Dissertation

submitted to the

Combined Faculties for the Natural Sciences and for Mathematics
of the Ruperto-Carola University of Heidelberg, Germany

for the degree of

Doctor of Natural Sciences

presented by

Sebastian Jester, MPhys
born in Luxembourg

Oral examination: 17$^{\text{th}}$ of October, 2001

High-resolution multi-wavelength study of the jet in 3C 273

Abstract

The jet in 3C 273 is one of only a few extragalactic optical synchrotron jets which are large and bright enough to be studied in detail. We present new broad-band observations of this jet at the unprecedented common resolution of $0\rlap{.}''3$ which have been obtained with the *Very Large Array* at radio and the *Hubble Space Telescope* at infrared, optical and ultraviolet wavelengths. These observations reveal a flattening of the high-frequency spectrum of the jet, a surprising feature which cannot be accounted for by any single-population synchrotron model. Both the observed flattening of the high-frequency spectrum and the X-ray emission from the jet can be explained by a model in which two distinct electron populations contribute to the jet's emission.

We fit spatially resolved synchrotron spectra for the jet and determine the run of the maximum particle energy. The decrease of the maximum particle energy along the jet is much slower than expected from the observed synchrotron emission. We find no evidence for localised acceleration or loss sites. We show that relativistic beaming and/or sub-equipartition magnetic fields cannot remove the discrepancy between light-travel time along the jet and the shorter lifetime of electrons emitting optical synchrotron radiation. We consider this further evidence in favour of a distributed electron acceleration process.

Zusammenfassung

Der Jet in 3C 273 ist einer von nur wenigen extragalaktischen Jets, die optische Synchrotronemission aufweisen und dabei groß und hell genug sind, um eine detaillierte Untersuchung zu erlauben. Diese Arbeit stellt neue Beobachtungen vor, die mit den Radioteleskopen des *Very Large Array* sowie im infraroten, optischen und ultravioletten Bereich des Spektrums mit dem *Hubble Space Telescope* gewonnen wurden und bei einer bislang unerreichten Auflösung von $0\rlap{.}''3$ kombiniert werden. Diese Beobachtungen zeigen, daß das Spektrum des Jets zum Ultravioletten hin flacher wird — ein überraschendes Ergebnis, das zeigt, daß die Strahlung des Jets nicht mehr als Synchrotronstrahlung einer einzigen Elektronenverteilung interpretiert werden kann. Sowohl der flache Verlauf des hochfrequenten Spektrums als auch die Röntgenemission des Jets können durch ein Modell erklärt werden, in dem zwei verschiedene Elektronenverteilungen Synchrotronemission abstrahlen.

Die neuen Beobachtungen werden zu Synchrotronspektren kombiniert, die eine Bestimmung der Maximalenergie der strahlenden Teilchen ermöglichen. Der Abfall der Maximalenergie entlang des Jets erfolgt sehr viel langsamer, als es sich direkt aus den beobachteten Synchrotronverlusten ergeben würde. Es finden sich trotzdem keine Hinweise auf räumlich begrenzte Gebiete, in denen die Teilchen nachbeschleunigt werden. Auch relativistische Zeitdehnungs-Effekte können die Diskrepanz zwischen der Lichtlaufzeit entlang des Jets und der kürzeren Lebenszeit der Elektronen, die optische Synchrotronstrahlung abstrahlen, nicht erklären. Diese Ergebnisse festigen die Schlußfolgerung, daß entlang der ganzen Länge des Jets von 3C 273 Teilchen nachbeschleunigt werden.

Bomi
und meinen Eltern

Contents

Chapter 1

Introduction

In 1918, Heber Curtis noticed a "curious straight ray" extending from the center of the elliptical galaxy M87 [20]. This was the first discovery of an extragalactic jet. Jets are collimated outflows transporting mass, energy, momentum as well as angular momentum and electromagnetic fields outwards from a central object. They are now observed emerging from objects spanning a variety of length and mass scales, but with the common property of harbouring an accretion disc:

- Young Stellar Objects losing mass through jets which are overdense with respect to the external medium, with typical velocities of a few hundred kilometers per second. They emit thermal radiation, fuelled by their internal energy or by energy dissipated in a variety of shock phenomena. The outflow might be an important factor in the removal of angular momentum from the accretion disc [79].

- Stellar-mass Black Holes or Neutron Stars in a binary system, onto which material is accreted from a companion star. These are observed either as X-ray binaries or as *Microquasars* [69], with thermal emission from the accretion disc and non-thermal emission from the jets.

- Supermassive Black Holes at the centres of galaxies, fuelling an Active Galactic Nucleus (AGN). As implied by the superluminal motion observed by Very Long Baseline Interferometry (VLBI) on milli-arcsecond scales, these jets can move at relativistic speeds at least near the core source launching them. The jets terminate in a shock which is observed as bright *hot spot*, embedded in fainter radio lobes, and the object is observed as radio galaxy (Fig. 1.1) or radio-loud quasar. These jets are composed of plasma which is underdense with respect to the external medium and detectable are through their non-thermal emission.

A detailed understanding of the formation of these jets, their connection to the accretion disc from which they are launched, and the physics governing their internal structure and observable properties has not been achieved[1] and is the subject of ongoing research. A study of any object with jets is therefore expected to bear importance for the understanding of all such objects. In the present study, we consider the synchrotron emission from the kiloparsec-scale jet of the quasar 3C 273.

---

[1]The emission of gamma-ray-bursts is also thought to arise from jets; however, these jets are probably not related to continuous accretion, but to cataclysmic explosion events.

Figure 1.1:   Cygnus A is the prototypical radio galaxy and the brightest radio galaxy for observers on Earth. The linear extent from hot spot to hot spot is of order $100h_{70}^{-1}$ kpc.

1.1   Extragalactic jets

Although the first few jets were detected at optical wavelengths, the vast majority of the few hundreds of extragalactic jets known today were all detected as radio jets, and just over a dozen of them show optical emission. The first optical jets thus present the extremes of the entire population, large and optically bright enough to be detected by ground-based optical telescopes. The light from essentially all other optical jets could only be detected with the Hubble Space Telescope (HST).

Before the launch of HST, the study of jets (apart from the three known optical jets) was the domain of radio astronomers. When radio astronomy had become a field of study in the 1950s, the first radio surveys of the sky were carried out [92, *e. g.*]. The radio sources discovered in these surveys were resolved by the first interferometers into giant doubles with faint cores, the bulk of the radio luminosity being emitted by the lobes (*cf.* Fig. 1.1). In the "twin exhaust" model [9, see below] of these radio sources, energy is continuously provided from the core to the lobes by a collimated relativistic flow — a jet. As more sensitive radio interferometers were built which provided images with higher dynamic ranges, the predicted jets were indeed found connecting the core to the hot spots and lobes.

1.1.1   3C 273 — the prototypical quasar

The subject of this work is the jet of the prototypical quasar 3C 273. It had been discovered in one of the first radio surveys, and later lunar occultation observations [32] showed that it consists of two radio components separated by about $20''$ at position angle $\approx 220°$. The northeastern of these components, 3C 273B, was optically identified with a faint, blue stellar object showing strange emission lines, leading to the coining of the term *quasar* [94]. The strange appearance of its emission lines were ascribed to a large redshift of cosmological origin [29, 94].[2] The southwestern component of 3C 273 was identified with "an associated nebulosity" on optical plates [29, 32, 76] — another of the first few optical jets had been

---

[2]The cosmological nature of the redshift of quasars was put beyond any doubt only in 1978, when Stockton published a redshift survey of galaxies near 27 radio-loud quasars with $z > 0.45$ [96]. Eight of the quasars in this survey have galaxies with similar redshift associated with them.

detected, whose spectrum was described as a "weak, bluish continuum" [29].

1.1.2 Synchrotron emission from radio sources

The spectra of the first known radio sources were not thermal black-body spectra but power laws, *i. e.,* of non-thermal origin, with a fairly universal spectral index. It was first suggested for the Crab nebula that this emission might be synchrotron emission, and the predicted large degree of linear polarisation was indeed observed. A strong polarisation was also observed for the jet in M87 [2]. In the following years, synchrotron emission was established as emission mechanism for the jets, lobes and hot spots of extragalactic radio sources — and although the jet of 3C 273 was suspected to be an optical synchrotron source by its discoverers in 1964 [29], the synchrotron nature of its optical emission was only confirmed as late as 1991 [87].

Even before the first jets had been observed, it was clear that the energy radiated by the lobes had to be supplied continuously from the central source. As is well known from standard synchrotron theory, the typical life-time of electrons against synchrotron losses is of the order of or smaller than the light travel time from nucleus to hot spots, over scales of tens or hundreds of kiloparsecs, so that the lobes could not have been simply ejected from the central source. The "twin exhaust" model was developed to explain the continuous feeding of the lobes [9] and is now the standard model of extragalactic radio sources.

In the standard model [5], the energy is fed into the radio lobes by a jet. The jet is highly collimated bulk relativistic outflow originating in the core of the radio source, near the central engine, a super-massive black hole ($\approx 10^9$ M$_\odot$ in the strongest sources). The central engine feeds energy into the jets through a collimation mechanism connected to the presence of an accretion disk. Where the jet impinges on the denser intra-cluster medium, a double shock structure forms, consisting of a bow shock separating jet material from the external medium and a Mach disk at which the relativistic flow is decelerated and bulk kinetic energy is channelled into highly relativistic particles through a shock acceleration mechanism.

These particles emit the observed synchrotron radiation and the radio hot spot is usually assumed to coincide with the Mach disk. The optical synchrotron emission observed from hot spots at the ends of radio jets can be well explained by first-order Fermi acceleration at the jet-terminating shock [34, 63, 66, 68]. After flowing through the Mach disk, the jet material escapes towards the sides and forms the radio lobes. One should make a clear distinction between the emission from the hot spot itself, the lobes and the body of the jet. Although all emit synchrotron radiation, the physical processes accelerating the emitting particles may be quite different between these regions.

1.1.3 Models of particle acceleration

Highly relativistic electrons are required to explain the observed emissivities of radio sources. A sufficient density of electrons of such high energies cannot be provided by thermal processes. The first-order Fermi acceleration provides a natural explanation both for obtaining a power law and predicting an exponent close to values observed in the spectra of synchrotron sources.

The presence of synchrotron radiation from jets shows that they comprise both relativistic particles and magnetic fields. Relativistic particles will necessarily not be in a bound state of any description so that the jet material will be a plasma. The jets are thought to be "thin" (under-dense with respect to the galactic and intergalactic medium) and "light" (most

of their energy is kinetic). Apart from the electrons, there have to be charge-balancing positively charged particles in the flow to avoid the radio sources becoming electric dipoles.

There is no consensus on whether jets are composed of electron-positron pair plasma, or proton-electron hydrogen plasma. In the latter case, the total energy content of a typical jet must be orders of magnitude larger than the emitted energy because of the large proton rest mass, *i. e.,* proton-electron jets are rather inefficient. An electron-positron jet would be more efficient. However, there are other difficulties with this kind of jet: there has to be some sort of pair production region to make positrons at all. Once they are made, they must not annihilate with electrons within the jet or hot spot, to allow the electrons to illuminate all parts of the radio source. Pair annihilation would also produce pairs of 511 keV-photons which should be observable from the annihilation region, making it a gamma-ray line source. Positrons would emit synchrotron radiation in the same way as electrons. Because they gyrate in a magnetic field in the opposite sense to electrons, the polarisation properties of an electron-positron plasma are different from those of an electron-proton plasma, especially for any circular polarisation component.[3]

As mentioned, the shock at the front of the jet flow is thought to be the "working surface" for the first-order Fermi mechanism. This mechanism was first suggested by Enrico Fermi in 1954 [25] to explain the power-law energy spectrum of cosmic-ray particles impinging on the upper atmosphere and further developed theoretically by A. R. Bell [6, 7], among others.

The idea of the first-order Fermi mechanism at a non-relativistic strong shock is that particles are repeatedly crossing a shock and gain energy on each crossing. The energy gain is proportional to the initial energy. The second ingredient is a finite probability for each particle to escape from the shock region, so that larger energy gains become increasingly unlikely. The combination of the two results in a power-law distribution of electron energies whose exponent only depends on the compression ratio $r(> 1)$ of pre- and post-shock fluids as $p = \frac{r+2}{r-1}$. The power law exponent is -2 for a limiting-case strong shock which corresponds to a synchrotron radiation spectral index of -0.5. This result is remarkable because no detailed micro-physics are required to obtain a universal power law of synchrotron spectra, with an exponent close to observed values. The mechanism does require the injection of supra-thermal particles, *i. e.,* it is really more of a reacceleration mechanism than one selecting a few particles from a thermal distribution. The nature of the injection mechanism is a further unsolved puzzle in our understanding of radio sources.

In practice, the power law will not extend to infinite electron energies but there will be a cutoff. The cutoff energy is that at which the synchrotron losses during one acceleration cycle cancel out the energy gain during that cycle.[4]

1.1.4   Optical synchrotron emission from jets

Electrons with the highly relativistic energies required for the emission of high-energy (infrared and optical) synchrotron radiation have a very short lifetime which is much less than the light-travel time down the jet body in, *e. g.,* 3C 273 (already noted by Greenstein & Schmidt

---

[3]In principle, it should be possible to distinguish between the two using polarisation measurements. This has indeed been attempted for the (milli-arcsecond) jet of 3C 279 [106], but the result relies heavily on assumptions.

[4]Even if there are no synchrotron losses, the highest attainable electron energy is limited by inverse-Compton scattering of cosmic microwave background photons by high-energy electrons (a detailed discussion of this is presented in Sect. 5.1). In compact sources with a high synchrotron luminosity, the highest-energy electrons will experience Compton scattering off self-produced synchrotron photons (the synchrotron self-Compton effect).

[29]). Under the assumptions of the standard model, those particles which are responsible for synchrotron emission from the jets themselves can only be accelerated in the source's core. It cannot therefore account for optical synchrotron emission from the *body* of jets such as those in M87 and 3C 273, which extend over tens of kiloparsecs in some cases. Instead, the particles radiating high-frequency synchrotron emission must be accelerated *inside* the jet, not just at its terminating shock.

Observations of optical synchrotron emission from such jets [87, 88] as well as from the "filament" near Pictor A's hot spot [81, 86] suggest that both an extended, "jet-like" and a localised, "shock-like" acceleration process are at work in these objects in general and 3C 273's jet in particular [68]. The extended mechanism may also be at work in the lobes of radio galaxies, where the observed maximum particle energies are above the values implied by the losses within the hot spots [62] and by the dynamical ages of the lobes [10].

1.1.5 X-ray emission from jets

Additional problems are posed by observations at even higher frequencies: ROSAT observations showed X-ray emission from the jets in M87 [71, 72] and 3C 273 [89]. More recently, observations with the new X-ray observatory Chandra showed extended X-ray emission from the jets in PKS 0637−752 (which also shows some optical emission [14, 15, 95, 98]) and Pictor A (which is a true radio jet [109]) as well as other jets and hot spots. Chandra also supplied the first high-resolution X-ray images of the jets in M87 and 3C 273 [58, 84]. The X-rays from these objects also seem to be of non-thermal origin: they could at least partially be due to synchrotron emission in M87 and 3C 273 [58, 84, 89]. Inverse-Compton scattering could also produce X-rays. The photon seed field can be provided by the synchrotron source itself if it is sufficiently compact, for example in the hot spots of Cygnus A [108]. If the bulk flow of a jet is still highly relativistic on large scales, the boosted energy density of the cosmic microwave background radiation field can lead to the observed X-ray fluxes [98]. In all cases, the energy observed as X-rays is provided by those electrons also producing the radio-optical synchrotron emission, decreasing their cooling timescale even below the synchrotron cooling scale.

1.2 Aim of this work

The fundamental question posed by the observation of optical extragalactic jets is thus: how can we explain high-frequency synchrotron and inverse-Compton emission far from obvious acceleration sites in extragalactic jets? While information on the source's magnetic field structure may be obtained from the polarisation structure, the diagnostic tool for the radiating particles is a study of the synchrotron continuum over as broad a range of frequencies as possible, *i. e.,* from radio to UV wavelengths, and with sufficient resolution to discern morphological details. The shape of the synchrotron spectrum gives direct insight into the shape of the electron energy distribution, thus also constraining the emission by the inverse-Compton process at other wavelengths.

There is considerable effort to model the structure and dynamics of extragalactic jets [13, 26, 99]. Since they are composed of plasma and moving at relativistic speeds, a fully relativistic three-dimensional magnetohydrodynamic code tracing synchrotron cooling as well as shock acceleration including back-reaction of particles on shocks would be required for a
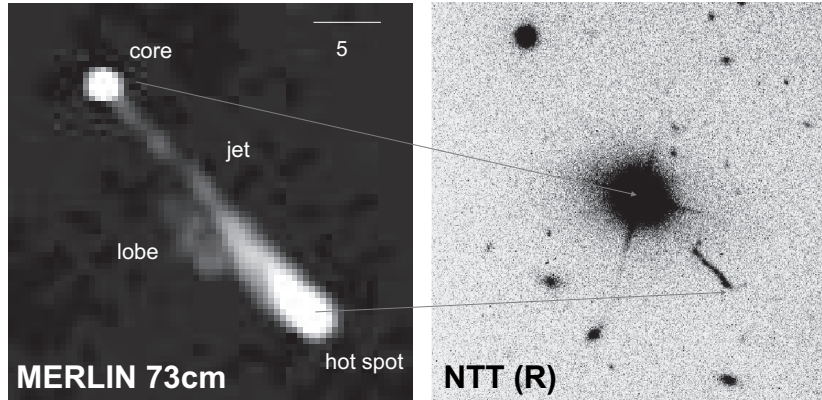
Figure 1.2:   MERLIN radio image (left) and ESO New Technology Telescope R-band image (right) of the jet of 3C 273. Taken from [90].

complete description. A realisation of such a scheme at sufficient resolution has not been achieved yet. Detailed observations are required to make progress.

Of all the known optical jets, there are only three with sufficient angular size and surface brightness to be studied in any detail (of which two were among the first to be detected): those in M87 (a radio galaxy), PKS $0521 - 365$ (an elliptical galaxy with a BL Lac core), and 3C 273 (a quasar).

We present a detailed study of the jet in 3C 273 using broad-band observations at various wavelengths obtained with today's best observatories in terms of resolution: the VLA (in combination with MERLIN data at $\lambda 6$ cm) and the HST. The aim of this study is the determination of the spectral shape of the synchrotron emission.

3C 273's radio jet extends continuously from the quasar out to a terminal hot spot at $21''\!.4$ from the core, while optical emission has been observed only from $10''$ outwards (Fig. 1.2).[5] On ground-based images, the optical jet appears to consist of a series of bright knots with fainter emission connecting them. So far, synchrotron spectra have been derived for the hot spot and the brightest knots using ground-based imaging in the radio [17], near-infrared K'-band [73] and optical I, R, B-bands [87] at a common resolution of $1''\!.3$ [64, 89]. This radio-to-optical continuum can be explained by a single power-law electron population leading to a constant radio spectral index[6] of -0.8, but with a high-energy cutoff frequency decreasing from 10^{17} Hz to 10^{15} Hz outwards along the jet.

Here, we present new observations, which constitute a unique data set in terms of resolution and wavelength coverage for any extragalactic jet — only M87 is similarly well-studied [84]. Using these observations at wavelengths 3.6 cm, 2.0 cm, 1.3 cm, 1.6 μm, 620 nm and 300 nm, we derive spatially resolved (at $0''\!.3$) synchrotron spectra for the jet. By fitting synchrotron spectra according to Heavens & Meisenheimer [34], we derive the maximum particle energy everywhere in the jet and aim to thus identify regions in which particles are either predominantly accelerated or lose energy. A detailed comparison of the observed and fitted spectral shapes will test the assumption that the entire observed continuum from radio to

---

[5] For the conversion of angular to physical scales, we assume a flat cosmology with $\Omega_{\rm m} = 0.3$ and $H_0 = h_{70} \times 70\,{\rm km\,s^{-1}Mpc^{-1}}$, leading to a scale of $2.7 h_{70}^{-1}$kpc per second of arc at 3C 273's redshift of 0.158 (see App. D).

[6] We define the spectral index α such that $f_\nu \propto \nu^\alpha$.

ultra-violet (and possibly even X-rays) can be explained as synchrotron radiation from a single electron population.

The derived shape for the electron energy distribution will be an invaluable help in determining the emission process responsible for the X-rays, and we will look for correlations between the fitted synchrotron spectra and the X-ray observations. Subtle deviations from the fitted spectral shape will be a guide for future investigations into the nature of the derivations, and the physical conditions giving rise to the observed emission. Only trough a detailed understanding of this jet, and its similarities with and differences to that in M87, can we hope to set the agenda for a study of other extragalactic jets.

Chapter 2

Observations and data reduction

2.1 Observations

2.1.1 Radio data

The jet has been observed at all wavelength bands available at the NRAO's Very Large Array (VLA), *i.e.,* at 90 cm, 20 cm, 6 cm, 3.6 cm, 2 cm, 1.30 cm and 0.7 cm. Observations were carried out between July 1995 and November 1997, to obtain data with all array configurations (thus covering the largest range of spatial frequencies). Total integration times are of order a few times 10,000 s in each band. At 3.6 cm, the best achievable resolution (set by the maximum VLA baseline of just over 30 km) is $0''\!.24$, with better resolution at shorter wavelengths. However, at 0.7 cm, only few antennas were equipped with receivers at that time, and the flux density of the jet is so low that only the hot spot is detected even at a fairly low resolution of $0''\!.35$. Observations were carried out in spectral line mode, in which the observing bandwidth of 12.5 MHz was split into 16 channels, each of which is correlated independently of all others. The spectral line mode was chosen because an image of the jet in 3C 273 with a dynamic range exceeding 200,000:1 had been obtained at a wavelength of 6 cm and with $0''\!.4$ resolution [80]. Such a high dynamic range is necessary to simultaneously image the quasar itself, which is a strong radio source, and the faintest features of the jet. Similarly high dynamic ranges were expected to be achievable at shorter wavelengths, that is, higher resolution.

A combination of the 6 cm data with observations obtained at the United Kingdom's Multi-Element Radio Linked Interferometer Network (MERLIN), which has significantly longer baselines than the VLA, was attempted in order to enhance the resolution and match them with the rest of the data set, to enhance the wavelength coverage. The MERLIN interferometric array consists of 8 radio telescopes, of which 6 are fitted with receivers at $\lambda 6$ cm. MERLIN has baselines up to 217 km, yielding a maximum resolution of $0''\!.04$ at $\lambda 6$ cm, about ten times better than the VLA at this wavelength. However, due to problems with the flux calibration between these data sets, the combined image could not yet be used (see Sect. 2.2.1). Therefore, the present analysis considers the data at 3.6 cm, 2 cm and 1.3 cm. The common resolution was fixed at $0''\!.3$, slightly inferior to the resolution of the data at 3.6 cm.

2.1.2 Optical and ultraviolet data

Optical and ultraviolet observations were made using the Planetary Camera of the second Wide Field and Planetary Camera (WFPC2) on board the Hubble Space Telescope on March
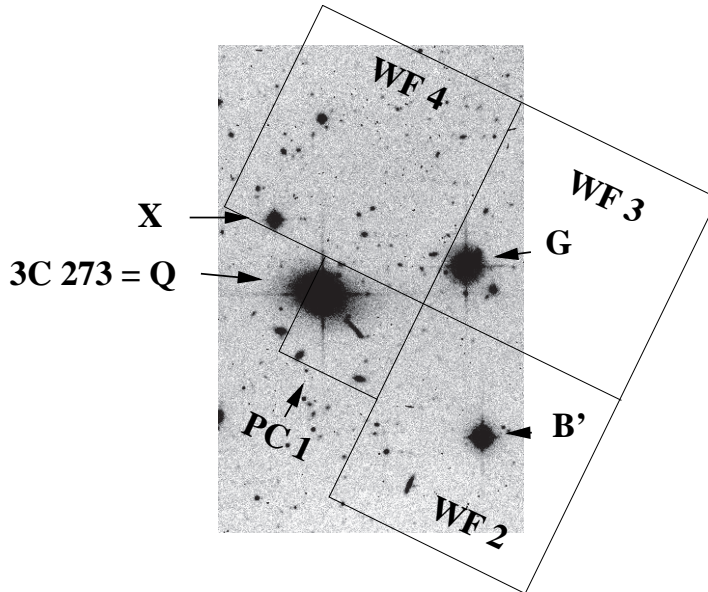
Figure 2.1: The field of view of the Wide Field and Planetary Camera 2 superimposed on a direct R-band image obtained at the ESO/MPG 2.2 m telescope at La Silla by H.-J. Röser and K. Meisenheimer. The stars are labelled as in [87].

23rd and June 5th/6th, 1995, through the filters F300W (centered near 300 nm) and F622W (centered near 620 nm, roughly corresponding to R_C in the Kron-Cousins-system).

The WFPC2 is one unit consisting of four separate cameras, the Planetary Camera (PC1) and three identical wide field cameras (WF2–4). The physical pixel size and scale of the Planetary Camera is half that of the Wide Field Cameras. Together, they map a contiguous area of the sky with the characteristic chevron shape (Fig. 2.1). The Planetary Camera has a 800×800 pixel Loral Charged Coupled Device (CCD) chip as detector. The camera has a (nominal) pixel scale of $0.04554''$ at $f/28.3$. The resulting size of the field of view is $36.4'' \times 36.4''$.

The pointing and roll angle of the telescope were chosen under the following considerations. The jet should be imaged at the centre of the Planetary Camera chip, in order to minimise any optical distortion effects. The quasar core has to be used as position reference when combining data taken by different instruments; the exact relative location of its image can be most easily determined if it is on the same chip as the jet image. An exact matching of these observations with those taken by other instruments is facilitated by having as many point sources imaged as possible, to allow the accurate determination of both shifts and rotations. In our case, there are only four point sources in the vicinity of the jet: the quasar core and three field stars [Table 3 in 87]. The pointing and roll were therefore chosen such that each star is imaged on a separate WF chip, the quasar is near a corner of the PC and the jet is near its centre. In addition, allowance is made for small pointing offsets, which help in removing artifacts introduced by the camera pixels.

The total observing times would ideally be such that the jet features have a similar signal-to-noise ratio in both filters. The exposure time in the red wavelength band was chosen as 10 ksec which enables imaging of the jet knots at a pixel S/N of around 25, while the inter-knot regions still have a S/N of 7–10. The jet flux decreases towards shorter wavelengths. As

the CCD camera becomes less sensitive in the UV region, the exposure time to achieve the same S/N becomes significantly larger. The total exposure time in the UV band was therefore fixed at 35.5 ks, which gives a typical pixel S/N of 10 in most of the knots, a value of 20 only in the brightest knot, and only marginal detection of the inter-knot regions. The corresponding *aperture* signal-to-noise ratios at $0\rlap{.}''3$ beam size will be larger by a factor of about 10. Despite the lower S/N compared to the optical data, the ultraviolet data are important in determining the shape of the synchrotron spectrum at the highest electron energies with the high spatial resolution afforded by the HST.

All exposures are grouped into three visits[1], distinguished by the first two digits of the file number. The total exposure time was split into single integrations of around 2500 s each, the longest possible duration of roughly one half of a full HST orbit[2]. The U-band images are read noise limited (see Sect. A.2), therefore the longest possible integration time was chosen, resulting in 16 individual exposures at different pointings. In the R-band, the total exposure time was split into four single exposures at different pointings. These are sky background limited, but more splits would only have increased the amount of raw data having to be processed. The exposure time was therefore chosen to be one orbit for each individual exposure. An observation log is given in Tab. 2.1.

2.1.3 Near-infrared data

The near-infrared NICMOS camera 2 ["NIC2"; see 11, 101] on board HST was used to image the jet through filter F160W (centered at $1.6\,\mu$m), on February 6th and March 3rd/5th, 1998, as HST proposal 7848. This wavelength is critical in precisely and accurately determining the cutoff frequency along the jet. The total exposure time was 34560 s. NIC2 has 256×256 pixels of size $\approx 0\rlap{.}''076$ on sky, making it well-suited for diffraction-limited observations at $1.6\,\mu$m (Rayleigh criterion for the HST's 2.4 m aperture, $0\rlap{.}''17$). The field of view is just under $20''$ squared, significantly smaller than that of WFPC2. In fact, the detector scale along the x axis is 0.9% larger than that of the y axis, because of a tilt of the detector plane with respect to the camera focal plane. In addition, there is a global scale change with time because of expansion of the dewar assembly which moves the detector plane [Sect. 5.4 in 22]. Both effects are ignored in the reduction process and instead accounted for by appropriate placement of photometry apertures (see Section 3.1). The array consists of four quadrants which are read out separately from each other.

For these observations, the telescope was rotated such that North is approximately along the positive x direction of the detector, and correspondingly East along the positive y direction.

---

[1] A "visit" is the term for "a group of exposures to be executed together". The telescope pointing is established and guide stars are acquired at the beginning of a visit.

[2] HST observing time is allocated as a number of actual spacecraft orbits of duration 97 min. Depending on the telescope's orientation and the location of the target on sky, it is observable for the full orbit or, as was the case here, for only half of the orbit due to obscuration by the Earth.

Filename	Filter	Date	Exp. time [s]	POS_TARG x ["], y ["]	
u2nz0101	F622W	06/06/95	2300	0.000	0.000
u2nz0102			2500	−0.501	0.000
u2nz0103			2	−0.501	0.000
u2nz0104			2600	0.000	0.501
u2nz0105			2600	0.501	0.000
u2nz0106	F300W	06/06/95	2600	0.000	0.000
u2nz0107			2600	−0.501	0.000
u2nz0108			2500	−1.002	0.501
u2nz0109			8	−1.002	0.501
u2nz010a			2600	0.501	−0.501
u2nz0201	F300W	23/05/95	2300	0.334	−0.334
u2nz0202			2600	0.835	0.167
u2nz0203			2500	−0.668	−0.334
u2nz0204			8	−0.668	−0.334
u2nz0205			2600	−0.167	0.668
u2nz0206			2600	1.336	0.668
u2nz0301	F300W	05/06/95	2300	0.668	0.835
u2nz0302			2500	−0.835	0.334
u2nz0303			8	−0.835	0.334
u2nz0304			2600	−0.334	−0.167
u2nz0305			2600	0.167	0.835
u2nz0306			2600	1.169	−0.167

Table 2.1: Observation log of proposal 5980. The POS_TARG column refers to commanded offsets from the reference pointing between exposures

2.2   Data reduction and calibration

2.2.1   VLA data

Calibration and imaging

The process of calibrating the interferometric data obtained at the VLA and deriving images of the surface brightness distribution on sky was carried out by R. E. Perley (NRAO, Socorro). He supplied files containing total-intensity images as well as polarimetric information (polarised flux, degree of polarisation, polarisation angle). Details of the data reduction and a full study of the jet at radio wavelengths will be presented elsewhere [R. Perley, *in prep.*]. Figure 2.2 shows the three images used here.

Error sources

While the error sources for the HST images are well-known and quantified, the transformation of interferometric fringes to images introduces uncertainties which are not easily quantified [80]. An absolute flux calibration is not straightforward for radio data since the measurement
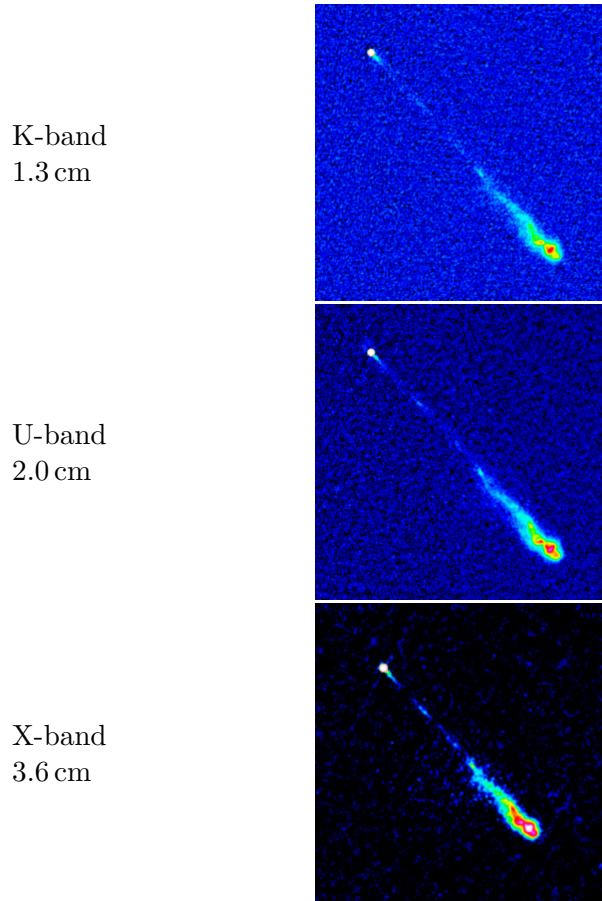
Figure 2.2: VLA images employed in this work

of absolute radio fluxes is difficult. Therefore, most observations rely on a calibration relative to standards established by Baars et al. [3], which is probably correct to 1–2% up to 23 GHz [R. Perley, *priv.comm.*]. The second problem for radio data is related to the fact that the brightness distribution on the sky cannot be inferred uniquely from the fringe patterns it produces through an interferometer because the deconvolution involved in the inverse process is not unique. Thus, although the noise on the actually detected signal is well-known, there is no good estimate on how this translates to noise in the image plane. In addition, there may be artifacts present in the image, *i. e.,* errors in the sense that the inferred brightness distribution does not correspond to the true brightness distribution. The usual way to quote the quality of a radio map is the "dynamic range", defined as the ratio of the peak surface brightness to the RMS noise of a blank region of sky. This RMS noise is the single quantifiable noise estimate for radio maps and a lower limit to the true precision. Table 2.2 quotes the dynamic ranges for the images used here. From the table, it is clear that the dynamic ranges which were actually achieved fell considerably short of the expectations.

Combination with MERLIN data

In order to enhance the resolution of the VLA data set at $\lambda 6$ cm, a combination was attempted with MERLIN data. The VLA interferometric data set comprising the calibrated data was

VLA band	λ	Peak flux	RMS noise	Dynamic range
	cm	mJy	mJy	
C*	6	35.9	6.0×10^{-4}	80,000
X	3.6	33.0	4.5×10^{-4}	75,000
U	2.0	28.3	2.6×10^{-4}	110,000
K	1.3	23.4	4.0×10^{-4}	59,000
Q*	0.7	20.9	2.5×10^{-3}	9,000

*image not used for spectra

Table 2.2:   Dynamic ranges for the VLA images

PLot file version 1 created 18-DEC-2000 12:27:06
3C273 IPOL 4994.000 MHZ C273MV-.14.IVTC.1

Peak flux = 1.2385E-01 JY/BEAM
Levs = 1.239E-03 * (3.125, 4.419, 6.250, 8.839,
12.50, 17.68, 25, 35.36, 50, 70.71, 100)
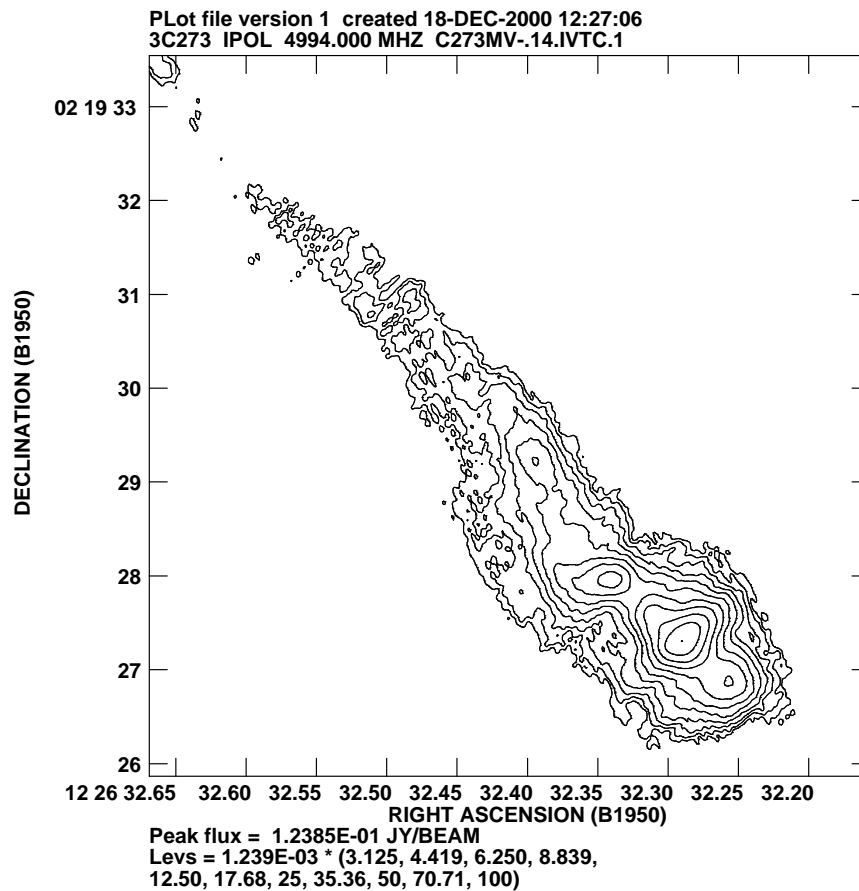
Figure 2.3:   Map obtained from a combination of MERLIN and VLA data at $\lambda 6\,\mathrm{cm}$.

merged with a coeval MERLIN data set provided by Simon Garrington and Tom Muxlow (Jodrell Bank Observatory). The joint data set was then imaged using a maximum entropy deconvolution method to obtain a map of the jet comprising the information on all the angular scales sampled by either telescope array.

The combined map is shown in Fig. 2.3. Although the map looks morphologically plausible, a comparison of the derived jet profile with the original VLA map shows a discrepancy between the two (Fig. 2.4). The total flux of the combined map is about 20% larger than that
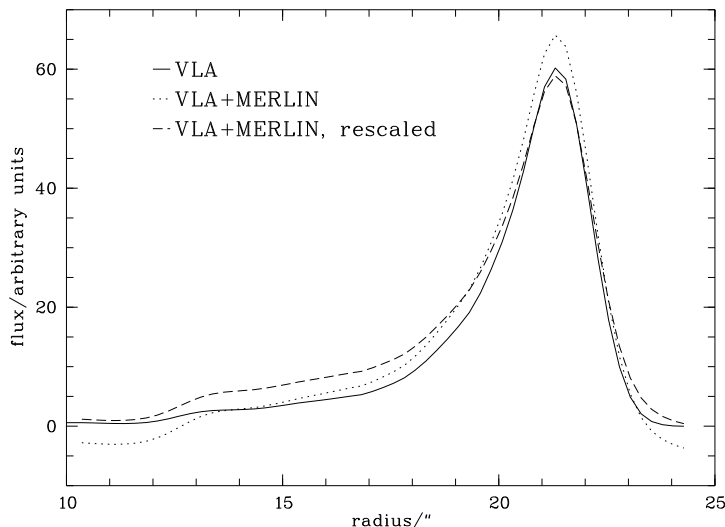
Figure 2.4: Comparison of VLA and VLA+MERLIN data sets at 1.″3 resolution. The jet profile derived from the combined data set is discrepant from that obtained from the VLA observations alone.

detected by the VLA alone, and the background flux level is not zero, but slightly negative. An attempt of correcting both background offset and flux normalisation error was made by adding a constant to the combined image to make the average background level zero, and then scaling the jet image to match the total flux contained in both. As Fig. 2.4 shows, after this normalisation attempt, the jet profile at 1.″3 resolution is still discrepant between both data sets.

The reasons for the discrepancy are not known in detail at present. They are most likely related to the strong north-south *sidelobes* of the MERLIN array. A two-element interferometer is sensitive only to structures perpendicular to its baseline. Unlike the VLA with its three-armed configuration, MERLIN has predominantly north-south baselines. Therefore, it best samples structures oriented in an east-west direction for sources observed on the meridian. For sources at large declinations, the Earth's rotation provides for a rotation of the baselines on the sky, so that all parts of the (u, v) plane (the "k-space" for celestial coordinates) are observed and structures at all position angles on sky can be detected (aperture synthesis technique). However, since 3C 273 is located near the celestial equator, the geometry of the MERLIN array leaves structures in extending in a north-south direction ill-constrained. Periodic positive and negative images of the jet appear to the north and south of the true image (Fig. 2.5).

Detailed investigations will be required to remedy the flux discrepancy between the VLA-only and the joint VLA-MERLIN data set. MERLIN data are also available at $\lambda 18$ cm, a combination with VLA data will only be attempted once the $\lambda 6$ cm data are understood.

2.2.2 WFPC2 data reduction steps

The WFPC2 data reduction for this project has already been described in detail previously [41], we therefore give an outline here and present details in App. A.

Images taken with CCD cameras have to undergo certain data reduction and calibration
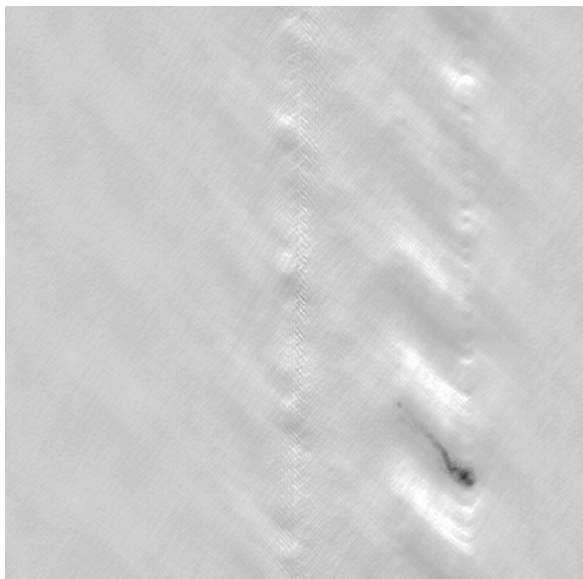
Figure 2.5: The MERLIN image at $\lambda 18$ cm shows the positive and negative images of the jet caused by the array's sidelobes.

steps to remove instrumental effects like variable detector response etc. After execution of a science program, the STScI provides pipeline-calibrated science frames along with the raw exposures. As some of the calibration information may have been updated since the pipeline calibration, the images have been recalibrated. The recalibration was done under IRAF, using the task `calwp2`, part of the STSDAS (Space Telescope Science Data Analysis System) package provided by the STScI. This task calls subroutines applying the individual calibration steps, using information saved with the data frames to determine the appropriate calibration files. The best calibration files were obtained from the STScI archive.

The value for background noise measured on the calibrated images agrees with the values expected from photon statistics, as calculated from the predicted sky background level, read noise and dark current (Sect. A.2). One of the PC chip's charge traps [105, 107] lies inside the jet image, in column 339. This has no observable effect on the faint UV image, but the effect had to be corrected on the well-exposed red-band images. This was done by replacing the affected portion of each image by the corresponding pixels from an offset image, as a correction according to Whitmore & Wiggs [107] unduly increased the noise gin the corrected part of the image (Sect. A.3.1).

The images were initially registered using the commanded offsets to the nearest pixels. This alignment is sufficient for the rejection of cosmic rays as these only affect a small number of adjacent pixels. Cosmic rays were removed using a standard κ-σ algorithm, rejecting all pixel values deviating more than 4σ from the local (low-biased) median in a first pass, and neighbouring pixels with more than 2.5σ deviation in a second pass (details in Sect. A.3.3). The number of pixels treated this way agrees with the expected cosmic ray hit rate for the images.

