 Power consumption while running

The power demand of an MLP motor during positioning was measured in the lab and is summarised in Table 5.2 alongside datasheet values for the other major components in the conceptual control circuit design. From this we obtain a total estimated power of ~ 0.2 W for a motor running in a nominal coarse-positioning mode.

Table 5.2: Estimated power consumption for a single MLP motor and its control board, using measured and datasheet values. Power supplies are assumed to be located away from the fibre field and therefore are not included.

Component	Power (RMS)
Microcontroller	66 mW
Amplifiers (quiescent)	99 mW
Actuators, reactive ^a	34 mW
Actuators, real ^b	6 mW
Total	205 mW

^aThe reactive power component of the actuators is the power ‘borrowed’ by their capacitance and then returned to the amplifiers; it is dependent on the piezo material’s dielectric loss factor.

^bThe real power component of the actuators is the power they actually consume, either through motion or heat loss.

5.5.2 Average power consumption

The average power dissipated by the entire system is dependent on the duty cycle of the positioner, which in turn depends on the intended operating mode of the instrument.

PART II: A NEW TILTING SPINE MOTOR

Taking the 4MOST instrument's maximum reconfiguration time requirement (1 min excluding metrology) and nominal observation block duration (20 min), we obtain a duty cycle of $\frac{1}{20} \Rightarrow 5\%$. Doubling this for a margin, and assuming we have 2500 spines, we get an average power dissipation behind the focal surface of

$$\begin{aligned} P_{average} &= 2500 \times 0.205 \times 0.1 \\ &= 51 \text{ W} \end{aligned} \tag{5.2}$$

Left unmanaged, this much heat is likely to be a problem. In contrast, the focal surface power dissipation for the motors only (control electronics elsewhere) is given by the real power component of the actuators in Table 5.2. Hence,

$$\begin{aligned} P_{average} &= 2500 \times 0.006 \times 0.1 \\ &= 1.5 \text{ W} \end{aligned} \tag{5.3}$$

Here we clearly see the benefits of placing control electronics in temperature-controlled racks away from the focal surface. In any case, the tilting spine technology is made more flexible by providing the option of either format.

5.5.3 Power consumption reduction

Various measures could be taken to reduce the estimated power consumption of the control electronics while the motor is running. The use of standard linear drive amplifiers, while a simple solution, is a cause of significant inefficiency for two reasons: i) they can have high current overheads (quiescent current); and ii) capacitive loads like piezo actuators have their major reactive power component wasted as heat.

Further engineering effort may be able to improve efficiencies by employing different classes of amplifier, or by using traditional linear amplifiers with added charge-recovery systems. This is considered beyond the scope of this work, which aimed only to prove the concept of distributed control. Indeed, if the existing approach of placing all spine control electronics in actively-cooled racks off of the fibre field is used then power reduction would not be necessary.

5.6 Software control principle

Since the MLP motor design has introduced a completely new candidate control architecture that allows spines to move in any direction, we shall review the basic calibration and control approach for both the traditional *component* moves approach as well as the new *blended* moves approach.

The fibre measurement procedure, which involves processing large fibre field images, will not be discussed, as this remains unchanged by the new motor design. Likewise, details of the considerations for controlling more than one spine (e.g. detailed collision avoidance strategies) are also considered beyond the scope of this thesis, despite some advantages actually being gained in this area.

5.6.1 Component moves control

Echidna spines have traditionally been able to travel in only four (roughly orthogonal) directions. Because of this, most spine positioning operations are actually compound moves in two component directions. The MLP motor design increases the number of available component vectors to six, which offers some increased efficiency in the route a spine can take to its target but is otherwise handled in the same way.

5.6.1.1 Calibration

Recall that the historical approach with tilting spine positioners is to have two types of moves: coarse moves with large step sizes, and fine moves with small step sizes. For now let us assume that we are happy in our choice of waveform frequency and amplitude for the coarse and fine modes. This topic will be covered later.

The positioner control software stores calibration data for every spine, which it subsequently uses to calculate the number of steps needed to traverse a certain distance. This is simply an array of values giving the single step size of the fibre and its angle of movement for every one of the available component move directions. These will be with respect to some meaningful coordinate system, such as the Cartesian distance from a fixed reference on the focal surface, but for all intents and purpose it is irrelevant at this level.

PART II: A NEW TILTING SPINE MOTOR

A standard component move calibration routine for a single spine is as follows. This is done separately for both the coarse and fine positioning modes.


1. Ensure the spine is close to its home position (centre of patrol area) and will not collide with any mechanical limits or obstructions.
2. Measure the current fibre location.
3. Move the spine in a certain component direction, at the relevant voltage and frequency, for a chosen number of steps (e.g. 10 steps); moving several steps is more indicative of a real move and also reduces the effect of measurement noise.
4. Measure the new fibre location.
5. Record the observed angle of motion with respect to the chosen coordinate system, and also the mean single step size.
6. Repeat for the remaining component directions.

Figure 5.8 shows a good example of the results of a coarse mode calibration routine for an MLP motor prototype. Note how imperfections in the motor are evident: vectors are not exactly aligned, and the step sizes are not all equal. It is these imperfections that make individual spine calibration so important.

Spine calibration also accounts for the physical orientation of the spine motor. This means that misalignment of the spine motor assemblies when they are fixed to the positioner structure do not matter.

Dir. no.	Move ang.	Step size
0	212.2	42.4
1	32.7	42.8
⋮	⋮	⋮
N	332.7	38.8

(a) Example data array



(b) Plot of vectors

Figure 5.8: Calibration data is stored for every spine in the positioner. Example stored values (a) are shown in blue. Non-uniformities in the behaviour of the spine can be seen in the plot of the data (b), e.g. in the slight asymmetry of directions 2 and 3, or the variation in step size magnitude across all directions. D0–D5 are component direction numbers.

5.6.1.2 Move calculations

Armed with calibration data for a spine, and comfortable with the assumption that spine behaviour is constant, it is a fairly trivial task for the control software to decide on the component moves required to position a spine on a target. Only two moves should ever be required to (theoretically) get within half of a step magnitude of the target. For an MLP motor, the software decides which direction numbers (0–5) are required and how many steps should be taken for each.




Figure 5.9: Multiple combinations of component moves are likely to be available, so the control software must decide which is the best solution.

There may be multiple move combinations that are valid (Figure 5.9), so the software may choose the best solution based on one of the two following rules. Remember that avoiding collisions with other spines is not being considered here.

- **Optimise for speed:** The solution involving the lowest total number of steps is considered to be the best, because it will complete sooner. This is the approach used in the spine test rig software for the purpose of this research.
- **Optimise for accuracy:** The solution that will theoretically arrive closer to the target is considered to be the best, because this means the spine should satisfy a given error tolerance sooner. One can imagine a range of tricks that exploit calibration non-uniformity, such as deliberately overshooting a target and back-tracking in the opposite direction. These will not be explored, as our aim here is produce better hardware rather than fix things with software.

The decision to switch from coarse positioning mode to fine positioning mode is based on the distance of the spine from the target. This was nominally set at $10\times$ our target RMS positioning accuracy of $5\ \mu\text{m}$ (i.e. a $50\ \mu\text{m}$ radius) for the purpose of testing, but naturally could be optimised for a particular instrument.

5.6.2 Blended moves control

The blended move principle has come about as a result of the MLP motor's compatibility with small and low-cost drive electronics, meaning that every actuator on every spine's motor can be driven at practically any voltage or frequency. This requires a totally new approach to spine calibration and move calculations.

5.6.2.1 Calibration

The calibration method that was created for MLP motors operating in the blended move mode was designed to be as simple and robust as possible. The calibration routine for a spine is similar to the component mode approach, except that the spine is driven in more than just six directions. These directions are based on the 'drive angle' concept explained previously (refer to Section 5.2.2, p. 165); the calibration routine moves the spine at increasing drive angles throughout the entire 0–360° range, in a specified number of equal divisions.

For example, the spine calibration data array may have ten entries recording the move direction and step size at drive angles of 0–360° in 36° increments. A more comprehensive calibration data array might have twenty entries in 18° increments, but at a cost of increased system memory. Twenty 18° increments were used for the purpose of testing the MLP motor prototypes.

A complete blended move calibration routine for a single spine is as follows. As with the component moves, this is done separately for both the coarse and fine positioning modes.

1. Ensure the spine is close to its home position (centre of patrol area) and will not collide with any mechanical limits or obstructions.
2. Measure the current fibre location.
3. Move the spine at a certain drive angle, at the relevant voltage and frequency, for a chosen number of steps (e.g. 10 steps).
4. Measure the new fibre location.
5. Record the observed angle of motion with respect to the chosen coordinate system, and also the mean single step size.
6. Repeat for the remaining drive angle increments in the range 0–360°.

PART II: A NEW TILTING SPINE MOTOR

Drive ang.	Move ang.	Step size
9.0	158.2	25.5
27.0	178.6	25.4
⋮	⋮	⋮
351.0	138.3	28.1

Figure 5.10: Spine calibration for blended moves is stored for a specified number of increments of the spine drive angle. Examples of the stored values are shown in blue.

Figure 5.10 shows example data from a coarse mode blended moves calibration for an MLP motor prototype, and Figure 5.11 shows the same data as a range of line-fitted plots revealing the direction and step magnitude relationships. The link between the *drive angle*, which is relative to the physical location of the ‘p0’ actuator, and the *move angle*, which is relative to the measurement system, shows good linearity. The calibration routine automatically accounts for the physical orientation of the motor assembly in this regard. The periodic step size variation in the example is actually the result of a modified magnetic cup design that will be described alongside test results later, not an artefact of the actuator positions as one might think.

5.6.2.2 Move calculations

Two approaches were considered for inferring the drive angle and number of steps required for a given move: i) function-fitting; and ii) linear interpolation.

The function-fitting approach used the same fitting functions as for the plots in Figure 5.11 to fill in the gaps between the measured calibration values. This worked well, but presented difficulties later when an adaptive calibration system was investigated, which will be explained shortly. Inference of the move parameters from the fit required non-linear optimisation, which is neither robust nor computationally-negligible.

Simple linear interpolation was favoured over fitting, as it offered a more robust result. Here the two calibration entries closest to the desired move direction are used and the behaviour of the spine is assumed to change linearly between them. The errors introduced by this simplification are of order 1% and in any case are far smaller than the intrinsic scatter in spine behaviour, so were deemed acceptable.

PART II: A NEW TILTING SPINE MOTOR

It follows that the procedure the control software uses for deciding on move parameters is quite simple: the angle between the current fibre position and the target is found (the ‘move angle’), and then converted into a spine drive angle using linear interpolation between the stored calibration data points; the predicted step size is also inferred using linear interpolation, and the distance to the target is divided by the predicted step size in order to find the number of whole steps that will get the fibre closest to the target.

As with component moves, the decision to switch to fine positioning mode is based on the distance of the spine from the target, and was set at $50\ \mu\text{m}$ for test purposes.

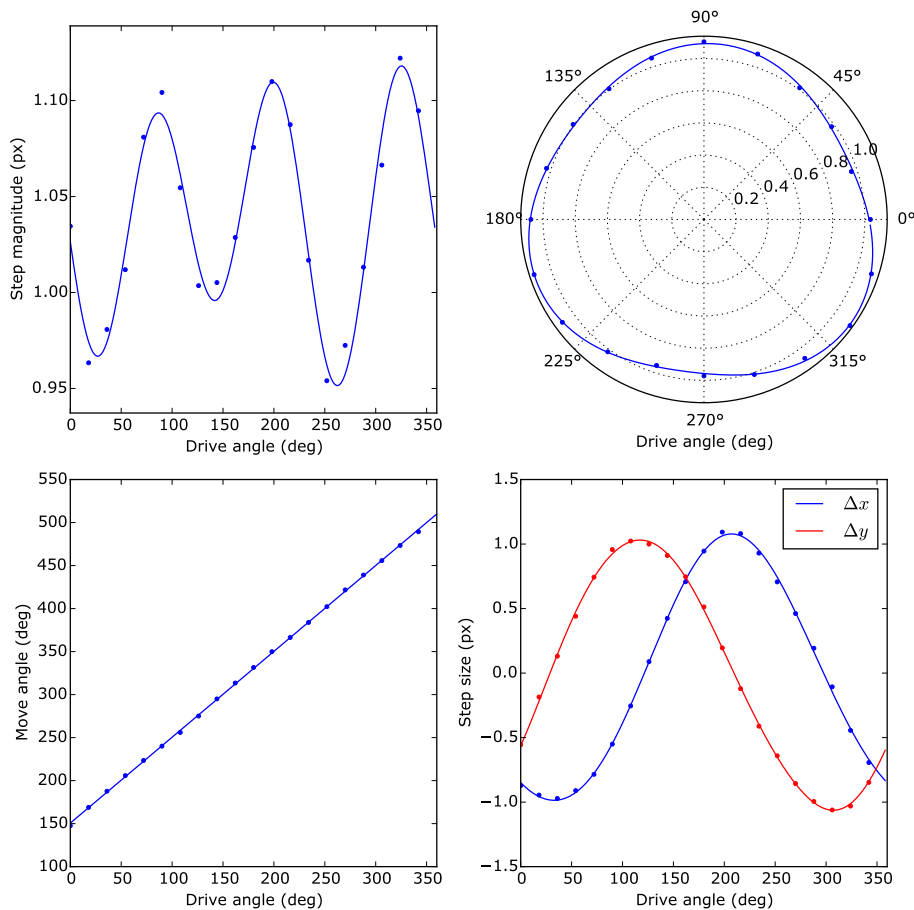


Figure 5.11: These plots show the relationship between the motor drive angle and resultant spine movement according to the metrology system. The bottom-left panel confirms a good linear relationship between the drive and move angles, whereas the top panels (which show the same data on different axis formats) reveal a variation in step size with the direction; this variation is actually an artefact of a new magnetic cup design on the motor, but in any case is represented well by a sinusoidal fitting function. The bottom-right panel displays the data in terms of the component x - and y -axis step sizes. Step size units are in measurement camera pixels ($\sim 25\ \mu\text{m}/\text{px}$).

