

Instrumentation for HARMONI and the mass ratios of ring galaxy systems



Jamie Allen
Pembroke College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas 2017

Abstract

HARMONI is the visible and near-infrared integral field spectrograph due to be one of the first light instruments on the European Extremely Large Telescope. The prototype HARMONI lens mounts were cryogenically tested, to measure the radial decentre of the lens as it was cooled to 130 K. Initial problems of the epoxy cracking the lens were solved by switching to a silicone RTV adhesive. The lens in the baseline mount design was found to move radially by $18\text{ }\mu\text{m}$: just below the tightest radial positioning tolerance.

The Oxford SWIFT Spectrograph spare camera was rebuilt using lens mounts based on the proposed HARMONI designs. A procedure to align spherical lenses to their mounts was developed that could position the lenses to within $10\text{ }\mu\text{m}$ both radially and axially. However, the performance of the camera did not meet expectations, as the lenses were moving in their mounts. This was attributed to either the adhesive or the cotton paper buffer between the lens and its mount; also, at least one bond had failed. These issues are thought to be solvable by using an automatic, rather than manual, adhesive dispenser, though this has yet to be proven.

A commercial off-the-shelf shutter was also tested to see if it would work in the HARMONI cryostat. Although it could be used at these temperatures, it could not be made fully reliable. Its lifetime of 9321 actuations was also much lower than the estimated 88000 actuations for the lifetime of HARMONI.

Ring galaxies are annulus-shaped galaxies thought to be formed from a disc galaxy when a companion galaxy passes directly through it close to its centre at a nearly-normal angle. The kinematics of three ring galaxies were measured from MUSE observations and used to estimate their dynamical masses. A simple analytic model was used to refine the fits to the kinematics, though this did not change the measured masses very much.

Stellar population synthesis models were fit to the integrated spectrum of the companions of two of the ring galaxies to estimate their stellar masses via their mass-to-light ratios. The mass ratios of the companions to the rings were both found to have a lower bound of $1/6$, though possibly increasing to $2/5$ when dark matter is taken into account. This value is comparable to other observational results of ring galaxy mass ratios, and is within the range of values used for simulations of ring galaxy formation.

Declaration

I declare that no part of this thesis has been accepted, or is currently being submitted, for any degree or diploma or certificate or any other qualification in this University or elsewhere. Except where explicit reference is made to the work of others, the work contained in this thesis is my own.

Chapter 2 is based on part of the work published in the Proceedings of SPIE: *Cryogenic testing of components for the HARMONI spectrograph*, Allen et al. (2014). Minor edits reflect the comments of co-authors.

Chapter 3 is based on work published in the Proceedings of SPIE: *Cryogenic testing of components for the HARMONI spectrograph*, Allen et al. (2014) and *Characterizing the performance of cryogenic lens mounts for the HARMONI spectrograph*, Allen et al. (2016). Mechanical design of the HARMONI cryogenic lens mounts was done by Jim Lynn. Additional minor edits reflect the comments of co-authors.

In Chapter 4 the mechanical design of the new mounting scheme and the alignment jig was performed by Jim Lynn. The development of the alignment process was a collaboration between myself and Jim. Further measurements of the lens positions were performed by Dr. Rob Barnsley.

Jamie Allen

(October 2017)

Acknowledgements

This thesis would not have been possible without the help of many people, too many to list here, and probably more than I can remember. So firstly, thank you to everyone I haven't mentioned here, your help was still much appreciated.

Many thanks to my supervisors, Niranjana and Kieran, for all of your hard work and guidance over the years. Thank you for helping coax me out of my shell and many apologies for retreating back there during the last few months of writing up.

Thank you also to the other members of the HARMONI group: to Jim for all of the friendly help, especially with the CMM; to Matthias and Fraser for answering all of my questions about SWIFT, and many insights during meetings; to Ryan for all the galactic help that I really needed; to Sarah for helping me get started with pPXF; to Rob for the extra measurements of the SWIFT camera; and especially to Vanessa for organising everything. Thank you also to Ian Tosh and Ian Lewis for the information about the KMOS lens mounts.

Thank you to Ashling for always being so knowledgeable and helpful.

Thank you to Jonathan Patterson for fixing all my Linux problems so quickly.

Thank you to the Oxford Astro grad students, especially Holly, Simon, Rupert, Ben, Marisa, Rob and Kimon for being fantastic office-mates.

Thank you to my teachers: Mr Edney, Mr Carter and Mr Brooks for helping to inspire me to study physics at university.

A massive thank you to my family, for an interesting and educational upbringing that made me who I am.

Finally, and perhaps most of all, thank you to Aimée. Thank you for the last decade, but especially for putting up with me over the last couple of months and for being there when I needed you.

Contents

I	Instrumentation	1
1	Introduction to instrumentation for the HARMONI spectrograph	3
1.1	Integral field spectrography	3
1.1.1	Methods	4
1.2	HARMONI	6
1.2.1	Design	7
1.2.2	Cryogenic environment	8
1.3	Interferometric testing of optics	9
1.3.1	Optical quality	11
1.4	Scope	14
2	Cryogenic shutter tests	15
2.1	Shutter	15
2.1.1	Requirements	18
2.2	Shutter controller	20
2.2.1	Programmable power supply	20
2.2.2	Voltage regulator	22
2.2.3	MOSFET	23
2.2.4	Solenoid failure	24
2.3	Methodology	25
2.4	Discussion	26
2.5	Conclusions	30
3	Cryogenic performance of the HARMONI lens mounts	31
3.1	HARMONI NIR cameras	32
3.1.1	Optical design	32
3.1.2	Mechanical design	33
3.2	Method	39
3.2.1	Mounting the blank	40
3.2.2	Applying the adhesive	41
3.2.3	Measurements	42
3.2.4	Analysis	44
3.3	Choice of adhesive	45
3.4	Design 1 - Double-beam spring	50
3.4.1	Version 1 - Curved pads	50

3.4.2	Version 2 - Flat pads	51
3.5	Design 2 - Z-spring	56
3.6	Cotton paper	57
3.7	Conclusions	59
4	Rebuilding the Oxford SWIFT Spectrograph spare camera	61
4.1	SWIFT camera	62
4.1.1	Optical design	62
4.1.2	Mechanical design	64
4.2	Assembling the camera	68
4.2.1	Mounting lenses	68
4.2.2	Final assembly	73
4.2.3	Discussion	74
4.3	Interferometric testing	75
4.3.1	Lens surface form	76
4.3.2	Lens radius of curvature	83
4.3.3	Whole camera	85
4.4	Problems	99
4.4.1	CMM measurements of lens alignment	99
4.4.2	Method	100
4.4.3	Discussion	102
4.4.4	Further work	107
4.5	Conclusions	113
II	Ring Galaxies	117
5	Introduction to collisional ring galaxy systems	119
5.1	Mass ratios of ring galaxy companions	123
5.1.1	Observational results	123
5.1.2	Simulations	125
5.2	Ring galaxy kinematic modelling	127
5.2.1	Kinematic model	127
5.2.2	Off-centre collisions	128
5.2.3	Mass outside the ring	132
5.3	Stellar population masses	133
5.3.1	Single stellar population models	134
5.4	Scope	135
6	Determining the mass of ring galaxies	137
6.1	Observations	137
6.2	Kinematic maps	140
6.3	Initial kinematic model	142
6.3.1	Estimating inclination	142
6.3.2	Fitting kinematics	144

6.4	Analytic simulations	147
6.4.1	Method	149
6.4.2	Results and discussion	151
6.4.3	Conclusions and improved fit	166
6.5	Improved kinematic model	167
6.5.1	Results and discussion	168
6.6	Mass estimates	174
6.6.1	Discussion	174
6.7	Future work	176
7	Masses of ring galaxy companions	178
7.1	Method	178
7.1.1	Spectrum fitting	178
7.1.2	Estimating mass to light ratios	183
7.1.3	Mass determination	184
7.2	Discussion	185
7.2.1	Spectrum fitting	185
7.2.2	Mass-to-light ratios	190
7.2.3	Companion masses	193
7.3	Ring galaxy system mass ratios	198
7.4	Conclusions	200
	Appendix A Ring galaxy simulations	202
	Bibliography	209

Part I

Instrumentation

Chapter 1

Introduction to instrumentation for the HARMONI spectrograph

1.1 Integral field spectrography

Spectroscopy has proven to be one of the most useful tools available in astronomy. As the observed wavelength of light depends on many properties, such as the source's chemical make up, temperature and relative velocity, it is very valuable to be able to measure its wavelength dependence as this opens up many new avenues of study. It can therefore allow direct measurements of these properties as well as indirect measurements of masses, distances and more.

Traditionally, spectra have been measured using long-slit spectroscopy. A major disadvantage of this method is that it yields measurements that are spatially one dimensional, as one of the two dimensions on the detector is used for the dispersion of the light. An additional disadvantage comes from the slit reducing the field of view, which increases the difficulty of accurately pointing the telescope.

For these reasons it is clear that spectral data containing two spatial dimensions is usually preferable to long-slit spectroscopy. At first this was only possible by taking

multiple exposures. These could be many long-slit spectra with different pointings to scan along the intended field of view or many narrow-band full-field images taken at different wavelengths using an interferometer.

An integral field spectrograph (IFS) allows these type of measurements to be taken in a single exposure, generating a so-called datacube. This is a three dimensional array of data containing fluxes as a function of two spatial dimensions and one spectral. Another way of thinking about the data is that it represents an image, but each element in the image has a measured spectrum. This image is said to be made up of spatial pixels, or spaxels, to differentiate them from the pixels that make up the detector.

The major advantage of using an IFS over scanning with a slit or an interferometer is that the measurement is made in a single exposure. This is especially important for time-critical measurements where a phenomenon is too short for multiple exposures, or if its behaviour is changing between exposures. However, the drawback is that either larger detectors and optics are needed to relay and record the extra information, or else the size of the field of view is compromised. This is particularly relevant to near infrared (NIR) astronomy due to the increased cost of detectors when compared to optical CCDs.

1.1.1 Methods

Integral field spectroscopy cannot be done by simply passing a 2D image through a disperser, as the field must first be reformatted. There are three main ways to achieve this: a lenslet array, optical fibres or an image slicer. Each of these methods, depicted in figure 1.1, then rely on software to construct a datacube from the raw images.

Lenslet arrays are the simplest way to reformat the field and were used in spectrographs such as the pioneering TIGER spectrograph (Bacon et al., 1995) and SAURON (Bacon et al., 2001). Here the lenslets are placed in the focal plane to create an array




Figure 1.1: Schematics of the three main methods used in integral field spectroscopy. Image credit: jwst-docs.stsci.edu.

of pupil images. These pupil images can be sufficiently spaced out to disperse them without any overlaps. A simple, high-throughput IFS can be achieved with these designs, but a relatively large fraction of the detector is unused and often long spectra are unobtainable as they would overlap.

The other two techniques work in similar ways, both using optics to manipulate the shape of the field to create a ‘pseudoslit’. This in turn is then fed into a spectrograph to be dispersed. Use of a slit means that a higher fraction of the detector can be used, which has the potential for longer spectra.

Optical fibres can be used by placing them in the focal plane and taking advantage of their flexibility to physically form this pseudoslit. Lenslet arrays can be used increase the efficiency of coupling the light with the bundle of fibres. Optical fibres have been used in instruments such as GMOS-IFU (Davies et al., 1997).

Image slicers are also used to form a pseudoslit. This is the method that is to be used in HARMONI, and has been used in the past for instruments such as the Oxford SWIFT Spectrograph (Niranjan Thatte, 2006) and MUSE (Bacon et al., 2010). The




Figure 1.2: An image slicer from the MUSE spectrograph showing how the slicer is formed of a stack of very thin mirrors. Image credit: Florence Laurent (2014).

slicer itself is a stack of thin mirrors placed in the focal plane, with an example from MUSE seen in figure 1.2. Each mirror is tilted at a different angle so that the field is spread out, with further optical elements used to align the image into a neat slit.

1.2 HARMONI

HARMONI is a future optical and NIR IFS (Thatte et al., 2016), and will be one of the first-light instruments on the 39 m diameter European Extremely Large Telescope (E-ELT) (Tamai et al., 2016). It is conceived to be a general purpose spectrograph for use on many of the E-ELT’s main science cases. Using Laser Tomographic or Single Conjugate Adaptive Optics (LTAO and SCAO), HARMONI will achieve close-to-diffraction scales, where the exposure time scales with the diameter as D^4 , and will offer unequalled sensitivity and spatial resolution.

Its monolithic field of view will consist of ~ 31300 spaxels, arranged in a $\sqrt{2}:1$ aspect ratio. The wavelength coverage will be $0.47\text{-}2.45\mu\text{m}$ at 3 different spectral resolving powers of ~ 3500 , 7000 and 18000 , although below $0.8\mu\text{m}$ only 3500 will be available. There will also be 4 different spatial scales, as summarised in table 1.1.

Field of view (")	Spaxel size (")
6.18×9.12	60×30
3.04×4.12	20×20
1.52×2.06	10×10
0.61×0.82	4×4

Table 1.1: The fields of view and spaxel sizes of HARMONI’s 4 spatial scales.

1.2.1 Design

The majority of the HARMONI optics are to be located in an upwards-looking cryostat that is placed on a field de-rotator on the Nasmyth platform of the E-ELT. Above this cryostat there will be cooled optics designed to relay the telescope focus into the cryostat as well as the calibration unit. On top of the cryostat there will be room for the sensors used for SCAO, the LTAO natural guide stars and the secondary guiding stars.

The inside of the HARMONI cryostat can be seen in figure 1.3. The light enters through a window at the top of the cryostat, with the optical train split into three main sub-systems: the pre-optics, the IFU, and the spectrographs.

The pre-optics take the relayed telescope focus and prepare the beam for the integral field unit, including allowing the selection of the spatial scale. There are also spaces for a fast shutter, cold-stop and other filters and masks.

The integral field unit creates the pseudoslit. This is achieved first by dividing the field into four quadrants using a field splitter. Then each of these quadrants are passed to their own image slicer, with the slices rearranged into a line, thus creating four pseudoslits.

There are four spectrographs, one per slit, which first collimate the beam before dispersing it using one of the volume phase holographic diffraction gratings that can be selected via a grating wheel. Finally, a fully-refractive camera is used to focus the spectrum onto the mercury cadmium telluride detector. This is the optical path for any NIR observations, and the spectrograph is where the path of optical observations




Figure 1.3: The inside of the HARMONI cryostat, which has a height of approximately 4 m.

is split off using a dichroic. The optical path will have its own fixed diffraction grating, refractive camera and CCD detector, but it will only be present in one or two of the spectrographs, resulting in a reduced field of view compared to NIR observations.

1.2.2 Cryogenic environment

As HARMONI is to take NIR observations, including the K-band, it will be recording spectral information from wavelengths with significant thermal radiation from sources close to or at room temperature. To minimise the background radiation that would otherwise compromise these observations, HARMONI will therefore be a cryogenic instrument, with the majority of the instrument's optics being held within a large cryostat at a temperature of ~ 140 K.

The conditions inside the cryostat, although necessary to meet the thermal background requirements, make the mechanical engineering of the instrument more diffi-

cult.

For example, the vacuum of the cryostat required to maintain the low temperatures introduces the need for all substances inside to have low outgassing properties. This was important in this work when selecting adhesives, since standard ones do not meet these requirements, although many have been developed that do.

An even bigger effect on the engineering is due to the temperatures themselves, with perhaps the largest of these due to the thermal contraction of materials as they are cooled. The problem is that every material has a different rate of contraction as described by its coefficient of thermal expansion (CTE). Therefore, as HARMONI cannot be manufactured from a single material, there will be different rates of contraction for its different parts. This can cause problems with any moving parts, such as joints becoming stiffer or looser, but these can be overcome with the correct engineering. Another problem is due to the precise alignment required in optics, with the thermal contraction much larger than the tolerances of many of the optical surface positions.

There are also other problems caused by low temperatures, such as many conventional lubricants no longer functioning properly and cold welding (Merstallinger et al., 2009). These can be overcome by carefully designing products with knowledge of their working environments, but can mean that if a product is designed to be operated at room temperatures it should be carefully tested before being used in cryogenic conditions and may require modifications.

1.3 Interferometric testing of optics

Interferometry is a powerful method that can be used to measure the form of a wavefront. It is therefore useful in measuring the performance of optical systems and can be used both to measure the form of an optical surface or to measure the optical

quality of a whole system.

The basic principle of interferometry is to use a beam splitter to allow light to follow two different paths simultaneously before they are recombined to measure their interference pattern. This pattern encodes information about the difference in optical path lengths of the two beams; if the beams have the same path for all but one optical surface then you can find the shape of one of those differing surfaces compared to the other. Therefore if one of these surfaces is manufactured to a very high optical standard it can be used as a reference surface in order to measure the surface form of the other. One way of performing these measurement is by using a Fizeau interferometer, shown in figure 1.4 in a setup that allows the measurement of the form of a flat test mirror.

A single interferogram can be used to measure the magnitude of path differences, but not the sign. For instance, if a surface is measured to be rotated about the x -axis, the interferogram could not be used to determine if the surface has been tipped up or down. Therefore additional information is needed to determine the sign of the path difference, which can be gathered in different ways, depending on whether a temporal or dynamic mode of the interferometer is being used.

A temporal mode requires multiple exposures. The reference surface is mounted on piezoelectric motors that can then be moved by a known amount between the exposures. If at least 3 exposures are taken then this gives the extra information needed to deduce the sign of the phase difference, as the amplitudes of the two beams are also unknowns. One of the biggest disadvantages of the temporal method is the extra time required for multiple exposures increases the amount of degradation of the measurement by any vibrations that are present.

To reduce the effects of vibration on the measurement it can be taken in a single exposure using a dynamic mode. One way of performing this is by introducing a relatively large known tilt to the reference flat. The fringe pattern from this tilt is




Figure 1.4: The layout of a Fizeau interferometer. The reference flat acts as both a beam splitter and reference surface, with its reflection combined with that from the test flat to produce the interferogram. In this work the reference flat is also replaced by a reference sphere that produces a convergent beam, as opposed to a collimated one. This was used to measure the form of spherical lenses. Image credit: wikipedia.com

then effectively modulated by the measurement being made. However, as the tilt is both much larger and known, it allows the sign information to be determined. The downside to most single frame interferometry methods is either a slight degradation to the accuracy or reduced resolution, which in high vibration environments is often a price worth paying.

1.3.1 Optical quality

Once the wavefront of an optical system has been interferometrically measured, there are a number of ways it can be analysed in order to quantify the optical quality.

These are all done by finding the wavefront aberrations — the difference between the measured and perfect wavefronts.

Perhaps the simplest way is by subtracting the most negative deviation from the most positive deviation. This is known as the peak-to-valley (PV) wavefront error and whilst being very quick to calculate, it is very susceptible to noise. A more robust peak-to-valley (PVr) (Evans, 2009) relies on fitting the first 36 Zernike polynomials, described below, to the wavefront aberrations. It is found by taking the PV of the fit and then adding to it 3 times the root mean square (rms) of the residuals. The PV of a wavefront is a very simplistic analysis of the aberrations as it only looks at the very worst parts of the wavefront. To get a better overall feel of the entire wavefront then the rms of the wavefront aberrations can be found.

The optical quality does not only depend on the magnitude of the aberrations, but also their shapes. One way to go about measuring this is by using a Zernike decomposition. The Zernike polynomials are a set of polynomials that are orthogonal on a unit circle, with the first 10 shown in figure 1.5. Therefore any circular wavefront can be decomposed into Zernike polynomials, so they are often used in optics, where circular apertures are common.

The 0th order piston term is usually unimportant in optical testing as it is uniform across the pupil and only proportional to the distance to the test optic. For this work any tilt and the defocus are also not considered as these can be cancelled out by adjusting the positions of the optics used for the tests. The more important terms therefore are the astigmatism, coma and higher orders. The easiest way to use these aberrations is to compare the measured aberrations with those predicted using optical design software, for example, Zemax. This can then be used to determine the optical aberrations that are produced by certain misalignments, meaning that the measured wavefronts can be used to predict the positions of the optical surfaces to aid in alignment.




Figure 1.5: The first 10 Zernike polynomials. Image credit: comsol.com

The effect that any aberrations have on an image can be determined by finding the point spread function (PSF). This defines how a point source would be seen as imaged using this optical system and can therefore be convolved with an object to find out the appearance of the image. The PSF can be found by taking the inverse Fourier transform, $\mathcal{F}^{-1}$, of the pupil function, P , and then finding the square of its magnitude:

$$PSF = |\mathcal{F}^{-1}(P)|^2. \quad (1.1)$$

The pupil function itself has the form:

$$P = A \cdot \exp(i\Theta), \quad (1.2)$$

where A is its amplitude and Θ is its phase. This phase can be found from the interferometrically measured wavefront, wf , via the equation:

$$\Theta = \frac{wf \cdot 2\pi}{\lambda}, \quad (1.3)$$

where λ is the wavelength of the interferometer.

1.4 Scope

In Chapter 2, an off-the-shelf shutter is tested in a cryogenic environment to assess its suitability for use in HARMONI.

In Chapter 3, prototype cryogenic lens mounts for HARMONI are tested to measure how well they would maintain the alignment of a lens when cooled down. Three different mount designs are tested and their pros and cons discussed. Work is also presented on the procedure to mount the lenses, including the effects of the adhesives as well as the shape of the pads that are bonded to the lens edges.

In Chapter 4 a method for aligning spherical lenses is developed and tested on the spare spectrograph camera from the Oxford SWIFT Spectrograph. The optical performance of the camera was thought to be lower than expected from the optical design, so it was measured before and after remounting the lenses to see if the new scheme offered any improvement. Unfortunately, problems were discovered in the newly mounted lenses whereby they were not being securely held, a problem seemingly caused at least in part by a failure with the adhesive. The positives of the alignment technique are discussed along with suggestions to improve the procedure to insert the adhesive.

Chapter 2

Cryogenic shutter tests

As HARMONI will use CCDs as the detectors for the visible-wavelength spectroscopy, it will require a fast shutter to accurately control the length of exposure times, as the CCD readout times will be tens of seconds. The shutter will also be used to protect the NIR detectors from bright sources during slewing, because their exposure to such sources can cause persistence. There are different places along the optical path that a shutter can be placed, namely very close pupil and focal planes, which can drive its size requirements depending on the beam width and the amount of space available. Other considerations include whether to place the shutter up or downstream of optical components such as the field splitter, which could mean multiple shutters are needed.

2.1 Shutter

As it is simpler to buy off-the-shelf components rather than using bespoke or modified components, a commercially available shutter was procured to test its performance at the HARMONI operating temperature. This shutter was a Melles Griot 04 UTS 268, which is a 63.5 mm clear diameter, normally-open shutter from their UltraThin range. As shown in figure 2.1, this is a six-bladed iris shutter that uses a 12 V rated solenoid to close it, with a spring used to open it again. The shutter was not designed for




Figure 2.1: The shutter that was tested was a 63.5 mm clear diameter, normally-open Melles Griot UltraThin shutter. This is an off-the-shelf iris shutter that was not designed for cryogenic operation.

cryogenic operation, with a minimum specified temperature of -40°C .

As this shutter will be placed in the HARMONI cryostat, it must therefore be proven to function reliably in cryogenic conditions for the lifetime of the instrument. Therefore the chief aim of these tests was to establish whether the shutter still worked at 130 K, if so then how quickly and reliably, and if not then what was causing it to fail and the severity and frequency of any problems.

As described in Chapter 1, cryogenic conditions can provide a challenging environment in which to operate mechanisms due to a number of issues. At these temperatures conventional lubricants no longer function properly due to their increased viscosities, which can cause difficulties for moving parts and often requires redesigning a part for it to work cryogenically. Another problem is due to different materials having different coefficients of thermal expansion, which can cause parts to move relative to each other as they cool and can cause mechanisms to jam and cease to work.

When HARMONI is operational, any problem with the shutter will require different steps to fix it depending on how serious the problem is, with different severity levels defined in table 2.1. The more serious problems will take longer to fix and therefore should happen less often, to prevent loss of observing time. When calculating how much observation time is lost, one of the biggest differences between levels is from level 2 to 3 as access to the shutter itself is required, and this will necessitate warming up the cryostat. Given the size and mass of the instrument, the time it will take to thermally cycle it is very long, upwards of seven days, and therefore it is a very undesirable procedure. A level 4 failure should only affect a single exposure and, though the length of time it should take to repair a level 3 failure will vary, it should be possible to determine whether it is necessary to warm up the cryostat very quickly after trying the different settings and resetting the system. A level 1 failure could mean a significant extra amount of observing time is lost over a level 2 failure if a new shutter needs to be purchased. This could be of the order of weeks when the lead time for the shutter is added to the extra time needed to dis- and reassemble the instrument.

The consequences of any shutter failure would depend on the severity of the failure. As this shutter would be placed at a pupil plane in the pre-optics, it would be in the common path for all types of observations. Therefore, if it completely failed it would prevent any visible observations and possibly all NIR ones as well. The worst case is if the shutter is stuck closed, as in the current design of the instrument this would prevent all observations from taking place, as the shutter will be in the beam path for all observations regardless of wavelength. At the other end of the scale, if the shutter is stuck open then whilst no visible observations are possible, infrared measurements will be, provided that the exposure time is long enough that the shutter is not needed to control it. If the shutter is stuck at an intermediate position then infrared measurements may still be possible if the aperture is large

Level	Description	MTBF (actuations)
1	Dead and needs replacing	25920
2	Stuck and requires fixing manually	25920
3	Stuck, but can be repaired remotely	17280
4	Isolated failure, no need to fix it	8640

Table 2.1: The different levels of failure for the fast shutter in HARMONI and their minimum required mean time between failures (MTBF). Levels 1-3 all result in the loss of further science observations if they are not fixed, with levels 1 and 2 requiring the cryostat to be warmed up to fix.

Severity	Thermal Cycle required?	
	Yes	No
Science prevented	25920 hr	17280 hr
Science degraded	19440 hr	12960 hr
Other failure	12960 hr	8640 hr

Table 2.2: The MTBF for different severities of failure. The severity depends on whether a failure to fix the problem results in a loss of science observations, a degradation of science observations or neither. The value depends on whether the HARMONI cryostat needs to be warmed up to fix the issue.

enough, though throughput could be significantly compromised. Therefore, unless visible wavelengths or very short exposure times are required, it may be possible to continue with observations to finish the observing programme before repairs need to be undertaken.

2.1.1 Requirements

The first requirement on the shutter, especially as it is not designed for cryogenic temperatures, is that it must function in the environment of the HARMONI cryostat and should do so for the lifetime of the instrument. To estimate the number of actuations the shutter will need to perform in total a number of assumptions need to be made. These include how many nights HARMONI will be used for as well as the number of observations and calibrations made.

The common requirements for E-ELT instruments are that the lifetime will be

a minimum of 10 years and that a first-light instrument will be assumed to be in observation mode for an average of 10 hours for 80 nights per year. This therefore leads to a value of 8000 hrs total. In this work it is assumed that the instrument will be exposed on average every 10 min during observations and approximately 50 calibration frames will be required for each observing night. This works out to an estimate of 88000 exposures required of the shutter. This number will represent a conservative estimate if the majority of observations use only the NIR detectors and on few objects, therefore reducing the amount the shutter needs to be used as protection during slewing. However, any overestimate is kept as a safety margin to increase the confidence in the shutter lasting the lifetime of the instrument.

The instrument common requirements also list requirements on the mean time between failures (MTBF) for the instrument. The required MTBF depends on the severity of the failure, that is, whether it prevents or degrades science observations, and whether the instrument needs to be warmed up or not. The values are shown in table 2.2, with the units of hours roughly equivalent to actuations due to there being 87600 hours in 10 years. Although these values are for the instrument as a whole, they were taken as underestimated values for the shutter initially, with the aim of providing updated numbers at a later date. However, during the tests it became clear that this further work was not required.

The values in table 2.2 can be compared to the levels of failure in table 2.1 to derive further requirements. As the shutter is in the common path for all observations it was assumed that any prolonged failure would lead to a loss and not a degradation of science observations and that level 4 failures would fall into the least severe category. Therefore the MTBFs vary between 8640 and 25920 actuations, as seen in table 2.1.

There are also HARMONI specific requirements that depend on the performance of the shutter. For example, the instrument shall support individual on-chip exposure times between 0.1 s and 900 s, with a goal of 3600 s. Also, the on-chip exposure time

Description	Requirement
Minimum temperature	<130 K
Lifetime	>88000 actuations
MTBF	>25920 actuations
Minimum exposure	0.1 s
Exposure repeatability	0.001 s

Table 2.3: A summary of the requirements on the HARMONI fast shutter.

shall be known to 0.1 seconds, or 1% of the exposure time, whichever is smaller.

All of the requirements on the shutter are summarised in table 2.3

2.2 Shutter controller

The shutter was designed to be closed by applying a 20 ms, 48 V electrical pulse across its terminals and then kept closed with a 6 V holding voltage. The shutter is then opened by a spring when the voltage across it is removed.

The electrical properties of the shutter will change as it is cooled down inside the cryostat, for example the resistances will decrease, meaning that a lower voltage will be needed to close the shutter for the same current levels. For this reason the power supply for the controller was used in a current limited mode to prevent any damage to the shutter caused by using too large a voltage when the shutter is cold.

2.2.1 Programmable power supply

The first controller that was tested was a piece of LabVIEW code that was written to control a programmable power supply to produce the voltage peak to open the shutter and then lower the voltage to hold it open. However, each output on our power supply was limited to 32 V and so two had to be used in series to reach the required peak voltage.

The code originally set the current limit for the voltage peak to close the shutter,




Figure 2.2: The voltage profiles used to close the shutter from the programmable power supply (*left*) and the MOSFET controlled circuit (*right*). The dashed lines on the left hand plot highlight the 40 ms minimum time between commands when programming the power supply. This was slower than the shutter closing time and the motivation for improving the shutter controller.

2 A, before setting the first output to 32 V and then the second output to 16 V. The current limit was then immediately lowered to 0.5 A before turning down the outputs for the holding current, and the the voltages were set to 0 V to reopen the shutter.

In preliminary tests this power supply was programmed to produce the pulse and holding voltages, but as the two channels could not be turned on simultaneously, this meant the pulse profile was stepped, rather than increasing straight up to 48 V. This was a problem because the time delay between commands being performed by the power supply was approximately 40 ms, as shown in figure 2.2, which is a lot longer than both the recommended pulse length of 20 ms and the closing time of the shutter of 0.01,s. The first step taken to reduce this problem was to try increasing the voltage before the pulse by as much as possible without closing the shutter. This did not solve the problem, so more work was performed to improve the shutter controller by building different circuits.




Figure 2.3: The circuit using the voltage regulator that can be used to produce two different output voltages. The values of R1, R2 and R3 can be changed to set the output voltages to their desired values. The transistor Q1 is used to switch between the two outputs and is controlled using an Arduino which can be programed to produce a series of exposures with the shutter.

2.2.2 Voltage regulator

The first circuit that was built to replace the programmable power supply to produce the required voltage profile was implemented using a LM317HV voltage regulator. A voltage regulator is designed to maintain a constant voltage level and the type used in this circuit was an adjustable 3 terminal device. The LM317HV's pins are input, the output and the adjuster pin. The output voltage from the regulator can be adjusted by changing the resistance between the output and the adjustment pins and between the adjustment pin and ground, and is given by:

$$V_{out} = 1.25V \left(1 + \frac{R2}{R1} \right) + I_{adj} R2, \quad (2.1)$$

for the circuit in figure 2.3 when the switch for R3 is open. I_{adj} is the adjustment

pin current and takes the value of $50\text{ }\mu\text{A}$ for the LM317HV. This circuit can be used to achieve the desired voltage pulse profile to close and hold the shutter using two different voltages by switching R3 in and out of the circuit using a transistor. The transistor was controlled using a simple program run on an Arduino Uno, where the timing of each step can be easily adjusted to change the exposure time and frequency.

As can be seen in equation 2.1, the lowest voltage that can be achieved with this regulator is 1.25 V . This means that a transistor was used at the input pin to control whether the circuit was on or off and therefore whether the shutter was closed or open. The LM317HV is only specified to 1.5 A , so two were needed in parallel to guarantee that the circuit worked with the 2 A pulse for the shutter.

In the end, this circuit proved too difficult to implement reliably, with the result that the regulators kept burning out. The fact that two regulators were required, even when using the high voltage versions of the regulators, was judged to be the reason for the problems. No formal tests were performed using this controller and work was turned to developing a new controller.

2.2.3 MOSFET

The second shutter controller circuit that was tested was built around metal-oxide-semiconductor field-effect transistors (MOSFETs), which are transistors where the gate is insulated from the rest of the circuit. MOSFETs are commonly used as electronic switches and in this case were used to switch resistors in and out of series with the shutter to change the voltage across it, as shown in figure 2.4. The MOSFETs were again controlled using an Arduino Uno, with the voltage across the shutter equal to 0 V when both MOSFETs were open.