A model of the sky background and "horizontal smear" (increased pixel values in rows containing saturated pixels from the quasar's core [Chap. 4 of 8]) was fitted in the part of the image containing the jet using second-order polynomials along rows (Sect. A.3.4). The

Filter	PHOTFLAM	λ_p	PHOTFNU
F300W	6.137×10^{-17}	2981.9	17.72
F622W	2.789×10^{-18}	6187.5	3.531

Table 2.3: Photometric conversion factors. PHOTFLAM is in erg s^{-1} cm^{-2} Å^{-1}, λ_p in Ångstrøm, PHOTFNU is in μJy.

coefficients of the polynomials were then smoothed in the perpendicular direction.

The photometric calibration of the exposures was done using throughput information provided in [8, Section 21.1.2]. A synthetic photometry approach is employed: The combined response of telescope, filter and detector is used to calculate the physical flux producing a count rate of 1 count per second in each filter. A constant flux *per unit wavelength* is assumed. The proportionality factor between count rate and spectral flux density is given in units of erg s^{-1} cm^{-2} Å^{-1} and referred to as PHOTFLAM. For the later comparison with radio astronomical measurements, which are commonly calibrated in Jansky (Jy, 1 Jy $= 10^{-26}$ W m^{-2} Hz^{-1}), the count rate is converted to Jy. The corresponding conversion factor is called PHOTFNU. The spectral flux per unit wavelength f_λ is converted to a flux per unit frequency f_ν by noting that

$$f_\lambda \mathrm{d}\lambda = -f_\nu \mathrm{d}\nu \tag{2.1}$$

(The $-$ sign arises because dν increases in the opposite sense to dλ.) From here,

$$f_\nu = -\frac{\mathrm{d}\lambda}{\mathrm{d}\nu}f_\lambda = \frac{\lambda^2}{c}f_\lambda \tag{2.2}$$

This is, strictly speaking, only valid at a single value pair for λ and ν. In order to obtain the correct spectral flux density conversion, the shape of the input spectrum would have to be convolved with the throughput curves over the filter passband. As the input flux spectrum is not known *a priori*, but is rather what we are trying to measure, this would have to be done in an iterative fashion: calculate a flux density assuming a spectrum constant in each passband, fit a spectral shape to measurements across the whole spectrum, then refine the flux density conversion using the spectral shape in each passband.

It will be sufficiently accurate here to convert PHOTFLAM to PHOTFNU using the *pivot wavelength* λ_p of each filter used. The pivot wavelength is a characteristic of the filter. It is defined such that Eqn. 2.2 is true with $\lambda = \lambda_\mathrm{p}$, and f_ν and f_λ being obtained by integrating a spectrum with constant flux per unit frequency or wavelength, respectively, over the bandpass transmission curve (see Sec. 18.2.2 of [105] and [12]). The conversion between the different units of PHOTFLAM and PHOTFNU introduces a numerical factor of 10^{33} which arises from the definition of 1 Jansky and the conversion from erg to W (1W $= 10^7$ erg).

The best values for the count-to-flux calibration, filter pivot wavelength and resulting flux in Jansky for 1 count per second are given in Tab. 2.3.

The throughput information is accurate to 2% if the most recent values are employed. The flux calibration is done by dividing the data frames by the exposure time (to give a count rate) and multiplying by the respective value of PHOTFNU.

With this calibration, the background uncertainty parameters (Sect. A.3.4) scale to the following values: the background scatter in the summed U and R-band images are 0.35 nJy and 0.14 nJy, respectively. The residual background levels are 0.07 nJy and 0.01 nJy per pixel, respectively.
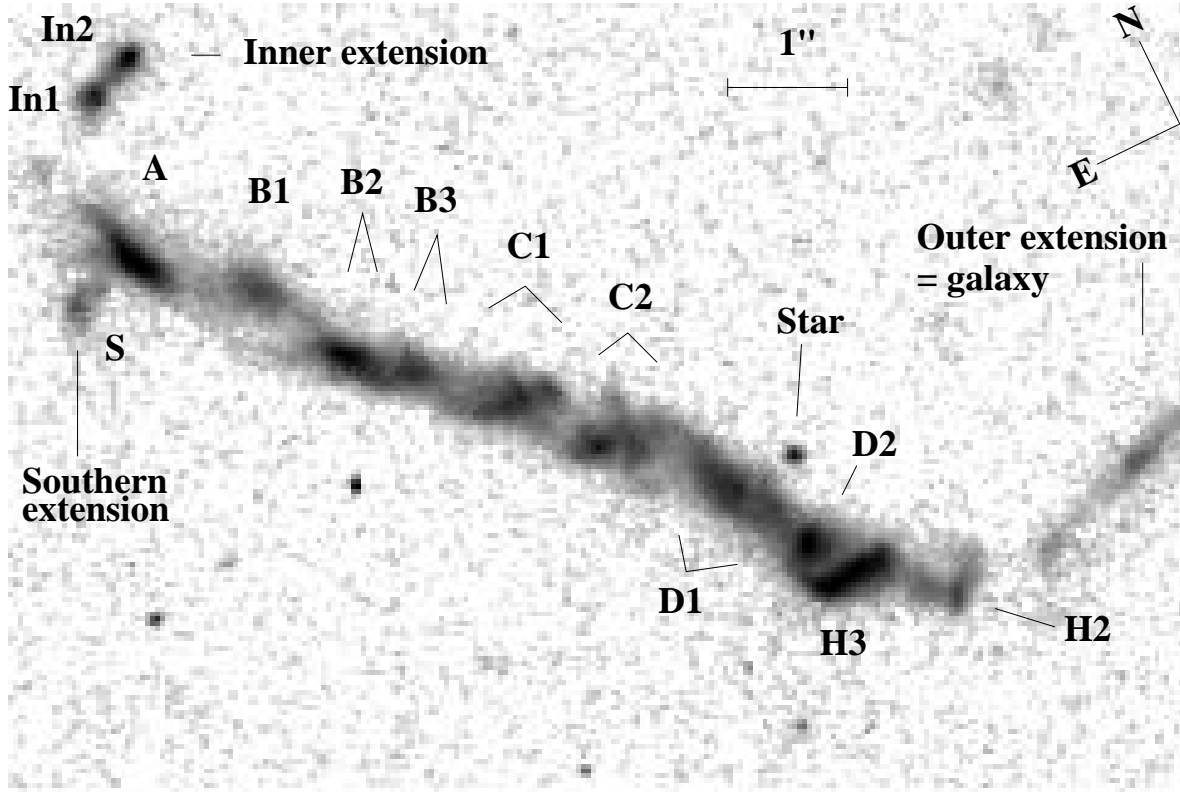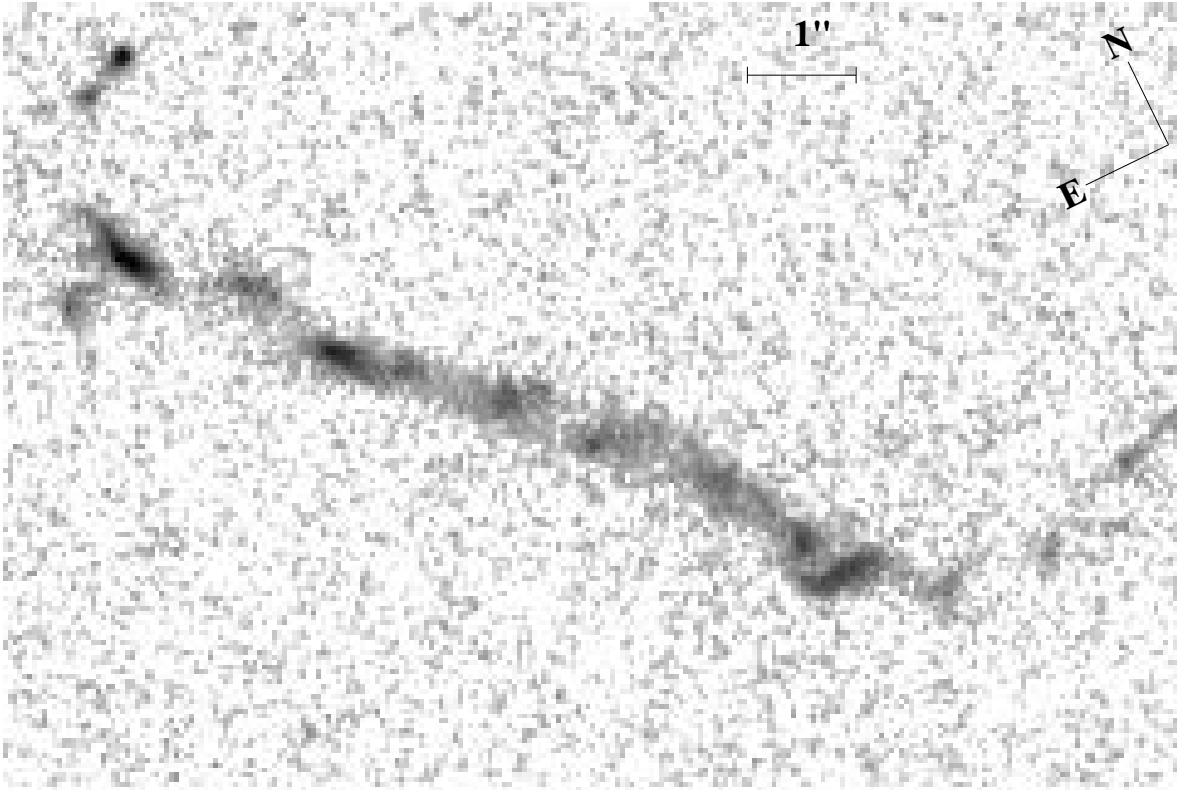
Figure 2.6:   The jet in red light (620 nm) after background subtraction. Logarithmic grey-levels run from 0 to 0.04μJy/pixel, $0\rlap{.}''08$ effective beam size, $0\rlap{.}''045$ pixel size. The quasar core lies $10''$ to the northeast from A. The labelling of the jet features as introduced by Lelièvre et al. [52] and extended by Röser & Meisenheimer [87], together with the hot spot nomenclature from Flatters & Conway [27] is also shown. Note that the labelling used by Bahcall et al. [4] is slightly different.

2.2.3   Maps of the optical brightness

The calibrated images are presented in Figs. 2.6 and 2.7. The morphology of the jet is identical in both images and appears rather similar to the morphology in high-resolution radio maps [4, 17]. The exception to this is the radio hot spot, being the dominant part in the radio but fairly faint at high frequencies (a detailed comparison follows in Sect. 3.2). Our images show structural details of the optical jet which were not discernible on earlier, shallower and undersampled HST WF images of $0\rlap{.}''1$ pixel size [4]. Based on our new maps, the term "knots" seems inappropriate for the brightness enhancements inside the jet, as these regions are resolved into filaments. The higher resolution necessitates a new nomenclature for the jet features (Fig. 2.6). For consistency with earlier work [27, 52, 87], our nomenclature is partly at variance with that introduced by Bahcall et al. [4].

The jet is extremely well collimated – region A has an extent (width at half the maximum intensity) of no more than $0\rlap{.}''8$ perpendicular to the average jet position angle of $\sim 222°$ (opening angle $\lesssim 5°$). Even H3 is only $1''$ wide along the longest axis (opening angle $\approx 2.5°$). The optical jet appears to narrow towards the hot spot, in the transition from H3 to H2.

Region A is now seen to extend further towards the core than previously known. It may be noteworthy that Lelièvre et al. [52] reported the detection of an extension of knot A towards

Figure 2.7: The jet in UV light (300 nm) after background subtraction. Logarithmic grey-levels run from 0 to 0.014μJy/pixel, $0''06$ effective beam size, $0''045$ pixel size.
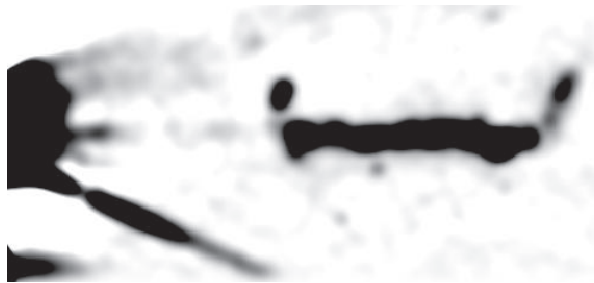
Figure 2.8: After modelling the background near the quasar and smoothing the 620 nm image to $0''25$, the faint inner $10''$ of the optical jet can be made out.

the quasar, whose existence at the reported flux level was not, however, confirmed by later work.

When considering a smoothed version of the summed WFPC2 image, a faint continuation of the optical jet can be made out (Fig. 2.8). Bahcall et al. [4] also reported a tentative detection of this inner jet. In order to establish a reliable detection, we have obtained deep optical images of the jet with the FORS1 instrument at the ESO Very Large Telescope (VLT). For these data, the accurate modelling and subtraction of the background due to the quasar PSF (with a seeing-limited FWHM of around $0''7$) turned out more difficult than anticipated and results will be reported elsewhere.

The criss-cross pattern visible in regions C1 and C2, and less clearly in B1-2 and D1, is reminiscent of a helical structure, possibly double [4], but could also be explained by oblique double shocks [31, *e. g.*].

The jet has three "extensions" (Fig. 2.6), none of which has been detected at radio wavelengths. The morphology of the *outer extension* supports the classification as a galaxy based on its colours made by Röser & Meisenheimer [87]. The nature of the other two extensions, however, remains unknown even with these deeper, higher resolution images. The *northern inner extension* was already resolved into two knots (In1, In2) on a FAINT OBJECT CAMERA image [102]. The two knots are extended sources and clearly connected to each other. The *southern extension* is featureless and an extended source.

Comparing the direct images, we can immediately estimate that the jet's colour slowly turns redder outwards from region A. The similarity of the jet images in both filters shows that there are no abrupt colour changes within the jet. A comparison of these images with those at other wavelengths is deferred to Sect. 3.2. A quantitative assessment of the jet's colour will be done through resolution-matched spectral index maps in Sect. 3.5.2.

2.2.4   NICMOS data

Like optical CCDs, today's near-infrared detectors make use of the photoelectric effect for the detection of light. The semiconductor material used in the NICMOS3 detector employed in the NICMOS camera is mercury cadmium telluride (HgCdTe), which has a band gap suitable for the detection of near-infrared photons. Unlike a conventional CCD, in which the accumulated charge is transferred out of the detector array for read-out, the NICMOS3 detector can be read out non-destructively. Apart from this, an IR detector has a bias level, dark current and readout noise just like an optical CCD, with analogous reduction steps. To keep the detector from detecting itself, it is cooled to liquid-helium temperatures.

Because of thermal emission from the sky (for the HST this is the solar-system dust emitting the zodiacal light) and the telescope itself, the background levels for IR observations are much higher than those for optical wavelengths. Ground-based telescopes suffer from still higher background levels than the HST: Firstly because the atmosphere is a much warmer emitter than zodiacal light, and secondly because of telluric absorption and emission ("airglow", mainly from from OH^- and O_2). The ground-based near-IR filter bands J, H, K at 1.2, 1.6 and 2.2 μm are designed to lie between the telluric features, but cannot avoid some of the numerous airglow lines. The present observations were carried out at a wavelength of $\lambda \approx 1.6\mu$m, at which the total background due to both sky and telescope is at a minimum for the HST [11]. They therefore constitute the deepest near-infrared exposure of any extragalactic jet so far.

Description of NICMOS data

Because of the possibility of non-destructive readout, *reset* and *read out* are two separate and independent operations for the NICMOS detector. While the reset clears all accumulated charge from the detector, it leaves the array at an uncertain bias level. The array is therefore read out immediately after the reset to obtain the so-called *zeroth read*, which is subtracted from subsequent reads to obtain the true signal. This means that the readout noise enters into each NIR image at least twice.

The total exposure time was split over three HST visits with 10 exposures each. Because

the field of view is not much larger than required to image the quasar and jet simultaneously, there is one short image of the QSO at the beginning of each visit. Each individual visit comprises five pairs of jet exposures, each pair placing the jet on a different part of the detector. One of the exposures per pair has exposure time 1024 s, the other 1280 s. There are additional interspersed sky exposures of 256 s.

Following a recommendation from STScI, the present images were taken in the so-called MULTIACCUM mode, with readouts 0, 0.3, 0.6 s after the reset, and then every 256 s. Using this exposure mode, it is in principle possible to obtain the count rate for each pixel by fitting the relation between observed counts and observing time. Cosmic ray rejection should also be facilitated, because a cosmic ray hit can be identified as discontinuity in the observed relation. However, with the exposure times and time step chosen, each data set consists of only 4 or 5 independent data points (there is hardly any signal in the first, very short readouts), so that full advantage of the MULTIACCUM capability was not taken.

A pipeline software `CALNICA`, provided by STScI, exists to carry out the reduction of all NICMOS data. This pipeline performs the following steps for the present data:

- correction of the zero*th* read for detectable signal (above 5σ) incurred in the 0.2 s elapsing between reset and read operation for each pixel

- subtraction of the zero*th* read from each subsequent read

- subtraction of an appropriate dark image from each read

- correction of detector non-linearity according to an empirical cubic relation

- correction of "bars" — small, noiseless bias changes in a pair of rows (one elevated, one lowered), replicated in all four quadrants

- flat-field correction

- linear fit to relation of counts vs. observing time, rejection of cosmic rays as outliers from the relation

A number of more or less subtle error sources which are not accounted for by this procedure meant the pipeline reduction was inappropriate for the present data. This necessitated the use of a tailor-made reduction procedure, whose outline follows the pipeline method. After a consideration of these error sources, we will turn to a description of the steps finally taken for the reduction of the NIC2 data.

Error sources for HST NICMOS

The main error source identified for our data is a spatial and temporal variation of the bias level. The true bias for each readout has three components [74]:

1. "normal" bias introduced at reset (different for each reset),

2. "shading", a readout bias which is variable with the time since last readout and the temperature,

3. and "pedestal", a random change of the overall bias level per quadrant.
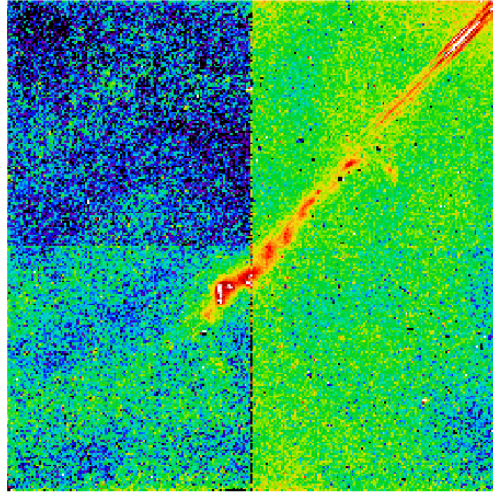
Figure 2.9: Example of the pedestal effect in NICMOS images. The background level differs clearly from quadrant to quadrant, and the flat-field imprint is visible.

The first component is taken care of by the usual zero*th* read correction. The "shading" means that in fact the readout bias varies over each quadrant, leading to a bias structure in each readout, while the overall level depends on the time since the last readout. In addition to this temporal bias variation, there is a temperature dependence, which may introduce an additional variation of the bias level from readout to readout. The variation of the shading from one read to the next is probably the cause of the third effect, the pedestal, which appears in the data as offset between the quadrants of one image (see Fig. 2.9).

The joint effect of the bias variations is that the pipeline calibration does not remove the true bias. Since the bias is additive and not subject to the sensitivity variations which are removed by flat-fielding, the pipeline-calibrated image contains an imprint of the flat-field structure: assume a total raw signal $I(x,y)$ was detected on the pixel with coordinates (x,y). This signal consists of the signal from the sky, $S(x,y)$, modulated by the sensitivity pattern $f(x,y)$, and the bias $B(x,y)$, which does not modulate with the flat-field. A wrong bias subtraction leaves a residual bias $\Delta B(x,y)$. After correction with the flat-field pattern $f^{-1}(x,y)$, which is the inverse of the sensitivity and assumed to be known, the calibrated signal $C(x,y)$ is

$$\begin{aligned}
C(x,y) &= f^{-1}(x,y)\left(f(x,y)S(x,y) + \Delta B(x,y)\right) \\
&= S(x,y) + f^{-1}(x,y)\Delta B(x,y),
\end{aligned}$$

so that the resulting image contains an imprint of the flat-field pattern, while the sky signal has been adequately corrected. As the sensitivity variations are rather large for NICMOS (factor of 5 at $0.8\,\mu$m, decreasing to near unity at $5\,\mu$m; *cf.* Fig. 2.10), the residual signal contained in the flat-field imprint is considerable. In addition, "pedestal" offsets between the detector quadrants remain. The offsets could in principle be removed *a posteriori*, by estimating a constant pedestal level per quadrant and equalising the quadrants, *e. g.* It is desirable, however, to correct the bias to the best available knowledge before flat-fielding,
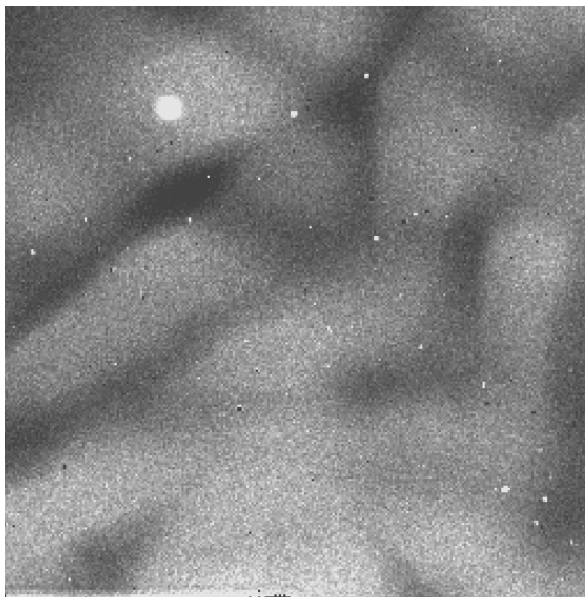
Figure 2.10: NIC2 flat-field frame. Oversensitive pixels appear dark, undersensitive pixels bright. Bright specks are "grot", pixels with drastically reduced sensitivity.

thus minimising the flat-field imprint, and before the estimation of count rates from the slope of the counts-vs.-time relation. One readout with a wrong bias level in the four or five reads per image may well lead to a systematically wrong slope, thus necessitating a multiplicative correction rather than an additive one.

A number of correction algorithms for this effect are publicly available. Their essence is to assume that the true signal should accumulate linearly with time, or that the true background is flat, and subtract a constant times the flat-field image from the calibrated data to optimise the result according to the chosen criterion. All of the available algorithms failed to remove both the offsets between quadrants and the flat-field imprint.

As the variation of bias level with temperature is systematic and reproducible, the NICMOS group have provided a tool to generate temperature-dependent dark files, with a bias appropriate for the readout temperature [75]. Using a temperature-dependent dark accurate to $0.05\,\mathrm{K}$ should significantly alleviate the bias problems. However, even using separate darks for each readout to $0.01\,\mathrm{K}$ did not improve the quality of the reduced images.

In addition to the bias variations, column number 128 is known to be a "bad column" with a bad bias level and elevated noise. A number of small patches of few pixels with reduced sensitivity is known as "grot". Finally, the sensitivity of the NICMOS pixels is wavelength-dependent, in the sense that oversensitive pixels are less sensitive toward longer wavelengths, while undersensitive pixels have sensitivity increasing with wavelength. This colour dependence of the flat-field means that the calibration flat-field taken with an internal lamp may not be appropriate for the sky or object. The effect of the colour dependence for the sky background is most severe at wavelengths above $1.8\,\mu\mathrm{m}$ [74]. Its effect is difficult to disentangle from the flat-field imprint left by the bias problems.

To remove these effects, in particular to remove the bias offsets between the quadrants, a completely different reduction procedure was implemented.

Reduction steps taken

In order to remove the calibration errors, some assumptions had to be made about the data and calibration errors, to constrain the number of free parameters. In particular, we assumed that

- the internal lamp flat-field is correct,

- the true sky is flat,

- the dependence of the bias level on the time since the last readout is identical for all images,

- and any further bias variation between readouts is constant within each quadrant of every readout.

We assume that the flat-field is correct because of the lack of information of what the true flat-field otherwise would be. The sky exposures which have been taken are short compared to the science exposures and hence noisy, and of course suffer from the same bias uncertainties as the object exposures. They provide therefore no further information. The assumptions about the bias have to be made to allow any correction at all.

The reduction algorithm's essence is the joint estimation of all background signals (dark current, sky signal, bias) after filtering the object (and cosmic ray) signal out of the frames. This is possible because the exposures were taken at 15 different pointings. Hence the object moves around the detector, while the background remains unchanging. Each readout has a different bias. All readouts are therefore grouped into image cubes of readouts with identical exposure time. For each detector pixel, the background signal is estimated as the lower quartile value of the distribution of pixel values across the readouts (30 with exposure times up to 1024 s, 15 with exposure time 1280 s). The lower quartile value is chosen rather than the median to be sure to exclude object and cosmic-ray signal (see Fig. 2.11). The background value is subtracted from the pixel in each readout.

This step has removed the dark current, sky background and all bias components which vary across the detector, but not from readout to readout. To remove the quadrant-to-quadrant variations, the residual background in each *quadrant* of every readout is estimated and subtracted separately. The background estimation is done using the routine `MMM` within the `IDL` graphics package. This initially rejects positive outliers (signal, cosmic rays) from the pixel value distribution by κ-σ-clipping and uses the mean as sky estimator, or $3\times$median - $2\times$mean if the mean is larger than the median.

As a drawback of this approach, the knowledge of the sky and dark signal cannot be used to estimate the noise associated with these components any more. Instead, the noise will be estimated in the photometry by considering the scatter of calibrated pixel values. After the treatment, the images contain only the object counts and cosmic rays. The quadrants are reassembled into images, and the count rate is estimated and cosmic rays are rejected using the program `fullfitbam` [59], which is derived from that used by the pipeline software `CALNICA`.

Since the cosmic ray rejection by `fullfitbam` is not perfect, a cosmic-ray rejection using `COSMIC/MEDIAN` (as for the optical data) is performed. For this step, all known bad pixels (including column 128 and the "grot" mentioned above) and negative outliers are flagged so that they, too, are replaced by the cosmic-ray rejection routine. The resulting images are not
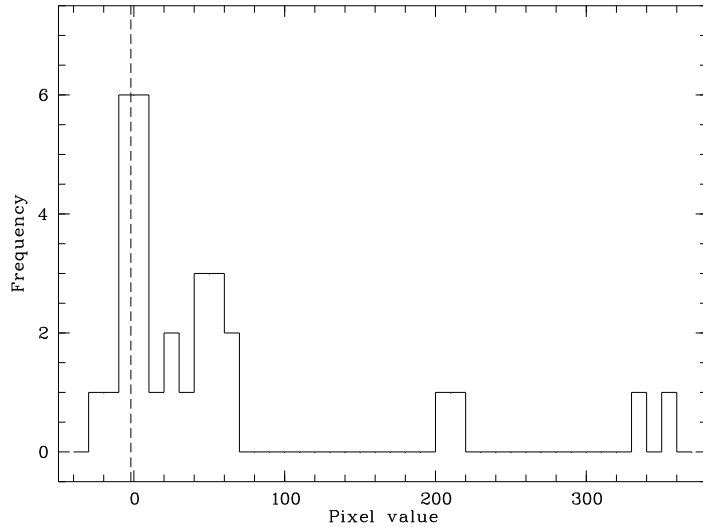
Figure 2.11: Sample histogram of the pixel value distribution in one read over thirty exposures. The dashed line shows the lower quartile value chosen as background value for this pixel.

perfectly flat individually, suggesting that there may be a residual problem with the flat-field. As mentioned above, no attempt is made to correct this.

2.2.5 Maps of near-infrared brightness

Figure 2.12 shows the reduced and summed images. The overall appearance of the jet's morphology at $1.6\,\mu$m is very similar to that at visible wavelengths. The only significant difference (apart from the second star) is the precursor H1 to the hot spot H2 which is visible on the infrared image, but not at higher frequencies. The jet is equally well collimated. All optical extensions are present, the southern extension S here being much weaker than knot A. In1 is clearly brighter than In2, indicating a marked colour difference between the inner extension's two knots.

It is clear from Fig. 2.12 that there is signal from one of the quasar's diffraction spikes superposed on the jet emission. In addition, an IR-bright object is located within the jet, close to the faint star just north of the jet detected on the optical image (Fig. 2.6). Both represent a contamination of the true jet emission that needs to be modelled and removed.

Removal of the diffraction spike

In 20 of the total of 30 exposures, one of the quasar's diffraction spikes caused by the spider holding the secondary mirror passes through the jet image. The signal from this diffraction spike has to be modelled and removed. In principle, this should be done on each frame individually, as the structure of the diffraction spikes changes with time because of movements of the NICMOS cold aperture mask with respect to the telescope aperture [44]. However, the variations cannot be modelled well, nor is the spike visible with sufficient signal-to-noise on individual images to allow a straightforward fitting and subtraction. Instead, the spike is modelled on the sum frame made from the remaining 10 exposures, in which the telescope has been rotated by 4° compared to the previous ones (Fig. 2.12, right). To do the modelling,
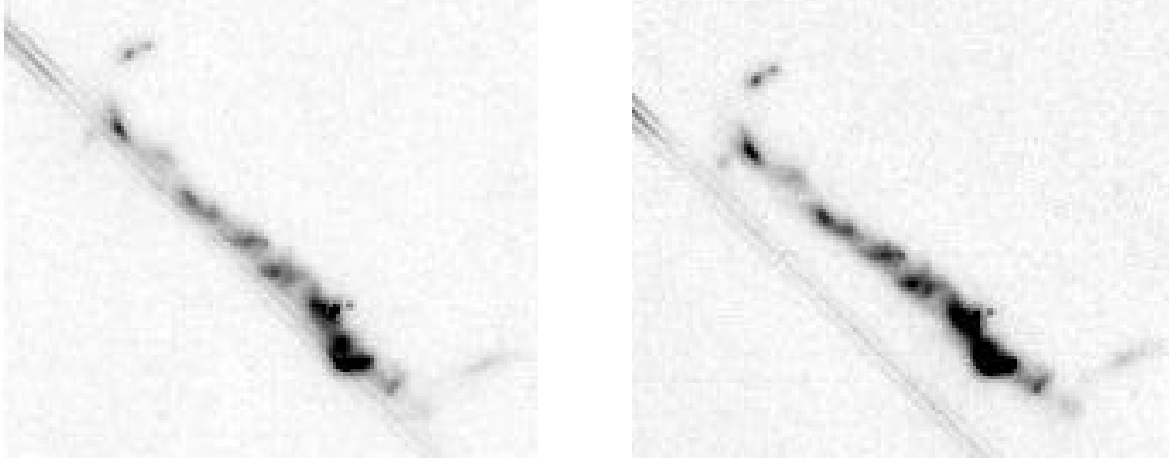
Figure 2.12:   Sum of reduced NICMOS images, rotated to approximate the orientation on sky. Left, spike along the jet. Right, telescope rotated by 4°, spike clear of the jet
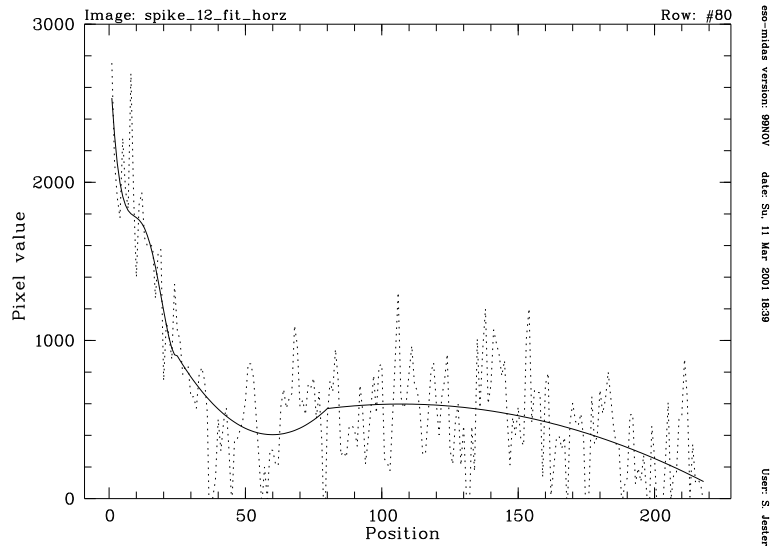
Figure 2.13:   Example of NICMOS diffraction spike modelling. Dashed line, diffraction spike image after subtraction of jet images in Fig. 2.12 from each other. Solid line, polynomial fit in three sections, of degree four, two, and two, respectively.

the images in Fig. 2.12 are scaled, rotated and subtracted from each other, removing the jet image and leaving only the spike. The resulting image is rotated so that the spike is along rows of pixels, and the eight rows corresponding to the spike image are extracted. Each row is then divided into three sections, each of which is modelled with a polynomial of degree four, two, and two, respectively, the polynomials' levels being matched at the boundaries (see Fig. 2.13 for an example). The resulting diffraction spike model is rotated back to the original position angle. A scaled version of the model is then subtracted from each individual image. The result of subtracting the spike model from the "contaminated" sum frame is shown in Fig. 2.14.

The background subtraction is done on the individual frames, using the same fitting
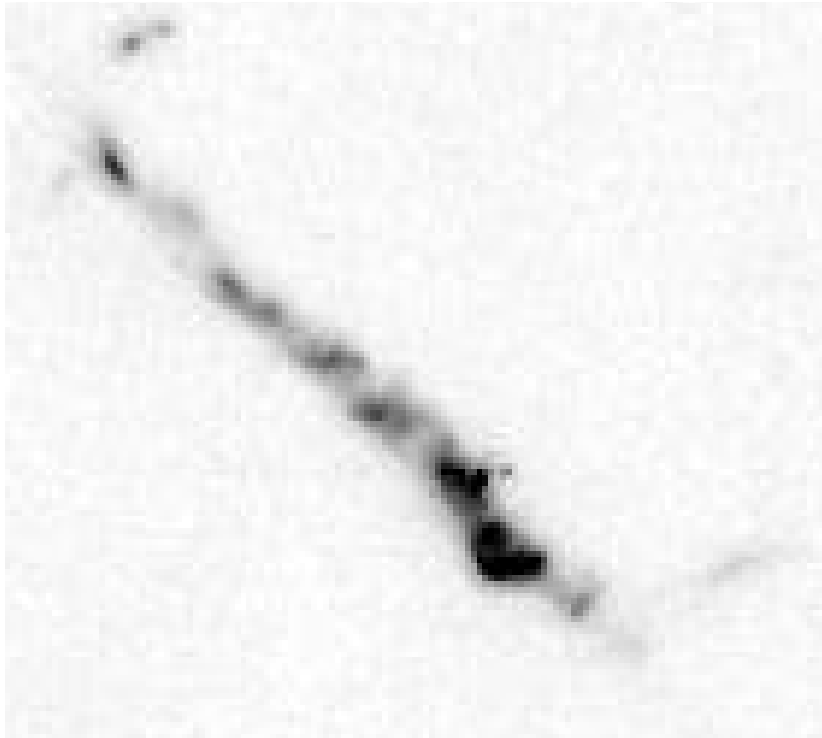
Figure 2.14: Result of subtracting the spike model from the "contaminated" frame in Fig. 2.12.

method as for the optical images (*cf.* Sect. A.3.4). As mentioned above, the NICMOS field-of-view is much smaller than that of WFPC2, therefore we have to rely on the jitter file information for image alignment, again tieing together the different visits with the quasar as reference point.

Removal of the star

Figure 2.15 shows the location of an IR-bright star-like object conspicuous inside the jet. It is situated just south of the faint star which was already detected on the R-band image (Fig. 2.6). To assess whether it is a foreground star (of very late spectral type) or a jet feature, we fit a Gaussian along the row and column intersecting at the object's brightest point, on both sum frames (Fig. 2.12). The FWHM fitted on the first and second sum image along the detector rows is 2.6 and 2.0 pixels, respectively, and along the columns, 2.3 and 1.8 pixels, respectively, the PSF width being 2.2 pixels.[3] This suggests that the object is a star. To check that the increased IR flux is not correlated with any jet features, in which case it would be the most interesting part of the IR jet, we compare the IR image with a radio map at $\lambda 3.6$ cm (Fig. 2.16). There are no features corresponding to the star on either the total or polarised intensity radio map. The source is thus most likely a foreground star, and its flux has to be modelled and remove to avoid a contamination of the synchrotron spectra.

To model the star, all individual images are sampled on a 10× finer grid and summed up, applying the shifts from the jitter files to each image. This is done once for each set of

---

[3]The discrepancies arise because of the shift-and-add procedure and because a Gaussian is a bad approximation to a diffraction-limited PSF, which is a sinc function.
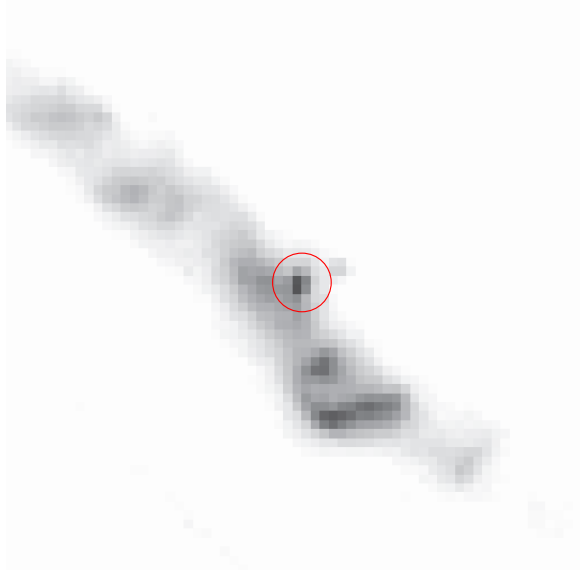
Figure 2.15:   Location of IR-bright object (star) inside the jet
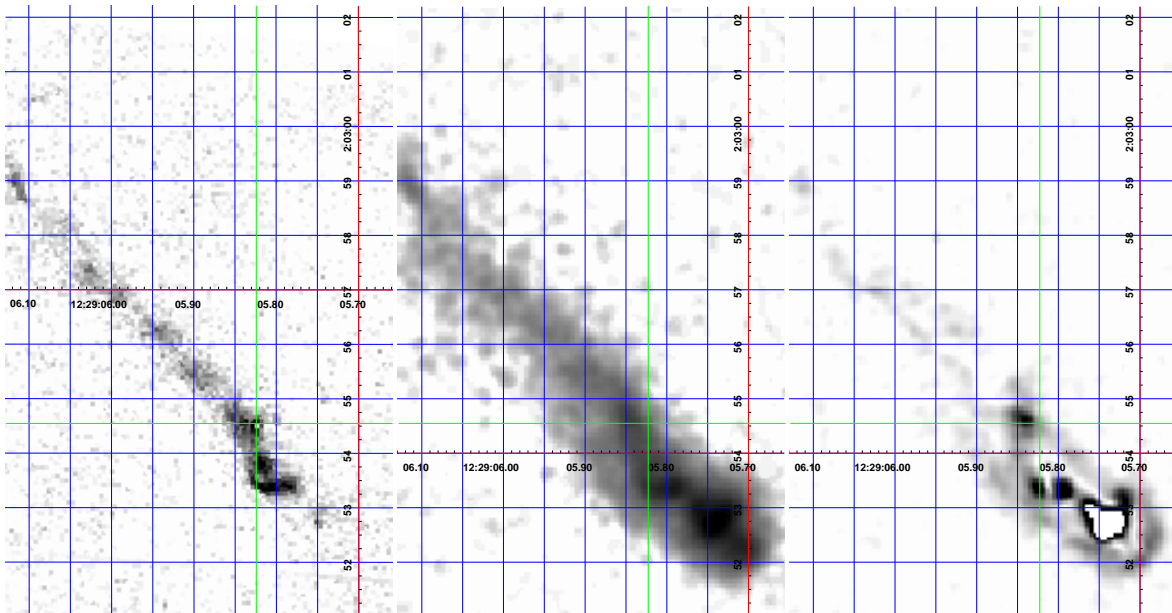
Figure 2.16:   Comparison of NICMOS image (left) with radio data at $\lambda 3.6$ cm, showing total intensity (middle) and polarised intensity (right). The location of the IR star is marked with green cross-hairs. There are no corresponding features on either radio map.

images with a different rotation angle. The star is fitted on the summed images as a Gaussian, allowing a sloping plane to account for the flux due to the jet. An appropriately scaled version of the Gaussian is then subtracted at the star's location on each frame, to yield the input frames for the photometry. The total signal removed was $2.0\,\mu$Jy, slightly under 2% of the jet's total infrared flux of $(129.29 \pm .03)\,\mu$Jy.

Chapter 3

Photometry

The aim of the project is the determination of the synchrotron spectrum over the entire jet. To this end, beam-matching aperture photometry is performed at a grid of positions covering the jet, using the `MPIAPHOT` software. Having completed the data reduction, we consider the steps necessary to obtain flux measurements, which are performed independently and individually on each of the detector frames at all wavelengths.

As we will take the ratio of fluxes measured on different exposures, great care has to be taken to refer all flux measurements to identical positions on the sky. Just how careful the alignment has to be depends on the accuracy of the required measurement, either direct flux measurement or here, the ratio of two flux measurements. Any misalignment should not contribute more than about 5% to the flux error. The analysis of all jet images is only possible on images whose resolution is matching each other. This final resolution will be $0\rlap{.}''3$, slightly inferior to that of the data set with the lowest resolution, the 3.6 cm radio data imaged with $0\rlap{.}''25$ resolution (*cf.* Sect. 2.1.1).

The flux measurements are done using the "weighted summation scheme" [60, 65, 87, 100] instead of the traditional aperture photometry. Rather than integrating all counts within a fixed-size aperture and discarding all counts outside it, this scheme places a Gaussian weighting function at the aperture position, and the *weighted counts* are summed up. This corresponds to measuring the flux on an image which is the result of the convolution of the original image with the Gaussian. This method has the following advantages:

- The width of the weighting function can be chosen to maximise the S/N of a point source flux measurement on a single exposure. In this case, the weighting functions should have the same width as the image PSF [100].

- When considering a set of images, the weighting function's width can be chosen for each image such that the result is a convolution to the desired common resolution, *i. e.,* the weighted summation does not require the computation of an explicit convolution.

- The weighting function can be centered anywhere to a fraction of a pixel so that a flux measurement at an arbitrary position is possible.

- There are no edge effects which occur at the boundary of a fixed-size "top hat" aperture due to bad sampling of the aperture's boundary by the image pixels.