5.7 Adaptive calibration of spine performance

The behaviour of traditional Echidna actuators is known to be non-constant [24]. Changes in the behaviour of an actuator can have various causes that manifest on different time scales, but in all cases will lead to positioning errors if an appropriate calibration strategy is not in place. The AAO has previously proposed an adaptive calibration strategy to account for “the step size, step direction, and memory effects all [evolving] with time.” Their suggested approach was to use a simple first in, first out (FIFO) buffer of motor performance information to continually update the calibration of each spine during normal positioning operation.

The approach taken in implementing adaptive calibration in the laboratory test rig system will be explained in the following pages. Before this, however, the causes of drifting motor behaviour will be summarised.

5.7.1 Causes of changing spine behaviour

Different factors affect spine behaviour by different amounts and on different time scales. The following is a summary of the main causes. During this work it was also discovered that spine performance can have a spatial as well as temporal dependence (see third point).

- **Temperature variations:** The sensitivity of piezo actuators, as well as the mechanical properties of glues etc., are temperature-dependent. Seasonal temperature swings will affect motor performance, but are generally slow enough to be fixed by nightly spine recalibration. This will not, however, fix temperature changes caused by the use of the instrument throughout the night.
- **Actuator hysteresis:** Piezo actuators exhibit varying levels of hysteresis. This can cause a ‘memory effect’, where the outcome of future moves is dependent on the direction of past moves. This is discussed later.
- **Spine position in field:** The step size and direction of a spine move can have a dependence on the spine’s current position within its patrol area. This was discovered during early tests and will be an important part of the chosen adaptive calibration approach. Perhaps this shouldn’t be surprising, as imperfections in the shape and finish of the spine’s ball will cause this, as will any imbalance in the spine assembly. Spines can also rotate slightly over many positioning cycles, adding a second order time-dependent term.

PART II: A NEW TILTING SPINE MOTOR

- **Positioner burn-in:** New positioners will probably exhibit some amount of ‘bedding in’ in the actuators and the mechanical interfaces of the assembly. This phenomenon could still be present to some degree after an instrument is commissioned, so must be accounted for.
- **Positioner ageing:** Similar to the burn-in of a brand new actuator, there will be (much slower) changes in its behaviour throughout its lifetime. Similarly, the reliance of the stick–slip mechanism on subtle tribological factors means that gradual wear of frictional surfaces will also have an effect.
- **Other effects:** Not to be overlooked are the unknown factors that can unsettle a system. Adaptive calibration can guard against many such changes.

In summary, the adaptive calibration system aimed to make the positioning performance of the new MLP motor robust to: i) known variations that cannot be fixed by infrequent (e.g. nightly) re-calibration; and ii) unpredictable short-term performance changes that are otherwise impossible to predict. In other words, what we want is for the calibration of the motors to look after itself, continually learning from its own ‘mistakes’ without any time penalty to observations.

5.7.2 Adaptive calibration of component moves

The work in this section will be published in the paper “Learning from history: Adaptive calibration of ‘tilting spine’ fiber positioners”, as part of the proceedings of the Astronomical Data Analysis Software and Systems XXV conference in Sydney, Australia, 2015 [56].

An adaptive calibration scheme was initially created for the traditional component moves mode. This was relatively straightforward, as spine calibration data only ever matters for a fixed number of movement directions.

The approach taken was based on the FIFO stack method proposed previously by the AAO [24], but with an important extension in that it also accounted for position-dependent performance changes.

All adaptive calibration algorithms run continually during the normal operation of the positioner, so there is no time penalty for the up-to-date performance data that the system provides.

5.7.2.1 The calibration grid

The concept of a calibration ‘grid’ was introduced. This divides the patrol area of every spine into a regular grid of cells, each with its own local calibration data. An example is shown in Figure 5.12.

The number of calibration cells for the purpose of testing was chosen empirically. Too few cells would fail to characterise spatial variation in motor behaviour, whereas too many would lead to data becoming stale before it is used. More cells also require more memory, scaling linearly. For the 69-cell example in Figure 5.12, an instrument such as 4MOST (~ 2500 fibres) would require a modest ~ 50 MB of memory for double-precision numbers and a FIFO depth of 3.

The calibration grid also accounts for the fact that the radial component of step sizes must reduce at the edge of the patrol area simply due to spherical geometry⁶. Although this could be easily modelled in the software, the effect has previously been treated as negligible since it amounts to a variation of only a few percent for the current spine design. Nevertheless, it is automatically accounted for with this adaptive calibration technique.

5.7.2.2 Updating calibration data

With the calibration grid in place, the adaptive calibration algorithm simply keeps a record of the inferred mean step sizes from the most recent moves in the local cell, adding them to a separate FIFO stack for each component direction. The stored data for each cell is illustrated in Figure 5.13. Testing has shown that a stack depth of three past moves works well, adapting to sudden changes quickly without being too sensitive to outlying data.

5.7.2.3 Querying calibration data

Before any move, the cell under the mid-point of the move is queried and the median step size and direction values are calculated from the FIFO stack for the relevant move direction. These are then used to decide on the best parameters for the move.

Taking the median value has the benefit of excluding any erroneous outliers in the data that may have resulted from a ‘bad’ move in the past.

⁶Imagine a spine tilted so much that it is orthogonal to the axis of the fibre measurement camera; now its radial step size would appear to be zero.

PART II: A NEW TILTING SPINE MOTOR




Figure 5.12: The patrol area of a spine is a circle (green line) of a certain radius (e.g. ~ 10 mm) from the home position of the fibre (red dot). This area is split into a number of ‘calibration cells’, each of which has its own historical record of motor calibration vectors (right) that were inferred from recent nearby moves.

Dir.	Move	Step
Dir.	Move	Step
Dir. no.	Move ang.	Step size
0	212.2	42.4
1	32.7	42.8
⋮	⋮	⋮
N	332.7	38.8

↕
FIFO

Figure 5.13: The adaptive calibration data arrays for component moves are the same as for non-adaptive calibration, but with a third dimension representing the FIFO stack. Note that a separate stack is actually maintained for each move direction.

5.7.2.4 Early test results

Laboratory testing of early prototype motors showed a significant reduction of the overall positioning error with adaptive calibration enabled. The system's learning is evident in Figure 5.14, which is taken from tests of the unsuccessful shear plate motor design. Figure 5.15 reveals a definite pattern in the spatial variation of spine behaviour, highlighting the need for the calibration grid.

The effect of the adaptive system is less drastic (both spatially and with time) for the latest MLP motor design, because its performance has proved to be generally more stable. However, the algorithms remain very effective at accounting for longer-term drifts in behaviour.

5.7.2.5 Incompatibility with compound moves

One problem with adaptive calibration for component moves is that the standard approach to positioning a spine involves a combination of *two* moves in different component directions. Without a measurement of the spine position after the first of these, it is impossible to decouple the combined positioning errors and therefore impossible to update the calibration data accurately.

It follows that the use of adaptive calibration in component move mode comes at a time cost after all, because spines can only ever complete single moves before another position feedback image must be captured by the measurement camera in order to assess the error. The general effect is that the number of iterations required to achieve a given accuracy is increased, although of course this is traded-off with a general increase in accuracy because of the calibration data being better.

For the sake of the tilting spine test rig, adaptive calibration for component moves was only ever enabled for the fine positioning mode, because this is where we care about accuracy the most. As will become clear later, positioning performance was proved to be excellent despite this limitation.

PART II: A NEW TILTING SPINE MOTOR




Figure 5.14: Running mean of fine move error for 400 full positioning cycles with a prototype shear plate motor and adaptive calibration enabled. This is the radial distance between where the control software expect the fibre to end up and where is actually ended up, so is a reflection of the overall quality of the calibration accuracy. The error falls quickly at first, as the cells are updated for the first time, then reduces steadily until hitting the motor’s repeatability limit.

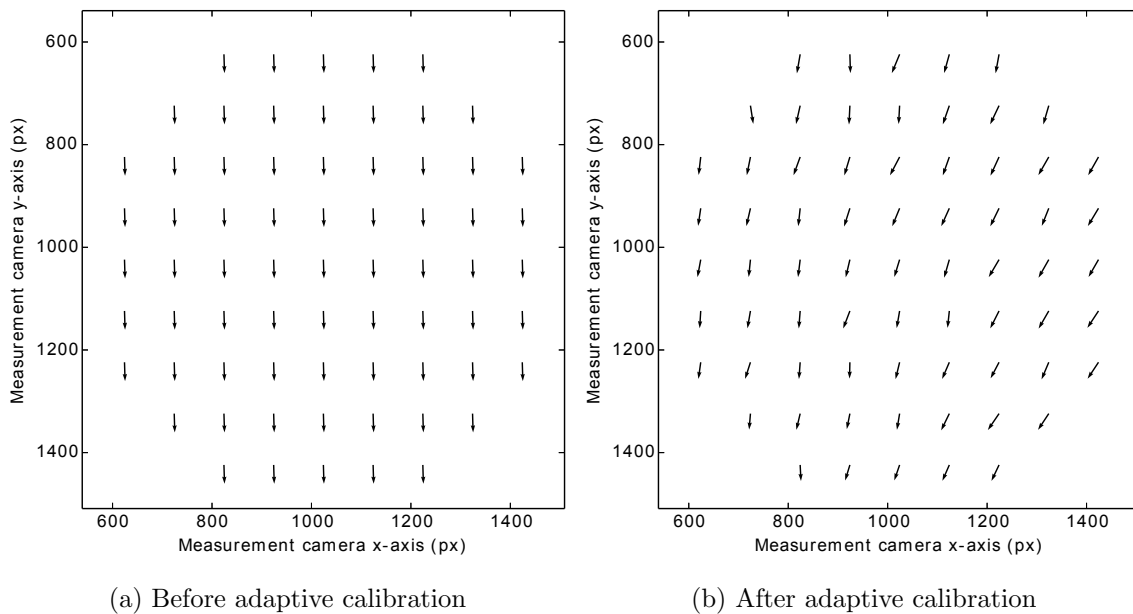


Figure 5.15: The data in the calibration grids is noticeably different after the adaptive calibration algorithm has been give some time to update. These plots show the calibration vectors for a single move direction with a prototype shear plate motor, revealing a definite spatial bias in the behaviour of the spine.

5.7.3 Adaptive calibration of blended moves

Adaptive calibration for blended moves requires a different tack to that for component moves, because moves can happen in any direction. This raises the question of how best to use the information from a past move in one direction to predict the motion of the same spine in a slightly different direction.

The solution adopted was to follow the same procedure used for component moves, but to bin the movement angles based on the number of divisions specified for the initial calibration routine in the 0–360° range (refer to Section 5.6.2.1, p. 178). This is illustrated in Figure 5.16 for clarity. Again, the patrol area is divided into a grid of cells, each with its own FIFO stack of calibration arrays.

θ	Drive	Move	Step
θ	Drive	Move	Step
θ range	Drive ang. (θ)	Move ang.	Step size
$0 \leq \theta < 18$	8.3	157.7	25.9
$18 \leq \theta < 36$	31.8	183.9	25.8
\vdots	\vdots	\vdots	\vdots
$342 \leq \theta < 360$	359.0	146.5	27.4




Figure 5.16: The adaptive calibration data arrays for blended moves are similar to those for non-adaptive calibration, but with a third dimension representing the FIFO stack. Drive angles are also binned, to accommodate arbitrary move directions while still ensuring the entire 0–360° range is covered.

Before a new move, the relevant calibration cell is queried and its calibration stack is returned. As with the non-adaptive case, the motor drive angle and expected step size are inferred using linear interpolation between the two closest data points for each of the stacked arrays. The median value is then used for the move calculations.

The memory required to store the calibration arrays depends on the number of angular divisions used, with 20 divisions seeming to work well in lab tests. In this case the required memory is increased by a factor of 5 compared to the component moves system, corresponding to ~250 MB for a 2500-spine instrument. This remains a reasonable demand for today’s computer systems.

Adaptive calibration will be used as standard in all of the blended move testing reported in the next major section. The system has so far performed well for more than 1000 h of testing.

5.8 Choosing optimal waveform parameters

The best way to decide on the drive waveform frequency and amplitude for the coarse and fine positioning modes could be the subject of a paper in itself. These parameters are important, as they ultimately set the performance of an entire positioning system.

The coarse mode frequency and amplitude will set the maximum *speed* of the spine, whereas the fine mode amplitude (less so the frequency) will set the positioning *accuracy* of the spine.

With a traditional ‘shared waveforms’ architecture, the fine mode amplitude would usually be set at the lowest voltage that guarantees the motion of all spines. The new architectures on offer with the MLP motor would allow different frequency and amplitude selection strategies to be used. Three approaches are outlined below.