Unfortunately, a lot of power needs to be dissipated when the shutter is being held closed therefore high current resistors were required, and several were used in parallel to reduce the risk of them burning out.




Figure 2.4: The circuit used to control the shutter. The MOSFET Q2 is used to switch on the 48V pulse current across the shutter and needs to be switched off after turning on Q1 to reduce the voltage to the 6 V holding current. These MOSFETs are switched using a 5 V logic output from an Arduino Uno.

A flyback diode, D1 in figure 2.4, was used to prevent any damage to the electronics from back electromotive force caused by the shutter's solenoid.

2.2.4 Solenoid failure

During the building and testing of the new circuits there was a problem with the circuit: the voltage across the shutter was held at 48 V for several seconds rather than the recommended 0.2 ms. This caused the solenoid to fail and seize up, and as a result the solenoid needed replacing. The shutter worked again after the repairs, but when analysing the results it should be taken into account that the shutter body and the solenoid are of different ages and have undergone different workloads in their lifetime.




Figure 2.5: The setup of the equipment showing the shutter mounted on top of a ring as well as the specially made photodiode holder with the diodes (*circled*) placed at 0, 0.5, and 1 times the radius of the shutter's aperture.

2.3 Methodology

The shutter was mounted on a ring inside a cryostat to allow space underneath to place three photodiodes. These were used to detect light illuminating the inside of the cryostat through a window above the shutter. One photodiode was placed at the centre of the shutter's aperture, one at the edge and the other at a distance halfway between them, with the setup shown in figure 2.5. This allowed the change in light levels to be measured as the shutter blades passed over them and then be compared to each other to determine the opening and closing times. The photodiodes were held in place using a specially fabricated holder with a small hole in the top to restrict the viewing angle of the diode. This design was chosen for two reasons. Firstly, it helps to block out scattered light, and therefore reduces the noise. Secondly, it effectively decreased the field of view of the photodiode, increasing the resolution of the transition between light and dark as the shutter was actuated.

The shutter was thermally cycled at a range of different temperatures to test how its properties varied by opening and closing the shutter between 50 and 100 times, at temperatures between 110 K and 300 K. To ensure that the duty cycle was not too

high, and to prevent excessive heating of the solenoid, the shutter was kept closed for 1 s before being left open for 4 s to cool down before repeating. The voltage across each photodiode as well the shutter was recorded by a Labjack at its maximum rate of 300 Hz, which was fast enough to resolve the shutter opening and closing. The temperature of the shutter was measured using Pt100 temperature sensors affixed to the body of the shutter using aluminium tape.

To ease the analysis of the data, the measured photodiode voltages were normalised and their differentials found. The time at which the edge of a blade passes a particular photodiode was taken to be the time when the light levels changed at the fastest rate, which was found by measuring the position of the peak in the time differential of each section. As the blades all converge towards the centre though, the time for the photodiode positioned there is when the shutter has fully closed and therefore when the differential drops to zero. The difference between these times can then be used to determine the opening and closing times for each exposure and these were averaged to find the best estimate for the opening and closing times.

Even when the shutter failed to fully close it did so after closing but rebounding open slightly, so whether this happened during an exposure was determined by comparing the minimum light level from the first half of the exposure to the second half as they should be the same if the shutter did close fully. Therefore for each series of exposures the fraction of exposures which closed perfectly could be determined alongside the opening and closing times.

2.4 Discussion

Figure 2.6 shows the fraction of the time the shutter closed successfully as a function of temperature for both the methods of controlling the shutter. When programming the power supply the shutter works well at room temperature, where it is within




Figure 2.6: The fraction of time the shutter fully closed and then opened again as a function of temperature. The white points show the results from one series of exposures at a given temperature using the programmable power supply and with the MOSFET controller used for the blue points.

its specifications, and is fine down to 220 K. At temperatures lower than this, the shutter starts to regularly fail to close fully. However, at the lowest temperatures its performance improves again, with a success rate of 97% below 130 K.

At temperatures between 130 K and 180 K it can be seen that there is a wide range of results for a given temperature, with those with a higher success rate generally coming from later runs. This improvement was made by altering the voltage profile by increasing the voltage before applying the main pulse to close the shutter. This though was not enough to make the shutter 100% reliable at these temperatures, prompting the development of better means of controlling the shutter.

The new MOSFET-based shutter controller however made a large difference. This increased to reliability to 98.7% from 1100 individual exposures averaged across the whole temperature range, with a much more uniform performance. The increased reliability of the shutter is likely due to the ability to create the voltage pulse using just one command, rather than increasing it in steps. The Arduino also increased the

Controller	Actuations	Level 4		Level 3		Level 2	
		N	MTBF	N	MTBF	N	MTBF
Power supply	5930	1295	3.6	7	661	3	1542
MOSFET	1100	12	91	2	542	0	—

Table 2.4: The number of each level of failure for both shutter controllers and the mean number of successful actuations between failures (MTBF).

control of the shape of the pulse due to the much higher rate at which commands could be processed.

For both controllers, the overwhelming majority of failures were classed as level 4, as can be seen in table 2.4. In these it was found that the problem was that the shutter did not quite close fully, leaving a small gap in the centre, usually of the order 1 mm. This is much smaller than the aperture of 63.5 mm, and as the shutter would be placed in a pupil plane it should fall within the radius of the central obscuration of the telescope. This raises the intriguing idea that the shutter would actually function well enough in HARMONI. In reality however this behaviour represents an unacceptable risk to observations.

Surprisingly, the probability of a level 3 failure, where the shutter jammed for more than one actuation before freeing itself, seems to be comparable for both controllers. There were no level 2 failures that required the cryostat to be warmed up for the MOSFET controller, which was an improvement over the power supply controller. Of these 3 failures 2 only required warming up with the other had to be freed by hand. These are very small number statistics though, so the only conclusion that can really be drawn is that all types of failure occurred far too often for both controllers when compared to the requirements in table 2.2.

Figure 2.7 shows both the closing and opening times of the shutter as a function of temperature. Both exhibit a similar trend; with the times increasing with increasing temperature between 120 K and 200 K, though there is a much larger scatter for the closing times. This scatter is probably lower for the opening times because the only

things affecting them is the spring strength and friction of the shutter, whereas there are all of the additional effects of the shutter controller added in to the closing times.

It can be seen that the MOSFET shutter controller had a noticeable effect of the closing times of the shutter, with all of the results at the lower end of the range produced by programming the power supply. This is due to the better control over the pulse profile and the ability to increase the voltage much more quickly rather than be limited to an increase of 32 V. This is important because the 40 ms between the steps of the profile is much longer than the shutter takes to close. The new controller did not change the opening times, as was expected, though the slope does appear to be different from the previous tests. However, the slope of the graph was found to vary with each cooldown, giving rise to the observed scatter in opening times.

The opening and closing times can be used to estimate the shortest exposure possible using this shutter. For all the tests with the MOSFET controller between 130 K and 140 K the average sum of the opening and closing times is 0.055 s with a standard deviation of 0.003 s. Therefore the shutter would easily be able to meet the minimum exposure time requirement of 0.1 s, but the repeatability of the exposure time of 0.001 s is more of a problem. If the reliability of the shutter could be improved then it is likely that the times would be more predictable. However, as the shutter finally broke before this was achieved, no work was done on whether the exposure repeatability of 0.01 s could definitely be met.

The shutter failed completely just after the statistics from table 2.4 were compiled. The number of recorded actuations was 9321 in total, but only 1181 from after the replacement of the solenoid. The specific cause of the problem was not isolated, but it was decided that neither of these totals is close enough to the estimated lifetime requirement of 88000 to warrant repairs and further testing. It is concluded that the shutter would not last long enough or work reliably enough in the HARMONI cryostat without modifications.




Figure 2.7: The length of time the shutter takes to close (left) and open again (right) as a function of temperature. The preliminary results from programming the power supply are in white with the results from the MOSFET shutter controller in blue.

2.5 Conclusions

Although the shutter was found to retain unexpectedly good functionality down to 120 K, it was never made fully reliable below 220 K. A controller made using MOSFETs was found to increase the reliability compared to programming a power supply to produce the voltage pulse needed to close it, but ultimately it suffered a terminal failure before this work could be completed.

More positively, the shutter was fast enough to meet the requirements for the minimum exposure time, and if it could have been made more reliable then the repeatability may have also met the tightest requirements.

In the end however, it was concluded that not only was the reliability below what was necessary, but the lifetime appeared to be as well.

Chapter 3

Cryogenic performance of the HARMONI lens mounts

Due to the cryogenic environment within HARMONI, certain design challenges arise; this is particularly true for the spectrograph camera. As the instrument is cooled down, different components can contract at different rates because the coefficient of thermal expansion (CTE) of an object varies depending on its material.

For example, the current design uses aluminium for the lens mounts, which contracts much faster than the glass lenses themselves. This differential thermal contraction therefore causes each lens mount to exert an inward radial force on each lens, which if not controlled could lead to deformation or permanent damage to the glass of the lens. Also, the net force on the lens needs to be zero so as not to displace the lens and move the camera out of optical alignment. The camera is particularly affected by this due to the tight tolerances caused by it being relatively fast and wide-field, as well as the large difference in CTE between the aluminium and glass of the lenses and their mounts respectively.

The work presented in this chapter prototyped and tested three different lens mount designs using optical blanks. This was firstly to demonstrate that they work,

i.e. they hold the lens securely without damaging it when it is cooled down, and also to measure the performance of the mounts to ensure they are capable of meeting the requirements for the HARMONI cameras. Another aim of this work was also to simply gain experience in mounting lenses, with the eventual aim of developing a procedure which can be used for the final HARMONI designs and to ensure that ease of use could be factored into the final decision on which mount design to choose.

3.1 HARMONI NIR cameras

3.1.1 Optical design

Each of the four NIR cameras in HARMONI takes the 20° fan of collimated beams from a disperser and images it onto two 4096×4096 detectors. The camera has a focal length of 349 mm with a focal-ratio of $f/2.67$ along the slit. The target performance for the camera was to have an rms wavefront error of below 250 nm across the field for all wavelengths.




Figure 3.1: The optical design of the HARMONI infrared cameras.

The optical design work for the HARMONI infrared cameras has been carried out by David Freeman and Martin Lee, with the the current all-refractive design shown in figure 3.1 and with the optical prescription in table 3.1. It has 6 lenses that are of the order 200 mm in diameter, each with an edge thickness of at least 11 mm to allow

the lenses to be edge-mounted.

Lens	Left ROC (mm)	Right ROC (mm)	Centre thickness (mm)	Diameter (mm)	Preceding air-gap (mm)	Glass
1	484.582*	708.044	65	241.54	—	CAF2
2	10417.03	-390.835	22.998	248.12	5.422	S-TIH53IR
3	382.223	326.185*	65	257.74	5.555	CAF2
4	467.014	4081.51	45	230.86	224.841	ZNS
5	-332.627	246.494	64.638	193.96	44.057	F_SILICA
6	-239.044	-674.044	15	145.42	29.561	ZNS

Table 3.1: The optical prescription of the HARMONI infrared camera. The aspheric surfaces are marked with an asterix.

For this design the tolerances are quite tight. They vary from lens to lens, but the toughest values that will have to be satisfied are detailed in table 3.2. These tolerances are on the final position of the lens, meaning that they represent a budget that has to be split between the alignment and any movement induced by cooling the instrument down to the operating temperature of 130 K.

Parameter	Tolerance
Axial position	20 μm
Radial position	20 μm

Table 3.2: The most difficult tolerances to meet for alignment of the lenses in the HARMONI NIR camera.

3.1.2 Mechanical design

For the mechanical design of the lens mounts themselves, as well as ensuring the lenses meet the positioning tolerances specified in table 3.2, the lens mounts also have to meet requirements for how securely they hold the lenses in place. The most demanding conditions the mounted lenses will have to survive are during shipping from the UK to the observatory and in the event of them being hit by an earthquake.

Therefore they are also required to survive a 20 cm drop should this happen during handling of the components and also a 3.6 g acceleration in any direction.

The three mechanical designs that have been prototyped and tested for use in HARMONI were produced by Jim Lynn. Each of them works using the same principles, which have heritage from the lens mount designs in FMOS (Froud et al., 2006) and KMOS (Masters et al., 2010). This involves adhering the mount to the edge of the lens and using compressible flexures to reduce the inward force of the mount on the lens to minimise the chance of the lens suffering any deformation or damage.

The FMOS lens mounts were revised after testing revealed that the original design caused the lens to crack when cooled down. To prevent this damage from happening again, the changes made to the revised FMOS lens mounts were:

- minimising the length of the area of the mount in contact with the lens to minimise the differential contraction between the two materials
- changing the material that is adhered to the lens to one with the same CTE, again to minimise the differential contraction between the two materials
- adding a layer of cotton paper between the lens and the mount to act as a buffer to reduce the stresses on the lens, as developed by O'Brien & Atwood (2003).

As all 3 changes were made simultaneously, it was not known how much each one helped, but the overall effect was that revised design did not damage the lens. In the HARMONI designs, the lens is held in place by using an adhesive to bond six pads to its edge, whose material is chosen such that its CTE is as close as possible to that of the lens. These pads are attached to an aluminium mount ring by flexures which have no preload. That is, they do not exert any force on the lens at room temperature. If possible, it would be desirable to not have to use cotton paper as this complicates the mounting procedure by adding an additional stage when preparing

the lenses and increases the handling of the lens, which should be minimised to reduce the probability of scratches and dust accumulation on the optical surfaces.

The radial alignment of the lens, i.e. how well centred it is in its mount, is maintained by the symmetrical distribution of the flexures around the lens edge. These flexures are then toleranced such that they have a spring constant similar enough that there is no significant net radial force on the lens which could cause it to decentre. The lens is prevented from moving in the axial direction relative to the mount ring by the adhesive holding it in place.

Glass	CTE ($10^{-6} \text{ }^{\circ}\text{C}^{-1}$)	Pad material	CTE ($10^{-6} \text{ }^{\circ}\text{C}^{-1}$)
Fused silica	0.55	Invar	1.2
Zinc selenide	7.1	Titanium	8.6
Calcium fluoride	18.9	Beryllium copper	17.8

Table 3.3: The types of glass found in the HARMONI infrared spectrograph cameras and the corresponding CTE-matched pad material.

Table 3.3 shows the corresponding CTE-matched material that will be used to make the pads for each type of glass in the infrared camera, and the CTEs for each. The reason that the CTE of the pads is matched to the lens is so that there is as little relative contraction across the edge of the lens as possible, to try and minimise the axial stress placed on the lens. For each of the designs the pad has a hole though it to enable a needle to be inserted through a similar hole in the mount ring so that adhesive can be injected into the bond line between the pad and the lens. The hole in the pad itself is a tight fit to prevent adhesive being forced out around the sides of the needle, whereas the hole in the mount ring is a looser fit to make it easier to insert the needle.

3.1.2.1 Design 1 - Double-beam spring

Design 1 is the baseline design for the lens mounts and has double-beam shaped flexures that are integrated into the CTE-matched pad itself, as seen in figure 3.2.



Figure 3.2: Design 1 for the HARMONI spectrograph camera lens mounts using double beam springs for the flexures. *Left*: version 1 with the longer curved pads. *Right*: version 2, the current baseline design, with shorter flat pads to make inserting the adhesive easier.

These pads can then be placed into T-slots in the mount ring and then screwed in to hold them in place. To change the thickness of the bond line with this design, the pads can be unscrewed and shims placed in the joint between them and the mount ring to move them further away from the lens.

In the first version of this design, the bonding surface of the pad was curved to match the radius of the lens diameter, meaning the whole pad had a constant bond line thickness, and the pads were also relatively long and thin. With this style of pad it was difficult to inject the adhesive into the bond line, and it did not reach the ends of the pad as it instead started to leak out of the edges of the pad as can be seen in figure 3.3.

The shape of the pads was changed so that the bonding surface of the pads was straight leading to the thickness of the bond line increasing towards the ends of the pad. Added to this, the aspect ratio of the pads was altered to make the pads squarer, with both changes aimed at making it easier for the adhesive to get to the ends of the pads. This results of the tests using pads with these modifications are also seen in figure 3.3, where the adhesive did manage to get further along the pads, but still not as far as the ends.




Figure 3.3: The extent to which adhesive was able to be inserted for different pad styles. *Left*: pads curved to match the lens diameter. *Middle*: shortened, flat pads. *Right*: shortened, flat pads with grooves added to help the adhesive move into the corners of the pad.

After this, grooves were added to the bonding surface of the pads, with two different patterns tested. The first pattern was a single groove through the middle of the pad surface along its length, with the second having a groove along each diagonal to form a cross reach out to the corners. The results of a test using both of these patterns are shown in figure 3.3. It can be seen that the grooves were effective in allowing the adhesive to spread farther, with the single groove along the length of the pad helping the adhesive reach the end of the pad, though it still could not reach the corners. The more effective pattern, and the one that was adopted into the baseline design, was the cross-shaped pattern which was effective in letting the adhesive move towards the corners.

3.1.2.2 Design 2 - Z-spring

Design 2 is the same as Design 1 in all aspects apart from the flexure, which instead of being a double-beam spring is a Z-shaped spring as illustrated by figure 3.4. The




Figure 3.4: Design 2 for the HARMONI spectrograph camera lens mounts using z-springs for the flexures.

flexures have a slot machined into the diagonal bar, which is designed to make it buckle in the middle, as opposed to bending at the ends where it joins the pad and the T-slot, so that the lens does not rotate as it is cooled down. As a further measure to ensure this does not happen, when the lens is mounted, alternate flexures are flipped round due to the asymmetric nature of the flexures.

As with Design 1 the first version of this design had long curved pads though this design was dropped before a second version was made due to the fragility of the flexures.

3.1.2.3 Design 3 - Moustache flexure

Design 3 is closer to the original designs seen in FMOS and KMOS insofar as the moustache-shaped flexure form part of the mount ring instead of the pads, as can be seen in figure 3.5.

The pads are very different in this design, being circular in shape rather than rectangular, and rather than fitting into T-slots in the mount ring they have a thread protruding from their backs which is screwed into place. This thread is also the means




Figure 3.5: Design 3 for the HARMONI spectrograph camera lens mounts using moustache-shaped flexures.

by which the bond line thickness is adjusted, screwing it in or out until the gap is the desired size. This distance is fixed when the adhesive sets, as when this happens the pads can no longer be turned.

A full set of tests could not be performed with this design due to damage caused during the first cool down by the adhesive. This is described further in section 3.3

3.2 Method

The first tests, using Design 3, were performed by bonding a 200 mm diameter glass flat into a prototype lens mount. Further tests with Designs 1 and 2 were performed using glass blanks to reduce costs after the flat sustained damage in the initial tests. The mounted blanks were then cooled down to a temperature of 120 K and then warmed back up to room temperature, with frequent measurements taken of the gap between the lens mount and the blank at three different points around the circumference for a range of temperatures. The mount's functionality was determined by a visual inspection after being warmed up and its performance measured by calculating

how far off-centre the blank moved using the measurements taken of the gap between the mount and the blank as a function of temperature.

3.2.1 Mounting the blank

The blanks were aligned using a co-ordinate measuring machine (CMM) for the metrology, either Oxford University Physics Department's Workshop's flatbed CMM, which is capable of making individual point measurements with an accuracy of below 5 μm , or a portable Faro Edge arm, which has an accuracy of below 30 μm . Both of these are shown in figure 3.6. The Faro arm is quicker and easier to operate and so was used after it was acquired as the extra precision of the flatbed CMM was not required for these tests.

For the tests in which cotton paper was used, this was applied to the edge of the blank before mounting it. This was done first by cutting a piece of cotton paper to the size of the edge of the blank. Next the adhesive was applied in a 1.5 mm bead down the length of the centre of the strip of paper before it was pressed against the edge of the blank. The paper is then strapped in place and left for a week for the adhesive to fully cure before the blank was mounted.

To set up the equipment for aligning the blank in the mount, first the mount ring was placed on 10 mm parallels and screwed in place, ready to be measured by the CMM. The parallels were used so that the mount was axially aligned with the middle of the blank whilst ensuring that the blank is not tilted with respect to the mount ring. When using Design 3 the pads were inserted into the mount ring at this point and fully screwed back, for Designs 1 and 2 this is done after alignment.

The blank was then placed in the middle of the mount before the mount ring was measured to initialise the co-ordinate system of the CMM with the origin at its centre. The top of the mount ring is probed to set the $x - y$ plane, with the edge then probed to set the $x - y$ co-ordinates of the origin. Finally, the angular position of the




Figure 3.6: The CMMs that were used to align the lenses within their mounts. *Left:* the Physics Workshop’s CMM shown here being used to align a 200 mm glass blank ready for cryogenic testing. *Right:* the Faro Edge arm CMM that was also used, image credit: Faro.com.

x -axis is set by measuring one of the holes on the mount ring. With the co-ordinate system set, the edge of the blank was then probed to determine the position of its centre and therefore how misaligned it was. The blank was then pushed towards the centre of the mount ring using screws that were inserted through the threaded holes that were also used to screw the pads into place. The edge of the blank was then re-probed to measure its new position and find out whether it needed adjusting again. Once the blank was accurately centred to within $20\text{ }\mu\text{m}$ it was clamped in place ready to be bonded.

3.2.2 Applying the adhesive

For Design 1, first the pads are inserted into their T-slots and screwed in place ready to set the bond line thickness to 0.2 mm . This is done by inserting thickness gauges into the gap between pad and blank to measure the current thickness, before adding or removing shims between the mount ring and the pad as needed, before remeasuring to confirm the new thickness. When the thickness of each bond line has been set then

the adhesive can be applied by injecting it through a blunt needle inserted through the hole in the mount ring and the back of the pad using a caulking gun. It was found that the best way to use the caulking gun was to gently use a constant pressure on the trigger as this was most likely to minimise the adhesive leaking out of the sides of the pad, allowing it to reach the ends and the corners.

As described in section 3.1.2.1, the pad shape for Design 2 and the first version of Design 1 meant that injecting the adhesive wasn't an effective way of filling the whole bond. Therefore, the bond line thickness was set as above, with the adhesive then applied directly to the bonding surface of the pad before it was inserted into the mount ring. This allowed air bubbles to remain in the adhesive after it had set, which did not prove to be a problem in the cryogenic testing that was then performed on the mounts, but prompted the redesign of the pads to make injecting the adhesive easier.

For Design 3, a 0.2 mm thickness gauge was inserted into the bond line and then the pad gently screwed in until it just touches the gauge. This makes the process of setting the bond line thickness simpler than having to insert shims, and the adhesive can then be injected as above.

3.2.3 Measurements

To determine the performance of the lens mounts, the gap from the glass blank to the lens mount was measured at three points equally spaced around its circumference. These measurements were taken to determine its relative radial movement between temperatures from 300 K to 120 K by cooling the lens mount down in the cryostat seen in figure 3.7. The amount of contraction of the lens mount relative to the lens was determined for each measurement and was expected to be the same at each position if the lens was staying perfectly centred in the lens mount; any differences could then be used to calculate how far the lens moves and in what direction.




Figure 3.7: The set up used to test the lens mounts cryogenically. The lid of the cryostat has three windows to enable measurements to be taken at three positions around the edge of the lens mount.

To take these measurements, three alignment blocks were attached to the mount ring at 120° intervals. These blocks, as shown in figure 3.8, gave a reference point for the mount close to the edge of the lens so that the distance between the two could be measured through a window in the cryostat lid. A measurement of the radial position consisted of photographing each measurement block with a DSLR camera with a Canon 60 mm macro lens mounted on a linear stage. The stage was used to avoid changing the focus of the camera thus ensuring a constant plate scale for each image, which was calibrated by imaging a rule before and after each cool down.

The temperature was measured by fixing three Pt100 sensors in place using aluminium tape. One was placed on the mount ring next to a flexure, one on the face of the glass blank close to the same flexure with the third in the centre of the glass blank. This way the temperature difference between the mount and the blank could be measured, with the temperature of the mount ring taken as the reference temperature. While the lens was in the cryostat, the rate of change of temperature was limited to 0.1°C s^{-1} using a Lakeshore temperature controller. This was to prevent the lens

from suffering any damage due to thermal shocks and also to limit the temperature gradient between the blank and the lens mount. Only the coldest measurements were left to thermally equilibrate, with the temperature differences between the mount and the blank centre for different measurements typically below 20 K.




Figure 3.8: The measurement blocks used during the cryogenic tests. The distance between the measurement block and the lens was measured as a function of temperature, with three of these measurements taken around the lens edge that are combined together to determine the relative position of the lens.

3.2.4 Analysis

To analyse all of the measurements, first one of the room temperature measurements was taken as the reference image for each of the three positions. Then for each measurement, two cross-correlations were performed between each image and the corresponding reference image: one to align the mount and one to align the lens. The difference in the displacement of the image required to achieve each alignment then gives the contraction between the mount and the lens. The three contractions were then used to determine the offset of the lens in the mount relative to the reference image.

The cross-correlations were performed using the microscopy image analysis software ImageJ (Rasband, 2012), with the cvMatch Template plugin used to align the images with sub-pixel accuracy using a bi-cubic interpolation routine. The first step was to open all of the images of a given measurement position in ImageJ and to combine them into a single stack with the initial image taken at room temperature used as the reference image. The relative rotations of all of the images were then determined by measuring features on the glass blank, before they were all derotated to match the orientation of the reference image.

A mask was then initialised on the corner of the alignment block on the reference image and a cross-correlation performed to align every other image to this reference, with the translations recorded. Next, a mask was initialised on a feature of the glass blank next to the alignment block and another cross-correlation performed, with these translations also recorded. With the distances to align each image with both the lens mount and the glass blank known, the difference between these was found to give the contraction between them for every image in the stack.

This was done for all of the images taken to plot how the mount contracted at each of the the three positions with temperature. The overall movement of the glass blank was then determined by calculating the centre of a circle which passes through the three newly determined point for each temperature.

3.3 Choice of adhesive

For the first tests, a fused silica optical flat was mounted using Design 3 with Invar pads, and fixed in place using Eccobond 286, a two component epoxy which is detailed in table 3.4. This adhesive was chosen as it was also used by a cryomagnetism group in the department to bond quartz windows to titanium mounts at 4K, so it was thought that it would be suitable for these tests. When the flat was mounted, the

cotton paper buffer layer was omitted as a design with the the shorter, CTE-matched pads was thought to be sufficient to protect the glass from cracking. However, after a thermal cycle had been performed on the flat, an inspection of the bond between the lens and the mount revealed that the lens had suffered small cracks at the edge of the pads, as are shown in figure 3.9.

The amount of damage varied between the pads, from a small crack at either end of the pad to a scallop of glass being detached from the edge of the flat, but all of them were affected. During a later test, a time lapse video was taken of the bond line to see how the cracks developed, and it was found that the cracks started at the very edge of the bond where the lens curves away from the pad and propagate inwards. As the cracks grew inwards they also moved around the circumference of the bond, though there was still only one case of the cracks from either end joining together to remove a scallop of glass and effectively detach the pad from the flat.

To test what was causing the damage to the glass — be it the adhesive, the pad or the mount — new pads were fabricated and were adhered to the edge of the flat. These new test pads were connected to the mount ring by a thermal strap to cool them down, but there was no structural link by which the mount ring could exert a meaningful force on the glass. There were three different configurations tested: a pad adhered straight onto the glass, a pad that had a layer of cotton paper adhered in between it and the glass, and also a layer of adhesive applied directly to the edge of the flat without a pad to see if that was enough to cause the same damage.

The flat was then thermally cycled again where it was found that all of the new test pieces had caused damage to the lens edge, as shown in figure 3.10. As all configurations caused cracking to the glass, this implies that there does not need to be any inward force from the mount ring to damage the glass and that just the adhesive itself is a problem. Therefore other adhesives were sourced to test in the same manner to see if any of the three configurations resulted in no damage being

made to the glass.




Figure 3.9: During the first test of Design 3, the glass was cracked at the edges of the bonds by the adhesive. This view is looking through the surface of the flat towards its edge where the pads are bonded. All six of the pads caused some cracking of the lens, ranging from a small amount at the edge of the pad (*left*), to right across the face of the pad (*right*).

The second adhesive that was tested was Lord 3135, which was chosen based on its use by the atmospheric physics group for use in space missions, and is another two part epoxy detailed in table 3.4. The results from the tests of the three configurations were qualitatively the same as those using Eccobond 286, with all samples cracking the edge of the flat, as shown in figure 3.10. This showed that this adhesive was also capable of damaging the glass on its own, meaning that another adhesive was required.

The tests with the same three configurations were repeated again using both Eccobond 286 and Lord 3135, but this time on the polished face of the flat rather than on its ground edge. After thermally cycling the assembly, no damage was found to have been caused to the glass, with the results shown in figure 3.11.

These results can be explained by the polished surface having fewer microcracks and imperfections in it than the ground edge of the flat. These microcracks are weaknesses in the glass and serve as the progenitors from which the larger cracks can then propagate, meaning that the polished surface can therefore withstand higher stresses. The damage to the lenses in the HARMONI camera could therefore be




Figure 3.10: The results of tests using three different adhesives on the ground edge of the flat. These are: Eccobond 286 (*left*), Lord 3135 (*middle*) and CV 1142 (*right*). The top row shows the damage caused by just a layer of the given adhesive and the bottom row shows the results of using cotton paper as a buffer between the lens and mount pad. All tests with Eccobond 286 and Lord 3135 caused cracking to the lens whereas none of the tests using CV 1142 showed any damage — what looks similar to cracking at the edge is actually refraction through the clear adhesive.

prevented by a higher degree of polishing on their edges, though of course finding a suitable adhesive would be preferable solution to reduce the cost of optics.

Work done for the KMOS lens mounts (private communication) showed that that to prevent damage to the glass, a silicone RTV was a more suitable type of adhesive to use. Although these have a higher CTE, which should make damage more likely, they also have a lower stiffness, as when they are cured they are soft and rubbery as opposed to hard and glassy like the epoxies tested. This means that even though the adhesive contracts more relative to the glass, a lower stress is imparted on the glass. In KMOS the adhesive used was Dow Corning 6-1104 CV, but as this is not available for purchase in UK an adhesive with similar properties, CV 1142, was procured. This lower stiffness of CV 1142 is only the case for its room-temperature state, as it undergoes a glass transition after the majority of the relative contraction has occurred at 150 K. The cure time and CTE are compared to those for the two epoxies in table 3.4.

Adhesive	Type	Cure time (days)	CTE (ppmK^{-1})
Eccobond 286	Epoxy	1	36
Lord 3135	Epoxy	1	32
CV 1142	Silicone RTV	7	320

Table 3.4: The properties of the different adhesives that were tested.

The same three configurations were tested on the ground edge of the flat using CV 1142. This time however, as seen in figure 3.10, there was no visible damage caused by any of the tests. This showed that the adhesive did not crack the glass on its own, nor did an adhered Invar pad. This suggested that a cotton paper buffer layer was not necessary, but it was decided to perform some of the full tests using mounted glass blanks both with and without it, in case there was any damage when the number of pads and cool downs was increased.




Figure 3.11: The results of adhesive tests that were repeated on the polished surface of the lens as opposed to the ground edge. This image is taken through the glass flat and is looking at the undersides of each bond. None of the adhesives were found to cause any damage to the lens.

3.4 Design 1 - Double-beam spring

3.4.1 Version 1 - Curved pads

Version 1 of Design 1 was used to mount a glass blank mounted and thermally cycled four times, with the results of these tests shown in figure 3.12. The total decentre measured was approximately $23\text{ }\mu\text{m}$, and it can be seen that the results were quite repeatable, with the flat decentring to within $\pm 5\text{ }\mu\text{m}$ of the same position for each cool down before returning to its original position upon warming it back up.

This decentre is larger than the $20\text{ }\mu\text{m}$ tolerance on the radial positioning of a lens, which would pose a problem for the HARMONI cameras as if they were aligned at room temperature they would move by enough to not meet the optical quality specifications after cooling them down. If this design of mount were to be used, then given that the movement is repeatable, it would have to be cooled down and tested before it was aligned to measure the magnitude and direction of the decentre. This information could then be used to align the camera in such a way that the lenses are properly aligned only after cooling them down.

The line in figure 3.12 shows the amount of contraction that was predicted using the CTEs of both the glass blank and the aluminium blank. This line should be treated as a rough estimate only, as the CTE of the glass was not known as a function of temperature and so all predictions were made using room temperature values by assuming the different materials had similar temperature dependencies. The data appears to match the theory well, though how well is uncertain due to the assumptions that had to be made.

As the mount took just over a day to cool from room temperature down to 130 K , the assembly was left overnight to cool down when no measurements were taken. As these were the first full tests completed, all of the cool downs were initiated at the same time of day, leading to the overnight period occurring at the same temperatures

as can be seen by the lack of data between 200-250 K. When testing other mount designs after this, the timings of the different cool downs were planned such that this lack of data for a region did not occur.