A suitable choice of the scaling of the weighting function ensures that the sum of counts over the weighted aperture obtains the correct flux measurement. The scaling factor can

be chosen to yield the total flux of one given source from a single measurement, provided the surface brightness profile of the source is known. This is usually only the case for point sources. Its profile $f(x, y)$ is represented by a Gaussian

$$f(x, y) = \frac{S}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

of width σ and total flux or brightness S. For this profile, the summation with a weighting Gaussian with scale factor W is the integral

$$F \;=\; \int_\infty^\infty \int_\infty^\infty \frac{S}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{\sigma^2}} \times W e^{-\frac{x^2+y^2}{\sigma_w^2}} \, \mathrm{d}x \mathrm{d}y \tag{3.1}$$

$$=\; S \frac{\sigma_w^2}{\sigma^2 + \sigma_w^2} W \stackrel{!}{=} S. \tag{3.2}$$

So the choice

$$W = 1 + \left(\frac{\sigma^2}{\sigma_w^2} \right) \tag{3.3}$$

makes sure that the weighted sum is indeed the total flux of the point source. For an extended source, the result corresponds to the flux per Gaussian beam with width given by $\sigma_{\mathrm{eff}}^2 = \sigma_{\mathrm{w}}^2 + \sigma^2$ with unit response at the centre, centered at the aperture position, which is identical to the convention used for radio astronomical maps.

Using this weighting function, the total flux from an *extended* source can only be calculated by adding up the contributions from individual apertures if these are placed in a regular grid covering the entire source. In analogy to radio interferometric images, one obtains an image of the source at the chosen resolution, and sampled with a pixel size corresponding to the grid point spacing. To compute the total flux, the flux per beam (determined by the photometry) is divided by the effective beam area $A_{\mathrm{beam}} = \pi\sigma_{\mathrm{eff}}^2$ to yield a pixel-by-pixel measure of the surface brightness of the source, averaged over each aperture. The surface brightness is then multiplied by the pixel area, given by the grid spacing, and added up over all pixels to yield the total source brightness. This corresponds to integrating the source's surface brightness distribution over the source area.

When a determination of the total flux is not required, but only a comparison of measurements at different wavelengths at a given position in a source, the apertures can be placed and spaced arbitrarily as only surface brightness *ratios* matter which remain unchanged by the weighting procedure.

While only one image per wavelength has been created from the VLA data, there are 30 individual NICMOS frames, 4 from WFPC2 at 620 nm, and 14 from WFPC2 at 300 nm. The reason for obtaining photometry on the individual frames and averaging the results, rather than doing photometry on a summed frame, are twofold. Firstly, the information about the scatter of photometry results across the frames is conserved, so this information can be used to check the error estimate obtained otherwise. Secondly, non-integer pixel offsets between individual frames and the need to correct for geometric distortion always necessitate special care to avoid errors from rebinning or interpolating between observed data points. Since the photometry routine allows placement of apertures at arbitrary positions with sub-pixel accuracy, incorporating arbitrary offsets including the correction of geometric distortion is simple, and adding up photometry results is straightforward and free of errors. The next session discusses the positioning of apertures, making use of the knowledge about HST telescope offsets summarised in Sect. B.2.

3.1 Positioning of apertures

For the examination of spatial variations of the synchrotron spectrum, we consider the systematic errors on the observed spectrum which may be introduced when combining two images taken through different filters, or by different instruments and telescopes. The main danger in a determination of spectral properties lies in a misalignment between the images which would introduce spurious spectral gradients, the largest error being introduced near steep flux gradients; referring the flux to different effective beam sizes will lead to wrong spectral index determinations as well. Following considerations given in Sect. B.1, we deduce that a 5% limit on the flux error due to misalignment requires aligning the images to better than 10% of the effective PSF full width. The $0\rlap{.}''2$–$0\rlap{.}''3$ effective resolution aimed for therefore requires a knowledge of the *relative* alignment of all images in the data set to better than 20 mas. The *absolute* telescope pointing does not need to be known for this purpose as we can tie all positions to the quasar core as origin. As shown in Sect. B.1.2, the error in the PSF determination is negligible when the smoothing Gaussian is much wider than the PSF (as is the case for the optical and ultraviolet images). A 10% error in the determination of the PSF width still only leads to a 5% flux error when the smoothing PSF is of the same width as the intrinsic PSF. By definition, the PSF is a perfect Gaussian for the radio images. For the HST images, the smallest smoothing occurs for the NICMOS images with resolution of $0\rlap{.}''17$, which needs a smoothing PSF of $\sqrt{0\rlap{.}''3^2 - 0\rlap{.}''17^2} \approx 0\rlap{.}''25$. For this smoothing, a 3% error in the intrinsic PSF full-width produces a 5% flux error. We neglect this error as well as the effect on the photometry caused by the input PSF not being a Gaussian, as is the case for the NICMOS images (see Fig. 3.1).

The grid of photometry aperture positions is defined by offsets on sky relative to the quasar position, in a left-handed coordinate system with y-axis parallel to the North direction at the quasar. Coordinates in this system then have to be transformed into detector coordinates for each individual exposure, allowing for the geometric distortion in the case of imaging data. The image generated from the VLA radio synthesis data is provided as `FITS` image and includes information on the projection between sky coordinates and image coordinates. Positioning the apertures is therefore slightly different for the VLA data, and we will first consider the placement for the HST data.

We are given the detector position (x_Q, y_Q) of the quasar on a frame pointed at celestial coordinate (α_Q, δ_Q). We are also given a position vector relative to the quasar of a desired aperture in the plane of the sky (X_A, Y_A). What is the pixel position (x'_A, y'_A) corresponding to this aperture on a frame pointed at (α_A, δ_A)?

The first step will be the correction of the quasar's position for geometric distortion, similarly the last step will be the inverse correction for the derived aperture pixel position. The intermediate steps involve subtracting the telescope offset from the desired aperture offset to obtain the offset on chip between quasar and aperture position.

These steps require the following items of information:

1. the position of the quasar image on a reference frame,

2. the offset between the "current frame" and the reference frame (which need not be distinct frames),

3. the geometric distortion correction.
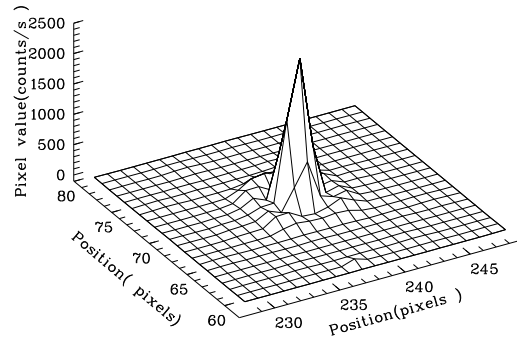
Figure 3.1:   The surface plot of the observed NICMOS Point Spread Function shows the central peak and first diffraction ring. Pixel size is $0.''076$.

Before describing in detail the calculations, we consider in turn how the three items above are obtained.

3.1.1   Measurement of the quasar position

The quasar position is straightforward to obtain. For the WFPC2 images, the position is determined on the short exposure of each visit/filter, as described above (Sect. B.2.2). For NICMOS, one short exposure of the quasar has been obtained at the beginning of each of the three visits. The quasar position is measured on these by the usual `CENTER/GAUSS` method in `MIDAS`. One may be concerned here that fitting a Gaussian is not appropriate to determine the centroid of a diffraction-limited PSF (a sinc function), of which the NICMOS PSF is quite a good example (Figure 3.1). To check the accuracy of this centroid measurement, we created a synthetic NICMOS PSF for the nearest integer pixel position using `TinyTIM` [46], in order to obtain the correct field-dependent PSF. We then used a cross-correlation to find the shift between the observed and the synthetic PSFs. The positions found in this way agree with the `CENTER/GAUSS` to better than 0.5 pixels. The measured position is then corrected for detector distortion as below (Sect. 3.1.3), to obtain the "ideal" pixel position of the quasar image.

3.1.2   Offsets between frames

A detailed discussion of the alignment procedure has been presented previously [41]; details omitted in the description here can be found in App. B. For the WFPC2 images, offsets between visits are obtained using point source measurements on the short exposures in each telescope visit. Information from the "jitter files", engineering data provided with the observations, is used to determine offsets within one visit.

Unlike for WFPC2, there are no centroidable point sources on any of the NICMOS jet frames. We therefore need to rely on the jitter file information to obtain all the offsets. As discussed in Sect. B.2.2, these provide the telescope pointing in right ascension and declination, corrected for the location of the respective instrument in the focal plane. The jitter file offsets are thus given in celestial coordinates. In principle, rotations between the reference frame and the "current frame" could also be taken into account. Fortunately, the rotations within one visit are negligibly small, leaving only translations.

3.1.3 Geometric distortion

The geometric distortion correction is done using the `traxy`/`tranback` routines (for the forward/backward correction) from the `drizzle` package, running under `IRAF`. This uses the wavelength-dependent cubic distortion correction for WFPC2 as determined by Trauger et al. [103], and the quadratic NICMOS coefficients obtained by Cox et al. [19]. The routine converts "ideal" detector coordinates into physical detector coordinates. For NICMOS, this routine also includes the effects of the difference in plate scale along the x- and y-direction. This requires using the geometric mean of the x and y plate scales when calculating offsets on the detector (see below). Since the plate scale difference is small ($< 1\%$, see Sect. 2.1.3), we neglect the change in aperture area induced by it and only ensure the correct positioning.

3.1.4 Calculation of HST aperture positions

We are now provided with sufficient information to transform offsets on sky relative to the quasar to any of the observed frames. The geometric distortion correction is straightforward, we therefore only need to consider the calculation of the offset between the quasar pixel position (x_Q, y_Q) and the aperture pixel position (x'_A, y'_A). In principle, the calculation of arbitrary offsets needs to take into account the effects of projecting the spherical celestial coordinate system onto the plane system of the telescope's focal plane and hence the detector. In particular, at different telescope pointings, the tangent point between celestial sphere and image plane has shifted. Therefore, the general transformation between different images is not a linear coordinate transformation. However, we can use some approximations for the following reasons:

1. 3C 273 is located near the celestial equator, so that the projected coordinate system is nearly flat,

2. all required shifts and fields of view are small, at most of the order of the jet length of $20''$, so that we may expect to be able to use a single local projected coordinate system,

3. the jitter files provided with HST data provide the telescope pointing not of the telescope's optical axis, but already corrected for the position of the respective instrument in the focal plane.

We will therefore make the assumption that the celestial coordinates $\mathbf{C} \equiv (\alpha, \delta)$ are related to detector coordinates $\mathbf{X} \equiv (x, y)$ by the simple linear transformation

$$\begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix} = \Sigma \begin{pmatrix} \Delta\alpha \\ \Delta\delta \end{pmatrix},$$

Σ being given by

$$\Sigma = \frac{1}{S} \begin{pmatrix} -\zeta\cos\theta & \sin\theta \\ \zeta\sin\theta & \cos\theta \end{pmatrix}, \tag{3.4}$$

where S is the detector scale in arcseconds/pixel (the geometric mean of the different values in x- and y-direction for NICMOS, see above), $\zeta \equiv \cos\delta_0$, the cosine of the declination (to account for the foreshortening of circles of constant declination towards the poles), and θ is the position angle of the detector's y-axis, measured in the usual sense (North through East). The next-order neglected terms are of order 3 parts in 10^6.
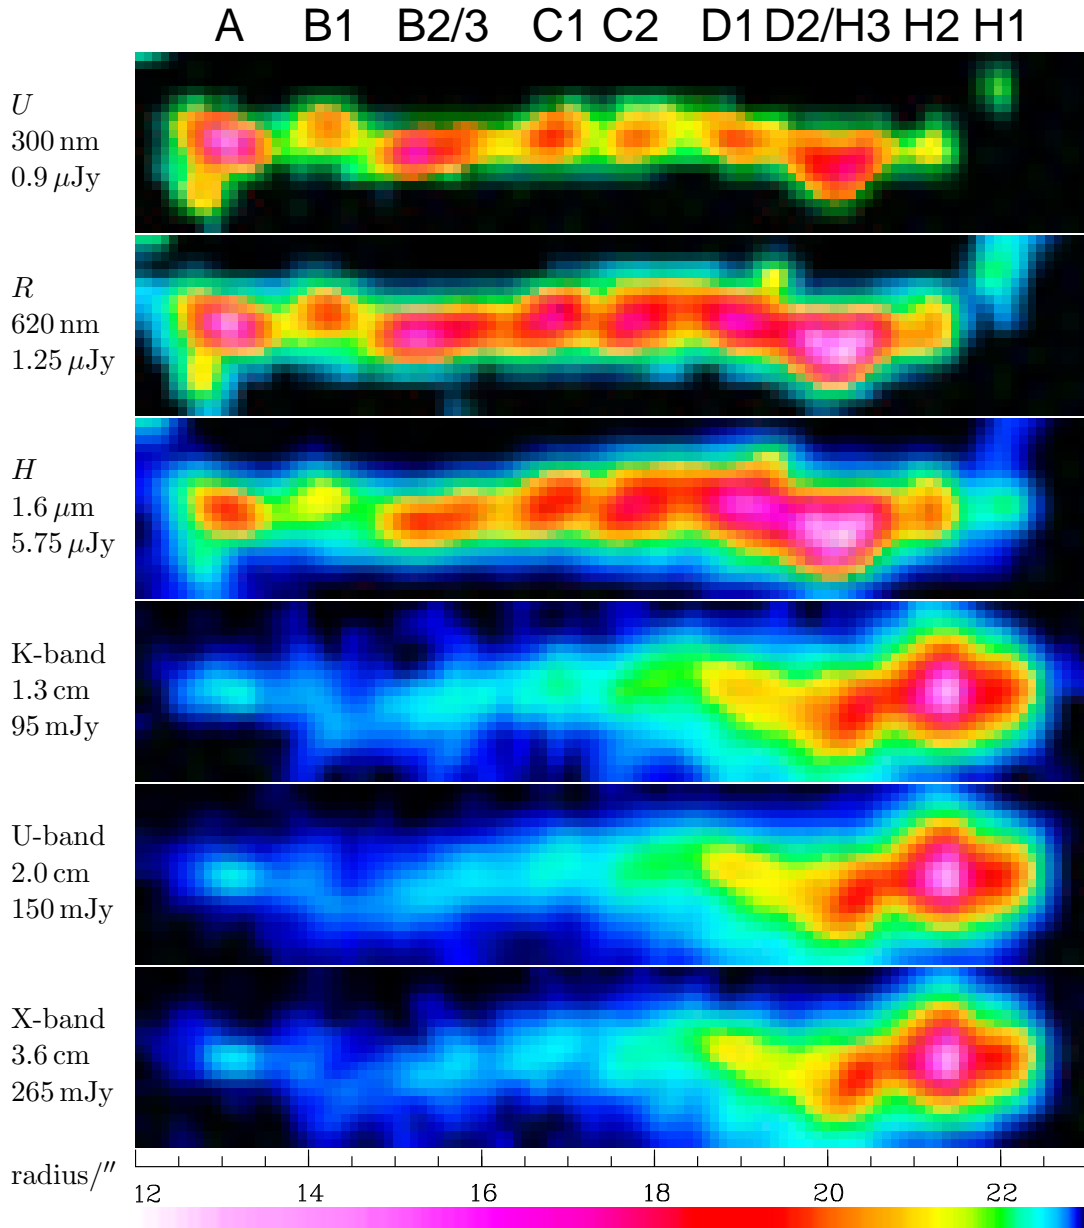
Figure 3.2:   Photometry of the jet in 3C 273 at 0″.3 effective beam size. Clipped to show only measurements with aperture signal-to-noise ratio > 5. Colour-levels are logarithmic, running from 0 to the peak flux/beam as given. The offset of 0″.2 between radio and optical hot spot position can be made out clearly.

The desired aperture offset from the quasar is converted to an offset in celestial coordinates (α, δ) from which the telescope offset for each frame is subtracted. The result is the offset between the quasar pixel position (x_Q, y_Q) and the aperture pixel position (x'_A, y'_A), only expressed in celestial coordinates. The pixel offset is then obtained using Eqn. 3.4.
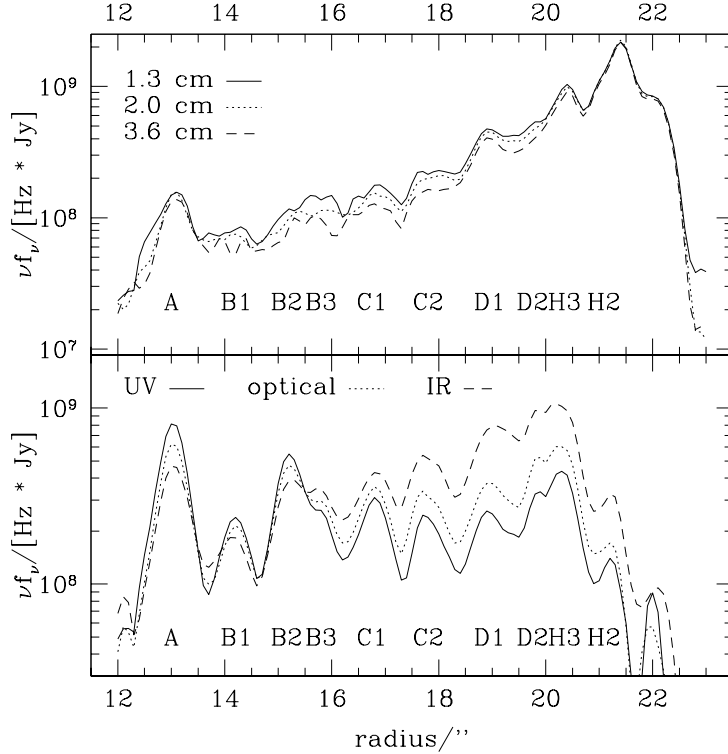
Figure 3.3: Plot of surface brightness along the jet's ridge line, *i.e.,* showing the brightest point per column from Fig. 3.2.

3.1.5 Calculation of VLA aperture positions

All VLA images are made employing a sine projection of the celestial sphere with the quasar position as tangent point, and in the "North up, East left" convention. Again using the approximations above, with similarly small errors of a few parts in 10^6, the (α, δ) are straightforwardly converted to pixel offsets using Eqn. 3.4 with $\theta = 0$.

After the photometry aperture positions have been determined for all individual frames, we proceed with the photometry as such.

3.2 Jet images at matched resolution of 0.''3

In order to examine the synchrotron spectrum everywhere in the jet, we use a rectangular grid of aperture positions. The grid extends along position angle $222°2$, starting at a radial distance of $r = 12.''0$ from the quasar and extending to $r = 23.''0$. Perpendicular to the radius vector, the grid extends to $\pm 1.''0$. Individual grid points are spaced $0.''1$ apart, yielding a good sampling of the $0.''3$ effective resolution, so that there are 111 radial grid positions and 21 on lines perpendicular to the radius vector, *i.e.,* 2331 in total. In order to compare the images at different wavelengths, the photometry results are reassembled into the images shown in Fig. 3.2. Figure 3.3 shows the brightness profile along the ridge line, that is, the brightest pixel per column from the images in Fig. 3.2. Before considering the spectral information qualitatively, we compare the jet morphology at different wavelengths.
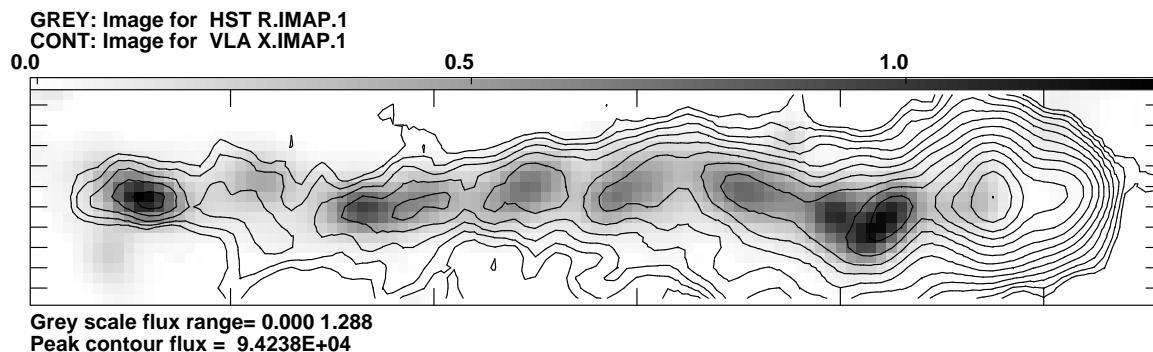
Figure 3.4:   Comparison of optical (grey-scale) and radio ($\lambda 3.6$ cm, contours) images at common resolution of 0.''3. Flux values are given in μJy/0.''3 beam. Contours are logarithmic in $\sqrt{2}$.
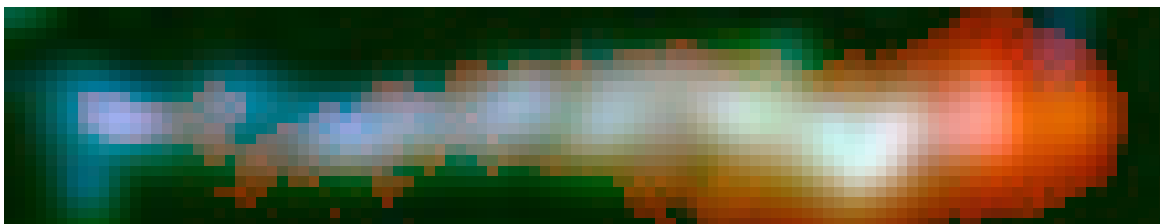
Figure 3.5:   Colour composite image at 0.''3 resolution, generated from a radio image (red channel), near-infrared (green) and optical (blue).

Comparing images at different wavelengths

The comparison of morphological features of the jet at different wavelengths can be done in a number of ways. The most straightforward one is a simple comparison of the images in Fig. 3.2. To ease the comparison, one may produce an overlay of a contour plot over a greyscale image, such as Fig. 3.4. On the other hand, differences in morphology between different wavelengths lead to spectral index gradients, so that a spectral index map (considered further down in Sect. 3.5.2) ultimately reveals morphological differences. As an attempt to compare the appearance of the jet at more than two wavelengths, one may combine three images by assigning them to the red, green and blue channels of a single colour image, such as Fig. 3.5.

The direct images (Fig. 3.2) and ridge line tracings (Fig. 3.3) show that the brightness profile tends to invert from short to long wavelengths: at 300 nm, the jet is brightest at the inner end, while towards longer wavelengths, the outermost regions are brightest, until finally at radio wavelengths, the hot spot is the brightest location. The optical/UV jet appears to terminate entirely there. In contrast to this, the radio and infrared images show an elongated precursor (H1) to the hot spot H2.

Figure 3.6 shows an overlay of a radio map and the 620 nm image of the radio hot spot H2. The optical counterpart to the hot spot lies 0.''2 closer to the core than the radio hot spot. This offset has been predicted by Meisenheimer & Heavens [63] based on their theory of hot spot emission [34].

While there is an overall correspondence between radio and optical features in the entire jet body [4, 17], there is a severe discrepancy in region B1. This is one of those parts in the
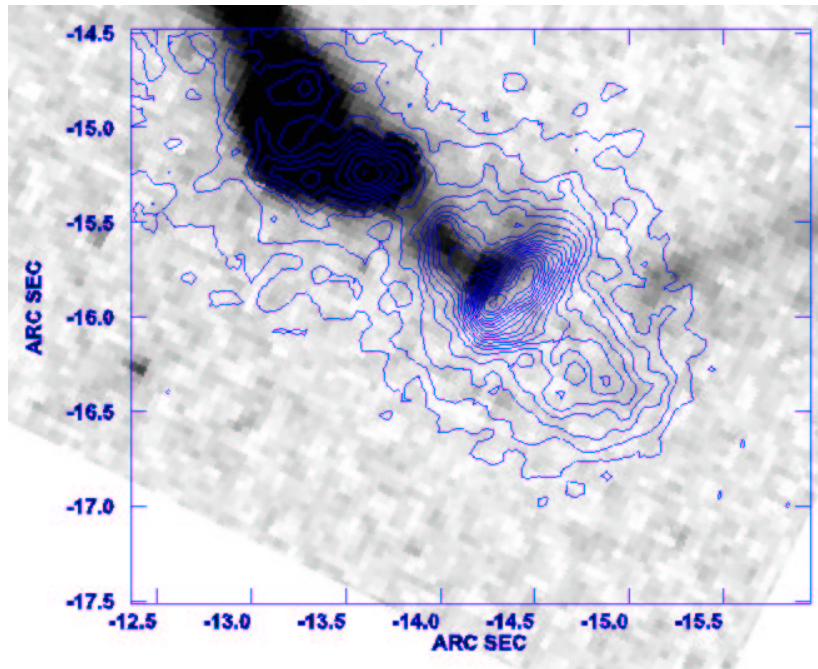
Figure 3.6: Comparison of hotspot position at optical and radio wavelengths. Greyscale shows the 620 nm image at original resolution, contours show the 1.3 cm radio map at 0.″125 resolution. Coordinates are relative to the quasar core, using which both images were aligned. The radio hot spot is clearly offset by 0.″2 from the optical hot spot, while the radio contours closely agree with the optical image in the preceding region (D2/H3). This offset has been predicted by Meisenheimer & Heavens [63], based on their model of the hot spot as a cylinder of diameter and height near 2 kpc at an orientation of 45° to the line of sight.

jet where its apparently double-stranded nature is most conspicuous. At high frequencies, the northern strand is considerably brighter than the southern, while on the radio images, the situation is exactly the opposite. This is also the only location in the jet which may be classified as truly edge-brightened (*cf.* Sect. 3.3).

Figure 3.7 shows that there is polarised extended emission south of the jet at radio wavelengths, a feature already apparent on the infrared image. This was first identified by Röser et al. [85] on radio maps and by Neumann et al. [73] on infrared images, who concluded that this extended emission corresponds to the backflow (material having passed the hot spot and flowing into the lobe, sometimes also termed cocoon), while the optical jet traces the emission of the jet channel as such. We expect this difference to be reflected in the spectrum of both components.

This immediately raises the concern that we observe the jet channel through a layer of cocoon material, leading to a "contamination" of the channel's radio emission by the cocoon. If this contamination is significant, the channel's radio emission should be corrected for it. We therefore consider the likely geometrical structure of the jet.

3.3 Inferring the jet volume

For a number of applications, *e. g.*, for the minimum energy considerations later in Sect. 4.3 it is necessary to know the jet volume. We consider what information about the geometry
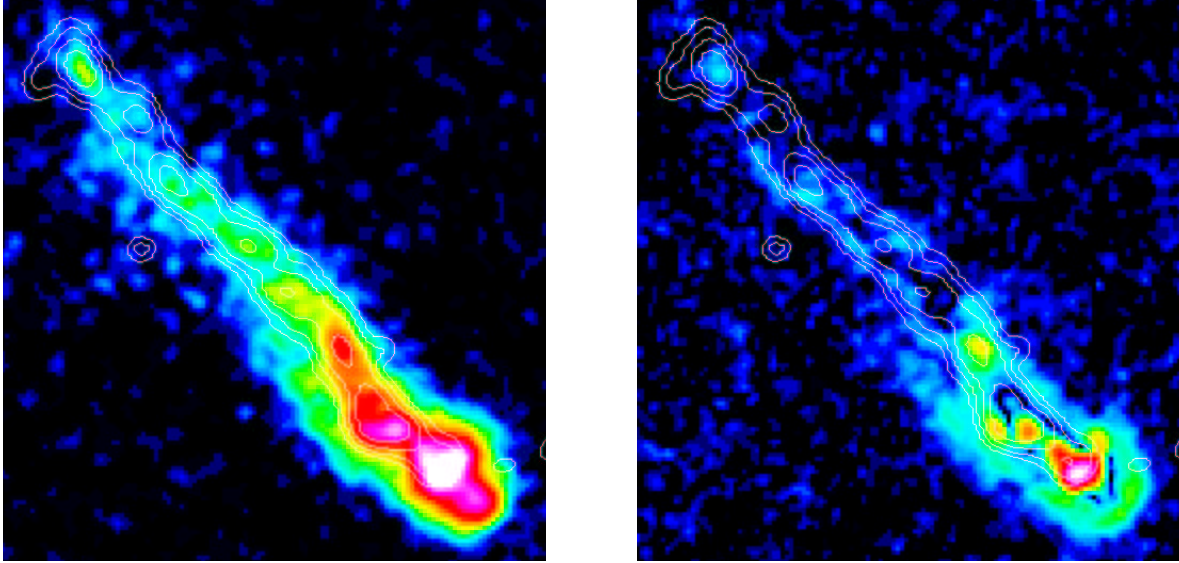
Figure 3.7: Comparison of polarisation and morphology. Colorscale image of radio flux at 0.″24 resolution, λ3.6 cm, showing total flux (left) and polarised flux (right). Contours show optical flux smoothed to the same resolution (single WFPC2 image, λ620 nm).

of the emitting volume, and possibly the magnetic field geometry, can be inferred from the observed brightness distribution (Fig. 3.2).

3.3.1   Geometry

As noted above, the optical jet probably delineates the jet channel as such, while the radio emission contains contributions from the surrounding material. The geometry of the jet channel is therefore best constrained by considering the optical morphology. Important information, in particular about the magnetic field configuration, is also contained in the polarisation of the jet flux. As no high-resolution optical polarisation maps are available, we consider the radio polarisation map at 3.6 cm (Fig. 3.7).

In all following considerations, we assume that the magnetic field strength and electron density are constant along all lines of sight through the jet. As further simplification, beaming effects are neglected. This is appropriate even in the presence of beaming as long as the beaming does not change over the jet, which may happen either because the jet velocity changes (along the jet) or because the flow direction changes due to a complicated flow pattern (along and transverse to the jet). Assuming the flow is uniform, any relativistic beaming will be constant over all parts of the jet.

As simplest case, we can assume that the magnetic field is fully tangled, so that the emission is isotropic. This would be appropriate for a turbulent flow. Changes in surface brightness are then simply due to changes in the path length through the source. The simplest jet-like geometry is that of a cylinder, possibly hollow (a hollow structure has been inferred for the radio emitting volume in the jet of M87 [77, 83]). Figure 3.8 shows the brightness profile of hollow cylindrical shells of varying thickness. Any hollow cylinder will always exhibit edge-brightening. This means that the brightest points lie away from the centre line, as long as the jet is resolved transversely, which is the case here. Since the jet images do not this show edge-brightening, with the possible exception of region B1, the jet cannot be a hollow
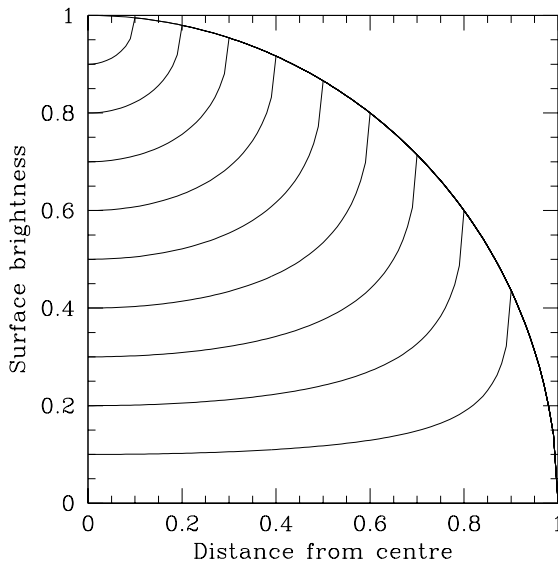
Figure 3.8: Brightness profile of hollow cylinders of varying thickness, showing only one half of the cylinder's diameter. Profiles are shown for cylindrical shells with outer radius held constant at unity and (from bottom to top) shell thickness 0.1, 0.2, . . . 0.9, 1.0 outer radii. Unity surface brightness corresponds to the brightness observed on the axis of a completely filled cylinder. All shells share the same outer profile. This diagram is also appropriate for cuts through spherical shells, such as planetary nebulae.

cylinder of uniform isotropic emissivity.

For a homogeneous magnetic field in the jet, there is a dependence of the emissivity on the angle between the line of sight and the magnetic field direction; the origin of this angular dependence is illuminated by Meisenheimer [61], useful formulae are tabulated in [51]. Could the observations be explained if there is a homogeneous field geometry, leading to anisotropic emissivity? Laing [48] has computed brightness profiles for various field geometries of cylindrical jets. For example, he considers a helical field confined to a cylindrical shell of thickness much less than the jet radius. Such a configuration will always produce an edge-brightened appearance. In contrast to the isotropic-emissivity case, which always produces a profile symmetric about the jet axis, the helical field configuration can lead to an asymmetric appearance. However, unless the jet lies in the plane of the sky, even the helical field will always lead to an edge-brightened appearance, assuming the axis of the jet has been identified correctly and the jet is resolved transversely (*cf.* Fig. 1 in [48] and Fig. 43 in [61]). The same geometric effect will apply to any field configuration, so that we can conclude that the emitting region is not confined to a thin shell at the surface a cylinder (still with the exception of B1).

The true jet volume can therefore be approximated by a filled cylinder, whatever the field geometry. The small-scale, double-helix-like structure seen on the optical images (Fig. 2.6) and the $0\rlap{.}''2$ optical spectral index map [42] suggests that the true internal structure of the jet is more complicated — so complicated that a more accurate model than the simple one assumed here requires a detailed understanding of the internal structure, composition and flow parameters governing the fluid dynamics of the jet. Any model with more parameters than a filled cylinder is not constrained by the available data. In principle, a detailed comparison of the observed intensity, degree of polarisation and projected magnetic field direction with the
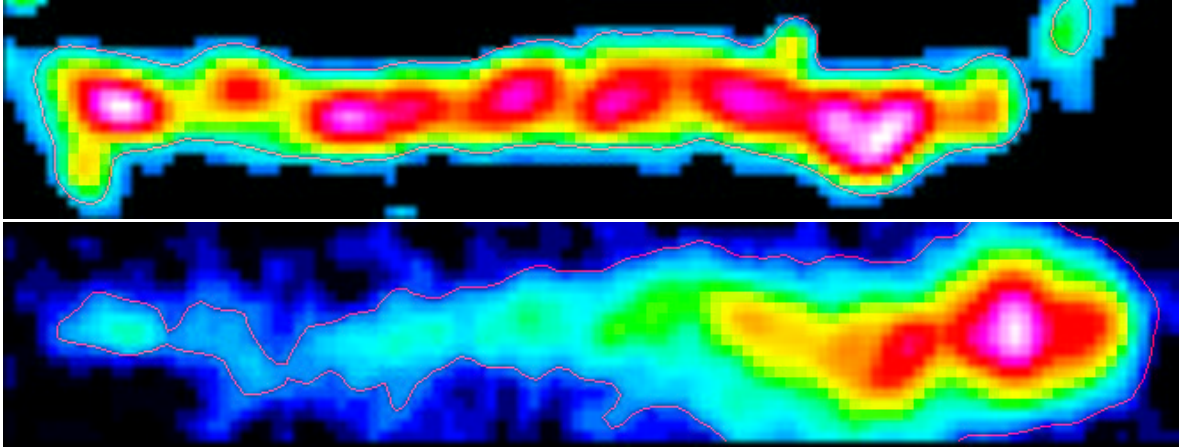
Figure 3.9: Sample isophote on an optical ($\lambda 600$ nm, top) and radio ($\lambda 1.3$ cm, bottom) jet image.

predictions by Laing [48] should yield some insight into the field geometry. This is, however, beyond the scope of the present work, as is an attempt to infer the internal structure of the jet. (A comparison of the optical (unpolarised) flux and the radio polarisation (Fig. 3.7) suggests that parts of the optical jet are edge-brightened in polarised radio flux, but regions A and D1 are clearly centre-brightened with a good correspondence between optical flux and polarised radio flux.)

Therefore, we will assume for simplicity that the jet is a filled cylinder extending along position angle $222^\circ\!.2$ whenever a value of the jet volume is required.

3.3.2   Width of the jet

Having chosen a model for the geometry of the jet, we need to determine the parameter describing the model, namely its radius, to infer the jet volume. In principle, the model parameter could be obtained by a fitting procedure using the observed brightness distributions. But since the assumed cylinder model is so simple, it appears sufficient to estimate the cylinder's radius from the direct images. However, this estimate is not as accurate as may be assumed. To illustrate the pitfalls of "straightforward" width estimations, we consider ways to determine the width of a jet *without* fitting a model.

When considering the isophotal width, the jet appears of constant width at high frequencies, but widens considerably at radio frequencies (Figures 3.9 and 3.10). However, the radio jet's brightness increases considerably outwards (Fig. 3.3), while the optical jet's brightness is approximately constant or even decreasing. A widening of the isophotes will already be caused by the brightness increase alone if the jet is not well resolved, as a larger part of the PSF wings is visible above the background noise. To disentangle this confusion between rising central brightness and actual widening caused by the finite resolution without doing an explicit deconvolution, one may define a dimensionless width parameter w_{iso}. w_{iso} is the ratio of the measured isophotal width to the width a Gaussian with the observed central brightness would have at that level. The isophotal width is determined column by column, and "central brightness" here refers to the maximum brightness along that column. In effect, w_{iso} is a measure of how resolved the jet is transversely, with $w_{\mathrm{iso}} = 1$ corresponding to the jet being unresolved. A similar width parameter w_{FWHM} may be defined as the ratio of the observed width at half the maximum brightness along the column to the resolution FWHM. Figure 3.10
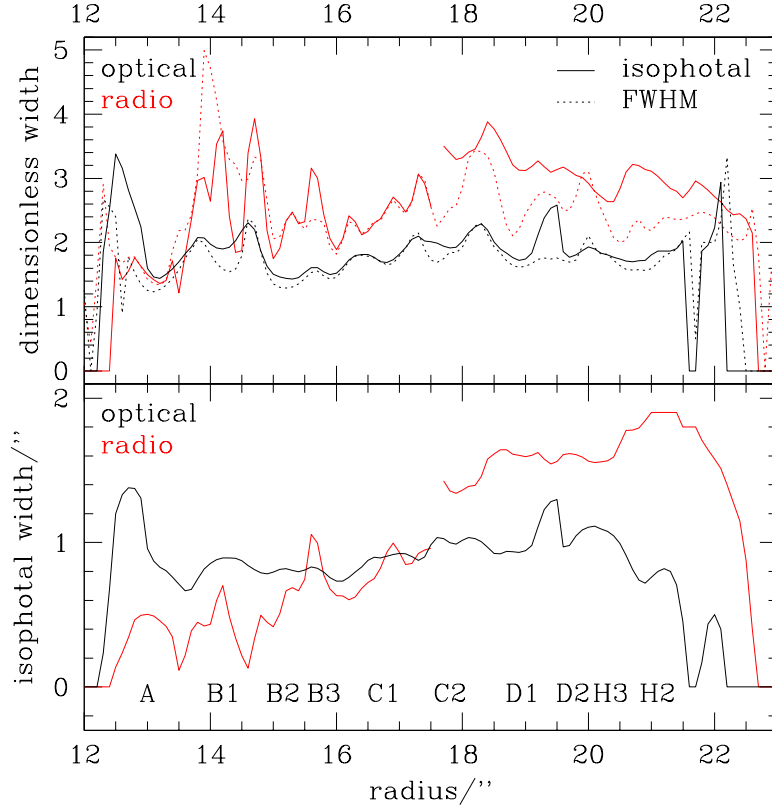
Figure 3.10: Comparison of different criteria for jet width. Black, width determined on optical image (the apparent widening between D1 and D2 is due to the star north of the jet; features beyond $r = 21\rlap{.}''5$ are due to the galaxy ("outer extension") and do not belong to the jet); red, width determined on radio map (the gap and jump at $r = 17\rlap{.}''6$ is due to the narrow feature of the isophote extending to the south of the jet). Bottom, isophotal widths for the isophotes shown in Fig. 3.9. Top, dimensionless isophotal width w_{iso} for these isophotes (solid lines, see text) and ratio w_{FWHM} of width at half maximum brightness to resolution FWHM (dashed lines).

compares the two parameters. It turns out that for the optical jet, both dimensionless width parameters agree rather well, the exception being those regions affected by the presence of the southern extension S at radii $12''$–$13''$ and the star near $19\rlap{.}''25$. At the resolution of $0\rlap{.}''3$, both w_{iso} and w_{FWHM} for the radio jet are always larger than the corresponding values for the optical jet, despite the radio jet's narrow isophotal width. The narrow appearance at this resolution is thus only due to the its low surface brightness. It appears that region A ($r \approx 13''$) is in fact unresolved transversely even on the highest-resolution radio maps, while it is surrounded by diffuse emission on the red-band image.

One might hope to determine the true extent of the jet on the images or radio maps with the highest available resolution. While this is possible on the HST images for the entire jet, the resolution of $\approx 0\rlap{.}''1$ for the optical and $\approx 0\rlap{.}''2$ for the infrared is not reached with sufficient signal-to-noise for most part of the jet by the VLA. Therefore, the current radio data do not permit a comparison of the jet width at the resolution reached by HST.

Since the width of the optical jet can be determined very well, this is used as value for the width of the jet channel, which is thus described as a cylinder of constant radius $0\rlap{.}''7/2 = 0\rlap{.}''35$
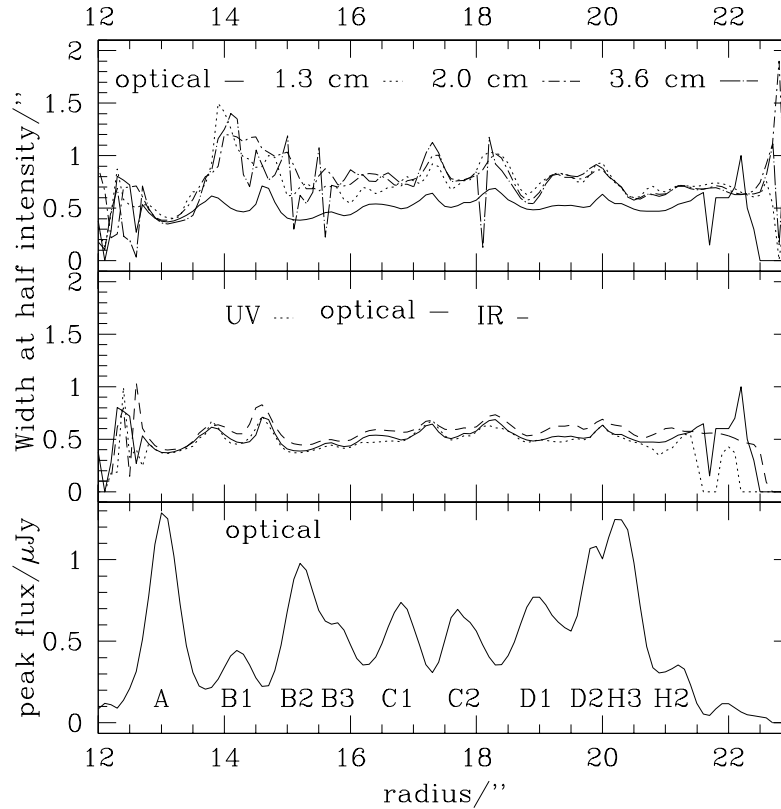
Figure 3.11:   Comparison of jet width at different wavelengths. The width is determined column by column on the images in Fig. 3.2 as the full width at half the maximum intensity along the column. Bottom panel shows optical flux profile (maximum intensity) for reference. Middle panel shows width at $\lambda 300\,\mathrm{nm}$ (short dash), $\lambda 620\,\mathrm{nm}$ (solid) and $\lambda 1.6\,\mu\mathrm{m}$ (long dash). Upper panel shows again the optical width (solid) for reference and width at $\lambda 1.3\,\mathrm{cm}$ (dotted), $\lambda 2.0\,\mathrm{cm}$ (dot-short dash) and $\lambda 3.6\,\mathrm{cm}$ (dot-long dash).

for regions A1–D2/H3 ($r = 12''$–$21''$). This width agrees with the hot spot diameter given by Meisenheimer et al. [68], so to lowest order, we can extend the cylindrical model to cover the entire jet out to the hot spot precursor H1. Since the morphology of the backflow is even more difficult to infer from the much fainter backflow emission, and since the object of study is the jet channel, we will not consider it in as much detail as the jet itself. However, to assess any likely contribution of the backflow emission to the flux observed from the jet, the radio emission surrounding the optical jet channel is described as cylindrical shell with inner radius $0\rlap{.}''35$ (enclosing the optical jet channel without a gap) and outer radius $0\rlap{.}''8$, close to the isophotal width of the radio jet there and roughly twice that of the optical jet (Figs. 3.10 and 3.11).