1. **Fixed amplitudes:** This is the traditional method of control for Echidna spines and the approach used in the MLP motor tests. A waveform amplitude is chosen for each of the coarse and fine modes that will produce the desired step sizes. The voltage is always the same, even though the step size may change slightly over time. Spine calibration accounts for any changes in step size.
2. **Fixed step size:** It would be possible to take the inverse approach to fixing the waveforms amplitudes, instead specifying the desired step sizes. The calibration routine would be modified to adjust the coarse and fine mode amplitudes so that the specified step sizes were maintained. Although not implemented, this would be a useful extension to the blended move adaptive calibration algorithm, as it is far more meaningful to be able to maintain a given positioning accuracy than some arbitrary voltage value.
3. **A more advanced scheme:** The entire coarse/fine concept could be discarded in place of a system that makes use of the step size vs. voltage relationship across the entire operational range. Adaptive calibration could be centred around keeping an up-to-date model of this relationship, meaning that any step size in the available range could be chosen. The gains of this method may not outweigh the added complexity it would bring, but may be worth exploring.

Since the goal of this work is to achieve a 5 μm positioning accuracy, it was felt that the existing ‘fixed amplitudes’ scheme would be sufficient. The selection of optimal waveform parameters will be discussed alongside real results in the next section.

5.9 Section summary

This section has been a minor diversion to explore the hardware and software development that resulted from the advent of the MLP motor design. This was done in parallel with much of the motor performance testing, which is reported next.

Small low-voltage drive amplifiers have been identified and tested in the tilting spine test rig. The use of these amplifiers as part of a miniature, stand-alone control circuit has been investigated, and a revolutionary new ‘waveform-per-actuator’ architecture has been proved feasible in terms of hardware size and cost. This architecture permits any spine to move in any direction at any time, and allows every motor to be tuned according to its specific needs.

The approach to spine calibration and software control for traditional component moves has been described and new approaches have been developed for the blended moves mode. An implementation of adaptive calibration has been presented for both component and blended moves, which automatically accounts for changes in spine behaviour with position as well as with time.

PART II: A NEW TILTING SPINE MOTOR

6 The multilayer plate motor: Refinement, testing and analysis

This section presents the actual results from detailed performance testing of the MLP motor design, including a number of refinements and modifications that were made along the way. Much of this work occurred in parallel with the hardware and software developments covered in the previous section.

The testing and development reported here spanned a period of around twelve months. The final three months of this time were spent at the Australian Astronomical Observatory (AAO) laboratories in Sydney, Australia, where the AAO staff kindly provided the technical resources (test equipment etc.) needed to complete the planned experiments.


Refinements and results are reported here in rough chronological order, finally arriving at proof that the MLP motor is able to satisfy all performance goals set out at the start of this research.

6.1 Prototype motors and spines


The tests reported in the following sections involved various prototyping iterations of the MLP motor design. In some cases, the design of the spine used to test the motor was modified.

Photos of all the prototype assemblies from throughout this research are shown in Figure 6.1. All prototypes are numbered and will be referred to accordingly. Table 6.1 gives an overview of the prototype motors' characteristics.

PART II: A NEW TILTING SPINE MOTOR



(a) Prototype MLP motors



(b) Prototype spines

Figure 6.1: A total of eight MLP motor prototypes were built (a), some identical and some with incremental changes. The spine assemblies used in testing the motor were also changed at times (b)^a.

^aMuch appreciation goes to the AAO's Rebecca Brown, Jurek Brzeski and Nick Staszak for their assistance in producing this hardware.

PART II: A NEW TILTING SPINE MOTOR

Table 6.1: Summary of all MLP motor prototype iterations shown in Figure 6.1a. Many of their tests are described in detail in the following pages. Note the identical positioning accuracy achievement for all those used in closed-loop tests.

No.	Description	Approx. run time	Achieved accuracy	Comments/notes
1	Initial prototype	—	—	Accidentally destroyed by 190 V signal (software bug)
2	As Prototype 1	300 h	2.8 μm	Good performance; mechanical wear issues
3	Increased max. patrol radius	50 h	2.8 μm	Same good performance; same mechanical wear issues
4	As Prototype 3	—	—	Superseded by Prototype 5; assembly not completed
5	Modified cup design	2500 h	2.8 μm	Wear issues solved; sustained good performance
6	Miniature version (for 5 mm ball)	—	—	Manufacturing issues (discussed later)
7	As Prototype 5	15 h	—	Calibration tests only; consistent with Prototype 5
8	Modified actuator geometry	1 h	—	Failed attempt at reducing signal count to two

6.2 Wear resistance

Many of the early tests of Prototype 2 and Prototype 3 revealed a degradation in closed-loop positioning performance over time when runs of several hundred positioning cycles were carried out. Figure 6.2 shows one such example, with an abrupt change in behaviour after 400 targets. This was seen on several occasions, and was fixed every time by cleaning the contact surface of the cup.

The fact that wiping the cup remedied the problem indicated that some sort of contamination was to blame, most likely caused by debris from mechanical wear. This wear would have been from the cup itself, since this was made from an unhardened stainless steel and the ball was a hardened variety. Wear was to be expected in a mechanism such as this, but it was disappointing to see it cause problems after only a few hundred runs.

It turned out that wear-induced performance degradation had been observed by the AAO during their development of the 4MOST/AESOP cup assembly. This is not seen in the original FMOS–Echidna positioner, because here the spine balls rest on three small, hard hemispheres instead of on the rim of an enclosed cup. The

PART II: A NEW TILTING SPINE MOTOR




Figure 6.2: The mean final positioning error was often seen to suddenly worsen, usually after 400–500 positioning cycles.

enclosed cup design was introduced after the commissioning of FMOS–Echidna to better contain the magnetic flux in the spine mounting, but had obviously come at an unforeseen cost.

It became apparent that the wearing of the cup also introduced a period of ‘running in’ before spine behaviour stabilised, presumably while the ball wore away any micro-protrusions around the cup’s face. This phenomenon became the primary culprit for a gradual performance *improvement* seen in early MLP motor testing, which is reported in the next section.

Fortunately, the AAO had developed a solution to the wear issue: A hybrid cup design had been created that re-introduced the hard hemispheres while still enclosing the magnetic circuit sufficiently. Figure 6.3a shows this. Three tungsten balls were now embedded into the rim of the cup, raising the spine’s ball slightly off the face. In addition to this modification, the spine balls had been coated with titanium nitride to further increase wear resistance. This had become the baseline design for the 4MOST/AESOP positioner.

It was a simple matter to build a new MLP motor prototype using the latest AAO cup design, several assemblies of which were generously provided by AAO engineers. The MLP motor assembly procedure was essentially unchanged by this. Figure 6.3 shows close-up views of ‘Prototype 5’, which was the first to use this new format.

A side-effect of the new cup design was that the spine step size had a slight directional dependence relative to the positions of the small tungsten balls. This is generally accounted for when calibrating the motor, so is of no real consequence but

PART II: A NEW TILTING SPINE MOTOR




Figure 6.3: An improved magnetic cup assembly (provided by the AAO) was incorporated into the MLP motor prototypes with minimal hassle. In this design, the spine ball sits on three small tungsten balls (circled) that are embedded into the inside face of the cup. This provides the necessary magnetic circuit while allowing the use of a hardened precision contact surface at three equally-spaced points.

is worth noting. This phenomenon was the cause of the shape previously seen in the blended move calibration plots (refer to Figure 5.11, p. 180).

Continued testing showed no evidence of the issue observed with the previous prototypes, and there continues to be no such effect.

6.3 Closed-loop positioning performance

The speed and accuracy of closed-loop positioning was dramatically improved from the first MLP prototype tests due to a combination of factors relating to the hardware and the software. One major factor was the discovery that the initial prototypes performed much better after a ‘running in’ period.

Results for both component and blended move modes will be presented, although most testing and analysis took place before blended moves were made possible.

6.3.1 Component moves

Even before the addition of the improved cup assembly described in the previous section, some excellent results were seen with the original motor (Prototype 2) after it had been ‘run in’ for a few thousand positioning cycles. This, along with the use of adaptive calibration, vastly improved the early closed-loop positioning results. Figure 6.4 shows an RMS positioning accuracy of $2.9\mu\text{m}$ being achieved in just six moves, with all cycles achieving $<10\mu\text{m}$ errors. This already meets our performance goals, and at a drive amplitude range of only $\pm 8.0\text{V}$.

PART II: A NEW TILTING SPINE MOTOR

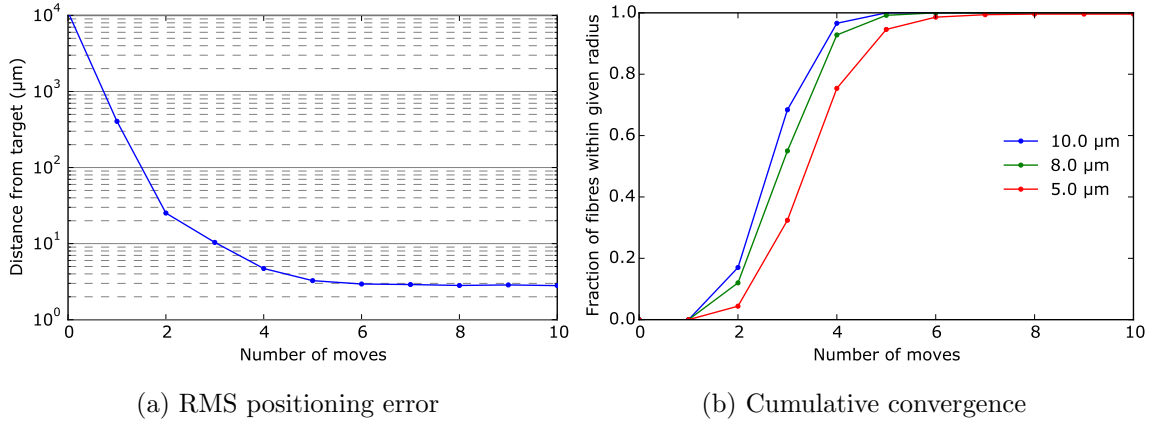


Figure 6.4: The initial MLP motor prototype showed excellent performance after ‘running in’. Details: Prototype 2; Spine 1; 500 random targets; 8.0 V/75 Hz coarse; 3.4 V/75 Hz fine; $R_{patrol} = 11.2$ mm; $Zd = 180^\circ$; component moves; adaptive calibration on (fine mode).

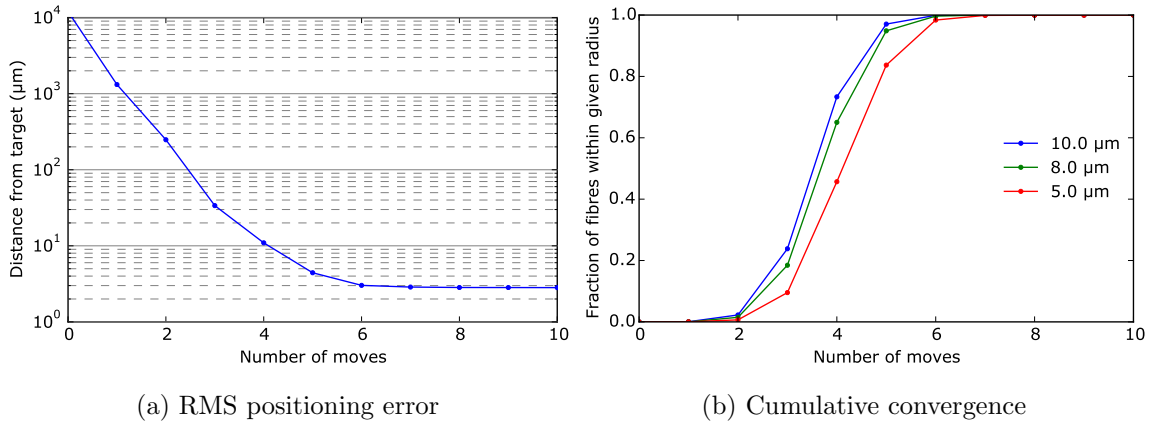


Figure 6.5: A new MLP motor prototype with a modified cup design showed very similar performance to the initial prototype and was able to sustain it for longer. Details: Prototype 5; Spine 2; 3500 random targets; 8.0 V/75 Hz coarse; 2.8 V/5 Hz fine; $R_{patrol} = 12.0$ mm; $Zd = 90^\circ$; component moves; adaptive calibration on (fine mode).

Prototype 3, which had the same format as Prototype 2, was found to achieve almost identical performance (2.8 μm RMS final positioning accuracy) after a similar running-in period. The same evidence of wear was also seen.

The wear issue prevented good performance being achieved for more than ~ 500 successive targets. Prototype 5, with the modified cup design, solved this problem and also eliminated the need for running in. Figure 6.5 shows similar, although ‘slower’, performance over 3500 cycles. Note the bump in Figure 6.5a at move two, corresponding to the 50 μm switching radius from coarse to fine mode. This suggests that drift in the coarse mode behaviour is to blame, because adaptive calibration was only enabled for fine moves. Figure 6.6 appears to confirm this.

PART II: A NEW TILTING SPINE MOTOR




Figure 6.6: Throughout 3500 successive targets (a period of 21 h), coarse moves (upper points) become generally less accurate while fine moves (lower points) maintain accuracy. This suggests that the adaptive calibration is working, since it was only enabled for fine moves. M1–M6 are move numbers; RMS distances for each move are indicated by the triangular markers. These results are for the same test as Figure 6.5.

