Durham E-Theses

Astronomy with integral field spectroscopy:: observation, data analysis and results

Turner, James Edmund Henry

How to cite:

Use policy

Astronomy with Integral Field Spectroscopy: Observation, Data Analysis and Results

James Edmund Henry Turner

A Thesis submitted to the University of Durham
in accordance with the regulations for
admittance to the Degree of Doctor of Philosophy.

Department of Physics, University of Durham

September 2001

27 JUN 2002

To my parents
and Louisa.

Astronomy with Integral Field Spectroscopy: Observation, Data Analysis and Results

James Edmund Henry Turner

September 2001

Abstract

With a new generation of facility instruments being commissioned for 8 metre telescopes, integral field spectroscopy will soon be a standard tool in astronomy, opening a range of exciting new research opportunities. It is clear, however, that reducing and analyzing integral field data is a complex problem, which will need considerable attention before the full potential of the hardware can be realized. The purpose of this thesis is therefore to explore some of the scientific capabilities of integral field spectroscopy, developing the techniques needed to produce astrophysical results from the data.

Two chapters are dedicated to the problem of analyzing observations from the densely-packed optical fibre instruments pioneered at Durham. It is shown that, in the limit where each spectrum is sampled by only one detector row, data may be treated in a similar way to those from an image slicer. The properties of raw fibre data are considered in the context of the Sampling Theorem and methods for three dimensional image reconstruction are discussed. These ideas are implemented in an IRAF data reduction package for the Thousand Element Integral Field Unit (TEIFU), with source code provided on the accompanying compact disc.

Two observational studies are also presented. In the first case, the 3D infrared image slicer has been used to test for the presence of a super-massive black hole in the giant early-type galaxy NGC 1316. Measurements of the stellar kinematics do not reveal a black hole of mass $5 \times 10^9 M_\odot$, as predicted from bulge luminosity using the relationship of Kormendy & Richstone (1995). The second study is an investigation into the origin of [Fe II] line emission in the Seyfert galaxy NGC 4151, using Durham University's SMIRFS-IFU. By mapping [Fe II] line strength and velocity at the galaxy centre, it is shown that the emission is associated with the optical narrow line region, rather than the radio jet, indicating that the excitation is primarily due to photoionizing X-rays.

Finally, a report is given on the performance of TEIFU, which was commissioned at the William Herschel Telescope in 1999. Measurements of throughput and fibre response variation are given and a reconstructed test observation of the radio galaxy 3C 327 is shown, demonstrating the functionality of the instrument and software.

Preface

The work described in this thesis was undertaken whilst the candidate was a student of Dr. J. R. Allington-Smith at the Department of Physics of the University of Durham. No part of the content has been submitted for any other degree.

The project discussed in chapter 2 was undertaken as part of a collaboration with Prof. R. Davies and Dr. H. Kuntschner of Durham University, Prof. J. Mould, formerly Director of Mount Stromlo and Siding Spring Observatories and Prof. T. deZeeuw of Leiden University. The long-slit and aperture spectroscopy data were reduced by Dr. Kuntschner, who also performed kinematic measurements using FCQ. The remaining data reduction and analysis was done by the author.

Chapter 5 results from a collaboration between the author, Dr. C. Done, Dr. J. R. Allington-Smith and Dr. R. Haynes of Durham University and Drs. S. Morris and S. Chapman, both formerly of the Dominion Astrophysical Observatory, Canada. The work discussed is, however, the author's own. This chapter has recently been submitted to Monthly Notices of the Royal Astronomical Society for publication.

The concept of densely-packed fibre spectra was introduced by Dr. R. Content of Durham University, whose design for the Thousand Element Integral Field Unit motivated the new data reduction paradigm formulated in chapters 3 and 4.

Contents

List of Figures

Chapter 1

Introduction

1.1 Background

Unlike the slit of a traditional optical spectrograph, astronomical objects are rarely one-dimensional. Whilst aperture and slit spectroscopy have played crucial roles in our understanding of the Universe, they cannot reveal the full spatial complexity of many of its consituents. Historically, relatively well-ordered systems, such as archetypal elliptical galaxies, have been modelled with some success using constraints from slit spectroscopy. The limitations of simplistic geometric assumptions are, however, increasingly apparent. More complex objects, such as active galaxies, can be difficult enough to interpret in two projected dimensions, without discarding available information. In order to gain a more complete picture, *imaging spectroscopy* in two spatial dimensions is required.

In addition to their basic dimensional limitation, slit spectrogaphs have several technical drawbacks. They always exclude a substantial proportion of the light from a point source; in poor seeing, widening the slit improves efficiency at the expense of spectral resolution, but the seeing disc must still be truncated so that spectral line profiles are independent of the source. Aligning a narrow slit accurately with a point source can be awkward and time consuming, especially at high resolutions with adaptive optics (AO). For extended sources, errors in slit placement due to limited accuracy or incorrect prior assumptions can give misleading results, with no direct indication that there is a problem. Even when the set-up is optimized, care must be taken that spurious differences in velocity or brightness are not inferred where multiple intensity peaks along the slit are centred differently across its width. Most forms of imaging spectroscopy overcome these limitations.

During the last century, techniques for imaging spectroscopy have included slit scanning at several positions across the object, Fabry-Perot interferometry and, more recently, Fourier Transform (Michelson) interferometry. All three methods have the disadvantage that of only two out of three dimensions (two spatial and one spectral) can be resolved simultaneously, whilst the third dimension is scanned using multiple exposures. For many purposes, this is clearly inefficient and limits the number of elements which can be observed in the third dimension. Furthermore, since factors such as sky background and seeing conditions can vary on short timescales at ground-based observatories, it is difficult to achieve consistency between exposures unless the total integration time is short or the number of scan steps is small.

The underlying problem for these forms of imaging spectroscopy is a reliance on two-dimensional detectors for capturing three-dimensional data. The ideal solution would be a detector array which can instrinsically resolve colour. Such devices are currently in development, in the forms of superconducting tunnel junctions (STJs, eg. Perryman et al. 1999) and transition edge sensors (TESs, eg. Romani et al. 1999). These technologies, however, have fundamental drawbacks of their own. Whilst they offer time-resolved detection of individual photons and broad spectral response, they are limited to low spectral resolution and require extreme cryogenic cooling. In addition, current prototype arrays are still very small. It will not, therefore, be possible to replace general-purpose spectrographs with colour-sensitive cameras in the forseeable future.

In order to resolve images simultaneously in three dimensions with conventional detectors, integral field spectroscopy (IFS) is required. An integral field unit (IFU) is used to reformat images spatially so that they can be dispersed without the spectra from different points overlapping. Dedicated computer software is then used to convert the dimensionally scrambled raw data back into three-dimensional images. Whilst the first such systems were developed during the 1980s (Vanderriest, 1980; Gray et al., 1982; Courtes, 1982; Barden & Wade, 1988), IFUs are only just becoming widely available at the time of writing. The development of the relatively complex optics which are needed has been motivated by efficiency considerations as much as direct scientific requirements. Three main hardware techniques are outlined in §1.2.

With international competition for observing time on eight-metre and larger telescopes, there is an increasing emphasis on observational efficiency. For spectroscopy,

this means employing multi-object, integral field and multi-field systems to gather as many useful data per exposure as possible (eg. see Eisenhauer et al., 2000). Such designs can involve complex data reduction and observing methods, as well as intricate hardware. For integral field spectroscopy in particular, very little pre-existing software is available to satisfy sophisticated requirements for data reduction.

The purpose of this thesis is to explore some of the scientific capabilities of integral field spectroscopy and develop the techniques needed to produce astrophysical results. Focusing on Durham University's own instrumentation, some new results, data analysis methods and software will be presented. It is hoped that the discussion will also provide a useful overview of many of the issues involved in working with integral field spectroscopy.

1.2   IFS Techniques

There are three basic types of IFU, as illustrated in fig. 1.1. The following subsections describe each one in more detail. This thesis deals with data from the optical fibre and image slicer designs.

1.2.1   Microlens arrays

The microlens array, lenselet array or micropupil array IFU was pioneered by TIGER (eg. Bacon et al., 1995) on the Canada-France-Hawaii Telescope. Images are divided by the microlenses and refocused into small pupil images at the spectrograph detector. By dispersing these spots at a suitable angle, spectra can be formed in the space between the spatial elements.

This is the simplest of the three IFU types and can therefore achieve a high optical throughput. Its main drawback is the geometry, which allows only short spectra between the spatial elements. In addition, any overlaps between adjacent spectra will occur at different wavelengths, so it is necessary either to leave dead space between the spectra or remove the contamination by deconvolution during processing; either solution reduces the overall efficiency in terms of detector usage or signal-to-noise. Since telescope images are reformatted in a two dimensional

Figure 1.1: Three basic types of integral field unit: (top) microlens array IFU, (middle) optical fibre IFU, (bottom) image slicer (diagram courtesy of J.R. Allington-Smith).

pattern, rather than a slit, microlens array IFUs are better suited to integration in complete instruments than retro-fitting to existing spectrographs.

Despite the limitation on spectrum length, instruments such as OASIS (Emsellem, 1999) and SAURON (Bacon et al., 2001a) have used microlens arrays very productively for integral field spectroscopy. Because large microlens arrays are easier to fabricate than large fibre bundles or image slicers, the design has proved a good choice for sampling short spectral ranges over a relatively large field. With improvements in the other technologies, however, this advantage will diminish and the issue of efficient detector use will become more important.

1.2.2   Optical fibres

Optical fibres have been used for both multi-object and integral-field spectroscopy since around 1980 (Hill et al., 1980; Vanderriest, 1980), but have not been widely used for the latter purpose until recently. In either application, they reformat the

field of view into a one-dimensional row of points, which is dispersed perpendicular to its length. It is therefore possible to obtain full-length spectra, as for a long slit.

Bare fibre bundles can suffer from substantial throughput losses for two reasons. First, fibres are circular and their cores are surrounded by thin cladding and/or buffer material, so their *fill factor* in an array is significantly less than one. Second, when light enters fibres at a specific angle, it emerges at a small range of angles around the original (eg. Heacox, 1986; Clayton, 1989)—a phenomenon known as *focal ratio degradation* (FRD). Unless the spectrograph is enlarged to accept focal ratios faster than that of the telescope, a significant amount of light can be lost. FRD is proportionally worse for slow beams (large f-ratios) than fast ones.

The performance of fibres can be improved by combining them with microlens arrays. At the bundle input, square or hexagonal lenses can be packed continuously, providing a fill factor close to 100%. Lenses also allow the fibres to be fed at a lower focal ratio than that of the incident beam, minimizing FRD effects. At the fibre outputs, a second set of lenselets restores the focal ratio to that of the spectrograph. This is the technique which has been used successfully in Durham University's own fibre-based IFUs (Haynes et al., 1998a,b, 1999; Murray et al., 2000; Allington-Smith et al., 2000). Because fibre-microlens IFUs can reformat images into one dimension, have adjustable focal ratios and allow arbitrary beam routing, they are relatively easy to retro-fit to long-slit spectrographs.

One basic disadvantage of fibres is that their characteristics vary, depending on such factors as stress, manufacturing defects and alignment with microlenses. Thus different IFU elements will have significantly different throughputs and possibly slight differences in their spectral profiles. Whilst undesirable, the former effect can essentially be removed by flat fielding and the latter by careful design and, if necessary, beam switching fields during observing. Fibres which are broken during handling are more of a nuisance, but can be dealt with by interpolation or dithering observations, as for a CCD. Overall, fibre systems are labour intensive to build, but are cheap, versatile and effective. Scaling to much larger numbers of elements may require more automated manufacturing methods.

1.2.3   Image slicers

Image slicers have been in use throughout much of the last century, to avoid blocking out light from point sources with a spectrograph slit (Bowen, 1938); machining optics suitable for imaging purposes is, however, still a technical challenge today. The main component of this type of IFU is a slicing mirror, which consists of a stack of long, narrow reflective surfaces, each angled differently. Thus the incident image is separated into thin slices, whose positions diverge along different paths. A second set of mirrors reimages the slices end-to-end to form a one-dimensional slit. A variant design which has recently become popular is the 'Advanced Image Slicer', conceived at Durham University to overcome several practical drawbacks of previous implementations (Content, 1998).

IFUs based on image slicers have the great advantage that they can safely be cooled to cryogenic temperatures, like most other reflecting optics. They are therefore currently the only solution for IFS at wavelengths longer than $\sim 2\mu$m. Of the three basic classes of IFU, image slicers also provide the most efficient detector use, allowing maximum spatial and spectral coverage. As explained in chapter 3, fibre-lenselet IFUs can also approach this level of efficiency, but with notable complications for data reduction.

Image slicers have a few minor technical disadvantages with respect to fibre and microlens array designs. They can suffer slightly from optical aberrations, such as off-centre defocus. Also, since they do not scramble the light within each image element, the shapes and positions of spectral profiles at the detector depend on the incident image, as with a slit. In principle, however, the second dimension of spatial information can be used to remove these differences during processing. The main disadvantage of slicers at present is that they are difficult and expensive to manufacture—when this situation changes, they are likely to become the dominant technique for integral field spectroscopy.

1.3   Advantages of IFS

Integral field spectroscopy provides a variety of advantages over other methods for spatially resolved spectroscopy, solving most of the problems discussed in §1.1:

- As a form of imaging spectroscopy, IFS allows one to associate spectral features with spatial image structure and study complex kinematics or morphology.

- Both spatial dimensions and wavelength are resolved simultaneously. This is much more efficient for small to moderate field sizes than other forms of imaging spectroscopy. It also gives more self-consistent data. (Where larger fields are needed, it may sometimes be better to consider interferometry for imaging single emission lines or use a slit for detailed spectroscopy.)

- Prior assumptions as to where to take slit measurements are removed. IFS has already uncovered a number of examples where important features are not apparent from broad-band imaging (eg. Mediavilla et al., 1998; Bacon et al., 2001b; Davies et al., 2001).

- Slit losses and alignment errors are avoided, especially at high resolution with adaptive optics. As well as improving sensitivity and reliability, this reduces observing overheads: acquisition times are shorter and there is no need to adapt the set-up for different seeing conditions. As a result, spectral resolution is almost independent of the seeing.

- Image processing operations can be performed on IFS datacubes; important examples are correcting atmospheric dispersion and deconvolving an adaptive optics PSF from the data.

- Not only do IFS and AO work well together technically, but their capabilites are required for many of the same projects, such as studying the centres of galaxies or shock fronts. This new combination is therefore a powerful one, especially in the near-infrared, where AO is more effective and dust obscuration less significant than in the visible.

- IFS has more potential for serendipitous discovery than instruments which are targeted at precise image slices or wavelengths. It may therefore quickly help to formulate the next generation of questions in astronomy.

1.4   Astronomical applications

Integral field spectroscopy will have an important impact on many areas of astronomical research; the following are just a few of the potential applications:

Galaxy dynamics and composition: black hole detections with adaptive optics (eg. Bacon et al., 2000), correlations between nuclear black hole mass and bulge properties, kinematic detection of bars (eg. Emsellem et al., 2001), triaxiality of elliptical galaxies and properties of low versus high luminosity ellipticals (eg. Peletier et al., 2001), redshift evolution of the fundamental plane, interactions between galaxies.

Star formation in molecular clouds: kinematic structure of shock fronts and outflows using adaptive optics (eg. Tedds, 1999).

Active galaxies: interaction of AGN with their host galaxies, properties of the emission surrounding AGN (eg. Durret et al., 1994), emission line counterparts of radio jets and other features (eg. chapter 5), testing unified models.

Gravitational lensing: spectroscopic association of lens images (eg. Adam et al., 1989), reconstruction of lensed objects, redshift measurements using integrated light, modelling lenses (eg. Mediavilla et al., 1998).

Binary systems: spectra of components resolved with adaptive optics, coronography with graded throughput fibre IFUs (eg. Arribas et al., 1998*b*).

Local planets and moons: distribution of mineral composition (eg. Martin et al., 1995), atmospheric composition.

1.5   Instrumental projects

A large proportion of the spectrographs now being developed or commissioned for 8m telescopes include some type of IFS capability. Table 1.1 lists many of the IFUs which have been constructed to date or are currently in development. Several instruments which are in design study or prototyping stages have not been included, as final details have not been published.

Table 1.1: Some planned or completed integral field units.

Instrument	Lead group	Year	Design	λ	Field(s)	Sampling
2D-FIS	IAC	1994	Fibres	0.3–1μm	12×9″	1.2″
3D	MPE	1994?	Slicer	1.2–2.5μm	6×6″	0.4″
ATLANTIS	IAC	future	Fibres	1–2.4μm	1–8″	0.03–0.25″
CIRPASS	Cambridge	2001	Fibre-lens	0.85–1.8μm	2×0.6″–13×4″	0.05–0.36″
DensePak I	KPNO	≤1986	Fibres	0.4–1μm	22″	3.2″
DensePak II	KPNO	1986	Fibres	0.4–1μm	16×19″	2.7″
FIS	AAO	1981	Fibres	0.4–1μm	17×22″	1.3″
GIRAFFE	ESO	2001	Fibre-lens	0.4–1μm	15× 2×3″	0.5″
GMOS-IFU	Durham	2001	Fibre-lens	0.4–1μm	5×7″, 5×3.5″	0.2″
GNIRS-IFU	Durham	2002	Slicer	1–2.5+μm	3×4″	0.15″
HEXAFLEX	IAC	1989	Fibres	0.4–1μm?	13″/60″	1.5″/7.3″
HEXAFLEX II	IAC	1994?	Fibres	0.4–1μm?	8×10″	0.9″
INTEGRAL	IAC	1997	Fibres	0.4–1μm	8×6″–34×29″	0.6–3″
KTS I	Kyoto	1998?	Microlens	0.4–0.7μm	9×15″	1.3″
KTS II	Kyoto	future	Microlens	0.36–0.9μm	14″	0.39″
MOS/ARGUS	CFHT	1993	Fibres	0.4–1μm	13×8″	0.4″
MPFS	Russ. Ac. Sci.	1990	Microlens	0.4–1μm	~8–16″	0.5–1″
NIFS	ANU	2002?	Slicer	1–2.5+μm	3″	0.1″
OASIS	Obs. Lyon	1997	Microlens	0.4–1μm	2×1″–15×12″	0.04–0.4″
PIFS	Cal. Tech.	1998	Slicer	1–5μm	5×10″	0.67″
PMAS	Potsdam	2001	Fibre-lens	0.35–0.9μm	8″	0.5″
SAURON	Obs. Lyon	1999	Microlens	0.48–0.54μm	9×11″/33×41″	0.26/0.94″
SLIFID	Obs. Meudon	1986	Fibres	visible	7.5″/16″	0.33″/0.7″
SMIRFS-IFU	Durham	1997	Fibre-lens	1–2μm	6×4″	0.62″
SPIFFI	MPE	2002?	Slicer	1–2.5μm	0.8–8″	0.025–0.25″
SPIRAL A	Cambridge	1997	Fibre-lens	0.5–1μm	3.5×3″	0.5″
SPIRAL B	AAO	2000	Fibre-lens	0.5–1μm	22×11″	0.7″
TEIFU	Durham	1999	Fibre-lens	0.4–1μm	2× 7×4″	0.25″
TIGER	Obs. Lyon	1987	Microlens	0.35–1μm	7″/10″	0.4″/0.6″
UIST-IFU	ROE	2002?	Slicer	1–5μm	3×7″	0.24×0.48″
VIMOS-IFU	Marseille	2000	Fibre-lens	0.4–1μm	27″/54″	0.33″/0.67″

1.6   Thesis outline

This thesis begins with a project to test for a super-massive black hole in the centre of Fornax A, using the 3D infrared image slicer (Weitzel et al., 1996) along with optical slit spectroscopy. With the help of dedicated reduction software, provided by the instrument group at Max-Planck-Institut für extraterrestrische Physik (MPE), the stellar kinematics have been mapped as a function of isophotal radius over the central $6''$ of the galaxy.

Chapter 3 introduces the complex problem of reducing data from "densely packed fibre" IFUs, as pioneered by the Astronomical Instrumentation Group at Durham University. It is explained that fibre IFUs can have almost the same limiting efficiency as image slicers in the visible, given a design and data reduction paradigm which differs from established multi-object methods. This paradigm is discussed in writing for the first time and the issue of recovering three dimensional images from the raw data is explored conceptually in some depth.

In chapter 4, a prototype software package for reducing data from TEIFU and future densely-packed fibre IFUs is presented. This implements the techniques developed in chapter 3 and provides some basic facilities for analysis of IFS datacubes. The TEIFU data reduction framework is documented, with some discussion of generic practical issues. The source code of the software package can be found on the accompanying compact disc.

Chapter 5 presents an investigation into the origin of [Fe II] emission in NGC 4151, undertaken using Durham University's SMIRFS-IFU. The [Fe II] and Paβ lines have been mapped in intensity and velocity, allowing their relationships with the Seyfert galaxy's extended narrow line region and radio jet to be studied. The distribution of [Fe II] with respect to these features provides a strong indication as to its excitation mechanism. Some discussion of data reduction methods and instrumental performance is included.

A report on TEIFU commissioning is given in chapter 6. The fibre response pattern of the IFU is characterized and measurements of absolute system throughput are given, based on a standard star observation. The IFU's efficiency is also compared with other fibre feeds for the same spectrograph. A reconstructed observation of the radio galaxy 3C 327 demonstrates the successful implementation of TEIFU's

innovative hardware and software designs.

Finally, an overview and summary of conclusions is given in chapter 7.

Chapter 2

Testing for a super-massive black hole in NGC 1316

Using the 3D infrared image slicer on the Anglo-Australian Telescope, together with optical slit and aperture spectra, we have measured the stellar velocity dispersion at the centre of NGC 1316 as a function of radius. Our measurements fall within the range of previous values from the literature, showing a velocity dispersion of $\sim 220\,\mathrm{km\,s^{-1}}$, with no significant gradient over the central few arcseconds. It is concluded that NGC 1316 does not contain a central black hole of mass $\sim 5 \times 10^9 M_\odot$, predicted by the relationship of Kormendy & Richstone (1995). The discrepancy is most likely due to the presence of a young stellar population, recently dated by Goudfrooij et al. (2001) at 3 Gyr. An approximate upper limit of $8 \times 10^8 M_\odot$ on the mass of a central black hole is, however, compatible with the predictions of Gebhardt et al. (2000) and Ferrarese & Merritt (2000).

2.1   Introduction

The demographics of quasars as a function of redshift suggest that now normal galaxies were active in the past; thus it has been argued that nearby galactic nuclei should contain remnant massive black holes (Haehnelt & Rees, 1993). Kormendy & Richstone (1995) outline the first few detections of massive black hole candidates by dynamical means, using the results to demonstrate a correlation between black hole mass and host bulge luminosity. Faber et al. (1997) imaged nearby galactic nuclei with the HST, finding that the steep power-law brightness profiles of more luminous

bulges ($M_V < -22$) give way to shallow *cuspy* cores at the centre, in contrast to the global power-law profiles of their fainter ($M_V > -20.5$) counterparts (both cases are found between these magnitude limits). They argue for a model in which cores are formed by black holes during mergers, spiralling towards the centre of mass and disrupting the power law by heating and ejecting stars. This process reproduces the observed cores for a mass ratio, $M_\bullet/M_{bulge}$, of 0.2% (Quinlan & Hernquist, 1997), consistent both with the Kormendy & Richstone sample and with an argument based on the energy density of quasar light.

NGC 1316, which hosts the radio source Fornax A, is the most luminous early type galaxy in the nearby Fornax cluster. Based on supernova light curves, Goudfrooij et al. (2001) deduce a distance of 22.9 Mpc, giving $M_B = -22.4$. From the data in Kormendy & Richstone, this suggests a nuclear black hole mass as high as $\sim 5 \times 10^9 M_\odot$. The radius inside which the gravitational influence of such an object would dominate the bulge kinematics is $r = GM_\bullet/\sigma_{bulge}^2$, or about $4''$ at 22.9 Mpc; under normal seeing conditions, one would therefore expect to observe a central rise in the dispersion.

Recently, two groups have independently discovered that black hole mass is more tightly correlated with bulge velocity dispersion than luminosity (Gebhardt et al., 2000; Ferrarese & Merritt, 2000). Such a strong connection is somewhat unexpected, since the radius over which kinematics are dominated by a black hole is generally smaller than the scale of the bulge dispersion measurements; moreover, the relationship is not $M_\bullet \propto \sigma^2$ due to gravitational interaction, but (roughly) $M_\bullet \propto \sigma^5$. This seems to indicate that the formation of galaxies and central black holes are closely linked. Adopting an approximate dispersion of $235 \mathrm{km\,s^{-1}}$ for the central region of NGC 1316 gives $M_\bullet \sim 3 \times 10^8 M_\odot$ ($r \sim 0.2''$), an order of magnitude smaller than the luminosity-based estimate.

Located at the cluster periphery, NGC 1316 is of lenticular type. Its extensive dust lanes (faintly visible in fig. 2.1) are patchy and misaligned with the projected axes. Schweizer (1981) reported the first detailed ground based photometry and spectroscopy, showing an unusually sharp, high surface brightness core, with ripples and shells in the outer regions. These characteristics suggest that the galaxy is a recent merger product which may now be undergoing core formation.

Surveying the Fornax cluster, Kuntschner & Davies (1998) found that the object

has the strongest Hβ absorption of any of the metal rich galaxies, suggesting that star formation in the merger has left a modest fraction of young, low M/L stars superimposed on the progenitor population (cf. de Jong & Davies, 1997). Moreover, Goudfrooij et al. (2001) have observed a number of unusually bright star clusters, for which they have derived ages of 3.0 ± 0.5 Gyr by comparing Hα and Ca II triplet absorption measurements with the models of Bruzual & Charlot (eg. Bruzual & Charlot, 1993). A young, bright stellar population may therefore account for the discrepancy between black hole masses inferred from luminosity and velocity dispersion, as well as for the deviation of NGC 1316 from the Faber-Jackson and fundamental plane relations (eg. Kuntschner, 2000; D'Onofrio et al., 1997). That being the case, the value of $M_\bullet$ would be overestimated by the relationship of Kormendy & Richstone.

Measurements of the galaxy's stellar kinematics have previously appeared in a number of papers, but with considerable systematic offsets between different authors (see §2.4). The published values are largely based on optical measurements using Fourier Quotient techniques, leaving open the possibility that the velocity distribution is affected by dust along some lines of sight. Here, we present integral field spectroscopy of stellar CO absorption at 2.3 microns, avoiding dust obscuration, slit losses, misalignments and placement assumptions. In addition, tip-tilt seeing correction provides sub-arcsecond resolution.