To avoid the use of another approximation, the volume belonging to each photometry aperture (or pixel in Fig. 3.2), *i. e.,* the effective jet volume sampled by each photometry aperture, is calculated by explicitly convolving the model assumed for the jet (filled cylinder of radius $0\rlap{.}''35$) with the observing beam of $0\rlap{.}''3$ FWHM at the location of each aperture, assuming that the symmetry axis of the cylinder lies along the radius vector at position angle $222\rlap{.}°2$. These obtained values are tabulated in Tab. 3.1. The effect of the inclination of the jet to the line of sight will be discussed below in Sect. 4.3.2.

δ_{y} ''	Jet volume $(h_{70}^{-1}\,\mathrm{kpc})^3$	Backflow volume $(h_{70}^{-1}\,\mathrm{kpc})^3$
0.0	1.3	1.8
0.1	1.2	1.9
0.2	.98	2.0
0.3	.62	2.3
0.4	.28	2.4
0.5	.084	2.3
0.6	.015	1.9
0.7	.0016	1.3
0.8	.0001	.71
0.9	.00003	.26
1.0	.0000007	.06

Table 3.1: Effective volume of the jet sampled by photometry apertures at distance δ_{y} from the symmetry axis of the model. *Jet volume*, contribution from the jet channel visible in the optical, assumed as filled cylinder. *Backflow volume*, contribution from the backflow, modelled as hollow cylinder wrapped around the jet channel. These values assume a fully side-on view.

3.3.3 "Backflow" material

Using Tab. 3.1, we can now also estimate the contribution of the backflow material along the line of sight to the central part of the jet. The effective volume contributed by the backflow volume is much larger than that of the jet channel. Therefore, the backflow's volume emissivity must have no more than about 1% of the jet's volume emissivity. Otherwise, the backflow emission would completely dominate the jet emission and the profile would not appear centrally peaked, or fall off more slowly than observed. In the central part of the jet, the backflow will then also contribute only about 1% of the jet emission. The same constraint from the observed brightness profile implies that the contribution of the backflow to the central jet flux cannot be appreciable even if the true width of the backflow is different from the assumed $0\rlap{.}''8$.

3.4 Jet morphology: summary

In summary, the jet appears very similar morphologically at all wavelengths, with discrepancies in three regions (Fig. 3.2):

1. In B1, there is a bright optical filament arching to the north, while the radio emission is arching to the south.

2. The precursor to the radio hot spot is clearly detected in the infrared, but not at optical or UV wavelengths.

3. There is a "backflow" component south of the jet, detected in the radio and infrared [73] but not at optical wavelengths. The backflow component has a much lower emissivity than the jet material even at long wavelengths and therefore does not make a significant contribution to the central surface brightness of the jet at any wavelength.

   4. There is an offset of 0″2 between the optical and radio peak of the hot spot (Fig. 3.6).
      This feature was predicted by Meisenheimer & Heavens [63], based on their model of
      the hot spot as a cylinder of diameter and height near 2 kpc at an orientation of 45°to
      the line of sight.

3.5 Spectral indices

To describe the observed spectra quantitatively, we construct spectral index maps from the
observed brightness distributions.

3.5.1 Definition of spectral index

Independently of whether a spectrum actually does follow a power law over any range of
frequencies, a local two-point spectral index can be defined between any two surface brightness
measurements B_1, B_2 at frequencies ν_1 and ν_2, respectively, as

$$\alpha = \frac{\ln \frac{B_1}{B_2}}{\ln \frac{\nu_1}{\nu_2}}. \tag{3.5}$$

   The error on the spectral index is computed from the values of the noise σ_1, σ_2 in the
respective input images:

$$\sigma_\alpha = \frac{1}{|\ln \frac{\nu_1}{\nu_2}|} \sqrt{\frac{\sigma_1^2}{B_1^2} + \frac{\sigma_2^2}{B_2^2}}. \tag{3.6}$$

This formula shows that the spectral index error depends on the S/N of the input images
and the "baseline" between the wavelengths at which the observations are made and a given
relative flux error results in the same relative spectral index error. We compute the pairwise
spectral indices between maps neighbouring in frequency space.

3.5.2 Spectral index maps

A detailed discussion of the optical spectral index map at 0″2 resolution has been presented
elsewhere [41, 42]. Figure 3.12 shows the spectral index maps generated from the jet pho-
tometry at 0″3 resolution. We describe in turn the features of the maps.

Optical spectral index

The optical spectral index declines globally outwards from -0.5 near the onset of the optical
jet at A to -1.6 in D2/H3. This trend does not continue into the hot spot. The general
steepening is in agreement with previous determinations of the knots' synchrotron spectrum
which showed a decline of the cutoff frequency outwards [64, 89]. The global run of the
optical spectral index α_{BRI} determined at 1″3 resolution [87] agrees very well with our new
determination of α_{opt} at much higher resolution (Fig. 3.13). Discrepancies arise in region A
(13″) and C2/D1 (18″–19″); we defer a discussion to Sect. 5.2.
   The spectral index varies very smoothly and gradually along the jet, without any sharp
jumps. These smooth variations of the show that the physical conditions in the jet change
remarkably smoothly over scales of many kpc. There are large but smooth variations of
the spectral index over the jet's projected extent of about 10″, while the surface brightness
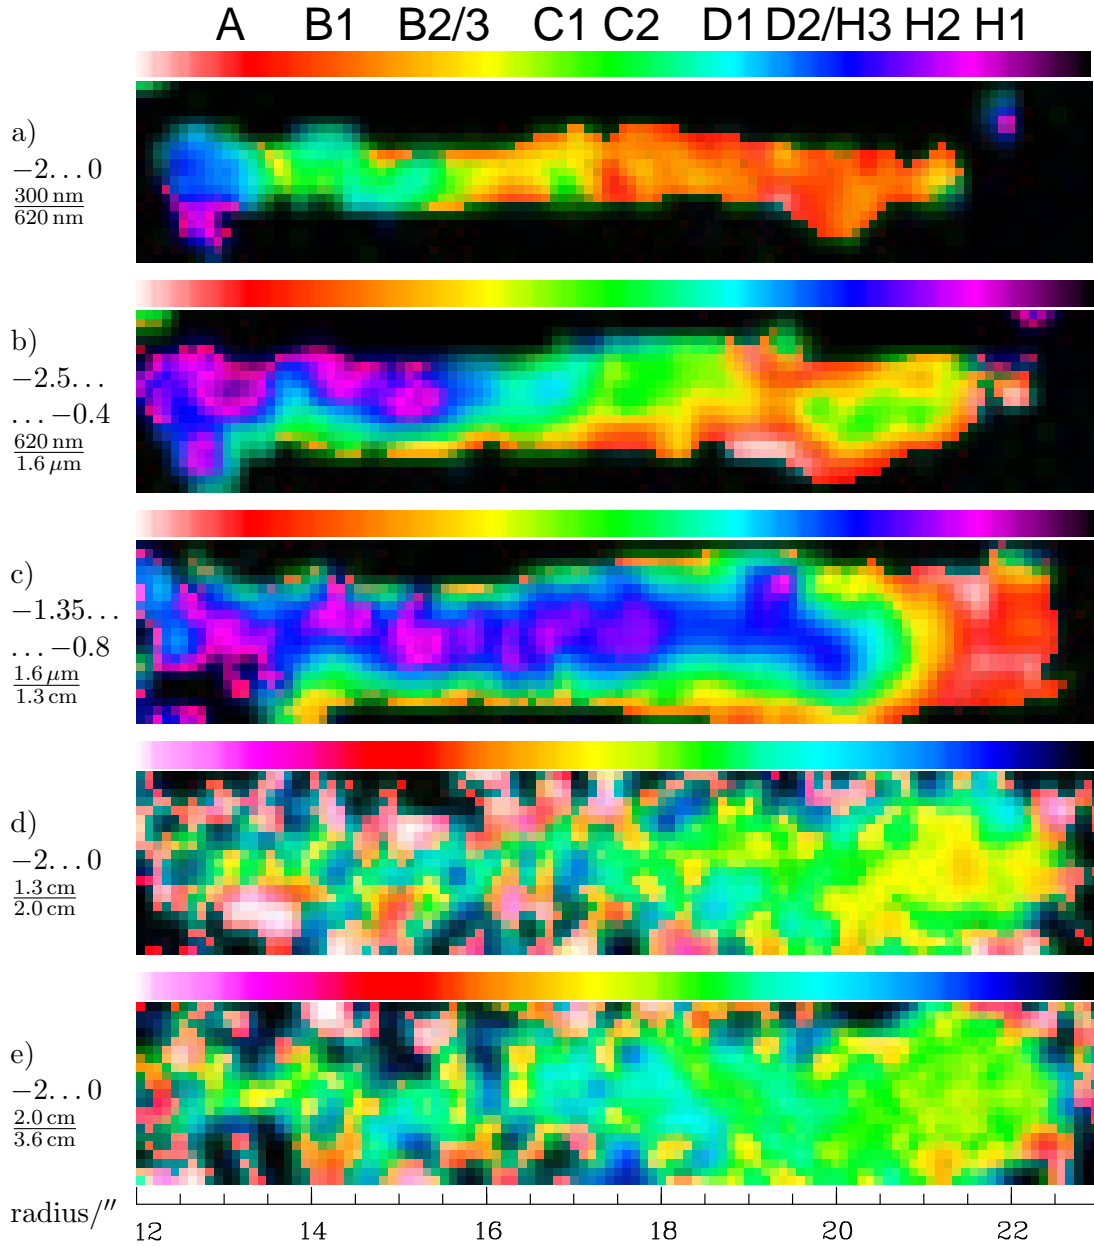
Figure 3.12: Spectral index maps at 0″3 resolution generated from the photometry data in Fig. 3.2. Images are combined pairwise in order of increasing wavelength. Linear colour scales (shown above the respective images) have been chosen to stress variations within one map.
a, optical spectral index (range: $-2\ldots0$)
b, optical-infrared ($-2.5\ldots-0.4$)
c, infrared-radio ($-1.35\ldots-0.8$)
d, radio $\lambda1.3\,\mathrm{cm}$-$\lambda2.0\,\mathrm{cm}$ ($-2\ldots0$)
e, radio $\lambda2.0\,\mathrm{cm}$-$\lambda3.6\,\mathrm{cm}$ ($-2\ldots0$)
The variations of both radio spectral indices are mainly due to low signal-to-noise and the associated imaging uncertainties for the inner part of the jet (*cf.* Fig. 3.14 below).
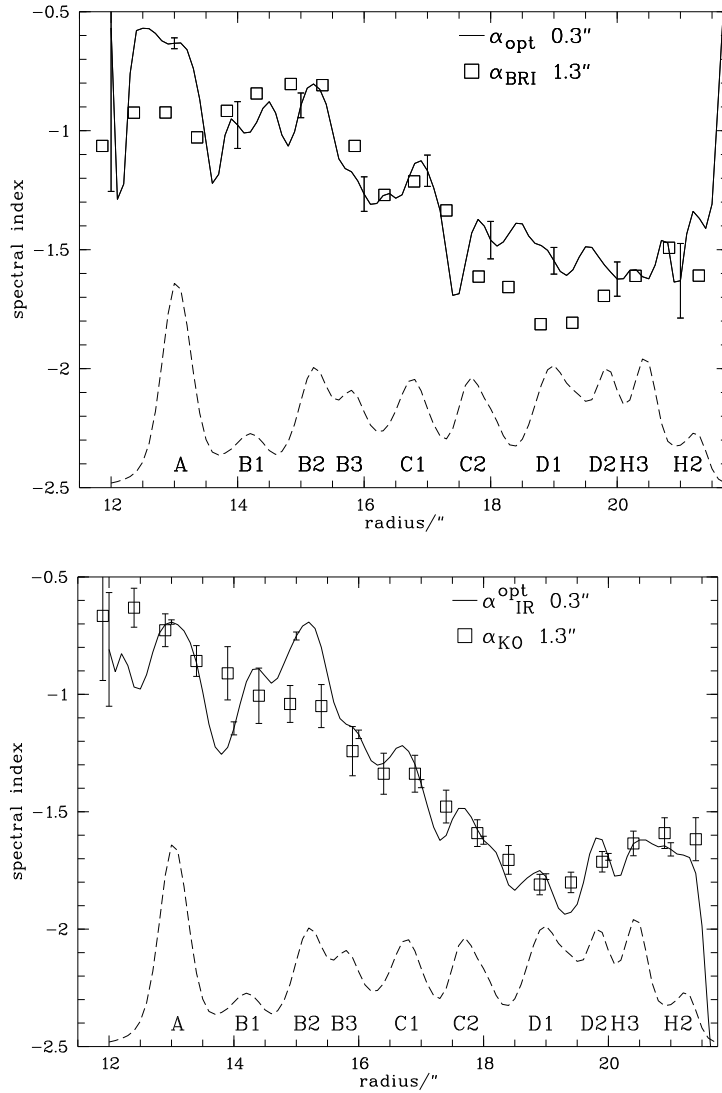
Figure 3.13: Comparison of high-frequency spectral indices with previous observations at $1''3$ resolution. Top, comparison of optical spectral index α_{BRI} at $1''3$ resolution from Röser & Meisenheimer [87] with optical-UV spectral index α_{opt} at $0''3$ from this work (error bars are similar for α_{BRI} and α_{opt}). Below, comparison of infrared-optical spectral index α_{KO} from Neumann et al. [73] and $\alpha_{\mathrm{IR}}^{\mathrm{opt}}$ from this work (only every tenth error bar shown).

remains fairly constant. Conversely, there are large *local* variations of surface brightness without strong changes in the spectral index. Thus, we do not find the strict correlation between optical surface brightness and optical-ultraviolet spectral index that has been found in the jet in M87 [64, 82].

Considering the spectral index changes from knot to knot, there is a marginally significant flattening of the spectrum in the transitions A-B1, B1-B2, C1-C2, and moving out of C2. From the strong losses leading to the observed high synchrotron luminosity a corresponding steepening of the optical spectral index is expected. In contrast to the expectation, the

behaviour of α_{opt} is consistent with the *absence of losses.* In any case, the overall steepening of the spectrum (from region A down do D2/H3) is less rapid than that *within* individual regions (*e. g.*, A and B2). There is no significant steepening from D1 out to the connection between H3 and H2, despite large variations in surface brightness.

The criss-cross morphology in C1, C2, B1, B2 and D1 (*cf.* Sect. 2.2.3) is reflected on a spectral index map at $0\overset{''}{.}2$ resolution [42] as a band of one colour crossing a second one. This is most clearly seen in knot C1 which has a green band of $\alpha \approx -1.1$ across an orange region of $\alpha \approx -1.4$, supporting the interpretation of two emission regions appearing on top of each other. Such features indicate that the simple cylindrical model for the optical jet assumed here will need to be modified in studies making use of the full resolution of below $0\overset{''}{.}1$ available at optical wavelengths with the Hubble Space Telescope.

Optical-infrared spectral index

There is an excellent correspondence of the optical-infrared spectral index $\alpha_{\mathrm{IR}}^{\mathrm{opt}}$ with the spectral index α_{KO} between $2.1\,\mu\mathrm{m}$ and the optical as determined by Neumann et al. [73] at $1\overset{''}{.}3$ resolution (Fig. 3.13). The overall run *along* the jet is a steepening with increasing distance from the quasar, similar to the optical spectral index. However, *transverse* to the jet, the optical-infrared spectral index map has a divided appearance (Fig. 3.12 b): the inner part, out to D2 at $r \approx 19''$, shows a strong spectral index gradient $\alpha \approx -.4 \ldots -1.6$ across the jet. In contrast, the outer part appears symmetric about the jet axis.

Unlike for the optical spectral index, there is a correlation between the optical-infrared $\alpha_{\mathrm{IR}}^{\mathrm{opt}}$ and the jet's surface brightness, in the sense that brighter regions show a flatter spectrum ($|\alpha_{\mathrm{IR}}^{\mathrm{opt}}|$ smaller, see Fig. 3.13). This correlation is most clearly visible on the spectral index map for the outer end of the jet. The bright regions D2/H3 together with the hot spot H2 appear as an island of $\alpha_{\mathrm{IR}}^{\mathrm{opt}} \approx -1.6$ surrounded by regions with steeper spectrum. The correlation between local maxima in surface brightness and $\alpha_{\mathrm{IR}}^{\mathrm{opt}}$ is also present in the inner part of the jet, although the spectral index maxima are displaced from the brightness peaks due to the transverse spectral index gradient.

This spectral index gradient is suggestive of a misalignment between the optical and infrared images, *e. g.*, a rotation between the two about a point close to D2/H3. It could also have been caused by an overestimation of the diffraction spike signal which has been modelled and subtracted (see Sect. 2.2.5). Since Neumann et al. [73] did not detect a significant change of the infrared-optical spectral index transversely to the jet at $1\overset{''}{.}3$, and although the alignment procedures described above (Sect. 3.1) should have ensured that such an error should not have occurred, we reconsidered this possibility to avoid the introduction of spurious gradients.

After a detailed investigation (for particulars, see App. C), we concluded that the misalignment necessary to produce such a gradient was far greater than compatible with the alignment precision established previously. Neither can the gradient firmly be linked to the diffraction spike subtraction or any obviously detectable misalignment. The gradient is qualitatively consistent with being caused by the detection of the steep-spectrum backflow component south-east of the jet also at $1.6\,\mu\mathrm{m}$, which had been detected at $2.1\,\mu\mathrm{m}$ by Neumann et al. [73], although the backflow component would be expected to make a significant contribution only at the outer end of the jet (see Sect. 5.2 for a detailed discussion of a possible "backflow" contribution; *cf.* also Sect. 3.3.3). In the absence of better knowledge, we rely on the data with the offsets established to the best of our knowledge. The clarification of this matter has to await new observational data.

Infrared-radio spectral index

The infrared-radio spectral index is nearly constant at $\alpha_{\text{radio}}^{\text{IR}} \approx -0.9$ along the centre of the optical jet, with some indication of flattening in optically bright regions. It steepens markedly to ≈ -1.2 away from the centre line. These features are identical to those identified by Neumann et al. [73] on a spectral index map at $1\overset{''}{.}3$ resolution generated from observations at 73 cm and $2.1\,\mu$m. There is no spectral index feature easily identified with the hot spot, as is the case on the radio spectral index map. On the other hand, the precursor is identifiable as region with slightly flatter infrared-radio spectral index.

The bright star to the north of the jet stands out as a fleck with flatter spectral index than its surroundings; the second star, which has been modelled and subtracted (Sect. 2.2.5), is not equally conspicuous, suggesting that the modelling and subtraction have been accurate.

Radio spectral indices

The two radio spectral indices behave erratically out to a radius of about $19''$. These variations are not significant: the radio jet is detected at low signal-to-noise ratio at the inner end, and the deconvolution involved in the reconstruction of the brightness distribution on the sky from the observed interferometric data is only accurate (in the sense of achieving an image representing the true brightness distribution) for high signal-to-noise. For the outer part of the jet, the radio spectral index maps show a steepening of the hot spot spectrum compared to the remainder of the jet. The run of the spectral index between 6 cm and 3.6 cm, for which only lower-resolution data at $0\overset{''}{.}5$ are available, agrees with the spectral index run determined at the wavelengths considered here.

3.5.3   Characterisation of spectra along the jet

Figure 3.14 shows the run of spectral indices along the radius vector at position angle $222\overset{\circ}{.}2$. Using this figure as a guide, we consider the shape of the synchrotron spectra encountered at $0\overset{''}{.}3$ resolution.

Any synchrotron spectrum arising from a single electron population and whose shape is dominated by losses always has a convex shape, so that higher-frequency spectral indices are steeper than those at lower frequencies. In contrast to this expectation, the observed spectrum shows regions in which the optical spectral index is flatter than the infrared-optical (*e. g.*, regions C2 to D2/H3, from $r \approx 18''$ to $r \approx 20''$ in Fig. 3.14). In A ($r \approx 13''$) and B2 ($r \approx 15''$), both these spectral indices are flatter than the radio-infrared spectral index, which is also surprising.

In spite of these difficulties, we attempt to find a common way of characterising the spectra everywhere in the jet, assuming that the spectra can be described as due to a single electron population. It is important to assess how far this simple assumption is consistent with the data, and for which parts of the jet. The derivations from the assumed simple picture will then be a guide to more refined models and testable predictions. Because of the small number of data points ($N = 6$), any model for the spectra will have a small number of degrees of freedom. The best-fit parameters of the model spectrum are therefore not all interpreted with the usual meaning of a *measurement* of physical parameters, but instead mainly as a *description* of the observational data.

We consider different descriptions of synchrotron spectra, starting with the smallest number of degrees of freedom.
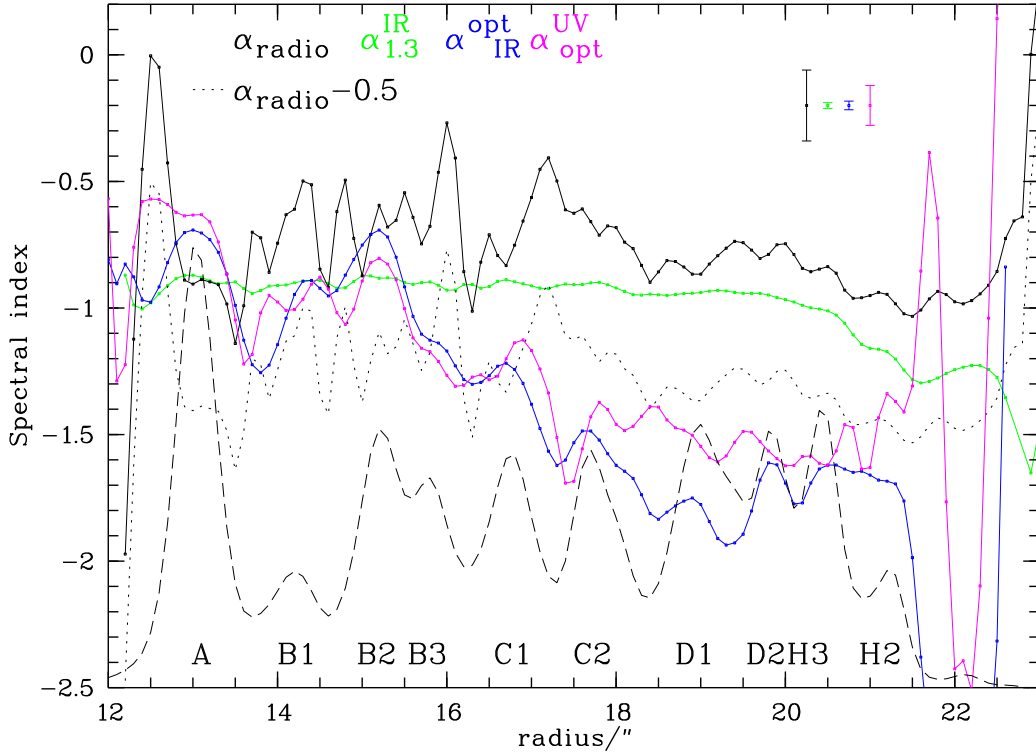
Figure 3.14: Run of the spectral indices along the jet (cut along radius vector at position angle 222°2). Spectral indices are taken from Fig. 3.12, with the exception of the radio spectral index, which was obtained by a least-squares straight line fit to the radio data. For reference, the dashed line shows the flux profile at 600 nm (flux scale 0–2 μJy). Brightness peaks show a clear correlation with peaks of $\alpha_{\mathrm{IR}}^{\mathrm{opt}}$ (blue), but there is no clear correlation with α_{opt} (magenta).

For sake of clarity, only typical 2σ error bars are shown. The radio error bar shows the uncertainty arising from the conservative 10% error accounting for the uncertainty of the radio flux scale as well as the inaccuracy associated with the deconvolution process. Assuming sky RMS noise for the radio points makes these error bars similar to the optical-infrared error bar (blue).

Simple power law

A power law of the form $S_\nu = S(\nu_0)\,(\nu/\nu_0)^\alpha$ describes the spectrum with just two free parameters, the spectral index α and the flux normalisation S_0. For a single, straight power law to be adequate everywhere, spectral indices at different frequencies need to be equal to each other within the errors, although they may in principle vary across the source. If at all, only the spectral indices at the brightness peaks at $r = 13''$ (A1) and at $15''$ (B2), in which no cutoff is detected (see below), may only be consistent with a single power law from radio to ultraviolet if the errors have been significantly underestimated.

Power law with break

The next simplest description would be a broken power law, introducing a change in spectral index at some break frequency ν_{b}. When describing synchrotron spectra as power laws, the

inherent assumption is that the electron energy distribution is a power law. The only physical break in a synchrotron spectrum is then a break from α_{low} to α_{high} by $1/2$ [34]. With the magnitude of the break fixed by synchrotron theory, the only additional free parameter is the break frequency ν_{b}. Slightly larger breaks may be possible by assuming certain source conditions, for example systematically varying magnetic fields [J. Kirk, *priv.comm.*].

From Fig. 3.14, it is clear that there is a break significantly larger than 0.5 between the radio spectral indices of approximately -0.8 to the infrared-optical and optical-UV spectral indices of below -1.5. A break larger than 0.5 cannot be justified here. Since our aim is to describe the spectra of the entire jet with a single model whose parameters vary along the jet, instead of having to assume different models for different parts of the jet, a simple broken power-law is also inadequate.

Power law with break and cutoff

The next physical extension of a simple model for a synchrotron spectrum is the inclusion of a high-energy cutoff to the electron energy distribution. Such a cutoff will then lead to a curved high-energy cutoff in the emitted synchrotron spectrum. It introduces a fourth free parameter for describing the spectrum, the cutoff frequency ν_{c}.

Heavens & Meisenheimer [34] have computed the shape of synchrotron spectra for hot spot emission. In their model, the hot spot is a localised particle acceleration region, identified with the Mach disk of the terminal jet shock. At the shock, electrons are accelerated by the first-order Fermi mechanism to a power-law energy distribution extending up to a maximum Lorentz factor γ_{max}. This maximum Lorentz factor is that for which acceleration gains during one cycle are balanced by synchrotron losses. Upon escaping the acceleration region, electrons then travel downstream, now only losing energy by emission of synchrotron radiation. Since the synchrotron cooling rate is proportional to the particle energy squared (Eqn. D.4), the maximum energy decreases with increasing distance from the shock, and higher-energy particles loose a significant fraction of their energy faster than lower-energy particles. If the downstream region is observed unresolved, the resulting net spectrum is the sum of power laws with different high-energy cutoffs, resulting in a low-energy power law, a break at which the spectral index steepens by 0.5, and a high-frequency cutoff corresponding to the cutoff at the shock. These spectra are described by the simple parameter set just described.

For hot spots, the relevant physical parameters emission volume, downstream electron diffusion length and hence magnetic field, maximum particle energy are model parameters and can thus be inferred from the spectrum. Although the same physical model cannot be applied to the jet itself, the *shape* of the synchrotron spectrum from the entire jet can also be described by the same set of parameters as a hot spot spectrum [64, 67, 71, 89]. This is possible because the jet, like the hot spot, can be envisaged as a loss region into which a power-law distribution of electrons is fed continuously, and with superposed emission from electron populations with different cutoff frequencies. We also follow this approach here, since this model can account for the variations of the spectra along the jet by variations of the model parameters.

The next step is thus an analysis of the observed flux distribution by fitting synchrotron cutoff spectra.

Chapter 4

Analysis

4.1 Fitting synchrotron spectra

Synchrotron spectra according to Heavens & Meisenheimer [34] are fitted to the observed flux per $0\rlap{.}''3$ beam at $3.6\,\text{cm}$, $2\,\text{cm}$, $1.3\,\text{cm}$, $1.6\,\mu\text{m}$, 620 nm and 300 nm. Fits are performed for those apertures out of the total of 2331 which have a signal-to-noise of at least 5 at 620 nm or at $2\,\text{cm}$, to include only points with significant flux in at least one of the two wavelength regions observed. The fits are performed using existing software [61]. In employing fits here which have been developed to describe hot spot spectra, we make use of the property that the spectra arising from the hot spots (via an integration over particles at different distances downstream from a localised acceleration region) are identical to spectra arising from ageing in a continuous injection model (mixing particles with different ages), appropriate for a jet [34]. However, the parameters of the physical hot spot model cannot be applied when considering the emission from the jet body. We therefore only use the fits here to extract the parameters describing the shape of the synchrotron spectrum, and in particular the cutoff frequency ν_{c}, which is one of the four free parameters in the fit:

1. low-frequency spectral index α_{low},

2. ratio of cutoff energy to break energy of the electron population $= \sqrt{\nu_{\text{c}}/\nu_{\text{b}}}$; at ν_{b}, the spectrum steepens from α_{low} to $\alpha_{\text{high}} = \alpha_{\text{low}} + 0.5$,

3. observed cutoff frequency ν_{c},

4. flux normalisation.

With four free parameters and six data points, there are only two degrees of freedom for the fit. However, the only fit parameter that will be directly interpreted is the cutoff frequency, and this is accurately determined by the three observed high-frequency flux points where a cutoff is present in the radio-ultraviolet spectrum. Where no such cutoff is observed, a lower value for the cutoff frequency is obtained.

Figure 4.1 shows an example spectrum. The first two parameters describe the overall shape of the spectrum, while the last two shift the shape in the $\log \nu - \log S_{\nu}$ plane. Actually, the shape of the spectrum is evaluated by varying the physical parameters of the model (electron injection index and extent of downstream region) to yield corresponding variations in the observed spectrum (α_{low} and $\nu_{\text{c}}/\nu_{\text{b}}$, respectively). The physical model also takes into
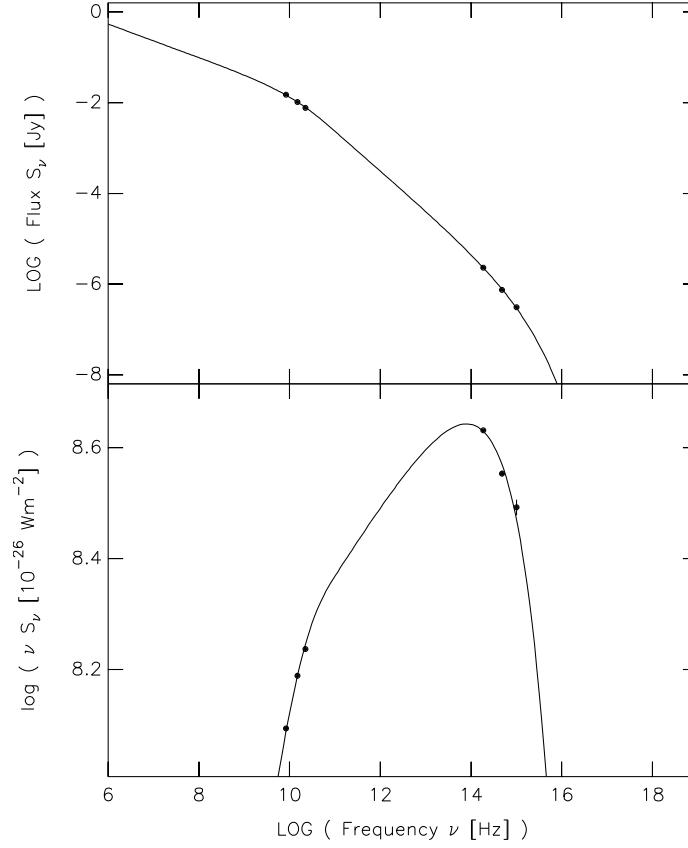
Figure 4.1:   Sample synchrotron spectrum, fitted to an observed spectrum from region C1 ($r = 16''8$).
Top, the plot of S_ν against $\log \nu$ shows the observed and fitted surface brightness; below, the plot of νS_ν
against $\log \nu$ shows the spectral energy density per frequency decade.

account the orientation of the magnetic field with respect to the line of sight (in this case, fixed
at $45°$[16, 68]). This allows to compute also the polarisation properties of the synchrotron
emission, although this information is not used here. The computation of the spectra is done
by

1. computing the electron distribution with a given injection index and cutoff energy (in
   the hot spot picture, arising from a certain distance downstream of the acceleration
   region),

2. integrating the contributions with different cutoff energies (different downstream dis-
   tances), resulting in an integrated electron energy spectrum with a break and cutoff,

3. convolving with the synchrotron Green's function to obtain the observed spectrum for
   a cell with a given orientation of the magnetic field,

4. averaging the contributions from different cells according to the prescribed magnetic
   field, which is described in cylindrical geometry. Here, the magnetic field has only a
   toroidal component and the emisison is summed up over the contributions of 360 sectors.

The fit is performed using a χ^2 minimisation technique.

4.2 Fit results

As mentioned above (Sect. 3.5.3), those regions in which the infrared-radio spectral index is steeper than the optical spectral index are impossible to describe by a single-component electron spectrum employed here. The discrepancy can be explained by assuming that one of the two spectral indices reflects the true synchrotron spectrum, while the other is contaminated by flux not due to the same population as the remainder of the jet.

To assess the likely reason for this discrepancy between the observations and the expectations from synchrotron theory, we perform two separate fits which differ in the determination of the cutoff frequency:

- **Fit A**, in which the steepness of the infrared-optical spectral index is ascribed to an infrared-bright but optically quiet "backflow" component. The high-frequency cutoff is determined by the optical and ultra-violet flux points alone, unless the resulting spectrum overpredicts the observed infrared flux.

- **Fit B**, in which the discrepancy is ascribed to an additional component only visible at ultra-violet wavelengths. The high-frequency cutoff is fitted to all observed data points. The UV residual of this fit is the minimum UV flux that needs to be contributed by a separate electron population or synchrotron mechanism.

In practice, secondary minima of χ^2 cannot be avoided if all parameters are varied simultaneously. Therefore, only one of the parameters α_{low} or $\nu_{\mathrm{c}}/\nu_{\mathrm{b}}$ is fitted, the other being kept fixed.

Suitable initial values for the parameters are required to optimise convergence of the fit. Since both the radio spectral index the radio-infrared spectral index are fairly constant from $r = 12''$ to $r \approx 21''$, we can describe the spectra from the entire jet with a low-frequency spectral index varying similarly little. The radio-infrared spectrum can be described by the part of the spectrum lying between cutoff and break by choosing α_{low} approximately 0.5 flatter than the observed radio-infrared spectral index of $\alpha_{\mathrm{IR}}^{\mathrm{opt}} \approx -0.9$. This enables to fit any radio spectral index steeper than α_{low} and leads to a fairly constant break frequency near the observed radio frequencies. In effect, these fits correspond to spectra in which only the cutoff frequency varies along the jet, but whose shape is otherwise unchanged.

Thus, using $\alpha_{\mathrm{low}} = -0.3$ as initial value, a first fit is performed to determine a rough value for $\nu_{\mathrm{c}}/\nu_{\mathrm{b}}$. Since we expect a cutoff near 10^{15} Hz and a break at frequencies near 10^9 Hz, we do not expect the ratio $\nu_{\mathrm{c}}/\nu_{\mathrm{b}}$ to lie significantly above 10^6 where a cutoff is observed, *i. e.*, at $r \geq 15''.5$. An upper limit of 10^7 is employed at radii $r \leq 15''.5$, where the observed infrared-ultraviolet spectrum is flatter than the radio-infrared spectrum.

Using the optimal value for $\nu_{\mathrm{c}}/\nu_{\mathrm{b}}$ from this fit, a small variation of α_{low} is allowed within the the interval $-0.45 < \alpha_{\mathrm{low}} < -0.35$ in Fit A, and $-0.45 < \alpha_{\mathrm{low}} < -0.2$ in Fit B. $\nu_{\mathrm{c}}/\nu_{\mathrm{b}}$ is varied a further time to yield the best-fit set of parameters.

In contrast to the remainder of the jet, the hot spot shows an offset between optical and radio hot spot position (Sect. 3.2) and a change of the radio-infrared spectral index (Sect. 3.5.2). Therefore, the spectra from the hot spot regions H2 and H1 at radii beyond $r = 21''$ are fitted differently by allowing α_{low} to vary between -0.8 and -0.3. This fit is referred to as **Fit HS**.

To select an appropriate initial value for the cutoff frequency of the fitted spectrum, its cutoff is matched with the observed highest-frequency points, which are presumed to show

Figure 4.2: Location of apertures for which spectra are shown in the following.

the cutoff. Where there is no cutoff in the observed spectrum, an artificial high-frequency data point is introduced to allow the fit to proceed at all. This point is chosen so that it has a spectral index relative to the observed UV point of -1.2 and is introduced at a frequency 10^{18} Hz, 1000 times higher than the frequency of UV emission. It is assigned a large error so that it does not influence the goodness-of-fit. The results obtained for ν_c for these points are then only lower limits to the true cutoff frequency.

Because of the wealth of independent data points, we consider the fitted spectra in detail only at 16 selected locations. Figures 4.3–4.5 show the spectra fitted at the locations indicated in Fig. 4.2. At the chosen resolution of $0\rlap{.}''3$, the inter-knot regions are fairly well-resolved from the knots themselves. We therefore show spectra fitted both in the inter-knot regions and at the brightness peaks.

The series of spectra highlights the development of the spectra along the jet: an increase in luminosity coupled with a decrease in cutoff frequency. It is seen that the ultra-violet flux dominates the jet luminosity in regions A, B1 and B2 ($12'' \lesssim r \lesssim 15\rlap{.}''5$; *cf.* Fig. 3.3). Beyond this, the infrared emission dominates the high-frequency energy output. The peak of the spectral energy density as determined by Fit B moves to longer and longer wavelengths outward along the jet. The radio emission dominates only at the radio hot spot, at radii $r \gtrsim 20\rlap{.}''5$. There is no optical or ultraviolet emission detected from the precursor H1 to the hot spot, so that the cutoff frequency is again an upper limit.

Compared to earlier studies at lower resolution, the contributions from individual knots are now clearly separated from each other. While previously only knot A showed a spectrum without a cutoff, it is now seen that there is no cutoff at the brightness peaks in B1 and B2, either. There is, however, a cutoff in the transition A-B1, between these peaks. The presence of a cutoff in the southern part of B1 can already be inferred from the fact that it is optically much fainter (*cf.* Sect. 3.4). The differences between the spectra fitted at the brightness peaks and the regions connecting them are nowhere else as drastic.

As noted above, the high-frequency spectral indices in these inner regions are flatter than the radio and radio-infrared spectral indices, and the high-frequency part of the observed spectrum is not fitted well by the assumed single-electron population spectrum. The high-frequency points thus may contain emission contributed from a different electron population or by a different emission mechanism. These possibilities are discussed in detail in Sect. 5.2, together with the differences between Fits A and B in further regions of the jet.