The coarse move drift was later remedied by scheduling an automatic coarse mode recalibration after every 500 targets (~ 6 h)¹. Following this fix, a much longer test was run over an uninterrupted period of 25 days. Figures 6.7 and 6.8 show that excellent positioning performance is maintained, with an RMS positioning accuracy of 3.2 μm after six moves and 99.994% of those cycles achieving ≤ 10 μm errors.

The 50 000-cycle test could almost count as an instrument lifetime in terms of target numbers, but longevity will be discussed separately in a later section.

¹Recall that adaptive calibration is not possible for compound moves without incurring a time penalty (refer to Section 5.7.2.5, p. 185).

PART II: A NEW TILTING SPINE MOTOR

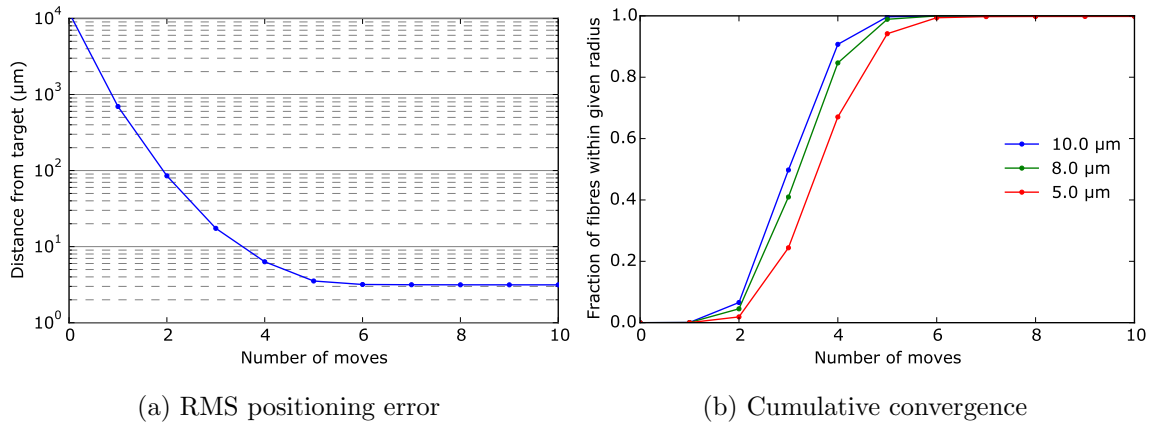


Figure 6.7: A long-term test showed excellent performance throughout, with periodic coarse mode recalibration improving the accuracy of the first two moves. Details: Prototype 5; Spine 2; 50 000 random targets; 8.0 V/75 Hz coarse; 2.8 V/5 Hz fine; $R_{patrol} = 12.0$ mm; $Zd = 180^\circ$; component moves; adaptive calibration on (fine mode); 500-move coarse mode recalibration period.




Figure 6.8: Throughout 50 000 successive targets (a period of ~ 25 d), performance is shown to be stable. There is no evidence of the previous wear issue with the modified cup design. Accuracy actually improves after the first week of testing, which may be the result of the adaptive calibration algorithm, but is more likely due to the lab being vacated over the 2015 festive season^a. M1–M6 are move numbers; RMS distances for each move are indicated by the triangular markers. These results are for the same test as Figure 6.7.

^aThis plot may also be evidence of a healthy Australian work–life balance.

6.3.1.1 Position-dependence

Recall from the introduction to this chapter (Table 1.4, p. 127) that ideally a spine would show identical performance regardless of location in the patrol area. One way to check this is by plotting the number of moves required to achieve a certain accuracy based on the target location. Figure 6.9 shows such a plot. A slight spatial bias does appear to be present, with spines achieving $\leq 10\ \mu\text{m}$ errors in fewer moves when the target is closer to the spine's home position (the centre of the patrol area). This may be caused by a spine imbalance, tension in the fibre, or a slight deformation of the spine's ball.

While the observed spatial variations may not be ideal, the overall performance of the positioner for any random target is almost always within our goal of a $\leq 10.0\ \mu\text{m}$ error in six moves or fewer. Therefore this phenomenon is considered insignificant.




Figure 6.9: This MLP motor and spine combination was able to achieve a $\leq 10.0\ \mu\text{m}$ error in fewer moves when targets were closer to the centre of the patrol area. These results are for the same test as Figure 6.7 ($Zd=180^\circ$).

6.3.1.2 Speed

Until now there has been no mention of the time taken to position a fibre, because achievable accuracy in a certain *number of moves* has been considered more important. The number of moves is so crucial because positioning time is likely to be dominated by the camera exposure and image processing time of the fibre measurement procedure, rather than by the spine move time itself.

Figure 6.10 shows the cumulative amount of time that the fibre is actually moving for each target in the 50 000-target test reported previously. This has been called the ‘positioner busy’ time. To give some context, this is required to be <60s for the 4MOST/AESOP positioner. This single-spine test easily satisfies that requirement, but of course the figure would increase for a positioner comprising many spines that must all finish moving before a feedback image is captured.

From the plot we see that every positioning cycle completes its sixth move in <15s, giving a mean move duration of ~ 2.5 s. Position measurement for thousands of fibres is likely to take longer than this², so spine move time is not the critical factor. Furthermore, the fact that spine move calculations, collision avoidance planning, and software–hardware communication overheads are not included, further adds to the per-move time cost. The number of moves to reach a certain positioning error should therefore be our main concern.




Figure 6.10: The cumulative amount of time spent waiting for a spine to reach its destination (not including move calculations) is relatively small when compared to the position feedback and control software overheads that would occur for every closed-loop iteration. The initial coarse moves will naturally be the longest, so speed is mostly influenced by the coarse mode frequency. M1–M6 are move numbers. These results are for the same test as Figure 6.7.

²AESOP design documentation allows 5.0s for the fibre measurement process.

6.3.2 Blended moves

The building of the hardware and software to enable blended moves came late in the development process, so much of the important analysis had been done for component moves, as we have seen. Much of this still stands, so will not be repeated.

The blended moves principle was proven and shown to provide faster convergence on targets. This is in addition to the time-saving advantages offered in terms of simplified collision avoidance. Figure 6.11 shows the results of a short run of 100 random targets, with blended moves performing very well indeed. A 1 V increase in the coarse mode amplitude was required to maintain similar step sizes to component moves, because the ‘drive angle’ algorithm does result in a reduction of the differential actuator displacements in the component directions (refer to equation 5.1, p. 165).

Successfully proving the blended moves concept was a significant achievement, as this feature makes the entire tilting spine technology more flexible. It also enabled full adaptive calibration to be used on both coarse and fine moves.

All results reported from here onwards will use blended moves unless otherwise specified.

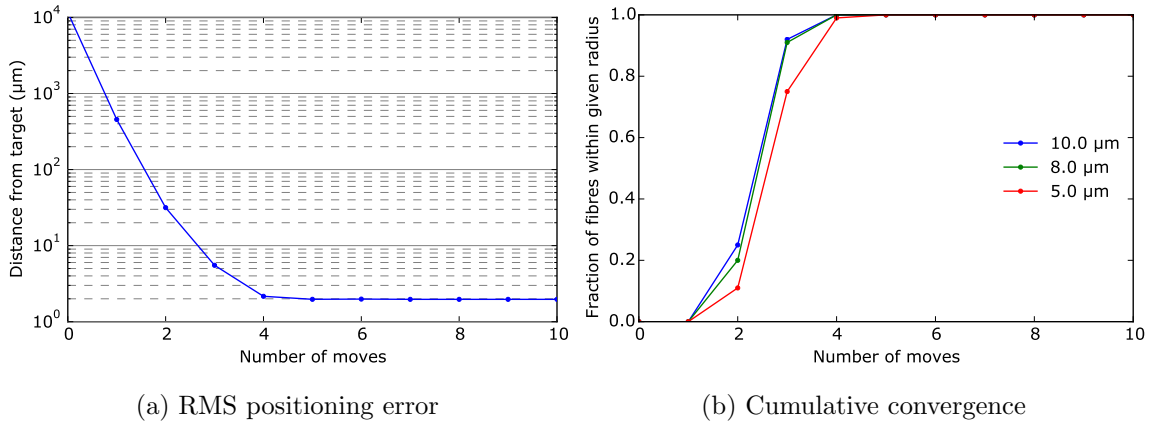


Figure 6.11: As hoped, blended moves showed much faster convergence. Details: Prototype 5; Spine 2; 100 random targets; 9.0 V/75 Hz coarse; 3.5 V/5 Hz fine; $R_{patrol} = 12.0$ mm; $Zd = 180^\circ$; blended moves; adaptive calibration on.

6.4 Gravity vector response

A tilting spine positioner will often be placed at the prime or Cassegrain focus of a telescope, and as such must perform well regardless of the telescope's pointing angle. Tilting spine positioners are designed for this: Each spine is balanced about its pivot point, and magnet strengths are high enough to hold all spines securely in their motor assemblies regardless of the local gravity vector.

Excellent results have already been presented for spines that are both vertical and horizontal in orientation, but these were also calibrated at those orientations. Therefore, a further experiment was carried out to ensure that *changing* the gravity vector under the *same* calibration, as would occur in reality, did not present a problem.

Figure 6.12 shows that suddenly switching the spine orientation from pointing downwards, to pointing horizontally, to pointing upwards, has no significant effect on the positioning accuracy.

We can conclude that performance is sufficiently stable even under extreme changes of the local gravity vector. Further comfort can be found in the fact that real telescope zenith angles rarely exceed 60° .




Figure 6.12: Changing the spine's orientation from $Zd=180^\circ$ (left panel) to $Zd=90^\circ$ (centre panel), to $Zd=0^\circ$ (right panel) had no significant effect on the final positioning accuracy in this test. RMS values are indicated by the triangular markers. The same calibration data was used throughout. The coloured markers from top to bottom represent moves 1–6.

6.5 Further characterisation

Much of the early spine testing used waveform frequencies and amplitudes that were chosen empirically, usually after a limited amount of trial-and-error testing and reviewing the results of calibration routines under different settings.

More complete data on how a spine responds to different waveform parameters was obtained by running many calibration routines in succession and cycling through different settings. This created a better picture of the conditions under which the spine performed best.

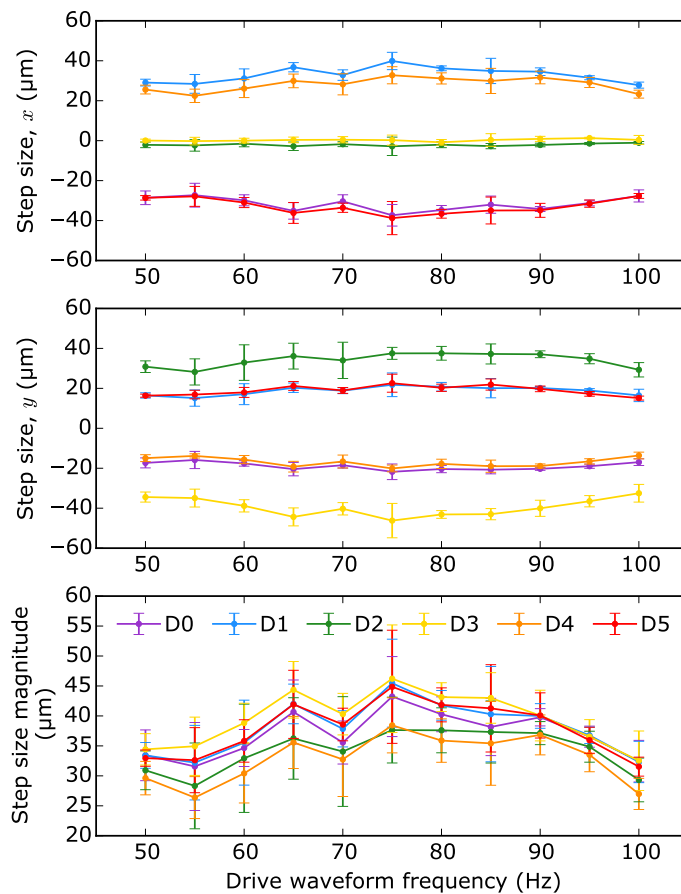


Figure 6.13: The mean coarse step size of a spine varies considerably with drive frequency, as does repeatability. These plots show the step size in x and y as a means of assessing direction stability (top panels), and the overall step size magnitude (bottom panel). Details: Prototype 5; Spine 2; 227 calibration cycles per frequency; amplitude 8.0 V; component moves; error bars show one standard deviation (top panels) or RMS scatter (bottom panel); the legend for D0–D5 applies to all panels and indicates the component direction number.

PART II: A NEW TILTING SPINE MOTOR




Figure 6.14: The mean fine step size of a spine is quite stable at low frequencies. These plots show the step size in x and y as a means of assessing direction stability (top panels), and the overall step size magnitude (bottom panel). Details: Prototype 5; Spine 2; 227 calibration cycles per frequency; amplitude 2.8 V; component moves; error bars show one standard deviation (top panels) or RMS scatter (bottom panel); the legend for D0–D5 applies to all panels and indicates the component direction number.

6.5.1 Frequency selection

Resonance in the spine and/or motor assembly means that certain drive frequencies will result in less repeatable behaviour than others. The first characterisation test involved fixing nominal coarse and fine mode voltages and running calibration routines throughout a range of frequencies. Fine mode frequencies were kept low (a few hertz) on the assumption that this would avoid resonant modes and inertial effects, thus keeping scatter at a minimum for more precise moves. The results confirm this.

Note that calibration was for component moves rather than blended moves in order to save time. This also allows the directional performance differences to be more easily seen when plotted.