Figure 3.12: The results of testing Design 1 using the first version with curved pads. *Left:* The amount of contraction measured between the flat and its mount as a function of temperature. Each colour represents a set of measurements taken at one of three positions around the edge of the lens mount. The fact that there is not equal contraction between each of the positions shows that the glass blank moved towards the positions with more contraction. The straight line shows the expected contraction, although it is based on the room temperature CTEs due to a lack of temperature dependent data for the glass. *Right:* The position of the lens as a function of temperature, as derived from the measurement of the contraction between the lens and its mount. It can be seen that there is a total displacement of $23\mu\text{m}$ between room temperature and 130 K.

3.4.2 Version 2 - Flat pads

After redesigning Design 1, version 2 was assembled, with the results for 11 thermal cycles shown in figure 3.13. It can be seen that the amount of contraction is higher at one position than for the other two, which results in a total offset in the radial position of the blank of about $18\mu\text{m}$ between room temperature and 120 K and is directed towards the position with more contraction.

When it is compared with figure 3.13 it can be seen that there is less scatter at low temperatures, meaning that the results are more repeatable and therefore more

confidence can be had with the results. There is however a noticeable scatter in the contraction measurements at each position, especially for temperatures between 200-300 K. This is due to some of the measurements being taken as the assembly is cooling down and some as it is warming back up. This means that the lens is at different temperatures for a given temperature of the mount: it is warmer than the mount during the cool down measurements and colder whilst everything is warming up. This does not affect the derived position of the blank as the effect is the same for each measurement position and hence the same scatter is not visible in those results.




Figure 3.13: The initial results using Design 1. *Left:* The amount of contraction measured between the lens and its mount as function of temperature. Each colour represents a set of measurements taken at one of three positions around the edge of the lens mount. The fact that the magnitude of the contraction was larger at one of the positions shows that the lens moved towards that position as it was cooled down. The straight line shows the expected contraction, although it is based on the room temperature CTEs due to a lack of temperature dependent data for the glass. *Right:* The position of the lens as a function of temperature, as derived from the measurement of the contraction between the lens and its mount. It can be seen that there is a total displacement of $18\mu\text{m}$ between room temperature and 130 K.

This decentring of the blank is higher than desired as it approaches the $20\mu\text{m}$ tolerance for the radial alignment of the HARMONI lenses, leaving very little budget for the alignment of the lenses in their mounts and then in the camera barrel. Whether this is a big problem or not depends on the final optomechanical design of the camera

barrels, which has not been done yet. If they are designed such that the mounted lenses can be moved in the barrel and have to be aligned then any misalignments between the lenses and their mounts can be recorded and then removed during the second alignment process given that the current radial alignment procedure is limited by the adjustment precision rather than the measurement of the lens' position. If however, the mounted lenses are aligned in the barrel using just a mechanical fit then this would ease the alignment process at the possible expense of accuracy if the residual misalignments cannot be compensated for. A decision on this will be made depending on factors such as which is greater out of the accuracy with which the lens can be adjusted in the mount and the accuracy the mount can be aligned in the barrel. With the current alignment plans, as described in Chapter 4, the budget for aligning the lens in the mount is 5-10 μm , with the final camera barrel and alignment procedure not yet designed.

If this performance is ultimately not good enough for the HARMONI camera then, as with version 1, the decentre of the blanks does however appear to be consistent in both the magnitude and direction of the movement, with the measurements of the full decentre varying by about $\pm 3 \mu\text{m}$. This means that there is again the possibility that this offset could be measured before assembling the camera to dramatically reduce this misalignment. This could lead to a large amount of extra work in the assembly of the HARMONI spectrographs, though it may only have to be done for the lenses with the tightest tolerances.

It was found when performing further tests that the blank had not been aligned perfectly when mounting it, and was miscentred by about 100 μm , so it was thought that this may be the reason why the blank was moving. The blank was removed from the mount and cleaned up ready to be remounted more accurately and tested again. This time the a more rigourous alignment procedure was followed - making sure the position of the blank was remeasured after inserting the adhesive to ensure it

was still well centred, and using more accurate shims to try and control the thickness of the bond line as well as possible. Importantly, the pads were kept in the same position around the blank's edge as when it was first assembled — to test whether they instead were causing the movement.




Figure 3.14: The results from the first cool down of Design 1 after realignment. *Left:* The amount of contraction between the lens and its mount at three positions around the lens edge. The straight line shows the expected contraction, although it is based on the room temperature CTEs due to a lack of temperature dependent data for the glass. *Right:* The position of the lens as a function of temperature. It can be seen that there is a maximum displacement of $14\mu\text{m}$ as it is being cooled down, but that reduces to a displacement of $7\mu\text{m}$ at 130 K.

The initial results of this remounted blank, shown in figure 3.14, were promising - showing a maximum decentre of about $14\mu\text{m}$ before moving back to give a final displacement of only $7\mu\text{m}$. However, further cool downs, whose results are shown in figure 3.15, show worse performance: moving in the same direction as this first cool down, but by a magnitude of $18\mu\text{m}$ - the same as the results with the first assembly. This seems to suggest that a misalignment of $100\mu\text{m}$ during the mounting of the lens does not have a significant effect on how much the lens will move when cooled down, which allows a looser tolerance between the mechanical and optical axes of the HARMONI lenses.

The reason why this first cool down gave much better results than subsequent tests




Figure 3.15: The results from further testing of the realigned Design 1, excluding measurements taken during the first cool down. *Left:* The amount of contraction between the lens and its mount at three positions around the lens edge. The straight line shows the expected contraction, although it is based on the room temperature CTEs due to a lack of temperature dependent data for the glass. *Right:* The position of the lens as a function of temperature. It can be seen that there is a total displacement of $18\text{ }\mu\text{m}$ between room temperature and 130 K.

is as yet unknown, it is the only time that a first cool down has given significantly different results from the others. It could be that the properties of the adhesive changed slightly the first time it went through the glass transition, but it seems odd that this effect has not been observed for every assembly if this is the cause. A more likely explanation could be that the assembly suffered a slight amount of damage between the first and second cool downs, but CMM measurements do not show the blank as having moved in the mount since being aligned which does not provide any evidence for this scenario.

The fact that when the blank was remounted it moved by the same distance, but in a different direction indicates that the problem does not lie in the blank being slightly miscentred to begin with, and the fact that the flexures were all kept in the same positions suggests that neither is it that there are differences between the flexures causing a net force on the blank when it is cooled. This means that the cause remains unknown for now and we have no way to predict in which way a lens will




Figure 3.16: The results from the testing Design 2. *Left:* The amount of contraction between the lens and its mount measured at three positions around the lens edge. The straight line shows the expected contraction, although it is based on the room temperature CTEs due to a lack of temperature dependent data for the glass. *Right:* The position of the lens as a function of temperature. It can be seen that there is a total displacement of approximately $40\text{ }\mu\text{m}$ between room temperature and 130 K.

move when it cooled down before we have measured it, meaning that further testing is required to understand this effect.

3.5 Design 2 - Z-spring

The results from Design 2 are shown in figure 3.16, but are much worse than for Design 1. There is a lot of decentre of the blank from warm to cold temperatures: at around $40\text{ }\mu\text{m}$ total displacement. This is much greater than for the other flexure design and also tightest tolerances in the current camera design.

As the flexures have holes machined out of them and are designed to buckle when a radial load is applied they are more fragile than the flexures in Design 1. Some flexures were observed to have been permanently deformed and be slightly buckled even before cooling the assembly down, as seen in figure 3.17. This means that it is likely that there was a variation in the spring constant for different flexures, which




Figure 3.17: One of the z-springs used in design 2. It can be seen that diagonal bar of the spring is not perfectly straight and is slightly buckled. This happened before the flexure was cooled down, showing that the design is not robust enough and is no longer being considered for use in HARMONI.

would explain the decentre measured in these tests.

These pads should not have been put under any significant stress prior to being cooled down, so the most likely time the damage could have been caused was during the manufacturing process, rather than their transport to the lab or during the mounting process. Either way, these flexures do not offer any advantages over Design 1 whilst also not being robust enough and so there are no plans to pursue this design any further.

3.6 Cotton paper

As described in section 3.3, when using CV 1142 there was no damage caused to the flat both whilst using cotton paper and without it, which was the case on both a ground surface and a polished surface. This suggested that it is not necessary for the lens mount design to incorporate it, but it was decided to make sure there were full-scale prototype test both with and without it for more realistic testing.




Figure 3.18: Marks on the edge of the glass blank that occurred after cooling it down without a cotton paper buffer layer. It was thought that they were cracks in the glass, but they did not grow with subsequent cool downs and did not appear to be cracks on closer inspection after removing the blank from the lens mount.

However, in the first full-scale test in which it was not used there appeared to be small amount of damage at the ends of the pads on the edge of the adhesive, as shown in figure 3.18. As the tests were performed on a glass blank, not an optical flat, the bond could not be seen through the glass as it was in the adhesive tests and thus it was difficult to be sure that it was a crack. It was also difficult to closely inspect the possible damage as it was right next to the pads and the mount ring could not be moved without risking damaging the flexures.

The marks were where cracks were expected to form if they were present, but they were not conclusive, so the mount was cooled down again to see if it caused them to grow. Another visual inspection couldn't detect any changes after this cooldown, so the mount ring was removed from the assembly to enable a closer view and the pads were also removed. The closer view revealed that the appearance and texture of the marks were not consistent with cracks after all.

It appears that cotton paper is not necessary when CV 1142 is used as the adhesive, but further small survival tests may be desired to increase the confidence of this conclusion. These tests would also need to be performed if a different adhesive was

sought as an improvement on CV1142, for example if one could be found with similar mechanical properties but with a shorter cure time to speed up mounting procedure.

3.7 Conclusions

When bonding cryogenic lenses in place the choice of adhesive is very important. The first tests of prototype lens mounts resulted in the glass cracking, a problem that was isolated to the adhesive. Two epoxies were tested and found to crack the glass when cooled down, before a silicone RTV was found that did not cause any damage. The epoxies were thought to have damaged the glass due to their high CTEs, but the silicone RTV had a higher CTE but a lower stiffness which allowed lower stresses to be placed on the glass to prevent damage.

Tests show that with CV 1142 a layer of cotton paper between the lens and its mount is not needed to prevent damage to the glass. This was suggested by the adhesive selection tests and backed up by the lack of damage suffered during full-scale prototype testing. This should remove a step from the mounting of the lenses which should speed up the process and reduce the risk of damage to the lens.

Version 2 of Design 1 is the baseline design for the lens mounts in the HARMONI spectrograph cameras. The lens has been measured to move by $18\text{ }\mu\text{m}$ when cooled from room temperature to operating temperature, and although this has been shown to be repeatable, it is not known a priori in which direction it is going to happen. The current camera alignment plan allows lenses to be aligned to within $5\text{ }\mu\text{m}$, which if this were maintained after mounting in the camera barrel would mean that the lenses most sensitive to a radial shift would need to be cooled down to measure the decentre when cold so this could be accounted for when mounting the lens in the barrel. To avoid these additional cold tests the procedure for mounting the lenses in the barrel would need to reduce the misalignment errors to below $2\text{ }\mu\text{m}$.

The performance of Design 2 is worse than Design 1; with a $40\text{ }\mu\text{m}$ decentre of the lens measured. This is attributed to slight permanent deformations in the flexures, which are more fragile than for Design 1 due to their need to buckle under stress, and so this design is no longer being considered for use in HARMONI.

Chapter 4

Rebuilding the Oxford SWIFT Spectrograph spare camera

The Oxford SWIFT Spectrograph (Niranjan Thatte, 2006) is an *I*- and *z*-band IFS which was designed and built by the University of Oxford to work on the 200" Hale Telescope at Palomar Observatory with high-order adaptive optics.

However, the as-built image quality of the spectrograph was not as high as expected from the theoretical optical design — a problem that has, as yet, not been solved. One idea for the source of the aberrations was the spectrograph camera, of which there are two in the instrument and a spare third camera back in Oxford. To test this hypothesis, the spare SWIFT camera was disassembled and reassembled using the new lens mounting schemes designed for the HARMONI spectrograph cameras described in Section 3.1.2.

The aims for this work were twofold: to improve the image quality of the SWIFT camera in an attempt to solve the image quality problems, and also to come up with a reliable process for, and gain experience in, mounting powered lenses and assembling them into a camera. The SWIFT cameras are not designed for cryogenic use, so this second aim is supposed to work in conjunction with the cryogenic testing presented

in Chapter 3. When these two projects are considered together it should mean that we have confidence in both our ability to accurately assemble a camera and that it will then maintain its alignment when cryogenically cooled, and thus have confidence in our ability to build cameras for HARMONI that meet their requirements in use.

4.1 SWIFT camera

The camera has a focal length of 117.6 mm and accepts a collimated beam from the diffraction grating with an elliptical pupil of 45 mm $\times$ 22.5 mm. The elliptical pupil is due to anamorphic magnification in the pre-optics of the SWIFT spectrograph, which is used to allow Nyquist sampling of the spectrum whilst maximising the number of spatial elements that can be fit onto the detector. This gives it a focal ratio of 2.6 along the slit and 5.2 across it. The input field of view is approximately $7^\circ \times 14^\circ$ which maps onto a detector that is 30 mm $\times$ 60 mm in size. It is designed to work at wavelengths from 650 nm to 1000 nm, with the optical quality and throughput still good enough to enable testing using a 1064 nm infrared interferometer.

4.1.1 Optical design

The camera design consists of 6 lenses with spherical surfaces, 2 of which form a cemented doublet, and these are separated into two groups of three lenses each. The design is pictured in figure 4.1 with the parameters of each lens found in table 4.1.

The general tolerance on the positioning of a lens is $\pm 40 \mu\text{m}$ on the decentre of the lenses and $\pm 50 \mu\text{m}$ on the spacing between adjacent lenses. However, there is a looser tolerance on the distance between lenses 3 and 4 — between the two groups.




Figure 4.1: The optical design of the camera, with the sandwiched diffraction grating. It is shown cut through the plane of dispersion with 15 wavelengths between 650 nm and 1000 nm plotted. For scale, the largest lens has a diameter of 123 mm.

Lens	Left ROC (mm)	Right ROC (mm)	Centre thickness (mm)	Preceding air-gap (mm)	Glass
1	-67.0498	-258.9939	7.50	—	S-TIH6
2	-749.3723	106.4174	15.00	7.162	S-LAM2
3	293.1195	124.7324	40.00	0.50	S-LAM2
4	288.8594	487.2714	29.00	55.25	S-LAM3
5	119.6143	97.9265	32.10	0.50	S-LAM3
6	-97.9265	97.1095	8.00	0	S-TIH6

Table 4.1: The parameters of each lens in the SWIFT camera, with ROC denoting radius of curvature. Lenses 5 and 6 form a cemented doublet.

4.1.1.1 Damaged lenses

When the camera was inspected, it was found that 3 of the lenses had suffered some damage. As the camera had been in storage since it was manufactured, it should not have experienced any shocks, and so a possible explanation for any damage to the glass could be the stresses it was placed under due to the mounting procedure.

Lens 1 was found to have been very badly damaged, as shown in figure 4.2, and was unusable. This is the lens which is most likely to have been damaged due to mounting stresses, which is discussed in Section 4.1.2.1.

Lens 4 had a small crack, seen in figure 4.3, which was not large, or central, enough to render it unusable, however a replacement was still ordered alongside one




Figure 4.2: The damage to the original Lens 1 in the spare camera. For scale, the diameter of the lens is 86 mm.

for Lens 1.

There was also some damage to lens 5/6 - the cemented doublet. It is shown in figure 4.4, but the cause of the damage has yet to be ascertained. It is in the form of a series of striations along one side of the lens, which appear to be located at the interface between the two lenses. This lens was not been replaced due to the expense that would be required, meaning that it cannot be known how big an effect it has. When the camera was interferometrically tested these striations showed up as lines on the wavefront maps at some parts of the field of view. It is hoped that where these lines cannot be seen this damage has no effect.

4.1.2 Mechanical design

4.1.2.1 Original design

The mechanical design of the original assembly can be seen in figure 4.5. As the optical design features two groups of lenses with relatively loose tolerances on their position relative to each other, the original mechanical design had two camera housings to




Figure 4.3: The crack at the edge of the original lens 4 in the spare camera. For scale, the diameter of the lens is 119 mm.

which the lenses were attached. These housings could then be mounted into the barrel to form the final camera.

The lenses were held in place by clamping them against the camera housings by the outer edge of the optical surface. The annular surfaces to which the lenses were clamped were shaped to exactly match the spherical surface of the lens itself, with the inner diameter such that it falls outside the clear aperture of the lens. Inserts were then screwed into the housing with similar annular surfaces shaped to match the opposite surface of the lens. By screwing these inserts in tightly enough, the lens was held in place, though they placed a certain level of stress on the lenses themselves.

The alignment between the two housings was achieved using manufacturing tolerances on each housing and the camera barrel itself, which meant that the housings had a simple push fit into the camera barrel that could then be screwed in place.

Given that the spare camera was found to have damaged lenses, and that it had been stored in a cupboard for its lifetime, it is possible that the force necessary to hold the lenses in place using this lens mounting scheme was the cause of the damage. This is not certain, and it would probably mean that the surface of the lens did not




Figure 4.4: The striations visible in lens 5/6. They appear to be located at the boundary between the two lenses, though their exact cause is unknown. For scale, the diameter of the lens is 108 mm.

exactly match its partner surface due to damage or a foreign particle. It does however, show one advantage of a stress-free mounting scheme — if the stress of mounting is indeed high enough to crack a lens then it would definitely be high enough to deform the surface of the lens, which could affect the image quality of the optics.

4.1.2.2 New design

The mechanical design work of the new mounting scheme was again done by Jim Lynn, with the aim of reusing as many components of the original assembly as possible to reduce costs. However, a side effect of this is the constraints it places on sizes of the new components if they are to fit into the same volumes as their predecessors.

Two mounting schemes were used: Designs 1 and 3 from Chapter 3. This is due to the difficulty in working in the restricted volumes and also allowed us to gain experience with two slightly different types of mount. Both of these designs consisted of a circular mount ring with flexures between it and the lens, which was adhered in place. This means that the mounts themselves were originally designed to be used




Figure 4.5: The original mounting scheme of the camera. For scale, the diameter of the largest lens, lens 3, is 123 mm.

at cryogenic temperatures, but for these tests they were only intended to be used at room temperature.

There were some differences in these designs when compared to the HARMONI designs in Chapter 3. These were due to the space constraints brought about by reusing the old camera barrel and therefore needing the new camera housings to fit into the same volumes as the old ones. These designs are shown in figure 4.6. The HARMONI mount pads fit into a T-slot in the mount ring and move away from the lens as they are loosened, whereas the redesigned mount pads are simply screwed to the inside edge of the mount ring and therefore move towards the lens. This made it very difficult to insert shims between the pad and the mount ring without touching the lens, and therefore increasing the difficulty of setting the bond line thickness.




Figure 4.6: The new lens mounts. They are based on the prototype HARMONI cryogenic lens mounts, but modified to fit the space constraints of the SWIFT camera. *Left:* The 97 mm diameter lens 2. *Right:* The 86 mm diameter lens 1.

4.2 Assembling the camera

Assembling the camera is done in two main steps with this lens mounting scheme. Firstly the lenses are aligned and inserted into the lens mounts, before the lens mounts are aligned and inserted into the camera housings. The housings can then be inserted into the barrel and screwed into place to complete the camera.

4.2.1 Mounting lenses

The lenses were mounted into their mount rings using a two stage process. First the lens was precisely aligned to the mount rings and clamped in place, before the second step of inserting the adhesive that was used to hold them in place.

4.2.1.1 Aligning lenses

A method for aligning and mounting powered, spherical, lenses has been developed and tested, with the philosophy that the lens surface should be touched as little as possible to minimize the chance of it being damaged. This work was the result of a

collaborative effort from myself and Jim Lynn. The aim was to place the lens down on three setting balls positioned outside of its clear aperture, measure its position and then adjust only the mount to move it into alignment with the lens.




Figure 4.7: The jig which was used to align the lenses when remounting the optics in the spare Oxford SWIFT Spectrograph camera. The mount is held between four alignment blocks using slip gauges and sits on a three point kinematic mount. The lens sits on three setting balls which are placed outside of its clear aperture. For scale, the width of the baseplate is 300 mm.

A jig was designed to be used with the department's CMM, which is capable of making an individual measurement with an accuracy of less than $5\text{ }\mu\text{m}$. The design, seen in figure 4.7, uses a three-point kinematic mount on differential adjusters to move the mount in the axial direction and also to adjust the tip and tilt of the mount relative to the lens. The mount ring was positioned on the jig between four alignment blocks, with stacks of slip gauges placed between the mount ring and the blocks. The radial position of the mount could then be changed by adjusting the thickness of the slip gauge stacks to ensure that the lens was precisely centred within it.

First, the CMM was used to probe the three setting balls to find their positions, before the lens was placed on top, and clamped in place. The position of the centre of curvature of the lens's bottom surface was calculated using the positions of the

three setting balls and the surface's radius of curvature. The centre of curvature of the lens's top surface was then found using the CMM. This was done by probing the surface using three concentric rings of points each. This meant the position of the optical axis, being the line connecting the centre of each spherical surface, was then known, ready for the mount to be moved into position.




Figure 4.8: The springs that were used for clamping the lens in place using its edge during the alignment of the mount. For scale, the diameter of the lens is 123 mm.

There were two methods implemented for clamping the lens in place. The simpler of the two methods was holding the lens edge by three springs connected to the base of the jig, as seen in figure 4.8. The other design involved fixing a ring above the lens on two posts. Three ball-spring plungers were then screwed into the ring to push the lens onto the setting balls.

The first alignment was to make sure that the mount was perpendicular to the optical axis. This was measured by probing the flat surface of the jig just outside the the lens mount and comparing it to the angle of the optical axis. It could then be adjusted by turning the appropriate differential adjustor to tilt the jig in either the $x - z$ or $y - z$ plane before remeasuring to verify the new position. The tip and tilt of the jig were adjusted and verified until they both fell within the tolerance of 0.01° .

The second alignment was to position the mount the desired distance along the optical axis. This was measured by again probing the flat of the jig just outside the mount, but then comparing it to the positions of the centres of curvature of the lens surfaces. This was adjusted using the same differential adjusters on the kinematic mount, but this time by moving them all by the same amount. Again the alignment was adjusted and verified until the alignment was within the tolerance of $50\text{ }\mu\text{m}$.

As both the axial and angular alignments were done using the same adjusters, it was possible that performing one alignment affects the other. This meant that an iterative procedure of performing an angular alignment then an axial alignment was needed to make sure that both alignments were within tolerance, though this typically only took a couple of cycles to complete.

The final alignment was to ensure that the optical axis passed right through the middle of the mount ring. This was done by probing the outer edge of the mount ring to determine the distance from its centre to the optical axis. The thicknesses of the stack of slip gauges between the mount ring and the jig's alignment blocks were then adjusted to move the mount in the desired direction before verifying that the alignment is within the tolerance of $10\text{ }\mu\text{m}$, to leave some tolerance for assembling the lenses in the camera barrel.

4.2.1.2 Adhering lenses

When the mount was in place and all alignments had been verified to be within their tolerances, the procedure for setting the bond line thickness and inserting the adhesive outlined in Section 3.2.1 was followed to fix the lens in place.

The procedure for setting the bond line depended on the design of mount being used.

For Design 3, the bond line is adjusted by screwing the pad in and out using the slot in the back of the pad. The bond line thickness is set to 0.2 mm by placing an

appropriate thickness gauge into the gap and very gently screwing the pad in until it just touches the gauge. When this was first attempted the force of the pad touching the thickness gauge was enough to move the lens more than the $40\text{ }\mu\text{m}$ radial tolerance on the lens position, meaning that it had to be realigned. However, this problem was reduced with more practice and making sure the pad was screwed in very slowly and stopping as soon as it touches the gauge.

For Design 1, setting the bond line was more challenging. It was achieved by inserting shims in between the mount ring and the pads to decrease the width of the bond line. However, due to the space restraints of the redesign of the SWIFT camera this meant taking the mount ring off the jig every time shims had to be inserted or removed. Due to the design of the jig, this could be done without disturbing the alignment if extreme care was taken however, it meant that only the springs could be used to clamp the lens in place and replacing the mount ring was difficult without disturbing the alignment — as the inner radius of the pads was only 0.2 mm larger than the outer radius of the lens.

Once the bond lines were set the alignment was reconfirmed before injecting the adhesive. For these tests the adhesive was Nusil CV1142, a silicone RTV, which was injected using a caulking gun and blunt ended needle. The first time this was attempted the lens was found to have moved by around $100\text{ }\mu\text{m}$ afterwards, which was mostly due to the needle coming into contact with the lens edge, so the following steps were introduced to minimise the chance of it happening again. Firstly, a collar was fitted onto the needle to prevent it from being inserted far enough into the back of the pad to touch the lens edge. Secondly, the order in which the pads were injected always traced a star-shaped pattern rather than just being injected round in order to minimise the chance of the lens being pushed in one direction. Lastly, if the bond line opposite the one being injected was empty, a thickness gauge was placed in the gap to try and prevent the lens being pushed that way.

However, despite these precautions there were occasions where it proved particularly difficult to actually inject the adhesive into the bond line, and the lens appeared to have moved by around $50\text{ }\mu\text{m}$ after the glue had been inserted. In these cases the adhesive was still pliable enough to redo the radial alignment using the slip gauges, but this is not ideal and shows the method needs refining before being used for the HARMONI cameras.

Once the lens is aligned and adhered, the final measurements are recorded so that the precise position of the lens in the mount is known, and the adhesive is left to fully cure, a process that takes about 7 days for CV1142. Due to limited time available to mount the lenses, they had to be gently moved off the jig after 3 days before being left to finish curing to enable all of the lenses to be mounted in the available time.

4.2.2 Final assembly

The process of assembling the mounted lenses into the camera housing is a simpler procedure as there is only one alignment that needs to be done.

The camera housing was measured with a Faro Arm CMM to set up the coordinate system before the lens was put approximately in place and its mount measured. The lens mount was then moved so that the now-known position of the optical axis was within $40\text{ }\mu\text{m}$ of the centre of the camera housing before being remeasured. This movement was done by mounting a micrometer next to the edge of the lens mount and adjusting it to push the mount in the desired direction.

This worked for all of the lenses apart from lens 2, which is mounted inside its housing rather than at the end, meaning that the micrometer could not reach the edge of the lens mount. Therefore the adjustments needed to be done by hand, necessitating many more iterations before alignment was achieved.

Once all lenses for a housing were assembled, the housing could then be pushed into the camera barrel and screwed in place.

When the camera was fully assembled, it could be tested on the interferometer to measure its optical quality and verify the alignment of all the the lenses in the camera. This process is further described in Section 4.3.

4.2.3 Discussion

The differential adjustors used for the axial and angular alignment of the lens mounts have a pitch of $25\text{ }\mu\text{m}$, meaning that the axial position of the lens mount can be positioned to within $10\text{ }\mu\text{m}$ peak to valley. The angular precision depends on the baseline between the adjustors, which for these tests was 200 mm , giving an angular precision of approximately 0.003° for both tip and tilt. The set of slip gauges used had $10\text{ }\mu\text{m}$ steps in thickness, allowing the radial position of the mount to be set to within about $5\text{ }\mu\text{m}$ ptv in both the x - and the y -direction.

The jig was easy to use, with the most fiddly step adjusting both the axial position and the tip and tilt with the same set of adjustors. It would be more convenient to decouple these two steps, however once the alignment is close only one or two iterations between the two steps were necessary to get them both in alignment. Overall, once the program for the CMM has been written, a lens can be fully aligned and glued in approximately 2-3 hours.

The easiest step of the alignment was the radial adjustment using the slip gauges, where each block of slip gauges should be adjusted by the amount of misalignment, rather than having to calculate how many fractions of a turn were needed.

The biggest problem with the alignment procedure was the occasional displacement of the lens when the adhesive was injected. The solution to this problem could be to improve the way in which the lens is clamped in place alongside upgrading the equipment for injecting the adhesive. The steady flow rate produced by an automatic glue dispenser should avoid such a high maximum pressure placed on the adhesive, minimizing the chance of it moving the lens. The springs to hold the lens in place

were not particularly effective as they made it more difficult to place the lens on the setting balls and did not offer a large enough lateral force to keep the lens in place — it was easy to move the lens just by knocking it. The clamping ring offered a better clamping force, but was incompatible with the modified Design 1 of the lens mount due to having to remove the mount ring to insert shims. This should not be a problem with the HARMONI cameras as they will not have the same tight volume restrictions. It is also possible that this is less of a problem with the HARMONI lenses which are larger and heavier.

Another aspect which is not ideal is the adhesive cure time. If only one lens could be mounted per week then it would take half a year to adhere the lenses for 4 HARMONI infrared cameras. Therefore to decrease this to a more reasonable time frame the process would need to be parallelized by having, say, 4 jigs to mount 4 lenses per week.

4.3 Interferometric testing

A 30 cm diameter, $1.064\,\mu\text{m}$ 4D Accufiz interferometer was used in conjunction with the 4Sight companion software to test the optical performance of the camera and selected lenses. Two different test setups were used: one to test the whole camera in double-pass, and another to test the form of an individual optical surface. The test setups themselves were mounted on an optical table fitted with a pneumatic isolation system to reduce vibrations. This was necessary to allow a series of measurements to be taken in the temporal rather than dynamic mode, therefore giving the interferograms a higher resolution, as described in Section 1.3.

4.3.1 Lens surface form

The form of as many of the surfaces as was possible was measured for the lenses both before and after mounting. These tests were performed on the first lens in the original camera as well as on every feasible surface with the new mounting scheme. The aim was to see if there was any discernible change in optical performance due to the shape of the lens being deformed by any stresses induced by mounting the lenses.

4.3.1.1 Method

The interferometer was set up with a 130 mm radius transmission sphere to create a convergent beam into which the lens could be inserted. It was set up on a rail system, ensuring that the rail ran directly below and parallel to its beam and could be used to move a stage along the z -axis. The lens itself was placed in a three-jaw chuck in a mount capable of tip and tilt adjustments, and this was secured to a z -stage on the rail. The heights of the interferometer and the lens mount were set such that their centres were the same distance above the table. This setup can be seen in figure 4.9.

To measure the surface form of the lens it was positioned in the confocal position: where its centre of curvature coincided with the centre of curvature of the transmission sphere. Therefore, a concave surface needs to be mounted further away than the focus of the beam and a convex surface closer than the focus. Schematics for the setups used to testing both types of surface can be seen in figure 4.10.

During testing it was found that the easiest way to align the lens in this way was first to align the lens so that the transmission sphere centre of curvature lies on the lens surface itself: the ‘cats-eye’ position. This was first done roughly by placing the lens at 130 mm, the transmission sphere radius of curvature, away from the interferometer. This, combined with the constraint of the rail meant that the alignment was close enough that a return signal from the lens was visible using the interferometer alignment camera. The tip and tilt of the lens could then be varied to




Figure 4.9: The setup to test the surface of a lens using an interferometer. For scale, the width of the interferometer's beam is 300 mm.

move the return signal into the alignment using the box indicated on screen, which meant that the alignment was then good enough to see a fringe pattern on the live interferogram. These fringes could be minimised by adjusting the tip and tilt of the lens along with its z -position, which would have precisely aligned the lens in the cats-eye position, though this was not necessary for this stage of the work.

To confirm the z -position of the lens, a piece of card was inserted into the beam at one side of the lens. If the shadow of the card was cast on opposite edges of the interferogram simultaneously then the lens was indeed at the cats-eye position, however, if the shadow was cast on only one edge then the lens would have already been at the confocal position, albeit accidentally at this stage.

If the lens was placed in the cats-eye position then it needed to be moved into the confocal position, which was done by simply sliding the stage along the rail whilst watching the image of the live interferogram. For a concave surface the lens was moved further away, with a convex surface moved closer. The alignment of the rail to the interferometer should have been sufficient to keep the interferogram on the screen




Figure 4.10: Schematics of the setup used to test the surface of a concave (*left*) and convex (*right*) lens using the interferometer. Both lenses must be placed such that their centre of curvature is coincident with the centre of curvature of the interferometer transmission sphere. This results in a concave lens needing to be positioned further away than a convex lens.

at all times, but if however, the interferogram drifted out of the field of view then the information from the alignment camera was used to adjust the tip and tilt of the lens to realign it. As the lens was moved from the cats-eye to the confocal position the interferogram appeared to get smaller until it appeared point-like, before increasing in size to a maximum when correctly positioned. When fringes could be discerned on the interferogram these could be minimised to precisely position the lens, with the tip and tilt adjusters used to remove vertical and horizontal fringes and the stage moved in the z -direction to remove circular fringes.