Using V and I band images from the HST Planetary Camera, Shaya et al. (1996) modelled NGC 1316 with a constant stellar mass-to-light ratio, $(M/L)_R = 2.2 \pm 0.2$, and a dark mass at the centre. They found the available kinematic measurements to be consistent with a central mass of $2 \times 10^9 M_\odot$, but recognize that the results could also be explained by an increase in radial anisotropy. More sophisticated modelling techniques now incorporate anisotropy as well as rotational flattening, and can accept higher-order measurements of the velocity distribution as constraints (Cretton et al., 1999). With samples of the velocity distribution over two spatial dimensions, it is possible to break the degeneracy between changes in orbital structure and a central increase in M/L (van der Marel et al., 1998). Our infrared data are therefore complemented by extensive long-slit mapping at optical wavelengths, allowing us to model the kinematics on a larger scale and account for dynamical anisotropy.

This chapter focuses on whether there is direct evidence for a super-massive black

hole in NGC 1316, based on stellar velocity measurements from our integral-field and major-axis optical spectra of the centre. The results are compared with data from previous publications. A future paper, in preparation, will include the work of co-authors in dynamical modelling and line-strength analysis, and will consider the evolutionary history of NGC 1316 in further detail.

2.2 Observations and data reduction

2.2.1 Optical spectroscopy

A central spectrum of NGC 1316 (2.3 × 3.85″, PA = 90°) was obtained with the 3.9m AAT as part of a survey of the Fornax cluster (Kuntschner & Davies, 1998; Kuntschner, 2000); for basic reduction, the reader is referred to Kuntschner (2000). Spatially resolved long-slit spectra were taken with the MSSSO 2.3m telescope, at PAs of 5, 15, 50, 60, 95, 105, 140 and 150° (see fig 2.1); observational details are provided in table 2.1. Tasks such as bias subtraction and flat-fielding were performed using IRAF. Cosmic rays were removed using the program *cleanest*, in N. Cardiel's RED$_{m}^{uc}$E software[1]. Wavelength solutions were determined, with a typical RMS $\lesssim 0.05$Å, from Ne-Ar lamp spectra bracketing most of the science observations. Finally, the sky was subtracted. The spectra were rebinned spatially for analysis, with a constant signal-to-noise ratio of 35 per Ångstrom.

Figure 2.1: NGC 1316, marked with the MSSSO slit positions and the region observed using 3D (box).

---

[1]available at http://www.ucm.es/info/Astrof/reduceme/reduceme.html.

Table 2.1: The instrumental set-up

Telescope	AAT (3.9m)	MSSSO (2.3m)	AAT (3.9m)
Dates	6-8 December 1996	28-30 Oct 1997	12-14,17-18 March 1998
Instrument	RGO spectrograph	DBS spectrograph	MPE 3D image slicer
Spectral range	4243 - 5828 Å	4335 - 5487 Å	2.16 - 2.41μm
Grating	600 V	1200 blue	K_{long} grism
Dispersion	1.55 Åpixel^{-1}	0.555 Åpixel^{-1}	5 Åpixel^{-1}
Resolution	$\sim$ 4.1 Å	$\sim$ 2.3 Å	$\sim$ 13 Å
Spatial Scale	0.77″pixel^{-1}	0.91″pixel^{-1}	0.4″pixel^{-1}
Slit Width	2.3″	2.5″	0.4″
Detector	Tek1k #2	Site1.7x0.5k #7	Rockwell NICMOS-3 256^2
	(24μm^2 pixels)	(15μm^2 pixels)	(40μm^2 pixels)
Gain	1.36 e^-ADU^{-1}	1.0 e^-ADU^{-1}	
Read-out-noise	3.6 e^- rms	5 e^- rms	25 e^- rms
Seeing	$\sim$1″	$\sim$1–2″	$\sim$0.8″

Table 2.2: Log of observations: MSSSO template stars

Name	Type	Comment
Night1		
HD001461	G0V	Lick/IDS std
HD003567	F5V	Lick/IDS std
HD004656	K5III	Lick/IDS std
HD037160	K0III	Lick/IDS std
HD217877	F8V	Lick/IDS std
HD219449	K0III	Lick/IDS std
HD219617	F6VI	Lick/IDS std
HD221148	K3III	Lick/IDS std
Night2		
HD014802	G1V	Lick/IDS std
HD022879	F9V	Lick/IDS std
HD217877	F8V	Lick/IDS std
HD219449	K0III	Lick/IDS std
HD219617	F6VI	Lick/IDS std
HD221148	K3III	Lick/IDS std
HR8634	B8V	spec. std

2.2.2   3D observations

MPE's 3D cryogenic image slicer (Weitzel et al., 1996), hosted on the AAT, was used to observe the central 6×6" of NGC1316 in the K band. These measurements complement optical data at larger radii, providing a spectroscopic map of the centre with reduced sensitivity to dust obscuration. Stellar kinematics can be studied via the first overtone lines of CO absorption, around $2.3\mu m$. The 3D optics incorporate a tip-tilt mirror, which helps maintain an image quality close to the sampling limit of 0.8" FWHM. Further instrumental details are given in table 2.1.

The useable data consist of 31 cycles of two object and two sky frames, each integrated for 100s; the total time on source was therefore 103 minutes. In total, 7 runs of 7 cycles were attempted, with a number of frames lost due to cloud. Each run was accompanied by an observation of a nearby B dwarf star (see table 2.3); this has no prominent features at $2.3\mu m$, so serves as a reference for correcting telluric absorption lines at the appropriate zenith distance, several of which coincide with galaxy features. K and M giants were also observed, again with suitable atmospheric references, as templates from which to measure the galaxy's velocity distribution.

Basic reduction was performed using the 3D package provided by MPE for the GIPSY data processing system (van der Hulst et al., 1992). The calibration procedure is discussed in Weitzel et al. (1996). Briefly, the stages are as follows for each cycle: frames are corrected for non-linear photometric response; sky is subtracted pixel-for-pixel from the object; the result is divided by a dome flat; bad pixels are re-estimated; exposures offset by half a pixel in wavelength are interleaved to achieve critical sampling; the spectrogram is interpolated onto a regular grid in wavelength, using an arc lamp calibration for reference; image slices are resampled spatially and stacked in an $(\alpha, \delta, \lambda)$ cube.

The datacubes produced for individual cycles are resampled such that their centroids are spatially aligned, and added together for each run. The results are divided through by normalized atmospheric spectra, weighted with their relative integration times and combined into a single cube for analysis. A wavelength-integrated image from the final datacube is shown in fig. 2.2.

Table 2.3: Calibration stars for 3D data

Star	Type	Comments
HR1190	B9 V	atmospheric calibrator for galaxy
HR4961	K5 III	velocity template
HR4299	M0 III	velocity template
HD97585	A0 V	atmospheric calibrator
HD129135	G0 V	atmospheric calibrator

Figure 2.2: Wavelength-integrated image of NGC 1316 from the final 3D datacube (covering the region marked on fig. 2.1).

2.3   Kinematics

Velocity profiles were determined using version 8 of the Fourier correlation quotient (FCQ) method (Bender, 1990). For the aperture and slit spectra, we derived measurements of recession velocity, dispersion and the first order deviations from a Gaussian distribution, as in Bender et al. (1994). For the 3D data, a simple Gaussian profile was fitted to the broadening function since we are mainly interested in the magnitudes of v and σ and the infrared observations are noisier.

By visual inspection of the residuals from the FCQ fit, we determined the best template star for the MSSSO data to be HD221148 (K3III); other templates give

σ values which agree within 6 km s^{-1}. For the AAT spectrum, we averaged the results of 8 different template stars (G,K giants). For the 3D observations, HR4299 (M0III) fits better than HR4961 (K5III), probably due to more accurate atmospheric correction rather than spectral type.

The optical wavelength regions used in the analysis were 4856.1–5486.2 Å for MSSSO data and 4875.6–5597.1 Å for the AAT spectrum. These ranges narrowly exclude $H\beta$, which is too strong to be fitted and is also biased towards the young stellar population. In the infrared, a section from 2.201–2.373 μm was used, covering the first three ^{12}CO bandheads, at 2.294μm, 2.323μm and 2.352μm.

Parameters related to continuum removal and Fourier filtering were optimized using Monte-Carlo simulations—the stellar template was used to generate simulated spectra corresponding to particular values of σ, h_3, h_4 and S/N. For each combination, 30 different noise realizations were created. This experiment demonstrates that we can recover the velocity dispersion within random errors of $\sim$5 km s^{-1}. The 3D values are, however, systematically overestimated by $\sim$10 km s^{-1}.

Because the 3D observations are relatively noisy and σ turns out not to vary greatly over the field of view, the data are not best exploited through 2D mapping of the galaxy centre. Instead, we have fitted elliptical isophotes to the wavelength-integrated image, and combined the spectra in four isophotal bins, with mean semi-major axes of 0.5″, 1.3″, 2.1″ and 2.9″. Hence we measure σ as a function of distance in arcseconds along the major axis, providing a long-wavelength counterpart to the optical measurements. In doing this, the drop in intensity from one ellipse to the next is compensated well by the increased number of binned spectra, such that the signal-to-noise ratio of the final result falls gradually from the centre. The rotation velocity has relatively little impact on the dispersion measurements, reaching 10–20 km s^{-1} along the major axis at the corners of the field. We see no evidence in our slit data for an offset between the optical and kinematic nuclei, as suggested by D'Onofrio et al. (1995). For the final analysis, we integrated data in four bands, producing the spectra in fig. 2.3.

Figure 2.3: Spectra from the 3D datacube, binned azimuthally for analysis, and the corresponding velocity template. The galaxy spectra have been de-redshifted.

2.4   Results and comparison with other data

Our optical and infrared observations are in good agreement, particularly after the 3D dispersion values are adjusted for the systematic errors found by Monte-Carlo simulation. Within the limits of error, the small velocity gradient across the 3D field of view is not detectable in the unbinned datacube. Table 2.4 lists the velocity dispersions derived from the combined infrared spectra, σ_{meas}, along with the values after adjustment for systematic errors, σ_{adj}. The second column shows the average radii of the constituent spectra when projected onto the major axis. This does not account for the 0.8″ seeing, which would effectively move the central point to $\langle r \rangle \simeq$ 0.63″, shifting the others by small amounts. The central dispersion values obtained at different slit angles in the optical (1–2″ seeing) range from 211 to 226 km s^{-1}, with a mean of 217 km s^{-1}. This level of self-consistency confirms that the observations have been taken accurately.

The 3D dispersion values show a small upturn towards the centre, which could be real, but is not clear evidence of an $r^{-1/2}$ increase. The major-axis optical measurements are almost flat as a function of radius, but with slightly more scatter. Fig. 2.4 shows major-axis dispersion values from all three sets of data, alongside measurements from other literature. The radii at which the points are plotted are those reported by the authors; most of them are clearly not adjusted for finite pixel sizes or seeing, but either way there is no significant gradient.

An obvious feature of the plot is that the data sets fall in two systematically offset groups. Our three sets of measurements sit with those of Arnaboldi et al. (1998), around 215km s^{-1}, whilst the values of D'Onofrio et al. (1995), Longhetti et al. (1998) and Gorgas et al. (1990) are $\sim$ 40km s^{-1} higher. Early central estimates (Jenkins & Scheuer, 1980; Schweizer, 1981) agree with the latter groups, but used

Table 2.4: Velocity dispersions from the binned 3D spectra.

$r_{min}-r_{max}$	$\langle r \rangle$	σ_{meas} (km s^{-1})	σ_{adj} (km s^{-1})
0.0–0.8″	0.52″	237 ± 3	227
0.8–1.6″	1.25″	231 ± 4	220
1.6–2.4″	2.07″	226 ± 5	214
2.4–3.2″	2.86″	227 ± 6	215

Figure 2.4: Velocity dispersion as a function of distance along the major axis—results and comparison with previously published values. The 3D values have been adjusted for systematic errors of $\sim$10 km s^{-1}, estimated from Monte-Carlo simulations.

relatively crude analysis techniques. We have gathered a number of self-consistent observations in two wavebands with three instruments, using modern techniques and a number of templates; hence we are confident in our absolute calibration. Our estimates fall comfortably within the range of the previous data.

2.5   Conclusions

Whilst we have accurately determined the velocity dispersion in the central few arcseconds of NGC 1316, there is no evidence for a BH of mass $\sim 5 \times 10^9 M_\odot$, as predicted by the mass–luminosity relationship for early type bulges. This disagreement is most likely due to the presence of a 3 Gyr old stellar population, formed during a major merger, which raises the total luminosity with respect to the bulge mass, Fundamental Plane and correlation found by Kormendy & Richstone.

A rough upper limit on the mass of any central black hole can be established using $M_\bullet = \sigma^2 r/G$. Adopting an RMS radius of 0.63$''$ for the central IFU measurement

(including binning & seeing) gives $M_\bullet < 8 \times 10^8 M_\odot$. The recently discovered $M_\bullet/\sigma$ relation indicates that we are probing just outside the radius of influence of a $3 \times 10^8 M_\odot$ black hole, consistent with the hint of an upturn seen in our 3D data. We would therefore expect to be able to make a positive black hole detection using a spectrograph with adaptive optics, such as SINFONI (Mengel et al., 2000).

Although we have not presented two-dimensional kinematic mapping, integral field spectroscopy has ensured that we avoided slit alignment problems and has allowed azimuthal binning of the whole field for an improved signal-to-noise ratio. Since we are interested in the object centre, this is more efficient than using a slit, even if the more complex IFU optics result in a modest drop in throughput.

Chapter 3

Interpreting densely packed fibre spectra

The Astronomical Instrumentation Group at Durham University has pioneered the use of optical fibre IFUs which produce closely spaced, overlapping spectra at the detector. By minimizing the fibre separation so that it is similar to the pixel spacing, the number of elements which can be sampled using the available detector area is maximized. The complex problem of recovering three dimensional images from the resulting data is discussed here in detail for the first time. With reference to sampling theory and some basic simulations, it is shown that observations with one detector row per fibre may be interpreted in a similar way to those of a long-slit spectrograph or image slicer. One side effect of this approach is that spatial sampling becomes irregular. Methods for extracting and reconstructing sample sets from the raw data are discussed, with a brief review of available literature and algorithms.

3.1 Introduction

With instrument development focusing on the new generation of eight metre class telescopes, efficiency optimization is more important than ever before. The overall performance of a system depends on many factors—optical throughput, detector sensitivity, the number of data per exposure, the fraction of these which are useful, the fraction of time spent observing targets, techniques for data reduction and, of course, the scientific impact of results. In spectroscopy, the parameter with most room for improvement is usually *multiplex gain*. This measures the gain in useful data per exposure due to having multiple elements of some kind—one-dimensional image slices for an IFU or object apertures for a multi-object instrument.

Despite making notable contributions in astronomy, early IFUs have been limited in application by their small fields of view. Those with larger fields are mainly specialized instruments with short spectral ranges (eg. Bacon et al., 2001a). Any increase in multiplex gain therefore corresponds to a genuine improvement in efficiency, and not just a higher figure of merit. An IFU is intended to provide a more complete picture of its target than a one-dimensional slit, but this objective can only be met with an adequately sized field.

As construction techniques improve, optical fibre arrays can be made larger, achieving greater multiplex gains with respect to single slit designs. The stage has recently been reached where the number of fibres in a bundle can match the linear dimension of a large CCD array (eg. Le Fevre et al., 1998). It is therefore important to avoid redundancy in sampling fibre spectra, maximizing the information gathered by a detector of fixed size. Subject to the constraints of a particular optical system, it is possible for a fibre IFU to utilize a detector almost as efficiently as a camera or image slicer (Haynes et al., 1998b). Such a system departs significantly from previous designs and needs a reconsidered approach towards data reduction—formulated here and in chapter 4.

3.2   Integral-field versus multi-object data

In the past, fibre IFUs have often been viewed as an extension of multi-object fibre systems, inheriting the design philosophies associated with them. In some respects, this point of view is misleading, because the similarity between the two instrument types is more technical than functional. Spectra from an integral field have less individual significance than those from separate objects. The required role of an IFU is to separate telescope images into slices in one direction, allowing three-dimensional information to be encoded on a two-dimensional detector. However, optical fibres (and simple lenslet arrays) divide the field in both spatial dimensions, producing point spectra instead of slit spectra. This is an added complication, because the role of sampling in the second spatial dimension is then duplicated by the detector.

Once a field of view has been divided into discrete points, there is a tendency to regard those points as elemental—it is convenient to think of the IFU as a pixel array. Treating fibres as separate elements allows IFU design and data reduction to follow

the multi-object paradigm. This has two key advantages—conceptual simplicity and the ability to re-use existing software. However, by treating the IFU more like an image slicer, with less distinction between fibres at the detector, it is possible to increase the multiplex gain by a substantial factor. Recognizing this, Content and colleagues at Durham University have pioneered an alternative paradigm for densely packed fibre systems.

In multi-object spectroscopy, it is important to keep fibres well separated at the detector, avoiding cross-contamination between the spectra of unrelated objects. In contrast, imaging always involves crosstalk between spatial elements, due mainly to atmospheric seeing or the telescope diffraction pattern. Where adjacent spectra from an IFU merge together at the same wavelengths, the effect is just an extra convolution term in one spatial dimension. For moderately large overlaps, this represents a second order degradation of the seeing resolution (Allington-Smith & Content, 1998). Hence it is possible to place spectra close together, recovering a continuous, almost slit-like pattern at the output (fig. 3.1).

3.3   Design considerations

The spot profile created by an imaging system in response to a point source of light, or plane wave, is known as its point spread function, or PSF. This determines the extent to which adjacent image features blend together, and hence the level of detail which can be resolved. Since the size and resolution of any analogue image are both limited, the information content is finite and can, in principle, be recorded digitally without loss. This is essentially a qualitative statement of the *Sampling Theorem*.

Image resolution is usually quantified in terms of the full width at half maximum (FWHM) of the PSF. In astronomy, a commonly cited result of the Sampling Theorem is that the FWHM must cover twice the distance between measurement points, for the information in images to be sampled critically (§3.5). Undersampling, with a lower pixel density, creates artifacts in the data, whilst oversampling is inefficient because it provides no extra information.

For a detector array to be used optimally, it must be filled by the incident image and the magnification must be such that the PSF is sampled critically—a scenario which occurs most often in direct imaging. With fibre spectra, this means

Figure 3.1: Constant wavelength image cross-sections at the detector of a multi-object fibre spectrograph (top) and an integral field spectrograph with densely packed fibres (bottom). Uniform illumination and transmissions are assumed.

avoiding dead space between fibres and minimizing the width of each spectrum so that the Sampling Theorem is just satisfied. In practice, oversampling is sometimes unavoidable: where a spectrograph is used for both fibre and slit spectroscopy, the image scale has to accommodate different requirements, and even with a dedicated fibre spectrograph, the pixel scale is limited by the maximum sizes of detector pixels and optical elements. It is nevertheless important to limit redundancy as much as possible, leading to designs which are quite different from multi-object fibre systems. In the limit, a fibre IFU can approach the detector utilization of a camera or image slicer, albeit with more complex requirements for data reduction.

For faithful reproduction of an IFU's observation space, each of the three dimensions must be sampled (at least) critically. In x and y, resolution is set primarily by

the atmospheric seeing or telescope PSF, which should cover two fibres in FWHM at the IFU input. There is also some blurring in x due to crosstalk between fibres in the spectrograph (§3.6). In λ, resolution is determined (for a given dispersion scale) by the spot profile from each fibre at the detector, which should have a two pixel FWHM in the wavelength direction. There is no strict requirement, however, for individual fibres to be critically sampled in the *spatial* direction; adjacent spectra are already correlated if there are two fibres per FWHM at the IFU input. It turns out that the Sampling Theorem can effectively be satisfied in x, y and λ even where the two-dimensional spectrogram is itself undersampled spatially (§3.7).

With a fibre separation of one pixel at the output, there are just enough pixels to preserve critical sampling of the seeing by fibres—the IFU is optimally matched to the detector. If images are oversampled at the IFU input, the fibre output spacing can be even less than one pixel. However, critical sampling in the spectral direction requires a two pixel FWHM; for axially-symmetric fibre beams, this would lead to overlaps on the scale of the seeing, degrading the resolution in x unacceptably. Thus for a given detector to observe the maximum possible critically sampled volume, fibres must project elliptical image profiles, requiring anamorphic magnification in the spectrograph. A very similar problem arises with image slicers, again because the widths of spatial elements define the spectral resolution.

Anamorphic optics are not always straightforward or practical to implement, and are not used by any existing fibre IFU. However, a few alternative strategies can still make good use of the detector. First, one can undersample spectrally but shift the detector by half a pixel between subsequent integrations, providing the missing intermediate samples whilst doubling the spectrum length. This is the technique adopted by MPE for their 3D image slicer (Weitzel et al., 1996). Second, for some applications it may be acceptable to undersample the instrumental profile in λ— such as when observing spectral lines with intrinsic widths of two pixels or more. Durham University's TEIFU, which motivated the present study, uses this approach (Haynes et al., 1998b). Finally, it may be cost effective just to use a fibre spacing of two pixels at the output, which is still reasonably efficient. Where the overall optical design of a system does not permit a FWHM of less than two pixels, the issues of spectral undersampling and anamorphism do not arise.

At this stage, it is clear that densely packed fibre IFUs pose quite a different data reduction problem from previous fibre systems. With fewer than two detector rows per fibre, or even fewer than one, it is necessary to deal with data pixel-by-pixel, rather than fibre-by-fibre. In some sense, a fibre output pitch of one pixel marks the centrepoint of a transition between discrete fibre designs and image slicers—a slicer being equivalent to an instrument with many fibres per pixel. Because dense packing results in similar spatial sampling rates at the IFU and detector, the data sets produced are more complex than those of instruments whose spatial elements are defined only by fibres or only by pixels.

Having dealt previously with multi-object spectra, it is tempting to follow a similar approach for integral field spectroscopy, keeping data reduction as simple and compatible as possible. However, a typical multi-object design has a much lower information density than is possible for an IFU, not only because spectra cannot overlap, but also because it is difficult to configure a large number of separate fibres at the focal plane. Hence whole object spectra are usually well resolved in both directions, allowing profile-weighted extraction, which compensates to a small extent for the reduced number of independent data. For an IFU, such an arrangement is far from ideal if it limits the size of the fibre bundle. For example, critically sampled Gaussian profiles have to be separated by $\sim$5 pixels in order to keep crosstalk below 1%, reducing the maximum field of view by 80%. Moreover, imaging each spectrum with more pixels than are needed increases the level of read noise, which can be critical in certain applications.

Merely allowing spectra to overlap does not have a great impact on data reduction methods (§3.8). If the cross-sectional FWHM is greater than two pixels, it is still possible to work with separate one-dimensional spectra, although the details of extraction and three-dimensional image reconstruction are slightly different from the case of separate spectra. The situation becomes more complicated with one detector row per spectrum, but this added complexity is relatively unimportant once suitable reduction techniques have been developed. The following sections explore the properties of densely packed fibre data in some detail, whilst chapter 4 presents a software implementation of the reduction methods discussed here.

3.4   Sampling theory

The Sampling Theorem gained widespread recognition in the context of electronic communications, owing largely to the writings of C. E. Shannon. It is often named after one or more of Shannon, Nyquist and Hartley of Bell Labs (Nyquist, 1928; Shannon, 1949), Whittaker of Edinburgh University (Whittaker, 1915) and Kotel'nikov of the Russian Academy of Sciences (Kotel'nikov, 1933). The theorem deals with sampling on a regular grid, such as a CCD array. A simple result allows exact reconstruction of functions from critically spaced measurements. Where a function is sampled irregularly, reconstruction is generally more difficult—research into efficient techniques is ongoing (eg. Strohmer, 1993; Marks, 1993). Spectroscopy with densely-packed fibres is peculiar in that it can produce irregular measurement points via two regular sampling stages. The classical Sampling Theorem provides a good starting point for further investigation.

For simplicity, consider a one-dimensional function, $f(x)$. Sampling at a particular point, x', can be represented via the sifting property of the Dirac delta (impulse) function:

$$f(x') = \int_{-\infty}^{\infty} \delta(x - x') \, f(x) \, dx. \tag{3.1}$$

Physically, $\delta(x)$ can be any function of unit integral which is zero except over a negligible width at the origin; thus, multiplication with $f(x)$ produces an impulse whose area, or strength is the value at that point:

$$f(x) \, \delta(x - x') = f(x') \, \delta(x - x'). \tag{3.2}$$

The process of sampling at regular intervals, a, is equivalent to multiplication by a *comb* or *shah* function, defined as

$$\text{III}\left(\frac{x}{a}\right) = |a| \sum_{n=-\infty}^{\infty} \delta(x - an). \tag{3.3}$$

The multiplicative $|a|$ is just for consistency, so that the unit area of δ is preserved whilst the co-ordinate is rescaled. Clearly, $\text{III}\,f$ is a series of impulses, whose areas correspond to the sampled values—it contains the same information from $f(x)$ as would point-like measurements:

$$\text{III}\left(\frac{x}{a}\right) f(x) = |a| \sum_{n=-\infty}^{\infty} f(an) \, \delta(x - an). \tag{3.4}$$

Fig. 3.2 represents this graphically, where the strength of each impulse is shown by its height, in the usual way.

Figure 3.2: Sampling a continuous waveform with a shah function.

The shah function is a useful tool because its Fourier transform is also a shah function. Working with frequencies in reciprocal units of x, rather than $2\pi/x$,

$$\mathrm{III}\left(\frac{x}{a}\right) \overset{\mathcal{F}}{\rightleftharpoons} a\,\mathrm{III}(au). \tag{3.5}$$

According to the Convolution Theorem, multiplication of functions in the spatial domain is equivalent to convolution of their transforms (and vice-versa). Thus the spectral representation of sampled data is

$$\mathcal{F}\left(\mathrm{III}\left(\frac{x}{a}\right)f(x)\right) = a\,\mathrm{III}(au) * F(u). \tag{3.6}$$

This is the frequency spectrum of the original, continuous $f(x)$, replicated infinitely at intervals of $1/a$ (fig. 3.3); the delta function has a constant spectrum over all frequencies, so convolution with each $\delta(u - u')$ simply shifts the origin to u'.

A function is said to be band-limited if its Fourier transform is zero above some cutoff frequency, u_c. In fig. 3.3, the replicas of $F(u)$ do not overlap if $u_c \leq 1/2a$. In that case, the original frequency spectrum is separable from the transform of the sampling pattern, and $f(x)$ can be recovered without error; otherwise, the spectrum is not determined uniquely. The Sampling Theorem can therefore be stated as follows: *a band limited function can be reconstructed completely from a set of regular*

Figure 3.3: Frequency spectrum of a uniformly sampled function.

samples, providing that the shortest component wavelength covers at least twice the measurement interval. For a given sampling, the maximum recordable frequency of $1/2a$ is known as the *Nyquist frequency.* Conversely, the critical sampling rate for a particular bandwidth is often called the *Nyquist rate.* (Sometimes, however, the former term is used for the sampling rate or its critical value.)