The assumed spectral shape permits very good fits to the observations for B3 and C1, with little difference between Fits A and B. The differences between both fits become significant for the remainder of the jet, in regions C2 and D ($17\rlap{.}''5 \lesssim r \lesssim 20\rlap{.}''5$). Fit A, in which the fitted spectrum does not pass through the near-infrared flux point in these regions, has a significantly lower luminosity in the range 10^{11} Hz–10^{14} Hz than Fit B.
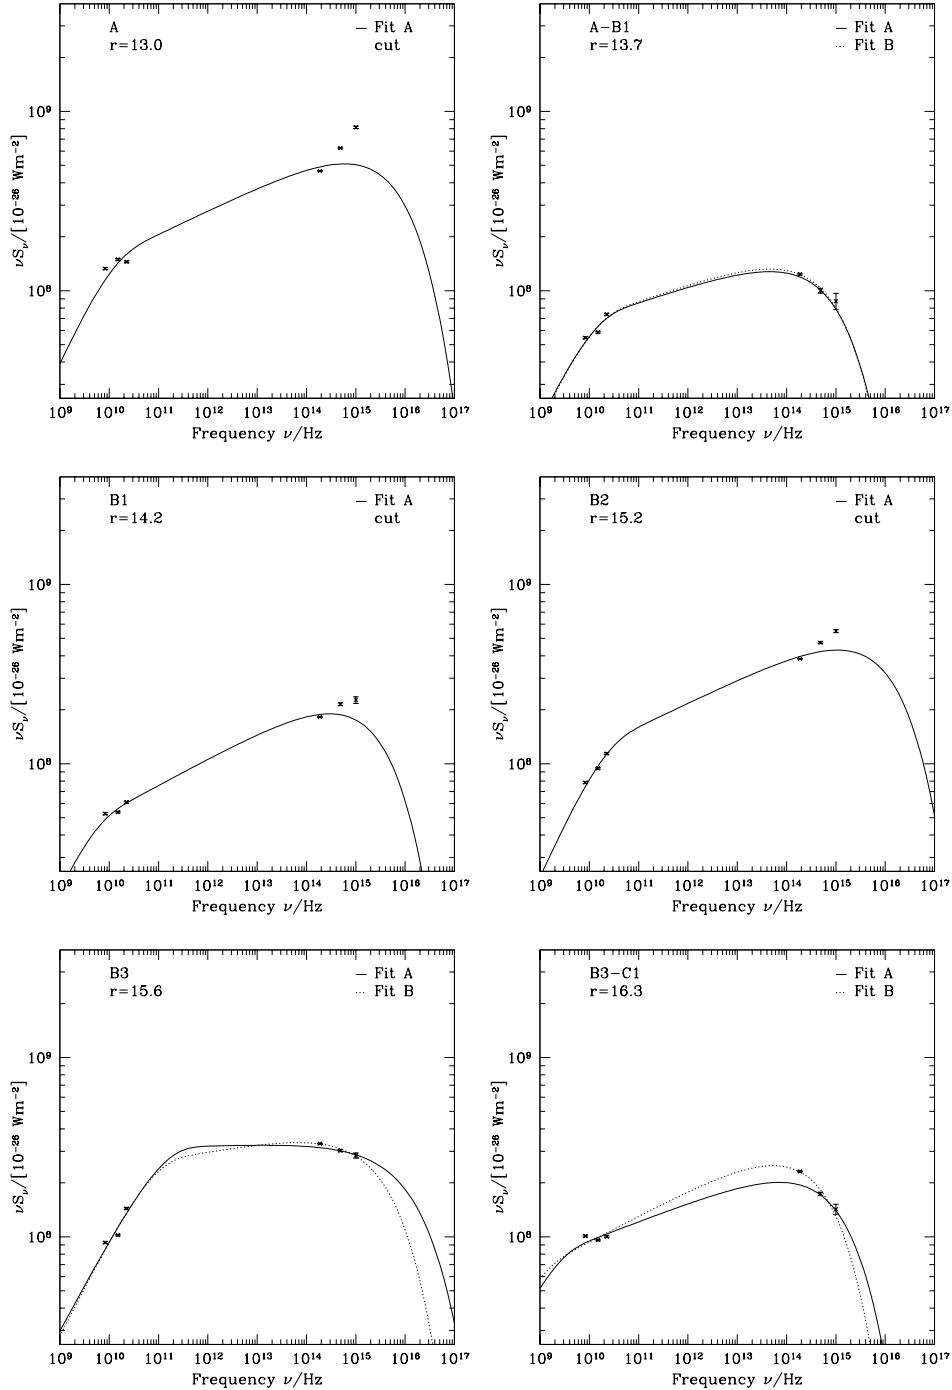
Figure 4.3: Observed data points with fitted spectra for the points shown in Fig. 4.2. To account for the observed flattening of the spectrum towards the ultraviolet, Fit A assumes a contamination in the infrared, so that the cutoff is determined by the optical-ultraviolet spectral index there, while the cutoff in Fit B is determined by the infrared-optical spectral index. Those spectra which needed the inclusion of an artificial high-frequency data point to fit a cutoff are labelled "cut"; for these, Fit A and Fit B are identical.
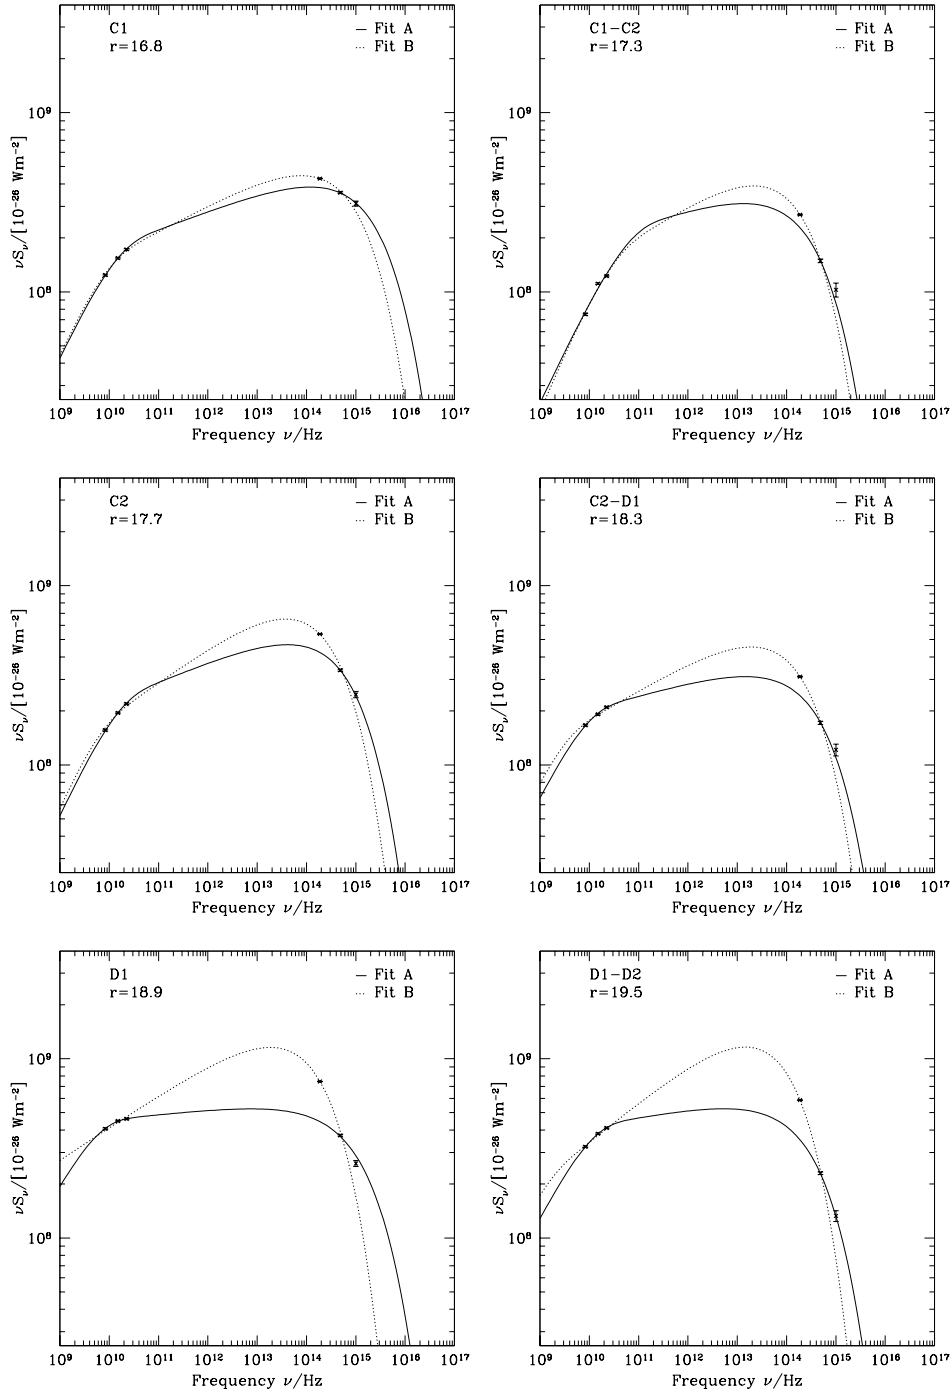
Figure 4.4:   As Fig. 4.3, continued

The synchrotron emission at the cutoff frequency is emitted by those particles with the largest energy. The maximum particle energy can therefore be computed from the cutoff frequency through the synchrotron *characteristic frequency*, the frequency around which most of the synchrotron emission of an electron of energy $\gamma m_{\mathrm{e}} c^2$ in a magnetic field of flux density
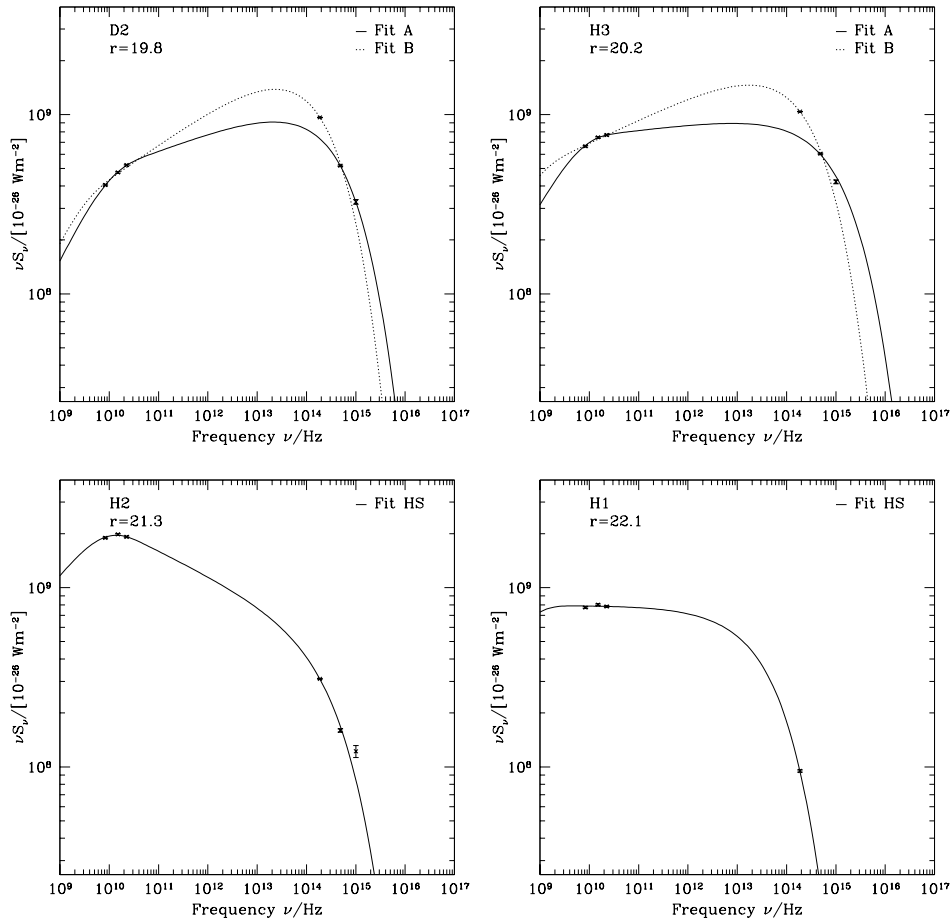
Figure 4.5: As Fig. 4.3, continued

B is emitted:

$$\nu_{\mathrm{c}} = 4.2 \times 10^{15} \left(\frac{\gamma}{10^7}\right)^2 \left(\frac{B \sin \psi}{\mathrm{nT}}\right) \mathrm{Hz} \tag{4.1}$$

(see Sect. D.1.1). This calculation requires the knowledge of the magnetic field in the source.

In the absence of any other estimate, the magnetic field for the jet can only be estimated by making use of the minimum-energy assumption. We therefore present the derivation of the value of the minimum-energy magnetic field for the type of spectra described here.

4.3 Minimum energy estimates for synchrotron sources

A given synchrotron luminosity can be produced by a large number of particles in a weak magnetic field, or a small number of particles in a strong magnetic field [56, *e. g.*]. Observationally, the two cases cannot be discerned. However, the total energy stored in the form of relativistic particles and magnetic field energy is very large for the extreme cases. Between them, there is a firm minimum value for the energy density necessary to generate a given synchrotron luminosity. This energy density *can* be deduced from observations with the help of

assumptions about the source geometry and the range of frequencies over which synchrotron emission is emitted.

Pacholczyk [78] presents a clear derivation of the minimum energy density of a synchrotron source for the assumption of a power-law electron energy distribution in the source (basic facts from synchrotron theory are recapitulated in Sect. D.1). Since all derivations in the literature only deal with simple power laws, we consider in detail how to obtain the minimum energy estimate for the case of an electron energy distribution with a break. In the following, we will present a summary of the derivation and outline suitable approximations.

Taking the minimum-energy magnetic field as estimate for the true magnetic field introduces a local correlation between the magnetic field strength and the energy density in particles. This correlation provides the most efficient way to produce a given synchrotron luminosity and may be envisaged as being caused by magneto-hydrodynamical processes in which particles are both tied to magnetic field lines and create them by their motion. As long as no detailed microphysical feedback process can be identified which maintains this correlation on all scales, the interpretation of any result derived with the help of the minimum-energy argument must not be stretched beyond what is warranted under this assumption.

The Heavens & Meisenheimer [34] model for the emission from hot spots does enable an independent determination of the hot spot field strength from physical parameters of the source giving rise the observed spectrum. Although the physical hot spot model does not apply to the body of the jet, its synchrotron emission can be described by the same spectral parameters as that from hot spots (Sect. 4.1). Since the present observations do include the hot spot region, the minimum-energy field for the hot spot can be compared to that inferred from the physical hot spot model using both the present and previous spectral fits [68], to assess its validity for the hot spot.

We write the electron energy distribution function as $n(\gamma)\mathrm{d}\gamma$, the number density of electrons with Lorentz factor in the range $\gamma \to \gamma + \mathrm{d}\gamma$. If the electron energy distribution follows a single straight power law, $n(\gamma) = n(\gamma_0)\left(\frac{\gamma}{\gamma_0}\right)^{-p}$, the source has a spectral flux density of the form $S(\nu_0)\left(\frac{\nu}{\nu_0}\right)^{\alpha}$, with $\alpha = -(p-1)/2$ (Eqn. D.18). If the electron population's minimum and maximum Lorentz factors are given by γ_{min} and γ_{max}, respectively, synchrotron emission is assumed to be observable in the range $\nu_{\mathrm{min}} \to \nu_{\mathrm{max}}$, ν and γ being related through Eqn. 4.1. In the following, we shall always use this relation to convert electron energies into observing frequencies.

4.3.1 Derivation

The derivation of the minimum energy density proceeds by expressing the energy density contributed both by magnetic fields and relativistic particles in terms of the unknown magnetic field and observed quantities and then minimising the total with respect to the magnetic field.

The magnetic energy density in the source is simply $\frac{B^2}{2\mu_0}$. The energy density in the source due to electrons is obtained through the integral

$$U_{\mathrm{el}} = \int_{\gamma_{\mathrm{min}}}^{\gamma_{\mathrm{max}}} \gamma m_{\mathrm{e}}c^2 n(\gamma)\mathrm{d}\gamma \qquad (4.2)$$

$$= m_{\mathrm{e}}c^2\gamma_0^p n(\gamma_0) \int_{\gamma_{\mathrm{min}}}^{\gamma_{\mathrm{max}}} \gamma^{-p+1}\mathrm{d}\gamma. \qquad (4.3)$$

To account for the energy in other relativistic particles (positrons or ions), it is customary to write the total energy density in particles $U_{\mathrm{part}} = (1 + k)U_{\mathrm{el}}$, so that k is the ratio of energy
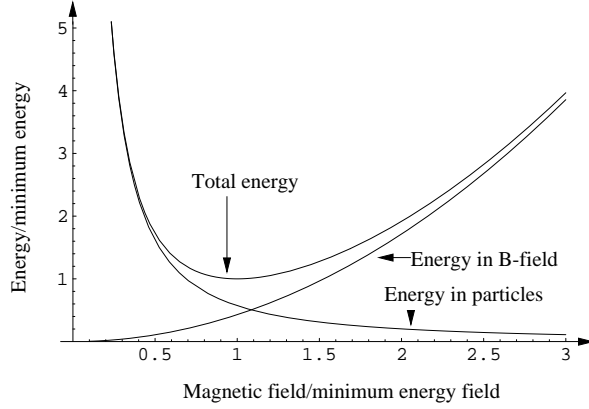
Figure 4.6: Illustration of the minimum-energy condition for synchrotron sources. The total energy content of the source (relative to its minimum value) is shown as function of magnetic field (the source conditions are arbitrarily chosen such that the minimum energy condition is fulfilled at unity). The lower two lines show the separate contributions due the magnetic field energy and that stored in particles.

in other particles to energy in electrons. Since $n(\gamma_0)$ is unknown, it is solved for in terms of the source's observed luminosity L, by expressing L in terms of $n(\gamma)$, using Eqn. D.4:

$$L = \phi V \int_{\gamma_{\min}}^{\gamma_{\max}} \left(-\frac{\mathrm{d}E}{\mathrm{d}t}\right) n(\gamma)\mathrm{d}\gamma \tag{4.4}$$

$$= \phi V \int_{\gamma_{\min}}^{\gamma_{\max}} \frac{4}{3}\sigma_{\mathrm{T}} c U_{\mathrm{mag}} \gamma^2 n(\gamma)\mathrm{d}\gamma \tag{4.5}$$

$$= n(\gamma_0)\phi V \frac{4}{3}\sigma_{\mathrm{T}} c U_{\mathrm{mag}} \gamma_0^p \int_{\gamma_{\min}}^{\gamma_{\max}} \gamma^{-p+2}\mathrm{d}\gamma, \tag{4.6}$$

where ϕV is the volume occupied by emitting particles in the source, so that ϕ is the filling fraction.

On the other hand, the source's luminosity can be inferred from the observed brightness. In the general case, relativistic Doppler effects and the dependence of the emissivity on the angle between line of sight and local magnetic field have to be taken into account.

The simplest case for the field configuration is a fully tangled field, resulting in an isotropic emissivity (*cf.* Sect. 3.3). Retaining only the Doppler factor, the observed luminosity is

$$L_{\mathrm{obs}} = 4\pi d_{\mathrm{L}}^2 \int_{\nu'_{\min}}^{\nu'_{\max}} \mathcal{D}^3 S'(\nu')\mathrm{d}\nu', \tag{4.7}$$

where $S'(\nu')$ is the source-frame brightness at source-frame emission frequency ν' related to the observed frequency by ν by $\nu = \mathcal{D}\nu'$, d_{L} is the luminosity distance, given by Eqn. D.22 [51]. The influence of the Doppler beaming on the value of the maximum particle energy inferred here turns out to be negligible [64, 71]. We therefore assume $\mathcal{D} = 1$ in the following.

All the Lorentz factors in Eqns. 4.3 and 4.6 are expressed in terms of the magnetic field and the frequencies appearing in Eqn. 4.7 through Eqn. 4.1. By equating Eqn. 4.7 and Eqn. 4.6, the unknown $n(\gamma_0)$ can be solved for and inserted into Eqn. 4.3. The result is an expression for U_{el} in terms of observables and the unknown magnetic field. The total energy density in the source is then

$$U_{\mathrm{tot}}(B) = (1+k)U_{\mathrm{el}}(B) + \frac{B^2}{2\mu_0},$$

with $U_{\mathrm{el}}(B)$ a falling function of B and $U_{\mathrm{mag}} \propto B^2$ a rising function of B. Minimising $U_{\mathrm{tot}}(B)$ yields the minimum energy density, and the minimum-energy field B_{min} (Fig. 4.6). We discuss here the assumptions made in using this magnetic field estimate.

Firstly, the ratio k of energy in other relativistic particles to energy in relativistic electrons is unknown. If the jet is an electron-positron jet, $k = 1$ since positrons are accelerated in the same way as electrons. If the charge-balancing particles are protons or even heavier ions, the value of k depends on details of the injection and acceleration process (*cf.* Sect. 1.1.3). A typical number found in cosmic-ray particles is $k \approx 100$. We choose $k = 1$ which is correct to 30% for all values of k between the extremes $k = 0$ and $k = 100$. Secondly, the source extent is only observable in the plane of the sky, so symmetry arguments are usually used to infer the source extent along the line of sight, and indeed have been used here (Sect. 3.3). Thirdly, the filling fraction ϕ is unknown, that is, the part of the source's volume actually taken up by relativistic particles. For the magnetic field, no filling factor is usually written in. Finally, homogeneity of both magnetic field strength and electron density across the source have been assumed throughout.

Further assumptions are usually made in choosing $\nu_{\mathrm{min}}, \nu_{\mathrm{max}}$. To discuss these, we evaluate the above integrals and substitute for p in terms of α. Equation 4.3 becomes

$$U_{\mathrm{el}} = \begin{cases} m_{\mathrm{e}}c^2 n(\gamma_0)\gamma_0^2 \times \ln\left(\frac{\gamma_{\mathrm{max}}}{\gamma_{\mathrm{min}}}\right) & \text{if } \alpha = -0.5 \\ m_{\mathrm{e}}c^2 n(\gamma_0)\frac{\gamma_0^{1-2\alpha}}{1+2\alpha}\left[\gamma_{\mathrm{max}}^{1+2\alpha} - \gamma_{\mathrm{min}}^{1+2\alpha}\right] & \text{if } \alpha \neq -0.5 \end{cases} \tag{4.8}$$

Similarly, we obtain for Eqn. 4.6 that

$$L = \begin{cases} n(\gamma_0)\phi V \frac{4}{3}\sigma_{\mathrm{T}}c U_{\mathrm{mag}}\gamma_0^3 \times \ln\left(\frac{\gamma_{\mathrm{max}}}{\gamma_{\mathrm{min}}}\right) & \text{if } \alpha = -1 \\ n(\gamma_0)\phi V \frac{4}{3}\sigma_{\mathrm{T}}c U_{\mathrm{mag}}\frac{\gamma_0^{1-2\alpha}}{2+2\alpha}\left[\gamma_{\mathrm{max}}^{2+2\alpha} - \gamma_{\mathrm{min}}^{2+2\alpha}\right] & \text{if } \alpha \neq -1 \end{cases} \tag{4.9}$$

Finally, assuming that there is no Doppler boosting, Eqn. 4.7 becomes

$$L_{\mathrm{obs}} = \begin{cases} 4\pi d_{\mathrm{L}}^2 S(\nu_0)\nu_0 \ln\left(\frac{\nu_{\mathrm{max}}}{\nu_{\mathrm{min}}}\right) & \text{if } \alpha = -1 \\ 4\pi d_{\mathrm{L}}^2 S(\nu_0)\frac{\nu_0^{-\alpha}}{1+\alpha}\left[\nu_{\mathrm{max}}^{1+\alpha} - \nu_{\mathrm{min}}^{1+\alpha}\right] & \text{if } \alpha \neq -1 \end{cases} \tag{4.10}$$

Note that equating Eqns. 4.9 and 4.10 will lead to the integral terms (logarithms and terms in square brackets, respectively) dropping out, since both integrals are identical up to a magnetic-field dependent normalisation factor.

For synchrotron sources without high-frequency emission, neither ν_{min} nor ν_{max} is usually known. For typical values of the spectral index $\alpha \approx 1$, the choice of ν_{max} is unimportant, but the choice of ν_{min} dominates the value found for U_{el} in Eqn. 4.8. It has been customary in the radio community to set $\nu_{\mathrm{min}} = 10\,\mathrm{MHz}$. As argued by Myers & Spangler [70], using a fixed value for γ_{min} would be more physical. This again is not straightforward: for once, ν_{min} needs to be expressed in terms of B and γ_{min} in the above equations, while conversely γ_{max} is expressed in terms of B and ν_{max}, resulting in a more complicated dependence of U_{el} on B. On the other hand, choosing the correct value for γ_{min} may be as difficult as choosing ν_{min}. We return to this point further down.

Having observed the high-frequency cutoff of the synchrotron spectrum and determined a spectral break for most part of the jet, we can make use of this information and find suitable approximations. The break in the electron spectrum of course means that the above integrals have to be performed over the broken power-law. This case is considered in full detail in the following section.

4.3.2 Minimum energy estimates for broken power laws

To account for the break in the electron spectrum at Lorentz factor γ_b, we write the electron distribution function as

$$n(\gamma) = \begin{cases} n(\gamma_0)\left(\frac{\gamma}{\gamma_0}\right)^{-p_{\text{low}}} & \text{for } \gamma < \gamma_b \\ n(\gamma_0)\left(\frac{\gamma}{\gamma_0}\right)^{-p_{\text{high}}} & \text{for } \gamma \geq \gamma_b \end{cases} \tag{4.11}$$

Similarly, the observed spectrum is approximated by

$$S(\nu) = \begin{cases} S(\nu_0)\left(\frac{\nu}{\nu_0}\right)^{\alpha_{\text{low}}} & \text{for } \nu < \nu_b \\ S(\nu_0)\left(\frac{\nu}{\nu_0}\right)^{\alpha_{\text{high}}} & \text{for } \nu \geq \nu_b \end{cases} \tag{4.12}$$

The spectra determined in Sect. 4.1 all show their break at $\nu_b \approx 10^{10}\,\text{Hz}$ (see discussion below), from $\alpha_{\text{low}} \approx -0.5$ to $\alpha_{\text{high}} \approx -1$, corresponding to $p_{\text{low}} \approx 2$ and $p_{\text{high}} \approx 3$, respectively. Even the separate hot spot fit does not have a significantly steeper best-fit spectral index. All integrals can therefore be approximated by setting $\alpha_{\text{low}} = -0.5$ and correspondingly $\alpha_{\text{high}} = -1$. It is useful to write the electron energy spectrum in terms of $\gamma_0 = \gamma_b$, and correspondingly the observed spectrum in terms of $\nu_0 = \nu_b$. With this choice, Eqns. 4.8, 4.9 and 4.10 below the break become

$$U_{\text{el,low}} = m_e c^2 n(\gamma_b)\gamma_b^2 \ln\left(\frac{\gamma_b}{\gamma_{\text{min}}}\right), \tag{4.13}$$

$$L_{\text{low}} = n(\gamma_0)\phi V \frac{4}{3}\sigma_T c U_{\text{mag}}\gamma_b^3 \left[1 - \frac{\gamma_{\text{min}}}{\gamma_b}\right], \tag{4.14}$$

$$\text{and } L_{\text{obs,low}} = 8\pi d_L^2 S(\nu_b)\nu_b \left[1 - \left(\frac{\nu_{\text{min}}}{\nu_b}\right)^{\frac{1}{2}}\right]. \tag{4.15}$$

With $\nu_b \approx 10^{10}\,\text{Hz}$, we can confidently use the approximation $\gamma_{\text{min}}/\gamma_b = (\nu_{\text{min}}/\nu_b)^{-\frac{1}{2}} \ll 1$ in Eqns. 4.14 and 4.15, because these integrals converge for $\alpha \approx -0.5$. Since Equation 4.13 depends only logarithmically on this ratio, the exact choice of ν_{min} is unimportant here as well. In light of this, we will follow tradition and choose $\nu_{\text{min}} = 10\,\text{MHz}$.[1] Similarly, the integrals for $\alpha_{\text{high}} = -1$ become

$$U_{\text{el,high}} = m_e c^2 n(\gamma_b)\gamma_b^2 \left[1 - \frac{\gamma_b}{\gamma_c}\right], \tag{4.16}$$

$$L_{\text{high}} = n(\gamma_b)\phi V \frac{4}{3}\sigma_T c U_{\text{mag}}\gamma_b^3 \ln\left(\frac{\gamma_c}{\gamma_b}\right), \tag{4.17}$$

$$\text{and } L_{\text{obs,high}} = 4\pi d_L^2 S(\nu_b)\nu_b \ln\left(\frac{\nu_c}{\nu_b}\right). \tag{4.18}$$

Again, we note that $\gamma_b/\gamma_c \approx 10^{-2.5} \ll 1$.

Making use of the approximations and substituting for Lorentz factors $\gamma^2 = K\nu B^{-1}$, where $K^{-1} = 4.2\times10^{10}\,\text{T}^{-1}\,\text{Hz}$ is the numerical constant from Eqn. 4.1, we obtain the following

---

[1] Neumann [71] chose $\nu_{\text{min}} = 408\,\text{MHz}$ as this is the lowest frequency at which the jet has been detected; compared to 10 MHz, this would decrease the minimum-energy field by 20% and the minimum energy density by 30%, far less than the influence of the choice of ϕ or k.

equations by adding up the contributions below and above the break:

$$U_{\rm el} \;=\; m_e c^2 n(\gamma_{\rm b}) K \nu_{\rm b} B^{-1} \left(1 + \frac{1}{2}\ln\frac{\nu_{\rm b}}{\nu_{\rm min}}\right), \tag{4.19}$$

$$L \;=\; n(\gamma_{\rm b})\phi V \frac{4}{3}\sigma_{\rm T} c U_{\rm mag} \nu_{\rm b}^{\frac{3}{2}} K^{\frac{3}{2}} B^{-\frac{3}{2}} \left(1 + \frac{1}{2}\ln\frac{\nu_{\rm c}}{\nu_{\rm b}}\right), \tag{4.20}$$

$$L_{\rm obs} \;=\; 8\pi d_{\rm L}^2 S(\nu_{\rm b})\nu_{\rm b} \left(1 + \frac{1}{2}\ln\frac{\nu_{\rm c}}{\nu_{\rm b}}\right). \tag{4.21}$$

Equating Eqns. 4.20 and 4.21 yields

$$n(\gamma_{\rm b}) = \frac{12\pi d_{\rm L}^2 S(\nu_{\rm b})}{\phi V}\frac{\mu_0}{\sigma_{\rm T} c} K^{-\frac{3}{2}} (B\nu_{\rm b})^{-\frac{1}{2}} \tag{4.22}$$

The integration limits have dropped out and at fixed synchrotron luminosity, $n(\gamma) \propto B^{-\frac{1}{2}}$. Formally, this result can be used to derive the total number density of relativistic electrons in the source:

$$n_{\rm el} = \frac{d_{\rm L}^2}{\phi V}\frac{12\pi\mu_0}{\sigma_{\rm T} c} K^{-1} B^{-1} \left(\frac{\nu_{\rm b}}{\nu_{\rm m}}\right)^{-\frac{1}{2}}. \tag{4.23}$$

We stress that this is a formal result with limited physical meaning because this number depends strongly on the chosen value for the lower frequency cutoff, which is dominated by low-energy electrons for any monotonically falling electron distribution. At this place, the local correlation between magnetic field and relativistic particles becomes explicitly apparent.

After inserting Eqn. 4.22 into Eqn. 4.19, we obtain the total energy density in the source:

$$\begin{aligned}
U_{\rm tot}(B) &= (1+k)\frac{12\pi d_{\rm L}^2}{\phi V}\frac{\mu_0 m_e c}{\sigma_{\rm T} K^{\frac{1}{2}}} S(\nu_{\rm b})\nu_{\rm b}^{\frac{1}{2}} \left(1 + \frac{1}{2}\ln\frac{\nu_{\rm b}}{\nu_{\rm min}}\right) B^{-\frac{3}{2}} + \frac{B^2}{2\mu_0} \\
&= C_{\rm source} C_{\rm nat} f\left(S(\nu_{\rm b}),\nu_{\rm b}\right) B^{-\frac{3}{2}} + \frac{B^2}{2\mu_0}, \tag{4.24}
\end{aligned}$$

subsuming the source-related constants in the first fraction into $C_{\rm source}$, the term containing natural constants only into $C_{\rm nat}$, and the observational terms into $f\left(S(\nu_{\rm b}),\nu_{\rm b}\right)$. $U_{\rm tot}(B)$ attains a minimum at

$$B_{\rm min} = \left(\frac{3}{2}\mu_0 C_{\rm s} C_{\rm n} f\left(S(\nu_{\rm b}),\nu_{\rm b}\right)\right)^{\frac{2}{7}}, \tag{4.25}$$

and the minimum energy density can be written as

$$U_{\rm tot}(B_{\rm min}) = \frac{7}{3}\frac{B_{\rm min}^2}{2\mu_0}, \tag{4.26}$$

so that at minimum energy, the energy in particles is 4/3 of the energy stored in the magnetic field. Since similar amounts of energy are then present in relativistic particles and magnetic fields, the condition is often referred to as *equipartition*.

The value for the minimum energy density, and hence the minimum-energy field, depends on the location of the spectral break because the luminosity is dominated by the steep-spectrum region between $\nu_{\rm b}$ and $\nu_{\rm c}$. The break is not actually detected in the present data set and the break frequency $\nu_{\rm b}$ has therefore been constrained to lie near the observed data points, i.e., $\nu_{\rm b} \approx 10^{10}\,{\rm Hz}$ (Sect. 4.2). We consider whether this constraint significantly influences the resulting value for the minimum-energy field.
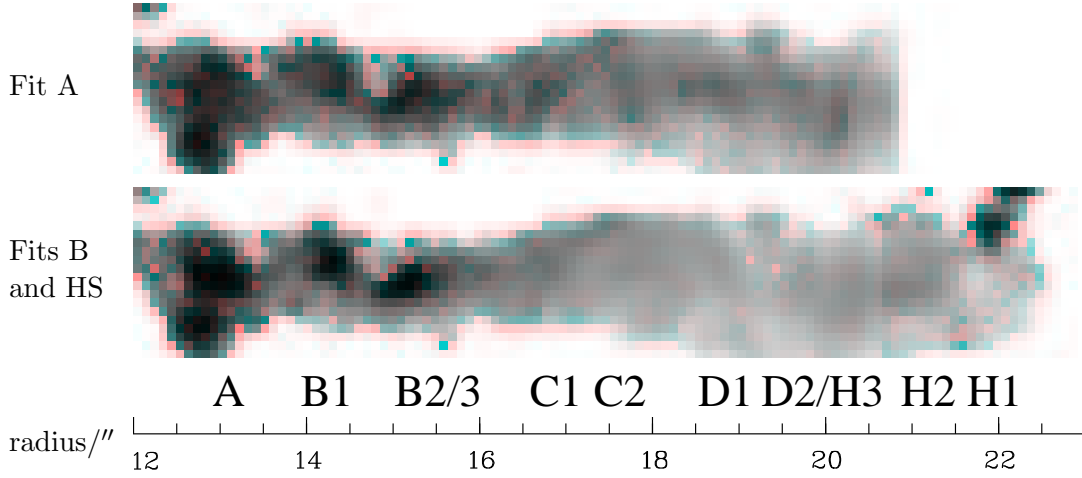
Figure 4.7: Maps of the cutoff frequency. The greyscale is logarithmic and runs from $10^{13}\,\mathrm{Hz}$ to $10^{17}\,\mathrm{Hz}$. The values fitted in A and B2 are lower limits. As expected, Fit A bears a closer resemblance to the optical-ultraviolet spectral index map, while Fit B is dominated by the infrared-optical spectral index map (*cf.* Fig. 3.12).

At least the hot spot spectrum does show a break around $10^9\,\mathrm{Hz}$ in the spectra presented by Meisenheimer et al. [64]. The spectrum for the remainder of the jet needs to flatten at some low frequency, as otherwise the luminosity would grow without bounds. If the true break frequency was as low as $10^6\,\mathrm{Hz}$ instead of $10^{10}\,\mathrm{Hz}$, the luminosity would only change by a factor of about two. The influence of the value of the cutoff frequency is therefore negligible.

If the jet lies not in the plane of the sky but is inclined to the line of sight by an angle i, all lines of sight passing through the jet, and hence the total jet volume, are longer by a factor $1/\sin i$ compared to the side view. This treatment ignores edge effects at the end of the jet. As can be seen from Eqns. 4.25 and 4.26, the minimum-energy field varies as $V^{-\frac{2}{7}}$. A line-of-sight angle $i \approx 45°$ has been inferred for the flow into the hot spot from independent considerations of the jet's polarisation change there and the hot spot's morphology [16, 68]. If the jet is at the same angle, the values in Tab. 3.1 need to be scaled up by $1/\sin 45° \approx 1.4$. Hence, the minimum-energy field needs to be scaled down by about 10%, and correspondingly the maximum energy up by 10%. Since we do not expect the minimum energy to be accurate to this level, the presented values will not include this correction.

4.4 Run of ν_{c}, B_{min}, and γ_{max} along the jet

For each photometry aperture (or pixel in Fig. 3.2), the values for $\nu_{\mathrm{b}}, S(\nu_{\mathrm{b}})$ and ν_{c} from the fitted spectra together with the appropriate volume from Tab. 3.1 and the assumed value of ν_{min} are used to calculate B_{min} from Eqn. 4.25. With this knowledge, the maximum particle energy γ_{max} can be inferred from the fitted value of ν_{min} (Eqn. 4.1).

We consider the determined values of the cutoff frequency, minimum-energy field and hence maximum particle energy and their relation to the jet morphology. In particular, we will be interested whether it is possible to identify localised acceleration regions in the jet.

We present maps of the fitted cutoff frequency ν_{c} for the three different fits in Fig. 4.7, and its run along the radius vector at position angle 222°2 in Fig. 4.8. No cutoff is observed
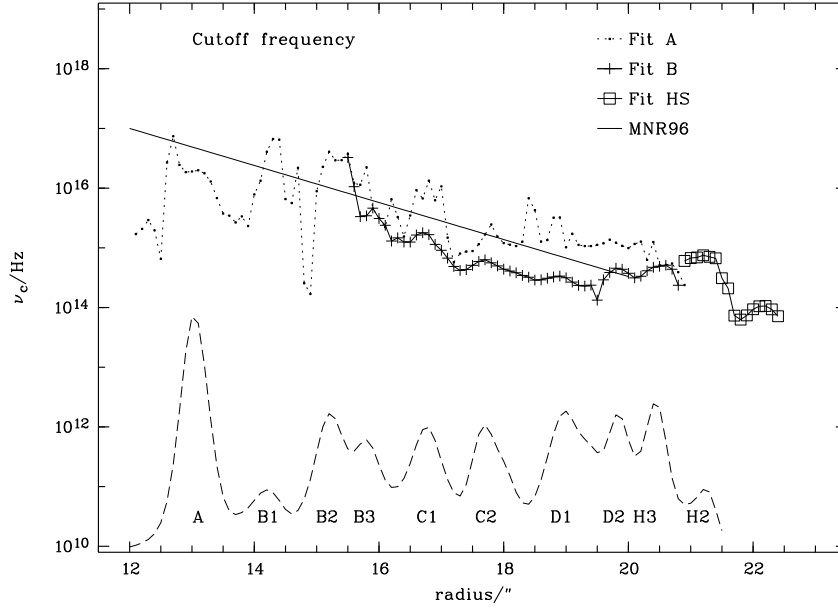
**Figure 4.8:** Run of the fitted cutoff frequency ν_c along the jet. The thin line represents the run of the cutoff frequency at $1\rlap{.}''3$ resolution, parameterised as $\nu_c = 10^{17}\,\mathrm{Hz}\exp(-(r-12'')/1\rlap{.}''4)$ by Meisenheimer et al. [64]. Fit A has a slightly higher ν_c than Fit B. The values near $r \approx 13''$ and $r \approx 15''$ are only lower limits to the true cutoff frequency. Error bars are not shown because the errors on the fitted cutoff frequency are correlated in a complicated manner with the observational flux errors and the assumed spectral shape. Variations in the cutoff frequency are expected to be significant where the variations in the high-frequency spectral indices are significant.

in the regions at $r = 13''$ (A1) and at $15''$ (B2) (*cf.* Sect. 3.5.3), so the fitted values there are only lower limits to the true cutoff frequency. The overall trend is a decrease in ν_c with increasing distance from the core. All variations are rather smooth. As expected from the overall similarity of the spectral indices determined at $0\rlap{.}''3$ resolution and in earlier work at $1\rlap{.}''3$ resolution (Fig. 3.13), the cutoff frequency determined here agrees well with the run of the cutoff frequency $\nu_c = 10^{17}\,\mathrm{Hz}\exp(-(r-12'')/1\rlap{.}''4)$ determined by Meisenheimer et al. [64].

Only small discrepancies arise in the value of the cutoff frequency between Fit A and Fit B (for a discussion, see Sect. 5.2 below). Where there *is* a difference, the value of ν_c in Fit A, in which the cutoff is determined mainly by the optical-UV spectral index, is a factor of 3–10 larger than in Fit B, in which the steeper (by $\Delta\alpha \approx 0.2$) infrared-optical spectral index determines the cutoff frequency. A minimum of ν_c is only reached at the hot spot in Fit HS, which is characterised by a sharp drop in ν_c from $10^6\,\mathrm{Hz}$ at $21\rlap{.}''2$ (optical hot spot position) to $3 \times 10^5\,\mathrm{Hz}$ at $r = 21\rlap{.}''6$.

There are local peaks in the values of the fitted cutoff frequency. Only those peaks are significant which correspond to significant peaks in the high-frequency spectral indices (Fig. 3.14). As we are interested in the variation of the particles' energy along the jet, we first consider the derived minimum-energy field and then discuss the variations of the particle energy.

The variation of the derived minimum-energy field perpendicular to the jet axis is dominated by the assumed geometry. We therefore present only its run along the assumed jet
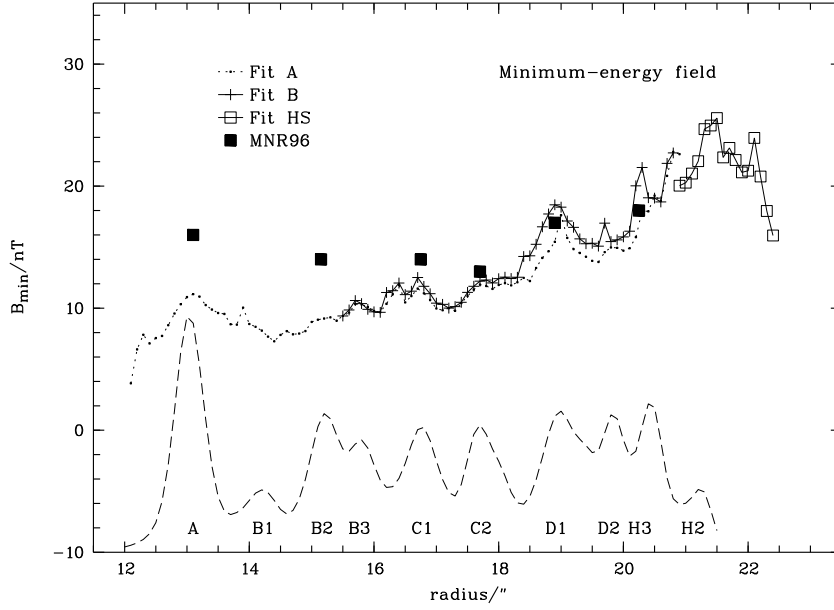
Figure 4.9: Run of the minimum-energy field B_{min} along the jet. The solid points show the values determined for individual regions by Meisenheimer et al. [64]. The overall run corresponds to that of the jet luminosity. The most recent value for the magnetic field determined for the hot spot is $(39^{+24}_{-10})\,\mathrm{nT}$.
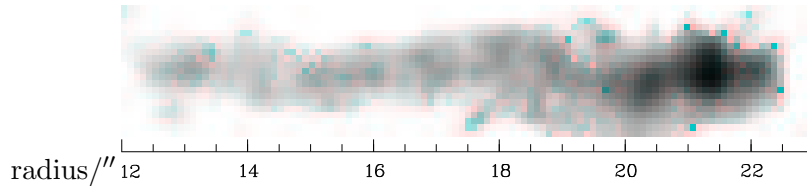

Figure 4.10: Map of the bolometric surface brightness (logarithmic greyscale). The luminosity of regions A and B2 could dominate the luminosity if the synchrotron spectrum extends up to X-rays (*cf.* Sect. 5.3).

axis in Fig. 4.9. It starts at just below $10\,\mathrm{nT}$ at the onset of the optical jet, up to about 20–$25\,\mathrm{nT}$ in the hot spot. This corresponds qualitatively to the run of the bolometric surface brightness (Fig. 4.10). This is expected because for the assumed constant spectral shape and for a constant volume, the minimum-energy field scales as (synchrotron luminosity)$^{2/7}$ (Eqns. 4.24–4.26). It also means that the minimum-energy field could be much higher in regions A–B if the cutoff there is much higher than the lower limit used here.

The minimum-energy field of $\approx 25\,\mathrm{nT}$ determined for the hot spot is lower the magnetic field value determined by Meisenheimer et al. [68], both from spectral fits ($(39^{+24}_{-10})\,\mathrm{nT}$) and from the minimum-energy argument ($(35^{+8}_{-4})\,\mathrm{nT}$) (*cf.* Fig. 4.9). Apart from a slightly higher cutoff frequency determined here, the fit parameters are identical within the errors. The discrepancy of the minimum-energy field value is explained by the differnt values assumed for the ratio k in Eqn. 4.24, chosen as $k = 1$ here but $k = 10$ in [68].

As discussed below, only the order of magnitude of the minimum-energy field matters in the considerations here. We therefore do not discuss its behaviour in detail. This also implies that the dependence of B_{min} on the value of the break frequency can be neglected here
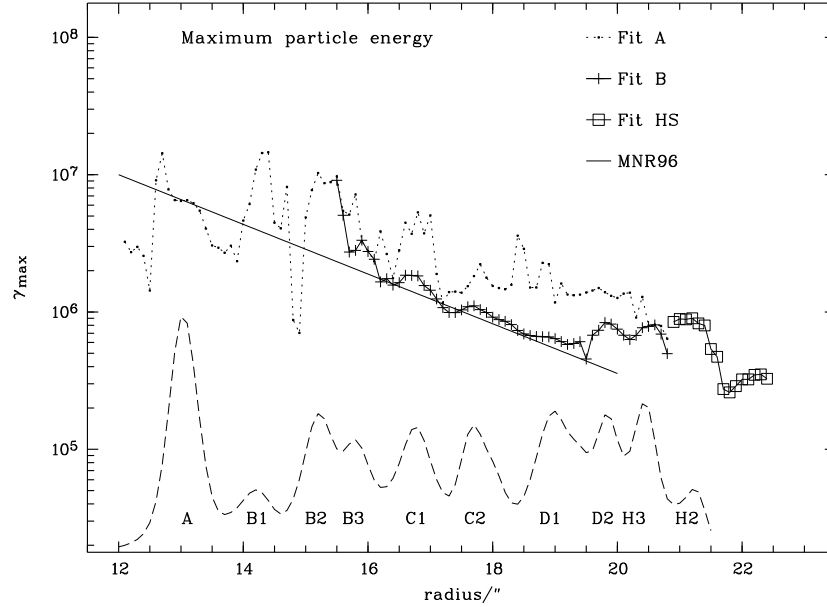
Figure 4.11: Run of the maximum particle Lorentz factor $\gamma_{\max}$ along the jet. The thin line represents the run of $\gamma_{\max}$ parameterised as $\nu_{\rm c} = 10^7 \exp(-(r - 12'')/2''\!.4)$ by Meisenheimer et al. [64]. The overall run of $\gamma_{\max}$ is identical to that of $\nu_{\rm c}$, with the exception of region H1 at $r > 22''$, where $\nu_{\rm c}$ drops but $\gamma_{\max}$ rises outwards.