Once the lens was accurately aligned, the interferometer could be set up to take measurements. First the focus was checked by placing an object with a hard edge, such as a piece of card, in the beam of the interferometer close to the plane of the lens surface without actually touching it. The focussing mechanism was then adjusted until the edge of the object on the interferogram appeared as sharp as possible. Two settings that needed to be correctly specified were the wedge factor and the terms to remove. The former is the inverse of the factor by which optical aberrations are multiplied due to the experimental setup; in this case it is 0.5, as illustrated in




Figure 4.11: When measuring the form of a surface using an interferometer, the path difference is twice the size of any aberration due to light having to pass the aberration both before and after the reflection.

figure 4.11. As these tests are measuring spherical lens surfaces, the Zernike terms that need to be removed from the aberration map are the tip, tilt and any defocus. This is because these terms are caused by residual misalignment: in the angle of the lens for the tip and tilt, and in the focus of the system for the defocus. Therefore any remaining aberrations are the deviations from a perfect sphere.

Next, the settings of the interferometer's detector were chosen using the detector's live exposure histogram to ensure there was no under- or over-exposed pixels in the image. To ensure the best possible fringe contrast, the gain was kept as low as possible, with the exposure time increased to get the right exposure for the image. The trade-off was performed like this as the vibration isolating optical table allowed the use of the temporal mode of the interferometer. As the temporal mode uses several images to deduce the direction of the slope of the wavefront, if the table is stable enough for it to be used then it is also stable enough to increase the exposure time of the detector rather than the gain.

If there were any interferometer returns coming from surfaces other than that being tested then these were masked out to avoid any confusion before the measurements were finally taken. A series of 10 wavefront measurements were captured and then averaged together to minimise the random uncertainties from any thermal currents in the air or residual vibrations and form the final aberration map. The interferometer software was then used to calculate the PVr values for the aberration maps, as well as their residual rms values.

These tests were first performed on the original, broken, lens 1 whilst it was still mounted before being repeated after unmounting the lens. Unfortunately the camera was disassembled before further mounted lenses could be interferometrically tested, however, all surfaces that could be tested with the available equipment were tested before and after mounting with the new mounting scheme. The surfaces that could not be tested were those on lens 4 and the left-hand surface of lens 3 as these were all convex surfaces with radii of curvature greater than that of the transmission sphere, therefore they could not fit inside the sphere.

4.3.1.2 Discussion

The PVr and rms values of the aberration maps for the lenses before and after mounting can be found in table 4.2. Here it can be seen that there was no significant increase in these quantities, and therefore degradation of form of the surfaces, other than the left-hand surface of lens 2. The rms increased by a factor of about a half, and the PVr increased by a factor of two-thirds, both of which imply that there has been a small amount of deformation to the lens. However, given the lens is a rigid body and lens 2 is a thin lens with a centre thickness of 15 mm, if there were any significant deformation this would likely affect both surfaces. As the pre- and post-mounting values for the right-hand surface were very similar, like all other tested surfaces, there were no worries that the performance of the camera would be significantly degraded.

Lens	Surface	PVr (μm)		RMS (μm)	
		Unmounted	Mounted	Unmounted	Mounted
1	L	0.103	0.098	0.016	0.015
	R	0.670	0.621	0.083	0.092
2	L	0.231	0.385	0.037	0.055
	R	0.174	0.162	0.030	0.029
3	R	0.573	0.605	0.080	0.094
5/6	L	0.925	0.835	0.119	0.109
	R	0.872	0.879	0.182	0.183

Table 4.2: The PVr and rms values for the aberration maps of all of the lens surfaces that could be tested with the interferometer. The surfaces were tested before and after mounting the lenses to see judge how much stress the mounts were placing on the lenses.

As a comparison between the new and old mounting schemes, the aberration maps for the right-hand surface of lens 1 are shown both mounted and unmounted for both the old assembly and the new assembly. As stated in Section 4.1.1.1, the lens 1 that was used in the original assembly was badly cracked and had to be replaced, therefore there are two different lenses used to create this figure. It can be seen that the old mounting scheme put a lot of stress on this lens, as proved by the very large aberrations that were measured when it was mounted — approximately a factor of 40 higher than for the other measurements — this could suggest a possible cause for the damage to this lens.

It can be seen that there is a large amount of astigmatism, which seems to be visible even after removing the force from the lens mount, although this can not be known for certain as it was not interferometrically tested before it was mounted. As the lens was damaged during storage and also measured to have been deformed by the mount it seems reasonable to propose that the mounting caused the damage, although this cannot be proven. If this did indeed happen, it must have been the case that the shape of the surface of the lens mount was not perfectly matched to the curvature of the lens. This could have been due to an error during the manufacturing process that was not spotted, damage to the lens mount, or possibly a foreign object




Figure 4.12: The different surface forms of the right-hand surface of the different, 86 mm diameter, lens 1s. Shown is the lens used in the old mounting scheme (*top*) and the one used in the new scheme (*bottom*), with both of them shown unmounted (*left*) and mounted (*right*). Note the different z-scale required for the old, mounted lens 1, with all units being μm .

in between the lens and its mount. Any of these could have caused the lens to be held more tightly in one direction than the other when it was clamped in place, causing the astigmatism and possibly the fracture.

The same surface with the new mounting scheme shows that the aim of mounting the lens stress-free appears to have been successful, with lens having the same form both when free and mounted. During the new alignment procedure the lens is clamped in place to prevent it from moving as the mount is aligned and adhered in place, but this appears to not have caused any permanent deformation.

4.3.2 Lens radius of curvature

As described in Section 4.2.1.1, when mounting the lenses the value of the radius of curvature of one of the surfaces needs to be accurately known. The new lenses 1 and 4 came with manufacturer's measurements of the radius of curvature, but no measurements of the final lenses were available for the original set, only their nominal values. Whilst the nominal value for each surface is accurate enough for the alignment given the requirements that were placed on them, the curvature were still measured, where possible, to verify its magnitude.

4.3.2.1 Method

When the lens is located at the cat's-eye position, the centre of curvature of the transmission sphere lies on the lens surface, whereas when the lens is located at the confocal position it is coincident with the lens centre of curvature. Therefore, if the lens is moved between these two positions then the distance it has travelled can be a very precise measurement of its radius of curvature if measured by a CMM on the optical bench. The interferometer software can then use the measured wavefronts to correct for any small inaccuracies in the positioning of the lens, meaning the positioning errors should be reduced far below the measuring errors.

The setup and method followed was the same as detailed in Section 4.3.1.1 for measuring the surface form of the lenses. When the lens was located at the cat's-eye position however, it was precisely aligned and the wavefront measured before a Faro Arm CMM was used to measure its position. This measurement was done by probing the chuck holding the lens rather than the lens itself, to avoid damaging the optical surfaces. When the lens was then moved to the confocal position, the procedure of taking an interferogram and measuring the position of the stage was repeated. The distance between the two position measurements was then calculated and fed into 4Sight along with the inteferograms, to produce the final estimates of the radii of

curvature.

As with the measurements of the surface form, the radii of curvature of lens 4 and the left-hand surface of lens 3 were too large to be measured using the available transmission sphere.

4.3.2.2 Results

The measured values of the radii of curvature of the lenses are shown in table 4.3. These are shown with and without the correction from the interferometer software, which was typically of the order $10\text{ }\mu\text{m}$, and they are compared to the nominal values and, where possible, the manufacturer's measurements.

The measurements match the corresponding manufacturer's measurements very well, with the surfaces of lens 1 just $48\text{ }\mu\text{m}$ and $3\text{ }\mu\text{m}$ different, which compares favourably to the precision of the Faro arm of about $30\text{ }\mu\text{m}$. Unfortunately, both surfaces of lens 4 were too flat to be measured with the available transmission sphere, and past measurements for the old lenses were not available for comparison.

Lens	Surface	Nominal	Manufacturer	Measured	Interferometer
1	L	-67.0498	-67.056	-67.056	-67.053
	R	-258.9939	-259.057	-259.033	-259.009
2	L	-749.3723	—	-749.685	-749.686
	R	106.4174	—	106.418	106.423
3	L	293.1195	—	—	—
	R	124.7324	—	124.743	124.743
4	L	288.8594	288.832	—	—
	R	487.2714	487.362	—	—
5/6	L	119.6143	—	119.566	119.565
	R	-97.1095	—	-97.085	-97.097

Table 4.3: The radii of curvature for each surface of all the lenses in millimetres. Shown are the nominal values from the optical design, the fabricated lens' actual diameter as measured by the manufacturer, the raw measurement from the CMM, and finally the best estimate when taking into account the interferograms.

4.3.3 Whole camera

4.3.3.1 Method

As the interferometer beam can not simply be reflected off the camera's surface to measure the wavefront aberrations, the setup required is more complex than in previous sections, as seen in figure 4.13. This time, a collimated beam produced by a transmission flat is used, with the camera placed facing the beam. The camera focusses the light to a point, and a return sphere is used to reflect the light back through the camera and to the interferometer. This is done by positioning the return sphere such that its centre of curvature is coincident with the focal point of the camera.




Figure 4.13: The setup used to interferometrically test the optical quality of the SWIFT camera. The camera is placed on its side on a turntable to rotate it about the x -axis, and it can be seen that the camera has been tilted to rotate the camera about the y -axis. For scale, the width of the interferometer's beam is 300 mm.

As the optical quality of the camera was to be measured across its field of view the camera needs to be rotated with respect to the interferometer to simulate the varying field angles. Therefore, the camera was mounted on a piece of breadboard which could have supports inserted underneath at one end to tip the camera down to

look towards the table as is the case in figure 4.13. This board was then placed on a platform which could be rotated by a turntable to allow the camera to be angled in the other direction.

After the angle of the camera with respect to the interferometer has been set, the return sphere has to be aligned to the camera. This is done precisely by mounting it on an $x - y - z$ stage, which, to minimise the amount of movement required between measurements at different field angles, was mounted on the same platform as the camera so that they rotate together.

The entrance pupil of the cameras in the SWIFT Spectrograph is 46×23 mm in size and 49 mm in front of the first surface of the camera. To simulate this, a pupil mask that consisted of a cylinder that fitted over the nose of the camera with an ellipse cut out of the other end was fabricated. The length was made such that the mask was correctly positioned in the z -direction.

As the back focal distance of the camera and the radius of curvature of the return sphere were known to be about 55 mm and 39 mm respectively, the return sphere was placed in approximately the right position by eye at first. This was sufficient to see a return signal so the $x - y - z$ stage of the return sphere could be used to align it using the alignment camera and then the live interferogram.

Once the whole system was aligned then the same procedure as described in Section 4.3.1.1 was used to set up the interferometer and take the measurements. The wedge factor was 0.5 again due to the double-pass nature of the experimental setup and the tip, tilt and defocus Zernike modes were subtracted from the final wavefront error maps as these corresponded to the residual x -, y - and z -misalignments of the return sphere.

This method for taking an individual measurement was used for an array of measurements across the $7 \times 14^\circ$ field of view. Measurements were taken every 3.5° from -7° to 7° on the x -axis of the field-of-view and from -14° to 14° on the y -axis. The

camera was placed on its side so the turn table was used to take the 9 measurements for a given angle on the x -axis, with this repeated for each of the 5 different values on the x -axis. The camera needed to be turned onto its other side in order to take the measurements along the negative x -axis.

The wavefront error maps were analysed to find their residual rms values, which were used as a means to compare the optical performance across the field of view. They could also be used to compare between the different assemblies and the predictions derived from the optical design using the Zemax software that was used during the design of the camera. The maps were also decomposed into Zernike polynomials to better understand what was causing any aberrations.

Since the pupil is elliptical in shape, the Zernike decomposition could not be done via a strict mathematical decomposition without stretching every one of the circular terms in one direction to match the shape of the pupil. Therefore a model consisting of a linear superposition of 8 of the first terms was fit to the maps. The terms used were piston, tip, tilt, defocus, astigmatism and coma.

To see the effects of the camera on its images, the PSF was estimated at each tested field point from the wavefront error maps. This was done by estimating the pupil function of the camera using equation 1.2, where the amplitude was assumed to be constant across the pupil and the phase was calculated by converting the wavefront maps from waves into radians. With the pupil function found, equation 1.1 could be used to estimate of the PSFs, which were binned up to the same scale as the Nyquist sampling of the SWIFT detectors for the best possible representation of the camera in the SWIFT Spectrograph.

4.3.3.2 Discussion of old assembly

To judge the optical performance of the camera across its original field of view, before disassembly, the rms wave front errors of both the measured wavefronts and those




Figure 4.14: The optical quality of the camera across its field of view, in degrees, derived from its measured wavefront (*top*) and its predicted wavefront (*bottom*). The rms of the wavefront are shown (*left*) alongside the rms of the residuals left over after the subtraction of astigmatism and coma (*right*). All z-scale units are μm .

predicted using Zemax are presented in figure 4.14. It can be seen that the measured on-axis performance is very good, and even exceeds what was expected from the optical design. The portions of the field furthest off-axis are worse than expected however, with the bottom edge of the field of view having an rms approximately 50% higher than expected from the Zemax model.

To try and determine the cause of the discrepancies between the measurements and predictions, the wavefront maps themselves were analysed, and they can be seen in figure 4.15. Each of the subplots represents one of the measurements that were

taken at evenly spaced intervals across the entire field of view. The z -scales vary from plot to plot so that the features of each map can be discerned, meaning that they are useful for a qualitative comparison of the types of aberrations rather than a quantitative comparison of the optical quality, which is better seen in figure 4.14. The equivalent, predicted, maps from Zemax are shown in figure 4.16. Again, the z -scales differ with these maps so that the shapes of both sets of maps can be directly compared to see how the types of aberrations differ.

There are a lot of visual similarities between the measured and the predicted wavefront maps. In the centre of the field of view they both have circularly symmetric patterns with a central ‘peak’ of deviation from a perfectly flat wavefront, which further from the centre develops into a bar-like shape tangential to the centre. However, at the farthest edges of the field, at the top and bottom, there is the biggest difference in the appearance of the wavefront. At the top of the field the bar-like shape lies at the bottom of the pupil rather than through the centre, whereas at the bottom of the field it runs through the centre, but has a much smaller amplitude. The same bar-like feature at the sides of the field of view is twisted into a spiral pattern that is well predicted by Zemax, albeit with a slightly lower relative amplitude.

To analyse the wavefront aberrations more quantitatively, the decomposed Zernike polynomial terms are represented in figures 4.17 and 4.18 for the measured and predicted wavefronts respectively.

When looking at the astigmatism measured in the camera, the behaviour of both the oblique and vertical terms in the lower half of the field is close to being as predicted, but with the vertical terms tending to be larger than expected. This appears to be the dominant cause of the poor image quality at the bottom of the field. At the top of the field though, there appears to be much less astigmatism, just a small amount of the vertical term that is less than suggested by the Zemax generated wavefront maps. These differences in astigmatism could be influenced by the forms of the




Figure 4.15: The interferometrically measured wavefront maps of the original camera assembly across its field of view. The top-left plot represents the wavefront errors in the top-left of the camera's field of view and so on as seen from the detector's point of view. From left to right the plots represent field angles of -7° , -3.5° , 0° , 3.5° and 7° , whereas from top to bottom they represent field angles of 14° , 10.5° , 7° , 3.5° , 0° , -3.5° , -7° , -10.5° and -14° . Each plot has a different z-scale to enable a quantitative comparison of the features in the wavefronts across the field of view, but the values for the peak-to-valleys of the plots range from $0.5\ \mu\text{m}$ at the centre to $5\ \mu\text{m}$ in the bottom corners.




Figure 4.16: The predicted wavefront maps of the camera across its field of view created using the Zemax optical design software. The top-left plot represents the wavefront errors in the top-left of the camera's field of view and so on as seen from the detector's point of view. From left to right the plots represent field angles of -7° , -3.5° , 0° , 3.5° and 7° , whereas from top to bottom they represent field angles of 14° , 10.5° , 7° , 3.5° , 0° , -3.5° , -7° , -10.5° and -14° . Each plot has a different z-scale to enable a quantitative comparison of the features in the wavefronts across the field of view, but the values for the peak-to-valleys of the plots range from $0.26\text{ }\mu\text{m}$ at the centre to $1.9\text{ }\mu\text{m}$ in the each corner.

lenses, especially if the measurements from lens 1 in figure 4.3.1 are repeated across more of the lenses. Unfortunately this was the only lens to be measured, and an analysis of the aberrations it was found to have, has not yet been performed using Zemax.

As expected, there is much less coma than astigmatism in the wavefronts, but the magnitude is higher than predicted and with different patterns formed across the field. The largest discrepancy is seen at the top of the field, where coma appears to be the dominant type of aberration causing the poor image quality for this part of the field. It is not known what the cause of this coma is. It can be produced by a decentring of one or more of the lenses, but this scenario is hard to achieve given the old mounting scheme for the camera, and could again be influenced by poor surface form of one or more lenses.

The rms of the residuals that were left after the Zernike decomposition are shown in figure 4.14 alongside the rms values of the wavefronts. They are very low, meaning that the majority of the poor optical quality is indeed caused by the astigmatism and coma discussed above.

The PSFs that were estimated from both the measured and predicted wavefronts are shown in figures 4.19 and 4.20. As expected from the analysis of the wavefronts, the measured PSFs at the centre of the field of view compare very well to the predictions, but deteriorate quicker than predicted further away from the optical axis. It can be seen that the PSFs appear to be stretched vertically. This arises due to the elliptical entrance pupil used in this work, meaning that the camera has a different horizontal compared to vertical resolution.

It is difficult to know why the optical quality is particularly bad at the extremes of the field, especially given the damage to lens 1. The camera was aligned by Dr. Matthias Tecza, and he felt that this camera was the best aligned out of the three that also included the cameras currently in the SWIFT Spectrograph, so it is thought that




Figure 4.17: The Zernike decompositions of the measured wavefronts in waves. *Top left:* oblique astigmatism; *top right:* vertical astigmatism; *bottom left:* vertical coma; *bottom right:* horizontal coma. Z-scale units are μm .




Figure 4.18: The Zernike decompositions of the predicted wavefronts in waves. *Top left:* oblique astigmatism; *top right:* vertical astigmatism; *middle left:* vertical coma; *middle right:* horizontal coma. Z-scale units are μm .




Figure 4.19: The estimated PSFs of the original camera assembly across its field of view derived from the measured wavefront aberration maps. The top-left plot represents the PSF in the top-left of the camera's field of view and so on as seen from the detector's point of view. From left to right the plots represent field angles of -7° , -3.5° , 0° , 3.5° and 7° , whereas from top to bottom they represent field angles of 14° , 10.5° , 7° , 3.5° , 0° , -3.5° , -7° , -10.5° and -14° . Each plot is 10 pixels across and has a z-scale normalised to each PSF to enable the change in shape across the field of view to be easily seen.




Figure 4.20: The predicted PSFs of the original camera assembly across its field of view derived from the measured wavefront aberration maps. The top-left plot represents the PSF in the top-left of the camera's field of view and so on as seen from the detector's point of view. From left to right the plots represent field angles of -7° , -3.5° , 0° , 3.5° and 7° , whereas from top to bottom they represent field angles of 14° , 10.5° , 7° , 3.5° , 0° , -3.5° , -7° , -10.5° and -14° . Each plot is 10 pixels across and has a z-scale normalised to each PSF to enable the change in shape across the field of view to be easily seen.

the alignment of the lenses shouldn't be the problem. It was hoped that the results obtained after remounting the lenses with the new mounting scheme would give an insight into what was causing these aberrations, but that comparison was not possible due to the problems discussed in Section 4.3.3.3.

4.3.3.3 Discussion of new assembly

When the camera had been fully assembled and was interferometrically tested for the first time it was found to have much larger aberrations on-axis than expected, as it was thought that each of the lenses had been aligned within their tolerances. The first measured wavefront can be seen in figure 4.21, and it can be seen that there is a large amount of vertical coma. This suggested that one or more of the lenses were vertically misaligned, with lenses 1 and 3 expected to be particularly sensitive in that regard according to the Zemax model of the camera.

Therefore, a process of trial and error was used to improve the on-axis performance using estimates for the misalignments of the lenses found by comparing the wavefront to the expected wavefront on Zemax. This process took more iterations than initially expected, at 20, before the wavefront was comparable to the predicted and old measurements. This wavefront can be seen in figure 4.22 along with measurements at the top and the bottom of the field. The equivalent wavefronts from the old assembly as well as those predicted using Zemax are also shown as a reference.

Whilst the on-axis wavefront is similar to the old assembly and the predictions, the off-axis performances were very different. The bottom of the field of view showed astigmatism like the old assembly, but with a much greater amplitude, whilst the top of the field had better optical quality than either the old measurements or the predictions.

These differences were assumed to be due to misalignments due to the trial and error procedure used to get the final alignment, but the results were more discouraging




Figure 4.21: The first on-axis wavefront map of the newly assembled camera. It is far from axisymmetric, with significant coma visible, which indicates some sort of misalignment. Z-scale units are μm .

when the camera was rotated onto its side before repeating the measurements. These results are shown in figure 4.23, where it can be seen that the interferogram changes when the camera is rotated. The only explanation for this change was that the lenses were not held firmly enough and were therefore moving when the gravity vector was changed. This also explains why the final alignment of the camera was more difficult than expected, as the lenses were not in the positions they were thought to be in, meaning that the interferometric tests had to be suspended until there was a better understanding of the actual positions of the lenses.




Figure 4.22: A comparison of the wavefront map of the new assembly (*right*) with old assembly (*centre*) and the Zemax model (*left*). The measurements were taken at different positions along the y-axis up the centre of the field: 14° (*top*), 0° — on the optical axis (*middle*) 0° , and -14° (*bottom*). Whilst each plot at the same position shares the same z-scale in μm , this is varied between positions to facilitate a comparison between the different assemblies and the model.

4.4 Problems

4.4.1 CMM measurements of lens alignment

As the optical quality of the new assembly was lower than expected and it seemed that the lenses were moving when the camera was rotated, it was clear that it was




Figure 4.23: The on axis interferogram with the camera standing upright and on its side. Varying gravity vector should not affect interferogram, so therefore there must be a movement of at least one of the optical surfaces. Z-scale units are μm .

not known where the lenses were in their mounts accurately enough to build a camera with sufficient optical quality. Therefore it was decided that it was necessary to use a CMM to measure the lenses in their mount rings with different gravity vectors to find out their positions and movements.

4.4.2 Method

The flatbed CMM used to align the lenses was not available at the time, so these tests were carried out using the Faro Arm CMM, despite its lower precision. Given that the tolerances on the radial positioning of the lenses were $40\ \mu\text{m}$ and the precision of the arm CMM should be better than $30\ \mu\text{m}$, the poor optical quality of the camera implied that the misalignments should be sufficiently large to detect with the arm CMM, but not necessarily to fully characterise.

The lens was clamped in place upright on the optical table so that the CMM could

access both optical surfaces. The gravity vector acting upon the lens was changed by rotating the lens mount about the optical axis, with six mounting holes meaning that the lens was mounted in six different positions by rotating it by 60° between each measurement.

To take the measurements of the lens' position, firstly the coordinate system had to be defined, which was done relative to the mount ring. The origin was set to be on the reference plane of the mount and in the centre of its ring. The z -axis was set to be along the desired optical axis and the y -axis set to point towards one of the holes in the mount. This was done by probing the plane of the mount reference surface and the edge of the mount to position the origin and z -axis. Then, the hole that defined the y -axis was probed and the axis set to point towards its position.

Next, the position of the lens' optical axis was found by measuring both of its surfaces. A sphere was fitted to the measurements of each surface, with the line connecting them calculated. As the position of a line can be defined as its offset and angle from the axes, the co-ordinate at which the optical axis passed the x - y plane was measured as well as the angles between it and the z -axis in both the $x - z$ and $y - z$ planes. These represented the decentre and axial misalignment of the lens with respect to its mount.

The uncertainties are difficult to estimate precisely given that they depend on how the measurement is taken: ones taken close together that require little movement of the CMM arm as well as those taken with each joint of the arm as close to 90° as possible will have a higher accuracy than ones taken at opposite sides of the arm's range or with the arm stretched out more straight. To try and account for this, two different components of the uncertainties were estimated and then combined for the final estimate.

As the measurements on each optical surface are taken with the arm in almost the same position these should have a high precision, so the uncertainty on the centre

of each surface was estimated statistically using the fit of a sphere to the measured data. However, moving between the two surfaces requires moving the arm to a very different position so that it faces in the opposite direction, so there will be larger and possibly systematic differences between the two measurements. Therefore the uncertainty on the position of the centre of the surface on the side that the mount was probed was assumed to be accurate, but the position of the other surface was assumed to have an additional $30\text{ }\mu\text{m}$ uncertainty term.

Finding the final uncertainty on the offset and angle of the optical axis with respect to the co-ordinate axes depends on the geometry of the lens centres, such as whether they are on opposite side of the $x - y$ plane or not. It was calculated by creating simulated data of the surface positions with noise drawn from a normal distribution with a standard deviation equal to the estimated uncertainties on the positions of the centres. The offsets and tilts were then calculated and their standard deviations determined and used as the uncertainties on their estimates.

4.4.3 Discussion

The measured decentres and tilts of lenses 1-5/6 are shown in figures 4.24-4.28 respectively. As the measurements were made with respect to the lens mounts, the decentres and tilts should be the same for each measurement and ideally have a magnitude of zero. However, looking at the graphs reveals that though most if not all of the lenses trace circle-like loops, which is what would be expected if they were sagging in their mounts due to gravity, which is pulling the lens in a different direction for each measurement. The uncertainties are small enough that there can be confidence that the lenses are sagging, but are too large to measure the deviations from a circle that would indicate the sag varies between the different orientations.

The estimates of the sag of each lens are found by measuring the radius of the loops traced by the measured decentres and can be seen in table 4.4. There is evidence




Figure 4.24: *Left:* The decentre of the lens 1 optical axis as it is rotated through 60° increments about the z-axis. *Right:* The angle of the optical axis with respect to its desired position in the y- and x-directions, that is in the z-x and z-y planes respectively. For both plots each reading was measured with respect to the lens mount, with the gravity vector effectively changed for each reading. Therefore, any measured movement of the optical axis is caused by the mount not holding the lens firmly in place rather than an off-centre lens being moved with respect to the world as the mount is rotated.

for sag in every lens other than lens 1, where the results are equally well explained by large uncertainties relative to the magnitude of the sag. The other lens where the sag is within two errorbars of zero is lens 5/6, which is due to both lens surface radii of curvature being close together on the same side of the $x - y$ plane making it difficult to precisely determine the decentre of the optical axis. The other lens with large uncertainties is lens 4, which appears to be due to the large lens radii of curvature limiting the precision of the measurements, however the sag appears to be significant even with such large uncertainties. The magnitudes of the measured sags are between 30 and 300 μm , which would be large enough to explain the poor optical performance of the camera, especially if each lens is moving by a different amount.

Also shown in table 4.4 is the decentre of the lens that would be measured without it sagging due to gravity, which is given by the centre of each loop traced by the decentre measurements. Lens 1 appears to have a decentre within the tolerance of




Figure 4.25: *Left:* The decentre of the lens 2 optical axis as it is rotated through 60° increments about the z-axis. *Right:* The angle of the optical axis with respect to its desired position in the y- and x-directions, that is in the z-x and z-y planes respectively. For both plots each reading was measured with respect to the lens mount, with the gravity vector effectively changed for each reading.

40 μm , with lenses 4 and 5/6 consistent with zero due to the large uncertainties. Unfortunately, it appears that lenses 2 and 3 have a significant permanent decentre above the tolerance. It is not known when this decentre occurred, but the most likely time appears to be when the lenses were taken off the alignment jig. They were supposed to have been left there to cure for 7 days each, but time constraints meant that they had to be moved after only 3. This should have been long enough for the majority of the cure to have occurred, and extreme care was taken to minimise the chances of this happening, but given that their alignment was measured to be better than this on the jig this appears to be the most likely explanation. It should be noted however, that as long as the decentre is known to better than 40 μm then this can be taken out during the alignment of the lens mounts in the camera housings.

The measured tilts of each lens are shown in table 4.4, and were found by calculating the average tilt for each measurement. Lenses 3 and 4 had tilts reasonably consistent with zero, and lens 1 also had comparatively low tilt, but with no tilt tolerance to work towards it is hard to ascertain the impact as this would require




Figure 4.26: *Left:* The decentre of the lens 3 optical axis as it is rotated through 60° increments about the z-axis. *Right:* The angle of the optical axis with respect to its desired position in the y- and x-directions, that is in the z-x and z-y planes respectively. For both plots each reading was measured with respect to the lens mount, with the gravity vector effectively changed for each reading.

further work to analyse the optical design. Lenses 2 and 5/6 had much larger tilts, although once again lens 5/6 suffered from high uncertainties due to the geometry of the optical surface centres.

Some of the lenses appear to show evidence of precession of the tilt as the lens is rotated, most notably lens 3 as seen in figure 4.26. Lenses 1 and 5/6 also show signs of precession, but the higher uncertainties make this less clear, and could be a result of systematic uncertainties between the measurement of the the two optical surfaces. Lenses 2 and 4 are consistent with the tilt being the same for each lens orientation. Precession of the lens should not be possible if the lens is held securely, but the amount of movement of the angle of the optical axis does not really correlate with the amount of sag, which might be expected if these are both due the lens being able ‘wobble’ in its mount.

Before these tests were performed it was known that the uncertainties were higher than desired and might prevent a firm conclusion being reached, which appears to be the case. The fact that 4 of the lenses seem to have abnormal measurements, from




Figure 4.27: *Left:* The decentre of the lens 4 optical axis as it is rotated through 60° increments about the z-axis. *Right:* The angle of the optical axis with respect to its desired position in the y- and x-directions, that is in the z-x and z-y planes respectively. For both plots each reading was measured with respect to the lens mount, with the gravity vector effectively changed for each reading.

Lens	Sag (μm)	Decentre (μm)	Tilt ($^\circ$)
1	30 ± 20	10 ± 10	0.05 ± 0.01
2	60 ± 20	140 ± 20	0.13 ± 0.02
3	60 ± 10	60 ± 10	0.02 ± 0.02
4	300 ± 100	100 ± 100	0.02 ± 0.02
5/6	300 ± 200	200 ± 200	0.15 ± 0.10

Table 4.4: The parameters of the alignment of each lens in its mount. The sag is how much the lens appears to move in the $x-y$ plane as it is rotated, with the magnitudes of the decentre and axial tilt also shown.

their initial alignment to their movement in the lens mounts, seems to suggest that there is a problem with either the design or the mounting procedure, rather than an isolated mistake.

If the problem is due to the design of the lens mounts then it is unlikely to be the mounts themselves, as two different designs were tested with similar results. The common factors are the layer of cotton paper that adhered to the lens edge as well as the adhesive itself. If the method is at fault however it could be the insertion of the adhesive, which proved difficult to perform manually and would benefit from a move to an automatic dispenser to ensure a more even and repeatable delivery which would




Figure 4.28: *Left:* The decentre of the lens 5 optical axis as it is rotated through 60° increments about the z-axis. *Right:* The angle of the optical axis with respect to its desired position in the y- and x-directions, that is in the z-x and z-y planes respectively. For both plots each reading was measured with respect to the lens mount, with the gravity vector effectively changed for each reading.

be less likely to move the lens. The position of the lens was corrected and remeasured after applying the adhesive though, so the problem could stem from moving the lens before it had fully cured, as previously discussed.

Ultimately, the arm CMM has a lower precision than required and its human-controlled nature reduces the effectiveness of repeating measurements to lower the uncertainties. Therefore it was necessary to retake the measurements with more precise equipment when possible.

4.4.4 Further work

Following the measurements taken with arm CMM, further measurements were taken on the flatbed CMM that was used to mount the lenses. This was done at a later date to improve the precision of the measurements of the position of the lens in its mount. This work was carried out by Dr. Rob Barnsley, and is summarised here for completeness, as it provides insight into the effectiveness of the methods described in this chapter to align and mount lenses.

Lens	Centre (μm)		Decentre (μm)	Residual (μm)	Hysteresis (μm)	z-angle (arcmin)
	x	y				
1	2.2 ± 3.1	6.3 ± 0.9	46.3	2.1	2.2	3.92 ± 0.03
2	-4.2 ± 2.0	-8.8 ± 2.8	117.2	2.1	3.0	6.91 ± 0.02
3	7.1 ± 2.7	-60.0 ± 2.2	35.7	3.2	41.4	1.46 ± 0.04
4	8.6 ± 0.6	-20.2 ± 2.0	28.0	5.0	8.7	1.24 ± 0.02
5/6	-4.7 ± 3.4	-32.4 ± 0.8	234.8	14.5	6.0	8.27 ± 0.10

Table 4.5: The parameters for the optical axis of each lens, including: the centre and radius of the circle fitted to the decentre measurements along with the rms of each set of residuals, the hysteresis of the lens and the average angle between the optical axis and the mount z-axis. These measurements were performed by Dr. Rob Barnsley.

The method followed was similar to the one described in section 4.4.2, whereby by both optical surfaces of each lens were probed to measure the positions of their centres and therefore the optical axis. This was done with the mounts rotated at different angles to effectively vary the direction of the gravity vector, with the movement of the axis decomposed into a decentre and a tilt. To cross-check these measurements the edges of the lenses were also probed where possible to find the position of the mechanical axis, and the same analyses performed.