PART II: A NEW TILTING SPINE MOTOR

Figure 6.13 shows an upper frequency range for coarse moves, revealing a significant variation in behaviour. Rather unfortunate is that a frequency of 75 Hz, used in many of the tests reported so far, shows poor repeatability. The silver lining here is that we have seen impressive results despite this. The best frequency appears to be in the range 90–95 Hz, where step sizes are relatively even and stable.

Figure 6.14 shows a lower frequency range for fine moves. As expected, the results are less variable than at high frequencies. A frequency of 5 Hz seems to be slightly more even between directions.

6.5.2 Amplitude selection

Selecting drive amplitudes is a more involved task than selecting the frequencies, as discussed previously (refer to Section 5.8, p. 188). Assuming that we are taking the traditional approach of fixing the voltage and not the step size, then the plots in Figure 6.15 can be used for guidance.

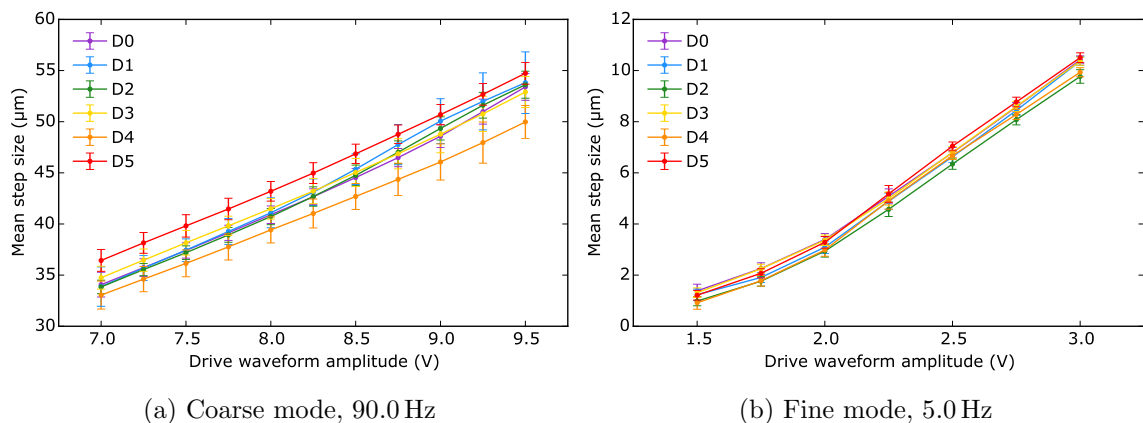


Figure 6.15: Coarse mode step sizes (a) and fine mode step sizes (b) have a generally linear relationship to drive voltage. Coarse move behaviour for this test is generally less repeatable above 8.5 V. Fine move behaviour appears stable down to very fine step sizes ($<2\mu\text{m}$), but will be more sensitive to variations in the motor and therefore less robust. Details: Prototype 5; Spine 2; 81 calibration cycles per voltage; component moves; error bars show RMS scatter; D0–D5 are component direction numbers.

The results show good linearity and repeatability with changing drive amplitude, although performance does appear to suffer if the coarse voltage is set much higher than $\sim 8.0\text{ V}$. Fine mode voltage selection depends on how important accuracy is for the particular application, as higher voltages will be more impervious to changing conditions. The complexity of this trade study is instrument-specific so will not be explored here. What we see is generally good linearity down to step sizes of only a few microns, which is a positive outcome.

6.6 Open-loop repeatability

Good open-loop repeatability is advantageous for two reasons: i) full closed-loop positioning will be faster, requiring fewer corrective moves; and ii) objects that move due to differential refraction and/or field rotation can be tracked across the focal surface without disturbing on ongoing observation.

Early tests of open-loop repeatability for the MLP motor focussed on the stability of step sizes across many successive calibration routines, often revealing very little scatter in the data. This paints a misleading picture, however, because in reality there are more variables to consider.

Figure 6.16a reveals the true ability of the tested prototype to move from a random point to another random point, based on the distance of the move. We see much larger relative errors for the finest moves, with $>100\%$ discrepancies for single steps that were supposed to be $\sim 5\mu\text{m}$ in size. Larger distances, made up of multiple steps, have significantly smaller relative errors. Two possible explanations for this are: i) hysteresis in the actuators being more significant for single steps (discussed later); and ii) spine inertia helping to maintain the ‘heading’ (direction) of longer moves.




Figure 6.16: The open-loop repeatability of any single move depends on the distance of the move. These plots encompass errors in direction as well as magnitude; they show the distance between the actual result of a move and the expected result. Results are for the MLP motor with blended moves (a) and the best available comparative data from the 4MOST/AESOP prototype positioner with compound component moves (b)^a. 110 000 data points are shown in (a), taken from a long-term closed-loop positioning test. Only short moves are plotted as these are most significant for object tracking. Note the periodic evidence of step multiples, and the population of coarse mode moves beginning to appear at a distance of $\sim 40\mu\text{m}$. Axis labels are small in (b) but match (a). It should be noted that the AESOP prototype positioner has not been optimised for open-loop moves.

^aPlot courtesy of Scott Smedley / Australian Astronomical Observatory

PART II: A NEW TILTING SPINE MOTOR

We also see that multiples of a certain distance are less prone to error, which may be due to spines behaving more predictably in particular directions, or may simply be an artefact of the calibration and/or move calculation algorithms. One can imagine a ‘perfect’ system with equal step sizes always; in this case the plots would be entirely discontinuous, with discrete populations of data points in integer multiples of the step size. It follows that the periodic pattern should perhaps be unsurprising.

6.6.0.1 Comparison with existing technology

Figure 6.16b shows the best available data for single-move errors in the prototype 4MOST/AESOP positioner, which uses the latest form of tilting spine technology. 4MOST has no requirement for open-loop positioning, so no attempts have been made to achieve good performance in this area, but the errors of single moves are still relevant to closed-loop positioning. There are no moves smaller than 10 μm for AESOP due to the way the control system is programmed, and there is a further difference in that the AESOP results are from multiple spines.

It would be unfair to make an absolute judgement on open-loop capabilities based on the current data. At most we could claim a marginal improvement with the MLP motor on the basis of lower maximum errors. Apart from this, average error magnitudes appear similar at around 5 μm for moves in the 10–20 μm range.

6.6.0.2 Object tracking capabilities

The requirements on object tracking are dependent on several factors, including the fibre core size, the observatory seeing conditions, the telescope field of view, and of course the airmass for a given pointing. As such, it’s hard to come up with a general measure of how useful the prototype motor is. What we can say is that it’s probably better to make one large correction than several small ones.

Let’s assume for the sake of an example that an instrument has a fibre positioning accuracy requirement of 10 μm , but that fibres are initially positioned to 5 μm accuracy via closed-loop control. Now let’s assume a single open-loop corrective move will be made during the observation, for the fibres that are most significantly affected by changes in differential refraction as the telescope tracks (at the edge of a field). If positioning errors add in quadrature, then our corrective move must have an absolute error of $\sqrt{10^2 - 5^2} \approx 8.7 \mu\text{m}$. Referring to Figure 6.16a, this permits a corrective move of up to $\sim 12 \mu\text{m}$ or $\sim 22 \mu\text{m}$ for 3σ or 2σ confidence, respectively.

6.6.1 Hysteresis and ‘memory effect’

It is known that Echidna spines show a ‘memory effect’, in that the result of future moves will always depend on the past moves [24]. This is one of the reasons that adaptive calibration was first proposed by the AAO, although not an ideal solution for such a high-order problem.

A simple test of the presence of a memory effect in the MLP motor was devised, comparing the results of two types of blended move calibration routines: i) the original approach to calibration, moving the spine in successive drive angle increments of 18° ; and ii) the same approach as (i), but following each measured calibration move by an unmeasured ‘backtrack’ in the opposite direction. The results (Figure 6.17) prove that small step sizes are heavily dependent on the direction of the previous move. This will have a major effect on open-loop repeatability, and indeed it can explain the large relative scatter seen for the smaller moves in Figure 6.16a.

A plausible explanation for the memory effect is hysteresis in the piezo actuators, where it is known that lower-voltage (or ‘softer’) piezoceramics exhibit worse hysteresis than their higher-voltage counterparts. The MLP motor’s use of three actuators under various operating conditions makes this phenomenon a challenge to understand analytically.

It is proposed that the effects of hysteresis could be reduced in two ways, although these have not been explored:

- **Hysteresis modelling:** It may be possible to model the result of a particular move as a function of the previous move(s). A suitable empirical model may become apparent with a carefully-constructed test, exploring all relevant variables. There is a pattern to the data in Figure 6.17b that hints at a cumulative effect from many past moves, hence this model would almost certainly need to be of a high order.
- **Actuator ‘reset’ procedure:** There may be a way to ‘ring out’ the actuators in the motor before doing a fine move, e.g. with a high-frequency sinusoidal signal that returns the piezoceramic material to a more neutral electromechanical state. Of course, such a procedure must have a negligible effect on the spine’s position.

Until the ‘memory effect’ phenomenon is better understood, it appears that tilting spine positioners will be inherently limited in their ability to make small corrective moves without position feedback.

PART II: A NEW TILTING SPINE MOTOR

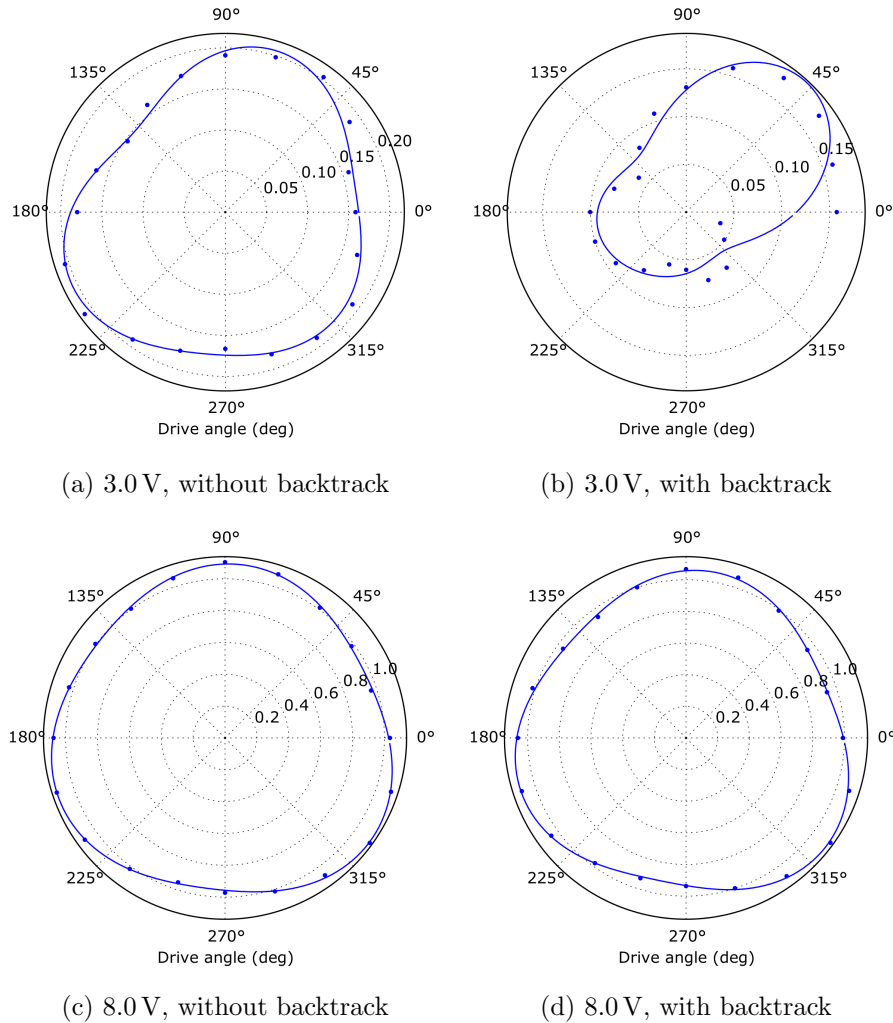


Figure 6.17: Blended move calibration routines, which usually increment the spine drive angle, have very different results if every move is followed by a ‘backtrack’ in the opposite direction. The effect is much more significant for fine moves (a \rightarrow b) than coarse moves (c \rightarrow d). This proves that the fine step size is highly dependent on the previous move(s). Note the definite pattern to the data points in (b), making a modelling attempt almost irresistible. Details: Prototype 7; Spine 2; polar axes show step size in pixels ($\sim 25 \mu\text{m}/\text{px}$).

6.6.2 Advantages of scatter

It’s worth noting that, with the exception of open-loop object tracking, an amount of random (or pseudo-random) scatter in a spine’s step size can improve the final positioning accuracy of an iterative system such as ours. In a perfectly repeatable system, positioning resolution truly is limited by the minimum achievable step size. With some additional random scatter, however, there is a chance that a fibre will end up closer to the target. This explains the observed MLP motor positioning accuracies of $<3 \mu\text{m}$, despite an average fine mode step size of $\sim 8 \mu\text{m}$.


6.7 Increasing spine length

A key aim of this work was to improve the motor design to a degree that would allow spines to be lengthened, reducing the optical losses caused by tilt-induced FRD.

Spine length has previously been limited to ~ 250 mm by the minimum achievable step size of the spines and by mechanical limitations involving resonance. The hope was that the new MLP motor would overcome the latter limitation due to its reduction of forces orthogonal to the spine (i.e. actuators now move along the spine axis).