Mathematically, $F(u)$ can be isolated by multiplication with a box function, of unit value between $\pm u_c$ and zero elsewhere. This is denoted $\Pi(u/2u_c)$, and has the Fourier transform $2u_c \operatorname{sinc}(2u_c x)$. Therefore, by the Convolution Theorem, $f(x)$ is reconstructed in the spatial domain by a summation of sinc functions:

$$f(x) \;=\; \left(\mathrm{III}\left(\frac{x}{a}\right) f(x) \right) * \frac{\sin(2\pi u_c x)}{\pi x} \tag{3.7}$$

$$\;=\; |a| \sum_{n=-\infty}^{\infty} f(an) \frac{\sin(2\pi u_c(x - an))}{\pi(x - an)}. \tag{3.8}$$

Whilst equation 3.8 specifies $f(x)$ at all intermediate points, this is only possible because the harmonic content of the function is limited; one can reconstruct an image on an arbitrarily fine scale without approximation, but the inherent optical resolution must be coarser than the pixel size to begin with.

If $f(x)$ is undersampled, with $u_c > 1/2a$, then $F(u)$ does overlap its satellite copies in the transform of data (fig. 3.4). High frequencies become indistinguishable from lower ones (fig. 3.5). Those greater in magnitude than half the sampling rate are wrapped around to the opposite end of the original spectrum, creating spurious power at short spatial scales. The resolution at which data can faithfully be reconstructed is lower than if the analogue image had been suitably band limited before sampling. Extra high-resolution information is not only unmeasurable, but

contaminates otherwise good data. In fig. 3.4, undersampling is so severe that even the lowest frequencies are not known accurately. The phenomenon of high frequency variations being measured at a lower rate is known as *aliasing*. Another common source of artifacts is reconstruction with a bandpass filter which is too wide—an effect sometimes erroneously referred to as aliasing, but not inherent in the data.

$$III(au) * F(u)$$

Figure 3.4: Aliasing in the spectrum of an undersampled function. The light curves represent the same spectrum as in fig. 3.3, whilst the dark curve shows the total power in the principal interval.

$$f_n(x)$$

Figure 3.5: An undersampled component and its lower frequency alias.

Where oversampling, one can choose to reconstruct by passing a maximum frequency anywhere between u_c and $1/2a$. The latter case is illustrative—each sinc function in equation 3.8 has successive zero points at the locations of all but the central sample. Thus the value at each measurement position is only taken from that single point, whilst at intermediate positions all of the measurements contribute.

3.5 Application to imaging

Section 3.3 introduced the usual criterion for sampling images—namely that two pixels are required across the FWHM of a point source profile. Without diverging into details of optical theory, this rule needs relating to the Sampling Theorem.

The far-field diffraction pattern of an aperture is described by the Fourier transform of its distribution. Since we are dealing with incoherent illumination, and hence intensity rather than amplitude, it is the square of the transform which is of interest. For a rectangular aperture, this is a product of sinc functions, whilst for a circular mirror it is the closely related Airy pattern,

$$I(\theta) \propto \left[\frac{J_1(\pi D\theta/\lambda)}{\pi D\theta/\lambda} \right]^2, \tag{3.9}$$

shown in fig. 3.6. Here D is the telescope diameter, λ the wavelength of incident light, and θ an angular displacement in radians from the direction of the source. $J_1(x)$ is the first order Bessel function of the first kind. A derivation can be found in Goodman (1968) or Hecht (1987).

Figure 3.6: Sketch of the diffraction pattern of a circular aperture.

The power spectrum of an instrument's PSF is known as its modulation transfer function (MTF). From the Convolution Theorem, this describes the sensitivity of the

instrument to different spatial frequencies. Since the diffraction PSF of a telescope is the squared Fourier transform of its aperture, the MTF is the self-convolution of the aperture shape. It follows that for a finite aperture, images must be band-limited.

It turns out that the limiting frequency for a circular mirror of diameter D is D/λ cycles per radian (Goodman, 1968, p.120). The FWHM of the Airy pattern is $1.02\lambda/D$ radians, as marked in fig. 3.6. In order to satisfy the Sampling Theorem, pixel spacing must therefore be no more than $\lambda/2D$ radians, or 0.49 times the FWHM, approximately confirming the two pixel criterion.

The above argument is applicable to diffraction-limited imaging, but the PSF of an astronomical image is typically dominated by atmospheric seeing rather than the telescope. For a typical length exposure, the seeing disc is the sum of patterns which change continuously with turbulence. Aberrations within the telescope and spectrograph also contribute to smooth out the Airy profile. According to the Central Limit Theorem (eg. Bracewell, 2000, chapter 8), the convolution of a number of PSF terms tends towards a Gaussian curve. Whilst convolution cannot create high frequency components, so the telescope image must remain band-limited, it is clear that the two pixel criterion cannot strictly represent the Nyquist scale for a seeing-dominated profile.

The Gaussian distribution and its transform can be written

$$e^{-\pi(ax)^2} \overset{\mathcal{F}}{\rightleftharpoons} \frac{1}{a} e^{-\pi(u/a)^2} , \tag{3.10}$$

which have a standard deviation of 0.399 and a FWHM of 0.939 for unit a. As is familiar from statistical applications, the distributions are infinite but their values die away quickly after a few standard deviations. If a Gaussian PSF is imaged with a FWHM of n pixels, then setting $na = 0.939$, the transform becomes

$$F(u) = 1.06ne^{-3.56n^2u^2}, \tag{3.11}$$

with u in cycles per pixel, and a standard deviation of $0.375/n$. For $n{=}2$, the critical sampling frequency of 0.5 cycles per pixel is thus 2.7 standard deviations from the centre, where the spectral power density is 2.8% of the maximum; the total power aliased back into the main spectrum is 0.7%. Even with a large telescope at short wavelengths and no adaptive seeing correction, the two pixel criterion is not far from providing critical sampling.

For the case of an 8m telescope observing at 5500Å, the cutoff frequency is 71 cycles per arcsecond. With 0.3 arcsecond pixels, designed to exploit natural seeing of 0.6 arcseconds, this corresponds to 21 cycles per pixel—the image is very much seeing dominated and strictly speaking is badly undersampled. Nevertheless, the overall MTF is low for any components which are aliased, and if two pixels per FWHM is not sufficient, a slightly larger number almost certainly will be.

3.6   Sampling with fibres

The process of observing with an integral field spectrograph can be described in the same terms as the Sampling Theorem. The key difference between slit spectroscopy and densely-packed fibres lies in the arrangement of data along the spatial dimension which is reformatted onto adjacent detector pixels. It is therefore best to consider a slice along this axis, rather than trying to include dimensional scrambling and a non-rectangular input grid in the notation.

An integral field unit captures light from an object with intensity distribution $I(x, y, \lambda)$ $Wm^{-2}\,arcsec^{-2}\,Å^{-1}$, via a telescope and atmosphere system with point spread function $P_t(x, y, \lambda)$. The latter represents the spatial response to a point source at a particular wavelength, such that $P_t(x, y)|_{\lambda'} = P_t(x, y, \lambda')$. Co-ordinate axes x and y are centred on the field, and rotated by some position angle with respect to the celestial system. Since light is integrated over each IFU element into an unresolved sample, the observation is convolved with both the optical PSF and the spatial sensitivity, $L(x, y)$, of a microlens or fibre:

$$I_*(x, y, \lambda) \;=\; t_t(x, y, \lambda) \iiiint_{-\infty}^{\infty} I(x'', y'', \lambda)\, P_t(x' - x'', y' - y'', \lambda) \\ L(x - x', y - y')\, dx'\, dy'\, dx''dy''. \qquad (3.12)$$

(Asterisk notation is avoided, since three-dimensional functions are being convolved in two dimensions.) The proportionality constant, t_t, incorporates system throughput and any other multiplicative factors; the details are unimportant since absolute flux can be calibrated observationally.

Consider sampling $I_*^r(x, \lambda) = I_*(x, y_r, \lambda)$ along a particular row of the IFU. It is convenient to assume that $x = 0$ is a fibre centre. The pitch, p_x, should be no more than half the width of the total PSF, $P_t * L$. Multiplication by a shah function

yields an expression describing the position and intensity of each beam; however, the series is not really infinite; in order to retain the shah function, the image is truncated using a box, so that all but the central N_r impulses are set to zero:

$$I_f^r(x, \lambda) = \Pi\left(\frac{x}{N_r p_x}\right) \text{III}\left(\frac{x}{p_x}\right) t_f\left(\frac{x}{p_x}, \lambda\right) I_*^r(x, \lambda)$$

$$= \Pi\left(\frac{x}{N_r p_x}\right) |p_x| \sum_{n=-\infty}^{\infty} t_f(n, \lambda) I_*^r(p_x n, \lambda) \delta(x - p_x n). \qquad (3.13)$$

In this case, the box edge falls half way between each end sample and the next of the series. Fibre throughputs are specified at the impulses by $t_f(n, \lambda)$.

Since limiting the image size is like multiplication with a box, the effect is to convolve the true power spectrum with a narrow sinc function. This imposes a finite limit on the information density and lowest resolvable frequency. In effect, frequency components too low to measure are synthesized from higher ones over the extent of the observed image.

Spatial truncation also leaks power outside the interval of the principal spectrum, because the smoothing sinc function dies away slowly—or equivalently, because $\Pi(x)$ has sharp edges. This can cause a sudden cut-off at the Nyquist frequency when the sampled image spectrum is low-pass filtered. For finite images, reconstruction by equation 3.8 generates high-frequency ripples (the Gibbs phenomenon)—spatially, an infinite summation is needed for the sinc terms to cancel properly. More practical reconstruction filters are tapered in the frequency domain rather than box shaped; this causes some smoothing by supressing high frequencies, but reduces the sensitivity to discontinuities and aliasing.

So far, this model describes a normal imaging system. However, detection is not performed by the fibres—an extra stage is needed, involving resampling the information. Upon leaving the fibre bundle, $I_f^r(x, \lambda)$ is both dispersed and blurred in the spectrograph. The finite profile of each fibre's beam acts as a low-pass filter for the sampled image, smoothing out the discrete pattern of the IFU. In this case, however, there is no freedom to choose the exact form of the filter—only its width.

If the beam profile at the detector is a function of continuous pixel co-ordinates, $P_s(i, j)$, then the detector sees

$$I_s^r(i, j) = t_s(i, j) I_f^r(x(i, j), \lambda(i, j)) * P_s(i, j) * \Pi(i)\Pi(j), \qquad (3.14)$$

where x and λ map to i and j with some rotation and distortion. If spectra are unresolved then their separations must be uniform, or at least well known, in order to make sense of the image. The throughput, t_s, combines wavelength-dependent variations with spatial vignetting, so corresponds to a particular grating angle. The box functions account for integration over each pixel. In reality, P_s varies slightly from fibre to fibre depending on focal ratio degradation and microlens alignment, but should be dominated by the smoothly varying spectrograph PSF.

The final observation for the rth row of the IFU is described in terms of the analogue spectrogram by

$$I_d^r(i,j) = \sum_{n_i} \sum_{n_j} t_p(n_i, n_j) \, I_s^r(n_i, n_j) \, \delta(i - n_i) \, \delta(j - n_j), \qquad (3.15)$$

where t_p is quantum efficiency for a given $\lambda(i,j)$, or grating angle.

From this point of view, there are two reasons why imaging with fibres generates more complicated data than direct imaging or spectroscopy: prior to resampling, the telescope image is not reconstructed fully in the x direction by beam profiles, nor is it corrected for fibre throughput variations. But for these limitations, data would appear just the same as for a long slit or image slicer. Neither factor is important for distinct multi-object style spectra, because beam intensities can be derived individually from the spectrogram, which amounts to reversing the optical resampling. In contrast, unresolved spectra must be extracted by interpolation;[1] any high frequency distortion of the output pattern must be dealt with explicitly. Such problems are not familiar from electronic imaging, where flat-fielding and fine-tuned interpolation can be performed before any regridding occurs.

3.7   Properties of the resampled data

In a fibre spectrograph, the beam from each fibre can be approximated by a Gaussian profile. This shape cannot suppress the sampling pattern effectively without also causing significant smoothing (eq. 3.11). In any case, throughput variations prevent adequate reconstruction of the input image inside the spectrograph. The system is therefore designed with a PSF which is small compared to the two fibre FWHM at the telescope focus, rather than one which would recover smooth images.

---

[1] If there is at least one pixel across each spectrum and P_s is well known and stable, then extraction might be performed by deconvolution, at some expense in noise amplification.

Figure 3.7 shows some simulated fibre IFU data, from a single detector row at constant wavelength. For illustrative purposes, the simulations have 32 pixels across each fibre, so that the plotting software can approximate the analogue spectrogram just by connecting the points with straight lines. In reality, we are interested in fibre pitches ranging from one to several pixels; these would give a subset of the simulated points, with some additional smoothing due to the pixel size. The values plotted do not include any noise.

Fig. 3.8 shows the frequency content of the simulated images in fig. 3.7. The beam width of 0.5 fibres (left) only blurs the IFU grid enough to suppress the first sampling replicas by $\sim$60%, whilst the 1 fibre FWHM (right) removes $\sim$97% of their power. These spurious peaks represent the difference between an image which is sampled then reconstructed with the IFU beam shape and an image which is merely smoothed with the same profile. In contrast, convolution with a narrow sinc filter would give the same result whether applied to the continuous or discrete image (the former would be reproduced). The spectra in fig. 3.8 also deviate from the telescope image spectrum due to the nonuniformity of fibre transmissions. Because variations in throughput are essentially random (apart from vignetting, which will not obscure image features), they cause relatively small fluctuations throughout the spectra, repeating at multiples of the sampling frequency.

Modulation on the scale of the fibre pitch is the most obvious consequence of a poor reconstruction filter. However, intensity is not the only parameter which is distorted with respect to the analogue image. At each point, the illumination is a weighted sum of the sample values from nearby fibres. The weights and their centroid vary nonlinearly with the relative fibre distances according to the PSF. Hence if the PSF is not wide compared to the fibre spacing, the mean sky co-ordinate of light will vary nonlinearly between fibres at the detector. An extreme example of such behaviour occurs in nearest neighbour interpolation, where the effective co-ordinate changes discontinuously at the midpoint between original samples; much the same thing will happen for any narrow enough PSF. Of course, differences between fibre throughputs also influence the co-ordinate mapping; a bright fibre dominates a larger region between its neighbours than does a faint one. For example, consider fig. 3.9 as a flat-field section—pixel a receives almost equal contributions from fibres 1 and 2, whilst pixel b is weighted towards 2 rather than 3.

Figure 3.7: Some simulated output patterns from a row of fibres, for two input images and two PSFs. Dotted lines represent illumination at the IFU input (a flat field in the lower plots), and solid lines the detected image. As a fraction of the fibre spacing, each spectrum has a FWHM of 0.5 on the left and 1 on the right. Fibre throughputs are the same in each plot, but for comparison the beam width does not conserve intensity.
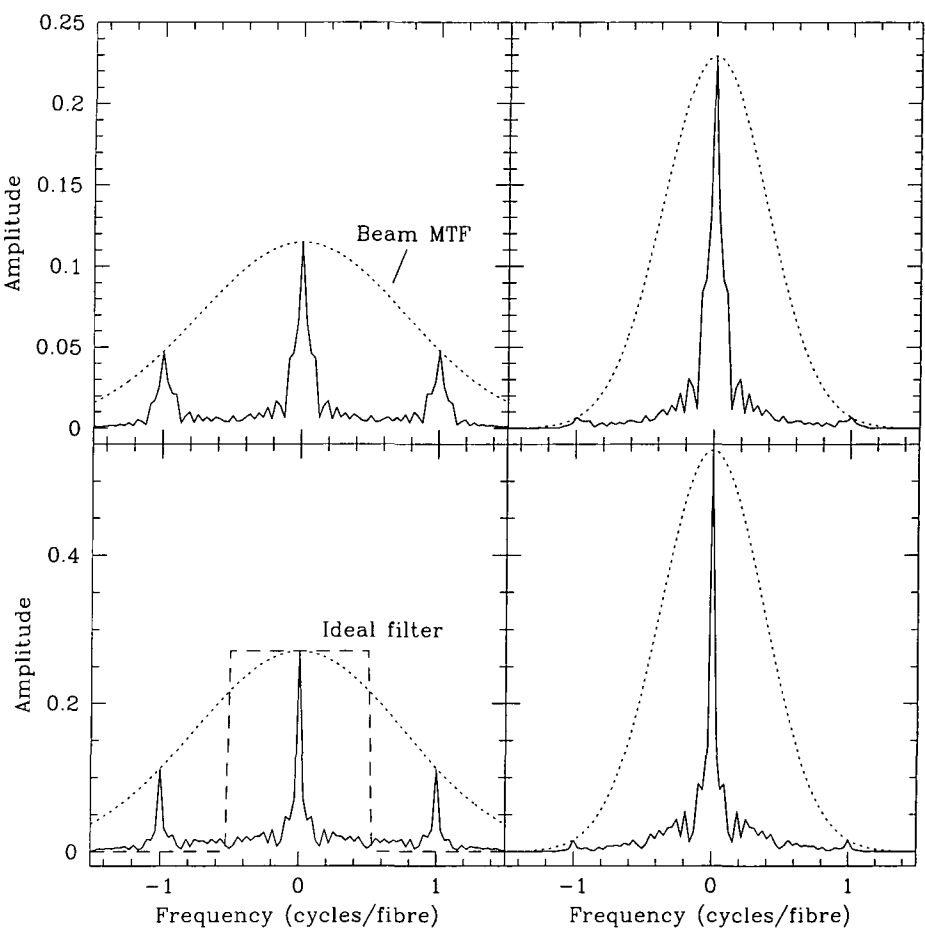
Figure 3.8: Fourier spectra of the simulated images in fig. 3.7, layed out correspondingly. Attenuation by the Gaussian PSF/MTF is shown normalized to the height of the zero frequency peak. The plots in fig. 3.7 are cropped slightly with respect to the data transformed here.

Figure 3.9: Sketch of some cross-sections of spectra at the detector.

On top of this co-ordinate distortion, power above the Nyquist frequency causes the spatial resolution of the data to vary. Convolution with a Gaussian beam shape increases the x co-ordinate variance of the light at each point (ie. suppresses high frequencies). Although convolution attenuates the Fourier spectra of continuous and sampled images in the same way, the residual sampling pattern can localize smoothing in the discrete case. For a narrow PSF, the original sample points are hardly smoothed by reconstruction at all, whilst the variance in x increases by the square of half the sample pitch at the midpoint between equally efficient fibres—as for linear interpolation.

If spectrograph images are stable between exposures, then the whole system may be flat-fielded, pixel for pixel, during reduction. This removes intensity modulation and throughput differences, but variations in co-ordinate mapping, resolution and signal-to-noise ratio are unchanged. Residual power at the sampling frequency is transferred back into the image spectrum, suppressing oscillation. However, the flat field has an impulsive spectrum whereas that of the image is extended (only marginally in fig. 3.8 because the image is a low-frequency wave). Flat fielding removes only the central portions of the sampling replicas, leaving behind surrounding structure which can affect image smoothness quite significantly.

Fig. 3.10 shows the effect of flat fielding the simulated fibre data. Distortion of the co-ordinate system becomes immediately obvious. For the narrower beam, it is clear that pixel values almost correspond to a nearest-neighbour interpolation between fibres. The pattern is slightly smoothed, however, because the beam shape has removed some of the high frequencies before flat-fielding. Practically, of course, intermediate pixels would become noise dominated as the beam width tends towards

zero, so the pattern would be less clean than in the simulation. The wider beam produces a flat-fielded image which is much smoother, but still rippled by co-ordinate distortion. In particular, the very low throughput fibre close to pixel 200 causes a noticeable error. Fortunately, the situation can be improved if information is available about individual fibre transmissions and the PSF.

Figure 3.10: The simulated images of fig. 3.7 (top) after flat fielding by division with fig. 3.7 (bottom). Once again, the beam profile has a FWHM of 0.5 fibres on the left and 1 fibre on the right.

The aim of integral field spectroscopy is to map a volume in x, y and λ. Where there is co-ordinate distortion, these variables are not represented well by pixel indices. Instead of reconstructing from uniformly spaced data, as in fig. 3.10, it makes sense to place each sample according to the mean co-ordinates from which light is measured. Spatially, a pixel is sensitive to the light falling on several fibres, each with a different weighting. Fig. 3.11 shows how this sensitivity is distributed, depending on alignment relative to the IFU output.

The centroid of a pixel's spatial sensitivity can be estimated from the positions of fibres at the detector, measured throughputs and perhaps a model PSF:

$$\langle x \rangle_{n_i,n_j} = \frac{\displaystyle\sum_{n_f} x_{n_f} t_{n_f} \int_{n_i-0.5}^{n_i+0.5} P_s(i - i_{n_f})\, di}{\displaystyle\sum_{n_f} t_{n_f} \int_{n_i-0.5}^{n_i+0.5} P_s(i - i_{n_f})\, di}, \tag{3.16}$$

Figure 3.11: Spatial sensitivity function of a detector pixel, given equal fibre throughputs and a beam FWHM equal to the pixel size. Column heights roughly indicate the relative contribution of each IFU element to the shaded pixel.

where x is the sky co-ordinate, n_i and n_f count pixels and fibres respectively, i varies continuously between n_i, t_{n_f} are fibre throughputs and P_s is the beam cross-section at (n_i, n_j). Since IFU elements at the input are symmetrical about their centres, their shapes are not needed in the calculation.

It is easy to see that an expression describing the flat-fielded spectrogram is obtained by replacing x_{n_f} with the intensity sampled by each fibre. Therefore if each pixel value is dominated by the two nearest beams, the co-ordinate adjusted points represent an approximately linear interpolation between fibre samples. Fig. 3.12 illustrates the effect of moving the simulated data to their effective positions on the sky; for comparison, the sky co-ordinate is mapped onto the original pixel index. Because fibre weightings are fixed instrumentally for each pixel, the adjusted points form the best available data for reconstruction; any improvement would require better image reconstruction in the spectrograph. Beam profiles cannot come close to the ideal sinc filter, but a wide PSF in equation 3.16 will smooth the linear interpolation, at some cost in resolution.

Figure 3.12: The flat-fielded data of fig. 3.10, before and after the x co-ordinate of each sample is adjusted. For clarity, the simulations are stacked top (FWHM 0.5) to bottom (1.0), instead of left to right. The adjusted points form the curves with linear segments, superimposed on the original curves of fig 3.10. New point positions are marked by ticks.

The tick marks in fig. 3.12 illustrate how adjusting the sample positions converts co-ordinate distortion into a variable point density. With a small number of pixels across each beam, samples are much more likely to fall close to fibre positions than to be equidistant from two fibres.

Resolution remains variable, as is the nature of linear interpolation. Compared to the telescope PSF, this effect is not large; nevertheless, it needs to be quantified because of its localization. It is useful to have a map of the instrumental contribution to the PSF, alongside noise and data quality arrays.

In an optical system, no distinction is normally made between the point spread function and the spatial sensitivity function of an image point, which are identical in form. However, the former maps a point on the sky to a region of the image plane, whilst the latter does the reverse. Although one speaks in terms of an observed PSF, it is really the spatial sensitivity which determines the resolution of data. A densely-packed fibre IFU breaks the equivalence between the two functions, so observation must explicitly be considered in reverse to find the resolution element. In effect, the sensitivity distribution of fig. 3.11 is the instrumental component of the observed image 'PSF'.

Individual PSF terms are characterized by their variances, since these add under convolution to give that of the product. If the variance of each pixel's spatial sensitivity is known, then the image resolution element can be estimated by adding the size of the seeing disc. Moreover, the Central Limit Theorem allows the final value to be converted to a Gaussian FWHM ($=2.355\sigma$).

The spatial sensitivity function of fig. 3.11 can be viewed as the convolution of the lens or fibre shape at the IFU input, $L(x,y)$, with a series of delta functions, normalized to fibre weightings. The input elements are hexagonal or square lenses, or circular fibres. Where an array is formed by tiling any one of these shapes, the variance of an element is the same along either of the array axes; in units of the diameter or distance between opposite faces,

$$\sigma^2_{\text{elem}} = \frac{\iint_{-\infty}^{\infty} [x^2 \text{ or } y^2] \, L(x,y) \, dx \, dy}{\iint_{-\infty}^{\infty} L(x,y) \, dx \, dy} = \begin{cases} 5/72 & \text{(hexagon)} \\ 1/12 & \text{(square)} \\ 1/16 & \text{(circle)} \end{cases} \qquad (3.17)$$

In the y direction, the variance of a pixel's sensitivity due to the instrument is just $\sigma^2_y|_{n_i,n_j} = \sigma^2_{\text{elem}}$, as long as the pixel only receives light from one IFU input row. Where spectra from the ends of different input rows overlap, or where the IFU has a spiral sampling pattern, the y variance is determined as for x below.

In the x direction, the instrumental variance is that caused at the IFU input plus the variance of the fibre weightings, corresponding to equation 3.16:

$$\sigma^2_x|_{n_i,n_j} = \sigma^2_{elem} + \frac{\sum_{n_f}(x_{n_f} - \langle x\rangle_{n_i,n_j})^2\, t_{n_f} \int_{n_i-0.5}^{n_i+0.5} P_s(i - i_{n_f})\, di}{\sum_{n_f} t_{n_f} \int_{n_i-0.5}^{n_i+0.5} P_s(i - i_{n_f})\, di}.$$

(3.18)

3.8 Extraction

Section 3.7 described how a set of samples in x, y and λ space can be derived from an IFU observation, by flat-fielding and characterizing the spatial sensitivity function of each pixel. In the special case of a multi-object style instrument, most pixels only receive light from the single nearest fibre, so groups of samples are coincident on the sky, forming discrete spectra. Generally, however, samples cannot be grouped into coherent spectra without interpolation. This is because the fibre spectra may be curved or tilted at the detector, so that their relative contributions at a given row or column change with wavelength—the spatial sample positions are wavelength dependent as well as nonuniform.

The issue of alignment changes affects image slicers and long slits, as well as densely-packed fibres. These instruments all produce data which are not arranged in one-dimensional spectra. However, the situation is more complicated with fibres, owing to the co-ordinate distortion discussed previously—even adjacent samples may drift quite differently with wavelength, along paths which are bumpy on small scales. Fig. 3.13 illustrates how smooth sub-pixel variations in alignment might affect the sample sky co-ordinates; a larger distortion or rotation of several pixels is likely, but is more difficult to sketch clearly.