These new measurements were taken with respect to a world co-ordinate system, which means that when the decentre of each position is plotted, an off-centre but securely-fixed lens would trace a circle centred on the origin. The lens sagging by a constant amount in the mount due to gravity would then cause that circle to be offset with respect to the origin. Lastly, more complex, angle-dependent movement would result in deviations from a perfect circle and from perfect rotational symmetry of the results.

An example of the plots drawn using these new measurements is shown in figure 4.29, which depicts the data from lens 3. A circle was fit to the measurements of the decentre of the lenses where the radius showed how far off-centre the lens was mounted, and the centre of the circle showed the sag of the lens as it is rotated. It can be seen how the additional precision of the flat-bed CMM allowed for the mea-

surement of the more complex movement of the lens that isn't described by a simple, constant sag. This is shown by the measurements' variation in the distance from the circle's centre as well as the angle between adjacent measurements being significantly different from 60° .

The positional parameters of the lenses that were derived from the newer measurements are shown in table 4.5. The centre of the circle denotes the average decentre of the lens due to the lens sagging in the mount, and it can be seen that there appears to be significant sag for each lens. Furthermore, the y -component is significantly larger than the x -component for each lens, which is consistent with the lenses not being held firmly enough and allowing gravity to pull them downwards with comparatively little lateral movement. Interestingly, lens 1 appears to move upwards in the mount, which is an unexpected result that appears to be statistically significant, but no explanation is offered here to fully explain it.

All of the lenses have measurable decentres, although lenses 3 and 4 have magnitudes less than the $40\text{ }\mu\text{m}$ tolerances. Lens 1 has a decentre only just outside of this value, but lenses 2 and 5/6 seem to be very poorly aligned. As both of these lenses have both of their radii of curvature to one side of the $x - y$ plane this can be explained by the lenses being knocked whilst still on alignment jig, which would rotate the lens about the lower surface's centre of curvature. As this alignment was checked after inserting the adhesive then this movement will have happened after then, such as when the lenses were moved off the jig before the full recommended cure time was completed. This theory seems to be supported by the fact that lenses 2 and 5/6 also have the highest measured tilts as well.

Lenses 1 and 2 both have low rms residuals and hysteresis, which is defined here as the distance between two measurements with 360° of rotation between them, and should therefore be $0\text{ }\mu\text{m}$ if the lens mount is functioning correctly. Lenses 3 and 4 both have rms and hysteresis values that are higher, with lens 3 having more hysteresis




Figure 4.29: *Left:* The x and y positions of the optical and mechanical axes relative to the centre of the mount with the orientation of the mount rotated. A circle is fitted to the measurements to help estimate the decentre and sag of the lens. *Centre left:* The decentres of the lens measured with the gravity vector rotated. *Centre right:* The residuals from the the circles fitted to the decentres of the lens. *Right:* The angles between both the optical and mechanical axes and the z -axis of the mount. These plots were produced by Dr. Rob Barrisley.

and lens 4 having higher residuals. These values along with lens 3 also having a low level of rotational symmetry between the decentre measurements suggest that the sag may not be uniform for each flexure, which could be due to the lens sticking slightly between different orientations. Lens 5/6 has the highest residuals, but relatively low hysteresis. It would seem that the most likely explanation is that the flexures are sagging uniformly, without the lower performance of the lens 3 and 4 mounts, and that the higher residuals are due to the decentre being particularly large and higher uncertainties due to the lens optical surfaces geometry.

When the values in tables 4.5 and 4.4 are compared, it can be seen that the preliminary measurements appear to slightly overestimate the sag of the lenses. This means that the mounts are performing slightly better than initially expected in this respect, but only lenses 2 and 4 have results that are inconsistent with each other. The ratios of the decentres of the lenses are more mixed, with lens 1 being underestimated previously and lens 3 overestimated, but otherwise the results are consistent. The tilt measurements fare better, with every lens giving consistent measurements with no bias seen. The preliminary results therefore stand up reasonably well to the more precise measurements, with the overestimates of the sags presumably deriving from systematic errors in preliminary measurements, but there are not enough lenses to be able to tell that for certain.

To look further into the causes of the flexure sag a direct imaging analysis was performed on lens 3. This was done by positioning the mount in an upright position such that one of the pads was directly below the lens before the lens was clamped to the mount so that it could not move. The mount was then inverted so that the flexure that had the highest load was now positioned above the lens before the clamps were released. This ‘most-loaded’ flexure was imaged both before and after releasing the clamps using a Canon 350D DSLR with a 1:1 macro lens to detect any movement of the lens.

The results of the direct imaging are shown in figure 4.30, which gives a qualitative view of what happened. It can be seen that the bond between the pad and the cotton paper appears to be imperfect, with a small gap visible that opens up further when the clamps are released. This indicates that the bond has failed, which appears to be consistent with idea that the alignment of the lenses was compromised when they were moved off the alignment jig. It is also consistent with a problem occurring during the application of the adhesive. As there were difficulties inserting the glue into the bond line for some pads, and as the lenses had to be repositioned after moving during this process, it is possible that some of the bonds did not form properly in the first place, which meant that the lenses were not held securely enough when finally assembled into the camera.

Another option is that the bond failing is what caused the more complex aspects of some of the the lens' movements, with the uniform sagging of every lens due to another factor. The reasoning for this line of thinking is that it seems unlikely that every bond of every lens has failed causing the sagging to appear consistent for 3 of the 5 lenses, as if only a fraction of the bonds failed then this symmetric pattern seems unlikely.

4.4.4.1 Implications for HARMONI

The largest problem for the HARMONI cameras is that currently we cannot be certain that the current methods will allow the HARMONI lenses to maintain their alignment and be mounted securely, especially given that the lenses will be much larger and therefore more massive than the SWIFT lenses. This is compounded by the fact that it is not known exactly where the problem is initiated: the biggest area of concern is the application of the adhesive, but it could be due to an incomplete curing of the adhesive instead.

The application of the adhesive was already seen as an area where improvement

could be made due to the difficulty in inserting the adhesive into such a narrow bond line whilst maintaining its uniformity. It could be improved via the use of an automatic adhesive dispenser, which would be able to apply a more consistent pressure of the adhesive to give more repeatable results and avoid accidental peaks in the application pressure that could lead to a loss of alignment. This should alleviate the problems whereby the lens was being moved due to the pressure from the adhesive, and may mean that the current lens clamping methods are sufficient to hold the lens still. Although I am confident that this would deliver a much more reliable mounting procedure, additional testing would be required to determine whether this one change is enough to solve the problems discussed in this section.

One other uncertainty is that it is still unclear whether improving the adhesive application would solve the problems of the uniform sag in the majority of lenses, as this may be due to the adhesive itself or the cotton paper. If it is due to the adhesive, then it should help on this front as it would enable much greater control of the bond line thicknesses and allow consistent, thin bond lines. If the problem is due to the cotton paper then this may not be a problem as the work described in Section 3.6 suggested that its inclusion is unnecessary and not using it would simplify the mounting procedure and reduce the number of areas in which a problem can develop. The reason that the cotton paper was used in the tests in first place was due to the fact that these conclusions had not yet been drawn.

4.5 Conclusions

A procedure for aligning and mounting spherical lenses has been developed and tested for use with the department's flatbed CMM. The axial and radial alignments can be made to a precision of approximately $10\text{ }\mu\text{m}$, with the angular alignment is dependent on the lens shape, but equivalent to about $10\text{ }\mu\text{m}$ over the distance between the two

surface's centres of curvature. The procedure is currently restricted to spherical lens surfaces.

There were some difficulties encountered when inserting the adhesive into the bond line once the lenses had been aligned to their mount rings, and this caused the lens to move by up to 50 μm . Improvements will be needed to this process to prevent the alignment from being lost, which could include looking at different ways to clamp the lens in place as the adhesive is being inserted as well as switching from a manual to automatic adhesive dispenser. This latter change would allow a more controlled and repeatable pressure to be exerted on the adhesive so should avoid the high pressures that could move the lens.

A big problem was discovered when the newly assembled camera was interferometrically tested. It was found that the lenses were not as well aligned as when they were measured on the CMM and there were also problems with the lenses moving in their mounts by up to 60 μm . Lenses 3 and 4 had the biggest problems with the lenses moving in their mounts, with their movements less predictable and with more hysteresis and it was found that at least one bond had failed on lens 3. The other lenses suffered from sagging in the mounts, but had more uniform, predictable movements.

It is suspected that the misalignments could be caused by moving the lenses earlier than desired and before the adhesive had fully cured due to time restrictions, which may also explain why at least one bond failed. The difficulties in inserting the adhesive may also have resulted in bonds that were more likely to fail. The uniform sag of some of the lenses does not seem consistent with some of the bonds failing and occurs across both designs of flexures tested. It is thought that the most likely reason is the cotton, which should not be necessary in the HARMONI spectrograph cameras, however this is not yet known for sure.

The original camera assembly was found to have two broken lenses, with lens 1 severely damaged. The surface form of this lens was found to have a lot of astigmatism

when interferometrically tested, which means that the lens cracked due to it being under a large amount of stress, whereas the new lens mounting scheme was found to not increase the stresses on lenses. The original camera was interferometrically tested across its field of view, with good on-axis performance, but large aberrations at the extremes of the field of view. The aberrations were dominated by coma at the top of the field of view and by astigmatism at the bottom, but the cause was not conclusively identified, with the form of lens 1 suspected to have a negative impact on the image quality.




Figure 4.30: The failure of a bond on lens 3, with the different components labelled (*top*). The lens was positioned so that there was a pad directly beneath it before it was clamped in place and turned upside down. These images show the bond with the lens still clamped in place (*bottom left*), and after the clamp was released (*bottom right*). The height of the zoomed region is approximately 1 mm. Images taken by Dr. Rob Barnsley.

Part II

Ring Galaxies

Chapter 5

Introduction to collisional ring galaxy systems

Collisional ring galaxies, hereafter ring galaxies, are a type of peculiar galaxy, one of several to be characterised by their annular form. They have a well-defined, high-contrast, elliptical ring that has strong nebular emission lines and a usually knotty appearance. Theys & Spiegel (1976) proposed that ring galaxies could broadly be split into three groups: symmetric RN and RE rings that did and did not contain an off-centre nucleus respectively, as well as RK rings, that were more asymmetric, with a distinct knot in the ring. Ring galaxies share a similar morphology with resonant rings in galaxies (Buta, 1999) and polar ring galaxies (Whitmore et al., 1990), although the characteristics mentioned above can be used to identify the ring galaxies from resonant and polar rings, which tend to appear more diffuse and also surround a larger and more central nucleus.

The most well known example of a collisional ring galaxy is undoubtedly the Cartwheel Galaxy, which was also the first such galaxy discovered (Zwicky, 1941). The Cartwheel is shown in figure 5.1, where it can be seen that it fits neatly into the RN group of Theys & Spiegel (1976), but also visible are the ‘spokes’ aligned radially



Figure 5.1: Hubble Space Telescope images of the three types of ring galaxies proposed by Theys & Spiegel (1976): the RN ring Cartwheel galaxy (*left*), the RE ring Arp 147 (*centre*) and the RK ring IIZw28 (*right*). Image credit: HubbleSite and the Hubble Legacy Archive.

between the nucleus and the ring. These spokes are not typical of RN galaxies and appear to be perhaps unique, and are thought to be the result of gravitational instabilities in the disc where they were well reproduced in the simulations of Hernquist & Weil (1993). Also shown are the galaxies Arp 147 and IIZw28, which in turn can be considered archetypal of RE and RK rings respectively.

Ring galaxies are a rare type of object that have been thoroughly catalogued by Madore et al. (2009) using data from Arp & Madore (1987), but there are relatively few well-studied examples. When compiling the catalogue Madore et al. (2009) estimated that they make up just 0.01% of galaxies at any given time using a sample of 10^5 galaxies. From the known systems it has been observed that the overwhelming majority of ring galaxies have a companion galaxy that can be found within a few ring diameters (Few & Madore, 1986).

Different scenarios were suggested as the cause for this type of galaxy, including a galaxy colliding with an intergalactic cloud of HI by Freeman & de Vaucouleurs (1974), but the currently accepted theory is a collision between two galaxies, with one

passing almost perpendicularly through the the disc of another, very close to its centre. This idea was first suggested by Theys & Spiegel (1976) when they presented an observational study of a sample of ring galaxies with supporting numerical simulations performed independently by Lynds & Toomre (1976) and Theys & Spiegel (1977). This theory also explains both the ubiquitous companion galaxies as well as the rarity of the systems, due to the specific geometry required of the collision in conjunction with the instability of the ring structure — thought to have a time scale of order 10^8 yr.

These first simulations represented the most basic collision scenario, whereby a point-like mass travels perpendicularly along the rotation axis of a target disc galaxy. This causes an inwardly-directed radial force that causes the stars in the target galaxy, which were initially in circular orbits, to perform radial epicyclic oscillations. Stars at different initial radii have different epicyclic frequencies, which causes annular density waves to propagate outwards through the galactic disc, giving rise to a ring galaxy. This effect can be seen in the model from Toomre (1978) in figure 5.2, where the initial density wave is especially clear. The simulations show that there should be a series of these densities waves that form, but the later waves get closer together until they overlap each other, decreasing their contrast until they are erased by collisional diffusion. These later density waves can also be seen in figure 5.2.

The density wave that propagates through the target galaxy's disc is present both in the stellar and gaseous distributions, although it is not necessarily located in exactly the same place for each medium. According to the simulations performed in Gerber et al. (1996), if the companion is at least as massive as the target galaxy then dissipation in the gas causes its density wave to propagate slower than the stellar wave, however, if the companion is one quarter the mass of the target galaxy then the gaseous ring can be found on the leading edge of the stellar ring. The crossover mass was assumed to lie somewhere between the ratios of one and one quarter. The




Figure 5.2: The radial positions of 40 different particles from a numerical model of a symmetric collision. The initial density wave can be clearly seen starting just after the time of the collision and propagating outwards. This image is reproduced from Toomre (1978).

effect of this increase in density of the target disc's gas is to trigger a massive increase in star-formation rates that leads to the crisp, bright rings that are observed as well as the strong nebular emission lines. In the work of Romano et al. (2008), the star-formation rates for a sample of 15 ring galaxies are presented, with a median value of $6.86 \text{ M}_{\odot} \text{ yr}^{-1}$, which compares to an equivalent value of around unity for the Milky Way (Robitaille & Whitney, 2010).

Ring galaxies are a worthy subject of investigation not only because of their rare and interesting nature, but also due to the light they can shine on other areas of astrophysics. The geometry of the collisions that form them is the simplest possible, meaning that models of galaxy interactions can first be tested by simulating ring galaxy formation before being used on more complex scenarios. This is the case for every element of these models, from the modelling of gas dynamics to star-formation recipes. In addition, the ring galaxies are a unique environment to study star formation. Although there are other galaxy types that have similar star-formation rates,

these objects do not also exhibit the simultaneous quenching of star formation in other regions of the same galaxy that occurs in the under-dense regions of the target galaxy's disc. These are just a couple of reasons why it is important that we can test numerical models of galaxy interactions on ring galaxies and that the results match observations. The work in the next two chapters will focus on estimating the masses of the galaxies involved in forming ring galaxies and comparing these values to simulations and previous observational work. This in turn will help to help establish whether this aspect of ring galaxy formation is well reproduced by the models.

5.1 Mass ratios of ring galaxy companions

Given the role of gravity in the formation of ring galaxies one of the most important aspects of the interactions is the masses of the galaxies themselves, and especially the ratio of masses between the two galaxies. There have been questions in the past of mass ratios seen in observations and simulations, such as by Davies & Morton (1982). There cannot be disagreement if there is to be trust in the understanding of these interactions. The importance of this is not only in understanding ring galaxies, but in using their relatively simple geometry as a testbed for models of galaxy interactions in general.

5.1.1 Observational results

As previously mentioned, Davies & Morton (1982) were concerned that their measurements of the mass of one of the Cartwheel's companion galaxies was very small compared to the mass of the Cartwheel measured by Fosbury & Hawarden (1977), at 5-10%. Their explanation for the discrepancy between this value and ratios from simulations was that the star formation rate could make a low-contrast ring appear brighter in observations than the simulations. This idea still appears to

Galaxy	Mass ratio(s)
Cartwheel ^a	0.03-0.08
NGC 922 ^b	0.05
Arp 10 ^c	0.25
Arp 141 ^d	1.92
Arp 143 ^d	0.99
Arp 147 ^d	1.75
Arp 148 ^d	2.00
IIHz4 ^d	0.29
IZw45 ^d	0.05
IIZw28 ^d	<0.12
VIIZw466 ^d	0.40
NGC 2793 ^d	0.05

Table 5.1: The measured ratios of the companion mass to the target galaxy mass for a sample of ring galaxy systems.

^a Higdon (1996)

^b Wong et al. (2006)

^c Bizyaev et al. (2007)

^d Romano et al. (2008)

be true, even though the discrepancy does not seem as large today as it did in 1982, for a couple of reasons. Firstly, the possible companion that they measured (G2 in the notation of Higdon (1993)) was assumed to be an elliptical galaxy, but was later discovered to have large, though faint, spiral arms that likely lead to a higher than expected mass. Secondly, it is not known for certain which companion of the Cartwheel actually caused the rings to form, with Higdon (1996) finding evidence of a HI bridge connecting the Cartwheel with the more distant companion G3, which appears to be twice as massive as G2. Also, there have been studies which give slightly lower than the $40 \times 10^{10} M_{\odot}$ of Fosbury & Hawarden (1977), such as Higdon (1993) and Amram et al. (1998), who estimate it at $34.6 \times 10^{10} M_{\odot}$ and $(6-24) \times 10^{10} M_{\odot}$ respectively.

Further observational measurements of the mass ratio in ring galaxy systems are shown in table 5.1, but due to the rarity of ring galaxies, the number of systems where both members have been measured is very small. Wong et al. (2006) and Bizyaev et al.

(2007) both measured mass ratios of ring galaxy systems and found the companions to be very small compared to the target galaxies, especially when compared to the values used in the earliest simulations, however the larger sample size of Romano et al. (2008) reveals that some rings may even be formed by companions twice as massive as themselves. Therefore, it is important that simulations can produce realistic looking rings not only for small companions, but across the whole range.

5.1.2 Simulations

Simulations have been a vital tool in understanding ring galaxies from their use in first explaining their origin (Lynds & Toomre, 1976) through to the present day. As ever their use lies in the power they give to make changes in astrophysical scenarios that can then be compared with observations to further our understanding. As astronomical masses are notoriously difficult to estimate with precision, simulations can give valuable insight into the effects of varying the masses once they are developed and trusted.

Table 5.2 shows the mass ratios that were used in a sample of simulations across a range of times from the proposing of the collisional formation theory to the present day. Most of the published results did not discuss the effect of varying the mass ratios, nor state that their choices were based on observational evidence other than producing realistic-looking rings, especially before such evidence was available. Therefore the range of ratios used should reflect the range that gave results that appeared qualitatively similar to observations of ring galaxies.

The effect of varying the mass ratios was specifically investigated by Gerber et al. (1996) and Athanassoula et al. (1997), with the former reporting that ring galaxies could be formed down at least to a ratio of 0.1 and the latter finding that whilst a ratio of 0.01 was too low, they could be formed at 0.05. These are encouraging results, especially when compared to the values in table 5.1, where there appear to

Simulation	Mass ratio(s)
Lynds & Toomre (1976)	0.67
Theys & Spiegel (1977)	1, 2
Huang & Stewart (1988)	0.4
Appleton & James (1990)	0.33
Gerber et al. (1992)	1
Hernquist & Weil (1993)	0.26
Gerber et al. (1996)	0.1, 0.25, 1
Athanassoula et al. (1997)	0.01, 0.05, 0.1, 0.2
Mapelli & Mayer (2012)	0.5, 1
Fiacconi et al. (2012)	0.5
Smith et al. (2012)	0.2

Table 5.2: The ratio of the companion mass to the target galaxy mass used in various simulations covering the time period from the first simulations to the present.

be several known ring galaxy systems with mass ratios as low as 0.05. Additionally, they found that the expansion velocity of the ring is expected to increase for higher mass ratios alongside the width and intensity of the ring.

Unfortunately, there are several difficulties in obtaining values that are directly comparable with both other simulations and with observations. As the simulations use different galaxy models and numerical methods, it is uncertain how many of the differences between the simulations occur due to differences in the resolution, the numerical recipes and different galaxy models, especially when there are different values for the disc-to-halo mass ratios of the target galaxies. For example, Gerber et al. (1992) and Mapelli & Mayer (2012) give different results for whether an off-centre nucleus is a real effect or caused by a projection of the nucleus being displaced out of the plane of the target galaxy’s disc, a disagreement that is put down to the lower resolution of the earlier work. The small values used by Athanassoula et al. (1997) seem to be due to a large fraction of the target galaxy’s mass being in the halo, as the mass fractions of the companion when compared only to the disc are a factor of 5 higher.

To try and minimise the effects of the differences between the different simulations

they are treated in this work as a sample of possibilities rather than trying to compare them with each other. Therefore the ideal is not to have all of the simulations use the same parameters and get the same results, but to have each of the results in table 5.1 be consistent with at least one of the simulations and not have rings that can not be explained by the simulations that have already been published.

5.2 Ring galaxy kinematic modelling

The following basic model was used by Fogarty et al. (2011) to fit the kinematics of integral field spectroscopy of Arp 147. In this work the same methods were repeated to measure the rotational velocity of the ring and help estimate the mass of the ring galaxy.

5.2.1 Kinematic model

The apparent elliptical shape of ring galaxies can be explained by them being intrinsically circular objects viewed at oblique angles. If this is assumed to be the case then the inclination, i , of the ring can be estimated by measuring the semi-minor and semi-major axes of the ellipse, b and a :

$$\cos i = \frac{b}{a}, \quad (5.1)$$

which can be found by using parameters given by fitting an ellipse to an image of the ring.

If the kinematics of the ring galaxy are determined by measuring the Doppler shift of an emission line across the galaxy then the mean velocity represents the galaxy's recession velocity, v_{rec} . This can be used to estimate the distance to the galaxy using Hubble's Law,

$$d = \frac{v_{rec}}{H_0}. \quad (5.2)$$

This distance can then be used to calibrate the plate scale of the images to determine the physical radius of the ring galaxy, r , from a .

The line-of-sight velocities, v_{los} across the ring can then be decomposed into recession, rotation and expansion velocities of the galaxy, v_{sys} , v_{rot} and v_{exp} , by fitting the velocity map using the equation:

$$v_{los} = v_{sys} + \sin i [v_{exp} \sin(\theta + \phi) + v_{rot} \cos(\theta + \phi)], \quad (5.3)$$

where θ is the angle measured around the centre of the ring and ϕ is the position angle of the ring's major axis.

The mass enclosed within the ring, M_{ring} , can then be estimated using its rotational velocity and radius by assuming equal gravitational and centripetal forces acting on the ring, which gives a result of:

$$M_{ring} = \frac{v_{rot}^2 r}{G}, \quad (5.4)$$

where G is the gravitational constant.

5.2.2 Off-centre collisions

The probability of a companion galaxy performing a perfect head-on collision with a target disc is obviously very low, but simulations suggest that significant deviations from this geometry can still produce ring galaxies. For example, the results of Appleton & James (1990) indicate that rings can still form with an impact parameter of up to 40% of the disc radius if the collision is still normal to the disc and with angles of up to at least 60° if it is central, and Huang & Stewart (1988) suggested it may even be possible to form rings with an angle of 80° . Deviations from this perfect geometry

also explain the different classes of ring galaxies: with the most symmetric collisions forming RN galaxies, and then as the asymmetry of the collision is increased this leads to RE and then RK galaxies, which is because the nucleus is displaced from the centre of the galaxy (Mapelli & Mayer, 2012).

There is therefore specific interest in modelling the kinematics of off-centre collisions, and this was done analytically by Appleton & Struck-Marcell (1996). These analytic models of the kinematics of ring galaxies were developed by Struck-Marcell & Lotan (1990) following the same approach as Arnold et al. (1982) in modelling the collapse of collisionless dark matter in an expanding universe. This approach uses catastrophe theory (Poston & Stewart, 2014), whereby idealised kinematics of the ring galaxy predict cylindrically symmetric waves with an infinite density, which can be interpreted as the edges of the density waves.

These analytic models represent the simplest of any type of galaxy collision and were compared to observational results in Struck (2010) with some success. Here he was able to predict the ratio of the radii of successive rings in a sample of the most symmetric rings found in Madore et al. (2009), but had mixed results when looking at selected individual ring galaxies.

The cylindrically symmetric models predict expanding density waves, even if the first wave rarely develops into a formal caustic, meaning that there is no crossing of different stellar orbits. This does not represent a problem when comparing to observations as the compressing of orbits is still sufficient to trigger the increases in star formation that are seen. Indeed it simplifies the expected kinematics, where the velocities can be decomposed into a single expansion and rotational velocity for the whole ring.

The more general kinematic equations from caustic theory, where the collision is off centre, are derived in Appendix 2 of Appleton & Struck-Marcell (1996). Here a particle with a precollision radius, azimuth and velocity in the target galaxy disc of

q , θ_0 and v_{θ_0} is perturbed by the collision that modelled using the impulse approximation — it is given a velocity perturbation of Δv directed towards the position of the collision. This is assumed to cause it orbit to now perform epicycles with an amplitude, frequency and phase of A , κ' and ψ' with an orbital guiding radius and frequency of q' and ω_{cir} .

Therefore, post-collision, the particles move on orbits with a radius of:

$$r = q' - A q' \sin(\kappa' t + \psi'), \quad (5.5)$$

and an azimuth of:

$$\theta = \theta_0 + \omega_{cir}(q') t, \quad (5.6)$$

with time after the collision, t . Their velocities can be decomposed into radial and tangential components which represent expansion and rotational velocities, which are given by:

$$v_{exp} = -A q' \kappa' \cos(\kappa' t + \psi'), \quad (5.7)$$

$$v_{rot} = v_{rot0} + \Delta v_{rot}. \quad (5.8)$$

If the target galaxy is assumed to have a flat rotation curve and a uniform amplitude of Δv across its disc, then Appleton & Struck-Marcell (1996) show that equations 5.5-5.8 can be solved numerically using the following equations:

$$q' = q \left[1 - \frac{\Delta v}{v_{rot0}} \sin(\alpha) \right], \quad (5.9)$$

$$\kappa' = \frac{\sqrt{2} v_{rot0}}{q'}, \quad (5.10)$$

$$A = \frac{\Delta v}{v_{rot0}} \sqrt{\frac{\sin^2(\alpha)}{\left[1 - \frac{\Delta v}{v_{rot0}} \sin(\alpha)\right]^2} + \left[\frac{v_{rot0} \cos(\alpha)}{q' \kappa'}\right]^2}, \quad (5.11)$$

$$\sin(\psi') = \frac{q' - q}{q' A}, \quad (5.12)$$

$$\omega = \frac{v_{rot}}{q'}. \quad (5.13)$$

From the position of each particle α is the angle subtended by the distance between the galaxy centre and the collision centre.

Therefore the kinematics of a ring galaxy formed from an off-centre collision can be modelled with the addition of just three inputs:

1. The original rotational velocity of the target galaxy,
2. The impact parameter of the collision,
3. The velocity perturbation caused by the collision, which encapsulates information about the collision speed and companion mass ratio.

5.2.2.1 Approximations and limitations

There are a number of approximations that are included in this model, which are necessary to give such simple equations to describe such complex interactions. Some have been mentioned above, such as the impulse approximation, the flat rotation curve and the uniformity of Δv . The impulse approximation relies on the interaction timescale being much smaller than the rotational timescale, which along with the flat rotation curve, seems to be a reasonable approximation for the majority of spiral galaxies. The magnitude of Δv across the disc depends on the morphology of the companion, with the uniform behaviour chosen by Appleton & Struck-Marcell (1996) to give the simplest solutions.

Other approximations include:

1. That the perturbation velocity is small enough that the orbits can be well described by epicyclic motion,
2. The kinematic approximation, that is that there is no self gravitation or dynamical friction after collision,
3. There is no significant movement perpendicular to the plane of the disc.

Approximation 1 gets worse with larger companion mass ratios, but should hold as long as the ratio is significantly below 1. Approximation 2 is similar to restricted simulations performed by Hernquist & Weil (1993), where it was concluded these could capture most of the essential dynamics, if not the details, such as spokes seen in the Cartwheel Galaxy. Approximation 3 only really holds for a normal collision, as oblique collisions are predicted to significantly warp the disc (Fiacconi et al., 2012), and also precludes the modelling of target galaxies with bulges.

It was decided that these limitations were worthwhile in order to minimise the complexity of the models being performed meaning that simple analytical simulations can be used in place of complex numerical ones. The reason behind this decision is that this work uses this model to get a feel for the broad differences in the behaviour of the kinematics of symmetric versus off-centre collisions to test and possibly improve the method in Section 5.2.1, not to make precise and absolute predictions.

5.2.3 Mass outside the ring

Using the results from Romano et al. (2008), where a small sample of ring galaxies were deeply imaged to reveal the disc exterior to the ring. They found that the rings had preferentially expanded to about halfway through the disc, suggesting that this is likely due to selection effects, as early on they are hard to identify as ring galaxies

due to their small size, and at later times they are too faint to be seen if they have expanded too far out. Their sample size of 9 systems gave the average radius of the disc out to 25th magnitude as 1.9 ring radii, with a standard deviation of 0.4.

To estimate how much of a ring galaxy's disc mass lies outside of the ring, it was assumed that it continued to have a flat rotation curve past this radius. As the mass of the galaxy internal to the radius of the ring was calculated using equation 5.4, a correction was made to the value of r being used to estimate the total mass. As the disc radius was 1.9 ± 0.4 times its ring radius this leads to the estimate of the mass of the galaxy, M_{gal} being given by:

$$M_{gal} = 1.9M_{ring}. \quad (5.14)$$

This correction is applied blindly to the ring galaxies in this work and is based on a small sample size, meaning that it is very uncertain. It is used to try and get a handle on the total disc mass of the target galaxy without assuming there is no mass outside of the ring. The assumption of a flat rotation curve means that the estimate of the mass outside of the ring is likely to have a greater dark matter fraction than for that of the mass inside the ring. This makes the mass estimate more difficult to compare to that of the companion galaxies, which were made using stellar population models and therefore do not include any dark matter. This discrepancy over the sensitivity of the two methods to dark matter was considered when estimating the uncertainties of the mass ratios.

5.3 Stellar population masses

Stellar evolution theory can be used to estimate the amount of light emitted by stars as a function of their mass, so these relations can therefore be reversed engineered to estimate how much mass there is, based on the amount of light. This is relatively

straightforward for a single coeval group of stars: a single stellar population (SSP), but is made much more complex when an entire galaxy is considered. This is because a galaxy is made of many different populations of stars, the ratios of which is not known a priori.

Finding the ratios of the different SSPs in a galaxy is complicated by degeneracies such as that between older and more metal-rich stars (Faber, 1972). However, it is possible to estimate them using SSP libraries and numerical methods, such as pPXF (Cappellari & Emsellem, 2004), (Cappellari, 2017). If the ratio of different SSPs is estimated, then the implied mass-to-light ratio of the entire galaxy, M/L_{tot} , can be found using:

$$M/L_{tot} = \frac{\sum_i w_i M_i}{\sum_i w_i L_i}. \quad (5.15)$$

Here, w_i represents the mass fraction of each template, i , and its associated mass and luminosity, M_i and L_i . This equation was taken from Cappellari et al. (2013a).

5.3.1 Single stellar population models

Calculating the spectral energy distributions of single stellar populations involves bringing together two main ingredients: the theoretical properties of a group of coeval stars, and empirical libraries of stellar photometry and spectroscopy. The models used in this work were calculated by Vazdekis et al. (2010).

The theoretical calculations come in the form of stellar isochrones (Girardi et al., 2000), which are the curves on a plot of luminosity against temperature for a single stellar population. One of the first pieces of input physics for these models is their chemical composition, their helium and heavy element mass fractions, as well as the opacities, which both affect how efficiently the stars are able to radiate energy. Also included are: the equation of state for the stars, the rates of the nuclear reactions

that release all of the energy and prescriptions to describe radiative transport by convection. These models are then calibrated by comparing the solar model with observational data from the Sun.

Stellar photometric libraries are used to obtain estimates for the colours and magnitudes of SSP models, with spectroscopic libraries used to predict their spectral properties. Vazdekis et al. (2010) use empirical libraries such as Sánchez-Blázquez et al. (2006) and Cenarro et al. (2001).

The end result is a collection of spectra for coeval groups of stars with known masses as a function of the initial mass function (IMF), age, metallicity and alpha abundance. These models are therefore a powerful tool for studying the stellar populations of galaxies, allowing investigation of properties such as the star formation history and the total stellar mass. The stellar mass appears to be the one of the more robust property that can be estimated from spectral fitting, with uncertainties of around 0.3 dex (Conroy, 2013).

5.4 Scope

In Chapter 6, IFS data of three ring galaxies are reduced and their kinematics measured. The model in Section 5.2.1 is then fit to the data in order to estimate the dynamical mass within the ring. Simulations based on the models in Section 5.2.2 are performed to derive improved models that are then fitted to the kinematics themselves. Final mass estimates for the ring galaxies are produced and their uncertainties discussed.