(Sect. 4.3.2). If the entire jet was operating at the minimum energy condition, the increase in the minimum-energy field would correspond to an increase of the energy density stored in the jet's magnetic fields and relativistic particles by about one order of magnitude. This conclusion is only valid in the absence of strong beaming.

The run of the maximum particle energy inferred from the minimum-energy field and the cutoff frequency is shown in Fig. 4.11. It is noted that this run is very similar to the run of the cutoff frequency in Fig. 4.8; this arises from the relation used to derive the maximum particle energy (Eqn. 4.1):

$$\gamma_{\max} = 10^7 \times \left(\frac{\nu_{\rm c}}{4.2 \times 10^{15} {\rm Hz}} \right)^{\frac{1}{2}} \left(\frac{B}{{\rm nT}} \right)^{-\frac{1}{2}} .$$

The maximum particle energy and cutoff frequency are both plotted logarithmically, that is, we are comparing their order of magnitude. Hence, the run of the two can only differ if the order of magnitude of the magnetic field changes along the jet. This is not the case, as Fig. 4.9 shows. This also implies that although the variations of $B_{\min}$ perpendicular to the jet axis are dominated by the assumed geometry, the variations in $\gamma_{\max}$ are affected not as badly. We therefore show a map of the maximum particle energy in Fig. 4.12 in addition to the development of $B_{\min}$ along the jet axis in Fig. 4.11.

Because of the discussed constancy of the order of magnitude of the minimum-energy field, extrema in $\nu_{\rm c}$ lead to extrema in $\gamma_{\max}$. Significant local extrema in $\gamma_{\max}$ are thus found in the following regions:

- A and B2 ($13''$ and $15''\!.5$), in which both high-frequency spectral indices peak simultaneously at a value significantly flatter than the radio-infrared spectral index
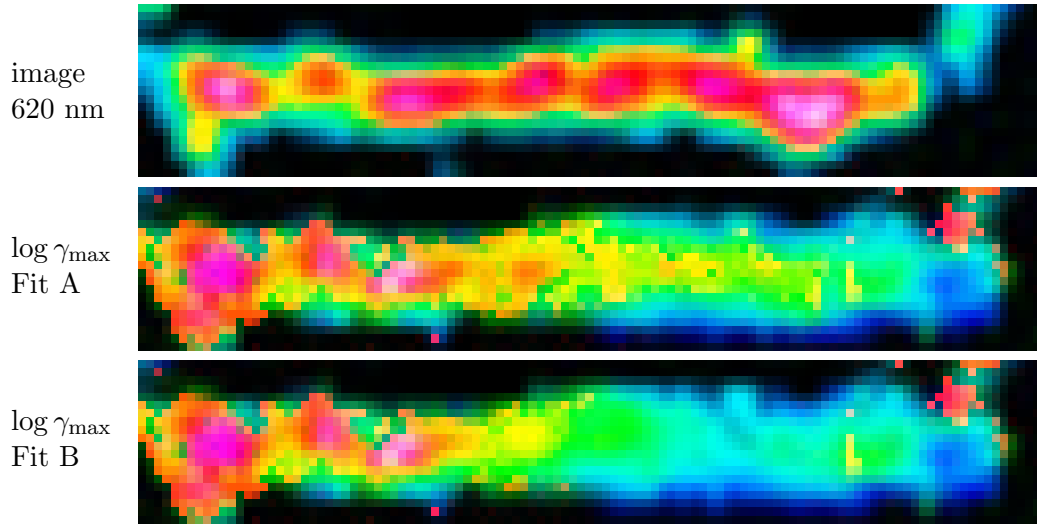
Figure 4.12:   Map of the maximum particle Lorentz factor γ_{max}. The results for Fit HS have been inserted in both. Only in regions A and B2 does γ_{max} show a strong correlation to features of the radio-optical jet morphology. The only differences between Fit A and Fit B occur in regions C–H3, where the discrepancies between the infrared-optical and optical-ultraviolet spectral indices are strongest.

- C1 and C2 (16$''$75 and 17$''$75), where both high-frequency spectral indices show a local peak; here, the infrared-optical spectral index peaks at the brightness maximum, while the optical-UV maxima are slightly offset radially outwards

- The radio hot spot H2, at which all high-frequency spectral indices and the radio-infrared index drop significantly. γ_{max} starts to drop from a value $\gamma_{max} \approx 10^6$ at the optical counterpart to H2 and reaches a plateau beyond the radio hot spot at one-third of the pre-hot spot value. This drop is caused by the absence of optical and ultraviolet emission beyond the radio hot spot in the precursor H1. The detailed run of the cutoff frequency in the hot spot is determined not only by the spectral evolution downstream of the Mach disk, which is located at the highest-frequency emission peak, but also by the effect of telescope resolution and integrating along the line of sight through a the cylindrical emission region inclined at 45°to the line of sight (*cf.* Fig. 3.6 on p. 37).

Thus, with the exception of regions A and B2, there is no strong correlation between γ_{max} and the jet's surface brightness at any wavelength. This is entirely unexpected, since the strong losses giving rise to the observed synchrotron luminosity should give rise to drastic changes of the cutoff frequency as the radiating particles move through regions in which they undergo strong losses.

At radii $> 17''$, the maximum particle energy stays nearly constant at a Lorentz factor of a few times 10^6. The Lorentz factor of the particles emitting optical synchrotron radiation cannot be much lower anywhere in the jet. As noted in all previous studies, synchrotron cooling timescales for these particles are of the order of a few hundred or thousand years [see 30, 64, 87,   *e.g.*]. This conclusion is most stringent for Fit A, in which significant synchrotron emission is presumed to be present at ultraviolet wavelengths, for which electrons have lifetimes of about 1 000 years for the derived equipartition magnetic fields of order 10 nT, with loss scales of a few hundred parsec for the unrealistic ideal case of free streaming at the

speed of light (Eqn. D.12). Since the cutoff frequency is slightly lower in Fit B, slightly longer lifetimes are possible, but not by more than a factor of two even if the true cutoff is assumed to lie closer to the infrared.

The observed spectra are thus consistent with the absence of any synchrotron cooling. The same conclusion could already be drawn from the smooth changes of just the optical spectral index, which should most strongly reflect the synchrotron losses the particles undergo [42].

The smooth changes of the spectral features along the jet indicated a correspondingly smooth variation of the physical conditions along the jet. Our detailed high-resolution study of the jet in 3C 273 set out with the aim to map those regions in which particles are preferentially accelerated. Since the observed cooling is much less drastic than expected, the conclusion is that the *entire* jet is the particle acceleration region. More precisely, the centres of particle acceleration need to be smaller than the beam size employed (0."2 for the detailed study of the optical spectral index map [42], corresponding to 550 kpc) and smoothly distributed along the jet. The detailed mechanism accounting for this observation is unknown; mechanisms which have been proposed include acceleration by reconnection in thin filaments [53] and turbulent acceleration [57]. Both processes manage to maintain the injection spectrum over distances much larger than the loss scales, although the latter so far only maintains cutoff frequencies in the range of 10^{12} Hz–10^{13} Hz, *i. e.,* below the values observed in 3C 273.

Since the cutoff frequency and maximum particle energy are very similar in both performed fits, these conclusions remain true, whichever of the high-frequency spectral indices indicates the true run of the synchrotron continuum.

Chapter 5

Discussion

5.1 Can beaming account for the lack of cooling?

Heinz & Begelman [35] proposed that sub-equipartition magnetic fields combined with mildly relativistic beaming could explain the lack of cooling in the jet of M87 – which is, however, ten times shorter than that of 3C 273. As an alternative to postulating re-acceleration, we consider whether low magnetic field values and beaming could lead to electron lifetimes sufficient to allow electrons to be accelerated at region A to illuminate the entire jet down to the hot spot over a projected extent of $27h_{70}^{-1}$ kpc (the argument would become even more stringent by demanding acceleration only in the quasar core). We consider the electron lifetime against synchrotron and inverse-Compton cooling off cosmic microwave background photons; the synchrotron Self-Compton process is negligible for electrons in the jet [89], as is Compton scattering off the host galaxy's star light at radii $\lesssim 30$ kpc.

The total energy loss rate of an electron with energy E due to synchrotron radiation and inverse-Compton scattering, averaged over many pitch-angle scattering events during its lifetime, is

$$-\frac{\mathrm{d}E}{\mathrm{d}t} = \frac{4}{3}\sigma_{\mathrm{T}}cU_{\mathrm{tot}}\beta^2 \left(\frac{E}{m_ec^2}\right)^2, \qquad (5.1)$$

where $U_{\mathrm{tot}} = U_{\mathrm{CBR}}(z) + U_{\mathrm{mag}}$ is the sum of the energy densities of the background radiation and magnetic field, respectively, and σ_{T} is the Thomson cross-section [55]. We integrate this equation from $E = \infty$ at $t = 0$ to $E(t)$, assuming $\beta = 1$ (appropriate for the highly relativistic electrons required for optical synchrotron radiation) and substitute for the electron's energy $E = \gamma m_ec^2$. Inverting yields the maximum time that can have elapsed since an electron was accelerated, given its Lorentz factor γ [104]:

$$t(\gamma) = \frac{m_ec^2}{\frac{4}{3}\sigma_{\mathrm{T}}cU_{\mathrm{tot}}\gamma}. \qquad (5.2)$$

This is the "electron lifetime", inversely proportional to both the energy density in which the electron has been "ageing", and the electron's own energy. This result is identical to that obtained with the usual *ansatz* $\tau = E/(\mathrm{d}E/\mathrm{d}t)$, which is, however, often wrongly interpreted as a half-life. The derivation presented here makes clear that it is in fact a firm upper limit.

As most of the electron's energy is radiated at the synchrotron characteristic frequency $\propto \gamma^2 B$, we can substitute for γ in Eqn. 5.2 in terms of the observing frequency and the

magnetic field in the source. Hence, Eqn. 5.2 becomes (in convenient units)

$$t_{\text{life}} = \frac{51\,000\,\text{y}}{B_{-9,\text{IC}}(z)^2 + B^2_{-9,\text{jet}}} \left(\frac{B_{-9,\text{jet}}}{\nu_{15}}\right)^{\frac{1}{2}}, \tag{5.3}$$

where $B_{-9,\text{jet}}$ is the magnetic flux density in nT of the jet field, the background radiation energy density has been expressed in terms of an equivalent magnetic field, $B_{-9,\text{IC}} = (1 + z)^2 \times 0.45\,\text{nT}$, and ν_{15} is the observing frequency in $10^{15}\,\text{Hz}$ [104]. Note that as the substituted $\gamma \propto B^{-\frac{1}{2}}$, setting $B_{-9,\text{jet}} = 0$ is now meaningless.

To be fully adequate for electrons in a relativistic jet at cosmological distances, the equation needs to be modified. Firstly, the frequency local to the source is $(1 + z)$ times the observing frequency because of the cosmological redshift. Furthermore, the radiating electron may be embedded in a relativistic flow with bulk Lorentz factor Γ with three consequences: relativistic time dilation and Doppler shift, and boost of the background radiation energy density. The relativistic time dilation enhances the electron lifetime in the jet frame by a factor Γ. The Doppler shift between the emission frequency ν_{int} in the jet frame (equal to the characteristic frequency) and the observation frequency ν_{obs} is given by $\nu_{\text{int}} = \nu_{\text{obs}}\mathcal{D}^{-1}$, where $\mathcal{D}(\Gamma, \theta) = [\Gamma(1 - \beta_{\text{jet}} \cos\theta)]^{-1}$, the Doppler boosting factor for an angle θ to the line of sight [see 40, e. g.]. A relativistic flow perceives the energy density of the background radiation field boosted up by a factor Γ^2 [21]. Inserting these yields

$$t_{\text{life}}(\mathcal{D}, z) = \frac{\Gamma\mathcal{D}^{\frac{1}{2}} \times 51\,000\,\text{y}}{(\Gamma B_{-9,\text{IC}})^2 + B^2_{-9,\text{jet}}} \left(\frac{B_{-9,\text{jet}}}{(1 + z)\nu_{15}}\right)^{\frac{1}{2}}. \tag{5.4}$$

The electron lifetime attains a maximum value at a certain value of the jet's magnetic field (Fig. 5.1) [104]. On either side of the maximum, the lifetime is decreased by larger losses suffered *before* it is observed: for a higher magnetic field in the jet, the synchrotron cooling is more rapid ($t \propto 1/U_{\text{tot}}$, Eqn. 5.2). A lower jet field requires an electron of higher Lorentz factor for emission at the given frequency, which also suffers more rapid losses ($t \propto 1/\gamma$, Eqn. 5.2). By differentiation of Eqn. 5.4, the value of the *maximum* lifetime is

$$t_{\text{max}}(B_{\text{IC}}, \mathcal{D}, z) = \sqrt{\frac{\mathcal{D}}{\Gamma}} \frac{29\,000\,\text{y}}{B^{\frac{3}{2}}_{-9,\text{IC}} \left[(1 + z)\nu_{15}\right]^{\frac{1}{2}}}. \tag{5.5}$$

Note that the largest possible value for the factor $\sqrt{\mathcal{D}/\Gamma}$ is $\sqrt{2}$. t_{max} is a firm upper limit for the lifetime of a synchrotron-radiating electron from a source at redshift z in a flow of bulk Lorentz factor Γ, whatever the magnetic field strength in the source. It is deduced only from the fact that synchrotron emission is observed at a certain frequency, and that electrons which can radiate at this frequency suffer drastic energy losses either by synchrotron or by inverse-Compton cooling between the time of acceleration and the time of emission. The only additional assumptions are rapid pitch-angle scattering and a homogeneity of conditions throughout the electron's lifetime.

The milli-arcsecond VLBI jet close to the core has a line-of-sight angle near $10°$ and a bulk Lorentz factor near 10 [1]. A line-of-sight angle $\theta \approx 45°$ has been inferred for the flow into the hot spot from independent considerations of the jet's polarisation change there and the hot spot's morphology [16, 68]. We have plotted the lifetime of an electron responsible for emission from the jet in 3C 273 observed at 300 nm as function of the jet's magnetic field
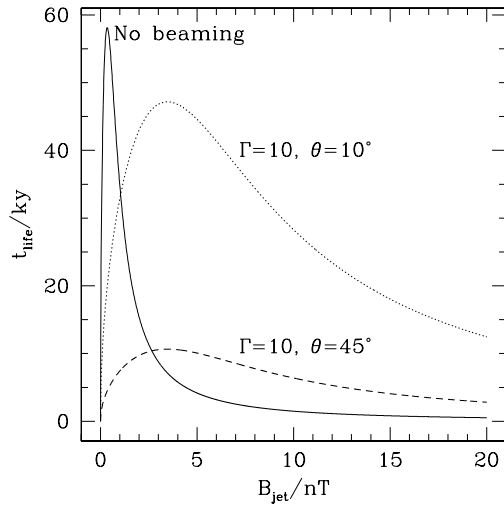
Figure 5.1: Maximum age against synchrotron and inverse-Compton cooling (off microwave background photons at 3C 273's redshift) of an electron radiating at observed UV wavelengths plotted against the jet magnetic field. Solid line, no beaming; dashed line, relativistic Doppler boosting with $\Gamma = 10$ and a line-of-sight angle $\theta = 45°$.

in Fig. 5.1 for the extreme cases of no beaming in the optical jet and beaming identical to that in the VLBI jet with $\Gamma = 10, \theta = 10°$, and for an intermediate case with $\Gamma = 10, \theta = 45°$ (though note that $\Gamma = 10, \theta = 10°$ is likely to be unrealistic as there is a difference in position angle between the VLBI jet at 244° and the arcsecond jet at 222°).

The equipartition flux densities derived above lie in the range of 8–40 nT (Sect. 4.4), leading to maximum ages of a few hundred years, less than the light travel time from one bright region to the next. In the absence of beaming effects, the largest possible lifetime for electrons in 3C 273 from Eqn. 5.5 is about 58 000 y. The "boosted lifetime" can be at most $\sqrt{2}$ larger than this. There is thus no combination of Γ, θ which enhances the electron lifetime to the order of magnitude $(10^5\,\mathrm{y})$ required for illumination of the entire jet in 3C 273 by UV-radiating electrons. Even if a significant fraction of the UV emission is not due to the synchrotron process but due to some different emission mechanism, the conclusions are not changed if the analysis is repeated for emission observed at 620 nm, for which the lifetime is only a factor of ≈ 1.4 longer.

Thus, the invocation of mild or even drastic beaming and/or sub-equipartition fields cannot resolve the discrepancy between the synchrotron loss scale and the extent of the optical jet of 3C 273, as has been possible for the jet in M87 [35]. As another alternative to invoking quasi-permanent re-acceleration, the existence of a "loss-free channel" in which electrons can travel down a jet without synchrotron cooling has been proposed by Owen et al. [77] for M87. As an extreme version of this case, we assume that the electron travels along the jet in zero magnetic field and is observed as soon as it enters a filament with magnetic field B_{fil}. The energy loss between acceleration and synchrotron emission is then only due to inverse-Compton scattering. The lifetime in 3C 273 is then $130\,000\,\mathrm{y}\ \times\Gamma^{-1}\sqrt{B_{-9,\mathrm{fil}}\mathcal{D}/\nu_{15}}$. Again, if the jet flow in 3C 273 is highly relativistic, the electrons suffer heavy inverse-Compton losses and the lifetime mismatch persists. In any case, it remains to be shown that the "loss-free channel" is a physically feasible configuration of a jet.

5.2   Is there an IR excess or a UV excess?

The observed spectra show a positive curvature in regions A, B1 and B2 (at the inner end of the optical jet) and in C2 and D (approaching the radio hot spot). This is a feature not expected for any synchrotron spectrum, in which losses deplete the high-frequency emitting population first. It also shows up as a discrepancy between the optical spectral index α_{BRI} determined at $1\rlap{.}''3$ resolution with data at 860 nm, 560 nm and 450 nm by Röser & Meisenheimer [87] and α_{opt} determined here using the new HST observations at 620 nm and 300 nm (Fig. 3.13). We now turn to consider the discrepancy between these spectral indices. We will also discuss what light the radio-to-UV spectra may shed on the mechanism for X-ray emission observed from the jet (Sect. 5.3).

Considering only the discrepancy between the ground-based optical spectral index α_{BRI} and α_{opt} one might ascribe it to flux variability of the jet. However, no variability has been detected so far[1]. Since the UV and optical observations presented here are contemporaneous, while the infrared observations were performed two years later, one would have to suppose a rather contrived variability to achieve an infrared-optical spectral index which is constant over a longer timescale than a variable optical spectral index.

Contamination by an error in the subtraction of the star which has been detected in this part of the jet might also contribute; however, this contamination has had little impact on the radio-infrared spectral index and is therefore expected to be negligible.

5.2.1   Possibility of contamination by a "backflow"

The indication that the radio "backflow" [85] may also be detectable at $2.1\,\mu$m [73] suggests that the flux from the jet at $1.6\,\mu$m may be contaminated by emission from the backflow as well. Therefore, we consider the possibility of contamination of the jet signal at this wavelength by the backflow.

The radio backflow arises south of the jet near C2 and D, those regions in the outer half of the jet showing a discrepancy between α_{opt} and α_{BRI}. It was detected because of its strong polarisation (up to $\approx 30\%$ compared to 5–15% for the jet), which leads to a discrepancy between radio (at $\lambda 73\,\mathrm{cm}, 18\,\mathrm{cm}$ and $6\,\mathrm{cm}$) and optical polarisation (at 560 nm) at radii $17''$–$19''$ from the quasar [85]. A radio spectral index map shows that the backflow has a radio spectral index of ≈ -1.15 between $73\,\mathrm{cm}$ and $18\,\mathrm{cm}$, significantly steeper than value of -0.8 determined for the jet. At these long wavelengths, the backflow contributes a small fraction of the total radio emission and continues smoothly into the lobe detected at the lowest frequencies (*cf.* Fig. 1.2). Similarly to the radio spectral index, the radio-infrared spectral index determined by Neumann et al. [73] steepens away from the jet, indicating that the backflow emits dominantly at radio wavelengths.

Figure 5.2 shows the infrared residuals of Fit A, in which the optical spectral index was assumed to represent the true synchrotron spectrum. The outward increase of the residuals qualitatively agrees with the steepening of the backflow spectrum towards the core observed in the radio spectrum [85]. However, the residuals from this fit are accounting for up to one-third of the observed infrared flux, which is a much larger fraction than expected for a backflow spectral energy distribution peaking in the radio. Moreover, a comparison of a map of the

---

[1]There is an effort by John Biretta (STScI) to repeat the presented HST observations at 620 nm to detect variability, with a baseline of 5 years.

Radio
polarisation
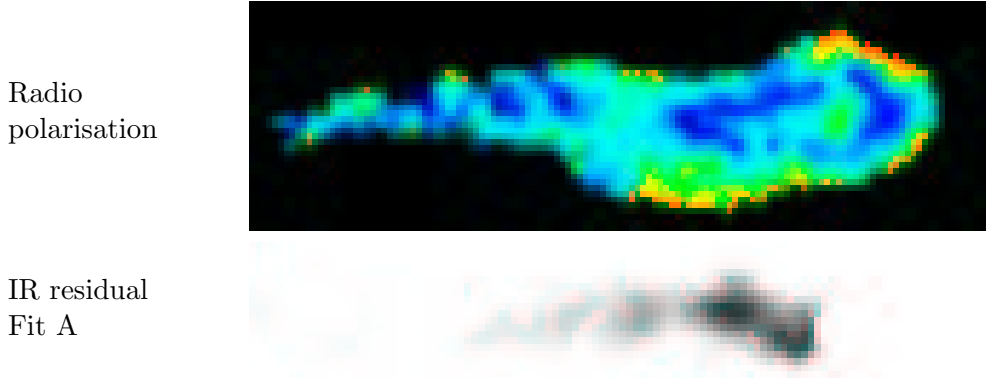
IR residual
Fit A

Figure 5.2: Comparison of radio fractional polarisation (top) and near-infrared fit residuals

residuals with a new high-resolution radio polarisation map (Fig. 5.2) shows that the largest residuals are not cospatial with those regions showing the highest fractional polarisation.

If such a large fraction of the infrared flux is due to the backflow, the agreement of the infrared-optical spectral indices α_{KO} and α_{IR}^{opt} (Sect. 3.5.2) implies that also a large fraction of the radio flux is due to the backflow, contradicting the conclusion drawn from the radio morphology in Sect. 3.3.3. This part of the backflow's emission would need to be less strongly polarised than that part which led to its detection.

Taken together, the implied large contribution to the total observed flux at all wavelengths longer than the infrared and its inferred low polarisation make this purported backflow very different from that originally observed, which has a steep spectrum (hence low flux at the short wavelengths considered here) and high polarisation. A backflow interpretation provides no explanation for the upward curvature of the observed spectra at the inner end of the optical jet, in regions A–B. A backflow observed at $\lambda\lambda 73\,\mathrm{cm}$–$6\,\mathrm{cm}$ [85] and $2.1\,\mu\mathrm{m}$ [73] cannot therefore account for the observed curvature of the infrared-optical-ultraviolet spectra and the discrepancy between the optical spectral index α_{BRI} [87] and the optical-ultraviolet spectral index α_{opt} presented here. We therefore consider the alternative explanation, additional emission in the ultraviolet.

5.2.2   Contamination in the ultraviolet

Polarimetry of the jet at 560 nm [87] had shown that this emission, like the radio emission, is due to the synchrotron process. Meisenheimer et al. [64] have presented fits of synchrotron continua to the observed spectral energy distributions of knots A, B, C, D and the hot spot H at $1\rlap{.}''3$ which have also been used by Röser et al. [89]. These observations were interpreted as emission from a single electron population. Figure 5.3 compares the run of the optical-ultraviolet spectral index α_{RU} determined at $0\rlap{.}''2$ resolution (figure taken from [42]) with the optical spectral index determined from the fits at $1\rlap{.}''3$ resolution [64]. It is noted that the optical spectral index predicted from these spectra (α_{fit}) in Fig. 5.3 is always steeper than the observed spectral index, that is below *both* α_{BRI} and α_{RU}. The reason is that already the infrared-optical spectral index α_{BRI} is slightly flatter than α_{KO}. At even shorter wavelengths, the spectrum flattens further, as evidenced by the flatter values of α_{RU}. Thus, all observational data indicate the important result of a gradual flattening of the observed spectrum towards higher frequencies.
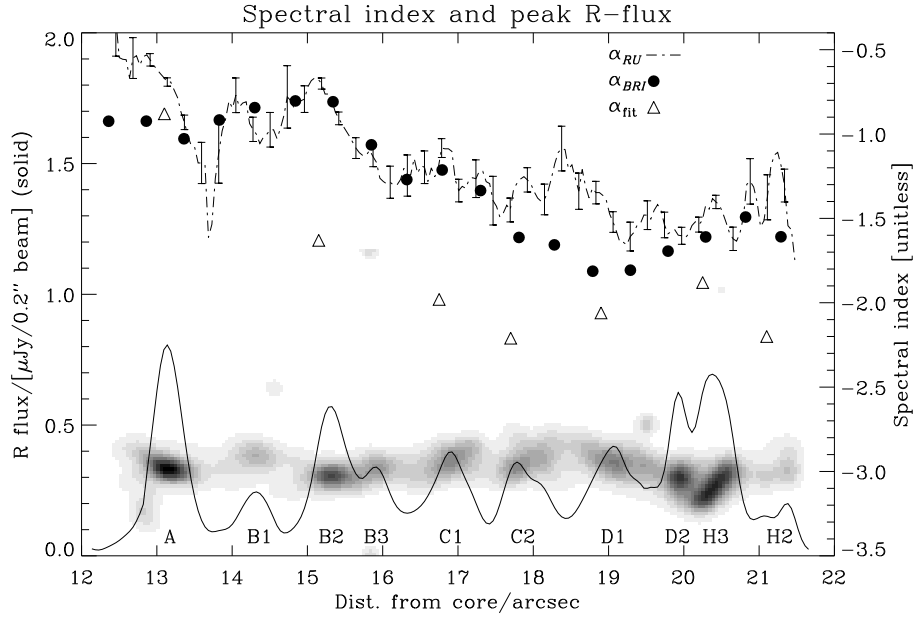
Figure 5.3: Run of the red-band brightness and optical-ultraviolet spectral index along the outer half of the jet in 3C 273, for a 0″2 beam. α_{RU} was determined from Figures 1 and 2, while α_{BRI} for a 1″3 beam is taken from Röser & Meisenheimer [87]. For comparison, we show α_{fit}, the corresponding spectral index obtained from synchrotron spectra fitted by Meisenheimer et al. [64] and Röser et al. [89]. While the observed spectral index agrees with older data, it is now clear that the fit is inadequate for the optical part of the spectrum. The flatter spectral index α_{RU} compared to both α_{BRI} and the fit may either be due to contamination of the infrared flux by a "backflow" component around the jet, or because of the presence of additional high-frequency flux, from a second high-energy electron population in the jet, *e. g.*, which is not included in the fit.

This flattening is most easily explained if the high-frequency emission from the jet is not entirely due to a single electron population, with contribution from an additional flux component. At 300 nm, already a significant fraction of the observed flux may be contributed by this additional component. Figure 5.4 shows the ultraviolet residuals of Fit B, in which the infrared-optical spectrum is assumed to be entirely due to synchrotron emission from a single electron population. An ultraviolet excess above the fitted infrared-optical cutoff is present in all parts of the jet. A smaller fraction of the optical emission will also be contributed by the same component. Unlike the backflow, which would be expected to contribute to the emission only in a limited part of the jet, the presumed additional UV component can be present in the entire jet and can account for the discrepancy between infrared-optical and optical-ultraviolet spectrum everywhere where it is observed.

As for the emission mechanism, the additional emission might be synchrotron emission as well, arising from a second, higher-energy population of electrons. Such a population was invoked by Röser et al. [89] to explain the X-ray emission observed from the jet. Alternatively, it might be contributed by an entirely different emission mechanism, by inverse-Compton scattering of cosmic microwave background photons, *e. g.*, another proposed explanation for the X-ray emission [58, 93]. Since both explanations link the observed ultraviolet emission to the X-ray emission, we consider the relation between both wavelength regions.

Figure 5.4: Ultraviolet residuals of Fit B, in which the infrared-optical spectrum determines the fitted cutoff. Significant ultraviolet excess over this fit is present in the entire jet. This excess might be due to the same emission mechanism responsible for the X-rays.

Figure 5.5: Adaptively smoothed X-ray image of the jet, obtained with the *Chandra* X-ray Observatory (adapted from Fig. 2 of Perlman et al. [84] with kind permission of the authors).

5.3 The X-ray emission from the jet

Figure 5.5 shows an X-ray image of the jet The strongest X-ray emission arises in regions A, B1 and B2 ($12'' \lesssim r \lesssim 15''.5$), where the UV excess above the radio-infrared spectra is largest. The emission mechanism responsible for the X-ray emission from the jet remains debated [58, 89, 93]. These authors agree that both synchrotron Self-Compton emission and thermal Bremsstrahlung appear unlikely processes as they require extreme source parameters, in particular unrealistically high photon or gas densities. The best currently available X-ray data have been presented by Marshall et al. [58], and all quoted X-ray flux measurements and spectral indices presented here are taken from that publication.

To explore the link between the X-ray emission and the ultraviolet emission and assess quantitatively whether both could stem from the same electron population, either from inverse-Compton or synchrotron emission, we consider whether the emission in both wavelength regions can be explained by a single simple model.

Now, we have to determine the spectra of two different components, the true synchrotron spectrum and the contamination spectrum. Although the synchrotron spectrum alone could have been characterised accurately based on the current data, the three high-frequency data points presented here are not sufficient to disentangle these two spectra.

From the optical jet polarimetry [87], it can be concluded that the synchrotron emission dominates in the infrared and the second component contributes significantly only in the ultraviolet. This implies that the contamination spectrum is flatter than the high-frequency spectrum of the jet.

The flattest plausible spectral index for a non-thermal contribution is $\alpha = -0.5$, the canonical value from shock acceleration theory [6, 7]. To assess whether the UV-contaminating component could reproduce the order of magnitude of the observed X-ray flux, the ultraviolet residuals from Fit B are extrapolated to the X-ray observing band with this spectral index. A similar comparison can also be made by extrapolating the observed X-ray flux from Marshall et al. [58] into the UV region with the X-ray spectral index also observed.
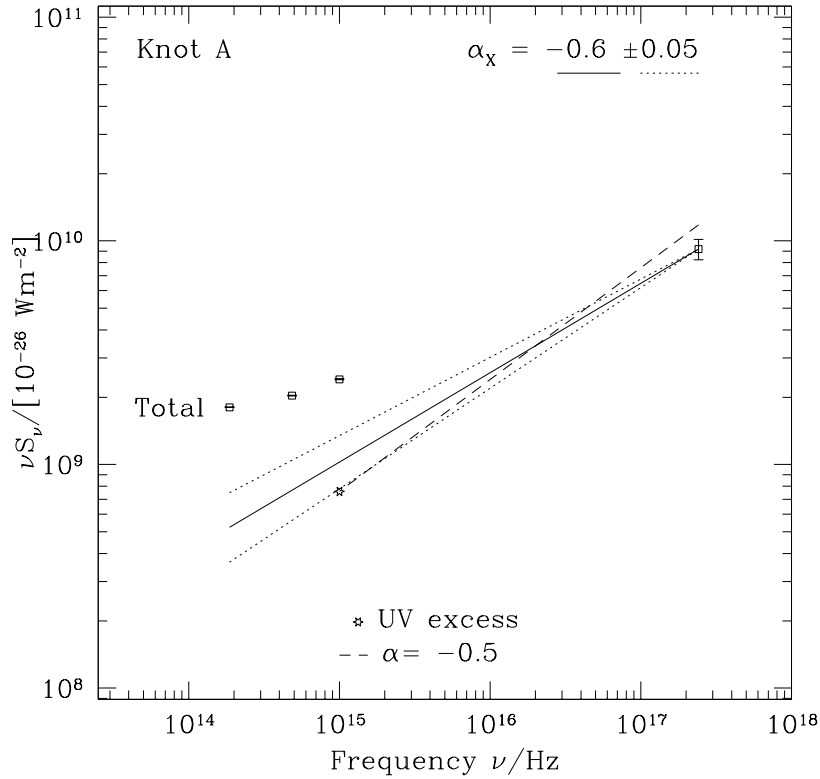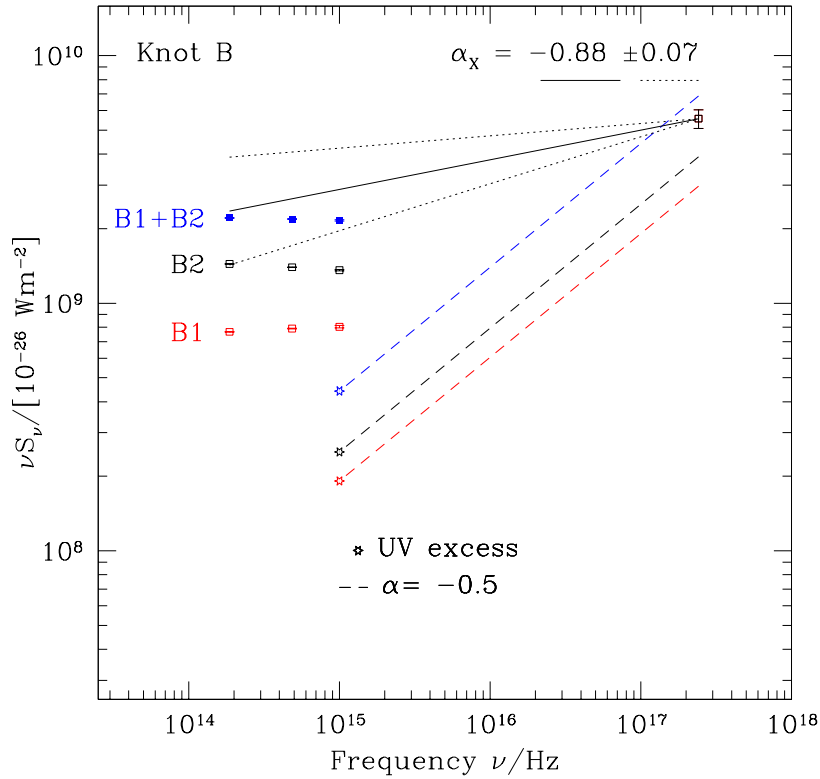
Figure 5.6: High-frequency flux for region A. The X-ray flux was measured in an aperture of radius $1''$ centered at $r = 13''$ [58], the surface brightness from the HST observations was summed over a corresponding area.

The X-ray images have a spatial resolution of $0''\!.7$, and no morphological details of the jet can be discerned at radii $r \gtrsim 15''\!.5$. Therefore, both extrapolations are carried out for the total emission in the three sections for which X-ray flux values have been determined [58]: regions A (Fig. 5.6), B (Fig. 5.7), and for the remainder of the jet (Fig. 5.8). We comment on each region of the jet in turn.

- Region A: The observed flux values as well as the ultraviolet residuals from Fit B were summed in a window corresponding to the X-ray aperture. X-ray emission and ultraviolet residuals are compatible with a single power law. The UV residuals account for 50% of the total flux, so a straight power law from infrared to X-rays cannot be ruled out definitely. However, a single radio-optical-X-ray power law appears unlikely, based on the steeper radio-infrared spectral index (Fig. 4.3).

- Region B: Contrary to the spectra presented by Marshall et al. [58], the HST data are only summed over regions B1 and B2 ($r \leq 15''\!.5$), that is, not including B3. The X-ray count rate decreases drastically at the transition between B2 and B3, and simultaneously the spectral energy density distribution turns over from ultraviolet-dominated to infrared-dominated. Unless this is entirely coincidental, the flatter spectrum of B1 and B2 is more likely to be related to the X-ray emission than the steep high-frequency spectrum of B3, whose X-ray emission is negligible compared to B1 and B2. The X-ray

Figure 5.7: High-frequency flux for region B. While the X-ray flux was measured in an aperture of radius 1″ centered at $r = 15″$ [58] which includes B3, only contributions from B1 and B2 are included in the infrared-ultraviolet flux points (see text). B1 and B2 are shown separately together with their sum.

and ultraviolet emission from these two regions may again be due to a single component.

- Remainder of jet: The inferred spectral index from UV excess to X-rays is ≈ -1, steeper than for either region A or B.

- Radio hot spot H2: No X-ray emission is observed here. This is compatible with the UV excess observed in this region (Fig. 4.5) only if the contamination spectrum steepens significantly below -1 at the transition from H3 to H2.

Based on these considerations, there is no strong argument against a common origin of the ultraviolet excess and the X-ray emission from the jet if he spectrum of the contaminating component steepens outwards along the jet.

In the inverse-Compton scenario, the spectral index of the scattered radiation is identical to the spectral index of synchrotron radiation from the scattering electrons, *i. e.,* the radio spectral index observed here. To reproduce the observed X-ray flux densities by inverse-Compton scattering of the microwave background radiation, it is necessary to assume that the jet flow is highly relativistic, so that the energy density of the background radiation field appears boosted up in the jet's rest frame (see Sect. 5.1). Model spectral energy distributions presented by Sambruna et al. [93] seem to indicate that the observed X-ray emission for the entire jet *can* be explained in this way; however, these authors were not aware of the flattening
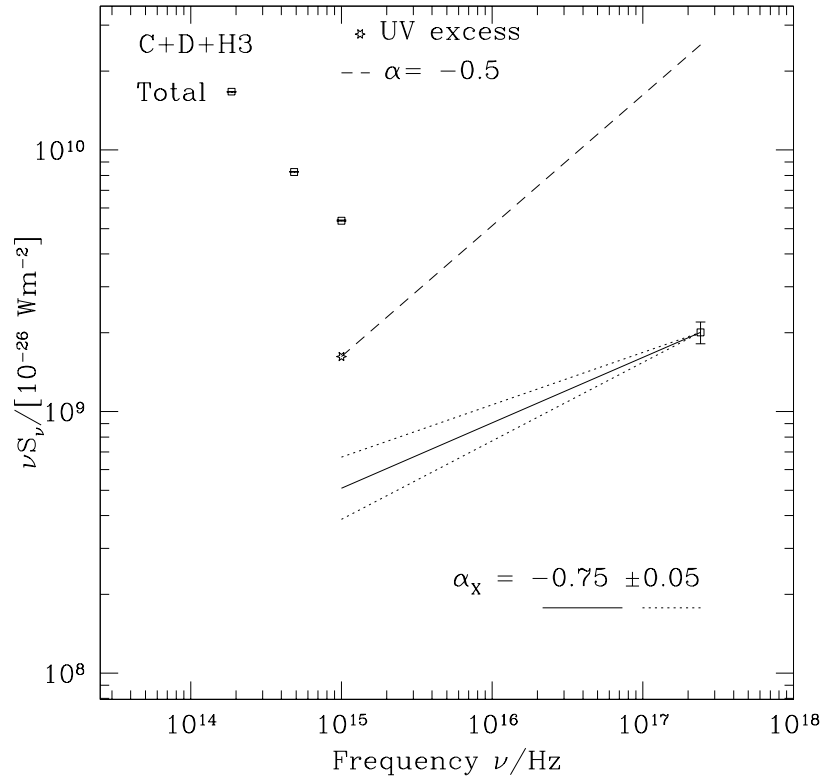
Figure 5.8: High-frequency flux for $r > 15{.}''5$.

of the infrared-optical-ultraviolet spectrum. An assessment of whether a steepening of the spectral index of the Compton-scattered radiation can be accounted for in the contamination scenario presented here is beyond the scope of the present work, as it requires a detailed study of the influence of changes in Doppler factor and magnetic field on the observed synchrotron and inverse-Compton scattered flux densities.

The outward flattening poses no problem if the contaminating component is synchrotron emission from a second, higher-energy electron population. The interpretation of the X-ray emission from the jet as synchrotron emission due to a second electron population [89] can thus explain both the flattening of the observed high-frequency spectra as well as the X-ray emission from the entire jet.

Chapter 6

Summary and outlook

One of the main unsolved questions in the study of extragalactic jets is posed by the observations of optical synchrotron radiation over scales much larger than typical synchrotron loss scales and far from the "working surfaces" of the standard model [9]. Another puzzle has been posed by the recent *Chandra* observations of X-rays from extragalactic jets, whose emission mechanism remains debated and which form a further sink of energy that has to be filled within the jet. The study of the physical conditions giving rise to the observed emission by keeping up the particle energy against strong synchrotron losses at optical wavelength and supplying the energy observed in X-rays can only be performed via a study of the synchrotron continuum, at the highest possible spatial resolution and covering the largest possible wavelength range.

The present work supplies a unique data set, in terms of spatial resolution and wavelength coverage, for the jet in 3C 273. A similar study has only been performed for the jet in M87 [82]. The observations serve as constraints for theories of particle acceleration within jets, far from obvious acceleration sites.

6.1   Observations

After a careful data reduction, new HST and VLA images at $\lambda\lambda$ 3.6 cm, 2.0 cm, 1.3 cm, 1.6 μm, 620 nm and 300 nm have been matched to $0.''3$ spatial resolution (Fig. 3.2). This enables a derivation of multi-wavelength synchrotron spectra of the jet in 3C 273 (Fig. 4.3). These data enable for the first time a direct determination of the spectra of individual jet knots and the regions of lower surface brightness connecting them.

Contrary to expectations for synchrotron emission, these spectra show a significant flattening in the infrared-ultraviolet wavelength range. The flattening is most pronounced near the onset of the optical jet, at radii $12''$–$15.''5$, where already the optical flux lies above an extrapolation of the radio-infrared power law and no cutoff is observed. This region is also the origin of the strongest X-ray emission. At radii $\approx 15.''5$, the spectral energy distribution of the brightness maxima turns over and peaks below optical frequencies (Figs. 3.3 and 4.3). In the outer part of the jet at $r \gtrsim 15.''5$, the flattening is only present in the optical-ultraviolet spectral index.

This negative curvature of the observed spectra means that the emission cannot be modelled as synchrotron emission due to a single electron population, as has been assumed in previous studies of this jet's emission [64]. Instead, additional emission must be present which

can lead to the observed flattening. The most likely explanation is an additional flat-spectrum component in the ultraviolet.

The optical-ultraviolet spectral index map generated at $0\rlap{.}''3$ resolution (Fig. 3.12) shows no strong correlation between spectral index and surface brightness. Such a correlation is present, however, on the radio-infrared and infrared-optical spectral index maps. Even these correlations are less pronounced than those found in the jet of M87 [82, 84], which are taken as evidence of localised particle acceleration in the knots of this jet. No evidence for similarly localised acceleration or loss regions is thus found in the jet of 3C 273.