If each IFU input row is narrow enough at the detector, any tilt or curvature will be practically constant across it. For a slicer, this means that x and λ vary linearly with the pixel index across each spectrum, so a row of pixels would describe a straight line in the x–λ plane. In that case, a regular spatial grid could be recovered by interpolating along each detector row—as long as both dimensions are critically sampled. With fibres, a row of pixels does not describe a straight line in x–λ, because of small-scale variations in $x(i,j)$. Resampling along rows is only safe if fibres are critically sampled in i and j. Otherwise, it is necessary to interpolate either between pairs of points, which is not ideal, or simultaneously in multiple dimensions.

Figure 3.13: Spatial distribution of samples for an IFU with densely packed fibres (left) and an image slicer (right). Wavelength variation of the sample points due to sub-pixel spectrograph distortion is sketched.

The traditional approach to fibre spectroscopy, where a spectrum is extracted for each fibre, differs considerably from the pixel-based method discussed in §3.7. In particular, a range of operations can be performed on extracted spectra before interpolating between them spatially to form one or more images. With slits or densely-packed fibres, however, spatial and spectral resampling are not completely separable; reconstructing a regular grid is more complex and has to precede any spectral analysis. Which paradigm is more appropriate for a particular instrument depends both on the separation of fibre beams and the extent to which they are individually resolved.

Consider three categories of fibre-based IFU, where, respectively, spectra are maximally packed, overlapping with resolved peaks, and well separated:

Maximally packed spectra: FWHM less than ~ 2 pixels in the spatial direction and comparable to the fibre separation (eg. TEIFU).

- Data are unevenly distributed within x-λ planes (or possibly all three dimensions), so deriving point spectra means resampling onto a grid.

- Since fibres generate distortion in the sky to detector mapping and the IFU output pattern is undersampled, interpolation should be performed in physical co-ordinates (eq. 3.16), after flat-fielding, rather than in pixels.

- There is no reason to interpolate spatially onto the original fibre input positions—one can choose any square grid of appropriate density, serving as a final reconstruction of the field of view. If spectra are placed on a

non-rectangular fibre grid, then further resampling will be necessary at a
later stage, probably degrading the resolution.

- Extracting fibre spectra by spatially integrating the interpolant around
  each one is like convolution with a boxcar prior to resampling. It is not
  generally necessary to do this. The reconstruction algorithm can fit an
  interpolant of the required resolution and the resulting datacube can be
  smoothed further if the astronomer deems it appropriate.

- Interpolation should ideally be weighted according to the signal-to-noise
  ratios of individual pixels. This is particularly true if there is more than
  one row per fibre—otherwise the overall noise will be marginally worse
  than for a multi-object style design, because pixels with slightly different
  signal levels are equalized by flat-fielding. This consideration may have
  to be balanced against the choice of efficient reconstruction techniques.

- In principle, one might extract fibre spectra by inverting a system of
  equations, relating the fibre intensities at each row to at least as many
  pixel values, via a known PSF. This would give the best possible reso-
  lution and avoid irregularity, but would degrade the signal-to-noise ratio
  because noise from neighbouring spectra cannot be subtracted along with
  the signal. Moreover, such deconvolution is sensitive to inaccuracies or
  changes in the PSF, and will fail if the number of rows per fibre drops
  too low. The dense-fibres paradigm instead accepts the overlap between
  spectra as one of many convolution terms (albeit one-dimensional).

Resolved, overlapping spectra: fibre separation of several pixels, comparable to
the FWHM in the spatial direction (eg. GMOS-IFU).

- As the pixel density increases with respect to the IFU output pattern,
  samples become bunched together in x, around fibre centres at the IFU
  input (fig. 3.12). Overall, resolution ceases to depend on the alignment
  of fibres with detector rows. Deriving spectra still requires some form of
  resampling, but this time it need not involve reconstruction in x and λ.

- If the fibre FWHM is at least two pixels, it is acceptable to interpolate
  onto fibre centres in pixel space. In that case, the IFU can be flat fielded
  after extraction. The degradation caused by this extra resampling stage
  will be minimal, because the interpolant can pass higher frequencies than

are present in the telescope image, and because each fibre is already close to the combined centre of the nearest pixels.

- The signal from each fibre is spread over several samples in cross-section. Although these extra samples provide no additional spatial information, their values are statistically independent. Before reducing the number of samples, their signal-to-noise ratio must be consolidated on the scale of the fibre pitch, reducing the high frequency content of the data. The simplest way to achieve this is to integrate across each fibre spectrum, producing one value per wavelength element.

- Extracting spectra this way is still like convolution with a boxcar, but degradation is minimized by the step-like influence of the fibres (fig. 3.10). The result is the same as for integration over a well-aligned pixel in the maximally-packed case. The effective size of the boxcar in x will vary slightly, depending on relative fibre throughputs.

- In the frequency domain, extraction has two effects. Integration over the fibre pitch multiplies the image spectrum (fig. 3.8) by a sinc function with zeroes at multiples of $\pm 1/p_x$, removing much of the power of the sampling replicas. Resampling at the fibres causes the dampened replicas to be aliased so that they coincide, in phase, with the central spectrum. This is equivalent to wrapping around the combined envelope of the sinc function and Gaussian MTF at the Nyquist frequency, so that it cancels out some of its own attenuation. In the limit of very narrow fibre images, the envelope would just be a superposition of sinc functions with unit value everywhere (eq. 3.8), ie. the original samples would be recovered.

- Because of throughput variations, extracted spectra will *not* correspond exactly to the sky co-ordinates of fibres. For example, pixels under a very faint fibre may actually receive most of their light from one of the two adjacent fibres. However, the point positions no longer depend on alignment and wavelength after extraction; perturbations in the fibre grid are smaller and fixed. The residual irregularity can be dealt with in one or two dimensions after spectral resampling.

- In summary, it makes sense to extract a spectrum for each fibre, rather than trying to reconstruct a smooth image from bunched samples which form linear segments. Fibre extraction is an effective way of consolidating

the full signal and removing artificial frequencies. The resulting data sets are also relatively compact. Where resolution is more important than signal-to-noise, deconvolving the fibre samples is more likely to work than in the maximally packed case.

Well separated spectra: fibre separation of several pixels, much larger than the FWHM in the spatial direction.

- All but the faintest pixels receive their light from the nearest fibre, so resampling onto the IFU grid does not require interpolation; each fibre spectrum can simply be summed in cross-section. If background light in the spectrograph is important, the signal-to-noise ratio can be optimized by fitting cross-sectional profiles in the usual way.

- It does not make sense to flat field individual pixels, since the illumination falls to a low level between fibres.

- Since the spectra barely overlap, extraction does not cause smoothing.

3.9 Image reconstruction from irregular samples

The following sections discuss how observations can be reconstructed in x, y and λ from the irregular sample sets produced by flat fielding and calibrating pixel coordinates. The reader who is more interested in understanding the data than in implementing reduction software may wish to skip to the conclusions.

3.9.1 Methods and requirements

Whilst a band-limited function can be recovered from regular samples by convolution, reconstruction from irregular samples is generally an inverse problem. Since the Fourier transform of the sampling pattern is no longer a comb function, it is not trivial to separate out the image spectrum. For particular nonuniform sampling patterns, such as interlaced regular grids, specific analytical reconstruction formulae can be found (eg. Bracewell, 2000). Sometimes a model may be fitted directly to the data. In general, however, one has to fit a convenient set of basis functions with as many degrees of freedom as match the original image resolution.

The basis functions used for reconstruction can take a number of forms, such as evenly spaced sinc or spline functions, Fourier series, trigonometric polynomials, or (more recently, eg. Graps, 1995) wavelets. Equally, there are various methods for determining the coefficients of basis functions from the samples. A system of equations may be solved directly or iteratively; the solution may interpolate the values, for critical sampling, or represent a best fit (approximation) if the bandwidth is overdetermined. Some example techniques are least squares fitting, iterative band-limiting or maximum-entropy methods and deconvolution of the sampling pattern from the discrete Fourier transform.

A band-limited function is fully specified by irregular samples, at least in one dimension, as long as the *average* sampling rate is greater than the Nyquist rate for a regular grid (Beutler, 1966; Feichtinger & Gröchenig, 1994). Formal proof of this criterion is more involved than in the regular case, but it is clearly equivalent to requiring that there are as many samples as harmonic frequencies (Fourier coefficients) over the extent of the reconstruction. The result might therefore be expected to hold in higher dimensions, although the author is not aware of a proof.

Despite the average spacing criterion, there is no guarantee that reconstruction will always converge or be numerically stable unless more stringent conditions are satisfied. For some methods and some basis functions in one or more dimensions, a solution is only certain to be found if the individual gaps between samples are smaller than the Nyquist interval (Strohmer, 1993; Feichtinger & Gröchenig, 1994; Strohmer, 1997). The mere existence of larger gaps does not, however, imply that a solution is unlikely. For a fibre IFU, only the average spacing can be expected to satisfy the Nyquist criterion; fortunately, deviations from a regular grid are not great, except close to damaged fibres, so highly unstable configurations are avoidable.

In astronomy, irregular sampling problems have mainly been discussed in the context of time series analysis. In particular, the literature deals with Fourier analysis of nonuniform samples in one dimension (eg. Deeming, 1975; Swan, 1982; Scargle, 1982, 1989). This subject is also mentioned in Numerical Recipes (Press et al., 1992). Techniques for resampling onto a regular grid are reviewed by Adorf (1995). Some robust iterative methods for reconstructing band-limited signals are presented by Vio et al. (2000). For more general one-dimensional data, interactive procedures for fitting splines, Legendre polynomials and other piecewise functions are routinely

implemented in data reduction systems.

In multiple dimensions, irregular sampling is less well understood and fewer algorithms are available. Some basic theorems no longer hold, and finding the closest samples to a particular point is no longer trivial. The Numerical Harmonic Analysis Group at the University of Vienna (http://tyche.mat.univie.ac.at) has been particularly active in developing both multidimensional irregular sampling theorems and practical image reconstruction methods. Feichtinger (1995) gives an overview of the issues involved. Since solving matrix equations directly becomes very slow with many samples, the Vienna group has presented fast iterative algorithms for reconstructing band-limited images (eg. Strohmer, 1993; Feichtinger & Gröchenig, 1994; Strohmer, 1997). Although these methods are computationally efficient and appear to be robust in practice, they formally assume that gaps in the data are smaller than the Nyquist interval. It is therefore not clear whether such procedures would be immediately suitable for resampling fibre IFU data.

Most surface fitting algorithms use some form of local, pointwise interpolation in the spatial domain. Formally, these methods are not ideal for band-limited data, because they do not model the underlying problem correctly. Nevertheless, their availability, robustness in some situations and relative simplicity makes them worth discussing. In computing and geophysics, many algorithms have been developed for triangulating scattered points in two or more dimensions. As an end result, however, bilinear interpolation over a mesh of points is not very satisfactory. There also exist smooth pointwise interpolation routines, which may provide reasonably accurate reconstruction where samples are not too deviant from a uniform grid.

A number of subroutines for generating smooth surfaces can be downloaded from Netlib (http://www.netlib.org) or as part of Starlink's PDA library (http://star-www.rl.ac.uk). Methods include both least squares fitting of splines with regular knots and piecewise polynomial interpolation with a node at each sample. Much of the code was originally published in Transactions on Mathematical Software. Similar subroutines are available in the commercial NAG library (http://www.nag.com). Code for fitting surfaces using Legendre and Chebyshev Polynomials is distributed with IRAF. Most of these methods have yet to be assessed in the context of IFU data reduction. Additional references for all kinds of sampling and reconstruction problems are given in Unser (2000).

It seems that there are very few algorithms for reconstructing smooth functions from scattered data in three dimensions. Searching Netlib reveals only one, namely QSHEP3D (Renka, 1988), which creates a smooth piecewise interpolant by fitting a quadratic function locally around each data point. Graphical toolkits such as VTK (http://www.kitware.com), Khoros (http://www-vis.lbl.gov/software/khoros.html), Matpack (http://www.matpack.de) and the commercial IDL (http://www.rsinc.com) may also be of some help.

3.9.2 Geometric considerations

Section 3.8 discussed why data are likely to be distributed irregularly within x–λ planes. An irregular distribution in y is also possible where spectra from the ends of adjacent input rows overlap, or where the input maps to the spectrograph slit in a spiral pattern. In the first case, most points will lie in x–λ planes, with a few points located unevenly in between; the problem is therefore three dimensional unless the deviant data are discarded. For the spiral pattern, most points are arranged in planes of different sizes and angles in x and y—a geometry which is badly suited to dimensional reduction. Either problem may be simplified by interpolating locally between pairs of points, forming a regular grid along the λ (or x) axis, but at the cost of significant and variable smoothing. The neatest solution is to reconstruct in x, y, and λ simultaneously, but this could prove computationally expensive and limit the choice of algorithms. Which is the best compromise will depend on the application, computer resources and reliability.

Usually, multi-dimensional interpolation involves independent variables of the same kind; working simultaneously with spatial and spectral co-ordinates is more complicated. Where local interpolation is used, the algorithm has to identify the nearest neighbours to each point, which is not a well-defined problem. Working in physical units is precluded unless the code is written or modified to be aware of the layout of pixels on the detector. In principle, it would make sense to calculate distances in terms of the resolution. Seeing varies, however, and if one dimension is better sampled than the others, a multi-purpose algorithm might estimate that the nearest neighbours are co-linear, causing it to fail. Another geometric problem is that, where there are dead fibres or gaps between spectra, adjacent pixels can sometimes receive light from the same mean co-ordinates. It is essential to remove

this duplicity, because interpolation will usually fail with coincident points. This can be done neatly by imposing a minimum threshold on intensity or resolution.

3.9.3   Noise

Reconstructing regular samples with a sinc function preserves the noise level of the original data, because the squares of the superimposed sinc values add up to one everywhere (see Bracewell, 2000). Smoothing interpolants give an even better signal-to-noise ratio at individual points. A key disadvantage of irregular sampling is that data are unevenly weighted with respect to the coefficients of a regular interpolant (ie. equally spaced knots or a harmonic Fourier series). This can lead to ill-conditioned systems of equations or large positive and negative weights in formulae, giving noisy and unstable results. The problem is worst where the number of data is similar to the number of unknown coefficients, particularly for bunched points, whereas oversampling can improve matters considerably (see Lauer, 1999).

Whilst oversampling is not an efficient option for a fibre IFU, noise amplification should be reasonably modest because points are just perturbed fractionally about a regular grid. With several samples across each fibre, the points become bunched in x but can be consolidated as discussed in §3.8. The whole problem could be avoided through local, pointwise interpolation, but this produces variable degradation of accuracy and resolution, depending on the distribution of samples—effects which are worst in the very cases where band-limiting reconstruction is noisiest. It may be possible to achieve better results by constructing special basis functions for a particular sampling pattern (eg. see Unser, 2000). This is a more complex problem and may be slow to compute, but could give the best results for datasets which are, on average, critically sampled.

3.9.4   Resolution and accuracy

When reconstructing a band-limited image, the bandwidth of the interpolant should match the instrumental resolution. With too few basis functions or knots, it will clearly be impossible to fit the data accurately, whilst too many parameters will lead to spurious high frequencies, determined by noise. Some iterative algorithms are able to adapt automatically to the data resolution, by incorporating bandwidth

as an additional variable (eg. Vio et al., 2000). As an inversion operation, the reconstruction problem is closely related to discrete deconvolution; the latter allows a greater bandwidth since a specific function is known to relate the measured values to a higher resolution image. Because overlapping spectra cause smoothing, the only case where it is permissible to solve for a knot at each fibre position is where the cross-sectional spectrum profile is used for deconvolving the original samples.

Since linear interpolation retains some artificial high frequencies from sampling, the question arises as to what happens to these when IFU data are treated as samples of a (roughly) linear interpolation between fibres (eq. 3.16). Actually, this operation is more naturally understood in the spatial domain. When a linear reconstruction is resampled regularly, aliasing tends either to suppress or increase the smoothing effect of the interpolant's frequency envelope—spatially, it is trivial to see that resampling at the original points recovers the original data whilst resampling at the midpoints is equivalent to having larger fibres and hence a lower resolution. Undersampling irregularly does not cause aliasing, even if it leads to inaccurate reconstruction, because the sampling pattern does not interfere coherently with signal frequencies (fig. 3.5). However, it is clear that the resolution of each sample (σ_x^2 in equation 3.18) is reduced, depending on alignment with fibres and their individual throughputs. For a fibre pitch of about one pixel, resolution will vary only slowly along the x axis with alignment. Designs with a fibre spacing of ~ 1.5 pixels are probably best avoided, because their resolution would vary rapidly. It may be desirable to smooth each row in x to a similar resolution, to avoid introducing variations in the y direction.

One obvious source of error with a fibre IFU is broken fibres. However, a small number of dead elements do not represent a very significant change in the sampling pattern; the average spacing is increased slightly and reconstruction is less well constrained at the gaps. For a regular grid, the transform is still a comb function of the intended period, but with some power leakage from the delta functions. Thus broken fibres affect the reconstruction accuracy slightly, but do not cause significant aliasing. With band-limiting interpolation, gaps can be filled in correctly if the sampling density is great enough, with some increase in noise. Local interpolation would keep the noise level low, but is inaccurate and will truncate any image peaks which happen to fall on the gaps. Vio et al. (2000) discuss a method for recovering lost samples from discrete data, which could be useful with a regular grid.

One way to compensate for undersampling or missing elements is by dithering, ie. shifting the IFU position on the sky by a small amount between observations and combining the results. This is discussed in detail in Lauer (1999). When a regular sequence of samples is offset with respect to the origin, its Fourier transform acquires a sinusoidal phase term, creating a phase difference between the central spectrum and its sampling replicas (fig. 3.3). When dithering is performed, an additional constant phase offset is applied, so that the replicas in the different images cancel out, eliminating aliasing. However, there is then a smaller phase difference between the principal spectra unless the spatial offset is exactly half the sampling pitch; hence the signal-to-noise ratio is degraded. This is an alternative explanation as to why reconstruction from bunched samples is noisy. Where undersampling is less of a problem, observations can instead be mosaiced, adding them after reconstruction or reconstructing with half the bandwidth. In that case, the principal spectra add in phase, minimizing noise, whilst sampling replicas add partly out of phase, so are still suppressed somewhat.

3.10   Conclusions

For IFUs with full length spectra, densely packed fibre designs maximize the spatial information which can be recorded by a given detector. As such, they provide a cost effective and relatively low-risk alternative to image slicers at visible wavelengths— see Haynes et al. (1998a,b, 1999) and Murray et al. (2000) for details. When slicer technology reaches maturity, fibre IFUs will continue to be suitable for retro-fitting to existing spectrographs or complex optical routing (eg. multiple fields).

The unique problems associated with reducing densely packed fibre IFU data have been discussed here in detail for the first time, with particular attention to the information content of raw data. Comparisons have been made with discrete fibre and image slicer designs, and some important differences from multi-object spectroscopy have been emphasized.

Where fibre centres are separated by about one pixel at the detector, spatial elements correspond naturally to detector pixels, rather than fibres. Reconstruction of three-dimensional observations from these samples has been discussed, with some pointers to the literature and available algorithms. Where fibre spectra are imaged

with FWHM of more than two pixels or much less than their separation, they can be extracted and processed individually before reconstruction. Due to crosstalk, their effective co-ordinates will not correspond exactly to the positions of IFU elements.

Future work will concentrate on software implementation of the ideas presented here, development of more general data reduction tasks for integral field spectroscopy and detailed investigation of reconstruction techniques. The first two tasks have been started in chapter 4, with the *imspec* package for IRAF.

Chapter 4

TEIFU data reduction: the *imspec* software package for IRAF

"We estimated then that the compiler would be done in about six months, and this six-month interval was to remain the estimated time to completion for the next two years." (John Backus, discussing the original ForTran compiler from 1956.)

The *imspec* IRAF[1] package is a prototype data reduction system for densely-packed fibre IFUs, specifically TEIFU. The software implements the pixel-based reduction paradigm which was introduced in chapter 3. Some basic analysis tools are also included. At present, the package focuses on specific instrumental requirements, rather than dealing explicitly with all the aspects of a complete reduction process. It is intended both as an end-user facility for TEIFU and as an experiment in working with high-density fibre IFU data. The following text summarizes the TEIFU data reduction process, with a brief description of each of the associated tasks.

4.1   Overview

This chapter discusses *imspec* as part of a framework for densely-packed fibre data reduction. Brief help pages are also available for the individual tasks. Source code is provided in the *imspec* directory of the attached compact disc (see directory listing in appendix A) and will be made available on the World Wide Web at a later date. The standard ISO 9660 file format is used. The code has been tested under Sun

---

[1]IRAF and associated documentation can be downloaded at http://iraf.noao.edu/.

Microsystems' Solaris and Mandrake Linux. Installation instructions can be found in the file *Readme*. Prior to compilation, the programs occupy ~ 0.5Mb of storage.

In the following outline, operations can be divided neatly into two categories—those which lead to the production of a 3D datacube (basic reduction stages) and those which take datacubes as input (analysis and manipulation tasks). The latter are essentially instrument-independent.

4.2 Basic design: data formats and handling

Although a 3D array is a natural representation of an IFU's observation space, not every IFU reduction process will result in a datacube. Some groups working on IFU software prefer to take velocity or line strength measurements from extracted 1D spectra, before interpolating spatially to form a 2D map. This can minimize storage and computation overheads and ensures, where undersampling occurs, that spatial aliasing does not introduce spectral variations before analysis is performed. On the other hand, some operations on IFS data deal with images rather than spectra—for instance, correcting atmospheric dispersion or deconvolution. It is therefore not always possible to avoid interpolating at an early stage. Moreover, the modularity provided by a datacube format is highly convenient and suited to interoperability. The final choice of approach depends partly on the instrument and application.

TEIFU fibre spectra (fig. 4.1) have a FWHM of ~ 1 pixel and a separation of ~ 0.93 pixels. Hence the fibres are unresolved by the detector and do not inherently provide a spatial grid for the data. Furthermore, owing to the design of the WYFFOS spectrograph camera, both spectra and wavelength contours are highly curved. The detector grid is therefore nonuniform when mapped onto a 3D observation space—more so than for a long slit or image slicer, for reasons explained in chapter 3. Extracting spectra involves interpolating in at least two dimensions. There is no real reason why the new sample points should correspond to fibre positions, especially when these are in a non-rectangular grid, so pixel values are interpolated directly onto a 3D datacube. With this approach, calibration involves finding x, y, and λ co-ordinates (sky position and wavelength) for each spectrogram pixel, rather than working in terms of fibres.

Since IFS combines imaging and spectroscopy, one potentially has to work with a

Figure 4.1: TEIFU arc lamp observation, showing the raw data format, from the September 1999 commissioning run. The 18 blocks of 56 fibre spectra are separated by gaps of ~ 3 pixels.

number of image formats of different dimensionalities. In addition, complex header information may be required to describe the mapping from one to another. Ideally, each data frame will have associated error and data quality arrays, which track such parameters as noise and resolution through each stage of processing. Given these requirements, as well as the goal of compatibility between software packages, it will be important in the future to have standard file formats and libraries for handling IFU data. At the time of writing, both the IRAF group in the US and the recently formed Opticon 3D Spectroscopy Working Group in Europe[2] are considering these issues. In the meantime, *imspec* uses its own semi-object-orientated approach, which internally treats each file format (eg. raw spectrogram, datacube, velocity map...) as a type of object, including image arrays and/or variables.

Object types are defined in *imspec.h*. Most also have associated subroutines in their own separate files. This saves replicating code, reduces the risk of errors and is helpful for buffering images. Like other modern instrumentation, IFS produces

---

[2]http://www.aip.de/Euro3D

increasingly large datasets, especially where datacubes are resampled on a fine pixel scale to aid visualization and preserve resolution in subsequent operations. Storage requirements are considerable, but reading or writing images one line at a time is very slow. Depending on how multi-dimensional data are accessed, it is not always possible for IRAF or the equivalent system to know how to optimize buffering. Building simple read-ahead mechanisms into routines for accessing image sections or spectra has cut some execution times from several minutes to a few seconds.

As a part of a Ph.D. project relating to a (currently) private instrument, the design inevitably has some limitations at present. A number of reduction processes which are not specific to IFS or TEIFU are not included. Infrastructure for handling headers is basic. The IFU geometry is described in a text file, rather than in headers, so the software is not suited to pipeline processing. A number of improvements for the future are listed in the file *todo*. Nevertheless, it is hoped that *imspec* will serve as a useful example for groups working on IFU reduction software, as well as a useable tool for reducing TEIFU observations.

4.3   Calibration and reconstruction in 3D

Calibration is simplified by the positioning of TEIFU at the Nasmyth focus of the William Herschel Telescope. In the absence of varying gravitational vectors, flexure is not a significant problem and has not been dealt with explicitly. Fig. 4.2 gives an outline of the basic reduction process. For each setting of the diffraction grating, a detector calibration image is generated, describing the mean x, y and λ co-ordinates, resolution and throughput for each pixel. The reduction process centres around this calibration image. In principle, the modular approach allows different reconstruction algorithms to be interchanged. Individual tasks in the flow diagram are grouped into higher level reduction stages by blue dotted boxes. Each of these stages is described in more detail in the following sections.

4.3.1   IFU description

This stage establishes how images are mapped to the spectrograph slit, and only needs to be executed once unless the hardware configuration in front of the spectrograph is altered. It is nevertheless useful to be able to change the IFU description

Figure 4.2: Flow diagram for basic TEIFU data reduction.

quickly; for TEIFU, a second input field was added between commissioning runs and the precise mappings were not known long in advance. Moreover, the mapping can be reversed if the number of mirrors in the optical train or the image orientation is changed. The information is stored in a text file, as illustrated in fig. 4.3.

The first line of the IFU description file, `pixshape hexagon`, specifies the shape of elements in the IFU input array(s). This is taken into account when calculating the instrumental resolution, ie. RMS distribution of the incident light, for each CCD pixel (see eq. 3.18). Options are `hexagon`, `square` and `point` (the last two are not needed for TEIFU, but were simple to implement). If this line is missing, the default is `point`, in which case the calculated resolution does not account for integration over microlenses at the input. There is no `circle` option at present, because the light collecting area of a bare fibre depends on the amount of dead space between fibre cores.

The keywords `pitch` and `FWHM` are initial values for the fibre separation and spectrum FWHM in detector pixels. Like `pixshape`, these may appear once on separate lines anywhere in the file. Note that calibration will fail if `FWHM` is too small compared to `pitch`.