Recall from the introduction that the relationship between tilt angle (and therefore spine length) and optical losses is approximately quadratic. With this in mind, two experiments were carried out with spine length increases of $\sim 1.2\times$ and $\sim 1.4\times$, corresponding to optical loss reductions of $\sim 35\%$ and $\sim 50\%$ respectively. Figure 6.18 shows successful positioning results in both cases, with accuracies reducing only by the factor of length increase. The longest spine is able to achieve $4.0\ \mu\text{m}$ RMS accuracy in six moves, and therefore still meets our ultimate positioning goal.

This is further proof that the MLP motor is superior to the existing technology; optical losses can be halved while maintaining excellent positioning performance.



(a) Spine 3: 309 mm (-35% optical loss)

(b) Spine 4: 354 mm (-50% optical loss)

Figure 6.18: Longer spines performed very well in two short closed-loop positioning tests of 200 targets. The test setup was otherwise equivalent to that of Figure 6.5 (p. 196). The positioning errors after each move scale very closely with spine length, as hoped.

6.8 Removal of ground connection

The original MLP motor design had four electrical connections: three ‘active’ drive waveforms connected to the outer electrodes of each actuator, and a ground connection to all three inner electrodes. The grounded electrodes are all electrically bridged via the metal cup.

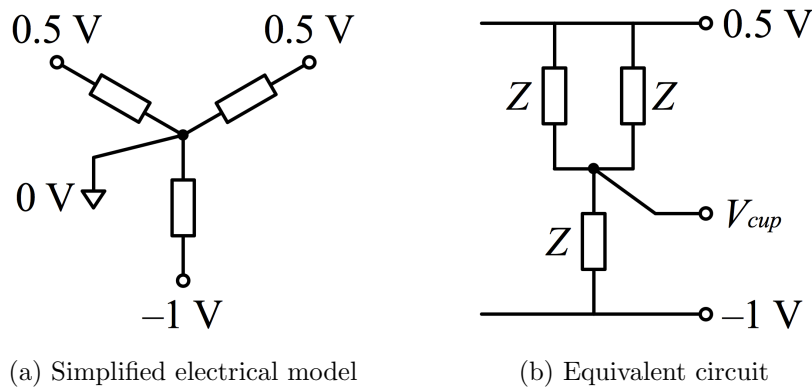


Figure 6.19: There is a synergy between the calculations for blended move actuator amplitudes and the equivalent electrical circuit of an MLP motor, which results in the potential V_{cup} being implicitly 0 V without the need for a ground connection. One of the four motor connections can therefore be eliminated. Actuators are modelled here as simple impedances, and the example voltages are for a blended move in a direction pointing up the page.

If we consider the equivalent electrical circuit for this arrangement, we find that the ground connection should not be strictly necessary. Figure 6.19 shows a simplified diagram for the case of a blended move, with the potential on the cup being *implicitly* zero without having to make the connection. This holds as long as the actuator impedances match, which they should within a nominal 10% tolerance. In any case, any minor asymmetry in this respect would only skew the step size vs. drive angle relationship, and would be accounted for by calibration.

Figure 6.20 shows blended move calibration data for the same prototype motor with the ground connected and disconnected. We see practically no difference in the results, confirming that the ground connection is superfluous for blended moves. With this the MLP motor has become even more attractive, requiring only three electrical connections to the easily-accessible outer electrodes of its actuators. This would make design, assembly and integration in a full-size positioner much easier.

PART II: A NEW TILTING SPINE MOTOR

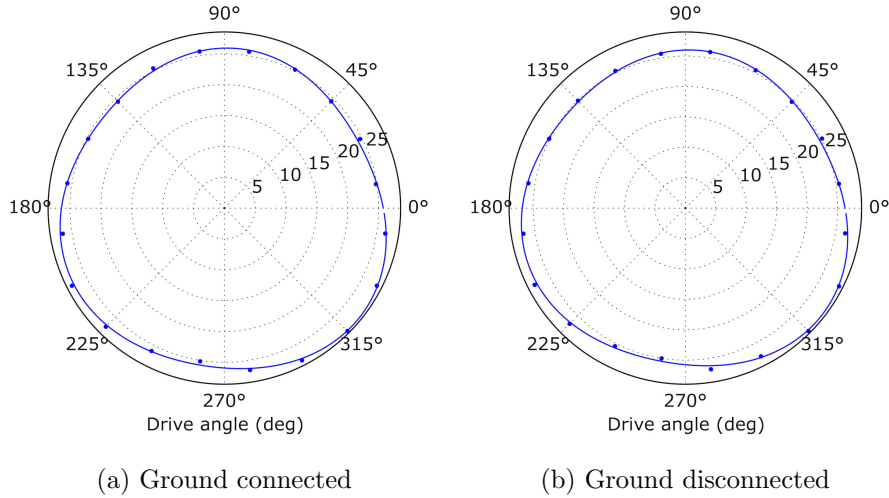


Figure 6.20: As expected, removing the ground connection on Prototype 5 had no discernible effect on its blended move calibration results. Polar axis shows step size in μm .

6.9 Longevity

The final test of the MLP motor conducted for this research concerned the mechanical robustness of the design and its ability to maintain good performance over a very large number of positioning cycles.

The lifetime of a MOS instrument is highly dependent on the observing strategy. The 4MOST instrument is a good example of high demands on a positioner throughout its planned 15 yr operational life, because of its relatively short ‘observation block’ size (exposure time) of 20 min. Every AESOP spine must be able to operate for

$$15 \text{ yr} \times 350 \text{ nights/yr} \times 10 \text{ h/night} \times 3 \text{ fields/h} = 157\,500 \text{ cycles} \quad (6.1)$$

At the time of writing, MLP ‘Prototype 5’ has completed $>160\,000$ closed-loop positioning cycles in total, with no signs of performance degradation. The most recent 100 000 targets were part of an eight-week-long test that brought together many of the core hardware and software developments described in the previous sections: the miniaturised drive amplifiers, blended moves, adaptive calibration, and the use of only three signal connections (ground removed). A zenith angle of 45° was chosen to simulate general telescope pointing.

PART II: A NEW TILTING SPINE MOTOR

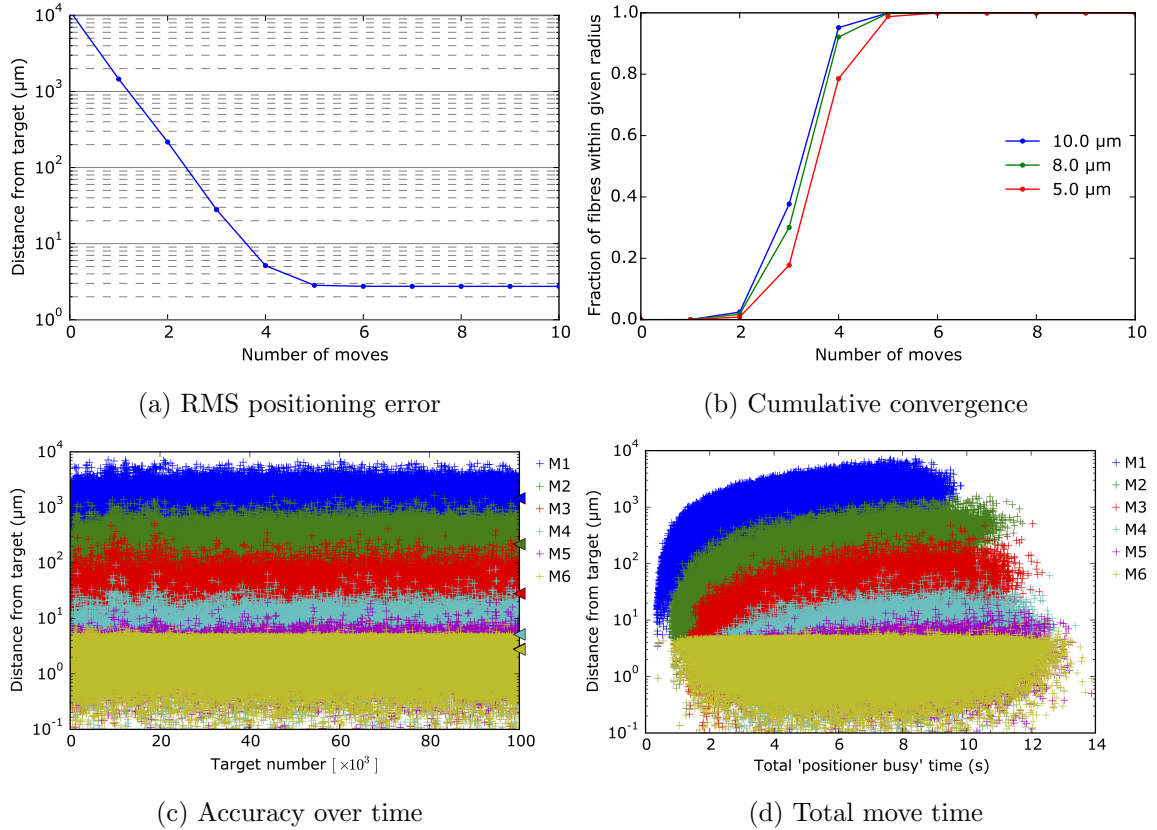


Figure 6.21: Throughout 100 000 successive targets (a period of ~ 60 d) the prototype MLP motor showed consistent and accurate closed-loop performance with an RMS positioning error of $< 2.8 \mu\text{m}$ after five moves, with negligible improvement thereafter. The maximum final error across all targets was $5.8 \mu\text{m}$. M1–M6 in (c) and (d) are move numbers. Details: Prototype 5; Spine 2; 100 000 random targets; 9.0 V/84 Hz coarse; 3.5 V/5 Hz fine; $R_{patrol} = 11.5 \text{ mm}$; $Zd = 45^\circ$; blended moves; adaptive calibration on.

Figures 6.21 and 6.22 show that performance is excellent, even after a number of positioning cycles that would far outlast some MOS instruments' lifetimes. $< 2.8 \mu\text{m}$ errors are achieved in five moves, at which point 98.76 % of cycles were within $5.0 \mu\text{m}$ of the target. 99.97 % of cycles were within $5.0 \mu\text{m}$ in six moves, and the maximum final error across all targets was $5.8 \mu\text{m}$. These results go a long way in proving the robustness of the MLP motor's design, particularly concerning the strength of the bond between the actuator layers (refer to Section 4.3.5, p. 4.3.5). The control and calibration algorithms have also been proven suitable.

This testing was done under normal laboratory conditions, so does not account for the environmental changes that would occur throughout a decade or more on a real telescope. This would be a natural next step in the verification of this technology.

PART II: A NEW TILTING SPINE MOTOR

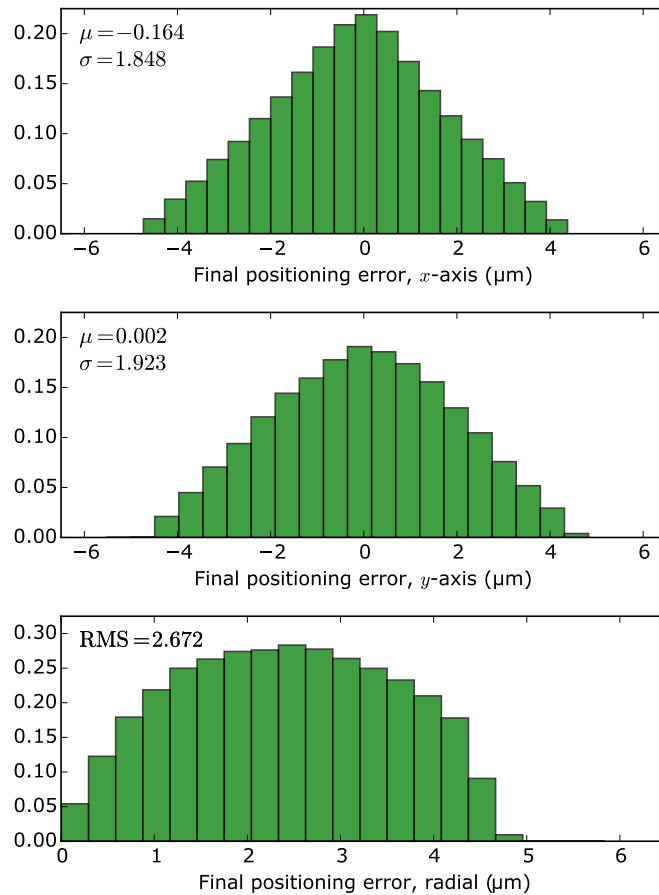


Figure 6.22: The final positioning error (once the spine was closer than a single step size) had an even distribution, suggesting that the control approach works well for any arbitrary move. These results are for the same test as Figure 6.21 (100 000 targets).

6.10 Section summary

In this section we have reviewed the most important results of extensive performance testing for various MLP motor prototypes. We have seen how the AAO's latest cup design has solved wear issues that were experienced with the first prototypes, allowing longevity tests that have so far exceeded 160 000 positioning cycles.

A closed-loop accuracy of 2.8 μm RMS has been seen across three different prototypes and has not worsened after longevity testing. This accuracy is achievable in just five moves when using blended moves. Our performance goal of $\leq 5 \mu\text{m}$ RMS error in six moves is nearly satisfied in *four* moves (5.1 μm).

PART II: A NEW TILTING SPINE MOTOR

Changes in the local gravity vector (zenith angle) have been shown to have no significant effect on positioning performance. One of the four electrical connections on the original MLP motor design (the most fiddly one) has been removed without affecting performance.