The presented observations enabled a further check of the hot spot model by Meisenheimer & Heavens [63] and Heavens & Meisenheimer [34]. The hot spot is distinct from the remainder of the jet, both by its spectrum and morphology. It is well described by the hot spot model proposed by Heavens & Meisenheimer [34], in which particle acceleration takes place at a strong shock and losses occur downstream of the shock. In particular, the predicted offset of $0\rlap{.}''2$ between the radio and optical hot spot positions has been confirmed using the new HST and VLA data at their maximum resolution of $0\rlap{.}''1$ or better. It thus constitutes the only localised particle acceleration region which could be identified in this jet. The energy gains are observable as increase of the radio luminosity rather than an increase of the maximum particle energy. This adds further to the confidence that the assumed hot spot model is correct [68].

Both the observed flattening of the high-frequency spectrum as well as the X-ray emission from the jet are explained by a model in which two distinct electron populations emit synchrotron radiation: all the radio luminosity and a significant of the high-frequency (optical) luminosity are radiated by one population, while the ultraviolet excess causing the observed flattening as well as the X-ray emission are radiated by a second, higher-energy population.

6.2   Synchrotron spectral fits

The observed spectra are fitted with model synchrotron spectra to extract physical information, in particular the maximum particle energy and its variation along the jet.

Model synchrotron spectra according to Heavens & Meisenheimer [34] have been used to describe the observed spectra in terms of a low-frequency spectral index α_{low}, a break frequency at which the spectral index steepens to $\alpha_{\mathrm{low}} + 0.5$, and a cutoff frequency ν_{c} beyond which there is no significant synchrotron emission. The cutoff is mainly determined by the infrared-optical spectral index. The cutoff in the synchrotron spectrum corresponds to a high-energy cutoff in the electron energy distribution.

The maximum particle energy, inferred from the fitted cutoff frequency by assuming an equipartition magnetic field, decreases from above 5×10^6 in the bright region A at $r \approx 13''$ and settles to a plateau of order 10^6 at $r \gtrsim 18''$. A further decrease to $\gamma_{\mathrm{max}} \approx 3 \times 10^5$ is only observed at the position of the radio hot spot, beyond which there is no optical, but still infrared emission.

This global run corresponds to the findings of previous studies at $1\rlap{.}''3$ resolution. The new observations show that also small-scale variations of γ_{max} are much less extreme than the variations in surface brightness, there are in particular no significant differences between the spectra observed at and between the local brightness peaks. The physical conditions within the jet must correspondingly vary smoothly.

An examination of the effect of relativistic beaming on the inferred synchrotron lifetime

show that the overall slow decrease of γ_{max} cannot be explained by an enhancement of the true synchrotron lifetime by beaming effects. We conclude that particle acceleration must take place distributed across the entire jet. Although the investigation of possible acceleration mechanisms is beyond the scope of this work, it is clear that particle acceleration is not limited to the bright regions. This does not preclude the possibility that the enhanced-brightness regions are shocks – but even if they are, re-acceleration between them is necessary to explain the observed spectral features.

The observations presented here have the important conclusion that the physical conditions within the jet of $3C\,273$ vary very smoothly down to the spatial scales of $0.9 h_{70}^{-1}$ kpc resolved here, so that the particle acceleration mechanism must act over the entire jet. They have furthermore established a flattening of the high-frequency spectrum. The jet's emission cannot therefore be assumed to arise from a single electron population and requires the presence of an additional emission component.

Future work to investigate the acceleration mechanism at work in this and other jets must be based on these results.

6.3 Future work

The motivation for the collection of the data set presented here was an accurate determination of the maximum particle energy and its variations across the jet (Sect. 1.2). Because of the presence of a second emission component, the three high-frequency flux points presented here are not sufficient to disentangle the contributions from the synchrotron spectrum and this second component. A significant improvement could already be achieved by raising the signal-to-noise obtained at 300 nm to the level reached by the longer-wavelength data.

To test all the predictions made by assuming two electron populations, observations at further wavelengths and making use of the full resolution afforded by the HST are needed with similar signal-to-noise ratio as the presented optical data. The lower-energy synchrotron spectrum needs to be constrained by observations at optical and longer wavelengths, while the higher-energy population's spectrum can be determined at shorter wavelengths. Particular insight can be expected by an accurate determination at higher resolution of the high-frequency spectra in those regions of the jet showing a difference between radio, optical and X-ray morphology (*cf.* Sections 3.4 and 5.3).

Such observations can be obtained with present (the WFPC2 employed here) or planned (the new, more sensitive *Advanced Camera for Surveys*) UV-sensitive cameras on board the HST, the latter being sensitive at wavelengths from 900 nm down to 200 nm. Further exposures near 450 nm and 800 nm, of comparable duration to the present optical data, would provide two additional flux points with high signal-to-noise.

To achieve a similar resolution at longer infrared wavelengths, it is will be necessary to employ a 10m-class ground-based telescope with adaptive optics (*e. g.*, in the 1–$5\,\mu$m range with the CONICA-NAOS system on the ESO's VLT). Such observations are needed to make a further step and determine the spectrum of the excess emission. Smaller space-based infrared telescopes, such as SIRTF (85 cm mirror, launch planned for 2002) or Herschel (3.5 m mirror, launch planned for 2007) with imaging capabilities at even longer wavelengths can be employed to test the predictions of more detailed models, even if they will not reach a similar resolution.

Even further in the future, the unprecedented sensitivity and resolution of the planned Atacama Large Millimeter Array (ALMA) will provide an important data point in the mil-

limeter wavelength region, which is the largest gap in the currently achievable wavelength coverage.

The inverse-Compton scenario for the X-ray emission makes strong predictions for emission at the highest energies: the order of magnitude of the expected ultra-high-energy cutoff frequency is given by $\gamma_{\max}{}^2\nu_{\mathrm{CMB}} \approx 10^{23}\,\mathrm{Hz}$, corresponding to very hard γ-rays. Because of the required sensitivity in an accessible energy range, such predictions can only be tested with the next generation of γ-ray telescopes (*e. g.*, GLAST).

Further planned work concerned with the jet in 3C 273 includes high-frequency polarimetry as well as an investigation of the differences between the optically bright outer part of the jet discussed here, and the optically much fainter inner $10''$ of the radio jet.

HST polarimetry of the jet in M87 has indicated that the electrons responsible for optical and radio synchrotron emission are not entirely colocated there [83]. To do a similar high-resolution optical polarimetry study of 3C 273 with the HST, we have embarked on a collaboration with Eric Perlman (University of Maryland, Baltimore County). We have already secured guaranteed observing time to perform adaptive-optics assisted near-infrared polarimetry of the jet at the ESO's VLT, which will constitute the first high-resolution near-infrared polarimetry of any extragalactic jet. The polarimetry is also expected to provide new insight in the emission mechanism of the jet's radio-quiet extensions, which will also be observed spectroscopically.

To investigate the differences between the inner and outer jet (Sect. 2.2.3), deep optical VLT imaging has been obtained. Further deep near-infrared observations will enable a comparison of the overall run of the synchrotron spectrum in the inner and outer jet halves and be the first step towards an explanation of the sudden onset of bright optical synchrotron emission half-way between core and hot spot.

Exciting new facts about the jet in 3C 273 have been revealed by the advances in optical and radio observing technology. From the surprising observational facts encountered here, more surprises are yet to be expected *en route* to an understanding of the physics of this object. Current advances in numerical work are also expected to shed further light on the structure and dynamics of jets. As an extreme object of its kind, 3C 273 will thus remain high on the agenda of astrophysicists studying jet phenomena.

Appendix A

WFPC2 calibration issues

A.1   CCD calibration

The images worked with are in units of counts or Data Numbers (DN). At the chosen gain, 7.12 ± 0.30 detected photo-electrons (and hence photons) correspond to one DN ([8], 4.12).

All data reduction steps are described in Section 26.2.2 of [105].

A.1.1   A/D correction

Photons incident on the CCD detector liberate electrons from the semiconductor material by excitation across the band gap from the valence to the conduction band. The charge accumulated in each pixel of the detector is measured and converted to a digital signal. As the analogue-to-digital (A/D) converter which converts the observed charge in each pixel to a DN is imperfect, some DN values are reported with a wrong frequency. This systematic error has been well characterised and is only a small effect. The A/D correction applies a statistical correction based on an empirical calibration.

A.1.2   Bias level removal

The charge accumulated in each pixel is measured by electronic amplification and subsequent conversion to a digital signal. In order to avoid the difficulties of measuring a charge near zero, each pixel has an electronic pedestal so that A/D levels are well above zero. The mean level of the bias has to be removed from the data frame to count only the "science charge". It is determined from the overscan region of the CCD. Overscans are additional readouts beyond the last pixel, *i. e.,* operation of the readout amplifier after the charge from all pixels in one row has been shifted through the read-out register. The bias level is systematically different for even and odd columns of the chip and hence determined and subtracted separately for both.

A.1.3   Bias pattern subtraction

In addition to the global bias level, there are pixel-to-pixel variations of the bias. These variations are measured by reading out the CCD array twice in rapid succession. The first read-out resets all pixels to the bias value. The data frame obtained in the second read-out (the so-called "zero-length exposure") is just the bias value. The zero-length exposure is

corrected for the bias level as described in Sect. A.1.2 and then contains only the bias pattern. In practice, the bias pattern is obtained by averaging a large number of bias level-corrected zero-length exposures, in order to avoid the introduction of extra noise. The correction frame is subtracted from the science exposure.

Bias subtraction is expected to be accurate to 0.1 DN RMS.

A.1.4   Dark image subtraction

Because of thermal excitations of electrons from the valence to the conduction band of the CCD semiconductor material, a small charge is detected every second even if there is no radiation falling on the detector. CCDs are cooled to keep the average electron energy down and thus reduce the dark current. A dark file is constructed from 10 long exposures with the shutter closed. It is scaled to the exposure time of the science frame and then subtracted from it.

Care has to be taken when the intrinsic noise of the science frame is very low. In this case, the shot noise from the scaled dark frame may dominate the total noise and thus introduce extra noise. The background of the U exposures is so low that this may be a problem. The noise from the dark image will therefore be included in the consideration of errors.

In addition, there will be some warm pixels with a slightly elevated dark current. These are a problem in principle, as they will positively bias the photometry of the jet. However, warm pixels are fairly easy to reject. They are the only image artifact limited to exactly one pixel; both point sources and cosmic ray events extend over more pixels. As we are combining four and fourteen offset images in the R and U band respectively, any warm pixels with large deviations will be removed by the cosmic ray rejection routine. Deviations of pixels with smaller deviations will be dominated by the noise.

"Cold pixels", which are pixels with a large negative reading, can be introduced by a hot pixel in the dark frame. When the dark frame is subtracted from the science frame, the hot pixel produces a negative spike. However, these are easily recognised and will pose no problems. The use of the best currently available dark files instead of those used in the initial pipeline has visibly reduced the number of warm pixels in the R-band exposure as well. The dark subtraction is expected to be accurate to 0.1 DN/hr (worse for warm pixels).

In addition to the electronic dark current, there is a so-called *dark glow* correlated with the cosmic ray rate during an exposure. It is believed to be due to luminescence induced by cosmic rays in the MgF_2 field flattener windows in front of the WFPC2 chips. In the Planetary Camera, it dominates the total dark current. Due to the geometry of the windows, the effect is stronger at the centre of the chips than at the edges. It is variable with cosmic ray flux and hence with time. Therefore, the calibration dark frames may have a different dark glow from the science frames and the standard dark subtraction may lead to a curvature in the background. This effect is small even from edge to centre of the chip ($2~e^-$/pixel) and only slowly-varying. Hence, it behaves essentially like an additional background component and will be corrected for by the background fitting (see Sect. A.3.4).

A.1.5   Flat-fielding

There are small-scale variations of the quantum efficiency of individual pixels, *i. e.,* some pixels record systematically more or less electrons for a fixed number of incident photons than others. This is a multiplicative effect (the true number of counts per pixel is related to

the detected number by a constant factor), in contrast to the previously considered ones which are all additive (the detected number of counts is offset from the true number). After their contributions to the read-out charge have been correctly subtracted, the remaining counts on the chip are only due to the sky and sources. The pixel sensitivity variations result in varying count levels in an image of a uniformly illuminated source. The image appears darker where pixels are less sensitive. Such an image is used as flat-field frame. For the HST, the flat-fields are combinations from pre-launch laboratory images and several on-orbit earth flats contemporaneous with the observations. The former ones establish the pixel-to-pixel sensitivity variations, while the latter ones provide information on large scales across the chip. The science frame is divided by a flat-field frame normalised to unity in order to scale the measured counts up or down to the true value.

The flat-fielding is accurate to better than 1% RMS on large scales, and to 0.3% on small scales.

A.2 Noise considerations

The noise in HST exposures has the following components:

1. Source photon Poisson noise

2. Sky photon Poisson noise

3. Dark current Poisson noise

4. Gaussian readout noise

5. Digitisation noise

The sky background level has to be estimated and subtracted from the science frame to obtain only the photon flux from the object. The signal to noise ratio (S/N) is then the ratio of object photon flux to the total noise contained in each aperture used to measure the flux. The background subtraction can only ever be correct to within the background noise.

Read noise and digitisation noise are constant for all exposures, regardless of the filter. The dark noise is also independent of the filter and only a function of exposure time. Its magnitude is about 0.005 e^-/s for the Planetary Camera. The digitisation noise results from the imperfect mapping of discrete count values to the continuous distribution of source and sky photon flux. Its magnitude *in electrons* is estimated to be $\frac{1}{\sqrt{12}}$gain, i.e., $\sqrt{4.1}e^-$ for the 7.24 e^-/DN ("gain 7") used in these exposures. For high background levels, this contribution is insignificant, for low levels, this noise component has been included in the specification of the read noise [8, Sect. 6.5.1]. It therefore need not be considered.

Readout noise is the noise in the signal chain of the CCD, such as the readout amplifier. For the Planetary Camera chip in gain 7 mode, its magnitude is $5.24 \pm 0.30 e^-$.

The total width of the background count distribution is the joint effect of all these noise sources. The level of the total background signal is the sum of the individual components, its width the width of the convolution of all contributing noise distributions. As a CCD is a counting device counting electrons, the error on each count reading is the Poisson noise on the total *electron* counts *including the dark current*, convolved with the Gaussian distribution of the read noise. Writing σ_{sky}^2 for the sky count variance, σ_{dark}^2 for dark e^- count noise and

σ^2_{RON} for the read noise, the total variance of the electron count distribution is given by the sum of individual variances,

$$\sigma^2_{e^-} = \sigma^2_{sky} + \sigma^2_{dark} + \sigma^2_{RON} + \sigma^2_{obj} + \sigma^2_{flat} + \sigma^2_{darkframe}, \tag{A.1}$$

with $\sigma_{sky} = \sqrt{N_{sky}}$, $\sigma_{dark} = \sqrt{N_{dark}} = \sqrt{I_{dark} \times t_{exp}}$ and $\sigma_{RON} = 5.2e^-$, writing N to designate electron counts. $\sigma^2_{darkframe}$ and σ^2_{flat} are the noise due to the dark and flat-field calibration files respectively. The noise *in DN* is then obtained by a simple scaling with the gain g:

$$\sigma_{DN} = \frac{\sigma_{e^-}}{g}. \tag{A.2}$$

Writing C for the sky count rate measured on the calibrated frames, which are in units of DN, equation A.2 and A.1 combine to

$$\sigma_{DN} = \frac{\sqrt{g(C+S) + \sigma^2_{RN} + I_{dark} \times t_{exp} + \sigma^2_{flat} + \sigma^2_{darkframe}}}{g} \tag{A.3}$$

where S is the number of source counts measured *above background*. The signal to noise ratio of a count measurement in any aperture is then

$$S/N = \frac{gS}{\sqrt{g(S+C) + \sigma^2_{RN} + \sigma^2_{dark} + \sigma^2_{flat} + \sigma^2_{darkframe}}} \tag{A.4}$$

where both source counts and variances are summed over all pixels in the aperture. It is evident from this formula that the S/N increases by adding up several frames, as the numerator grows linearly with observing time while the denominator consists of the square root of a sum of terms growing linearly with either observation time (dark current, sky and source signal) or with the number of exposures contributing (read noise, calibration frame variances). The ratio therefore grows approximately with the square root of the number of frames summed up if their observation time is similar.

The noise sources are of varying relative importance in the R and U filters, hence exposure sets in each filter will be considered separately. Both noise level and width of the noise distribution in individual frames and summed image will be considered and have to be satisfyingly accounted for.

A.2.1 Noise of R images

The global sky background level and noise were determined from the histogram of pixel values for each image. In the R filter, the expected background level for the ecliptic coordinates of 3C 273 is 0.02 e^-/s/pix, or about 50 e^-/pix for a 2500 s exposure (a typical value for this set of observations). The total dark charge accumulated in such an exposure will be 12.5 e^-/pix. The R exposures will therefore be sky background limited, but with still a sizeable contribution from read noise. Thus, after the calibration, which has removed the dark level, a background count level of $\approx \frac{50}{7.24} = 6.9$ DN is expected. The expected noise is $\frac{1}{7}\sqrt{50 + 12.5 + (5.24)^2} = 1.35$ DN. Table A.1 shows the background levels and noise of individual and summed exposures. The background level has been measured by locating the peak of the pixel value histogram. A typical histogram is shown in Fig. A.1. The noise has been obtained from statistics on an empty area of the frames.

Data	BG level		BG noise	
set	exp.	obs.	exp.	obs.
0101	6.57	6.75	1.32	1.59
0102	7.14	7.58	1.35	1.66
0104	7.43	7.93	1.37	1.65
0105	7.43	7.38	1.37	1.63
Sum	28.57	29.64	2.71	3.27

Table A.1: Expected and observed background counts and noise for R exposures
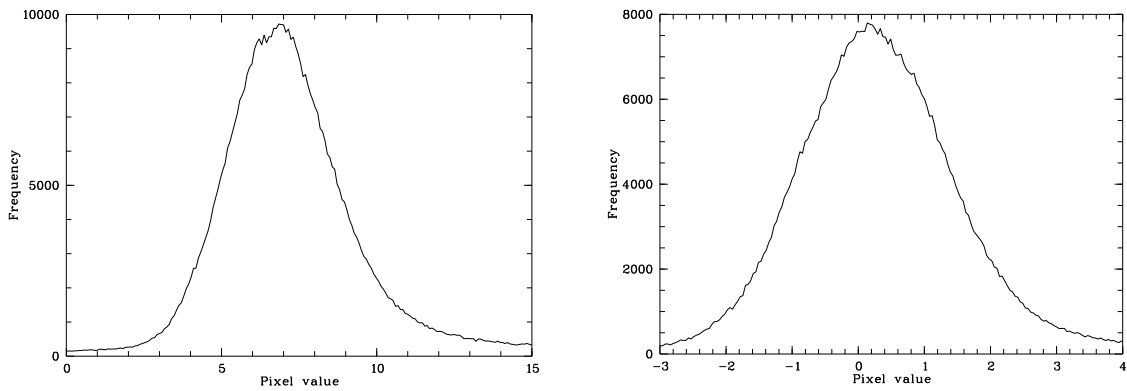
Figure A.1: Histogram of R-band (left) and U-band (right) background values

A.2.2 Noise of U images

In the U, the sky background level is much lower than in the R. In the Hubble Deep Field (HDF) exposures, the median background level for F300W was near 0.0005 e^-/s (see [33]), or 1.25 DN in 2500 s. The WFPC2 handbook [8] indicates negligible background levels in F300W. The observed background levels lie between near 0 and 1.1 (see Tab. A.2), i. e., they are within the expected range. The Poisson noise on sky counts is therefore negligible and these exposures will be limited primarily by the read noise. There is a contribution to the total noise from the Poisson noise on the dark current. The latter contribution is especially important when evaluating the S/N — it cannot be read off from the calibrated science frame (before background subtraction) any more, but the subtracted dark level must be taken into account as well.

Neglecting the background Poisson noise compared to the dark and read noise, the expected background noise in the U filter is $\frac{1}{7}\sqrt{7 \times 0.5 + 12.5 + (5.24)^2} = 0.9$ DN, nearly independent of the background. This agrees well with the observed values (Tab. A.2).

A.3 Image defects

A.3.1 The charge trap problem

There are a few pixels on each CCD acting as charge traps during detector readout. This means that not all the charge of pixels in higher rows is transferred down a column, but

Data set	BG level	BG noise
0106	0.22	1.01
0107	0.37	0.99
0108	0.53	0.96
010a	0.74	0.97
0201	0.33	0.98
0202	0.44	0.98
0203	0.74	0.99
0205	0.63	0.99
0206	1.12	1.01
0301	0.42	0.97
0302	0.38	0.99
0304	0.24	0.99
0305	0.17	1.00
0306	0.22	1.01
Expected sum	6.55	3.70
Sum image	5.73	4.32

Table A.2: Expected background counts and noise for U exposures. The last two rows show the level and noise expected for the sum frame and the actually measured value.
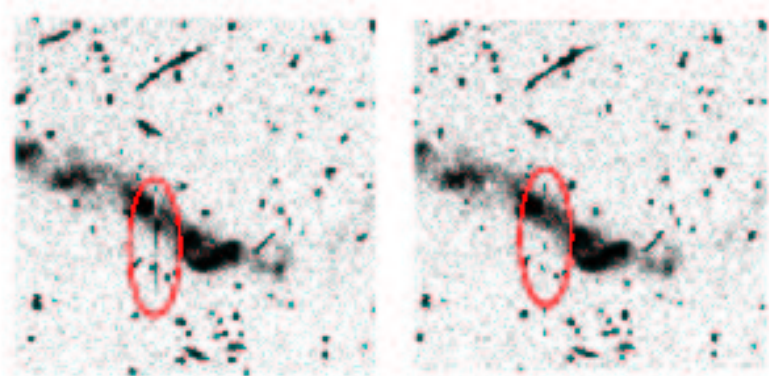
Figure A.2: Charge transfer trap on R-band image. *Left*, before correction, *right*, after correction.

a constant fraction remains in the pixel (see [105], 26.4.3). This leads to bright streaks appearing above objects or cosmic rays falling on a charge trap, and dark streaks above those on empty sky. For the signal levels of the U exposures, the traps pose no serious problems. They are clearly visible on the R exposures, though, as streaks (see Fig. A.2). These streaks are not completely removed by the cosmic ray rejection (this is described further down in Sect. A.3.3) because it is optimised to remove positive outliers, while part of the charge trap effect is producing a negative outlier (inside the jet in Fig. A.2).

Some of the traps, including the one located inside the jet, are well characterised and the fraction of trapped charge is known. In principle one can work out how much charge *should* have been registered from how much charge *was* registered. Such a correction was attempted

according to the prescription in [107]:

$$C_j^* = \frac{C_j - C_{j-1}(1 - T)}{T} - B \qquad (A.5)$$

C and C^* are the original and corrected pixel value *above the background level B*, the subscript is the pixel number (increasing in the opposite sense to the readout direction) and T is the transfer efficiency of the trap, which is 17% for this trap. Starting at the first pixel above the trap, the "leftover" charge from the previous pixel is subtracted and the remaining charge scaled up by the inverse of the transfer efficiency.
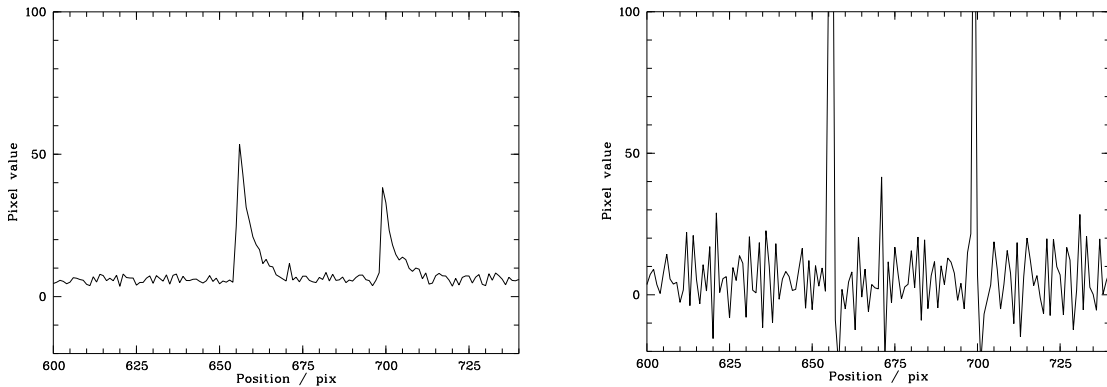
Figure A.3: Part of the column affected by the charge trap before (left) and after (right) the attempted correction. The two peaks are due to cosmic ray hits. Their exponential decay can be seen. They have been sharpened up by the charge transfer correction. However, noise peaks have also been sharpened up where they should not have been.

It can be seen in Fig. A.3 that the exponential tails created by cosmic rays clocked through the charge traps are sharpened up (near columns 650 and 700). However, this correction also leads to an large increase in the background noise in rows above the affected pixel without removing the charge trap effect completely. Hence, this way of correction was rejected as unsuitable.

Instead, the part of the affected column inside the jet was replaced by a suitably scaled copy from a different, offset exposure in which the same region of sky falls on a different column. This is possible because the exposures have integer offsets with respect to each other. Starting at the charge trap, 50 pixels along the column are replaced before CR rejection. The one drawback of this method is the lower S/N of the corrected pixels, as only three exposures contribute to the summed image at the affected sky coordinates instead of four. This disadvantage is judged to be tolerable compared to the error made by omitting the charge trap correction, which would have rendered parts of four columns of the jet image faulty and introduced errors in the photometry. Instead of giving one (arbitrary) image twice the weight in the summation in the affected regions, all three clean images could have been given equal weight by averaging over these.

A.3.2 Horizontal smear

All images are affected by a bright horizontal streak across the chip in rows with saturated pixels in the quasar core. It is only dimly visible in individual frames, but clearly visible in
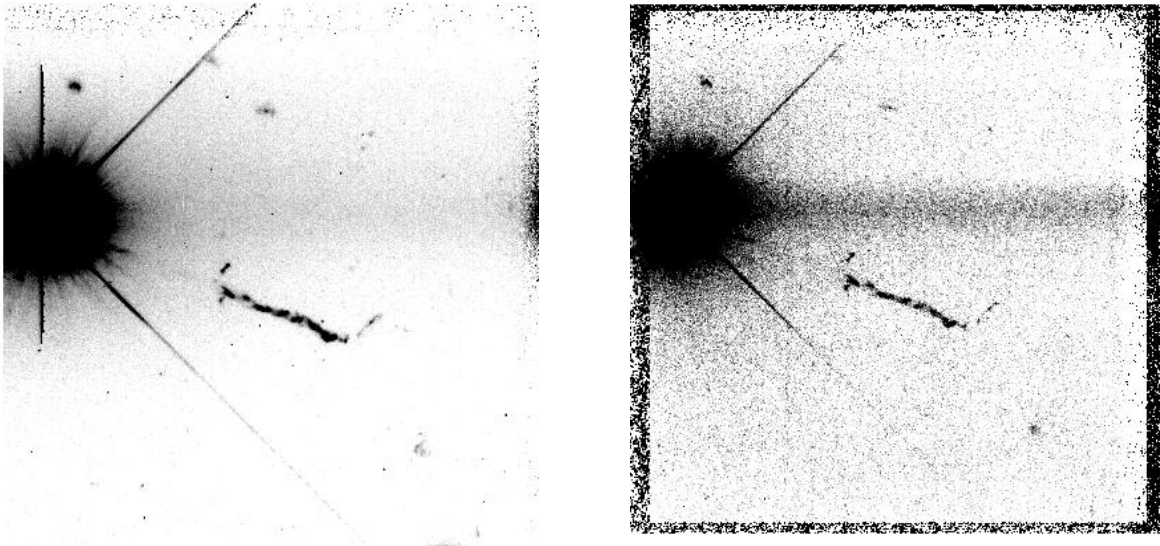
Figure A.4: R (left) and U (right) sum images. The greylevels are such that the expected background level is white and black corresponds to twice this level. The smear is visible as increased background level peaking at the columns containing the quasar core.

the summed frame (see Fig. A.4). This effect arises in the readout electronics, which continue to detect spurious charge after a saturated pixel has been read out. The amount of detected charge decays exponentially along a row and wraps around the edge of the chip to the next row. Each saturated pixel contributes a constant offset. In addition, the background above and below the streak is systematically different. Again, the effect can in principle be corrected [47], but the decay length and offset are not well-characterised and seem to vary with time. Hence, the correct parameters would have to be found by some sort of fitting. Instead of attempting such a fit, the streak will be removed by the background fitting, if the affected region of sky proves to be of interest at all.

A.3.3   Cosmic ray rejection

Hubble Space Telescope images are heavily affected by cosmic ray hits. The impact is more severe than for earth-bound observations, as the atmosphere provides some shielding. Fig. A.5 shows a typical image after all data reduction steps have been applied. It is clear that significantly more counts have been produced by cosmic ray hits than by source or sky photons. The jet as such can hardly be made out on the U-band image.

In order to remove the cosmic ray hits, a κ-σ-clipping algorithm was used. An implementation of this algorithm is provided locally as command COSMIC/MEDIAN [36], part of the photometry package MPIAPHOT running under the image analysis system MIDAS [24]. The images in each filter are registered so that the sky coordinate is known for each pixel of each individual exposure. In line with the discussion in Sect. B.2.2, these coordinates are relative to those of the first R-band image. Henceforth, "pixel" is used to refer to those pixels mapped onto identical positions on the sky, not identical physical detector pixels. For each pixel in each filter, the median value of all exposures is computed. This median is then adopted as the best estimate of the true pixel value. All pixels lying significantly ($\kappa \times \sigma$) above the median are considered to be cosmic rays.
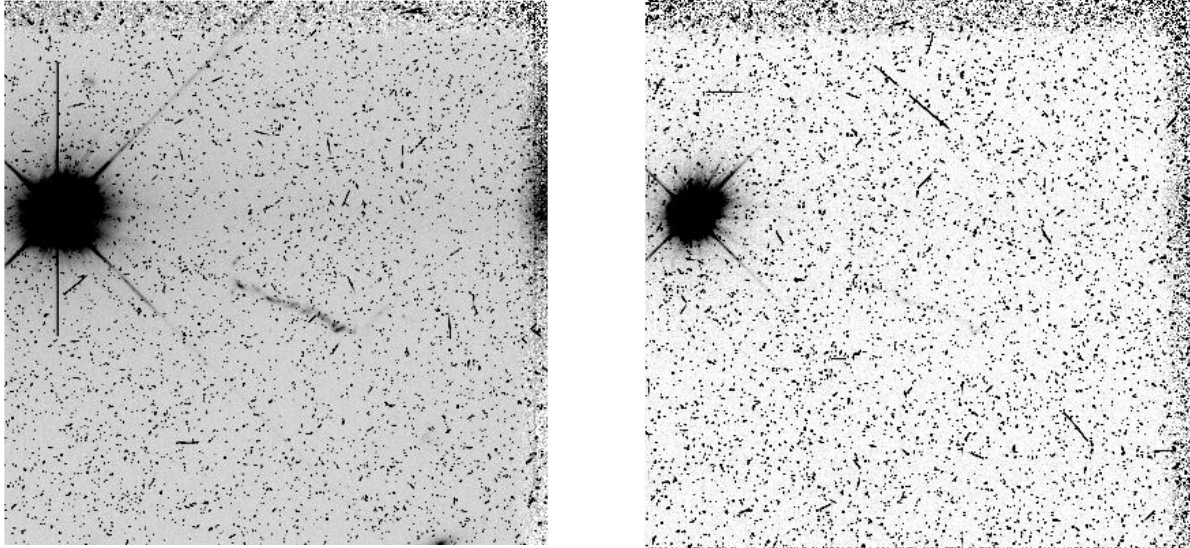
Figure A.5: Examples of individual data frames after reduction, but before cosmic ray rejection. Exposure time ≈ 2500 s. Left, R-band, right, U-band.

As cosmic rays always have a higher signal than unaffected pixels and never a lower one, a low-biased median of pixel values above background is generated. When an unbiased median is used, not all cosmic ray hits are correctly identified for sets comprising only a few exposures (like the present R-band images) because cosmic ray hits are so frequent that a pixel may be affected in more than one exposure. All pixel values deviating more than $\kappa \times \sigma$ from the median are replaced by the median, scaled according to the exposure time of each image. As cosmic rays affect more than one pixel, there is a second iteration. In this, pixels adjacent to those corrected in the first pass are considered again and corrected if they deviate by a smaller κ. The values for κ used in the first and second pass were 4 and 2.5, respectively. With this choice, about 35,000–50,000 pixels are corrected in each frame. This agrees well with the expected 30,000–50,000 affected pixels in a fiducial 2500 s exposure (different values according to [8] or the on-line Exposure Time Calculator [38]). Figure A.6 shows the same images as Fig. A.5 after the correction algorithm has been applied.

After the correction for cosmic ray hits, the rejection log was examined in order to verify that no pixels inside the jet image were affected by cosmic ray hits in all four R-band images.

A.3.4 Background fitting

Rather than fitting a globally constant background level, which would not reproduce the large-scale variations of sky and other background effects (see Sect. A.3.2), a more sophisticated fitting routine is employed allowing local variations. Polynomials are fitted to each row and the coefficients of these polynomials smoothed across a fixed number of rows. The degree of the polynomials is chosen as low as possible, usually second. Initially, only the region around the jet is considered as this is the prime interest of the investigation. Areas containing the jet pixels are masked out and disregarded by the fitting routine. The same sets of parameters (degree of polynomial, smoothing, mask, region) for the fit are used both for the R and U sum image.

To assess the quality of the fit, the following criteria are used:
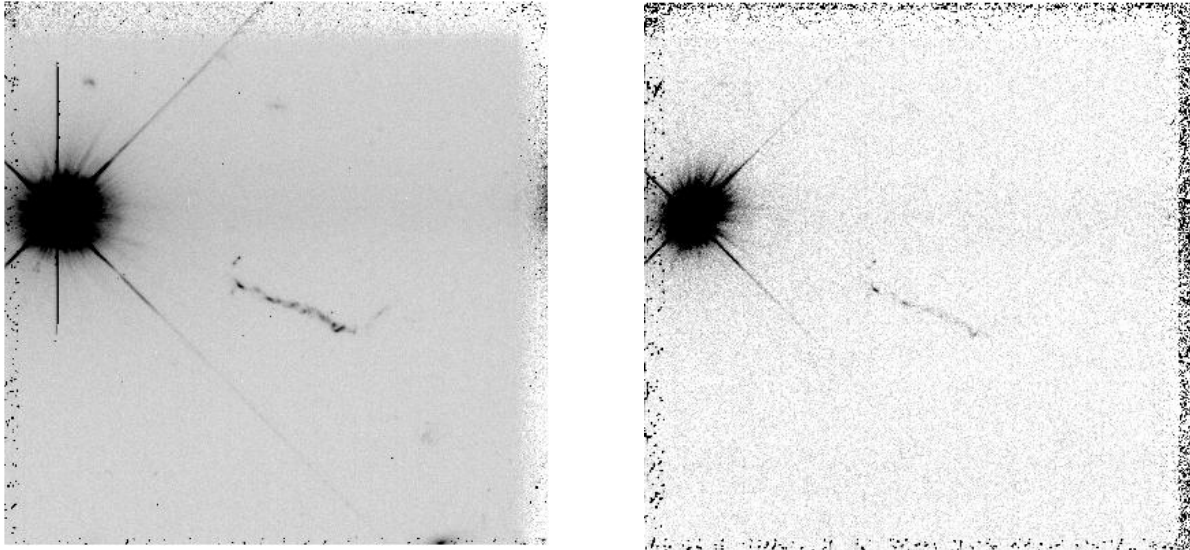
Figure A.6: The same images as in the previous figure, after application of the cosmic ray rejection routine. Left, R-band, right, U-band.

- Median of background subtracted area

- Flatness of the background

Naturally, the mean of the resulting background should be as close to zero as possible so that photometry of empty sky regions would lead to the expected zero flux. The flatness of the background is determined by adding sky counts in small regions of 3×3 pixels (corresponding to a 0.15" aperture) near the jet. The RMS scatter of the count levels in these small apertures limits the accuracy with which the jet flux in corresponding apertures can be measured, in addition to the intrinsic background noise.

In the R image, the background levels in these small apertures have a 1σ-scatter of 2.1 DN before and 0.43 DN after the fit. The overall background level, again determined by a Gaussian fit to the pixel value histogram as in Sect. A.2, is 0.032 DN per pixel; the median of background values in a region containing no source is 0.061 DN per pixel. For the U, the scatter values are 0.75 DN and 0.70 DN before and after fitting, respectively. The residual background level after the fit is 0.14 DN per pixel, the median in an empty region of sky is 0.044 DN per pixel. A comparison of these numbers for the two bands is only possible after the photometric calibration, as identical count rates correspond to different physical fluxes and the exposure times of both images are differing by a factor of three.

Appendix B

Alignment of images

B.1 Determination of the required accuracy

B.1.1 Misalignment of point sources

For an extended source, the greatest error in a flux measurement due to a misplaced weighting function is made in the steepest intensity gradients. The steepest possible gradient is that of the flanks of the PSF. The worst-case error for an extended source is then just the error made when measuring a point source flux with a misplaced weighting function. The result a flux measurement supposedly carried out at the peak of the point source image is the surface integral over the PSF multiplied by the weighting function. Both are assumed to be Gaussian bells. (This is appropriate for earth-bound seeing-convolved observations; the PSF of the optical and especially the near-infrared camera on board the Hubble Space Telescope has substantially more structure due to the primary mirror aberration and corrective optics and because it is diffraction limited.)

Suppose the flux of a point source whose Gaussian PSF has a width σ is measured with a weighting Gaussian of width σ_w. The resulting effective resolution is $\sigma_{eff} = \sqrt{\sigma_\omega^2 + \sigma^2}$. If the weighting function is offset a distance ϵ from the centre of the point source, the measured flux $f(\epsilon)$ corresponds to the integral

$$f(\epsilon) \propto \int_\infty^\infty \int_\infty^\infty e^{-\frac{x^2+y^2}{\sigma^2}} e^{-\frac{(x-\epsilon)^2+y^2}{\sigma_w^2}} \, \mathrm{d}x\mathrm{d}y$$

The true flux is $f(\epsilon = 0)$ and the fractional flux error is $\Delta f = \frac{f(\epsilon)-f(0)}{f(0)}$. After some straightforward but lengthy maths,

$$\begin{aligned}
\Delta f &= 1 - e^{\frac{\epsilon^2}{2\sigma_w^2}} \\
&\approx \frac{\epsilon^2}{2\sigma_w^2} \text{ if } \frac{\epsilon^2}{2\sigma_w^2} \ll 1 \quad\quad\quad\quad\quad\quad\quad (\text{B.1}) \\
&\Leftrightarrow \frac{\epsilon}{\text{FWHM}_{eff}} = \sqrt{\frac{\Delta f}{4\ln(2)}} \approx 0.6\sqrt{\Delta f}. \quad\quad\quad (\text{B.2})
\end{aligned}$$

With this formula, a 5% limit on the relative flux error is achieved by demanding that the misalignment be less than 10% of the effective PSF full width. With the $0\rlap{.}''2$ effective resolution, this means that the alignment error should be less than $0\rlap{.}''02$, equivalent to 0.44 Planetary Camera pixels.

B.1.2   Wrong PSF determination

If the assumed value for σ in the scaling factor (Eqn. 3.3) is not the correct value but in error by some ϵ, the scaling factor is wrong and hence the obtained flux value. The relative error on the flux measurement is

$$
\begin{aligned}
\Delta f &= 1 - \left(1 + \frac{(\sigma + \epsilon)^2}{\sigma_w^2}\right) \div \left(1 + \frac{\sigma^2}{\sigma_w^2}\right) \\
&\approx \frac{2\epsilon\sigma}{\sigma^2 + \sigma_w^2} \text{ if } \epsilon \ll \sigma &\text{(B.3)} \\
&\Leftrightarrow \frac{\epsilon}{\text{FWHM}} = \frac{\text{FWHM}_{\text{eff}}^2}{4\sqrt{2\ln 2}\text{FWHM}^2}\Delta f \approx \frac{\text{FWHM}_{\text{eff}}^2}{5\text{FWHM}^2}\Delta f &\text{(B.4)}
\end{aligned}
$$

For a 5% point source flux error, the PSF determination has to be accurate to about 1% if the weighting PSF has the same width as the image PSF and thus the effective FWHM is twice the intrinsic. It can be larger than that when the desired effective FWHM is much larger than the intrinsic FWHM, as is true for the HST data considered here.

B.2   Pointing accuracy of HST

The HST can be rotated about all three axes. The observer specifies a desired pointing (in right ascension and declination) and roll angle of the telescope. The actual pointing and roll are measured during the observation. The pointing specifies a sky coordinate to be imaged on a fixed position in the focal plane of the Science instrument used. This position, or its camera pixel coordinates, is called the reference aperture. This is an absolute pointing specification. In addition, an offset relative to the reference aperture can be specified. This is called `POS_TARG` and moves the image in the focal plane, *e. g.*, to enable a correction for chip defects.

The achievable accuracy differs for absolute and relative measurements.

B.2.1   Absolute pointing

The absolute pointing of the telescope is established using a catalogue of guide star positions. The guide stars are observed with the Fine Guidance Sensors (FGS). The FGS are interferometers which are the HST's astrometric instruments. In these observations, two FGS were used to point the telescope. One star is observed by each of the FGS, whose directions are separated by about 90°. From the star positions on the FGS, the position of the FGS in the spacecraft reference frame and from the spacecraft velocity, a fit is derived to the celestial coordinates of the guide stars. Knowing the position of the science instrument reference aperture relative to the FGS, the pointing of the reference aperture is calculated. (See "Pointing Calculations and Sources of Error" in Ch. 2 of [49].)

The limiting factor in determining the absolute pointing are the uncertainties of Guide Star Catalogue positions, and transformation between the FGS and the Science Instrument, in our case WFPC2. The GSC coordinates have typical RMS uncertainties of $0\overset{''}{.}5$ in each coordinate. The transformation from the FGS to WFPC2 introduces an error of order $0\overset{''}{.}1$ (see [105], 28.4, [8], 7.7.1 and [49], Ch. 1). The resulting total error of $0\overset{''}{.}51$ corresponds to 11 Planetary Camera pixels, *i. e.,* the measured absolute pointing scatters this much about

the true pointing. Hence, absolute coordinates as reported by the telescope should not be used for astrometry, to align images from different telescopes, *e. g.* Where it is required, an absolute pointing measurement should be obtained by independent astrometry of field objects and a fit to observed on-chip positions, taking into account any distortion and scale change effects.

A different issue to how accurately the pointing of any exposure can be known is the question of the pointing repeatability, *i. e.,* how far the executed pointings differ for two identical requested pointings. This is discussed in Sect. B.2.2.

B.2.2   Relative pointing

The relative pointing error limits the accuracy with which all the images can be aligned for summation and for calculation of the spectral index map. Even if we are not interested in the absolute position, the pointing of any single exposure relative to all others in the program has to be known to the required accuracy. The limiting issues are the pointing and roll angle repeatability and accuracy of `POS_TARG` offsets. For the present case, where one exposure takes up a complete orbit and the guide stars have to be reacquired at the beginning of each exposure, a value better than 5 mas (0.1 pix) was typical in Hubble Deep Field exposures ([8], 7.7.3). The dominant error is the pointing repeatability. The additional `POS_TARG` from the reference aperture position does not increase that error significantly.