The keyword `output`, on a line of its own, begins the section describing the fibre ordering at the spectrograph slit. This consists of lines which either specify a gap between adjacent fibres or contain a field number (for multiple fields) followed by a fibre number or range of consecutive fibre numbers. In the raw data, the first listed fibre has the lowest pixel index along the slit. For TEIFU, spectra are arranged in blocks of 56 regularly spaced fibres with gaps of ~ 3 pixels in between. Each block corresponds to two rows of one of the two input fields. Every input field for which there is an entry in the `output` section must be illuminated in the flat-field image used for calibration (§4.3.2).

For each input field, the positioning of fibres on the sky is specified by a section beginning `field N`, where N is a positive number. Each fibre listed in the `output` section must have a matching entry in an input field, containing a fibre number, x and y co-ordinates and a throughput. In the example file, throughputs have been set to 1 or 0, for good and broken fibres respectively; more precise values will improve the calibration accuracy. Co-ordinates are listed in units of the fibre spacing along each input row; these are scaled during reconstruction, according to the actual image

Figure 4.3: Format of the instrument description file

```
Pixshape hexagon

Pitch 0.93
FWHM 1.0

Output

2 1:56
gap
1 1:56
gap
2 57:112
```

⋮

```
Field 1

#            (normalized units)
# Element   delta_x     delta_y     throughput

    1        -7.361      13.750       1.00
    2        -7.361      12.750       1.00
    3        -7.361      11.750       1.00
    4        -7.361      10.750       1.00
```

⋮

```
Field 2

#            (normalized units)
# Element   delta_x     delta_y     throughput

    1         7.361      13.750       1.00
    2         7.361      12.750       1.00
    3         7.361      11.750       1.00
    4         7.361      10.750       1.00
```

⋮

scale in arcseconds. For TEIFU, y is defined as the direction parallel to input rows, contrary to the convention in chapter 3.

The output section of the IFU description file is fairly straightforward to enter manually. The field sections are more complicated, but can be generated with the aid of the task *ifudef*. This takes some geometric parameters, defining a regular grid of hexagonal or circular elements over a rectangular field, and appends a field section to the specified text file. The IRAF parameters are as follows:

```
   outfile = "test.def"      Output text file with mapping
    (field = 1)              Input field number to add / replace
    (rowax = "dec")          Row axis at PA 0 (ra/dec)
 (pattern = "raster")        Input sampling pattern
  (altlen = no)              Row lengths alternating by 1el?
  (stagger = "-")            Sense of even row offset relative to odd
 (startel = "+")             Centre -> element 1 sense along rows
 (startro = "-")             Centre -> element 1 sense across rows
  (slicel = 28)              Number of row elements (max)
  (nslice = 18)              Number of rows
```

More details can be found in the *ifudef* help pages. There is no facility at present for generating rectangular grid mappings.

Ideally, the IFU description would be written into image headers by instrument control software. This would aid pipeline data reduction and could automatically allow for changes in configuration, such as varying image magnification at the IFU input. Future facility-class instruments will operate this way, probably using standard libraries for handling the information.

4.3.2 Basic calibration

Basic calibration is performed using *ifucal*. This uses a raw flat-field image to trace blocks of spectra on the detector, determining x and y co-ordinates for each pixel according to eq. 3.16. The resulting calibration image is used for reference when reconstructing observations taken with the same hardware set-up as the flat field. Application of a flat field image to the data is a separate stage (§4.3.4). The main parameters of *ifucal* are as follows:

```
   flatim = "r291068"       flat field image suitable for tracing
```

Table 4.1: Values stored for each detector pixel in the 8 planes of the *imspec* calibration image.

1   Field number (0 for non-illuminated pixels).
2   Mean x co-ordinate of incident light.
3   Mean y co-ordinate of incident light.
4   Mean λ co-ordinate of incident light (initially set to the row number).
5   RMS distribution of light in x due to the instrument (not inc. seeing).
6   RMS distribution of light in y due to the instrument (not inc. seeing).
7   Blank; reserved for RMS distribution of light in λ.
8   Calculated flat-field intensity, from fibre throughputs and positions.

```
    calim = "testcal"        output calibration image
  (ifumap = "test.def")      instrument description file with I->O map
 (npoints = 10)              num of bands in which to find block centres
 (sumdist = 5)               half band (pix) to sum when finding centres
 (imbound = "51 1074 4 1024") useful image boundaries
(interact = yes)             plot measurements?
```

Blocks of spectra are located at npoints different wavelengths along the detector, using an intensity thresholding algorithm. This calculates the proportion of the image which is directly illuminated, from the number of fibres and approximate pitch given in ifumap. The scattered light level and mean fibre intensity are then estimated, and the largest dips with approximately the correct spacing are identified as inter-block gaps. The mean fibre intensity is re-estimated for each block, allowing the half-maximum points to be measured more accurately. From these, the centres and widths are calculated. Any dead fibres at block edges are excluded from the calculation. The thresholding method is found to be robust as long as dead or vignetted elements at block edges are accounted for in the throughput list. Sub-pixel accuracy is achieved, although this might be improved by cross-correlating each block with the listed fibre throughput pattern as a final stage.

In interactive mode, *ifucal* plots the block measurements at each wavelength as in fig. 4.4. Fibre pitch estimates are constrained to vary smoothly along the slit, using a quadratic polynomial fit. The centre and fibre pitch for each block are then fitted as a function of row number (wavelength) using quadratic polynomials (fig. 4.5).

After tracing each block of fibres, *ifucal* generates a calibration image, containing 8 planes of the same dimensions as the flat-field image. The 8 values stored for

Figure 4.4: Measurement of block positions in two flat-field images, from the June 1999 (top) and Sept 1999 (bottom) commissioning runs. The second TEIFU input field was completed between the two runs, raising the number of output blocks from 9 to 18. Horizontal bars indicate the width (not shown accurately) and half-maximum of each block, as plotted by *ifucal*. The September data show some vignetting at the block centres, which correspond to row ends at the input. A few dead fibres are apparent in both cuts.

Figure 4.5: Quadratic fits to the measured centre positions and fibre pitches for block 1 in the June 1999 flat-field image, as plotted by *ifucal*.

each detector pixel are listed in table 4.1. These are calculated from the measured positions and listed throughputs of the three fibres nearest each pixel (eq. 3.16), assuming Gaussian profiles of the specified FWHM. The approximation that more distant fibres make a negligible contribution is reasonably accurate for `FWHM` $<\sim$ `pitch`. Significantly larger FWHM values will cause the estimated flat field (plane 8) to be less accurate, although calculation of the spatial co-ordinates, which is more important, is relatively insensitive to all but the nearest fibres and does not depend strongly on the profile shape.

Currently, the execution time of *ifucal* is determined equally by disk access and integration of Gaussian profiles over each pixel; modifying the program to use a look-up table for the latter would allow the five nearest fibres to be taken into account with little penalty. Ideally, information as to how the PSF varies across the detector would also be incorporated in the calculation, although this might tend to make the software quite instrument-specific.

Fig. 4.6 compares a June 1999 flat-field image with the flat-field plane of the derived calibration image. Fibre throughputs were listed as 1 or 0 in the instrument description file, except where this resulted in misidentification of block edges; more detailed characterization has yet to be performed. Without all the fibre variations, the calibration image reveals a small oscillation in intensity due to the gradual change in alignment of pixels and fibres. The resolution in y oscillates similarly, with large peaks in the centre, where spectra from the opposite ends of IFU input rows overlap. As well as the information listed in table 4.1, the calibration image stores the polynomial fits to fibre positions in its header. The wavelength co-ordinates in plane 4 are initially set to detector row numbers, allowing quick-look reconstruction before full wavelength calibration.

4.3.3 Wavelength calibration

Wavelength calibration is performed with the help of standard *onedspec* tasks in IRAF. An *imspec* task called *ifuarc* extracts a few one-dimensional spectra per block from an arc lamp observation. This is acceptable, since spatial resolution is not required. The traces stored in the calibration image header are used for reference. Spectra are stacked in a two-dimensional image, with their original block numbers and offsets recorded in the header. Parameters are listed below:

Figure 4.6: Comparison of June 1999 TEIFU flat field (top) with the final plane of the detector calibration image derived from it (bottom). The inset figure shows overlaid cross-sections of the two images, for two of the blocks. The smooth curve is the calibration image; fibre throughputs have been approximated by 0 or 1.

```
    arcim = "r291069"      wavelength reference frame
    calim = "testcal"      calibration image
  outspec = "arcspec"      output image with extracted spectra
   nspecb = 4              number of spectra to extract per block
```

Wavelength solutions are determined for the sample spectra using *identify*, and written to their header using *dispcor* (both *onedspec* tasks). Another *imspec* task, *ifuwcal*, then reads the solutions and fibre offsets from the header, calculating a λ value for each pixel of the original detector image by interpolating linearly between the reference spectrum positions. The λ co-ordinates are written to plane 4 of the calibration image, in place of whatever values were stored there previously. The parameters for *ifuwcal* are just the names of two files:

```
  refspec = "arcspec"      image with calibrated reference spectra
    calim = "testcal"      detector calibration image to update
```

If possible, the above process may be simplified in future using a CL script. In addition, *ifuarc* should be made to avoid dead fibres automatically.

4.3.4   Pre-processing

This stage involves bad pixel and cosmic ray correction, bias subtraction and flat fielding. There are currently no dedicated tasks for these purposes. Because detector images are undersampled, cosmic ray removal is relatively difficult—it is advisable to take at least three observations of each target object, so that median filtering can be used. For TEIFU, the bias level can be estimated from the detector overscan region and gaps between blocks. Ideally, tasks should be made available for removing the spectral content from flat-field observations and for removing cosmic rays from single images. With some improvements to the calibration process, plane 8 of the detector calibration image could also be used for flat fielding.

Where nodding or beam switching between the two TEIFU fields is used for background estimation, observations can be sky subtracted during the pre-processing stage, pixel for pixel. This has the advantage of ensuring that the spectral PSF is the same for sky and object (see Allington-Smith & Content, 1998). Otherwise, it makes sense to perform sky subtraction after generating datacubes. Sky spectra can then be averaged spatially, improving the signal-to-noise ratio by $\sqrt{2}$.

4.3.5 Reconstruction

Datacubes are reconstructed from pre-processed observations using *ifurec*, whose parameters are as follows:

```
    obsfile = "327pp"         Input spectrograph image
      field = 1               IFU field to reconstruct
    outcube = "327rec"        Output datacube
   (calfile = "testcal")      Detector calibration image
 (threshold = 0.5)            Minimum throughput threshold
   (imscale = 0.25)           IFU image scale used (arcsec/elem)
  (spixsize = 0.0625)         Angular size of output pixels (arcsec)
  (wpixsize = 0.)             Output wavelength step (0 for default)
```

Pixels which receive too little light in the flat-field image, as specified by threshold, are excluded from the reconstruction. The cut-off value is measured as a proportion of the light incident on one fibre at the IFU input. Pixels with a spatial co-ordinate RMS greater than the fibre input spacing are also excluded; these correspond to some dead fibres and to regions between fibres which are not adjacent on the sky.

The current reconstruction algorithm begins by interpolating linearly between pixels in consecutive detector rows, approximately in the dispersion direction, onto a uniform grid of wavelength planes. This would not normally be a desirable approach in terms of preserving resolution, but simplifies the difficult reconstruction problem considerably and is more appropriate for the undersampled TEIFU spectra. At each regular wavelength step, a smooth two dimensional interpolant is constructed over the scattered points in x and y, using the subroutine IDBVIP from the Starlink PDA library. An algorithm based on solving for the coefficients of a band-limited interpolant might be more accurate and perform better at dead fibres, but no tests have yet been performed to establish whether the resulting execution times would be reasonable. Reconstructing a single TEIFU field currently takes perhaps 2 minutes on a shared Sun server, so degradation in performance should be avoided if possible.

Because wavelength contours are highly curved at the detector, the first and last few wavelength planes of the output grid only contain a small number of points. Interpolation may fail in those planes, in which case they are set to zero and a warning message is issued. This is not normally a cause for concern.

4.3.6   Post-processing

For maximum signal-to-noise, or where one IFU field is used only for background observations, sky subtraction can be performed at this stage. Reconstructed sky datacubes can first be smoothed spatially in whatever way gives good signal-to-noise, taking care not to combine spectra with substantially different PSFs (eg. from the middle and end blocks on the detector).

4.3.7   Quick-look imaging

Whilst observing, it is useful to have a task for reconstructing wavelength-integrated images directly from raw data, for target acquisition. This is much faster and simpler than generating whole datacubes, with 1000 images at different wavelengths. The spatial content of a quick-look image does not have to be very precise, so a single value may be estimated for each fibre after extracting spectra in pixel co-ordinates or by linear interpolation.

The *ifuimrec* task, used for TEIFU commissioning, was not of suitable quality for inclusion in the current version of *imspec*. It is expected that a new version will be available for future runs, based on some of the *ifuarc* code.

4.3.8   Flat field analysis

Measurement of individual fibre throughputs is a key problem which has yet to be addressed. These are required if the y co-ordinate for each pixel is to be calculated as accurately as possible (eq. 3.16) during basic calibration. In principle, values could be determined in the laboratory, but it is difficult to illuminate IFU elements individually and to simulate a telescope and spectrograph accurately. In particular, it is important that throughputs listed in the instrument description file take vignetting into account as well as the intrinsic fibre transmissions. A method for analyzing flat-field observations is therefore required—a non-trivial task, given that fibres are unresolved at the detector.

Before fibres can be characterized, it is essential that the calibration process is as accurate and reliable as possible. Ideally, some further refinements of the present code should be implemented. At that stage, the curvature of spectra on

the detector might be put to good use—since the spectral content of a flat-field image is redundant and most block centres are displaced by several pixels from one end of the detector to the other, there are effectively several samples across each fibre in the spatial direction. An iterative method could perhaps be used to find the best-fitting fibre intensities, modelling the PSF width as a function of detector co-ordinates using ray-tracing software.

4.4   Datacube manipulation

The *imspec* package includes two tasks for spatial mosaicing, which were originally written for the SMIRFS-IFU. These are less well coded and commented than the core tasks, but are included because they have proved useful. Headers must be in a simple format, including `CDELT` keywords. An outline of the mosaicing process is given in fig. 4.7. Other operations on datacubes can be envisaged, such as deconvolution of the adaptive optics PSF or correction of atmospheric dispersion, but are not provided for by *imspec*.

The first task, *mosaic*, is used to combine datacubes with deliberate pointing offsets, enlarging the field of view. This has five parameters:

```
   incubes = "pos1,pos2,pos3" List of x,y,lambda datacubes to add
   outcube = "mos1"           Output datacube
   offlist = "offsets.txt"    File containing RA and DEC offsets (arcsec)
  posangle = 50.              Position angle (degrees) E of N
(maskname = "mos1m")          Optional output S/N mask
```

Input cubes are interpolated onto a common grid using tapered sinc functions. Overlapping pixels are then added together. To recover correctly scaled images, the resulting cube must be divided by the two-dimensional output mask (`maskname`), which tracks the number of input pixels contributing at each spatial position. The division is not done automatically because output images may themselves be mosaiced using *comcubes* (see below). If the input cubes need to be weighted before summation, this should be done explicitly prior to running *mosaic*. All the cubes must be on the same wavelength scale.

The offsets between input datacubes, in right ascension and declination, are listed in arcseconds, in the text file `offlist`. This contains one line for each input file,

Figure 4.7: Flow diagram for mosaicing datacubes using tasks written for the SMIRFS-IFU.

in the same order as for the parameter `incubes`. The values are converted to pixel offsets in x and y for the appropriate IFU position angle (`posangle`), measured from North to the y axis. In addition, any pixel offsets specified in the third and fourth columns of `offlist` are added, allowing positions to be tweaked where appropriate. The output datacube is just large enough to contain all the input points. An asterisk at the end of a line in `offlist` causes the corresponding datacube to be excluded from the mosaic without affecting the output dimensions; the user can therefore inspect the alignment of a subset of the cubes. An example *offlist* for three datacubes looks like this:

```
# d_ra    d_dec  dx dy
 0.000   0.000
-3.646   1.263   -2 -1
 3.646  -1.263    0  2
```

The second mosaicing task, *comcubes*, is similar to *mosaic*, but calculates spatial offsets automatically. This requires that a single bright peak, eg. galaxy nucleus, is contained in all of the input cubes. The centroid of the peak is found and the datacubes are resampled accordingly. Parameters are as follows:

```
   incubes = "mos1,mos2,mos3" List of x,y,lambda datacubes to add
   outcube = "moscom"            Output datacube
  (inmasks = "mos1m,mos2m,mos3m") Optional list of x,y input masks
 (maskname = "moscomm")         Optional output S/N mask
 (centzmin = 5)                 First plane contributing to centring
 (centzmax = 1000)              Last plane contributing to centring
```

Comcubes can combine spatial masks, as generated by *mosaic*, to create an output mask showing the total number of observation pixels which have been added at each point in the final mosaic. Once processing of error and data quality arrays is implemented, it would be better to store masks as FITS extensions in the main image files. The parameters centzmin and centzmax specify the first and last wavelength planes which are combined to form an image for locating the spatial peak. For IFUs with relatively large fields of view, hence multiple image peaks, a cross-correlation technique may be more appropriate for aligning observations spatially.

4.5   Analysis

4.5.1   Kinematic mapping

Basic provision is made for velocity and line strength mapping, as shown in fig. 4.8. In general, kinematic maps can be derived by converting datacubes into a standard two-dimensional format, with a spectrum in each row, using *dctomult*. A spatial intensity threshold can be applied, to exclude spectra with poor signal-to-noise:

```
   infile = "n4151"        Input x/y/lambda datacube
  outfile = "n4151_2d"     Output MOS-type image
  (thresh = -0.025)        Faintest spectrum (-ve = fract of max)
 (thrunit = "counts")      Threshold units (mean)
```

Analysis is then performed with any velocity mapping task which can provide a simple table of results as output, such as *fxcor* in the *rv* package or *FCQ* (Bender,
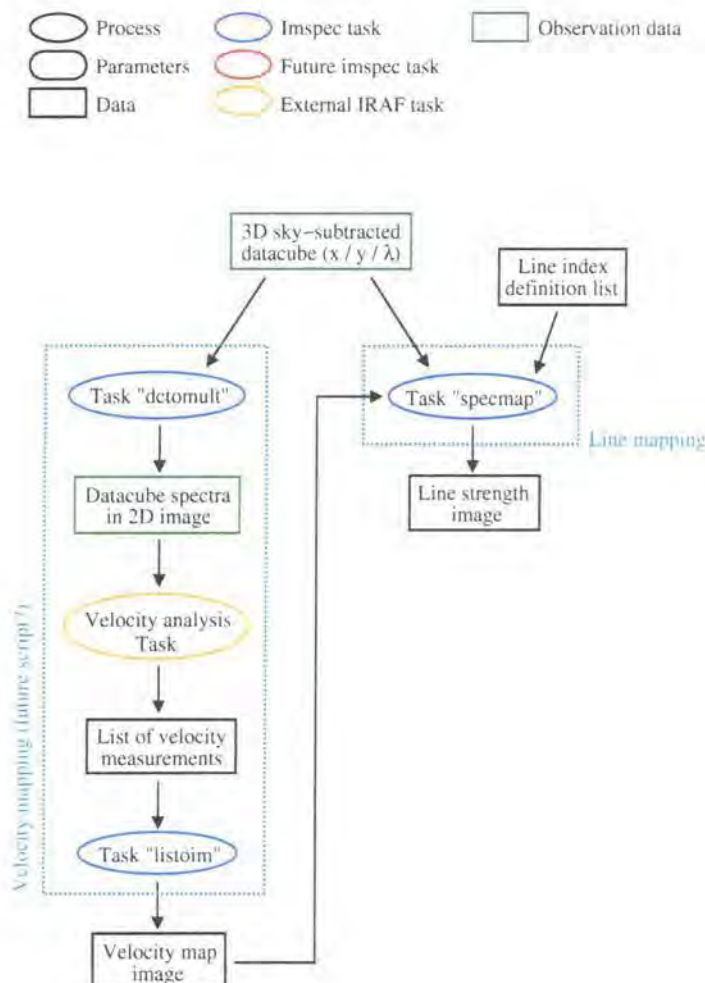
Figure 4.8: Flow diagram for TEIFU velocity and line strength analysis.

1990). For emission line velocities, *rvidlines* can operate directly on datacubes, as can tools which are currently being developed by the UK Starlink project.

Once velocity or line profile measurements have been made for each spatial point, *listoim* can be used to convert a text file of results into a two-dimensional image. The values to be mapped must be listed, one per line, in a column of the text file. The standard output from *rv* tasks is suitable for this purpose:

```
  infile = "vels.txt"     Input text file
  column = 2              Column containing values to map to pixels
 outfile = "velmap"       Output image corresponding to datacube
 (numcol = 1)             Column with row nums or image sections (0=none)
(mapping = "")            Row#->pix map ('nx,ny', filename or ''=infile)
```

The spatial mapping between listed values and image pixels is determined in one of three ways. The output file from *dctomult* contains header information defining the correspondence between the selected spectra and image pixels. If a filename is specified in `mapping`, *listoim* reads the appropriate header information from there. It is not, however, always necessary to specify a filename, as *rv* tasks write an image section to each line of output, in the format `filename[*,n]`. If `mapping` is blank, *listoim* attempts to read a filename from the column of `infile` specified by `numcol`. The required header is then retrieved from that file. Finally, the values in `infile` can be written in order to a rectangular array, by entering the output dimensions in `mapping` in the form `nxpix,nypix`. In the first two cases, it is permissible to delete lines from the text file, blanking the corresponding pixels. In the third case, there must be an entry for each pixel of the output array.

4.5.2   Spectral line mapping

Line strength mapping is performed using the final *imspec* task, *specmap*:

```
    incube = "n4151"        Input x/y/lambda datacube
     index = "PaB"          Line index to measure
    velmap = "velmap"       Reference velocity/redshift map image
    outmap = "PaBflux"      Output line map image
  (maptype = "flux")        Type of measurement
  (deflist = "Jlinedefs")   Text file with line index definitions
  (veltype = "cz")          Type of velocity units
  (fluxcal = no)            Are vals calibrated in flux / unit wavelength?
 (interact = yes)           Preview measurement parameters?
```

Line index names and wavebands are defined in the text file `deflist` and selected for mapping via `index`. A velocity map, `velmap`, matching the spatial datacube dimensions is required for reference; this can be in units of redshift, z, velocity, cz, or relativistic velocity, v, as specified by `veltype`. The output image, `outmap`, represents line flux or equivalent width, depending on `maptype`. The format of `deflist` is one index definition per line of text as follows:

```
# J-band line index definitions for specmap task

PaB     cec   1.2786-1.2804  1.2804-1.2832  1.2832-1.2848
[FeII]  cec   1.2496-1.2512  1.2548-1.2584  1.2588-1.2598
```

The first column gives a name to each index. The second consists of a string of letters corresponding to band types, *a* for absorption, *c* for continuum and *e* for emission. There must be at least one continuum band for each index; the present code allows a maximum of two, estimating the continuum at intermediate points by linear interpolation. If an index definition contains multiple feature bands, they are treated as sections of a single, larger band. The third and subsequent columns of `deflist` specify limits for each of the bands listed in the second column. The same wavelength units should be used as in the datacube header (it would be preferable for the units to be independent in future).

Measurement bands and values can be overplotted on individual spectra by executing *specmap* with `interact=yes` (fig. 4.9). An interactive loop then allows the user to inspect the results at arbitrary spatial pixels (see help pages). On exiting the loop, the complete set of results can optionally be saved to the image `outmap`.


Figure 4.9: Example plot from *specmap* in interactive mode. Green lines mark the continuum bands and estimate. Yellow lines mark the emission band and mean flux. The measured flux or equivalent width is printed at the bottom left. Wavelength and intensity values at the cursor can also be displayed.

4.6   Summary of progress and future development

The *imspec* IRAF package successfully implements the pixel-based paradigm for reducing fibre IFU data, introduced in chapter 3. Thus the practical potential of densely-packed fibre IFU designs is demonstrated. The present set of tasks enables calibration and reconstruction of TEIFU data, as well as making basic provision for mosaicing and analysis of datacubes. The latter capability has been used to help produce science results from SMIRFS-IFU observations.

In future, image reconstruction can be made more accurate on the scale of a fraction of a fibre, through detailed analysis of fibre throughputs, testing of different reconstruction algorithms and further refinement of the basic calibration stage. Some other key areas for future work are quick-look imaging, propagation of error images, automation of wavelength calibration, cosmic ray removal, flux calibration, flat field processing, improvements to the mosaicing tasks and generalization of the software for use with resolved fibre spectra or image slices.

IRAF is distributed by the National Optical Astronomy Observatories, which are operated by the Association of Universities for Research in Astronomy, Inc., under cooperative agreement with the National Science Foundation.

Chapter 5

The origin of [Fe II] emission in NGC 4151

The centre of NGC 4151 has been observed in the J-band with the SMIRFS-IFU on UKIRT. A map of [Fe II] emission is derived, and compared with the distributions of the optical narrow line region and radio jet. It is concluded that [Fe II] arises mainly through photoionization of gas by collimated X-rays from the Seyfert nucleus. The velocity field and strength with respect to Paβ are found to be consistent with this argument. The performance of the IFU is considered briefly, and techniques for observation and data analysis are discussed.

5.1 Introduction

5.1.1 [Fe II] emission in AGN

Seyfert galaxies exhibit strong emission lines from forbidden electronic transitions—notably those of oxygen, nitrogen and sulphur at visible wavelengths (Seyfert, 1943) and iron in the near-infrared (eg. Simpson et al., 1996). These transitions can occur where ions are excited by collisions with electrons yet can decay radiatively before subsequent collisions lead to de-excitation (eg. Peterson, 1997; Osterbrock, 1989). Hence the lines are associated with low-density regions of ionized hydrogen, where free electrons are abundant but the timescale for collisions is sufficiently long.

Some species, including neutral oxygen and Fe$^+$, have similar (re)ionization potentials to hydrogen, so do not survive where it is fully ionized. The most favourable conditions for [O I] and [Fe II] emission arise in partially ionized zones, the extent of which influences the observed intensity. Since H II regions around stars nor-

mally give way sharply to neutral gas beyond the Strömgren radius, their [O I] and [Fe II] lines are relatively weak. The central spectra of Seyferts, in contrast, indicate the presence of extensive partially-ionized hydrogen clouds as well as very highly-ionized species. These differences are explained somewhat naturally if the emission line clouds in Seyferts are photoionized by the 'power-law' continuum of the central AGN, which is spectrally flatter than a black body (Osterbrock, 1989; Mouri et al., 1990). Numerical simulations (eg. Ferland & Netzer, 1983; Mouri et al., 2000) lend credence to this scenario.