The MLP motor is able to position spines up to $1.4\times$ the existing length limit with an increase in positioning error no larger than the associated $1.4\times$ increase in step size, meaning our accuracy target is still achieved. Spines of this length (354 mm) cut tilt-induced optical throughput losses in half.

The motor has been proven to be robust, as has the control hardware and software. Blended moves work as expected, presenting a flexible collision-avoidance solution when controlling multiple spines.

The maximum operational voltage range for all testing was a mere $\pm 9\text{ V}$, easily accommodated by a $\pm 15\text{ V}$, $\pm 12\text{ V}$, or $\pm 10\text{ V}$ power supply, all of which are common in electronics systems.

PART II: A NEW TILTING SPINE MOTOR

7 Conclusions and prospects

This research had a rather broad set of goals, aimed at improving several aspects of tilting spine technology in order to make it more attractive for future astronomical instruments. Of high priority was reducing the drive voltage of the spine actuators.

A new motor has been designed that uses low-voltage multilayer plate (MLP) piezo actuators. The MLP motor remains true to the stick-slip principle of the traditional Echidna technology and requires no changes to the existing spine format. Moreover, it utilises the latest AAO spine mount design.

Closed-loop positioning has been demonstrated that outperforms the existing tilting spine technology in terms of both accuracy and speed. Drive voltage has been reduced by more than an order of magnitude, to just ± 9 V.

The MLP motor's demonstrated RMS closed-loop fibre positioning accuracy currently beats all other published methods for automated parallel fibre positioning, at $< 2.8 \mu\text{m}$ in only five moves. This has been shown to hold for $> 150\,000$ targets over 90 days of non-stop operation.

Echidna technology's lowest possible tilt-induced throughput loss has been halved by achieving a $\sim 1.4\times$ increase of the current spine length limit, to ~ 354 mm. This means that: a) survey SNRs improve, increasing scientific value; and/or b) spectrograph collimator speeds can be more efficiently matched to the nominal f-number of the instrument, reducing costs. Good accuracy is maintained ($4.0 \mu\text{m}$ RMS in six moves) and a further length increase looks plausible.

A new control architecture has been designed and tested, made possible by lower voltage demands. This provides every actuator on every motor with a customisable drive waveform, meaning that each spine can now achieve its full performance potential. Spines can also move in any direction, simplifying collision avoidance strategies.

The cost of implementing the new motor in a full positioner appears similar to the existing technology, with reasons to believe that the overall size of a full MLP motor control system will be smaller than it is at present. There is potential, therefore, to reduce the size and/or numbers of electronics cabinets required on-telescope.

PART II: A NEW TILTING SPINE MOTOR

An additional avenue of investigation concerned fully modularising the control electronics to allow a distributed control system with stand-alone spine assemblies. This was shown to be feasible, but likely to require additional cooling measures.

Adaptive calibration algorithms have been designed and implemented that automatically compensate for changes in spine behaviour over a range of timescales. Spatial variation in spine performance is also automatically accounted for. These algorithms learn from normal fibre positioning operations and therefore come with no time penalty to field reconfigurations.


The design of the MLP motor is such that mechanical crosstalk between neighbouring spines in a full positioner should be reduced as a matter of course, although this has not been quantitatively measured. The improved design, along with the fact that waveform phases can now be staggered to reduce vibration magnitudes, significantly reduces the threat that crosstalk poses to positioning accuracy.

Table 7.1 summarises the outcomes of this research against the development goals set out at the beginning. In this respect we can conclude that the work has been successful. Also recall our general wish to improve tilting spine technology to a level that could compete with the high-accuracy radial arm positioners of recent years; Figure 7.1 shows that this has been accomplished.


Table 7.1: Comparison of outcomes with the goals of this research.

General goal	Specific goal	Achieved
Reduce drive voltage range	± 24 V (pref. ± 15 V)	± 9 V
Improve RMS closed-loop accuracy	≤ 5.0 μm (pref. ≤ 3.6 μm) in six moves	< 2.8 μm in five moves
Increase max. spine length to reduce losses	None set; must maintain goal accuracy	$\sim 1.4\times$ increase, halving losses; 4.0 μm accuracy in six moves
Improve open-loop repeatability	None set	Marginal improvement evident
Reduce actuator forces orthogonal to spine	None set	Achieved by design

PART II: A NEW TILTING SPINE MOTOR



(a) RMS positioning error



(b) Cumulative convergence (error $\leq 5.0 \mu\text{m}$)

Figure 7.1: In five moves, the MLP motor is able to beat all other published technologies in terms of RMS closed-loop positioning accuracy (a) and consistency across multiple targets (b). This includes the Cobra [29] and ES-CH [32] radial arm ($\theta-\phi$) positioners, although we see that the latter of these can achieve smaller errors in its first four moves. In any case, the MLP motor is a clear competitor.

PART II: A NEW TILTING SPINE MOTOR

Improvements to open-loop capabilities were somewhat difficult to evaluate due to a lack of comparative data from the existing technology, but look similar. Analysis of the MLP motor’s general move errors as a function of distance has shown that small, single-step moves suffer a much larger relative error than multi-step moves. Moreover, the influence of previous moves on fine step sizes, probably from actuator hysteresis, has been shown to be substantial.

The MLP motor technology has so far been developed to a Technology Readiness Level (TRL) of 4, meaning that the technology has been validated in a laboratory environment. It is suggested that future work shall include accelerated life testing with relevant temperature and humidity levels to further prove the technology’s suitability for use in an observatory environment, therefore raising the TRL.

Future work could also include: testing of more prototype assemblies, possibly in a group, to increase confidence in consistent performance; investigation of the effect of ball–cup contact angle on performance; a miniaturisation exercise to match the ~ 7 mm pitch previously achieved with the existing Echidna technology [27]; and attempts to model the memory effects seen in small moves, to enhance open-loop repeatability for ‘blind’ object tracking during exposures.

Excepting a reasonable amount of further validation, there appears to be no technical reason why the MLP motor should not become the new standard in tilting spine fibre positioning.

This work has been recently been accepted for presentation at the 2016 SPIE Astronomical Telescopes + Instrumentation conference in Edinburgh, in a paper currently entitled “Echidna Mark II: one giant leap for ‘tilting spine’ fibre positioning technology” [57].

Appendix

A Piezo actuator background

This appendix summarises the nature of piezoelectric materials, how they are manufactured, and how their microscopic displacements are used in larger motor mechanisms to produce motion over a long range of travel.

A.1 Traditional materials and manufacturing processes

The piezoceramic industry has been producing bulk ceramic actuators for decades. Although a number of synthetic ceramic compounds exhibit piezoelectric properties, the most widely used is lead zirconate titanate, commonly referred to as ‘PZT’.

PZT has substantially better piezoelectric properties than any similar alternative, to such an extent that it currently has an exemption under the Restriction of Hazardous Substances (RoHS) directive, despite its lead content. Many manufacturers offer standardised PZT compounds alongside their own proprietary recipes, doped with ions of different elements and therefore optimised for different applications.

The process involved in making traditional piezoceramic actuators has remained relatively unchanged for some time. Ceramic compounds are mixed in powder form, then heated, wetted, formed, and finally sintered at a high temperature. This procedure is well understood in the context of mass production, meaning that many companies offer off-the-shelf PZT actuators in various shapes and sizes.

Bulk PZT actuators are often a low-risk and cost-effective choice when designing a micro-positioning system. Furthermore, they have properties that have been hard to match with other materials in the past. The existing tilting spine actuators are PZT devices.

A.2 Typical voltages and displacements

The actuator displacement that results from a given electric field strength depends on the properties of the piezo material. Specifically, the piezoelectric charge constant, d , gives the sensitivity of a piezo material when used in a particular mode, in metres per volt. When designing positioning systems for the lowest possible voltage, we generally want d to be as large as possible.

The highest piezoelectric charge constants in traditional ceramic actuators are of the order $100 \times 10^{-12} \text{ m V}^{-1}$, which means a theoretical 10 kV signal is required to produce a displacement of just 1 μm .

A.3 Actuator stacks

The small displacements achievable with longitudinal and shear mode piezo actuators can be increased by stacking multiple elements on top of each other. This has the effect of multiplying the displacement by n , where n is the number of actuators in the stack. The drawback here is that signals need to be connected to each actuator, which is labour-intensive and increases the bulk of the whole assembly. Furthermore, a minimum thickness of around 0.5 mm for monolithic actuators¹ limits the practical number of stacked actuators.

Tilting spine fibre positioner actuators must be cost-effective and compact in order to keep their appeal for massively-multiplexed instruments with densely-packed fields. Therefore, actuator stacking was not considered a viable option for improving the tilting spine motor design.

A.4 Piezo actuation styles

The way in which piezoelectric actuators are used in micro-positioning applications can be split into two broad categories: i) quasi-statically, for low-displacement applications; and ii) dynamically, for high-displacement applications. We are dealing with the latter. More specifically, the motor designs presented in this thesis are all of the dynamic ‘stick-slip’ type.

¹Minimum thickness specification obtained from Noliac A/S, for their shear plate actuators.

A.4.1 Quasi-static drives

The term quasi-static is used here to denote the use of a piezo actuator for relatively low-frequency motion throughout its natural displacement range. In other words, these actuators are used ‘as is’. Examples of quasi-static piezo applications include electron beam steering in scanning tunnelling microscopes, and mirror deformation for wavefront correction in adaptive optics systems; both of these applications, although not ‘slow’, use piezo actuators within their natural ranges of travel.

A.4.2 Dynamic drives

The small displacements produced by piezo actuator elements can be used to produce continuous linear or rotary motion by dynamically driving them as part of a larger mechanism. Examples of dynamic piezo applications include linear stages for laboratory positioning and auto-focus motors in camera lenses; both of these have long or unlimited travel ranges because of how they cyclically drive their actuators at high frequencies, producing many hundreds or thousands of microscopic ‘steps’ every second. The traditional Echidna fibre positioner motor is also in this category, because it must be able to provide relatively fast and long-distance travel (tens of millimetres in a few seconds) while remaining very accurate (a few microns resolution).

Table A.1: Dynamic-drive piezoelectric motor mechanisms generally fall into one of two categories, each with their own unique strengths and weaknesses.

	Discrete-stepping	Stick-slip
Repeatability	Better	Worse
Mechanical wear	Less	More
Complexity	Higher	Lower
Motor size	Larger	Smaller
Cost	Higher	Lower

While there are many dynamic piezo motor designs, almost all are one of two types: i) discrete-stepping motors; or ii) stick-slip motors. The merits of each are summarised in Table A.1. Discrete-stepping motors, such as the inchworm (Figure A.1), offer superior repeatability due to the continuous contact between their actuators and the surface being moved. Stick-slip designs (Figure A.2), with their reliance on breaking friction with high accelerations, are much less predictable by comparison. Nevertheless, stick-slip designs are more suitable for tilting spine positioners because of their simplicity, lower cost, and smaller footprint.

APPENDIX

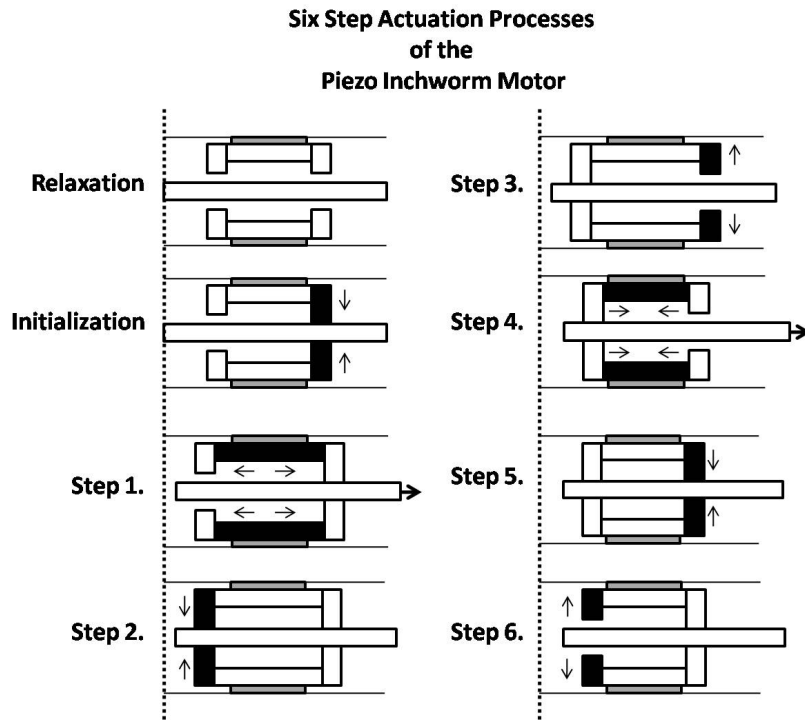


Image credit: User:Jatosado / Wikimedia Commons / Public Domain

Figure A.1: The ‘inchworm’ is a popular type of discrete-stepping piezo motor that produces a small and precise displacement step using several actuators.