In Chapter 7, the companions of two of the ring galaxies studied in Chapter 6 are investigated to try and find their masses. With a simple dynamical mass estimate not measurable, stellar population models are fit to their integrated spectra to estimate their stellar mass-to-light ratios. These mass-to-light ratios are then combined with

the measured magnitudes and redshifts of the galaxies to estimate their stellar masses. These mass estimates are compared to those from the ring galaxies in Chapter 6 with the results discussed in the context of mass ratio estimates from the literature, including both simulated and observation work.

Chapter 6

Determining the mass of ring galaxies

6.1 Observations

I successfully applied for observing time with the MUSE IFS on the VLT during its Science Verification (SV) run, and was allotted 1.5 hrs with the Program ID 60.A-9332(A). This was used to observe three different ring galaxy systems from Madore et al. (2009), of which two have been fully analysed, as detailed in table 6.1.

Galaxy	RA (h m s)	Dec (° ' ")	Exp. Time (s)
Arp 146	00 06 44.3	-06 38 08	1000
AM0058-220	01 00 46.0	-21 43 30	1000

Table 6.1: The sample of ring galaxies observed during the MUSE Program ID 60.A-9332(A) that were analysed for this work.

Due to the small amount of observing time, it was not possible to observe a large sample of galaxies. Instead the focus was to get data on three different systems with sufficient signal-to-noise to allow both a study of the emission line kinematics of the ring galaxy, and to perform stellar population modelling of the integrated light of

the companion. The criteria that were considered when choosing which systems to observe were the surface brightness, having a clear ring, being well placed in the Chilean night sky during the SV period and having a single companion galaxy close enough that the whole system could fit in half of the MUSE field of view to allow ‘nodding on IFU’. Most of these galaxies are not well studied as they have less well-defined or incomplete rings and smaller apparent sizes when compared to the more famous ring galaxies.

Two 500 s exposures were taken of each target, and the telescope was positioned such that the target was in opposite corners of the field of view for each exposure to assist with the sky subtraction. There were also a full set of calibration frames taken to be used in the datacube reduction, including observations taken of standard stars for photometric and telluric calibrations.

During the MUSE SV run there were also observations of the Cartwheel Galaxy in Program 60.A-9333(A), lead by Philippe Amram. As this is the most well-studied ring galaxy system, due to its size and brightness, my analyses were also performed on this dataset to compare my results to previously derived properties in the literature. It was also useful as a test on one of the more symmetrical ring galaxies, indicative of having a simpler interaction geometry, meaning that the simple kinematic model should be a better fit if the effects of the collision’s off-centre nature are smaller.

The data were reduced using version 1 of the MUSE Data Reduction pipeline with EsoRex following the method detailed in the MUSE data reduction cookbook. The first step was to combine raw calibration frames into higher signal-to-noise ratio (SNR) master frames, which was done for the bias and lamp flat field frames. The flat field master was used to trace the positions of each slice on every detector.

Next the wavelength solution for each detector was found using frames taken using Neon, Xenon and HgCd arc lamps, which again created a master frame from many raw images. This master arc frame was also used to measure the MUSE line-spread

function (LSF) as a function of position on each detector. During the SV run there was a problem with a small amount of vignetting in one of the spectrographs when the temperature was too low. This caused the arc lines in one slice to be too faint to be detected by the software at this stage in the calibration and had to be fixed by changing the arc frames for this spectrograph to ones taken on a warmer night.

The instrument geometry calibration is complicated and so was treated as a static calibration, which was downloaded with the data. Twilight flats were then used to perform illumination correction of the instrument.

These were all of the calibrations required to remove instrumental characteristics from the observational data, allowing them to be ‘pre-processed’. This consisted of turning the raw scientific and calibration observation images into pixel tables of reduced data and also allowed corrections to the wavelength solution based on sky emission lines.

‘Post-processing’ involved using calibration observations to remove the effects of the atmosphere from the data before creating the final products. First, pre-processed standard star observations were used for the flux calibration which created the response curve and telluric correction spectrum. The sky spectrum for the region of interest in the science frames was then estimated using the measured LSF data in combination with the sky spectrum extracted from the background of the scientific frames. Finally, the astrometric solution was computed using an observation of an astrometric field which allowed for the assignment of a world co-ordinate system to the data.

The final datacubes were then made by interpolating over the data in the pre-processed science pixel tables using the post-processed calibrations, and the different exposures were also aligned and stacked at this stage to create science and error datacubes.

The subtraction of the sky emission spectrum was first performed using the




Figure 6.1: A comparison of the strength of the sky emission lines with respect to features in the companion galaxy absorption spectrum (*left*) and the ring galaxy emission spectrum (*right*). The red lines show sky subtraction performed using pipeline routines based on the measured LSF with the blue lines performed by subtracting the background of one exposure from the other.

pipeline routines which estimated the spectrum at the position of the target using the measured LSF and the background spectrum. This produced residuals that were small enough when working with the emission spectra of the ring galaxies, but were too large compared to the absorption features in the companion galaxy, as shown in figure 6.1.

To get improved residuals, the ‘a-b’ sky subtraction method was used. The pipeline was set to produce a non-sky-subtracted datacube for each exposure, ‘a’ and ‘b’, of the galaxies, rather than combining them. Then each of these exposures was subtracted from the other to get two datacubes that had one positive image, one negative image and had no sky background. These two datacubes were then aligned and combined to produce the final reduced datacube with much smaller residuals.

6.2 Kinematic maps

A map of the line of sight velocities of the gas in the galaxies was produced using the H α emission line. To do this, first a signal-to-noise map was created for the H α

emission of the galaxy to which a threshold value of 15 was applied to create a mask that contained only regions with bright enough line emission to accurately fit a model spectrum and measure the position of its peak. The value of 15 was chosen using a trial and error method via a visual inspection of the extracted spectra.

For each spaxel in the masked datacube, the spectrum was least squares fitted with a Gaussian model spectrum the of form:

$$flux = peak \times \exp \left[\frac{-(\lambda - \lambda_{obs})^2}{width} \right] + background, \quad (6.1)$$

where λ_{obs} gives the position of the peak of the emission line along the wavelength axis. This wavelength can then be converted into a recession velocity for the gas in that pixel using the Doppler formula:

$$v = c \frac{\left(\frac{\lambda_{obs}}{\lambda_{em}} \right)^2 - 1}{\left(\frac{\lambda_{obs}}{\lambda_{em}} \right)^2 + 1}. \quad (6.2)$$




Figure 6.2: The kinematic maps derived from fitting a Gaussian profile to the H α line emission from the Cartwheel Galaxy (*left*), Arp 146 (*centre*), and AM0058-220 (*right*). The z-scale shows the recession velocities in km s^{-1} .

The kinematic maps for the ring galaxy sample are shown in figure 6.2 and all are qualitatively very similar. They all show a gradient in the line-of-sight velocity which can be understood as a sum of a rotational velocity (v_{rot}) and an expansion velocity (v_{exp}), as described in section 5.2.1. The line-of-sight velocity for points on the major axis of the ring is due to the galaxy's rotation and the line-of-sight velocity for points

on the minor axis of the ring is due to the ring's expansion, so any gradient along either of these axes is caused by a change in the magnitude of v_{rot} or v_{exp} respectively.

The biggest difference between the kinematic maps in figure 6.2 is the angle between the photometric major-axis of the ring and the axis of the kinematic gradient, which depends on the ratio v_{rot} to v_{exp} . The kinematic axis of the Cartwheel Galaxy is almost aligned with its photometric axis meaning that rotation dominates over expansion, whereas in Arp 146 the angle is close to 90° , meaning that it is dominated by expansion.

6.3 Initial kinematic model

The model described in section 5.2.1 is very simple, but should provide a reasonable estimate of v_{rot} provided that its main assumptions hold true, namely that the ring is closed, close to circular and that the velocity field can be well approximated by a constant v_{rot} and v_{exp} across and around the ring.

6.3.1 Estimating inclination

To measure the eccentricity of the ring, an ellipse was fitted to an image of the $H\alpha$ linemap, which had had the background masked out to leave only the galaxy. This fit was performed by minimising a flux-weighted chi-squared (χ^2) value by varying the x - and y -co-ordinates of its centre, a , b and its position angle. To determine the χ^2 value, the radius of the fitted ellipse as a function of the angle around the ring, $R_{ring}(\theta)$, is found for each spaxel in the image using:

$$R_{ring}(\theta) = \frac{ab}{\sqrt{(b \cos \theta)^2 + (a \sin \theta)^2}}, \quad (6.3)$$

which is then subtracted from the distance of the spaxel from the centre of the ellipse. This is finally multiplied by the flux in that spaxel, f , with the sum across all spaxels

used to determine the goodness of fit:

$$\chi^2 = \sum_{spaxels} [f \cdot (r - R_{ring})]^2. \quad (6.4)$$

Figure 6.3 shows that the rings are well fit by ellipses — the point source on the right-hand edge of the Cartwheel Galaxy can be better identified as a foreground star in the Hubble Space telescope image of figure 5.1. The parameters of these fits are shown in table 6.2 as well as their implied inclinations, estimated from a and b using equation 5.1.

Galaxy	RA (h m s)	Dec (° ' ")	a (")	b (")	PA (°)	i (°)
Cartwheel	00 37 40.415	-33 42 59.53	36.24 ± 0.02	26.45 ± 0.02	144.4 ± 0.1	43.15 ± 0.06
Arp 146	00 06 44.032	-06 38 05.19	9.06 ± 0.08	5.80 ± 0.04	222 ± 3	50.2 ± 0.5
AM0058-220	01 00 46.294	-21 43 30.99	7.56 ± 0.05	5.78 ± 0.06	210 ± 1	40.1 ± 0.8

Table 6.2: The parameters of the best-fit ellipses to the rings.




Figure 6.3: The ellipses that were fit to the rings of the Cartwheel Galaxy (*left*), Arp 146 (*centre*), and AM0058-220 (*right*). The intensity is shown logarithmically scaled and the section missing from the Cartwheel Galaxy is due to the fact that not all of the observations could be taken on that run.

6.3.2 Fitting kinematics

With an estimate of the inclination determined, equation 5.3 was used to estimate the real values of v_{rot} and v_{exp} , and not just their line-of-sight components. The results of this can be seen in table 6.3. The Cartwheel is the galaxy in the sample that has been best studied in the past, and the measurements of the rotational velocity agree with past measurements. Fosbury & Hawarden (1977) estimated v_{rot} and v_{exp} at 254 and 90 km s⁻¹ respectively, with Amram et al. (1998) giving values of 220 and 15 km s⁻¹. My measurements of 229 and 76 km s⁻¹ both lie nicely in between the two literature values. There appears to be a lot of variation between measurements of the expansion velocity, with my estimate lying closer to that of Fosbury & Hawarden (1977), but the rotational velocity is the parameter used to estimate the mass of the galaxy, and that appears to be better behaved.

Galaxy	v_{sys} (km s ⁻¹)	v_{rot} (km s ⁻¹)	v_{exp} (km s ⁻¹)
Cartwheel	8916 ± 2	229 ± 3	76 ± 3
Arp 146	21788 ± 5	95 ± 9	145 ± 10
AM0058-220	12507 ± 5	140 ± 9	85 ± 9

Table 6.3: The velocities derived from fitting the simple model to the kinematic maps

Comparing the fits, seen in figures 6.4–6.6, directly to the measurements, they looked reasonable and broadly fit the gradient across the galaxy. The Cartwheel Galaxy appeared well-fit, as expected from a collision which is expected to be close to axisymmetric, but a couple of features were prominent in the Arp 146 and AM0058-220 fits.

The residuals showed that although the general shape of the kinematic maps were reproduced, there seemed to be systematic variations. For the Cartwheel Galaxy especially, on the major axis the model gave velocities that were too low and on the minor axis the velocities were too high. This implies that the two components of




Figure 6.4: *Top left:* The kinematic map for the Cartwheel Galaxy. *Bottom left:* The fit produced using the initial kinematic model. *Right:* The residuals from the fit. All z-scale units are km s^{-1} .

the velocity are not constant as a function of θ , and that a more complex model is required to produce a more accurate fit.

There were also smaller scale features that were not well fit by the model, such as the right hand side of the AM0058-220 velocity map. These were probably perturbations to the velocity field caused by the off-centre nature of the collision which resulted in warping of the galactic disc, the specifics of which are unlikely to be replicated by any analytical model.

To reduce the effects of warping and irregular velocities across the width of the disc, the line-of-sight velocity was also averaged across the thickness of the ring, perpendicular to the fitted ellipses of section 6.3.1. These radially averaged velocities were then fit by the sinusoid of equation 5.3.

The results are shown in figure 6.7, which shows that the Cartwheel Galaxy and AM0058-220 especially are very close to sinusoidal, and although the residuals do show the patterns as from the velocity maps, it is clearer that their magnitudes are




Figure 6.5: *Top left:* The kinematic map for Arp 146. *Bottom left:* The fit produced using the initial kinematic model. *Right:* The residuals from the fit. All z-scale units are km s^{-1} .

small compared to the amplitude of the sinusoid.

Arp 146, however, was not very well fit by the model in its current form. The peak in figure 6.7 is too narrow, and it has a unusually flat region, which is probably due to warping and peculiar velocities in the ring.

The uncertainties of these results are dominated by how well these models simulate the large-scale morphology and kinematics of the galaxy system and how appropriate the assumptions made are, for example: the amount of warping present in the ring and whether the ring is closed or only appears so due to projection effects. It is also unknown how appropriate is it to assume that v_{rot} and v_{exp} are constant around ring. For example, theoretically v_{exp} is expected to be a function of radius, although the ring could be thin enough to neglect radial effects.

Some of these uncertainties cannot be addressed using these observations, such as whether the ring is closed or not in any particular case, meaning we have to rely on assumptions. However, questions about the properties of the velocity field can be




Figure 6.6: *Top left:* The kinematic map for AM0058-220. *Bottom left:* The fit produced using the initial kinematic model. *Right:* The residuals from the fit. All z-scale units are km s^{-1} .

investigated using simulations, as discussed in section 6.4.

6.4 Analytic simulations

The simple analytic description of off-axis collisions described in section 5.2.2 was investigated to see if its predictions were similar to the initial kinematic model used. The aims were to try to increase confidence in its use and see if any improvements could be made without making it significantly more complicated. Of particular importance was to determine if there were any ways in which the off-axis nature of a collision consistently modified velocity fields and whether they could be easily parameterised to improve fits to the observational data to better estimate the extracted parameters.




Figure 6.7: The line-of-sight velocity of the galaxies averaged across the thickness of the ring as a function of angle around the ring. The line shows the best fit produced by the simple kinematic model. *Top:* The Cartwheel Galaxy, *middle:* Arp 146 and *bottom:* AM0058-220. Error bars show standard deviation of velocities included in averaging.

6.4.1 Method

A ring galaxy was modelled as a 2D circle of approximately 118,000 particles randomly placed within a radius of 25 kpc with a uniform distribution and a constant v_{rot0} for each. The positions of each of these particles could then be calculated at an arbitrary time after the collision using equations 5.5 and 5.6, with v_{rot} and v_{exp} given by equations 5.8 and 5.7 respectively.

However, these velocity components are defined as perpendicular and parallel to the direction to the centre of the initial galaxy, the position of which is not known in the observational data. Therefore, the centre of the newly formed ring structure in the simulated galaxy was determined, as this is the only part of the galaxy visible in the observations and is not necessarily centred on the initial galaxy with an off-centre collision.

To find the location of the ring in the galaxy, an ellipse was fit to the density enhancement that forms the ring. This was done by binning the particles into a 2D histogram to give a density map of the galaxy, similar to the ones in figure 6.8 that show ring- or arc-shaped overdensities. A threshold was then applied to the density maps to create a masked image of the ring overdensity, which allowed an ellipse to be fit using the method described in section 6.3.1. The threshold value was chosen by hand such that the ring thickness-radius ratio appeared to match those seen in the observations. A circle, as is assumed in the kinematic fit, was originally fit to the ring to determine its position, but as the ring was not always perfectly circular, especially for very off-centre collisions, an ellipse was used instead to give a more accurate centre position.

Once the centre of the ring is known, then the velocity of each particle can be decomposed into components that are perpendicular and parallel to the direction to the centre of the ring. Velocity maps were then made for each velocity component in both reference frames to discover the effect of measuring the velocities with respect

to the ring rather than original galactic centre. The maps were created by averaging the velocities in the same bins used in the 2D density histograms.

The centre of the ring was also used to define the angle around the ring, θ , to investigate how it affects v_{rot} and v_{exp} to see if their constant form in equation 5.3 is accurate or whether a simple function of θ could be used to improve them. The velocities were fit with a sinusoid and its amplitude, phase and average were determined as a function of the initial conditions of the ring galaxy system.

The analytic nature of this simulation allows for quick and easy investigation of the effects of varying the impact parameter and perturbation velocity, Δv of the collision as well as the v_{rot0} , all as a function of time. Therefore a fiducial scenario was chosen as a reference and each of these parameters varied, one at a time, to measure their effect on the size and shape of the ring, and v_{rot} and v_{exp} .

The fiducial impact parameter of 4 kpc was selected as it is large enough to see the effects of being off-centre but small enough to produce recognisable rings that probably would not have too much warping in full simulations, and it is comfortably inside the estimated radius of the observed rings, shown in table 6.6. The upper limit was chosen to be 8 kpc, as above that the rings get less circular and require a more complex simulation to model accurately.

The reference v_{rot0} was set at 100 km s^{-1} , as it is approximately what is measured for both Arp 146 and AM0058-220. It was varied in steps, up to 200 km s^{-1} , with the contrast and completeness of the ring dropping with a higher v_{rot} and so the accuracy with which its position can be measured also drops.

The velocity perturbation reference was chosen to be 50 km s^{-1} , as above that parts of the galactic disc started to acquire epicyclic amplitudes larger than their orbital radii, meaning that model starts to break down and particles require a more complex orbital description. The lower end of the parameter space was chosen to be 30 km s^{-1} as the contrast and completeness of the ring fell the lower the Δv .

6.4.1.1 Assumptions

There were a number of assumptions made when choosing this model to compare to the observations. Firstly, as the models are predicting the movement of the stars, that the gas used to measure the kinematics behaves in a similar manner. Secondly, it was assumed that the density wave in the gas occurs at the same position as the density wave in the stars. This compares favourably to the simulations performed by Gerber et al. (1996) and Mapelli & Mayer (2012).

Other assumptions include that the perturbation velocity of the collision is independent across the target disc, which results in similar behaviour to the case of point-like companion galaxy. The perturbations are also assumed to be small enough to be well described by epicyclic motion, which should hold for lower mass and slower companions. The final results of this work do not suggest that this is a particularly bad assumption. Finally, the equations are made simpler by assuming a flat rotation curve for the target disc.

All of these assumptions were necessary to allow such simple models to be used and were therefore judged to be reasonable for this work.

6.4.2 Results and discussion

6.4.2.1 Ring morphology

Density maps of the formation and evolution of simulated ring galaxies are shown in figure 6.8. They show how quickly the density wave forms and propagates through the galactic disc. In the beginning there appear to be two arcs that form, one near the centre of the galaxy and one near the impact point on the x -axis, these are especially clear for the runs with a larger impact parameter. These join up to form a cuspy ring, which superficially looks similar to the simulations run by Mapelli & Mayer (2012), before tending towards a more circular shape.




Figure 6.8: Density maps showing the formation of ring overdensities for different collision scenarios with time uniformly increasing from top (25 Myr) to bottom (150 Myr). *Left*: The reference scenario, with a 100 km s^{-1} $v_{\text{rot}0}$ galaxy perturbed by a velocity of 50 km s^{-1} with an impact parameter of 4 kpc. *Centre left*: The effect of increasing the impact parameter to 8 kpc. *Centre right*: The effect of decreasing Δv to 30 km s^{-1} . *Right*: The effect of increasing the $v_{\text{rot}0}$ to 200 km s^{-1} . The side of each box is 20 kpc in length and the z-scales show the density of the target disc with the same normalisation for each subplot.




Figure 6.9: The centres of the rings as a function of time for impact parameters of 2, 4, 6 and 8 kpc. The time period shown is 60–200 Myr, with the rings appearing to move upwards as plotted.

As a result, the ring is centred close to halfway between the original centre of the galaxy and the collision point, as seen in figure 6.9. This plot also shows that the ring structure moves slightly with respect to the galactic disc, in the same direction that the galaxy rotates. This speed is quite small though, with a magnitude of $\sim 10 \text{ km s}^{-1}$.

The ring formed by an off-centre collision is not perfectly circular — the general form of the solutions to the caustic equations is a limaçon, which is formed by tracing a fixed point on a circle as it is rolled around another circle, and specifically they can range from a perfect circle to one with a cusp in it. Given that the observational data are well fitted by an ellipse when measuring their position and eccentricity, an ellipse was also fitted to the simulated data for consistency and to estimate how circular the real rings might be.

The results of this ellipse fitting is found in figure 6.10 and shows that the ratio of the length of the minor-axis to the major-axis is dependent on the impact parameter of the collision and ranges from approximately 0.8 for 8 kpc up to >0.95 for 2 kpc. It




Figure 6.10: The ratio of the ring's fitted major-axis to its minor-axis as a function of time. *Top left*: The effects of varying impact parameter. *Top right*: The effects of varying Δv . *Bottom*: The effects of varying v_{rot0} .

also depends on the time since the collision, but this is a much weaker dependence, with the rings slowly getting more circular with age after they have formed. As in figure 6.8, it is evident that for very off-centre collisions it takes longer for the rings to fully form, which is shown by the steeper gradient at earlier times in figure 6.10.

Any ellipticity would cause inaccuracies in the measured inclination of the ring if it is assumed to be circular. Given equation 5.1, an uncertainty in b/a , $\alpha_{\frac{b}{a}}$ leads to an uncertainty in the inclination, α_i , of:

$$\alpha_i = \frac{\alpha_{\frac{b}{a}}}{\sqrt{1 - \left(\frac{b}{a}\right)^2}}. \quad (6.5)$$

As v_{rot} is modulated by the sine of the inclination in equation 5.3, the uncertainty

of $\sin i$, $\alpha_{\sin i}$, is:

$$\alpha_{\sin i} = \alpha_i \cos i = \frac{\alpha_{\frac{b}{a}} \cos i}{\sqrt{1 - \cos^2 i}}, \quad (6.6)$$

which, taking $\alpha_{\frac{b}{a}} = 0.1$ from figure 6.10, would lead to a 10% uncertainty for v_{rot} of a ring viewed at an inclination of 45° .

6.4.2.2 Rotational Velocity

In the kinematic model of equation 5.3, v_{rot} was assumed to be constant for the whole ring. However, as the collision is modelled as a velocity perturbation directed towards the centre of the interaction, any non-zero impact parameter will cause an increase in v_{rot} for the half of the galaxy upstream of the encounter point and a decrease for the downstream half of the galaxy. This effect can be seen in figure 6.12, which shows maps of v_{rot} and v_{exp} for the galaxy as a whole and for just the ring structure.




Figure 6.11: Density maps of the simulated ring galaxy with an impact parameter of 6 kpc, v_{rot0} of 100 km s^{-1} , magnitude of Δv of 50 km s^{-1} and at a time of 160 Myr both without (*left*) and with (*right*) a cut on the density to predict the position of the ring. The z-scales show the normalised values of the density.

As v_{rot} is no longer constant with θ , it was expected that it would not be easy to parameterise — particles with different velocities would have different orbital periods and therefore move out of phase with each other. This is seen in the velocity field of the whole galaxy, where spiral patterns can be seen in the rotational velocity maps.




Figure 6.12: *Top*: v_{rot} with and without a cut on the density, with respect to the original galactic centre (*left*) and the centre of the ring (*right*). *Bottom*: v_{exp} with and without a cut on the density, with respect to the original galactic centre (*left*) and the centre of the ring (*right*). All z-scale units are km s^{-1} .

However, when the ring is isolated v_{rot} looks much more ordered and appears very similar at different times, which is due to different particles making up the ring at different ages.

When v_{rot} is plotted as a function of θ , as in figure 6.13, it can be seen that it is very close to sinusoidal in shape due to half of the ring being sped up and half of the disc being slowed down. Neither figure 6.12 nor 6.13 show any significant difference

in the average value of v_{rot} across the thickness of the ring, which should help ensure the model stays as simple as possible. It can also be seen that the effect of measuring v_{rot} around the centre of the ring, as opposed to the centre of the initial galaxy is very small. There is just a slight shift in the phase of this sinusoid of the order 15° , but these effects are fairly small compared to the magnitude of both v_{rot} itself and the amplitude of the sinusoid.




Figure 6.13: v_{rot} (*upper dataset*) and v_{exp} (*lower dataset*) as a function of the angle around the ring. For both graphs the impact parameter is 6 kpc, the v_{rot0} is 100 km s^{-1} , the magnitude of Δv is 50 km s^{-1} and the time is 160 Myr after the collision. The colours correspond to the radius of the data point in units of the ring's radius. *Left*: The velocities with respect to the centre of the original target galaxy. *Right*: The velocities with respect to the measured centre of the ring.

The average value of the sinusoid that was fitted to v_{rot} as a function of time and v_{rot0} is shown in figure 6.14. It can be seen that the measured magnitude of v_{rot} was the same when measured in both reference frames, meaning that the mass derived from measuring the rotation around the centre of the ring should be the same as that measured around the centre of the initial galaxy. Reassuringly, the only initial condition that was found to have an impact on the measured value was its initial value, with further plots given in Appendix A.

It was also not significantly affected by the collision, meaning that the measured velocity should represent the pre-collision velocity which is itself determined by the




Figure 6.14: The average of v_{rot} with time after the collision showing the effects of changing v_{rot0} of the initial galaxy. Solid lines shows results for v_{rot} measured around the centre of galaxy, dashed show v_{rot} around the centre of ring.

mass of the galaxy; therefore an accurate mass estimate should be feasible for the pre-collisional galaxy. However, this result is dependent on assumptions made in the model and so is likely to only hold for smaller Δv s, and not if the collision is too slow or the companion too massive. This conclusion is supported by Gerber et al. (1996), where it is shown that v_{rot} is expected to change drastically with an equal mass companion, but stay relatively stable with a companion that has a mass of 25% of the target galaxy. This compares well with the mass ratios calculated in Chapter 7.

The amplitude of the sinusoid that was fitted to v_{rot} is shown in figure 6.15 as a function of impact parameter, Δv and time. It was found to increase with an increasing impact parameter of the collision and increase weakly with an increasing Δv . This is as expected due to both parameters increasing the component of Δv that is tangential to rotational velocity field. The fact that the dependence on the magnitude of Δv appears so weak is somewhat surprising, but these models were only




Figure 6.15: The amplitude of the sinusoid describing the change in v_{rot} around the ring as a function of time after the collision. *Left*: The effects of changing the impact parameter. *Right*: The effects of varying the magnitude of Δv imparted by companion galaxy. Solid lines show results for v_{rot} measured around the centre of galaxy, dashed show v_{rot} around the centre of ring.

run for relatively low perturbations. This was due to the breakdown in the assumption that the stellar orbits perform epicyclic oscillations, meaning that significantly larger values could not be realistically tested and the strength of the relationship cannot be estimated in this regime. The amplitude can also be seen to decrease slowly with time after the collision, as the ring is formed from particles that had a lower range of rotational velocities. Given the complex dependencies of the amplitude on the impact parameter, Δv and time elapsed, it cannot be simply parameterised in terms of estimable quantities such as the two galaxy masses without more elaborate simulation work. Importantly, given the uncertain position of the original galactic centre in the observational data, the frame in which the velocities were measured in did not make a difference. This meant that the amplitude could be included in updated models as a simple free parameter to improve the quality of the fits. The amplitude was not found to be significantly affected by v_{rot0} , as is shown in Appendix A.

The phase of the sinusoid that was fitted to v_{rot} as a function of the impact parameter, Δv , v_{rot} and time is shown in figure 6.16. When v_{rot} was measured with




Figure 6.16: The phase of the sinusoid describing v_{rot} against time after the collision, showing the effects of changing the Δv of the collision. Solid lines shows results for v_{rot} measured around the centre of galaxy, dashed show v_{rot} around the centre of ring.

respect to the original galactic centre, the phase of the sinusoid was found to depend on all of the parameters, however when it was measured with respect to the ring, the phase was almost independent of all variables.

For example, when measured with respect to the galactic centre: changing the impact parameter from 2 kpc to 8 kpc caused a 20° decrease in the phase, an increase in Δv from 30 km s^{-1} to 50 km s^{-1} caused it to drop by 30° , doubling v_{rot0} from 100 km s^{-1} to 200 km s^{-1} increased the phase by up to 40° . From a time of 60 Myr to 200 Myr, and the phase increased by approximately 15° . When measured with respect to the ring centre the phase was constant to within 10° for all parameters apart from time, which saw the same slow increase of about 15° from 60 Myr to 200 Myr. This less variable phase with respect to the ring centre is good for the modelling of the observational data. It makes it easier for v_{rot} to be fitted with respect to the ring pattern, which can be measured, as opposed to the galactic centre, which cannot.

6.4.2.3 Expansion Velocity

Figure 6.12 also shows the v_{exp} both of the underlying galactic disc and of the ring, again with respect to the centres of both the galactic disc and the ring. When looking at the whole galaxy there is a clear, slightly asymmetrical, annular structure, with a rapidly expanding ring that is on the inner edge of the ring overdensity with the

galaxy outside of the ring still falling inwards. There is not a major difference in the overall form when switching from the frame of the galaxy to that of the ring, but phase of the pattern changes by 180° — the edge of the ring that is expanding the fastest switches sides.

The azimuthal dependence of v_{exp} is shown alongside v_{rot} in figure 6.13. The functional form is not as cleanly parameterised as v_{rot} in either reference frame, with a large scatter due to the radial dependence of v_{exp} clearly shown by the vertical colour gradient in the plot. The dependence is close to sinusoidal in the galaxy’s frame, only with the peak being broadened and flattened. In the ring’s frame it is also close to sinusoidal, but this time with a dip coincident with the maximum. The larger difference between the two frames is that there is phase shift of about 180° introduced, which is not the case with v_{rot} .




Figure 6.17: *Top*: The average v_{exp} with time, showing the effects of changing impact parameter (*left*) and v_{rot0} (*right*). The solid line shows v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.

The dependence of the average of the sinusoid fitted to v_{exp} as a function of the impact parameter and time is shown in figure 6.17. The average v_{exp} of the ring seems independent of the reference frame, with a slow increase as a function of time. The effect of increasing the impact parameter of the collision is to reduce the expansion velocity, meaning that the particles’ outward velocity should be at

a maximum for an axisymmetric collision. The results also suggest that a higher v_{rot0} , and therefore more massive target galaxy, could lead to a larger value of v_{exp} , although this does not appear to be certain for $v_{rot} > 100 \text{ km s}^{-1}$ for times greater than 100 Myr. This appears to be at odds with the preliminary observational results which show that the Cartwheel Galaxy, which is thought to have had the smallest impact parameter and highest v_{rot0} , has the lowest value for v_{exp} . This could be due to limitations of this model for at least one of these galaxy systems, which will be explored further in section 6.5. The effects of varying Δv appear to be negligible, with further plots given in Appendix A.




Figure 6.18: The amplitude of the change in v_{exp} around the ring with time, showing the effects of changing impact parameter (*left*), and v_{rot0} (*right*). The solid lines show v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.

The dependence of the amplitude of the sinusoid fitted to v_{exp} as a function of the impact parameter, v_{rot0} and time is shown in figure 6.18. The amplitude of the sinusoid fitted to the v_{exp} decreases as a function of time and also increases with an increasing impact parameter, both in a manner similar to the amplitude in v_{rot} . The biggest difference here is the discrepancy between the velocities measured with respect to the ring and the initial galactic centre. It can be seen that in the frame of the galactic centre, the amplitude does not seem to vary with v_{rot0} , and only weakly varies

with impact parameter. However, a different effect is seen in the velocities measured with respect to the ring, which now appear to have larger amplitudes, especially with v_{rot0} . This means that when fitting the expansion velocity of ring galaxies, if this expansion is in the frame of the galactic centre the amplitude is possibly much less important than for velocities in the frame of the ring. Given the position of the galactic centre, is not known in the observational data, it would appear this amplitude could be important to include to improve the kinematic models.




Figure 6.19: The phase of the sinusoidal pattern round ring against time showing the effects of changing v_{rot0} . The solid line shows v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.

The dependence of the phase of the sinusoid fitted to v_{exp} as a function of v_{rot0} and time is shown in figure 6.19. As with v_{rot} , it once again behaves differently between the the frame of the original galaxy and the ring. Although both have the same slow increase with time, there is a large difference of approximately 200° between them, as expected from figure 6.12. In the frame of the galaxy the phase increases with both decreasing impact parameter and decreasing v_{rot0} , but is apparently unaffected

by Δv . However, when measured with respect to the fitted ring the parameters seem to only depend on time, which is the same as is seen with the phase of the v_{rot} sinusoid.