The second mechanism which can generate enhanced zones of [O I] and [Fe II] emission is shock excitation of gas (Osterbrock, 1989). This appears to be the dominant cause in starburst galaxies, where shocks due to supernova remnants are important and supernovae have been associated directly with enhanced [Fe II] (Lester et al., 1990; Greenhouse et al., 1991; Forbes et al., 1993). Comparison of observed spectral line ratios with recent numerical models helps to distinguish between shocking and photoionization (Mouri et al., 2000), and indeed proves consistent with the former dominating in starbursts and the latter in Seyfert nuclei. Doppler broadening can also be used to check for kinematic disturbances.

Long-slit spectroscopy of Seyfert galaxies has revealed spatially extended narrow line regions (ENLR) in the visible, beyond the central NLR. Unger et al. (1987) traced the bright [O III] lines at 4960Å and 5007Å, generated in fully-ionized regions, along with Hβ (4861Å) in a sample of galaxies. The features were shown to have minimal line-of-sight or extra-rotational velocity structure, and ratios characteristic of high excitation levels, pointing once again to photoionization. Subsequent studies (eg. in this case Penston et al., 1990; Robinson et al., 1994) have tended to support this conclusion. The elongated, sometimes visibly conical morphology of ENLR (and NLR resolved with HST) points to a link with collimated radiation from the nucleus (eg. Unger et al., 1987; Mulchaey et al., 1996; Evans et al., 1993).

In the near-infrared, the origin of strong, extended [Fe II] emission in Seyfert galaxies has been somewhat controversial, with evidence for both shock excitation and photoionization occurring. It has been argued (see Greenhouse et al., 1991) that shocks must be important because iron in our interstellar medium is mainly condensed onto grains—shocks can destroy these via sputtering (Savage & Sembach, 1996). However, recent studies (Simpson et al., 1996; Mouri et al., 2000) indicate

that the ratio [Fe II]/[O I] in both Seyferts and starbursts is consistent with the usual depletion levels of iron and oxygen from the gaseous phase, and that certain other signs of grain destruction are absent. Even if metal enhancement is not the dominant factor, it is likely that shocks due to outflows along the radio jets of AGN do make some contribution to emission from partially-ionized zones (Simpson et al., 1996; Morse et al., 1996). In particular, Knop et al. (1996) find that whilst [Fe II] in NGC 4151 seems kinematically compatible with the ENLR, the 1.257 μm feature is broader than the nearby Paβ line

NGC 4151 is both a natural candidate in which to study [Fe II] excitation and an important test case for understanding the structure of active galaxies. It is one of the brightest nearby Seyfert galaxies (V $\simeq$ 11.5, $cz = 995\,\mathrm{km\,s^{-1}}$), yet eludes firm classification in the framework of unified AGN models (eg. Antonucci, 1993). The highly extended narrow line emission and radio jet (eg. Perez et al., 1989; Pedlar et al., 1993) are characteristic of a Seyfert 2 nucleus, whose collimation axis is almost perpendicular to the line of sight. However, the object also exhibits the strong broad lines and variability of a Seyfert 1 (eg. Maoz et al., 1991), suggesting quite a different orientation. NGC 4151 has therefore been the subject of many past projects involving slit spectroscopy and narrow-band imaging.

An unusual feature which facilitates the present investigation is the pronounced misalignment between the ENLR at position angle $\sim$ 50° (eg. Perez et al., 1989; Evans et al., 1993; Kaiser et al., 2000), and radio jet at 77° (eg. Pedlar et al., 1993; Mundell et al., 1995). Whatever its origin, this provides an opportunity to associate [Fe II] directly with the optical ENLR and/or radio jet, by tracing the emission along both axes. If a clear separation can be made, it will be possible to determine the relative influence of ionization by collimated X-rays and shocking by outflowing radio plasma. Such a task is inherently suited to integral field spectroscopy, which can form a complete, homogeneous picture without prior assumptions, slit alignment problems or the waveband limitations of a filter.

5.1.2   Instrumentation

The SMIRFS-IFU is a 72-fibre imaging feed for the CGS4 spectrograph at UKIRT. It was designed and built by the Astronomical Instrumentation Group at Durham University, as a prototype for TEIFU on the William Herschel Telescope (chapter 6)

and the GMOS-IFU on Gemini (Allington-Smith et al., 2000). Operating in the J and H bands with uncooled fused silica fibres, it was the first such system capable of imaging spectroscopy in the near-infrared. The instrument is discussed in detail in Haynes et al. (1998a, 1999) and Lee (1998). Further information about CGS4 can be found in Wright et al. (1993) and on the UKIRT Web pages.

SMIRFS is designed to intercept the telescope beam without modification to existing instrumentation, but consequently requires removal of the CGS4 calibration unit. In place of this module, the SMIRFS slit projection unit feeds light from the fibre bundle into the long-slit, circumventing the usual beam path; the IFU input is connected to another port of the UKIRT instrument support structure. As a result of this arrangement, any reference observations must be taken using the main telescope aperture.

At the image plane, a hexagonal microlens array provides optimal coupling to the slow telescope beam and avoids dead space between fibres (see Haynes et al. 1999 for more information). The field of view is 6 × 4 arcsec, with 0.62 arcsec spacing between adjacent fibre centres. A second set of microlenses at the output restore the correct focal ratio for CGS4. When using the long camera, each fibre projects a spot whose FWHM is ~1.8 pixels at the detector, giving a comparable spectral resolution to the standard 2 pixel wide slit. These spots are spaced two rows apart, placing the instrument on the borderline between the 'maximally packed' and 'resolved peaks' categories of chapter 3, although throughput variations obscure the regular pattern in practice. Since there is only one field, sky subtraction is performed by nodding the telescope off source.

The capability of the SMIRFS-IFU was previously demonstrated in a June 1997 commissioning run. Its throughput relative to CGS4 alone was measured to be ~ 50%; since then, minor adjustments have been made to improve performance. Despite difficulties due to flexure between the slit and detector, the instrument was successfully calibrated and used to image the circumnuclear starburst of NCG 7469 in [Fe II] and Paβ emission.

Table 5.1: Details of NGC 4151 observations, 15–17 February 1998

Instrument	CGS4 + SMIRFS-IFU	IFU field	$6 \times 4''$
Camera	300mm	Spatial FWHM	$\sim 1.2 \times 1.1''$
Grating	150 lines/mm	Spatial sampling	$0.62 \times 0.54''$
Diffraction order	3	Detector	256×256 InSb
Filter	B1	Gain	$6\,e^-$
Spectral range	1.25–1.31μm	Read noise	$23\,e^-$
Spectral FWHM	~ 4.2Å ($100\,\mathrm{km\,s^{-1}}$)	Dark current	$\sim 1\,e^-/s$ ($\sigma = 6\,e^-/s$)?
Spectral sampling	2.1Å	Single exposure	180s $\times$ 2 detector pos
		Central exp.	48 minutes on source

5.2   Observations

5.2.1   Science programme

NGC 4151 was observed with the SMIRFS-IFU in February 1998, as part of a sample of Seyfert galaxies, under a collaboration between researchers from Durham University and the Dominion Astrophysical Observatory. Details of the measurement parameters are presented in table 5.1. The selected region of the J-band includes both [Fe II] and Paβ, at rest wavelengths of 1.2567μm and 1.2818μm respectively. Previous studies have sometimes used [Fe II] at 1.6435μm and Brγ at 2.1655μm; such measurements are equivalent to ours, because the same upper energy level of Fe$^+$ is involved. Since, however, the longer-wavelength lines are much further apart, they cannot be observed together at high dispersion and their ratio is more sensitive to differential extinction along the line of sight. The J-band [Fe II] feature is also brighter, and the uncooled SMIRFS better suited to working in that regime.

The spatial resolution quoted in table 5.1 is for the final, combined data set; it was estimated by producing images of the unresolved [S IX] line and Paβ broad line. These FWHM values correspond to the Nyquist scale, tending to confirm that the seeing disc was undersampled—the implications are discussed briefly in §5.3.

The long axis of the IFU was orientated at 50° east of north, along the optical ENLR of the galaxy. The field of view, however, only covers the inner NLR at $< 4''$ (Winge et al., 1999), which is orientated less favourably at $\sim 60°$ (Kaiser et al., 2000). With arcsecond resolution, this gives ~ 1 FWHM separation from the radio axis. Hence a mosaic was made, of several telescope pointings around the nucleus.

Target centring was performed with the aid of a dedicated program, which derives a rough map of the illuminated input elements from the data. For the peripheral mosaic positions, it was necessary to rely on specified telescope offsets from the centre. Each position was observed in turn, following the repeated sequence object–sky–object; this maximizes the time on source without giving unequal object and sky exposures at individual pointings. The 180s integration time was just long enough for the background to dominate the read noise and dark current. The whole mosaic was repeated until the full exposure was reached.

In the absence of the calibration unit, sky emission lines were used as a wavelength reference. Some arc lamp observations were also taken through the telescope. The IFU was flat-fielded using the illuminated dome and the twilight sky, whilst a detector flat was made with the calibration unit, prior to the installation of SMIRFS. The bright F and A stars BS4572 and HD105601 were observed in order to calibrate the spectral response of the instrument and measure sky absorption features.

5.2.2   IFU set-up and characteristics

The IFU was aligned at the focal plane by projecting a laser beam back through the fibres onto the UKIRT secondary mirror. The CGS4 slit was set parallel to the row of fibre outputs, before adjusting the optics of the slit projection unit so that fibres were imaged two pixels apart in the spectrograph. The two pixel magnification was verified by illuminating the IFU with a mask in place, generating an image of the fibre slit at the detector with a gap for every input row. Since the mapping of the input to the slit reverses every row, to ensure that elements are adjacent both at the input and output (fig. 5.1), the distance between illuminated band centres at the detector alternates about the row separation. By fitting a straight line to the centroids, the magnification was measured as 2.002 ± 0.005 pixels/fibre.

The J-band throughput of the IFU was measured by comparing observations of the standard stars HD84800 and BS4069, taken with the CGS4 slit and the SMIRFS-IFU respectively. Examination of the PSF indicates that truncation of the seeing disc by the $4''$ slit did not cause significant light loss. The difference in airmass was small, giving an expected error due to atmospheric extinction of $< 1\%$. It was found that the throughput is 50%, consistent with the figure of 49% from the 1997 commissioning run (Lee, 1998).

Figure 5.1: Sketch of the SMIRFS-IFU input with a mask in place for determining the output magnification. The line with an arrow indicates the fibre ordering at the output slit. The dead fibre is shaded.

Variations in the fibre transmissions were quantified by extracting a flat-field spectrum for each element (see §5.3.2) and integrating in wavelength (fig. 5.2). The values therefore include crosstalk from neighbouring fibres. For Gaussian profiles with FWHM 1.8 pixels, the contribution from either neighbour is $\sim 10\%$ of the total. Ignoring further smoothing due to interpolation and any correlation between throughputs, the measured and true RMS values are thus related as follows:

$$\sigma_{\mathrm{data}} = \sqrt{(0.8\sigma_{\mathrm{IFU}})^2 + 2 \times (0.1\sigma_{\mathrm{IFU}})^2} = 0.8\sigma_{\mathrm{IFU}}. \tag{5.1}$$

The measured RMS, σ_{data}, indicates the signal-to-noise fluctuation; it is 8% overall or 6% excluding the dead fibre and some vignetted elements. The corresponding values of σ_{IFU}, which characterizes the fibre bundle, are $\sim 10\%$ and $\sim 7\%$ respectively. It is concluded that whilst careful flat fielding is important, inherent differences in the signal-to-noise ratio are reasonably modest.

The point spread function does not vary substantially between fibres. Whilst a previous estimate (Haynes et al., 1999) indicates that some profiles are up to 20% broader than others, the PSF is only known to $\sim 10\%$. Any variations at this level will contribute slightly to the Nyquist-scale noise in line width maps. The effect on line flux noise is negligible, since the broadening is a fraction of a pixel. Any

Figure 5.2: Throughput variation along the slit for the SMIRFS-IFU.

inhomogeneities introduced by the fibres can be smoothed out to some extent by combining observations with slight spatial offsets.

5.3 Data reduction and analysis

5.3.1 Overview

Data were reduced in IRAF, using several purpose-written scripts and programs alongside standard tasks. As discussed in chapter 3, the best method of dealing with fibre spectra depends largely on their separation and FWHM at the detector. For the SMIRFS-IFU, spectra are much more densely packed than in a multi-object design, precluding the use of existing extraction software. New tasks were also needed for reconstructing and working with the observed x–y–λ volumes in the form of datacubes, and for bad pixel correction. Because the slit maps to the field of view in a complex way, flexure in the spectrograph cannot be compensated by

recentring the target (this would be true anyway for off-centre mosaic positions). Observations must therefore be co-added after conversion to datacubes, rather than in their raw format. This requires effective bad pixel correction for integration pairs.

Chapter 3 presented an alternative reduction method for fibre-IFU data, inspired by the very dense packing of spectra, one per row, in TEIFU (see chapter 6). With the long CGS4 camera, the raw format of SMIRFS-IFU data is intermediate between that of TEIFU and multi-object designs. Since fibres are imaged with a FWHM of 1.8 pixels, the output pattern is almost critically sampled; interpolating in pixel co-ordinates does not cause the high level of distortion obtained with a width of ~ 1 pixel. Moreover, the two pixel fibre separation is great enough that moderate smoothing by the interpolant will be small in relation to the image resolution. At the same time, flexure-related errors and undersampling at the IFU input rule out much gain in precision through TEIFU-style reduction. Hence, as for most fibre-based instruments, the reduction is based on extracting 1D spectra.

Finding absolute fibre positions on the detector after flexure is complicated by vignetting at the slit ends, which prevents direct location of the first and last elements. With an already limited field of view, it was decided to rely on the single broken fibre for measuring offsets between frames, rather than masking off row ends at the input to create regular gaps. This is usually adequate, if somewhat awkward and instrument-specific.

With $0.62''$ elements at the IFU input, undersampling is potentially an important issue. Inspection of various observations (with and without SMIRFS) indicates that the seeing during the February 1998 run was $\sim 0.9''$, including integration over the input microlenses; this figure is relatively favourable for such coarse sampling. A very basic simulation of spatial frequency propagation through the instrument shows that, depending on the alignment of fibres, pixels and image features, aliased frequencies in the data can have amplitudes up to 10% of the mean. However, the Fourier spectrum is only affected significantly above $\sim 80\%$ of the Nyquist frequency, so moderate smoothing during extraction and reconstruction can easily eliminate much of the contamination. In principle, offset observations can be combined with phase shifts such that aliasing is removed (eg. Lauer, 1999), but in this case the main effect would be to amplify noise because the offsets are not known or controlled precisely (and the seeing varies). Nevertheless, simply co-adding datacubes after

alignment on a common grid does suppress residual alias components, which add out of phase like noise. At the end of the reduction process, aliasing artifacts are not a major source of error.

5.3.2 Calibration and extraction

Each observation consists of two integrations, offset by one detector pixel in the spectral direction to account for bad pixels. These raw frames are largely obscured by non-uniform dark current. The first step was to subtract a sky pair from an object pair, removing dark current in the process. However, this double use means that sky frames cannot be smoothed spatially to reduce their noise level before subtraction. The result was then divided by a detector flat-field image, obtained with the calibration unit before the IFU was installed.

A less accurate dark frame, taken earlier in the night, was also subtracted from the object separately. This reveals the fibre throughput pattern so that the dead fibre can be located. Ideally, this would be done using a sky minus dark image, but flexure of a fraction of a pixel can occur even between consecutive pointings. This is important for image reconstruction, but less so for sky subtraction—assuming Gaussian fibre profiles and an offset of 0.3 pixels, the sky residual due to throughput variations is likely to be around 3%, falling below the noise level.

A program was written to combine the pairs of offset, sky subtracted frames with bad pixel removal. This makes a crude noise estimate based on the RMS residuals between the two images; any input pixel which differs from the median of its neighbours by more than a specified number of standard deviations is excluded. The procedure can sometime truncate narrow arc lamp emission lines, whose gradients can be much greater than the noise level, but is safe for the galaxy spectra. Any obvious remaining blemishes were removed manually.

Spectra were straightened with respect to detector rows by applying a small pre-determined shift to each column. A curvature model was constructed by measuring the centroid of a suitable image at each wavelength and fitting a low-order polynomial to the values. Interpolation was performed using a cubic spline, which is good for dealing with slightly undersampled data. Inspection of corrected frames verifies that the curvature is practically constant along the slit and hence independent of

flexure. The maximum difference between columns is 0.7 rows.

Absolute fibre positions were found by summing the straightened object minus dark image in wavelength and fitting the dead fibre position using standard IRAF routines. The corresponding sky subtracted image was resampled so that fibres fell at the mid-points between odd pairs of rows. These pairs were combined with the *blkavg* task, producing a stack of extracted fibre spectra (boundary rows were discarded). Sometimes spatial gradients make the dead fibre difficult to identify, but the procedure can be repeated if a reconstruction appears distorted.

The stack of extracted spectra was divided by a fibre flat, created in the same way but averaged over its wavelength range. Sky emission lines in the object minus dark image were used for wavelength calibration—these allow correction of grating offsets between observations and also provide an absolute calibration. The spectra from every observation were interpolated onto a suitable standard grid, truncating the end few pixels which fell outside the common range. The error in determining the relative offsets is around 0.2Å. The error in the linear wavelength fit is 0.7Å, which is more than adequate for the present work. Separate solutions were not determined for the spectra within each observation; variation along the slit (±0.3Å) is smaller than the final random velocity errors in the mosaic and much smaller than the true gradient in the galaxy.

Finally, the stack was divided by a continuum-normalized stellar spectrum, to remove telluric absorption features. A Paβ line in one of the stars was interpolated over. Unfortunately, the mosaicing process makes it difficult to re-observe the standard stars frequently, so there was sometimes a non-negligible difference in airmass from the galaxy. Although the main absorption band at 1.269μm has been removed effectively in the final data, there may be some absorption residuals at around 2% of the continuum level.

5.3.3 Construction of datacubes and mosaicing

Sets of extracted spectra were interpolated spatially at each wavelength step, onto a square grid of 0.15″ pixels. This was done using the NAG library's implementation of the method of Renka and Cline (see Renka, 1983), chosen mainly for convenience. The algorithm fits a smooth, continuous surface of quadratic pieces, passing through

the original samples. Partial derivatives are determined by distance-weighted least-squares fitting to nearby points. Similar procedures are available freely from Netlib and Starlink on the Internet. As long as the points remain in their original grid, one could also resample each dimension in turn using a suitable 1D interpolant, although the present method is more general, placing the samples according to a list. The new grid is bounded spatially by extremal fibre centres and the images are stacked in a datacube with the same wavelength elements as the input spectra.

The constituent datacubes of each mosaic cycle were combined with appropriate offsets using the program *mosaic*, which has been included as a prototype task in the *imspec* package (chapter 6). This calculates the smallest common grid from a list of pointings and resamples the images at each wavelength using a tapered sinc function. The precise interpolation method is not important, since the datacube grids are much finer than the resolution of the data. A mask image is also produced, recording the number of input pixels contributing at each point.

Unfortunately, telescope pointing errors of a fraction of a fibre are evident. This is partly because a dichroic mirror is used to guide on the target at visible wavelengths whilst observing in the infrared; the two images drift apart as atmospheric dispersion changes with the airmass along the line of sight. Modelling this effect may systematically reduce errors, but was outside the scope of the current project. There may consequently be some image distortion at the points of overlap; mosaicing nevertheless reveals any wider-scale extension of spectral features within the noise limit of the data. For the NGC 4151 observations, emission turns out to be confined mainly to the central field; three mosaic positions with appreciable signal-to-noise, offset diagonally east and west, were finally used for mapping.

After constructing each separate mosaic, the resulting datacubes were combined using *comcubes*, also included in *imspec*. This is very similar to *mosaic*, but calculates the offsets between datacubes by locating the centroids of the intensity peaks after integrating in wavelength. The individual overlap masks are also combined to show the total number of data points contributing at each output pixel. The final datacube is divided by the mask, to scale the pixels correctly.

The spectral response of the instrument was determined by fitting piecewise polynomials to the continua of the standard star observations, dividing the results by black-body curves appropriate for their spectral types. This was achieved using

5.4   Results and comparison with other data

A wavelength-integrated image of NGC 4151 is shown in fig. 5.3(a). Whilst some
horizontal broadening is caused by the IFU, because spectra overlap, there is also
a real NW–SE extension of the continuum, following the stellar distribution seen in
optical images. Fig. 5.3(b) is an image of the Paβ broad line (including the narrow
line at ∼ 6% of the total), which gives an indication of the spatial resolution.

Figure 5.3:   (a) 1.25–1.3μm image of the nucleus of NGC 4151.
(b) Paschen β broad line image indicating resolution.

A central spectrum is presented in fig. 5.4; this includes the broad and narrow
Paβ components, [Fe II] and [S IX]. There may also be contributions from a weaker
1.2703μm [Fe II] feature, superimposed on the blue wing of the Paβ broad line, and
He I at 1.2791μm, to the left of the Paβ narrow line. The small bump to the right
of [S IX] is probably due to imperfect removal of atmospheric absorption features.

The combined [Fe II] and Paβ velocity map (fig. 5.5) traces the rotation of gas
about an axis projected at a position angle of ∼ 110°. It is broadly consistent both
with previous slit measurements of the same features (see Knop et al., 1996) and
with optical spectroscopy of the NLR (eg. Mediavilla & Arribas, 1995; Winge et al.,
1999). Estimates of the line widths are complicated by the Paβ broad line shape, lack

Figure 5.4: Central spectrum of NGC 4151.

Figure 5.5: Combined [Fe II] and Paβ velocity map.

of clear continuum and modest signal-to-noise ratio off centre. The line profiles have not been modelled with multiple components, but their overall deconvolved widths are close to the values found by Knop et al., who note that the width of [Fe II] relative to Paβ suggests a contribution from shock excitation. At the centre, the FWHM values are approximately $400\,\mathrm{km\,s^{-1}}$ for [Fe II] and $250\mathrm{km\,s^{-1}}$ for Paβ. Both lines are broader in the brightest parts of the surrounding extended emission (see below), reaching $\sim 600\,\mathrm{km\,s^{-1}}$ and $\sim 400\,\mathrm{km\,s^{-1}}$ respectively. Additional broadening in the direction of the radio jet is not evident.
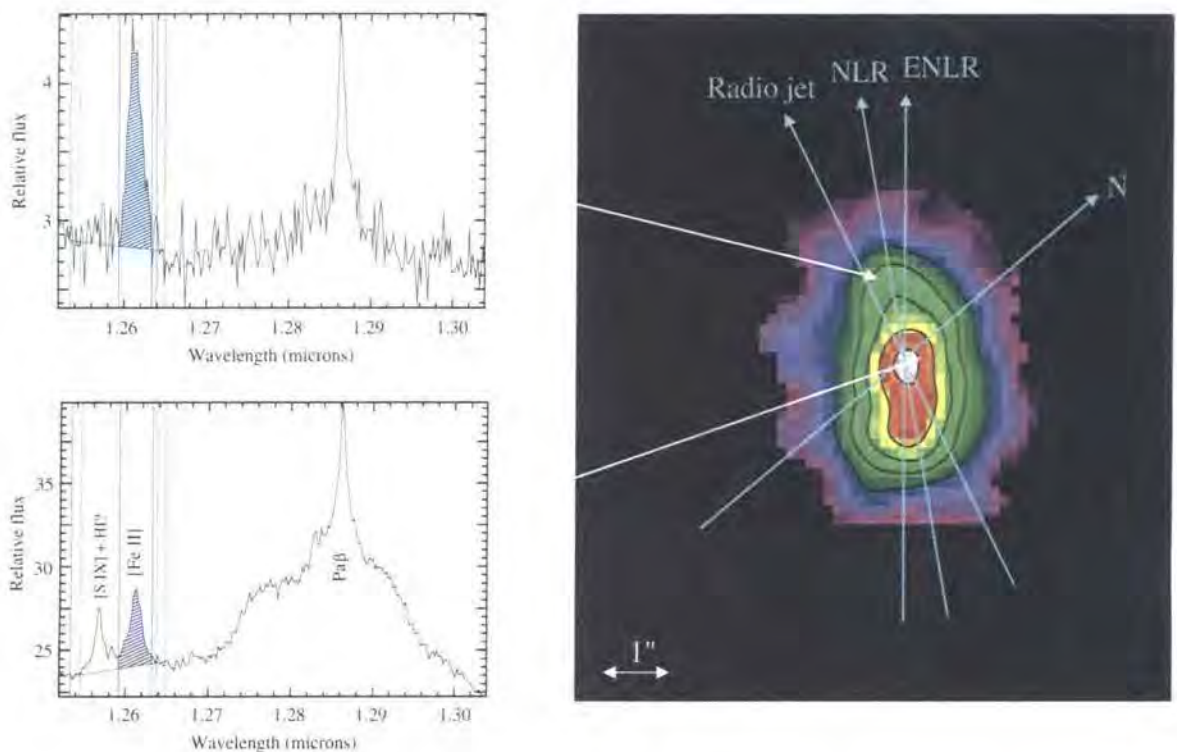

Figure 5.6: Map of [Fe II] flux, with measurements of sample spectra.

The spatial distribution of [Fe II] is mapped in fig. 5.6, alongside examples of individual spectral measurements. The image profile is extended along both axes of the IFU, with an intrinsic RMS co-ordinate of 0.3–0.4″ horizontally and 0.9–1.0″ vertically (ie. FWHM of 0.7–1.0″ and 2.1–2.4″ for a Gaussian distribution). The narrowest point is at the intensity peak, as might be expected for a biconical shape. Although these scales are comparable to the image resolution, it is clear that the [Fe II] is aligned closer to the optical line emission than to the radio axis. The best fit to the slope of the horizontal centroid has a PA of $\sim 54°$, part way between the

NLR and ENLR axes. For the data to fit the radio axis equally well, either end of
the outside ($\sim 25\%$) contour would have to be displaced by $\sim 0.8''$.

Fig. 5.7 compares the [Fe II] map directly with a narrow-band HST image of
[O III] line emission (from Kaiser et al., 2000) and a contour map of the radio jet
at 8 MHz (originally from Mundell et al., 1995). It can be seen that the region
over which [Fe II] is detected in the SMIRFS data corresponds to the extent of
the narrow line region. If there were a strong correspondence with the radio jet,
one might expect a more elongated distribution, resolving some knot structure at
77°, where the outflow interacts with gas. Although detection of the larger-scale
ENLR would have provided a better separation, the spatial association of [Fe II]
with photoionized gas is more convincing than with shocking by radio plasma. The
F502N filter used for imaging [O III] covers almost all of the NLR velocity range
(Hutchings et al., 1999; Kaiser et al., 2000), so represents the total emission well. It
is perhaps interesting to note that [Fe II] appears better correlated with the positions
of low velocity dispersion clouds, plotted in Kaiser et al., than the NLR as a whole.