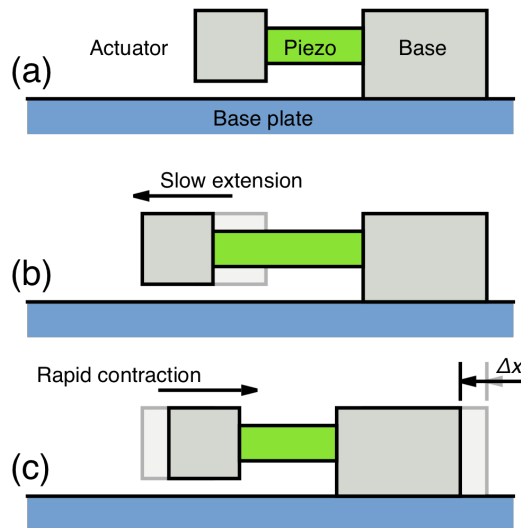


Image credit: User:Inductiveload / Wikimedia Commons / CC-BY-SA-2.5

Figure A.2: The stick-slip principle uses inertia in concert with the difference between static and kinetic friction of materials to produce a small displacement step using only one actuator.

References

- [1] I. S. McLean *et al.*, “MOSFIRE, the multi-object spectrometer for infra-red exploration at the Keck Observatory,” in *Ground-based and Airborne Instrumentation for Astronomy IV*, ser. Proc. SPIE, vol. 8446, Sep. 2012, p. 84460J.
- [2] F. Garzón *et al.*, “Results of the verification of the NIR MOS EMIR,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 91470U.
- [3] S. M. Croom *et al.*, “The Sydney-AAO Multi-object Integral field spectrograph,” *MNRAS*, vol. 421, pp. 872–893, Mar. 2012.
- [4] K. Bundy *et al.*, “Overview of the SDSS-IV MaNGA Survey: Mapping nearby Galaxies at Apache Point Observatory,” *ApJ*, vol. 798, p. 7, Jan. 2015.
- [5] J. S. Lawrence *et al.*, “Towards a spectroscopic survey of one hundred thousand spatially resolved galaxies with Hector,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 91476Y.
- [6] J. M. Hill *et al.*, “Multiple object spectroscopy: The Medusa spectrograph,” *ApJ*, vol. 242, pp. L69–L72, Dec. 1980.
- [7] J. M. Hill, “The History of Multiobject Fiber Spectroscopy,” in *Fiber Optics in Astronomy*, ser. Astronomical Society of the Pacific Conference Series, S. C. Barden, Ed., vol. 3, 1988, p. 77.
- [8] S. A. Shethman *et al.*, “The Las Campanas Redshift Survey,” *ApJ*, vol. 470, p. 172, Oct. 1996.
- [9] M. Colless *et al.*, “The 2dF Galaxy Redshift Survey: spectra and redshifts,” *MNRAS*, vol. 328, pp. 1039–1063, Dec. 2001.
- [10] D. H. Jones *et al.*, “The 6dF Galaxy Survey: final redshift release (DR3) and southern large-scale structures,” *MNRAS*, vol. 399, pp. 683–698, Oct. 2009.

- [11] D. G. York *et al.*, “The Sloan Digital Sky Survey: Technical Summary,” *AJ*, vol. 120, pp. 1579–1587, Sep. 2000.
- [12] D. J. Schlegel and SDSS-III Collaboration, “Overview of the Baryon Acoustic Oscillation Survey (BOSS),” in *American Astronomical Society Meeting Abstracts*, ser. American Astronomical Society Meeting Abstracts, vol. 225, Jan. 2015, p. 125.01.
- [13] J. M. Hill and M. P. Lesser, “Deployment of the MX spectrometer,” in *Instrumentation in astronomy VI*, ser. Proc. SPIE, D. L. Crawford, Ed., vol. 627, Jan. 1986, pp. 303–320.
- [14] I. Parry and R. M. Sharples, “AUTO FIB Current Status,” in *Fiber Optics in Astronomy*, ser. Astronomical Society of the Pacific Conference Series, S. C. Barden, Ed., vol. 3, 1988, p. 93.
- [15] D. Hamilton *et al.*, “A multiobject fiber spectrograph for The Hale Telescope,” *PASP*, vol. 105, pp. 1308–1321, Nov. 1993.
- [16] I. R. Parry *et al.*, “Autofib–2: an automated fiber positioner for the prime focus of the William Herschel Telescope,” in *Instrumentation in Astronomy VIII*, ser. Proc. SPIE, D. L. Crawford and E. R. Craine, Eds., vol. 2198, Jun. 1994, pp. 125–133.
- [17] I. J. Lewis *et al.*, “Autofib–2: commissioning results of a robotic multiobject fiber system for the William Herschel Telescope,” in *Optical Telescopes of Today and Tomorrow*, ser. Proc. SPIE, A. L. Ardeberg, Ed., vol. 2871, Mar. 1997, pp. 1318–1324.
- [18] I. J. Lewis *et al.*, “The Anglo-Australian Observatory 2dF facility,” *MNRAS*, vol. 333, pp. 279–299, Jun. 2002.
- [19] D. H. Jones *et al.*, “The 6dF Galaxy Survey: samples, observational techniques and the first data release,” *MNRAS*, vol. 355, pp. 747–763, Dec. 2004.
- [20] W. Saunders *et al.*, “6dF sees the light!” *Anglo-Australian Observatory Epping Newsletter*, vol. 97, pp. 14–17, May 2001.
- [21] P. Gillingham *et al.*, “Operational performance of OzPoz: the multifiber positioner on the VLT,” in *Ground-based Instrumentation for Astronomy*, ser.

- Proc. SPIE, A. F. M. Moorwood and M. Iye, Eds., vol. 5492, Sep. 2004, pp. 643–650.
- [22] P. R. Gillingham *et al.*, “Echidna: a multifiber positioner for the Subaru prime focus,” in *Optical and IR Telescope Instrumentation and Detectors*, ser. Proc. SPIE, M. Iye and A. F. Moorwood, Eds., vol. 4008, Aug. 2000, pp. 1395–1403.
- [23] A. M. Moore *et al.*, “Spine development for the Echidna fiber positioner,” in *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*, ser. Proc. SPIE, M. Iye and A. F. M. Moorwood, Eds., vol. 4841, Mar. 2003, pp. 1429–1439.
- [24] A. Sheinis *et al.*, “Advances in the Echidna fiber-positioning technology,” in *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation*, ser. Proc. SPIE, vol. 9151, Jul. 2014, p. 91511X.
- [25] W. Saunders *et al.*, “Target allocation yields for massively multiplexed spectroscopic surveys with fibers,” in *Modeling, Systems Engineering, and Project Management for Astronomy VI*, ser. Proc. SPIE, vol. 9150, Aug. 2014, p. 915023.
- [26] R. S. de Jong *et al.*, “4MOST: 4-metre Multi-Object Spectroscopic Telescope,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9147, Jul. 2014, p. 0.
- [27] W. Saunders *et al.*, “MOHAWK: a 4000-fiber positioner for DESpec,” in *Ground-based and Airborne Instrumentation for Astronomy IV*, ser. Proc. SPIE, vol. 8446, Sep. 2012, p. 84464W.
- [28] H. Hu *et al.*, “New type optical fiber positioning unit device for LAMOST,” in *Large Ground-based Telescopes*, ser. Proc. SPIE, J. M. Oschmann and L. M. Stepp, Eds., vol. 4837, Feb. 2003, pp. 548–555.
- [29] C. Fisher *et al.*, “Developing engineering model Cobra fiber positioners for the Subaru Telescope’s prime focus spectrometer,” in *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation*, ser. Proc. SPIE, vol. 9151, Jul. 2014, p. 91511Y.

- [30] M. Cirasuolo *et al.*, “MOONS: the Multi-Object Optical and Near-infrared Spectrograph for the VLT,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 91470N.
- [31] B. Flaugher and C. Bebek, “The Dark Energy Spectroscopic Instrument (DESI),” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9147, Jul. 2014, p. 0.
- [32] N. Fahim *et al.*, “An 8-mm diameter fibre robot positioner for massive spectroscopy surveys,” *MNRAS*, vol. 450, pp. 794–806, Jun. 2015.
- [33] J. Gilbert *et al.*, “Starbugs: all-singing, all-dancing fibre positioning robots,” in *Modern Technologies in Space- and Ground-based Telescopes and Instrumentation II*, ser. Proc. SPIE, vol. 8450, Sep. 2012, p. 84501A.
- [34] J. S. Lawrence *et al.*, “The MANIFEST fibre positioning system for the Giant Magellan Telescope,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Aug. 2014, p. 914794.
- [35] G. Dalton *et al.*, “WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope,” in *Ground-based and Airborne Instrumentation for Astronomy IV*, ser. Proc. SPIE, vol. 8446, Sep. 2012, p. 84460P.
- [36] G. Dalton *et al.*, “Project overview and update on WEAVE: the next generation wide-field spectroscopy facility for the William Herschel Telescope,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 91470L.
- [37] M. Balcells, “The WEAVE spectrograph on the WHT: plans for radial velocity and chemistry surveys of the northern and equatorial Milky Way,” in *EAS Publications Series*, ser. EAS Publications Series, vol. 67, Jul. 2014, pp. 227–230.
- [38] T. Prusti, “Gaia Mission Status,” in *EAS Publications Series*, ser. EAS Publications Series, vol. 67, Jul. 2014, pp. 15–21.
- [39] G. Gilmore *et al.*, “The Gaia-ESO Public Spectroscopic Survey,” *The Messenger*, vol. 147, pp. 25–31, Mar. 2012.
- [40] M. P. van Haarlem *et al.*, “LOFAR: The LOw-Frequency ARray,” *A&A*, vol. 556, p. A2, Aug. 2013.

- [41] M. A. W. Verheijen *et al.*, “Apertif, a focal plane array for the WSRT,” in *The Evolution of Galaxies Through the Neutral Hydrogen Window*, ser. American Institute of Physics Conference Series, R. Minchin and E. Momjian, Eds., vol. 1035, Aug. 2008, pp. 265–271.
- [42] T. Agócs *et al.*, “Final optical design for the WEAVE two-degree prime focus corrector,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 914773.
- [43] D. C. Abrams *et al.*, “The mechanical design for the WEAVE prime focus corrector system,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Aug. 2014, p. 91472K.
- [44] K. Rogers *et al.*, “The design of the WEAVE spectrograph,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9147, Jul. 2014, p. 6.
- [45] I. J. Lewis *et al.*, “Fibre positioning concept for the WEAVE spectrograph at the WHT,” in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Proc. SPIE, vol. 9147, Jul. 2014, p. 914734.
- [46] K. Shortridge, “AAOGlimpse: Three-dimensional Data Viewer,” Astrophysics Source Code Library, Oct. 2011.
- [47] J. Gilbert, “fibmeasure: Python/Cython module to find the center of back-illuminated optical fibers in metrology images,” Astrophysics Source Code Library, Mar. 2016.
- [48] S. Behnel *et al.*, “Cython: The best of both worlds,” *Computing in Science Engineering*, vol. 13, no. 2, pp. 31–39, March 2011.
- [49] J. Gilbert, G. Dalton, and I. Lewis, “Verification of commercial motor performance for WEAVE at the William Herschel Telescope,” in *Multi-Object Spectroscopy in the Next Decade: Big Questions, Large Surveys and Wide Fields*, ser. Astronomical Society of the Pacific Conference Series, 2016 (in press).
- [50] A. J. McGrath *et al.*, “Ukidna: the RAVE machine,” in *Ground-based Instrumentation for Astronomy*, ser. Proc. SPIE, A. F. M. Moorwood and M. Iye, Eds., vol. 5492, Sep. 2004, pp. 353–363.

- [51] A. McGrath *et al.*, “2dF grows up: Echidna for the AAT,” in *Ground-based and Airborne Instrumentation for Astronomy II*, ser. Proc. SPIE, vol. 7014, Jul. 2008, p. 70144K.
- [52] M. Akiyama *et al.*, “Performance of Echidna fiber positioner for FMOS on Subaru,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 7018, Jul. 2008, p. 70182V.
- [53] C. J. Chen, “Electromechanical deflections of piezoelectric tubes with quartered electrodes,” *Applied Physics Letters*, vol. 60, pp. 132–134, Feb. 1992.
- [54] S.-E. Park and T. R. Shrout, “Ultrahigh strain and piezoelectric behavior in relaxor based ferroelectric single crystals,” *Journal of Applied Physics*, vol. 82, pp. 1804–1811, Aug. 1997.
- [55] G. van Rossum and F. Drake, “Python reference manual,” 2001. [Online]. Available: <http://www.python.org/>
- [56] J. Gilbert and G. Dalton, “Learning from history: Adaptive calibration of ‘tilting spine’ fiber positioners,” in *Astronomical Data Analysis Software and Systems XXV (ADASS XXV)*, ser. Astronomical Society of the Pacific Conference Series, 2016 (in press).
- [57] Gilbert, J. and Dalton, G., “Echidna Mark II: one giant leap for ‘tilting spine’ fibre positioning technology,” in *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation*, ser. Proc. SPIE, vol. 9912, 2016 (in press).
- [58] F. Perez and B. E. Granger, “Ipython: A system for interactive scientific computing,” *Computing in Science Engineering*, vol. 9, no. 3, pp. 21–29, May 2007.
- [59] S. van der Walt, S. C. Colbert, and G. Varoquaux, “The numpy array: A structure for efficient numerical computation,” *Computing in Science Engineering*, vol. 13, no. 2, pp. 22–30, March 2011.
- [60] E. Jones *et al.*, “SciPy: Open source scientific tools for Python,” 2001–, [Online; accessed 2016-03-18]. [Online]. Available: <http://www.scipy.org/>
- [61] J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Computing In Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007. [Online]. Available: <http://dx.doi.org/10.5281/zenodo.15423>