Figure 6.20: v_{exp} as a function of radius, with the colour corresponding to the angle around the ring. The left hand plot shows v_{exp} measured with respect to the original galactic centre and the right hand plot shows it measured with respect to the centre of the fitted ring.

For the ring to form, v_{exp} has to decrease as the radius increases, as shown by figure 6.20. There is a fairly large amount of scatter compared to the change in velocity over the ring's width in this relation, with the fastest expansion at the outer edge of the ring not much slower than the slowest expansion at the inner edge of the ring. The relationship appears to very simple, with the ring being thin enough that it offers no motivation to fit anything other than straight line across its width.

The differences when the velocity is measured with respect to the ring, as opposed to the galactic centre, are small. The gradient is not significantly affected, though the scatter does slightly increase. The behaviour is more consistent around the ring: the colours are more stratified, meaning that the gradient varies less when measured as a function of angle than when measured with respect to the galactic centre.

The expansion rate of the ring density wave, as opposed to the particles that make up the ring, is shown in figure 6.21. It can be seen that the density wave expansion rate is predicted to be independent of both the impact parameter and Δv , depending




Figure 6.21: The radius of the ring as a function of time. *Left:* The effects of varying impact parameter. *Centre:* The effects of varying Δv . *Right:* The effects of varying v_{rot0} .

only on v_{rot} . A comparison with figure 6.17 shows that the expansion rate of the individual particles cannot necessarily be used to accurately predict the expansion rate of the ring as it is just a density wave. Therefore any ages estimated from the ring's velocity and radius are subject to a large uncertainty.

This behaviour is because the expansion of the ring depends on the outwards motion of particles, and in this model this is described by their epicyclic frequency. As can be seen in equation 5.10, the epicyclic frequency is a function of the target galaxy itself and not on the collision parameters, so it does not predict the ring expansion rate to depend on the collision's parameters.

This result seems at odds with more complex simulations, such as those performed by Gerber et al. (1996), and seems to show some of the limitations of an analytic simulation. This shows why these simulations are only being used for general trends seen in the v_{rot} and v_{exp} to help improve the fits and not any exact quantities or precise relations to estimate the collisional parameters.

6.4.2.4 Correlation between rotational and expansion velocities

Although the phase of the v_{rot} and v_{exp} sinusoids weakly depend on time in figures 6.16 and 6.19, they both appear to have the same relationship when measured with respect to the centre of the ring.




Figure 6.22: The difference between the phase of v_{rot} sinusoid and the v_{exp} sinusoid with the velocities measured with respect to the centre of the fitted ring. The left hand plot show the effect of varying the impact parameter, the middle Δv and the right the v_{rot0} .

When the difference between the two phases is plotted, as in figure 6.22, it can be seen that it is very close to independent of all of the variables tested — time, impact parameter, Δv and the v_{rot0} — at a value of approximately 110° . This means that this parameter can fixed at this value in updated fits to the kinematic maps.

6.4.3 Conclusions and improved fit

It has been shown that there are several ways in which the kinematic model of equation 5.3 can be improved using very simple equations with few additional free parameters. v_{rot} and v_{exp} can be measured with respect to the known centre of the ring without estimating the position of the original galactic centre. The simplest improvement that can be made is that their constant nature can be replaced by a sinusoidal form which is a function of the angle around the ring, with a fixed phase difference between the two. Therefore the new free parameters are the amplitudes of the sinusoids, A , and a single phase, ϕ :

$$v_{rot} = v_{rot} + A \sin(\theta - \phi), \quad (6.7)$$

$$v_{exp} = v_{exp} + A \sin(\theta - \phi - 110^\circ). \quad (6.8)$$

An additional term that could also be added is to let v_{exp} vary linearly with radius, by using its projected value at a radius of 0 kpc, c_{exp} , as a free parameter. This will probably not offer as great an improvement as the azimuthal dependence of the velocities because the scatter in v_{exp} due to the radial dependence was smaller than the amplitude of the sinusoidal changes. There are also uncertainties as although it is simple to implement and should give a better understanding of the velocity field in ring galaxies, it is perhaps more suited to describing stellar motions velocity field rather than gas dynamics.

The limitations of this model are discussed in section 5.2.2.1, but the most obvious case where this breaks down is for large perturbation velocities caused either by high mass or slow moving companion galaxies. This means that these relations may not hold if the companion is too massive, though this is probably not the case for the Cartwheel Galaxy, which should prove to be a good test of their accuracy.

6.5 Improved kinematic model

There were two refinements to equation 5.3 that were found in section 6.4 that could easily be added in to improve the quality of fit that it produced. The first was to change v_{rot} and v_{exp} to make them sinusoidal functions of θ , as seen in equations 6.7 and 6.8, whereas the second was to make the v_{exp} a linear function of radius.

Two improved models were then tested, with the first only adding the sinusoidal components and the second adding both improvements. In both models, the amplitudes of the sinusoids in the rotational and expansion velocities, A_{rot} and A_{exp} , were first allowed to vary independently from one other. However, this led to them having extremely large uncertainties. As the amplitudes of the curves in figure 6.13 are similar, they were then constrained to be equal to each other, which did not greatly affect the goodness of fit and gave much more believable uncertainties.

6.5.1 Results and discussion

The χ^2 values for each fit are shown in table 6.4. The uncertainties that were used were statistically derived and therefore tended to underestimate the true uncertainty, leading to larger χ^2 values. As such, comparing the relative values for the different models for each galaxy is the most important detail to consider. They show that the improved model gave a significant improvement in the quality of the fit for each galaxy, which appears to indicate that the off-centre nature of the collision could be an important point to consider when fitting the kinematics of these galaxies. However, the impact of having a radially varying v_{exp} was not very significant in any galaxy, which perhaps indicates that the rings are thin enough to neglect any radial variation.

Galaxy	Simple Model	Sinusoid	Expansion
Cartwheel	8.02	4.65	4.50
Arp 146	11.17	5.88	5.26
AM0058-220	1.62	1.17	1.17

Table 6.4: The reduced χ^2 of the different models fitted to the kinematic maps. The Sinusoid model adds the sinusoidal terms to v_{rot} and v_{exp} with θ , with the Expansion model adding the linear variation of v_{exp} with radius on top.

The improved fits to the velocity maps of each galaxy are shown in figures 6.23-6.25 and are compared to the measured velocity maps and the original fits. The residuals from the improved fits are qualitatively similar to the those from the original fits in terms of their shapes, differing mostly in their magnitudes — meaning that the areas where the fit was too high or too low are still in the same place. This shows that the major differences between the data and the models are small scale differences which are probably local velocity irregularities or warping caused by the collision.

The parameters derived from the fits for each model are given in table 6.5. When looking to derive the mass, the most important parameter is v_{rot} , but it can only be used to get an estimate in conjunction with the radius, which is derived from the fitted ellipse that is converted into a proper distance using the recession velocity and




Figure 6.23: *Top:* The measured velocity map of the Cartwheel Galaxy. *Below, in order:* The original model fit to the data, the improved model fit including the sinusoidal terms for v_{rot} and v_{exp} and lastly the improved fit also including the radial dependence on v_{exp} . Each dataset includes the velocity map (*left*) and the residuals (*right*). All z-scale units are km s^{-1} .

Hubble's constant.

For every galaxy, there is not a large change in the extracted rotational velocity when the model is changed from the original to either of the improved ones. This is




Figure 6.24: *Top*: The measured velocity map of Arp 146. *Below, in order*: The original model fit to the data, the improved model fit including the sinusoidal terms for v_{rot} and v_{exp} and lastly the improved fit also including the radial dependence on v_{exp} . Each dataset includes the velocity map (*left*) and the residuals (*right*). All z-scale units are km s^{-1} .

because the underlying model of an rotating and expanding ring is still the same, with the estimates of v_{rot} and v_{exp} coming from the positions of the peaks in the overall sinusoid. These do not move much when refining the model to better fit the fact that




Figure 6.25: *Top:* The measured velocity map of AM0058-220. *Below, in order:* The original model fit to the data, the improved model fit including the sinusoidal terms for v_{rot} and v_{exp} and lastly the improved fit also including the radial dependence on v_{exp} . Each dataset includes the velocity map (*left*) and the residuals (*right*). All z-scale units are km s^{-1} .

the curve is not a perfect sinusoid — due to the off-centre nature of the collisions. This increases the confidence with which v_{rot} can be estimated accurately, as for each galaxy the line-of-sight velocity when plotted against θ is very close to a sinusoid, as

seen in figure 6.26, making the estimates for v_{rot} very robust.

The only galaxy which shows a significant variation of v_{rot} between models is Arp 146, which shows an unusual feature on its leftmost edge in figure 6.24 and in the centre of its subplot in figure 6.26. This feature, which looks like a flattening of the rotation curve with θ , is too sharp and too local to be fitted using the simple models here and could therefore be slightly skewing the value of v_{rot} . However, as it is much smaller than the main sinusoidal appearance of the plot in figure 6.26, the value of v_{rot} does not change by a huge amount.

The derived recession velocities change by a significant amount when comparing the values between models, but this number is small when comparing to the uncertainty in the Hubble Flow due to peculiar velocities. This will be discussed further in section 6.6.

Galaxy	v_{sys} (km s ⁻¹)	v_{rot} (km s ⁻¹)	v_{exp} (km s ⁻¹)	A (km s ⁻¹)	ϕ (°)	C_{exp} (km s ⁻¹)
Cartwheel	8916 ± 2	229 ± 3	76 ± 3			
	8960 ± 10	229 ± 3	68 ± 4	140 ± 5	235 ± 3	
	8960 ± 10	229 ± 3	68 ± 4	134 ± 5	237 ± 3	19 ± 1
Arp 146	21788 ± 5	95 ± 9	145 ± 10			
	21791 ± 4	73 ± 9	121 ± 10	46 ± 10	16 ± 6	
	21792 ± 6	74 ± 9	123 ± 10	45 ± 10	19 ± 11	17 ± 7
AM0058-220	12507 ± 5	140 ± 9	85 ± 9			
	12557 ± 5	137 ± 14	101 ± 8	139 ± 30	290 ± 30	
	12563 ± 30	139 ± 14	101 ± 8	132 ± 30	280 ± 40	141 ± 4

Table 6.5: The fitted parameters for each model. For each galaxy, the top set of parameters is the original model, the middle set adds the sinusoidal terms to v_{rot} and v_{exp} and the bottom set also adds the radial dependence of v_{exp} .

The Cartwheel Galaxy has the largest rotational velocity, which implies that it is the most massive galaxy in the sample, but it also has the largest A . This was not expected as A is expected to increase with the impact parameter of the collision, which implies that the Cartwheel Galaxy had the least symmetrical collision, which is not expected to produce the most symmetrical ring.

Arp 146 has the smallest v_{rot} , implying it is the least massive galaxy, and the smallest A , implying it was produced by the most symmetrical collision. However, the feature on its leftmost edge in figure 6.24 and in the centre of figure 6.26, which could be the cause of the larger change in v_{rot} between the models, could also be causing A to seem smaller than it should do in reality.

AM0058-220 appears to be of a mass in between the Cartwheel Galaxy and Arp 146, and it is the only galaxy where the A is comparable to v_{rot} . This implies that v_{rot} is reduced to $\sim 0 \text{ km s}^{-1}$ at certain points on the ring, which sounds unlikely, but possible if the impact parameter of the collision and mass of the companion are large enough. Certainly, the incompleteness of the ring suggests that this was produced by the most off-centre collision in the sample.

One factor that could be influencing the apparent size of A is the phase difference between the sinusoids of the rotational and expansion velocities. The simulations in section 6.4 indicated that this was a constant 110° , but if this number is wrong then the average values of v_{rot} and v_{exp} would not be affected, but A would vary accordingly.

The radially averaged line-of-sight velocities with θ are shown in figure 6.26. The fits from the improved model are shown along with the fit from the original model. It can be seen that for the Cartwheel Galaxy and AM0058-220 the new model produces fits that are slight deviations from a perfect sine wave, to improve the fit. The peaks and troughs are in roughly the same place, so the rotational velocity is not changed by very much.

For Arp 146 however, the plateauing of the velocity at around $\theta = 180^\circ$ means that the best fit was a larger deviation from a perfect sine wave. This means that the rotational velocity was affected more by the new model, which may be a more accurate estimate of the intrinsic rotational velocity, or may be fitting small-scale features. The best way to handle it is therefore to use both models for the final

estimate for the rotational velocity and to increase the uncertainty accordingly.

6.6 Mass estimates

The masses of the ring galaxies were estimated using equation 5.14. Therefore the two parameters that were required for each galaxy were its rotational velocity and its radius, which is calculated from the apparent size and distance.

This only gives the mass within the ring itself rather than the mass of the whole galaxy. Romano et al. (2008) took deep images of a small sample of 9 ring galaxies and managed to detect the stellar disc outside of the ring. They found that the ring had typically propagated through half of the disc, postulating that this is due to ring galaxies being hard to identify when the ring is too small, and the contrast of the ring fading as it gets too far out. The ratio of the disk radius to the ring radius was found to be 1.9 ± 0.4 , which is a factor that I will multiply the radius by to get an estimate for the mass of the total galaxy, with the mass within the ring acting as a lower bound.

6.6.1 Discussion

Galaxy	d (Mpc)	r (kpc)	M_{ring} ($10^{10}M_{\odot}$)	M_{gal} ($10^{10}M_{\odot}$)
Cartwheel	131 ± 6	23 ± 2	28 ± 9	53 ± 20
Arp 146	321 ± 4	14.1 ± 0.7	1.8 ± 0.4	3 ± 1
AM0058-220	185 ± 6	6.8 ± 0.4	3 ± 1	6 ± 2

Table 6.6: The final estimates for the distance, radius and mass of each galaxy.

The results are presented in table 6.6, with the Cartwheel Galaxy a factor of ~ 10 more massive than other two. The result for the mass contained within the ring of the Cartwheel galaxy, at $(28 \pm 9) \times 10^{10}M_{\odot}$, is lower than the values of $40 \times 10^{10}M_{\odot}$ from Fosbury & Hawarden (1977), although still in agreement. This measurement

also agrees favourably with Amram et al. (1998), who estimate that it lies in the range $(6\text{--}24)\times 10^{10}M_{\odot}$, with their estimate lower than from this work. Work estimating the masses of the companion galaxies is presented in Chapter 7, including a comparison with the ring mass for each galaxy system.

The distances to the galaxies was calculated using equation 5.2 and taking H_0 from the Planck Mission (Collaboration et al., 2016). The uncertainty in v_{rec} is dominated by the peculiar velocities of the galaxy, and was therefore estimated by comparing the measured recession velocity of the ring galaxy with the measured recession velocity of its companion. The uncertainty was taken to be the difference in the two velocities, which should be slightly conservative, but a good estimate given the data available.

Due to the fact that the ring may not be intrinsically circular, as discussed in section 6.4, the uncertainties in the inclination and radius were both increased accordingly. The uncertainty in the inclination of the ring is given by equation 6.5, with the uncertainty on the radius taken to be 10% as the semi-major axis of the ellipse caused by the galaxies inclination could be aligned with the intrinsic semi-major or -minor axis of the ring.

Perhaps the largest source of uncertainties is how well the model represents reality. The estimate of the rotational velocity looks fairly robust due to the residuals being small compared to the sinusoid fitted in figure 6.26, but it is not known whether this is this an accurate estimate of the pre-collision galaxy. However, the relatively small mass ratios measured in Chapter 7 should increase the confidence placed in the analytic simulations, as this was one of the key assumptions made.

Another major unknown comes from the dark matter halo of the ring galaxy. Whilst efforts have been made to estimate how much of the disc lies outside of the ring, the dark matter halo will extend even further, with no way of measuring how much mass is unaccounted for. This will have to be taken into account when calculating the uncertainties on the mass ratios.

Galaxy	Date (mon-yr)	RA (h m)	Dec ($^{\circ}$ ')	Exp. Time (s)
IZw45	May-12	12 53.11	+36 49.10	600
IIZw28	Jan-13	05 01.73	+03 34.52	600
IIHz4	Jan-13	08 58.55	+37 05.19	600
Arp148	April-13	11 03.90	+40 51.00	1200

Table 6.7: The sample of ring galaxies observed using the Oxford SWIFT Spectrograph at Palomar Observatory that have yet to be analysed.

6.7 Future work

Observations of more ring galaxies have been taken with the Oxford SWIFT Spectrograph, which as yet have not been analysed and are summarised in table 6.7. A similar analysis to the work in this chapter can be performed to further increase the sample of ring galaxy mass ratios.

There are questions that have arisen from this work that have not yet been investigated. For example, it would be interesting to further investigate the ages of the rings, comparing the estimates derived from the measured expansion velocities to those derived from the rotational velocities by estimating their epicyclic frequencies. Another interesting idea is that the collisional geometry could be estimated, similar to the work by Fogarty et al. (2011). This could then be used in simulations along with the measured mass ratios, with the results compared to the observations.

In short, the small sample of ring galaxies that have been well studied can be increased, and does not necessarily require further telescope time. This work could lead to a better understanding of ring galaxies and therefore a better understanding of galaxies in general.




Figure 6.26: The line-of-sight velocity of the galaxies averaged across the thickness of the ring as a function of angle around the ring. The dashed line shows the best fit produced by the simple kinematic model, with the solid line showing the best fit produced by the improved kinematic model. *Top*: the Cartwheel Galaxy, *middle*: Arp 146 and *bottom*: AM0058-220. The error bars show the standard deviation of the velocities included in the average.

Chapter 7

Masses of ring galaxy companions

In this chapter, the aim is to estimate the stellar masses of the companions to Arp 146 and AM0058-220 using stellar population synthesis models described in section 5.3. These mass estimates are then compared to the mass estimates for the corresponding ring galaxies and the implications on the formation of the rings discussed.

7.1 Method

The stellar mass of each companion galaxy was estimated by fitting its integrated spectrum with a linear superposition of single stellar population (SSP) model spectra. The weights of each SSP model in the fit was then used to derive a mass-to-light ratio for the entire galaxy which could be multiplied by the the galaxy’s intensity to obtain the mass estimate. The uncertainties on this number were then rigorously estimated using bootstrapping techniques.

7.1.1 Spectrum fitting

The first task was to extract the integrated spectrum of the companion galaxy from the fully-reduced and calibrated datacube produced as described in section 6.1. This

was done by summing all of the flux in an elliptical mask that encompassed the galaxy for every wavelength channel in the datacube. The size of this mask was determined visually using an image made by summing the datacube along the wavelength axis ensure the measured flux was an accurate estimate of the galaxy's total flux.

Once the spectrum of the companion galaxy has been extracted it is then ready to be fit by the SSP models. The first stage of fitting the models to the data is to create the library of model spectra with which the data are to be fitted.

The stellar population models that were used were the MILES models (Vazdekis et al., 2010) that are based on a range of stellar libraries (Heap & Lindler (2007), Sánchez-Blázquez et al. (2006), Valdes et al. (2004), Cenarro et al. (2001), Rayner et al. (2009)). All spectra used had the same unimodal Salpeter IMF (Vazdekis et al., 1996) with a slope of 1.3, and a base alpha abundance. They were calculated based on the isochrones from Girardi et al. (2000). The differences between the different SSP models were the age of the stars as well as their metallicities. A range of ages logarithmically spaced between 0.0631 Gyr and 17.7828 Gyr were used for each metallicity, with the precise values of each shown in table 7.1.

During the spectrum fitting routine, the model spectra are combined into a single three-dimensional library array such that the axes represented wavelength, age and metallicity. The range of ages used to fit the companion to AM0058-220 was reduced to lower the run time of the fitting procedure, with ages below 1 Gyr excluded. This array can then be rebinned along the wavelength axis so that the data points are spaced logarithmically, as opposed to linearly. This is done in such a manner as to rigorously conserve the total flux. The final step in preparing the template spectra is to divide the whole array by the single median of all templates. This is done because each template represents the light emitted by a fixed total mass. The templates are normalised to simplify the fitting process and by dividing by a single median rather than the median of each template the final result is weighted by the mass of the stellar

Ages (Gyr)			log[M/H]
0.0631	0.4467	3.1623	0.2223
0.0708	0.5012	3.5481	0.0000
0.0794	0.5623	3.9811	-0.396
0.0891	0.6310	4.4668	-0.7052
0.1000	0.7079	5.0119	-1.3146
0.1122	0.7943	5.6234	-1.7129
0.1259	0.8913	6.3096	
0.1413	1.0000	7.0795	
0.1585	1.1220	7.9433	
0.1778	1.2589	8.9126	
0.1995	1.4125	10.0000	
0.2239	1.5849	11.2202	
0.2512	1.7783	12.5893	
0.2818	1.9953	14.1254	
0.3162	2.2387	15.8489	
0.3548	2.5119	17.7828	
0.3981	2.8184		

Table 7.1: The parameters of the SSP models that were fitted to the companion galaxy spectra.

populations, not the luminosity.

The companion galaxy spectrum and its variance are then loaded and first convolved with a Gaussian kernel to lower their spectral resolution to a value equal to that of the model spectra. Then they can be prepared in the same way to the model spectra, namely by logarithmically rebinning them along the wavelength axis and normalising them by dividing through by the median of the companion galaxy spectrum and its square respectively.

Now that the galaxy spectrum is defined in logarithmic space, it can easily be de-redshifted by applying a constant offset to each of the wavelength bins, this is done to minimise the velocity shift which has to be applied during the fitting routine to increase its reliability. A secondary reason for de-redshifting the spectrum is so that the rest frame wavelength is known, meaning that the area of the spectrum that overlaps the model spectra can be determined and the rest masked out. The

wavelengths affected by common nebular emission lines are also added to the mask to ensure that they do not bias the fit along with regions of the spectrum badly affected by telluric features.

With all spectra are prepared, the pPXF fitting routine (Cappellari & Emsellem, 2004) can be run to find the weightings of each model spectrum required to best fit the galaxy spectrum. There are a number of settings which can be used to change the way the fit is calculated and are discussed below.

When fitting the spectra, the LSF is represented by a Gaussian-Hermite polynomial, the order of which is a variable in the fitting routine. For all of the work done in this chapter the order was set at 2, meaning that only a simple Gaussian was used with the components corresponding to the velocity and the dispersion of the spectrum. This means that no bias parameter is set as this only applies to the orders above 2.

Two different Legendre polynomials can be applied to the spectrum to increase the goodness of fit by removing the continuum slopes. In this work a multiplicative polynomial was used, with no additive polynomial. This is in order to conserve the flux of the galaxy spectrum to ensure the final estimate for the mass-to-light ratio was accurate. The multiplicative order was set at 15 to improve the fit by reducing only the lower frequency residuals.

A spectrum cleaning routine is included that uses iterative sigma clipping to remove bad pixels that weren't already masked out, for example residual gas emission or telluric features. This routine requires realistic errors such that the reduced χ^2 is approximately equal to 1. As discussed in section 7.1.2, it was found that for the original fits the reduced χ^2 was much greater than 1, so clean was not used and emission lines and telluric regions not already masked out were added to the mask manually. The fit looked appropriate, so the noise was corrected by scaling it by the square root of the reduced χ^2 .

It is possible to include multiple components in the fit, each with different velocities. However, this work has been performed on an integrated spectrum for each galaxy assuming that all relative velocities internal to the galaxy are negligible. This will result in a broadening of the spectral features — especially in the companion of Arp 146, which appears to have a rotating disk — but this should not affect the derived mass-to-light ratios that depend on the overall shape of the spectrum rather than the width of the features.

The weights that give the best fit to the data can result in unphysical solutions, with only a handful of non-zero weights scattered about the array of ages and metallicities as depicted in figures 7.3 and 7.5. These zero-weighted templates imply no background star formation, but instead a series of isolated starbursts indicated by the non-zero templates. To better interpret the results, a degree of regularisation can be applied to the age-metallicity array (Cappellari & Emsellem, 2004). This works by placing an upper limit on the difference of two adjacent weights in this array so that the weightings vary more smoothly, but are still statistically consistent with the data and estimated noise. These regularised results can be interpreted as a longer period of steadier star formation, with fewer dramatic variations in the metallicities of the stellar populations.

Two different regularisations were performed with the derived results treated like limits by using both the maximum and minimum values. While the minimum regularisation value is 0, the maximum value has to be determined using an iterative procedure whereby it is increased until the χ^2 value increases from its minimum by $\sqrt{2 \times N}$, where N is the number of pixels included in the fit. This then represents the smoothest fit which is still statistically consistent with the data and therefore the most consistent star formation history, as opposed to the bursty history implied by the best fit, with no regularisation.

After the fitting routine has been run it outputs the results, which include: the

spectrum which best fits the data given the settings used, the velocity and dispersion of the stellar populations, the normalised weights of the SSP models that make up the fit as well as the reduced χ^2 value of the fit. These are then used to find the implied mass-to-light ratio of the galaxy as the next step towards estimating the total stellar mass.

7.1.2 Estimating mass to light ratios

To estimate the stellar mass-to-light ratios, the important results from the fitting procedure are the arrays of the weights of each model template. This is because each stellar population has a known individual mass-to-light ratio in both the V- and the R-band, which can then be combined to find the total mass-to-light ratio using equation 5.15.

As previously mentioned, best estimates of the mass-to-light ratios were found using both no regularisation and with the maximum consistent value. For each of these, a value is found for the galaxy in the V and R photometric bands.

With the mass-to-light ratio estimated, the next stage is to find its uncertainty. This can not be robustly estimated using only data from the fit itself, so instead a bootstrapping technique is used due to its relative accuracy given its simplicity.

For each regularisation setting, the intrinsic galactic spectrum is assumed to be that of the best fit, which then implies that the residuals from the fit represent the noise on the spectrum. This means that additional, synthetic, observations of the galactic spectrum with the same noise statistics can then be created by resampling the residuals adding them to the best fit. There was no restriction on the number of times a given portion of the residuals was allowed to be used in a single synthetic dataset. To preserve any smaller-scale correlation in the noise along the wavelength axis, the residuals are block resampled by randomly choosing sections of residuals 250 pixels in length and joining them to one another, as opposed to simply choosing a

random residual or drawing a value from a probability distribution that represents the residuals.

The fitting routine can then be repeated on each instance of the synthetic spectrum, taking care to use the same settings, and the mass-to-light ratios calculated to build up a sample of measured values. After repeating the fitting routine on 100 synthetic spectra, the standard deviation of these results was taken to be the uncertainty in the best estimate of the intrinsic value as this represents how much the measured mass-to-light ratio could vary given the known best-fit value.

7.1.3 Mass determination

To calculate the total stellar mass from the mass-to-light ratios they must be multiplied by the luminosity, which is found from the extracted galaxy spectrum.

The galaxy spectra were extracted aperture photometry performed on the final datacubes. The oval apertures used were positioned manually by inspecting images of the MUSE fields of view created by collapsing the datacubes along the wavelength axis to increase the signal to noise ratios. These images and the masks used can be seen in figure 7.1. The size of the masks was chosen by increasing them out to the radius of the minimum flux levels found between the adjacent galaxies. This was done as a compromise between capturing more flux from the companion galaxy without contamination from the ring galaxy. As can be seen in figure 7.1, the process of positioning the mask for Arp 146 was complicated by the companions galaxy's asymmetric morphology.

The luminosities of the companion galaxies were calculated in both the V and R photometric bands. This was done by multiplying the spectrum by each filter curve with the resultant product integrated along the wavelength axis to get the total flux, f , of the galaxy in units of $\text{ergs cm}^{-2} \text{s}^{-1}$.

To remove the area dependency of the flux the distance to the galaxy is needed,




Figure 7.1: The masks used for extracting the spectra from the companion galaxies to Arp 146 (*left*) and AM0058-220 (*right*).

with the luminosities, L , given by:

$$L = 4\pi d^2 f. \quad (7.1)$$

The distances in table 6.6 are used to convert the measured integrated fluxes into luminosities, which are then converted into stellar units. The mass-to-light ratios can then be multiplied by these measured luminosities to get the estimates for the total stellar mass of the galaxies in solar masses.

7.2 Discussion

7.2.1 Spectrum fitting

The fits to the de-redshifted spectra of the companion galaxies to Arp 146 and AM0058-220 are shown in figures 7.2 and 7.4 respectively, along with their residuals. The regions of the spectrum that were masked out during the fitting procedure are delineated by the vertical green lines, with the residuals inside these regions set to zero.

The fits to the spectrum of the companion galaxy to Arp 146 shown in figure 7.2




Figure 7.2: The spectrum and fits for companion galaxy to Arp 146 and corresponding residuals. The spectrum is split in two for clarity.




Figure 7.3: The normalised weights for each model in the fit to the spectrum of the companion galaxy in the Arp 146 system as a function of age and $[M/H]$. *Left*: the weights for the fit with maximum regularisation still statistically consistent with the data. *Right*: the weights for the best fit to the data. Note the different colour scales used for each graph for clarity.

look good, with the regularisation of the template weights visually making little difference. The spectral features are well reproduced outside of the regions affected by nebular emission, especially for the wavelengths below $5800 \text{ \AA}$, where the absorption lines have are stronger.

The weights of each model template are shown in figure 7.3, where the effects of the regularisation of the fit can clearly be seen. The best fit to the spectrum is composed of just 8 of the 300 templates, which can broadly be grouped into two populations — one with an age of $\sim 10^1 \text{ Gyr}$ with $[M/H] \approx -1.3$, and one with an age of $\sim 10^{0.4} \text{ Gyr}$ with $[M/H] \approx 0$. In contrast, when the weights are maximally regularised, there is one very broad population with much lower component weightings. Most stars appear to be old and relatively metal-rich, with the highest weight having $[M/H] \sim 0.2$ and an age of $\sim 10^{1.3} \text{ Gyr}$. The weights then gradually decline from the peak, with very few stars younger than $\sim 1 \text{ Gyr}$.

The best fit appears to suggest that there were two periods of star formation, one just over 10 Gyr ago and then a second one approximately 2.5 Gyr ago. However, the regularisation results show that this is not necessarily the case, with one long period

of star formation finishing about 1 Gyr ago being consistent with the data. These two results therefore give upper and lower bounds on the star formation, with other histories involving more shorter bursts being possible as well.

As previously alluded to, the regularisation has only a very small effect on the shape and size of the residuals, with there being a large amount of overlap in the plots. This is also seen in the rms values, which can be found in table 7.2, and show a difference of just 0.0002 on a normalised flux with an average of 1. Neither set shows any large-scale structure, only a reasonably constant scatter about zero with wavelength.

The fits to the spectrum of the companion galaxy to AM0058-220 are shown in figure 7.4, and again show the small effect of maximising the regularisation of the template weights. The spectral features are once more well reproduced in the unmasked regions, to a similar level to the one seen in figure 7.4.

The weights of each model template are shown in figure 7.5, where the effects of the regularisation of the fit are just as clear as with Arp 146. This time the best fit to the spectrum is composed of even fewer templates: just 4 out of the 156 that could have been included. This time however, there does not appear to be different groups of populations, with all of the templates having an age of $\approx 10^{1.2}$ Gyr. There are non-zero templates spanning the whole range of metallicities, but the only two significant values are $[M/H]$ of 0.2 and -0.4 — the higher end of the range. When the weights are maximally regularised, the pattern seen is very similar in appearance to the results for Arp 146, with peak value of $[M/H] \sim 0.2$ and an age of $\sim 10^{1.3}$ Gyr. The weights decline more quickly from the peak though, with the vast majority of stars being older than $\sim 10^{0.8}$ Gyr.

These results imply a shorter period of star formation than for the companion to Arp 146. The best fit would suggest that this was very short — appearing nearly instantaneous with the resolution of this work — although it could have lasted until




Figure 7.4: The spectrum and fits for companion galaxy to AM0058-220 and corresponding residuals. The spectrum is split in two for clarity.

$10^{0.8}$ Gyr ago and still be consistent with the data.

The regularisation for AM0058-220 has a similarly small effect on the shape and size of the residuals as with Arp 146. The rms values are presented in table 7.2, and show a difference of just 0.0003 in the flux that is normalised to have an average value of 1. The residuals again show very few large-scale deviations from 0, with a reasonably constant magnitude with wavelength.

Galaxy	χ^2	rms _{unregul}	rms _{regul}
Arp 146	10.14	0.0147	0.0149
AM0058-220	27.06	0.0163	0.0166

Table 7.2: The parameters measuring the quality of the fits to the galaxy spectra: the reduced χ^2 value of the best fit alongside the rms values of the residuals of both the best fit and the maximally regularised fit.

Table 7.2 contains the χ^2 values for the fits before they were corrected to 1 by scaling the measured errors on the galactic spectra. These show that uncertainties on the galactic spectra, which were derived from the variance cubes created during data reduction, were underestimated by a factor of between about 3 and 5. This is not unexpected, as the variance cubes were statistically derived and known not to be fully rigorous, but were used to estimate the relative weights to place on each pixel of the spectrum. This is important for regions of the spectrum affected by telluric features for example, where the data are known to not be as trustworthy as other parts of the spectrum so that undue importance is not given to them.