Figure 5.7: (left) HST image of NGC 4151 in [O III] emission, overlaid
with radio contours, from Kaiser et al. (2000). (right) SMIRFS-IFU
image of [Fe II] on the same scale. Dotted lines mark the [Fe II] axis.

The Paβ narrow line is slightly extended in both directions, and has a faint wing out to $\sim 1.5''$ South-West of the centre, where the NLR is brightest. Fig. 5.8 shows both the Paβ flux and [Fe II]/Paβ ratio. Since Paβ falls off fairly rapidly from the centre, the ratio is highly sensitive to noise artifacts. The datacube was therefore smoothed with a 0.9'' wide Gaussian kernel before producing the excitation map, preserving minimal resolution of spatial structure. The line ratio is ~ 0.7 at the centre and ≥ 2 along the radio jet at the edges. A dip at the centre is to be expected for photoionization, since the strong incident flux creates large fully ionized zones, favouring Paβ over [Fe II] (Mouri et al., 2000). The high areas could reflect an enhancement of [Fe II] emission by shocks, but the line images are more indicative of a drop in Paβ at those points, with respect to [O III], than of strong [Fe II]. To the South-West, a few highly-ionized clouds may just happen to give stronger Paβ away from the radio axis, whilst to the North-East, the NLR is angled slightly more towards the radio jet anyway.
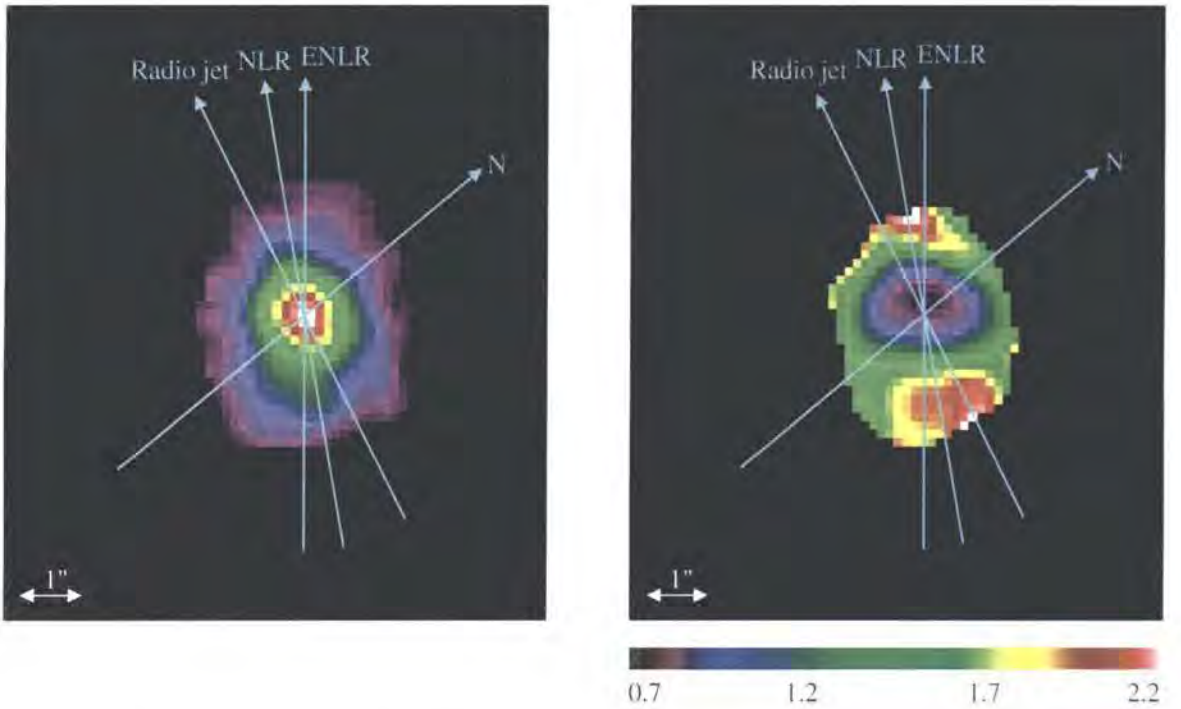

Figure 5.8: (left) Paβ emission (right) Ratio of [Fe II]/Paβ.

5.5 Conclusions

The distribution of [Fe II] emission at the centre of NGC 4151 has been mapped in two dimensions, using the innovative SMIRFS-IFU. Although the line is only detected over a few spatial resolution elements, it is evident that its image profile is better aligned with the visible narrow line region than with the radio jet. The study therefore suggests that [Fe II] arises primarily through photoionization of gas, by collimated X-rays from the Seyfert nucleus. The mean velocity field is consistent with this interpretation, indicating that [Fe II], $Pa\beta$ and the NLR all have similar bulk kinematics. Likewise, the variation in [Fe II]/$Pa\beta$ is compatible with photoionization dominating. As suggested by previous authors, shock excitation by outflowing radio plasma may make a secondary contribution to the emission. However, the evidence for this comes mainly from line widths, whereas tentative measurements suggest, if anything, that [Fe II] is broader along the NLR than the radio jet.

Despite the small size and prototypical nature of the IFU, these results are very promising, demonstrating the utility of integral field spectroscopy in investigating the properties of active galaxies. Without a two dimensional field, it would have been very difficult to determine the spatial orientation of the line emission. Narrowband imaging can also fulfil this requirement, given appropriate filters, but follow-up spectroscopy is then required to form a complete picture. The experience gained with the SMIRFS-IFU and associated data analysis has paved the way for subsequent projects, both technical and observational.

Reduction of the IFU data has highlighted a few important observational issues. In particular, location of spectra on the detector is difficult, and it is recommended that more than one reference gap between fibres should be available in future work. With a number of sources of error and inhomogeneity, the availability of multiple observations at each position has been very valuable.

Whilst the results presented here support the argument for [Fe II] originating in photoionized gas, there is no doubt that observations at high spatial resolution, as well as of the wider ENLR, will provide further insight into the nature of NGC 4151 and other Seyfert galaxies. With a bright, point-like nucleus, NGC 4151 is an ideal target for J-band observations with adaptive optics. This would provide a useful counterpart to the Hubble Space Telescope observations of [O III] emission, which

have revealed the cloud structure and detailed kinematics of the NLR. In conjunction with recent models of photoionized and shock excited gas, comparison of optical and infrared lines in individual clouds would provide important information about the physical conditions and kinematic structure of the NLR. The misalignment with the radio jet will again help determine its influence on the ambient gas.

Chapter 6

TEIFU commissioning

This chapter briefly considers the performance of the Thousand Element Integral Field Unit, as measured during commissioning in 1999. The total system throughput is determined to be 2.1% at 5000Å and 3.6% at 7000Å. These values are affected by the poor reflectivity of the telescope's Nasmyth derotator during the June 1999 observations, but are compatible with a laboratory measurement of 60% for the IFU alone. The typical variation between fibres is found to be $\sim 8\%$ of the mean. Direct comparison with INTEGRAL and Autofib confirms that the TEIFU bundle has a significantly greater throughput than the bare or individually connected fibres of the other systems. Finally, a reconstructed test observation of the radio galaxy 3C 327 is shown, demonstrating the functionality of TEIFU and the data reduction software.

6.1   Introduction

The Thousand Element Integral Field Unit (TEIFU) provides a high spatial resolution IFS capability for the William Herschel Telescope (WHT). The device was built by the Astronomical Instrumentation Group at Durham University, to complement the ELECTRA/NAOMI adaptive optics systems and to develop techniques for constructing large optical fibre arrays (Haynes et al., 1998b; Murray et al., 2000; Benn et al., 2001). Two fibre bundles feed the WYFFOS spectrograph (Bingham et al., 1994) at one of the WHT's Nasmyth foci. The system is designed for observing at ~ 0.5–1μm. With the standard modular fore-optics in place, each bundle of 500 fibres corresponds to a spatial input field of $7 \times 4''$ with $0.25''$ sampling; provision has also been made for scales of $0.125''$ and $0.5''$. Each fibre bundle is coupled to the telescope and spectrograph via microlens arrays, for optimal efficiency. The posi-

tions of the two IFU fields within the full AO field are varied using a motorized x-y-z translation stage and the Nasmyth derotator. A special pick-off mirror reverses half of the AO field of view, allowing the separation of the IFU fields to be varied on the sky without changing their physical separation (fig. 6.1). A central aperture in the mirror allows light from the guide star to pass through to the optical wavefront sensor of the AO system. The fields are always equidistant from the guide star.


Figure 6.1: Reformatting the AO field of view with the TEIFU pick-off mirror; this allows the effective separation of the two IFU fields to be varied and their distances from the guide star to remain equal (from Haynes et al., 1998b).

6.2   Commissioning

TEIFU was commissioned with ELECTRA in June 1999, with one of the two input fields completed. The second fibre bundle was installed and the system tested again without adaptive optics in September 1999. In addition to set-up and calibration, a small number of science test observations were made on each run.

The June commissioning run included the following work: optical alignment; focus tests; acquisition tests and x-y-z stage calibration using single and binary stars; flat field, arc lamp and spectrophotometric standard star observations; observation of the radio galaxy 3C 327, with one hour on source. Relatively poor (1-2") seeing conditions were corrected to $\sim 0.9''$ by the adaptive optics system. A similar set of tasks was performed on the September run, with additional trials of overlapping the two fields, beam-switching and physically interchanging the bundles.

During commissioning, it became apparent that the x-y-z stage was occasionally skipping steps on one axis, introducing some uncertainty in positioning the fields; this increases acquisition overheads and was most likely responsible for some failed observations during the September run. If TEIFU is to be used more routinely in future, it will be necessary for the control system to have more reliable information as to the positioning of the x-y-z stage. Other areas for improvement are modification of the x-y-z stage user interface to work in celestial co-ordinates and writing the co-ordinates to image headers. A second observing constraint was that the reflectivity of the Nasmyth derotator had obviously been degraded whilst out of use for some time. It is believed that the derotator has since been serviced.

The IFU itself performed effectively, with reasonable fibre throughput variations and a moderate proportion ($\sim 2\%$) of dead elements. The quick-look reconstruction software, whilst not accurate enough to produce images for analysis, allowed the pointing to be checked efficiently. With a comprehensive set of calibration measurements and some test observations, the commissioning was generally successful.

6.3 Characterization

6.3.1 Fibre throughput variations

A cross-section of the fibre flat-field pattern for both bundles is shown in fig. 6.2. The RMS intensity variation over the illuminated regions (the two fields) is 16% of the mean at 7000Å. Much of this variation is due to vignetting at the input. The level of vignetting in field 1 was somewhat lower on the June commissioning run, prior to adjustments to the fore-optics to accommodate the second field.


Figure 6.2: Constant-wavelength cross-section of the TEIFU output pattern for flat-field illumination (Sept. 1999).

The intrinsic throughput variation between fibres can be estimated as in eq. 5.1. At 7000Å, the fibre pitch is ~ 0.93 pixels and the FWHM is ~ 1 pixel. For Gaussian profiles, the average relative contributions at a given pixel from the nearest fibres are 73%, 13% and 13% when the central fibre is aligned with the pixel. For a pixel aligned mid-way between two fibres, the relative contributions are 49%, 49%, 1% and 1%. These values translate into an average smoothing factor due to crosstalk, $\sigma_{IFU}/\sigma_{data}$, of ~ 1.38.

Intrinsic fibre throughput variations are best represented by the June flat-field data. The RMS in the five blocks with no dead fibres, ignoring vignetting at the very edges, is 6%. Thus the actual variation in fibre transmissions, excluding dead fibres, is 8% of the mean transmission.

Fig. 6.3 indicates the positions of broken or very low throughput fibres in the two fields. A number of fibres in the second bundle were damaged simultaneously during construction. Images can be reconstructed from the remaining points, with reduced accuracy or signal-to-noise at the dead fibres, or observations can be dithered in steps of ~ 1 pixel, to fill in the gaps (although dithering is not implemented in the software at present).

Field 1   Field 2

Figure 6.3: Positions of dead fibres in the two TEIFU fields. For the Sept. 1999 run, North is upwards and East is to the left.

6.3.2   Absolute system throughput

The throughput of the atmosphere–detector system was measured for field 1 using the standard star HZ 44 on the June commissioning run. Oke (1990) gives the star's brightness in AB magnitudes as a function of wavelength, where

$$AB = -2.5 \log f_\nu - 48.60 \qquad (6.1)$$

and f_ν is the flux in erg s^{-1} cm^{-2} Hz^{-1}. From a few lines of working, these units can be converted to photons s^{-1} m^{-2} Å^{-1} using

$$f_\lambda(\lambda) = \frac{1}{h\lambda} 10^{-0.4(AB(\lambda)+56.1)}, \tag{6.2}$$

where h is in S.I. units and λ is in Ångstroms.

To find the proportion of incident photons detected by the system, the bias subtracted and wavelength calibrated detector image of HZ 44 was summed over two bands $\sim$ 50Å wide, centred at 5000 and 7000 Å. The total counts were converted to electrons using the gain listed in the image header and divided by the exposure time, telescope area and waveband width, to get the detected flux density.

For the observation used, the standard star was approximately centred in the IFU field, where the throughput is relatively uniform. From the image profile, it is estimated that $\sim$ 95% of the light in the telescope PSF was captured by the IFU. Based on a flat-field observation, the scattered light level between output blocks is $\sim$ 4% of the mean illumination within the blocks; neither scattered light nor the sky background are greater than the noise level in the 100 s exposure, and are not important in the calculation. The spectrograph was set up with the R316R grating.

The measured system throughputs for the June run were 2.1% at 5000 Å and 3.6% at 7000Å. The corresponding sensitivities for 1 e$^-$ s^{-1} Å^{-1} are 16.2 and 16.5 AB magnitudes, respectively. From the September run, the second field has $\sim$96% of the throughput of field 1. These figures are consistent with the predicted throughput of 60% for the IFU alone, which has since been confirmed in laboratory tests (Murray et al., 2000). The relatively low total efficiency of the atmosphere–detector system is due to the large number of elements in the optical train and non-optimal performance of a few components. The complete system consists of the atmosphere, three telescope mirrors, a three-surface reflective image derotator (estimated throughput 60%), the ELECTRA adaptive optics system (estimated throughput 50%), pick-off mirror, TEIFU fore-optics, TEIFU bundle and WYFFOS spectrograph. Recoating the derotator and upgrading the spectrograph with the proposed new camera should produce a substantial gain in sensitivity. In addition, the new (two-field) TEIFU fore-optics are anti-reflection coated.

6.3.3   Comparison with INTEGRAL and Autofib

Using observations from the Isaac Newton Group's data archive[1], a throughput comparison was made between TEIFU and the other WYFFOS fibre feeds, namely INTEGRAL (Arribas et al., 1998a) and Autofib-2 (Parry et al., 1994; Lewis et al., 1997). INTEGRAL is a bare-fibre IFU, also at the Nasmyth platform, whilst Autofib-2 is a robotic multi-object fibre positioner at the prime focus. The same spectrophotometric standard star used for TEIFU, HZ 44, was observed with INTE-GRAL on 19 July 1997 and Autofib-2 on 10–11 May 1996. The relevant observations were retrieved from the data archive and processed as for TEIFU. The spectral ranges of the datasets overlap in the region of 5000Å.

For meaningful comparison between the instruments, differences in throughput due to the spectrograph set-up must be estimated and removed. For the standard star observations, TEIFU was operated with the R316R reflection grating and a single order-sorting filter, whereas INTEGRAL and Autofib-2 were used with the R1200B grating and both the BG39 and GG495 filters. The header information as to exactly which filter was used with TEIFU apppears to be incorrect, but the latter two are the only possibilities, both of which have 85% transmission at 5000Å. Likewise, the R316R and R1200B gratings are equally (65%) efficient at 5000Å. The spectrograms from each of the three fibre feeds were summed over the same 50Å waveband, to the nearest pixel, before dividing by the precise bandwidth and adjusting for the differing exposure times and read-out gains. The throughput of the Autofib-2 fibre used to observe HZ 44, relative to the average for the instrument's 'large fibres', has not been established because the fibre number was not available in the header; thus the comparison with TEIFU is only approximate. The INTEGRAL observations were taken using the SB1 bundle.

Measurements of atmospheric extinction are not available for all three runs, but historical data from the Carlsberg Meridian Telescope suggest that differences are likely to be of the order of a tenth of a magnitude. The TEIFU and Autofib observations were taken close to the zenith, whilst the INTEGRAL data were taken through 1.2 airmasses. Observing logs indicate that the archive data were taken during clear weather.

---

[1]http://archive.ast.cam.ac.uk/ingarch

In addition to run-specific set-up differences, the three fibre instruments receive light via quite different optical trains. Autofib-2 feeds WYFFOS directly from the prime focus. For INTEGRAL, the beam path also includes the secondary and Nasmyth flat mirrors (the IFU rotates, however, so does not require an external image derotator). For TEIFU, the Nasmyth derotator, adaptive optics system, pick-off mirror and magnifying fore-optics introduce significant additional losses—none of these components are dispensable, but it is clear that performance could be improved, at least for the derotator. From an observational point of view, these factors all form part of the instrumental sensitivity, but from a technical viewpoint, it is also useful to compare the relative performances of the fibre systems alone.

The estimated total throughput with TEIFU on the June 1999 run, relative to that with INTEGRAL in July 1997, is $\sim$ 70%. Assuming 60% throughput for the derotator and 50% for ELECTRA, this means that the TEIFU fibre system (including fore-optics) transmits $\sim 2\times$ as much light as INTEGRAL. The latter result seems to confirm that a substantial gain is made by coupling fibres to the telescope and spectrograph using microlenses, whilst the former highlights the importance of optimizing other components.

The throughput of the June 1999 TEIFU set-up is again $\sim$ 70% of that for Autofib-2 with the fibre which was tested. From the ING Web page[2], the secondary and Nasmyth flat mirrors typically have reflectivities of 0.88. Accounting for these extra surfaces, as well as ELECTRA and the derotator, TEIFU has a throughput almost $3\times$ that of the Autofib fibre. It should be noted that Autofib-2 has recently been commissioned with a 'small fibres' option which should improve its efficiency in reasonably good seeing conditions.

The total throughput figures for the three fibre systems do not give a complete sensitivity comparison, since they do not account for the sky backgrounds in different sized fibres. The aim of the comparison, however, is to test the performance of the fibre-lenselet IFU design against bare and individually connected fibres. The study appears to confirm that a significant gain is made, as expected. The derived 1 e^- s^{-1} Å^{-1} sensitivities for INTEGRAL and Autofib-2 are both 16.6 AB magnitudes at 5000Å, within 0.5 magnitudes of the values quoted in their manuals.

---

[2]http://www.ing.iac.es

6.4   Test observations

During the June commissioning run, TEIFU was used to observe the radio galaxy 3C 327 in the wavelength range 0.43–1μm, with one hour on source. The data have been reconstructed into an x–y–λ datacube as described in chapter 4. A wavelength-integrated image and central spectrum are shown in fig. 6.4. Although the seeing was poor, the reconstruction demonstrates the ability of TEIFU and the *imspec* IRAF package to produce reasonable images and spectra. An [O III] velocity map is shown in fig. 6.5; this was derived from the 3C 327 datacube by Dr. Alasdair Allan during testing of the prototype Starlink *datacube* package.


Figure 6.4: Reconstructed TEIFU observation of 3C 327: wavelength-integrated image and central spectrum.

Figure 6.5: Alisdair Allan's velocity map of 3C 327. derived from the TEIFU datacube using prototype Starlink software.

Chapter 7

Conclusions

Integral field spectroscopy is a rapidly developing field, which creates many exciting new opportunities for astronomical research. From the present time onwards, IFUs will be a standard facility at major observatories, allowing a wide variety of scientific problems to be pursued from a new perspective.

This thesis has presented two studies which demonstrate promising applications of IFS. First, in conjunction with optical slit spectroscopy, the 3D infrared image slicer has been used to test for the presence of a super-massive black hole in the giant early-type galaxy NGC 1316. It is found that the galaxy does not contain a black hole of mass $5 \times 10^9 M_{\odot}$, as predicted from bulge luminosity using the relationship of Kormendy & Richstone (1995). The approximate upper mass limit of $8 \times 10^8 M_{\odot}$ is, however, compatible with the recently discovered correlation between black hole mass and bulge velocity dispersion (Gebhardt et al., 2000; Ferrarese & Merritt, 2000). As an indirect indicator of mass, bulge luminosity is sensitive to the properties of the stellar population, which in this case have been disrupted by a recent merger. In future, this type of work will benefit greatly from the availability of adaptive optics in conjunction with IFS.

In the second observational project, Durham University's SMIRFS-IFU has been used to investigate the origin of [Fe II] emission in the well-studied Seyfert galaxy NGC 4151. By mapping [Fe II] over the central few arcseconds of the galaxy, it is shown that the emission is spatially associated with the optical narrow line region, rather than the radio jet. The velocity field is also found to be consistent with that of the narrow line region. It is therefore concluded that [Fe II] excitation results primarily from photoionization by X-rays from the active nucleus. It is clear that IFS will be particularly valuable for studying the complex structure of active galaxies

and that adaptive optics will once again be important for resolving structural detail.

As IFUs are commissioned as common-user instrumentation, there is growing recognition that much work is needed to provide full support for their complex data reduction requirements. This has led to intensive discussions amongst interested parties and, notably, the formation of an 'OPTICON 3D Spectroscopy Working Group' in Europe. A key part of the work towards this thesis has been the development of a new approach to fibre IFU data reduction, needed to allow maximal packing of spatial samples onto the spectrograph detector. In chapter 3, it has been shown that densely-packed fibre spectra may be treated in a similar way to image slicer data, but with spatially nonuniform sampling. The problem of reconstructing three dimensional images from irregular sample sets has been reviewed qualitatively, with reference to available literature and algorithms. These ideas are implemented as part of a data reduction package for Durham University's TEIFU, described in chapter 4, whose source code has been made available on compact disc.

Finally, the performance of TEIFU has been characterized and compared with other fibre feeds for the WYFFOS spectrograph. It is deduced that the total system throughput on the June 1999 run was 2.1% at 5000 Å and 3.6% at 7000Å, with a typical intrinsic variation between fibres of $\sim 8\%$ of the mean. This transmission was somewhat lower than expected, due to degradation of the Nasmyth derotator whilst out of use, but is consistent with a laboratory measurement of $\sim 60\%$ for the IFU alone. Comparison with INTEGRAL seems to confirm that the fibre-microlens design provides a gain of ~ 2 with respect to bare fibres. A reconstruction of the radio galaxy 3C 327 demonstrates the successful implementation of hardware and software in a maximally-packed fibre system.

Appendix A

CD contents

In order to satisfy the requirements of the University, the files on the accompanying compact disc are listed below, in the Unix 'ls -R' format:

```
.:
imspec/     TEIFUdata/

./imspec:
comcubes.par    ifms.x        imspec.hd      oldnumsub.x    t_ifuwcal.x
dctomult.par    ifsm.x        imspec.men     pda/           t_listoim.x
doc/            ifsp.x        imspec.par     prototype/     t_specmap.x
fitspec.x       ifuarc.par    Jtextlib.x     Readme         todo
gensub.x        ifucal.par    lib/           smil.x         view.cl
getmap.x        ifudef.par    listoim.par    specmap.par    view.par
ifca.x          ifuimrec.par  logo.gif       t_dctomult.x   x_imspec.x
ifdc.x          ifurec.par    mkpkg          t_ifuarc.x
ifes.x          ifuwcal.par   mosaic.par     t_ifucal.x
ifim.x          imspec.cl     numsub.x       t_ifudef.x
ifml.x          imspec.h      objects.x      t_ifurec.x

./imspec/doc:
comcubes.hlp    ifuarc.hlp    ifuimrec.hlp   listoim.hlp    specmap.hlp
dctomult.hlp    ifucal.hlp    ifurec.hlp     locspec.key    specmap.key
fitspec.key     ifudef.hlp    ifuwcal.hlp    mosaic.hlp

./imspec/lib:
helpdb.mip      root.hd       rootimspec.hd

./imspec/pda:
pda_dsqf.f      pda_idlctn.f  pda_idtang.f   pda_side2.f
pda_idbvip.f    pda_idpdrv.f  pda_idxchg.f   pda_spdt.f
pda_idcldp.f    pda_idptip.f  pda_side.f
```

```
./imspec/prototype:
comfsub.f        comsub.x        t_comcubes.x    t_mosaic.x


./TEIFUdata:
june_instdescr.def    juneflat.fits         sept_instdescr.def
junearcNeAr.fits      junestar.fits         septflat.fits
```

Acknowledgements

Thanks are due to a large number of people for a variety of reasons. In particular, Jeremy Allington-Smith, who has been my Ph. D. supervisor for most of the last four years, has provided an excellent research opportunity in a new and promising area. Robert Content has provided valuable help and advice in coming to terms with his 'exotic' instrument designs. Harald Kuntschner, Reynier Peletier and Chris Done have also given me valuable guidance.

My doctoral research in Durham has been funded by a studentship from the Particle Physics and Astronomy Research Council.

Many thanks indeed are due to the staff at the Gemini Observatory, who have waited with great patience and understanding for me to finish my thesis.

Graham Murray has kept me well informed of progress in the lab, not to mention a number of other weird and wonderful subjects. Roger Haynes and David Lee have also helped in this respect. I remain grateful to John Lucey, who introduced me to astronomy research with great enthusiasm.

I am sad to hear that Claude Shannon, whose work underpins the theory in chapter 3, has died during the writing of this thesis.

Somewhere between the lists of colleagues and friends come my office mates: Harold, Dequing, Tony, Tom and others, with whom I share the binding experience of moving from relative civilization into room 4 of the Physics Department. I am also grateful to Harold and Graham for their help in submitting my thesis.

An enormous thank you goes to my wife, Louisa, who has spent the last two years waiting for me to finish my Ph. D., so that she can continue with her own study in music. I look forward to spending some time with you again...

It is only after 26 years of support from my parents that I have finally come to submit my Ph. D. thesis and can begin my education. I can't really say much more than that, except in the dedication. As for Charles, you should probably have a mention too.

Life used to be much more varied thanks to a large number of people in Durham who all deserve individual mention. If I try to list everyone here at this time of morning, however, I will doubtless miss many people out. They know who they

Bibliography

Adam, G., Bacon, R., Courtes, G., Georgelin, Y., Monnet, G. & Pecontal, E. (1989), 'Observations of the Einstein Cross 2237+030 with the TIGER integral field spectrograph', A&A **208**, L15–L18.

Adorf, H. M. (1995), Interpolation of irregularly sampled data series—a survey, *in* 'ASP Conf. Ser. 77: Astronomical Data Analysis Software and Systems IV', Vol. 4, pp. 460+.

Allington-Smith, J. & Content, R. (1998), 'Sampling and background subtraction in fiber-lenslet integral field spectrographs', PASP **110**, 1216–1234.

Allington-Smith, J. R., Content, R., Dodsworth, G. N., Murray, G. J., Ren, D., Robertson, D. J., Turner, J. E. H. & Webster, J. (2000), 'Integral field spectroscopy with the GEMINI multi-object spectrographs', Proc. SPIE **4008**, 1172–1180.

Antonucci, R. (1993), 'Unified models for active galactic nuclei and quasars', ARA&A **31**, 473–521.

Arnaboldi, M., Freeman, K. C., Gerhard, O., Matthias, M., Kudritzki, R. P., Méndez, R. H., Capaccioli, M. & Ford, H. (1998), 'The stellar dynamics and mass of NGC 1316 using the radial velocities of planetary nebulae', ApJ **507**, 759–773.

Arribas, S., Carter, D., Cavaller, L., del Burgo, C., Edwards, R., Fuentes, F. J., Garcia, A. A., Herreros, J. M., Jones, L. R., Mediavilla, E., Pi, M., Pollacco, D., Rasilla, J. L., Rees, P. C. & Sosa, N. A. (1998*a*), 'INTEGRAL: a matrix optical fiber system for WYFFOS', Proc. SPIE **3355**, 821–827.

Arribas, S., Mediavilla, E. & Fuensalida, J. J. (1998*b*), 'A new technique for performing two-dimensional spectroscopy of objects of large dynamic range: The binary system HD 167605 A+B', ApJ **505**, L43–+.

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.', A&AS **113**, 347+.

Bacon, R., Copin, Y., Monnet, G., Miller, B. W., Allington-Smith, J. R., Bureau, M., Marcella Carollo, C., Davies, R. L., Emsellem, E., Kuntschner, H., Peletier, R. F., Verolme, E. K. & Tim de Zeeuw, P. (2001a), 'The SAURON project - I. The panoramic integral-field spectrograph', MNRAS **326**, 23–35.

Bacon, R., Emsellem, E., Combes, F., Copin, Y., Monnet, G. & Martin, P. (2001b), 'The M 31 double nucleus probed with OASIS and HST', A&A **371**, 409–428.

Bacon, R., Emsellem, E., Copin, Y. & Monnet, G. (2000), Supermassive black hole searches with 3-D spectroscopy, in 'ASP Conf. Ser. 195: Imaging the Universe in Three Dimensions', pp. 173+.

Barden, S. C. & Wade, R. A. (1988), DensePak and spectral imaging with fiber optics, in 'ASP Conf. Ser. 3: Fiber Optics in Astronomy', pp. 113–124.

Bender, R. (1990), 'Unraveling the kinematics of early-type galaxies - presentation of a new method and its application to NGC4621', A&A **229**, 441–451.

Bender, R., Saglia, R. P. & Gerhard, O. E. (1994), 'Line-of velocity distributions of elliptical galaxies', MNRAS **269**, 785+.

Benedetto, J. J. (1992), Irregular sampling and frames, in C. K. Chui, ed., 'Wavelets—A Tutorial in Theory and Applications', Academic Press Inc.

Benn, C., Longmore, A., Myers, R., Gregory, T. & Davenhall, C. (2001), 'NAOMI - Adaptive Optics at the WHT', ING Newsl. **4**, 21–22.

Beutler, F. J. (1966), 'Error-free recovery of signals from irregularly spaced samples', SIAM rev. **8**(3), 328–335.

Bingham, R. G., Gellatly, D. W., Jenkins, C. R. & Worswick, S. P. (1994), 'Fibre-fed spectrograph for the 4.2-metre William Herschel Telescope', Proc. SPIE **2198**, 56–64.

Bosma, A., Smith, R. M. & Wellington, K. J. (1985), 'Rotation and velocity dispersion in the stellar component of NGC 1316 (Fornax A)', MNRAS **212**, 301–307.

Bowen, I. S. (1938), 'The Image-Slicer a Device for Reducing Loss of Light at Slit of Stellar Spectrograph.', ApJ **88**, 113+.

Bracewell, R. N. (2000), *The Fourier Transform and Its Applications*, third edn, McGraw-Hill.

Bruzual, A. G. & Charlot, S. (1993), 'Spectral evolution of stellar populations using isochrone synthesis', ApJ **405**, 538–553.

Clayton, C. A. (1989), 'The implications of image scrambling and focal ratio degradation in fibre optics on the design of astronomical instrumentation', A&A **213**, 502–515.

Content, R. (1998), 'Advanced image slicers for integral field spectroscopy with UKIRT and GEMINI', Proc. SPIE **3354**, 187–200.

Courtes, G. (1982), An Integral Field Spectrograph (IFS) for large telescopes, *in* 'ASSL Vol. 92: IAU Colloq. 67: Instrumentation for Astronomy with Large Optical Telescopes', pp. 123+.

Cretton, N., de Zeeuw, P. T., van der Marel, R. P. & Rix, H. (1999), 'Axisymmetric three-integral models for galaxies', ApJS **124**, 383–401.

Davies, R. L., Kuntschner, H., Emsellem, E., Bacon, R., Bureau, M., Carollo, C. M., Copin, Y., Miller, B. W., Monnet, G., Peletier, R. F., Verolme, E. K. & de Zeeuw, P. T. (2001), 'Galaxy mapping with the SAURON integral-field spectrograph: The star formation history of NGC 4365', ApJ **548**, L33–L36.

de Jong, R. S. & Davies, R. L. (1997), 'The shapes and ages of elliptical galaxies', MNRAS **285**, L1–L4.

Deeming, T. J. (1975), 'Fourier analysis with unequally-spaced data', Ap&SS **36**, 137–158.

D'Onofrio, M., Capaccioli, M., Zaggia, S. R. & Caon, N. (1997), 'The relative distances to the Virgo, Fornax and Coma clusters of galaxies through the D_n–σ and the Fundamental Plane relations', MNRAS **289**, 847–862.

D'Onofrio, M., Zaggia, S. R., Longo, G., Caon, N. & Capaccioli, M. (1995), 'Major axis kinematics of 15 early-type galaxies in the Fornax cluster.', A&A **296**, 319+.

Durret, F., Pecontal, E., Petitjean, P. & Bergeron, J. (1994), 'High spatial resolution integral field spectroscopy of extended ionized gas around three quasars', A&A **291**, 392–402.

Eisenhauer, F., Tecza, M., Thatte, N., Mengel, S., Hofmann, R. & Genzel, R. (2000), Near-infrared-spectroscopy with Extremely Large Telescopes: Integral-field- versus multi-object-instruments, *in* 'Proceedings of the Backaskog workshop on extremely large telescopes', pp. 292+.

Emsellem, E. (1999), 2-D kinematics of galactic nuclei: From TIGER to OASIS, *in* 'ASP Conf. Ser. 182: Galaxy Dynamics', pp. 45+.

Emsellem, E., Greusard, D., Friedli, D. & Combes, F. (2001), Dynamical evidence for the existence of nuclear bars: OASIS observations and N-body models, *in* 'ASP Conf. Ser. 230: Galaxy Disks and Disk Galaxies', pp. 235–236.

Evans, I. N., Tsvetanov, Z., Kriss, G. A., Ford, H. C., Caganoff, S. & Koratkar, A. P. (1993), 'Hubble Space Telescope imaging of the narrow-line region of NGC 4151', ApJ **417**, 82+.

Faber, S. M., Tremaine, S., Ajhar, E. A., Byun, Y., Dressler, A., Gebhardt, K., Grillmair, C., Kormendy, J., Lauer, T. R. & Richstone, D. (1997), 'The centers of early-type galaxies with HST. IV. central parameter relations.', AJ **114**, 1771+.

Feichtinger, H. G. (1995), Iterative methods for scattered data approximation of smooth signals, *in* 'Conf. SampTA-95'.

Feichtinger, H. G. & Gröchenig, K. (1994), Theory and practice of irregular sampling, *in* J. Benedetto & M. Frazier, eds, 'Wavelets: Mathematics and Applications', CRC Press, pp. 305–363.

Ferland, G. J. & Netzer, H. (1983), 'Are there any shock-heated galaxies?', ApJ **264**, 105–113.

Ferrarese, L. & Merritt, D. (2000), 'A fundamental relation between supermassive black holes and their host galaxies', ApJ **539**, L9–L12.

Forbes, D. A., Ward, M. J., Rotaciuc, V., Blietz, M., Genzel, R., Drapatz, S., van der Werf, P. P. & Krabbe, A. (1993), 'High-resolution imaging of forbidden Fe II 1.64 microns, Brackett-gamma, and H2 1-0 S(1) emission in the starburst galaxy NGC 253', ApJ **406**, L11–L14.

Gebhardt, K., Bender, R., Bower, G., Dressler, A., Faber, S. M., Filippenko, A. V., Green, R., Grillmair, C., Ho, L. C., Kormendy, J., Lauer, T. R., Magorrian, J., Pinkney, J., Richstone, D. & Tremaine, S. (2000), 'A relationship between nuclear black hole mass and galaxy velocity dispersion', ApJ **539**, L13–L16.

Goodman, J. W. (1968), *Introduction to Fourier Optics*, McGraw-Hill.

Gorgas, J., Efstathiou, G. & Salamanca, A. A. (1990), 'Line strengths in early-type galaxies', MNRAS **245**, 217–237.

Goudfrooij, P., Mack, J., Kissler-Patig, M., Meylan, G. & Minniti, D. (2001), 'Kinematics, ages and metallicities of star clusters in NGC 1316: a 3-Gyr-old merger remnant', MNRAS **322**, 643–657.

Graps, A. (1995), 'An introduction to wavelets', IEEE Comput. Sci. Eng. 2(2), 50–61.

Gray, P. M., Phillips, M. M., Turtle, A. J. & Ellis, R. (1982), 'Fibre optic development at the AAO', Proc. Astron. Soc. Australia **4**, 477–479.

Greenhouse, M. A., Woodward, C. E., Thronson, H. A., Rudy, R. J., Rossano, G. S., Erwin, P. & Puetter, R. C. (1991), 'Near-infrared forbidden Fe II emission of M82 supernova remnants - Implications for tracing the supernova content of galaxies', ApJ **383**, 164–173.

Haehnelt, M. G. & Rees, M. J. (1993), 'The formation of nuclei in newly formed galaxies and the evolution of the quasar population', MNRAS **263**, 168–178.

Haynes, R., Content, R., Turner, J., Allington-Smith, J. R. & Lee, D. (1998a), 'SMIRFS-II: multi-object and integral-field near-IR spectroscopy at UKIRT', Proc. SPIE **3354**, 419–430.

Haynes, R., Doel, A. P., Content, R., Allington-Smith, J. R. & Lee, D. (1998b), 'TEIFU: a thousand element integral field unit for the WHT fed by the ELECTRA AO system', Proc. SPIE **3355**, 788–797.

Haynes, R., Lee, D., Allington-Smith, J., Content, R., Dodsworth, G., Lewis, I., Sharples, R., Turner, J., Webster, J., Done, C., Peletier, R., Parry, I. & Chapman, S. (1999), 'Multiple-Object and integral field near-infrared spectroscopy using fibers', PASP **111**, 1451–1468.

Heacox, W. D. (1986), 'On the application of optical-fiber image scramblers to astronomical spectroscopy', AJ **92**, 219–229.

Hecht, E. (1987), *Optics*, second edn, Addison-Wesley.

Hill, J. M., Angel, J. R. P., Scott, J. S., Lindley, D. & Hintzen, P. (1980), 'Multiple object spectroscopy—the Medusa spectrograph', ApJ **242**, L69–L72.

Hutchings, J. B., Crenshaw, D. M., Danks, A. C., Gull, T. R., Kraemer, S. B., Nelson, C. H., Weistrop, D., Kaiser, M. E. & Joseph, C. L. (1999), 'High-velocity line emission in the narrow-line region of NGC 4151', AJ **118**, 2101–2107.

Jenkins, C. R. & Scheuer, P. A. G. (1980), 'The dynamics of three radio galaxies - NGC 741, 1316 and 7626', MNRAS **192**, 595–610.

Kaiser, M. E., Bradley, L. D., Hutchings, J. B., Crenshaw, D. M., Gull, T. R., Kraemer, S. B., Nelson, C. H., Ruiz, J. & Weistrop, D. (2000), 'The resolved narrow-line region in NGC 4151', ApJ **528**, 260–275.

Knop, R. A., Armus, L., Larkin, J. E., Mathews, K., Shupe, D. L. & Soifer, B. T. (1996), 'Infrared spectroscopy of Paβ and [Fe II] emission in NGC 4151', AJ **112**, 81+.

Kormendy, J. & Richstone, D. (1995), 'Inward bound—the search for supermassive black holes in galactic nuclei', ARA&A **33**, 581+.

Kotel'nikov, V. A. (1933), 'On the transmission capacity of 'ether' and wire in electrocommunications', Izd. Red. Upr. Svgazi RKAA.

Kuntschner, H. (2000), 'The stellar populations of early-type galaxies in the Fornax cluster', MNRAS **315**, 184–208.

Kuntschner, H. & Davies, R. L. (1998), 'The ages and metallicities of early-type galaxies in the Fornax cluster', MNRAS **295**, L29+.

Lauer, T. R. (1999), 'Combining undersampled dithered images', PASP **111**, 227–237.

Le Fevre, O., Vettolani, G. P., Maccagni, D., Mancini, D., Picat, J. P., Mellier, Y., Mazure, A., Saisse, M., Cuby, J. G., Delabre, B., Garilli, B., Hill, L., Prieto, E., Arnold, L., Conconi, P., Cascone, E., Mattaini, E. & Voet, C. (1998), 'VIRMOS: visible and infrared multiobject spectrographs for the VLT', Proc. SPIE **3355**, 8–19.

Lee, D. (1998), New techniques in astronomical spectroscopy for 8m telescopes, PhD thesis, University of Durham.

Lester, D. F., Gaffney, N., Carr, J. S. & Joy, M. (1990), 'A near-infrared spectroscopic study of the starburst core of M82', ApJ **352**, 544–560.

Lewis, I. J., Sharples, R. M., Parry, I. R., Jones, L. R., Watson, F. G., Barker, S. A. & Rees, P. C. (1997), 'Autofib-2: commissioning results of a robotic multiobject fiber system for the William Herschel Telescope', Proc. SPIE **2871**, 1318–1324.

Longhetti, M., Rampazzo, R., Bressan, A. & Chiosi, C. (1998), 'Star formation history of early-type galaxies in low density environments. II. Kinematics', A&AS **130**, 267–283.

Malumuth, E. M. & Kirshner, R. P. (1985), 'Dynamics of luminous galaxies. II - Surface photometry and velocity dispersions of brightest cluster members', ApJ **291**, 8–31.

Maoz, D., Netzer, H., Mazeh, T., Beck, S., Almoznino, E., Leibowitz, E., Brosch, N., Mendelson, H. & Laor, A. (1991), 'High-rate active galaxy monitoring at the Wise Observatory. III - The broad-line region of NGC 4151', ApJ **367**, 493–506.

Marks, R. J., ed. (1993), *Advanced Topics in Shannon Sampling and Interpolation Theory*, Springer-Verlag.

Martin, P., Pinet, P. C., Bacon, R. & Rousset, A. (1995), Near-infrared spectroscopy of the surface of Mars derived from TIGER spectro-imaging data, *in* 'ASP Conf. Ser. 71: IAU Colloq. 149: Tridimensional Optical Spectroscopic Methods in Astrophysics', pp. 298+.

Mediavilla, E. & Arribas, S. (1995), 'Two-dimensional spectroscopy of the Seyfert galaxies NGC 5728 and 4151 in the [OIII] $\lambda\lambda$ 4959, 5007 and Hα spectral ranges', MNRAS **276**, 579–591.

Mediavilla, E., Arribas, S., del Burgo, C., Oscoz, A., Serra-Ricart, M., Alcalde, D., Falco, E. E., Goicoechea, L. J., Garcia-Lorenzo, B. & Buitrago, J. (1998), 'Two-dimensional spectroscopy reveals an arc of extended emission in the gravitational lens system Q2237+0305', ApJ **503**, L27–+.

Mengel, S., Eisenhauer, F., Tecza, M., Thatte, N. A., Roehrle, C., Bickert, K. & Schreiber, J. (2000), 'New era of spectroscopy: SINFONI NIR integral field spectroscopy at the diffraction limit of an 8-m telescope', Proc. SPIE **4005**, 301–309.

Morse, J. A., Raymond, J. C. & Wilson, A. S. (1996), 'On the viability of fast shocks as an ionization mechanism in active galaxies', PASP **108**, 426+.

Mouri, H., Kawara, K. & Taniguchi, Y. (2000), 'Excitation mechanism of near-infrared [Fe II] emission in Seyfert and starburst galaxies', ApJ **528**, 186–200.

Mouri, H., Nishida, M., Taniguchi, Y. & Kawara, K. (1990), 'Excitation mechanism of forbidden [Fe II] 1.644 micron emission in Seyfert and starburst galaxies', ApJ **360**, 55–62.

Mulchaey, J. S., Wilson, A. S. & Tsvetanov, Z. (1996), 'An emission-line imaging survey of early-type Seyfert galaxies. I. The observations', ApJS **102**, 309+.

Mundell, C. G., Pedlar, A., Baum, S. A., O'Dea, C. P., Gallimore, J. F. & Brinks, E. (1995), 'MERLIN observations of neutral hydrogen absorption in the Seyfert nucleus of NGC 4151', MNRAS **272**, 355–362.

Murray, G. J., Allington-Smith, J. R., Content, R., Dodsworth, G. N., Dunlop, C. N., Haynes, R., Sharples, R. M. & Webster, J. (2000), 'TEIFU: a high-resolution integral field unit for the William Herschel Telescope', Proc. SPIE **4008**, 611–622.

Nyquist, H. (1928), 'Certain topics in telegraph transmission theory', AIEE Trans. **47**, 617.

Oke, J. B. (1990), 'Faint spectrophotometric standard stars', AJ **99**, 1621–1631.

Osterbrock, D. E. (1989), *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*, Mill Valley: University Science Books.

Parry, I. R., Lewis, I. J., Sharples, R. M., Dodsworth, G. N., Webster, J., Gellatly, D. W., Jones, L. R. & Watson, F. G. (1994), 'Autofib-2: an automated fiber positioner for the prime focus of the William Herschel Telescope', Proc. SPIE **2198**, 125–133.

Pedlar, A., Kukula, M. J., Longley, D. P. T., Muxlow, T. W. B., Axon, D. J., Baum, S., O'Dea, C. & Unger, S. W. (1993), 'The radio nucleus of NGC 4151 at 5 and 8 GHz', MNRAS **263**, 471+.

Peletier, R. F., Davies, R. L., Allington-Smith, J. R., Bacon, R., Copin, Y., Emsellem, E., de Zeeuw, P. T., Miller, B., Bureau, M., Verolme, E., Carollo, C. M. & Monnet, G. (2001), 'SAURON: integral-field spectroscopy of galaxies', New Astron. Rev. **45**, 83–86.

Penston, M. V., Robinson, A., Alloin, D., Appenzeller, I., Aretxaga, I., Axon, D. J., Baribaud, T., Barthel, P., Baum, S. A., Boisson, C., de Bruyn, A. G., Clavel, J., Colina, L., Dennefeld, M., Diaz, A., Dietrich, M., Durret, F., Dyson, J. E., Gondhalekar, P., van Groningen, E., Jablonka, P., Jackson, N., Kollatschny, W.,

Laurikainen, E., Lawrence, A., Masegosa, J., McHardy, I., Meurs, E. J. A., Miley, G., Moles, M., O'Brien, P., O'Dea, C., del Olmo, A., Pedlar, A., Perea, J., Perez, E., Perez-Fournon, I., Perry, J., Pilbratt, G., Rees, M., Robson, I., Rodriguez-Pascual, P., Rodriguez-Espinosa, J. M., Santos-Lleo, M., Schilizzi, R., Stasinska, G., Stirpe, G. M., Tadhunter, C., Terlevich, E., Terlevich, R., Unger, S., Vila-Vilaro, V., Vilchez, J., Wagner, S. J., Ward, M. J. & Yates, G. J. (1990), 'The extended narrow line region of NGC 4151. I - Emission line ratios and their implications', A&A **236**, 53–62.

Perez, E., Gonzalez-Delgado, R., Tadhunter, C. & Tsvetanov, Z. (1989), 'The complex narrow-line region in NGC 4151', MNRAS **241**, 31P–36P.

Perryman, M. A. C., Favata, F., Peacock, A., Rando, N. & Taylor, B. G. (1999), 'Optical STJ observations of the Crab Pulsar', A&A **346**, L30–L32.

Peterson, B. M., ed. (1997), *An Introduction to Active Galactic Nuclei*, Cambridge University Press.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992), *Numerical recipes in C. The art of scientific computing*, second edn, Cambridge University Press.

Quinlan, G. D. & Hernquist, L. (1997), 'The dynamical evolution of massive black hole binaries - II. Self-consistent N-body integrations', New Astronomy **2**, 533–554.

Renka, R. J. (1988), 'Algorithm 661: QSHEP3D; quadratic Shepard method for trivariate interpolation of scattered data', ACM Trans. Math. Software **14**(2), 151–152.

Renka, R. L. (1983), 'Algorithm 624: Triangulation and interpolation of arbitrarily distributed points in the plane', ACM Trans. Math. Software **9**, 440–442.

Robinson, A., Vila-Vilaro, B., Axon, D. J., Perez, E., Wagner, S. J., Baum, S. A., Boisson, C., Durret, F., Gonzalez-Delgado, R., Moles, M., Masegosa, J., O'Brien, P. T., O'Dea, C., del Olmo, A., Pedlar, A., Penston, M. V., Perea, J., Perez-Fournon, I., Rodriguez-Espinosa, J. M., Tadhunter, C., Terlevich, R. J., Unger, S. W. & Ward, M. J. (1994), 'The extended narrow line region of NGC 4151', A&A **291**, 351–391.

Romani, R. W., Miller, A. J., Cabrera, B., Figueroa-Feliciano, E. & Nam, S. W. (1999), 'First astronomical application of a cryogenic Transition Edge Sensor spectrophotometer', ApJ **521**, L153–L156.

Savage, B. D. & Sembach, K. R. (1996), 'Interstellar abundances from absorption-line observations with the Hubble Space Telescope', ARA&A **34**, 279–330.

Scargle, J. D. (1982), 'Studies in astronomical time series analysis. II—Statistical aspects of spectral analysis of unevenly spaced data', ApJ **263**, 835–853.

Scargle, J. D. (1989), 'Studies in astronomical time series analysis. III—Fourier transforms, autocorrelation functions, and cross-correlation functions of unevenly spaced data', ApJ **343**, 874–887.

Schweizer, F. (1981), 'Optical properties of the central region of NGC 1316 - a small bright core in a giant D galaxy', ApJ **246**, 722–739.

Seyfert, C. K. (1943), 'Nuclear emission in spiral nebulae.', ApJ **97**, 28+.

Shannon, C. E. (1949), 'Communication in the presence of noise', Proc. IRE **37**, 10–21.

Shaya, E. J., Dowling, D. M., Currie, D. G., Faber, S. M., Ajhar, E. A., Lauer, T. R., Groth, E. J., Grillmair, C. J., Lynd, R. & O'Neil, E. J. (1996), 'Hubble Space Telescope Planetary Camera images of NGC 1316 (Fornax A)', AJ **111**, 2212+.

Simpson, C., Forbes, D. A., Baker, A. C. & Ward, M. J. (1996), 'Forbidden Fe^+ emission from active galaxies', MNRAS **283**, 777–797.

Strohmer, T. (1993), Efficient Methods for Digital Signal and Image Reconstruction from Nonuniform Samples, PhD thesis, University of Vienna.

Strohmer, T. (1997), 'Computationally attractive reconstruction of band-limited images from irregular samples', IEEE Trans. Image Process. **6**, 540–554.

Swan, P. R. (1982), 'Discrete fourier transforms of nonuniformly spaced data', AJ **87**, 1608+.

Tedds, J. (1999), 3D Integral field H_2 spectroscopy in outflows, in 'CUP Astrophysics Series: H_2 in Space, Paris, 1999. Ed. Combes & Pineau.', pp. E31–+.

Unger, S. W., Pedlar, A., Axon, D. J., Whittle, M., Meurs, E. J. A. & Ward, M. J. (1987), 'The extended narrow-line region in radio Seyferts - Evidence for a collimated nuclear UV field?', MNRAS **228**, 671–679.

Unser, M. (2000), 'Sampling—50 years after Shannon', IEEE Proc. **88**(4), 569–587.

van der Hulst, J. M., Terlouw, J. P., Begeman, K. G., Zwitser, W. & Roelfsema, P. R. (1992), The groningen image processing system, GIPSY, in 'ASP Conf. Ser. 25: Astronomical Data Analysis Software and Systems I', Vol. 1, pp. 131+.

van der Marel, R. P., Cretton, N., de Zeeuw, P. T. & Rix, H. (1998), 'Improved evidence for a black hole in M32 from HST/FOS spectra. II. Axisymmetric dynamical models', ApJ **493**, 613+.

Vanderriest, C. (1980), 'A fiber-optics dissector for spectroscopy of nebulosities around quasars and similar objects', PASP **92**, 858–862.

Vio, R., Strohmer, T. & Wamsteker, W. (2000), 'On the reconstruction of irregularly sampled time series', PASP **112**, 74–90.

Weitzel, L., Krabbe, A., Kroker, H., Thatte, N., Tacconi-Garman, L. E., Cameron, M. & Genzel, R. (1996), '3D: The next generation near-infrared imaging spectrometer', A&AS **119**, 531–546.

Whittaker, E. T. (1915), 'On the functions which are represented by the expansions of the interpolation theory', Proc. Roy. Soc. Edinburgh, Sect. A **35**, 181.

Winge, C., Axon, D. J., Macchetto, F. D., Capetti, A. & Marconi, A. (1999), 'Hubble Space Telescope Faint Object Camera spectroscopy of the narrow-line region of NGC 4151. I. Gas kinematics', ApJ **519**, 134–152.

Wright, G. S., Mountain, C. M., Bridger, A., Daly, P. N., Griffin, J. L. & Ramsay Howat, S. K. (1993), 'CGS4 experience: two years later', Proc. SPIE **1946**, 547–557.