7.2.2 Mass-to-light ratios

The mass-to-light ratios for the galaxies, calculated using equation 5.15, are shown in table 7.3. As can be seen, four different values were calculated, one for each combination of V- and R-band as well as unregularised and maximally regularised.

The mass-to-light ratios for AM0058-220 are 3–4 times higher than for Arp 146,




Figure 7.5: The normalised weights for each model in the fit to the spectrum of the companion galaxy in the Arp 146 system as a function of age and $[M/H]$. *Left*: the weights for the fit with maximum regularisation still statistically consistent with the data. *Right*: the weights for the best fit to the data. Note the different colour scales used for each graph for clarity. The two plots use different colour scales for clarity.

Galaxy	$\text{mlr}_{v, \text{unreg}}$	$\text{mlr}_{v, \text{reg}}$	$\text{mlr}_{r, \text{unreg}}$	$\text{mlr}_{r, \text{reg}}$
Arp 146	1.75 ± 0.23	1.75 ± 0.06	2.24 ± 0.32	2.27 ± 0.07
AM0058-220	7.34 ± 0.30	6.23 ± 0.11	8.33 ± 0.32	7.21 ± 0.11

Table 7.3: The mass-to-light ratios of both companion galaxies for each combination of photometric band and regularisation.

depending on which values are compared, at approximately 6–8 times that of the sun. This is consistent with expectations based on a visual inspection of the spectra and template weightings, with the redder light and higher proportion of older stellar populations in AM0058-220 pointing towards the higher values.

Comparing like values computed from the V-band data with those from the R-band show a similar trend for both galaxies, with the mass-to-light ratios having a higher value in the latter. This is again an expected result, with the redder light containing a larger proportion from older stellar populations. However, the percentage increase is not the same for each galaxy, being $\sim 30\%$ for Arp 146, but only $\sim 15\%$ for AM0058-220. This is explained by the higher mass fraction of the younger stellar populations in Arp 146, meaning that increasing the amount of light from the older stars has a larger effect on the mass-to-light ratio.




Figure 7.6: Histograms showing the distributions of the residuals from the best fits to the spectra of the companions of Arp 146 (*left*) and AM0058-220 (*right*). Shown are both the measured residuals (*red*) and the synthetic residuals created by resampling the measured values (*blue*).

When the effects of regularising the weights on the mass-to-light ratios was analysed, it was found to have no significant effect for Arp 146 but to decrease the value for AM0058-220. This is likely due to the bimodal distribution in the age-metallicity array for the former. Therefore, new non-zero weights from the regularisation mostly fill in gaps between the two populations, whereas for the latter it had the effect of broadening the overall distribution in only one direction, with the youngest populations included being younger while the oldest populations are constrained by the age of the universe.

The uncertainties in the mass-to-light ratios, are also shown in table 7.3. As described in section 7.1.2, when using bootstrapping techniques to find the uncertainties on the mass-to-light ratios, sections of the residuals are resampled with replacement allowed to create synthetic residuals. When a histogram of the synthetic, resampled residuals is overlaid with a histogram of real residuals it can be seen that this technique results in a distribution very similar to the original, as illustrated in figure 7.6, meaning that the synthetic spectra which are created appear to be very statistically similar to the measured spectrum.

The histograms showing the probability distributions of the mass-to-light ratios




Figure 7.7: Histograms showing the distributions of the mass-to-light ratios of the companion galaxy to Arp 146 found using bootstrapping techniques. Shown are the distributions from the maximally regularised fits (*left*) and the best fits (*right*), with values measured using data from the V-band shown in green and from the R-band shown in red.

are shown in figures 7.7 and 7.8 for Arp 146 and AM0058-220 respectively. From both the histograms and the values in table 7.3 it can be seen that the choice of filter has very little effect on the magnitude of the uncertainties in the mass-to-light ratios, whereas increasing the regularisation of the weights does lower them. This is because the effect of the regularisation on the uncertainties is to place an upper limit on the allowed difference between adjacent weights, so the possible alternate solutions found during bootstrapping are more constrained and so have less scope to vary when compared to the unregularised fits, which typically have a handful of weights which are not correlated to their neighbours.

7.2.3 Companion masses

In determining the final stellar mass estimates for the companion galaxies, their luminosities were calculated, which involved multiplying the spectrum by the both V- and R-band photometric filter curves. The results of which are shown in figures 7.9 and 7.10 for Arp146 and AM0058-220 respectively.

To check the photometric calibration of the companion galaxy spectra their mag-




Figure 7.8: Histograms showing the distributions of the mass-to-light ratios of the companion galaxy to AM0058-220 found using bootstrapping techniques. Shown are the distributions from the maximally regularised fits (*left*) and the best fits (*right*), with values measured using data from the V-band shown in green and from the R-band shown in red.

nitudes and colours were measured and compared with survey catalogue photometry. The results of this work can be seen in table 7.4, where Arp 146 was compared with data from the Sloan Digital Sky Survey (Blanton et al., 2017) and AM0058-220 with Pan-STARRS1 (Chambers et al., 2016). It can be seen that all of the photometric values were within 0.04 magnitudes of the catalogued values, meaning that the overall photometric calibration is in very good agreement with these surveys. The comparison of the colours is not quite as good, but they are still both within 0.07 magnitudes, which lends confidence to the slope of the spectra as well. The photometric bands used were r and i, as opposed to V and R used for the mass-to-light ratios, as these were the ones for which the filter curves were entirely covered by the MUSE wavelength range. However, they still provide confidence that the heights and slopes of the spectra measured by MUSE are accurate.

The calculated luminosities for each companion galaxy are shown in table 7.5 and show how the companion to Arp 146 is brighter than the companion to AM0058-220 by a factor of approximately 1.68 in the V-band and 1.62 in the R-band. The fact that Arp 146 has the brighter companion is not a surprising result considering the




Figure 7.9: The spectrum of the companion galaxy to Arp 146 (*black*), with the effect of multiplying by the V-band filter curve shown in green and R-band in red.

Galaxy	Instrument	r	i	r-i
Arp 146	MUSE	16.68	16.22	0.46
	SDSS	16.71	16.18	0.53
AM0058-220	MUSE	15.99	15.60	0.39
	Pan-STARRS1	15.981	15.594	0.387

Table 7.4: Measured magnitudes and colours of the ring galaxy companions with comparisons to values from survey catalogues.

shapes of the two spectra, which suggest that it has had more recent star formation, resulting in a much larger fraction of younger, higher luminosity stars. This is also shown by the difference between the two photometric bands, with the lower ratio in the R-band also caused by the higher fraction of older stars.

The values for the companion masses, derived from the estimated mass-to-light ratios and the measured luminosities, are shown in table 7.6. The higher luminosity of the Arp 146 companion reduces the ratio of the masses between the companions when compared to the mass-to-light ratios, though it still has a lower value, with AM0058-220 being roughly twice as massive at $\sim 10^{10} M_{\odot}$.




Figure 7.10: The spectrum of the companion galaxy to AM0058-220 (*black*), with the effect of multiplying by the V-band filter curve shown in green and R-band in red.

Galaxy	L_v	L_r
Arp 146	0.195 ± 0.005	0.214 ± 0.005
AM0058-220	0.116 ± 0.008	0.132 ± 0.009

Table 7.5: The estimates for the luminosities of the companion galaxies of Arp 146 and Am0058-220 in $10^{10}L_{\odot}$

Given the measured luminosities of the galaxies do not depend on the regularisation of the fitted stellar models, the same effects can be seen as in the mass-to-light ratios. Therefore the regularisation does not change the estimated masses of the companion to Arp 146, but does for AM0058-220, meaning that the expected mass is the same for both bursty and more continuous star formation rate histories for the former but if the history was more bursty for the latter then the expected mass goes up. Once again, the measured mass is higher when measured in the R-band, and this expected to be the more accurate of the two measurements.

To obtain the final best estimate for the mass of each companion, the information from each of the photometric bands and regularisation settings is analysed. As the

Galaxy	$M_{v, unreg}$	$M_{v, reg}$	$M_{r, unreg}$	$M_{r, reg}$
Arp 146	0.34 ± 0.05	0.34 ± 0.01	0.48 ± 0.05	0.49 ± 0.01
AM0058-220	0.85 ± 0.07	0.72 ± 0.05	1.10 ± 0.06	0.95 ± 0.05

Table 7.6: The estimates for the total stellar mass of the companion galaxies of Arp 146 and Am0058-220 for each combination of photometric band and regularisation. All units are $10^{10}M_{\odot}$

unregularised and regularised fits are analogous to bursty and more continuous star formation respectively the different settings of this variable are treated as limits and therefore given equal weighting in determining the best estimate, with their difference put towards a first order estimate of the uncertainty on the mass. The contribution from the measurements using different parts of the spectrum is more complicated however. The R-band was favoured over the V- due to the reduced extinction associated with longer wavelengths of light internal to the companion galaxy. The second factor which favours redder light is the fact that the masses of galaxies are dominated by smaller, redder stars which can be greatly outshone by a smaller number of larger, bluer stars. This discrepancy between the masses and luminosities is smaller in redder light, meaning that the light is a more accurate way of estimating the mass.

For Arp 146, the star formation history appears to have a much smaller effect on the measured masses. This means that the unregularised R-band fits were taken as the best estimate as the magnitude is same within uncertainties of the regularised results and the higher uncertainties are more representative of the effect of the lack of knowledge about the star formation history.

For AM0058-220, there does appear to be a significant effect on the estimated mass caused by the star formation history, and so each limit is taken to be as likely as each other with the mean of the two results taken to be the best estimate. Half of the range is then taken to be representative of the uncertainties introduced by the different star forming scenarios.

One large uncertainty, which cannot be measured, is the dark matter fraction for

each galaxy, which arises due to the use of the stellar mass as opposed to a dynamical mass. Thus, no effort has been made to estimate the mass of dark matter in each galaxy. Realistically, it means that these values are only lower limits on the masses, although a statistical look at galaxy samples may be used to try gauge the magnitude of the uncertainties.

Cappellari et al. (2013b) suggest the dark matter fraction for early-type galaxies is 0.1-0.4 within a radius equal to r_e at a mass of $10^{10}M_{\odot}$. These estimates come from the ATLAS^{3D} project (Cappellari et al., 2011): a volume limited survey of 260 early type galaxies. Courteau & Dutton (2015) present data that suggest that for disc galaxies the fraction is higher, at 0.4-0.8 for radii less than 2.2 times the disc scale length at a mass of $10^{10}M_{\odot}$.

A visual inspection of the galaxies led to the companion to Arp 146 being classified as a disc galaxy and the companion to AM0058-220 as an early-type galaxy. This lead to them having an estimated dark matter fraction for the purpose of estimating the uncertainty of 0.6 and 0.25 respectively. These values were use to estimate an increased upper uncertainty of $0.7 \cdot 10^{10}M_{\odot}$ and $0.3 \cdot 10^{10}M_{\odot}$ respectively.

7.3 Ring galaxy system mass ratios

The final results for the companion mass estimates are shown in table 7.7, alongside the estimated ring galaxy masses and the calculated ratio between the two. It can be seen that in both systems the mass of the ring is estimated to be significantly higher than that of the companion, with both systems measured to have a lower limit to the mass ratio of about 1/6 and estimated uncertainties that are not consistent with mass ratios of order unity.

One major source of uncertainty on the mass ratios that has not been estimated is the mass of dark matter outside of the observed light distribution. This will be

Galaxy	M_{ring} ($10^{10}M_{\odot}$)	M_{comp} ($10^{10}M_{\odot}$)	$\frac{M_{ring}}{M_{comp}}$
Arp 146	3 ± 1	$0.5^{+0.7}_{-0.1}$	6^{+6}_{-3}
AM0058-220	6 ± 2	$1.0^{+0.3}_{-0.1}$	6^{+6}_{-1}

Table 7.7: The final estimates for the mass of the ring galaxies and their companions.

non-zero for both the ring and its companion, but the ratio can obviously not be measured.

As described in Section 5.1.1, there have been several previous attempts to measure the mass ratios in ring galaxy systems, notably for the Cartwheel Galaxy, but also for other, less well studied systems, such as the sample of 9 in Romano et al. (2008). For the Cartwheel the companion galaxy is not known for certain due to there being 3 candidate galaxies in the vicinity of the ring, but none of them were of a comparable mass to the ring itself, with one estimated to be just 5–10% by Davies & Morton (1982). This was seen as something of a problem when compared to simulations of ring galaxy formation, but was explained by the contrast of the ring being enhanced by the burst of star formation lowering the mass-to-light ratio in the ring, while the simulations were focussed on forming a ring purely as a physical structure in the galaxy.

In the work of Romano et al. (2008) though, the larger sample estimates that a wide range of mass ratios that can form ring galaxies. These are as low as around 5% for IZw45 and NGC 2793 but as high as around 200% for three systems in the Arp catalogue: numbers 141, 147 and 148. In that work the authors categorised the ratios into three bins: those with comparable masses (a ratio of < 2.5), those with moderate ratios (between 2.5 and 6.3) and high ratios (> 6.3). Interestingly, the galaxies mostly fell into the comparable and high ratio bins, but with so few data points no conclusions can be reliably drawn from this fact.

Given the rarity of ring galaxy systems, this work adds significantly to the small

sample of mass ratio estimates. Here, both estimates fall into the moderate ratio bin that was occupied solely by IIHz4 before, and suggest that the previous lack of moderate ratios was probably due to the small sample size. This would be encouraging if still true after the sample size is increased further, as from the simulations there is no reason to suggest a preference for certain values. There still hasn't been mass ratio of larger than about 200% of the ring measured, which implies a lower probability of a clean ring being identified in a collision with the higher interaction levels that the extra mass would generate.

Comparing these ratios with the mass ratios used in simulations, as summarised in table 5.2, it can be seen that they are very similar. The lower bounds presented here are similar to the lower ratios used in many simulations, and the asymmetric uncertainties suggest the ratios will fall closer to the average of about 0.5 used in simulations. This suggests that the simulations are using broadly the right mass ratios, but it would still be encouraging to see more work focused on lower values.

7.4 Conclusions

For the ring galaxy systems Arp 146 and AM0058-220, an integrated spectrum from companion galaxies was extracted and pPXF was used to fit a linear superposition of SSP models to them. The mass-to-light ratios of the stellar population models were then used to estimate the overall value for the companion galaxy and these were combined with the luminosities estimated from integrating the spectra to obtain final estimates to the stellar masses of the companion galaxies.

These masses were compared to the mass of their respective ring galaxies, which was estimated from the rotation velocity in the decomposed kinematics of the galaxy. The lower limit for this ratio was found be just one sixth of the mass for both systems, with a slightly higher value likely due to the stellar masses being compared with

dynamical masses. However, the uncertainties were estimated to be small enough that these results lend further evidence to the case that comparable mass companions are not necessary to form ring galaxies.

Appendix A

Ring galaxy simulations

This appendix contains the plots of the sinusoid parameters for the simulated rotation and expansion velocities from the work in Section 6.4.




Figure A.1: The average v_{rot} with time, showing the effects of changing b (*top left*), Δv (*top right*) and v_{rot0} (*bottom*). The solid line shows v_{rot} around the centre of galaxy, with the dashed showing v_{rot} around the centre of the ring.




Figure A.2: The amplitude of the change in v_{rot} around the ring with time, showing the effects of changing b (*top left*), Δv (*top right*) and v_{rot0} (*bottom*). The solid line shows v_{rot} around the centre of galaxy, with the dashed showing v_{rot} around the centre of the ring.




Figure A.3: The phase of the sinusoidal pattern of v_{rot} around the ring with time, showing the effects of changing b (top left), Δv (top right) and v_{rot0} (bottom). The solid line shows v_{rot} around the centre of galaxy, with the dashed showing v_{rot} around the centre of the ring.




Figure A.4: The average v_{exp} with time, showing the effects of changing b (top left), Δv (top right) and v_{rot0} (bottom). The solid line shows v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.




Figure A.5: The amplitude of the change in v_{exp} around the ring with time, showing the effects of changing b (top left), Δv (top right) and v_{rot0} (bottom). The solid line shows v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.




Figure A.6: The phase of the sinusoidal pattern of v_{exp} around the ring with time, showing the effects of changing b (top left), Δv (top right) and v_{rot0} (bottom). The solid line shows v_{exp} around the centre of galaxy, with the dashed showing v_{exp} around the centre of the ring.

Bibliography

- ALLEN, J. R., O'BRIEN, K., LYNN, J. D., THATTE, N. A., BRYSON, I., CLARKE, F., SCHNETLER, H. & TECZA, M., 2014. Cryogenic testing of components for the HARMONI spectrograph. In *Ground-based and Airborne Instrumentation for Astronomy V*, vol. 9147 of *Proc. SPIE*, 91479M.
- ALLEN, J. R., O'BRIEN, K., LYNN, J. D., THATTE, N. A., TOSH, I. A. J. & TACON, M., 2016. Characterizing the performance of cryogenic lens mounts for the HARMONI spectrograph. In *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II*, vol. 9912 of *Proc. SPIE*, 99124Q.
- AMRAM, P., MENDES DE OLIVEIRA, C., BOULESTEIX, J. & BALKOWSKI, C., 1998. The Halpha kinematic of the Cartwheel galaxy. *Astronomy and Astrophysics*, **330**, 881–893.
- APPLETON, P. N. & JAMES, R. A., 1990. *Self-consistent Simulations of Ring Galaxies*, 200–204. Springer Berlin Heidelberg, Berlin, Heidelberg.
- APPLETON, P. N. & STRUCK-MARCELL, C., 1996. Collisional Ring Galaxies. *Fundamentals of Cosmic Physics*, **16**, 111–220.
- ARNOLD, V. I., SHANDARIN, S. F. & ZELDOVICH, I. B., 1982. The large scale structure of the universe. I - General properties One- and two-dimensional models. *Geophysical and Astrophysical Fluid Dynamics*, **20**, 111–130.

- ARP, H. C. & MADORE, B., 1987. *A catalogue of southern peculiar galaxies and associations*.
- ATHANASSOULA, E., PUERARI, I. & BOSMA, A., 1997. Formation of rings in galactic discs by infalling small companions. *Monthly Notices of the Royal Astronomical Society*, **286**, 284–302.
- BACON, R., ACCARDO, M., ADJALI, L., ANWAND, H., BAUER, S. ET AL., 2010. The MUSE second-generation VLT instrument.
- BACON, R., ADAM, G., BARANNE, A., COURTES, G., DUBET, D., DUBOIS, J. P., EMSELLEM, E., FERRUIT, P., GEORGELIN, Y., MONNET, G., PECONTAL, E., ROUSSET, A. & SAY, F., 1995. 3D spectrography at high spatial resolution. I. Concept and realization of the integral field spectrograph TIGER. *Astronomy and Astrophysics Supplement*, **113**, 347.
- BACON, R., COPIN, Y., MONNET, G., MILLER, B. W., ALLINGTON-SMITH, J. R., BUREAU, M., CAROLLO, C. M., DAVIES, R. L., EMSELLEM, E., KUNTSCHNER, H., PELETIER, R. F., VEROLME, E. K. & DE ZEEUW, P. T., 2001. The SAURON project - I. The panoramic integral-field spectrograph. *Monthly Notices of the Royal Astronomical Society*, **326**, 23–35.
- BIZYAEV, D. V., MOISEEV, A. V. & VOROBYOV, E. I., 2007. Propagating Star Formation in the Collisional Ring Galaxy Arp 10. *The Astrophysical Journal*, **662**, 304–321.
- BLANTON, M. R., BERSHADY, M. A., ABOLFATHI, B., ALBARETI, F. D., ALLENDE PRIETO, C., ALMEIDA, A., ALONSO-GARCÍA, J., ANDERS, F., ANDERSON, S. F., ANDREWS, B. & ET AL., 2017. Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe. *The Astrophysical Journal*, **154**, 28.

- BUTA, R., 1999. Resonance Rings and Galaxy Morphology. *Astrophysics and Space Science*, **269**, 79–99.
- CAPPELLARI, M., 2017. Improving the full spectrum fitting method: accurate convolution with Gauss-Hermite functions. *Monthly Notices of the Royal Astronomical Society*, **466**, 798–811.
- CAPPELLARI, M. & EMSELLEM, E., 2004. Parametric Recovery of Line-of-Sight Velocity Distributions from Absorption-Line Spectra of Galaxies via Penalized Likelihood. *The Publications of the Astronomical Society of the Pacific*, **116**, 138–147.
- CAPPELLARI, M., EMSELLEM, E., KRAJNOVIĆ, D., MCDERMID, R. M., SCOTT, N. ET AL., 2011. The ATLAS^{3D} project - I. A volume-limited sample of 260 nearby early-type galaxies: science goals and selection criteria. *Monthly Notices of the Royal Astronomical Society*, **413**, 813–836.
- CAPPELLARI, M., MCDERMID, R. M., ALATALO, K., BLITZ, L., BOIS, M. ET AL., 2013a. The ATLAS^{3D} project - XX. Mass-size and mass- σ distributions of early-type galaxies: bulge fraction drives kinematics, mass-to-light ratio, molecular gas fraction and stellar initial mass function. *Monthly Notices of the Royal Astronomical Society*, **432**, 1862–1893.
- CAPPELLARI, M., SCOTT, N., ALATALO, K., BLITZ, L., BOIS, M. ET AL., 2013b. The ATLAS^{3D} project - XV. Benchmark for early-type galaxies scaling relations from 260 dynamical models: mass-to-light ratio, dark matter, Fundamental Plane and Mass Plane. *Monthly Notices of the Royal Astronomical Society*, **432**, 1709–1741.
- CENARRO, A., CARDIEL, N., GORGAS, J., PELETIER, R., VAZDEKIS, A. & PRADA, F., 2001. Empirical calibration of the near-infrared Ca II triplet I. The

- stellar library and index definition. *Monthly Notices of the Royal Astronomical Society*, **326**(3), 959–980.
- CHAMBERS, K. C., MAGNIER, E. A., METCALFE, N., FLEWELLING, H. A., HUBER, M. E. ET AL., 2016. The Pan-STARRS1 Surveys. *ArXiv e-prints*.
- COLLABORATION, P., AGHANIM63, N., ARNAUD78, M., ASHDOWN74, M., AUMONT63, J., BACCIGALUPI91, C., BANDAY103, A., BARREIRO69, R., BARTLETT, J., BARTOLO32, N. ET AL., 2016. Planck 2015 results. *Astronomy and Astrophysics*, **594**, A11.
- CONROY, C., 2013. Modeling the panchromatic spectral energy distributions of galaxies. *Annual Review of Astronomy and Astrophysics*, **51**, 393–455.
- COURTEAU, S. & DUTTON, A. A., 2015. On the global mass distribution in disk galaxies. *The Astrophysical Journal Letters*, **801**(2), L20.
- DAVIES, R. L., ALLINGTON-SMITH, J. R., BETTESS, P., CHADWICK, E., CONTENT, R. ET AL., 1997. GMOS: the GEMINI Multiple Object Spectrographs. In A. L. Ardeberg, ed., *Optical Telescopes of Today and Tomorrow*, vol. 2871 of *Proc. SPIE*, 1099–1106.
- DAVIES, R. L. & MORTON, D. C., 1982. A mass estimate for the companion to the Cartwheel galaxy. *Monthly Notices of the Royal Astronomical Society*, **201**(1), 69P–72P.
- EVANS, C. J., 2009. PVra robust amplitude parameter for optical surface specification. *Optical Engineering*, **48**(4), 043605–043605.
- FABER, S. M., 1972. Quadratic programming applied to the problem of galaxy population synthesis. *Astronomy and Astrophysics*, **20**, 361–374.

- FEW, J. M. A. & MADORE, B. F., 1986. Ring galaxies. II - Classification and statistics. *Monthly Notices of the Royal Astronomical Society*, **222**, 673–682.
- FIACCONI, D., MAPELLI, M., RIPAMONTI, E. & COLPI, M., 2012. Adaptive mesh refinement simulations of collisional ring galaxies: effects of the interaction geometry. *Monthly Notices of the Royal Astronomical Society*, **425**(3), 2255–2266.
- FLORENCE LAURENT, EDGARD RENAULT, H. A. D. B. P. C. J. K. M. L. H. N. W. S. Y. S. W. X., 2014. MUSE field splitter unit: fan-shaped separator for 24 integral field units.
- FOGARTY, L., THATTE, N., TECZA, M., CLARKE, F., GOODSALL, T., HOUGHTON, R., SALTER, G., DAVIES, R. L. & KASSIN, S. A., 2011. SWIFT observations of the Arp 147 ring galaxy system. *Monthly Notices of the Royal Astronomical Society*, **417**, 835–844.
- FOSBURY, R. A. E. & HAWARDEN, T. G., 1977. A0035 'the cartwheel' a large southern ring galaxy. *Monthly Notices of the Royal Astronomical Society*, **178**, 473–487.
- FREEMAN, K. C. & DE VAUCOULEURS, G., 1974. An interpretation of ring galaxies and the properties of intergalactic gas clouds. *The Astrophysical Journal*, **194**, 569–571.
- FROUD, T. R., TOSH, I. A. J., EDESON, R. L. & DALTON, G. B., 2006. Cryogenic mounts for large fused silica lenses.
- GERBER, R. A., LAMB, S. A. & BALSARA, D. S., 1992. A model for ring galaxies - ARP 147-like systems. *The Astrophysical Journal*, **399**, L51–L54.

- GERBER, R. A., LAMB, S. A. & BALSARA, D. S., 1996. A stellar and gas dynamical numerical model of ring galaxies. *Monthly Notices of the Royal Astronomical Society*, **278**(2), 345–366.
- GIRARDI, L., BRESSAN, A., BERTELLI, G. & CHIOSI, C., 2000. Evolutionary tracks and isochrones for low-and intermediate-mass stars: From 0.15 to $7M_{\odot}$, and from $Z = 0.0004$ to 0.03 . *Astronomy and Astrophysics Supplement Series*, **141**(3), 371–383.
- HEAP, S. R. & LINDLER, D. J., 2007. Hubble’s Next Generation Spectral Library (NGSL). In A. Vallenari, R. Tantalò, L. Portinari & A. Moretti, eds., *From Stars to Galaxies: Building the Pieces to Build Up the Universe*, vol. 374 of *Astronomical Society of the Pacific Conference Series*, 409.
- HERNQUIST, L. & WEIL, M. L., 1993. Spokes in Ring Galaxies. *Monthly Notices of the Royal Astronomical Society*, **261**, 804.
- HIGDON, J. L., 1993. *A multi wavelength observational study of the Cartwheel ring galaxy*. Ph.D. thesis, Texas Univ., Austin.
- HIGDON, J. L., 1996. Wheels of Fire. II. Neutral Hydrogen in the Cartwheel Ring Galaxy. *The Astrophysical Journal*, **467**, 241.
- HUANG, S. N. & STEWART, P., 1988. Ring galaxies with self-consistent haloes. *Astronomy and Astrophysics*, **197**, 14–28.
- LYNDS, R. & TOOMRE, A., 1976. On the interpretation of ring galaxies: the binary ring system II Hz 4. *The Astrophysical Journal*, **209**, 382–388.
- MADORE, B. F., NELSON, E. & PETRILLO, K., 2009. Atlas and Catalog of Collisional Ring Galaxies. *The Astrophysical Journal Supplement*, **181**, 572–604.

- MAPELLI, M. & MAYER, L., 2012. Ring galaxies from off-centre collisions. *Monthly Notices of the Royal Astronomical Society*, **420**(2), 1158–1166.
- MASTERS, R. J., LEWIS, I. J., TOSH, I. A. J., TECZA, M., LYNN, J., WATKINS, R. E. J., CLACK, A., DAVIES, R. L., THATTE, N. A., TACON, M., MAKIN, R., TEMPLE, J. & PEARCE, A., 2010. KMOS: assembly, integration and testing of three 0.8-2.5 micron spectrographs.
- MERSTALLINGER, A., SALES, M., SEMERAD, E. & DUNN, B., 2009. Assessment of cold welding between separable contact surfaces due to impact and fretting under vacuum. *ESA Scientific & Technical Memoranda*, **279**, 57.
- NIRANJAN THATTE, MATTHIAS TECZA, F. C. T. G. J. L. D. F. R. L. D., 2006. The Oxford SWIFT integral field spectrograph.
- O'BRIEN, T. P. & ATWOOD, B., 2003. Lens mounting system for cryogenic applications.
- POSTON, T. & STEWART, I., 2014. *Catastrophe theory and its applications*. Courier Corporation.
- RASBAND, W. S., 2012. ImageJ: Image processing and analysis in Java. Astrophysics Source Code Library.
- RAYNER, J. T., CUSHING, M. C. & VACCA, W. D., 2009. The infrared telescope facility (IRTF) spectral library: cool stars. *The Astrophysical Journal Supplement Series*, **185**(2), 289.
- ROBITAILLE, T. P. & WHITNEY, B. A., 2010. The present-day star formation rate of the Milky Way determined from Spitzer-detected young stellar objects. *The Astrophysical Journal Letters*, **710**(1), L11.

- ROMANO, R., MAYYA, Y. D. & VOROBYOV, E. I., 2008. Stellar Disks of Collisional Ring Galaxies. I. New Multiband Images, Radial Intensity and Color Profiles, and Confrontation with N-Body Simulations. *Astronomical Journal*, **136**, 1259–1289.
- SÁNCHEZ-BLÁZQUEZ, P., PELETIER, R. F., JIMÉNEZ-VICENTE, J., CARDIEL, N., CENARRO, A. J., FALCÓN-BARROSO, J., GORGAS, J., SELAM, S. & VAZDEKIS, A., 2006. Medium-resolution Isaac Newton Telescope library of empirical spectra. *Monthly Notices of the Royal Astronomical Society*, **371**, 703–718.
- SMITH, R., LANE, R. R., CONN, B. C. & FELLHAUER, M., 2012. Numerical modelling of Auriga’s Wheel - a new ring galaxy. *Monthly Notices of the Royal Astronomical Society*, **423**, 543–557.
- STRUCK, C., 2010. Applying the analytic theory of colliding ring galaxies. *Monthly Notices of the Royal Astronomical Society*, **403**, 1516–1530.
- STRUCK-MARCELL, C. & LOTAN, P., 1990. The varieties of symmetric stellar rings and radial caustics in galaxy disks. *The Astrophysical Journal*, **358**, 99–115.
- TAMAI, R., CIRASUOLO, M., GONZÁLEZ, J. C., KOEHLER, B. & TUTI, M., 2016. The E-ELT program status. In *Ground-based and Airborne Telescopes VI*, vol. 9906 of *Proc. SPIE*, 99060W.
- THATTE, N. A., CLARKE, F., BRYSON, I., SHNETLER, H., TECZA, M. ET AL., 2016. The E-ELT first light spectrograph HARMONI: capabilities and modes. In *Ground-based and Airborne Instrumentation for Astronomy VI*, vol. 9908 of *Proc. SPIE*, 99081X.
- THEYS, J. C. & SPIEGEL, E. A., 1976. Ring galaxies. I. *The Astrophysical Journal*, **208**, 650–661.

- THEYS, J. C. & SPIEGEL, E. A., 1977. Ring galaxies. II. *The Astrophysical Journal*, **212**, 616–619.
- TOOMRE, A., 1978. Interacting systems. In M. S. Longair & J. Einasto, eds., *Large Scale Structures in the Universe*, vol. 79 of *IAU Symposium*, 109–116.
- VALDES, F., GUPTA, R., ROSE, J. A., SINGH, H. P. & BELL, D. J., 2004. The Indo-US Library of Coudé Feed Stellar Spectra. *Astrophysical Journal Supplement*, **152**, 251–259.
- VAZDEKIS, A., CASUSO, E., PELETIER, R. F. & BECKMAN, J. E., 1996. A New Chemo-evolutionary Population Synthesis Model for Early-Type Galaxies. I. Theoretical Basis. *Astrophysical Journal Supplement*, **106**, 307.
- VAZDEKIS, A., SÁNCHEZ-BLÁZQUEZ, P., FALCÓN-BARROSO, J., CENARRO, A. J., BEASLEY, M. A., CARDIEL, N., GORGAS, J. & PELETIER, R., 2010. MILES SSP Models. In G. R. Bruzual & S. Charlot, eds., *Stellar Populations - Planning for the Next Decade*, vol. 262 of *IAU Symposium*, 65–68.
- WHITMORE, B. C., LUCAS, R. A., McELROY, D. B., STEIMAN-CAMERON, T. Y., SACKETT, P. D. & OLLING, R. P., 1990. New observations and a photographic atlas of polar-ring galaxies. *Astronomical Journal*, **100**, 1489–1522.
- WONG, O. I., MEURER, G. R., BEKKI, K., HANISH, D. J., KENNICUTT, R. C. ET AL., 2006. NGC922 - a new drop-through ring galaxy. *Monthly Notices of the Royal Astronomical Society*, **370**, 1607–1611.
- ZWICKY, F., 1941. *Hydrodynamics and the structure of stellar systems*, 137. California Institute of Technology, Pasadena.