When assessing relative pointing, there are again two contributors to the uncertainty: first, the repeatability of the pointing of the telescope as such, and second, uncertainties in the positions of the camera chips with respect to the vehicle frame.

The latter will not be a problem when assessing relative shifts as long as the relative positions of the four camera chips remain constant to the desired accuracy. Although there are long-term variations in the relative *positions* of the four WFPC2 chips of up to 100 mas, the relative *angles* have been constant to 0.01 degrees. On the time-scales between exposures in this program, no changes should occur.

The telescope pointing constancy during one exposure depends on the tracking mode of the HST. For the present data set, the most accurate mode, "Two FGS Fine Lock Guiding", was employed. It uses the same two guide stars in the FGS as the pointing calculation, one to establish the pointing and a second to fix the roll angle. Once the roll and pointing are established, they remain very stable during an exposure. However, even if the same pointing and roll are specified in two different exposures, there is a small non-repeatability as the roll guide star is not necessarily placed at the same position on the second FGS.

The telescope experiences so-called "jitter", vibrations due to thermal effects when entering or leaving the Earth's shadow, *e. g.* These typically result in 2–3 mas RMS variations of the pointing during an exposure (see Fig. B.1). Usually, the guide star lock is not lost. If the jitter becomes too strong, however, the guide star lock is endangered. The pointing control is temporarily transferred to the on-board gyroscopes if they detect a pointing deviation in excess of 20 milliarcseconds. This is called a "recentering event". During the recentering, data continue to be acquired. It may thus lead to motion of the target in the focal plane or blurring if it last for a significant fraction of the exposure time.

A flag indicates whether a recentering event has taken place for any given exposure. Recentering occurred in two exposures of the exposures in this set. They lasted 4 and 6 seconds, respectively, which is negligible compared to the total frame exposure time. The images and "jitter balls" show no signs of a pointing shift. If the jitter becomes even stronger,
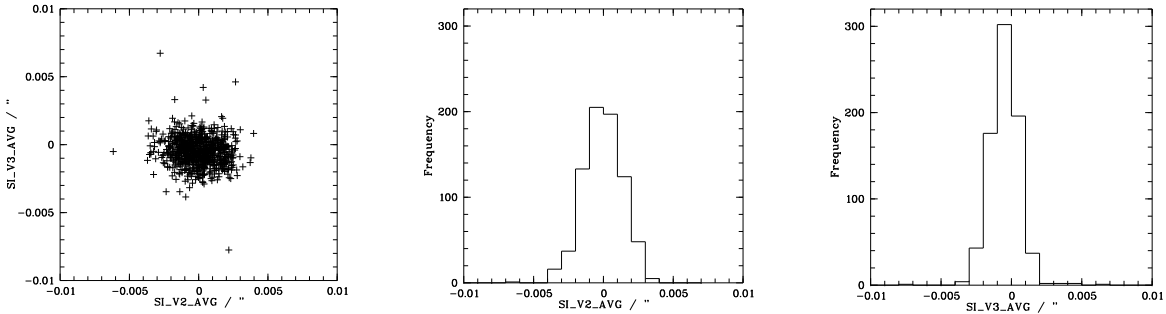
Figure B.1:   *Left,* Example of a "jitter ball". The vehicle motion parallel to the telescope axes V2 and V3 at the position of the science instrument reference aperture, averaged in three-second intervals, is plotted. *Right,* histograms of the jitter values along both axes. The zero-point is arbitrarily chosen as the average of the first few jitter values.

the guide star lock will be given up and the stars reacquired later. This will take about three minutes, during which the exposure is interrupted. The loss of lock is reported in the exposure. No loss of lock occurred during these exposures.

The telescope records its pointing and roll every second. The resulting Observatory Log file, also called jitter file, can in principle be used to establish the relative pointing between exposures. The accuracy of these measurements is about 10–20 mas, corresponding to 0.2–0.4 pixels.

The most accurate determination of relative pointing can be achieved if shifts of at least 4–5 objects can be measured for each exposure. A single object will be more inaccurate than the jitter file information because of the undersampling of the PSF by the camera pixels: at 300 nm, the diffraction limit of the HST's primary mirror is $0\rlap{.}''031$, while the nominal pixel size is $0\rlap{.}''04554$. The telescope PSF thus has a FWHM of 0.7 pixels. This undersampling leads to errors on the results of centroiding routines. Moreover, not all photons heading for one pixel actually liberate electrons which are counted in that pixel. This arises because some photons are scattered by the electrode structure of the front-illuminated CCD at wavelengths < 450 nm. Photons arriving near a pixel boundary are most likely to be scattered, which means that the detected signal depends on the sub-pixel positioning of the input signal, *e. g.*, a stellar image. This effect leads to some blurring of an originally sharp PSF [45].

The ability of the `MIDAS` [24] centring algorithm `CENTER/MOMENT` to find the true centre of a stellar image on the present data frames is tested in the following way: an oversampled Gaussian of varying width (FWHM 0.7–2.8 pixels representing both undersampled and well-sampled PSFs on the Planetary Camera chip) is shifted by $-\frac{2}{5}$ to $+\frac{3}{5}$ of a pixel in steps of $\frac{1}{5}$ pixel. The "image" is then resampled to the original, larger pixel size. A frame containing Poisson-distributed noise is added, to simulate sky background noise. With the available routines, it is not easily possible to simulate a true Poisson noise on the star counts. The central intensity of the Gaussian is held constant, so that the S/N of the central pixel is constant and results with differing width are comparable. This simulates the observation of an undersampled stellar image at various sub-pixel positions. The differences between the true position of the Gaussian and that recorded by `CENTER/MOMENT` routines are plotted in Fig. B.2. With the fixed central S/N, the wider Gaussians can be centered more accurately because more fit points are available in the wings. The test is repeated with a constant
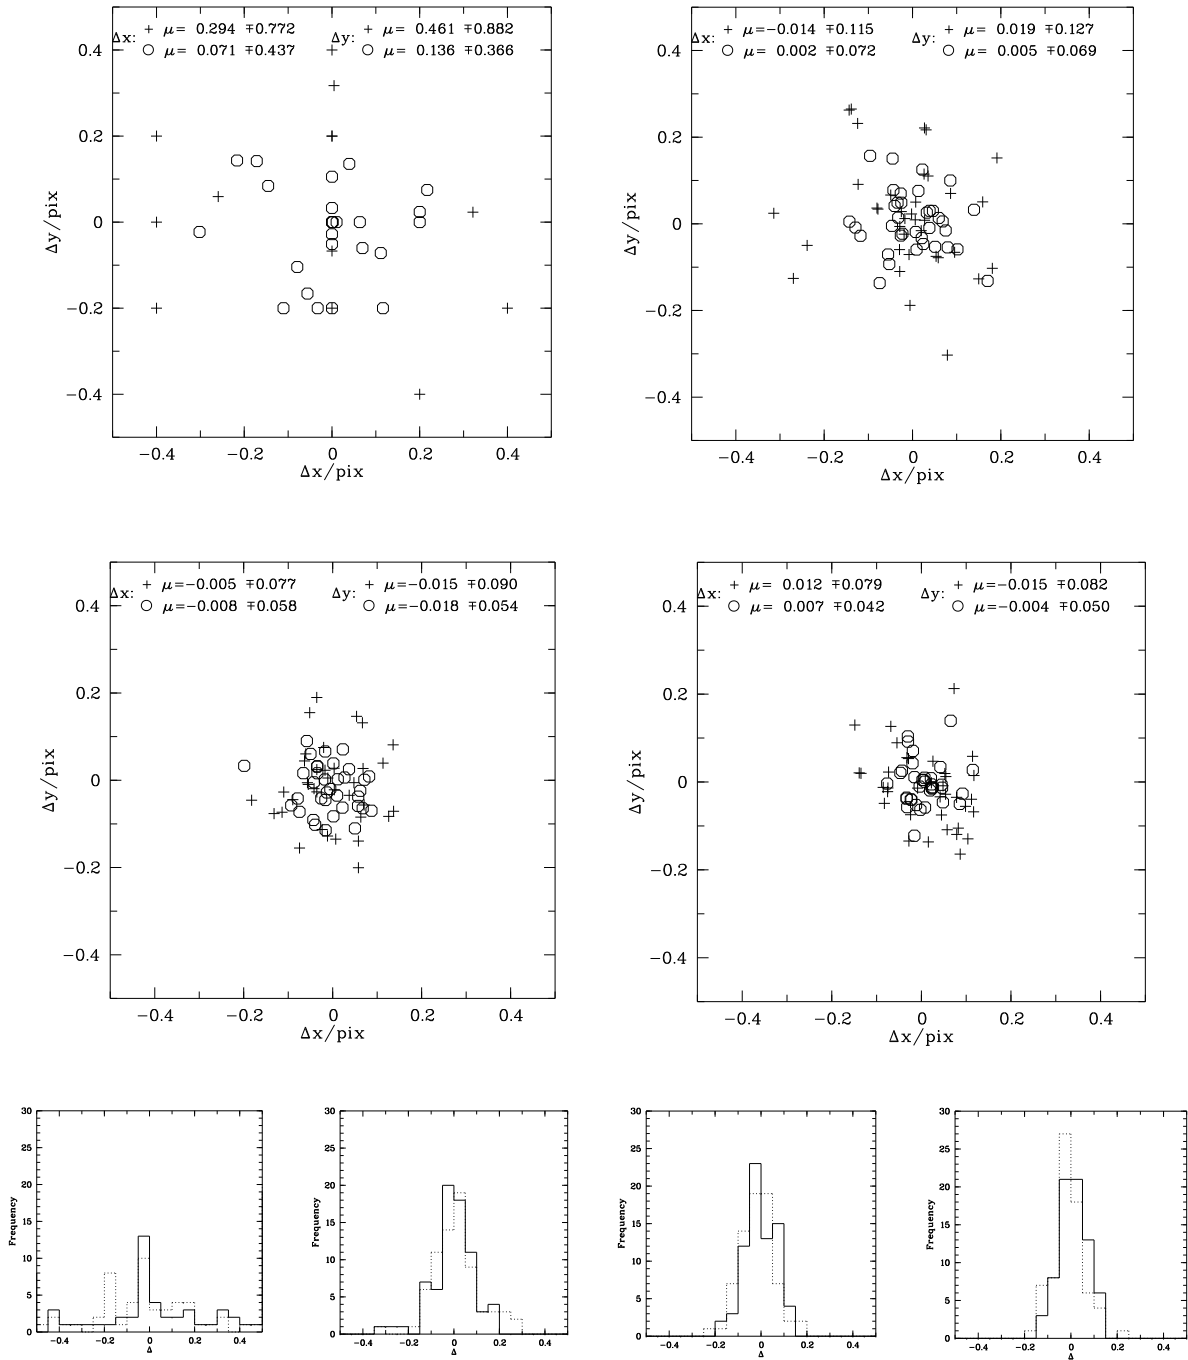
Figure B.2: Scatter plots of the difference between the true and the measured centroid coordinates. The dots and crosses refer to a central pixel S/N of ~ 45 and 30, respectively. The FWHM used is 0.7 pixels (top left), 1.0 pixels (top right), 1.3 pixels (bottom left) and 1.6 pixels (bottom right). The bottom four panels show the histograms of differences in both directions, in ascending FWHM order.

flux S/N summed over the whole PSF (central intensity inversely proportional to its width) and with a constant flux contained in the PSF (central intensity inversely proportional to its area). This changes the scattering behaviour as a function of width — a constant PSF

S/N has similar scatter for all widths, while a constant total flux means broader PSFs have a larger scatter. The overall magnitude of the effect remains the same, however. The values determined in this way agree with an assessment by STScI staff (centroiding errors *in excess of the statistical uncertainties* approximately uniformly distributed between -0.25 and 0.25 pixels (Lallo, *priv. comm.*)).

Additionally, one might imagine to study also the dependence of the fitted width on the sub-pixel centring. This would require a more accurate modelling of the telescope PSF. The real telescope PSF is much more complicated than either a Gaussian (most appropriate to a seeing profile for earth-bound observations) or the sinc function of an ideal circular aperture, because of the diffraction pattern caused by the secondary spider, the spherical aberration of the primary which has caused plenty of gnashing of teeth and the corrective optics used to remove these aberration effects. As we are not aiming to assess the capabilities of the HST *in principle*, we leave the topic noting that the sub-pixel positioning and the centring algorithm used do introduce an extra scatter.

The four centroidable objects are distributed on the four camera chips. In order to obtain shifts in the frame of the Planetary Camera, the relative location and rotation of the four camera chips has to be known.

The quasar core and at least two of the stars are saturated on the long exposures. Therefore, one short exposure has been included per visit at the same pointing as one of the science exposures. In this manner, there are non-saturated images enabling a centroiding measurements and an accurate determination of the telescope pointing differences between visits.

The same absolute pointing was requested for all exposures. Relative offsets between them were requested as `POS_TARG`. In this way, chip defects (*e. g.*, "bad rows/columns", hot pixels, charge traps, see below) can be corrected for as the same area of sky has always been mapped by at least two different regions on the chip. These relative shifts should be accurate to about 15 mas, corresponding to 0.3 pixels. Within one visit, the limiting factor will be the pointing repeatability at 5 mas level (see above). The total error in the alignment is thus ≈ 16 mas, corresponding to 0.35 pixels. This is comparable to the required alignment accuracy of 0.44 pixels.

The positions of the stars were measured using the task `imcenter` from the standard software IRAF (Image Reduction and Analysis Facility), a centroiding algorithm which finds the centres of gravity of the stellar images, independently in both coordinates (its internal workings are identical to `CENTER/MOMENT`). The positions thus obtained are corrected for the known and well-characterised geometric distortions of the WFPC2, using the `metric` task provided by the Space Telescope Science Institute (STScI). To transform the shifts on the different detectors to shifts in the systems of the PC, I used the plate scales and relative rotation angles of the WFPC2 chips as determined by the STScI Observatory Support Group for the date of the observations (Lallo, *priv. comm.*) (Tab. B.1). The positions of the stars and shifts obtained from them are listed in Table B.2. An arbitrary image in the set then has world coordinates $(0., 0.)$ assigned to the centre of the bottom left pixel. All images have a fixed scale (world coordinate change between neighbouring pixels) equal to the nominal pixel size, $0''\!.0455$. The relative shifts then define the world coordinates of the same pixel of all other frames.

There is an additional complication of this matter: a wavelength dependence of the pixel scale. The images presented here differ in wavelength by a factor of two, the pixel scales differ at the 0.1% level. This will affect the transformation of the stellar positions between the frames of the different cameras. However, the results will be consistent for exposures in

| Chip | Pixel scale | | Axis angle | | Rotation |
	x [''/pix]	y[''/pix]	x-axis	y-axis	wrt PC
PC1	0.04554	0.04552	134.69	224.69	0.00
WF2	0.09654	0.09963	224.22	314.22	89.53
WF3	0.09961	0.09951	314.52	44.52	179.83
WF4	0.09958	0.09967	45.09	135.09	270.40

Table B.1: Nominal scales and relative orientations of the WFPC2 detectors

| Data set | Measured shifts | | | | Jitter file | | Commanded | |
	x ['']	σ_x ['']	y['']	σ_y ['']	x ['']	y['']	x ['']	y['']
0109	−0.520	0.012	0.514	0.015	−0.507	0.505	−0.501	0.501
0203	−0.185	0.017	−0.310	0.013	−0.167	−0.308	−0.167	−0.334
0304	−0.330	0.007	0.352	0.016	−0.331	0.359	−0.334	0.334

Table B.2: Shifts of short exposures in visits 2–4, relative to visit 1, obtained from centroids of four objects, compared to commanded and jitter file shifts

any given filter. Care has to be exercised when aligning the exposures in two *different* filters. A full discussion is deferred to Sect. B.3.

B.2.3 Telescope roll

The pointing fixes the position of the reference aperture on the celestial sphere, *i. e.,* the direction of the telescope optical axis. This leaves the roll angle unspecified, *i. e.,* the rotation of the telescope about the optical axis.

The desired roll angle is specified in the telescope's frame U1,U2,U3 as the position angle from North of the U3 axis and stored in the keyword `ORIENT`. U1 is an axis parallel to the optical axis, U2 and U3 are fixed in the focal plane, with U3 pointing approximately along the PC main diagonal (from pixel 1,1 to 800,800). Science frames contain a predicted value for the roll angle both of the telescope and of the chip employed. The telescope's roll angle is stored as `PA_V3`, the position angle with respect to North of the V3 axis, which is an axis antiparallel to the U3 axis. The orientation of the chip on the plane of the sky is specified as `ORIENTAT`, which is the position angle of the detector's y-axis on the plane of the sky. Both values are merely calculated from the specified `ORIENT` using the camera position within the vehicle frame.

In addition, the "jitter file" contains a measured value of the telescope's roll angle, `ROLL_AVG`. This is again the position angle of the telescope's V3 axis *at the position of the reference aperture.* The difference between the angle from North at the V1 axis (parallel to U1) and that at the reference aperture position is only important for high declinations near the celestial poles, where the spherical coordinate system becomes non-linear. As 3C 273 is very close to the celestial equator, this correction will be unimportant. Similar to the pointing accuracy, the roll accuracy has to be such that no differences greater than 0''02 occur between different images. Assuming the quasar core can be placed at the same position on the chip to a high accuracy, different roll angles will lead to a different position of the outer end of

the jet at $20''$ separation. The desired roll angle accuracy is thus better than $0°057 = 3'26''$ which just produces a $0''02$ difference over 20" separation.

The roll repeatability (roll angle difference between exposures with identical roll requested) is at the $0°003$ (10") level [49]. The achieved roll angle was obtained from the jitter files. For the WFPC2 images, there were systematic offsets between visits, but only at the 6" level *i. e.,* well below the desired values. The images were thus taken to be all at constant roll angle. The NICMOS exposures were taken at one of two different roll angles (see Sect. 2.2.5), the first two visits at $81°51$ and the third visit at $85°51$. Changes within one visit are below the arcsecond level and thus negligible.

Apart from the roll repeatability, an error in the combination of images from different instruments might also arise from the lack of knowledge of the *absolute* rotation angle of the two cameras within the HST's focal plane. There is, *e. g.*, a difference of order $0°2$ in the quoted rotation angle of the WFPC2 camera within the focal plane between the current values in Lallo [50] and the values provided with the images, which are taken from Cox [18]. After consultation with STScI staff (M. Lallo, *priv.comm.*), we chose the rotation angle as provided with the observations, rather than the "current" value.

B.3   Geometric distortions of the focal plane

The focal surface of the Planetary Camera camera is not exactly a plane, so that the sky is not mapped onto the chip by a simple linear transformation. This has to be accounted for when measuring fluxes at different positions in the jet, and especially when comparing two images directly. For the project involving flux measurements on data frames taken by various different instruments, the choice method will be the transformation of the desired sky coordinates into the system of each data frame. This is discussed fully in Sect. 3.1. For the combination of the WFPC2 frames to the optical spectral index map presented in [42], the following procedure is sufficient.

As indicated in [103], there is a wavelength dependence of the sky-to-chip transformation. This is introduced by differential refraction in the MgF_2 field-flattener window in front of the camera. The filters themselves all introduce the same focus shift and thus do not contribute a further wavelength-dependence. Earlier investigations of the geometry of the focal surface did not include this wavelength dependence [37]. The transformation needed to remove the geometric distortion to sufficient accuracy over the whole chip is a bicubic transformation. Its coefficients are given in a form such that a corrected image has the same scale as an image taken at a wavelength of 555 nm.

To first order, the geometric distortion is a pixel scale change. When two frames at different wavelengths have been aligned perfectly using a point source, objects away from the alignment reference will not coincide. This scale change has to be accounted for, especially when performing an operation like constructing a spectral image map, which involves a division of the two images. A scale difference will introduce spurious gradients in the result. As discussed above (Sect. B.1.1), the difference in the apparent positions of two images of a source taken through different filters must be less than 10% of the size of an effective resolution element. If the quasar core images are made to coincide, the images of the head of the jet at $20'' \sim 450$ pixels separation have to coincide to better than 0.44 pixels. This means that the scale difference has to be less than $\frac{0.44}{450} \sim 0.1\%$ (assuming there are no other effects).

From the theoretical parameters given in [103] obtained by a raytracing algorithm for the
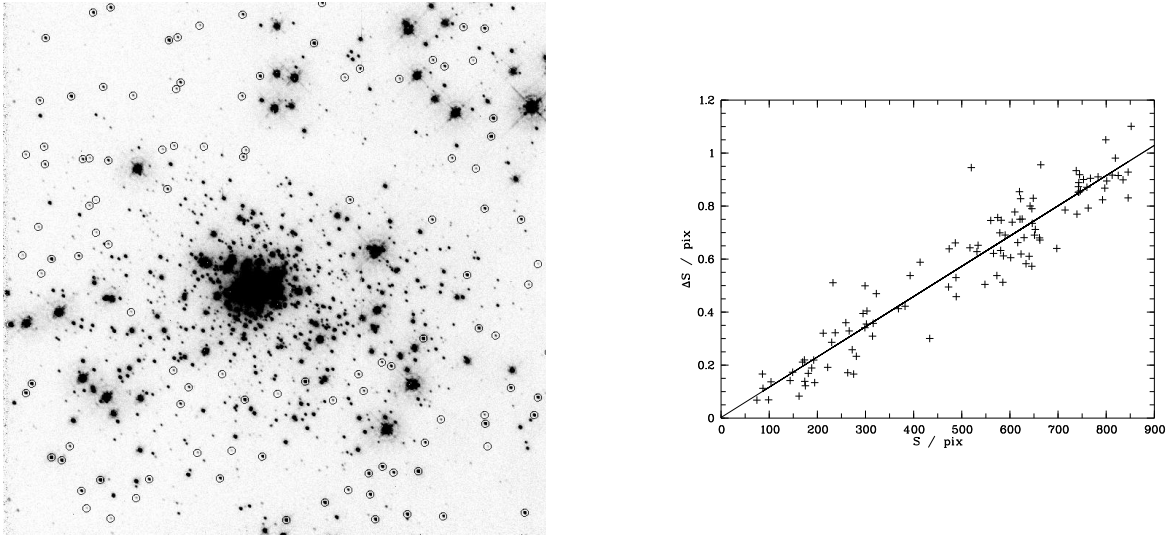
Figure B.3: *Left,* U-band image of R136. The stars whose positions were determined are encircled. *Right,* difference of star separations between U and R-Band ΔS plotted against separation S on R-band image, together with best-fitting straight line. The difference grows systematically with the separation, indicating a scale change. The slope of the line is 1.14×10^{-3}, the intercept with the ordinate is 0.2×10^{-2}

camera optics, the relative scale change is of the order of one part in 1000 between 300 nm and 600 nm which are the central wavelengths of the filters employed here. The effect is thus large enough to be significant. Note that to higher orders than the linear, the distortion not only changes the *position* of objects on the chip, but also reduces the area of sky mapped on one camera pixel near the corners of the chip. As the jet is imaged near the centre of the camera chip in all frames and the distortion parameters are negligibly small in this central region, the non-linear distortion term has no significant effect within one frame.

With the scale difference, using the quasar as position reference is the only way to align the exposures in the different filters. As only relative shifts are known accurately, it would otherwise be unclear where the common reference point of the two exposures might be.

The scale difference was confirmed experimentally by analysing two images of the stellar cluster R136 (Fig. B.3) taken through filters similar in wavelength to those used in the present work. The images were obtained from the archive of HST exposures[1]. For a number of well-isolated well-exposed stars, positions are determined on the Planetary Camera chip by fitting a two-dimensional Gaussian to the stellar images. The distances to an (arbitrary) reference star are computed in both filters. If there is a scale difference, the distances will be systematically larger in one filter compared to the other.

Such a scale difference does indeed exist, as can be seen from the plot in Fig. B.3. The parameters of the least-squares-fit line are tabulated in Tab. B.3. The slope of the best-fit line is the scale factor between the two plate scales. The measured separations on chip are systematically larger by $\sim 0.1\%$ in the U-band than in the R, *i.e.*, the scale of the U-band image is 0.1% *smaller* than that of the R-band image. The expected significant scale difference between the filters has thus been confirmed. Its effect has to be removed.

In order to remove the scale difference, the images were resampled by linear interpolation to a common pixel size. This was chosen as one tenth of the average plate scale of the two

---

[1]proposal 5589, PI John Trauger

	Slope and error		Intercept and error		Corr. coeff.
Original	1.14×10^{-3}	0.04×10^{-3}	0.2×10^{-2}	8.8×10^{-3}	0.94
After resampling	0	0	2.0×10^{-2}	6.6×10^{-3}	0.02

Table B.3: Parameters of the linear least squares fit of difference in separation ΔS against separation S on the R-band image. ΔS and S are uncorrelated on the processed image, showing that the scale difference has been removed.

frames. This reduces the mismatch by a factor of ten, to a negligible 0.01%, or 0.05 pixels. To verify the usefulness of this approach, the images of R136 have been subjected to the same treatment. On the resampled images, the scale difference has been removed. ΔS is now of the order of the centroiding accuracy in the undersampled U-band image (*cf.* Sect. B.2.2) for all values of S, with the exception of a few outliers which are most likely centring errors (not included in the fit).

In summary, the best way of stacking the images to a final version is the following: within one visits, rely on the jitter file information. All stellar objects are overexposed on the long exposures, so accurate centroiding is not possible. All long exposures within one visit are summed directly. This is possible because the offsets within one visit are in integer pixel steps. The alignment of the three U visits relative to each other is done with the centroiding measurements on the short exposures (Tab. B.2). The U visits are summed up only after the rebinning, as the offsets between visits are nearer $1/6^{\text{th}}$ of a pixel. Finally, the position of the quasar core as obtained on the short exposure is used to align the R and U images with each other — the scale difference between the filters leaves no other choice for the fixed common reference point.

Appendix C

Investigation of IR-optical spectral gradients

The infrared-optical spectral index map presented in Sect. 3.5.2 showed a gradient typical of a misalignment. We consider here whether the gradient could be spurious and due to a misalignment or a wrong subtraction of the diffraction spike.

C.1   Possibility of misalignment

To check the magnitude of any possible misalignment or rotation which might have introduced a spurious gradient, we shift and/or rotate the infrared and optical photometry positions with respect to each other.

C.1.1   Shifts

Shifting the photometry apertures on the infrared images with respect to the optical image does not lead to an absence of a gradient, it appears either at the inner or at the outer end. To asses whether there might be systematic errors in one visit leading to the observed gradient, we considered photometry performed separately only the images from each visit.

An attempt was undertaken to derive the shifts between the photometry results determined separately from individual visits. A shift perpendicular to the jet direction of order $0\rlap{.}''03$ was indeed detected; however, the *magnitude* of the shift between two images varied by factors of up to 2, depending on the pixel size used for the photometry. The attempt therefore proved inconclusive.

We conclude that there is an error in the determination of the pointing from the jitter files of the order of $0\rlap{.}''03$, which is the desired accuracy. Even if the relative pointing of all the near-infrared frames with respect to each other could be established by some other means, by cross-correlating the images, *e. g.*, it is impossible to derive an accurate value for the relative shift to images at other wavelengths. Since spectral index features are in essence due to morphological differences, morphological features at different wavelengths must not be used to perform an alignment, as this method might erase information that is present.

C.1.2 Rotations

Including rotations does make the spectral index map more symmetrical. However, the magnitude of the rotation about the centre of the index map needs to be of the order of a few degrees to produce a noticeable effect. This is far exceeding the nominal roll accuracy of the HST of $0°003$ (*cf.* "Telescope roll" on page 99). Since the rotation that needs to be applied to remove the gradient is much larger than any rotation that might arise from one of the known error sources, we cannot explain the gradient as being spurious and due to a rotation error.

C.2 Possibility of diffraction spike subtraction error

An error in the modelling and subtraction of the diffraction spike (Sect. 2.2.5) might be an error source for the spectral index. For this to be possible, the residual of the spike subtraction needs to be comparable to the jet signal. The total signal in the diffraction spike, scaled to the sum image of the first two visits, is $\approx 1.4 \times 10^6$ photoelectrons, compared to a jet signal (after the spike subtraction) of 7.4×10^6 photoelectrons. In region A ($13''$ from the core), the jet signal per pixel is typically around 5000 photoelectrons, and up to 7000 photoelectrons at the brightness peak. In comparison, the spike signal amounts to a few hundred photoelectrons per pixel. Even without a correction for the spike, it contributes at most 10% of the jet signal. A given relative flux error leads to the same relative error in the spectral index (Eqn. 3.6), so the spectral index error due to a wrong spike subtraction cannot be larger than 10%. At larger angles from the quasar, the diffraction spike signal decreases (Fig. 2.13) and the image appears cosmetically better (Fig. 2.14), so the residual error is expected to be significantly below 10%.

A spectral index map generated only from the third visit, which is not affected by the spike, is not qualitatively different from the spectral index map generated from all images; the only change is a slight decrease of the transverse gradient at the inner end. Any error due to the spike subtraction cannot therefore cause the observed gradient.

C.3 Conclusion

None of the considered error sources (shifts, rotations) can account for the spectral gradient. Residual pointing errors at the $0''03$ level are impossible to correct, since no other method to establish the telescope's true pointing than use of the "jitter files" is available. Further observations are needed to clarify the issue.

Appendix D

Physical background

D.1 Physics of synchrotron radiation

The observations of extragalactic radio sources clearly show that the emission is due to the synchrotron mechanism. As a convenient overview, important results from the theory of synchrotron radiation are summarised here. Full derivations will not be reproduced; they are given in [56], [78] or [91], *e. g.* SI units are used throughout.

D.1.1 Radiation of individual charged particles

A moving charge has an electric field whose strength has a $\frac{1}{r^2}$ dependence, exactly like a static charge. An accelerated charge, however, has an electric field E varying with the inverse of the distance. This means that the total energy of the field threading spherical shells of increasing radius r remains constant, *i. e.*, energy is radiated: the flux is given by the Poynting vector, $\mathbf{P} \propto \mathbf{E} \times \mathbf{B} \propto |\mathbf{E}|^{\mathbf{2}} \propto \frac{1}{r^2}$. Hence, the total energy = flux $\cdot$ area $\propto \frac{1}{r^2} \cdot r^2$ crossing any shell remains constant and energy can be transmitted to large distances.

A particle of mass m and electromagnetic charge q moving in a magnetic field $\mathbf{B}$ at velocity $\mathbf{v}$ with magnitude βc experiences the familiar Lorentz force $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$ perpendicular to both the field direction and the direction of its velocity, causing it to follow a helical path. The motion of the particle can be thought of as the superposition of two motions: motion of the guiding centre at constant speed parallel to the magnetic field and a circular motion around the guiding center. The speed of the former is the component of the particle's initial velocity $\mathbf{v}$ parallel to the field, $\frac{\mathbf{v} \cdot \mathbf{B}}{|\mathbf{B}|}$. The pitch angle ψ of the helix is then given by $\cos \psi = \frac{\mathbf{v} \cdot \mathbf{B}}{|\mathbf{v}||\mathbf{B}|}$. The circular frequency of the motion about the guiding center is called the relativistic gyrofrequency ω_r:

$$\omega_r = \frac{qB}{\gamma m} \tag{D.1}$$

where $\gamma = (1 - \beta^2)^{-1/2}$. It is related to the non-relativistic gyrofrequency ω_g through the additional factor γ in the denominator: $\omega_r = \frac{\omega_g}{\gamma}$. The acceleration towards the guiding center causes the particle to emit radiation. In the case of charges moving at relativistic velocities, the radiation is called synchrotron radiation, as it was first observed in particle accelerators of the synchrotron type [23]. The total energy loss (or cooling) rate of a relativistic particle

due to synchrotron radiation is

$$-\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right) = \frac{q^4 B^2}{6\pi\epsilon_0 cm^2}\beta^2\gamma^2\sin^2\psi \tag{D.2}$$

Because of the inverse dependence on the particle's mass squared, radiation losses by protons or other nuclei are much smaller than those by electrons. (The energy loss of positrons is identical to that of electrons, only the polarisation properties are different. For simplicity, we mention just "electrons" instead of "electrons or positrons"). The synchrotron radiation from radio sources will therefore be dominated by emission from electrons. All following formulae will therefore be written with the electron charge e and mass m_e. For an electron, some of the physical constants in Eqn. D.2 can be joined into the Thomson cross-section σ_{T}. As the square of the magnetic induction appears, the magnetic field can be expressed in terms of its energy density $U_{\mathrm{mag}} = \frac{B^2}{2\mu_0}$:

$$-\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right) = 2\sigma_{\mathrm{T}}cU_{\mathrm{mag}}\beta^2\gamma^2\sin^2\psi \tag{D.3}$$

This equation applies to particles with a fixed pitch angle. Averaging over an isotropic pitch angle distribution $p(\psi)\mathrm{d}\psi = \frac{1}{2}\sin\psi\mathrm{d}\psi$ yields

$$\begin{aligned}-\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right) &= \frac{4}{3}\sigma_{\mathrm{T}}cU_{\mathrm{mag}}\beta^2\gamma^2 \\ &= 1.058\times10^{-14}\left(\frac{B}{\mathrm{T}}\right)^2\gamma^2\beta^2\mathrm{W}\end{aligned} \tag{D.4}$$

This is appropriate for an ensemble average over an electron distribution isotropic in pitch angle, or as life-time average for one given electron assumed to undergo a large number of scattering events before losing a significant fraction of its energy. The lifetime of an electron with energy E against synchrotron radiation cooling is approximately

$$\tau_r \approx \frac{E}{-\frac{\mathrm{d}E}{\mathrm{d}t}} = \frac{\gamma}{-\frac{\mathrm{d}\gamma}{\mathrm{d}t}} = 1.66\times10^4\left(\frac{\gamma}{10^7}\right)^{-1}\left(\frac{B\sin\psi}{\mathrm{nT}}\right)^{-2}\mathrm{yr} \tag{D.5}$$

Since the energy loss is dependent on the *square* of the energy, this is not an e-folding time: a particle will lose *most* of its energy in this time. Note that the lifetime is inversely proportional to the energy. The lifetime can be turned into a synchrotron *loss scale*, the maximum distance a particle with initial Lorentz factor γ can travel in a field **B**:

$$\lambda_{\mathrm{loss}} \le c\tau_r \approx \frac{5}{\beta}\left(\frac{\gamma}{10^7}\right)^{-1}\left(\frac{B\sin\psi}{\mathrm{nT}}\right)^{-2}\mathrm{kpc} \tag{D.6}$$

As already shown by Liénard in 1898 [54], the radiation of a highly relativistic charge is beamed into the forward direction. Most of the power is emitted into a cone of half-opening angle θ given by

$$\theta = \gamma^{-1} = \frac{mc^2}{E} \tag{D.7}$$

where $\beta = {}^v\!/_c$ and E is the total energy of the particle. Because of the beaming, part of a particle's trajectory has to be along the line of sight to the observer for any radiation to
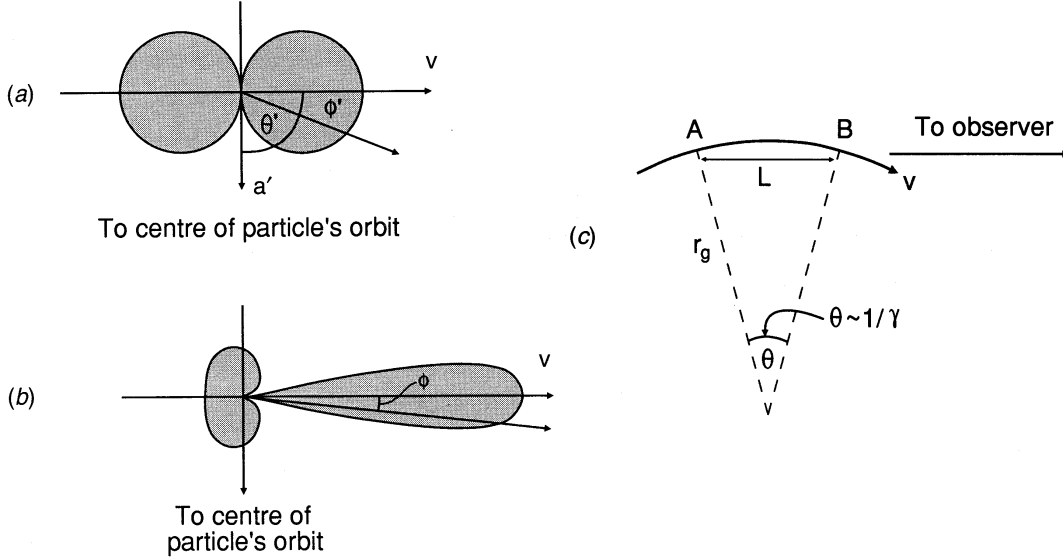
Figure D.1: Radiation pattern of a relativistic accelerated charge. (a) shows the pattern in the charge rest frame, (b) the pattern as seen by an external observer. (c) illustrates the condition for observation of radiation from the charge during a short part of the helical trajectory in a magnetic field. Figure 18.5 from [56]

be observed at all. When the particle moves in a magnetic field, the beaming cone sweeps through the line of sight only for that part of the trajectory during which the angle between the direction of motion and the line of sight is smaller than $1/\gamma$ (Fig. D.1). As the beaming implies that an electron emitting observed radiation moves along the line of sight towards the observer, ψ is just the angle between the magnetic field vector **B** and the line of sight.

The observer only sees a short pulse of radiation as the beaming cone sweeps through the line of sight. For this reason, the observed power spectrum is not that instantaneously emitted by the electron, but the Fourier transform of the short pulse (accounting for time dilation and Doppler shifting). There is therefore a maximum frequency of the radiation in the pulse much higher than the gyrofrequency ω_g, corresponding to the inverse of the pulse duration. Following the argument in [56] (p. 238), the pulse duration is roughly $\Delta t \approx \frac{1}{2}\gamma^2\omega_g$. The maximum frequency observed then turns out as $\nu \sim \gamma^2\nu_g \sin\psi$ for a particle at arbitrary pitch angle ψ.

An exact calculation of the spectrum of an ultra-relativistic electron in a magnetic field ([91] or [56], *e. g.*) shows that the spectrum is continuous, but with a sharp peak, and extends to larger frequencies than just described. There is a characteristic frequency ν_c which is useful in parameterising some of the integrals occurring in the detailed maths. It roughly corresponds to the frequency at which the total power spectrum peaks; in fact, the maximum is at $0.29\nu_c$ with

$$\nu_c = \frac{3}{2}\gamma^2\nu_g \sin\psi = \frac{3}{4\pi}\frac{e}{m}\gamma^2 B \sin\psi \qquad (D.8)$$

The presence of a significant synchrotron emission at a frequency ν_{obs} implies that there are particles with the "correct" combination of γ^2 and $B \sin\psi$ such that $\nu_c \approx \nu_{obs}$. With the critical frequency, Eqn. D.6 can be rewritten to give the maximum distance a particle radiating at a frequency $\nu \approx \nu_c$ can travel. One obtains a loss scale of a few kiloparsec for
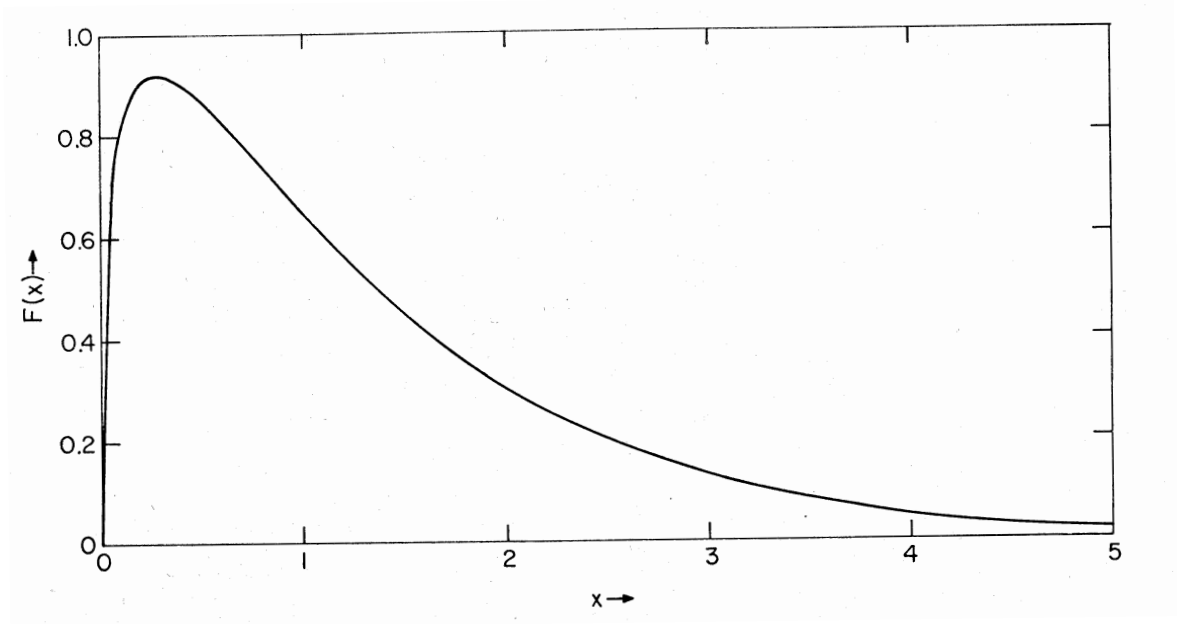
Figure D.2: Shape of the synchrotron spectrum emitted by a single charged particle

photons radiating ultraviolet light ($\nu \approx 10^{15}$ Hz) in a field of 1 nT:

$$\lambda(\nu) \lesssim 1.6\,\text{kpc} \left(\frac{\nu}{10^{15}\text{Hz}} \right)^{-\frac{1}{2}} \left(\frac{B \sin \psi}{\text{nT}} \right)^{-\frac{3}{2}} \tag{D.9}$$

The emissivity is different in the two polarisations, parallel and perpendicular to the magnetic field. The radiation is therefore polarised. The total synchrotron spectral power (in both polarisations) emitted by a single electron is

$$w(\nu) = \frac{\sqrt{3}e^3 B \sin \psi}{4\pi\epsilon_0 c m_e} F(x) \tag{D.10}$$

(elegantly derived in [91]), where $F(x)$ is related to a modified Bessel function[1] and $x \equiv \frac{\nu}{\nu_\text{c}}$. $F(x)$ is tabulated, *e. g.*, in [78]. A plot of $F(x)$ is shown in Fig. D.2.

For the qualitative analysis, we summarise here the dependence of the total emitted power, emissivity, characteristic frequency and lifetime as function of the magnetic field B, and dimensionless energy factor γ:

- Power loss:

$$-\left(\frac{\text{d}E}{\text{d}t} \right) = 1.1 \times 10^{-20} \left(\frac{B \sin \psi}{\text{nT}} \right)^2 \beta^2 \left(\frac{\gamma}{10^6} \right)^2 \text{W} \tag{D.11}$$

- Loss scale for electrons emitting at frequency ν and wavelength λ:

$$\begin{aligned} \lambda_\text{loss} &\lesssim 1.6\,\text{kpc} \left(\frac{\nu}{10^{15}\text{Hz}} \right)^{-\frac{1}{2}} \left(\frac{B \sin \psi}{\text{nT}} \right)^{-\frac{3}{2}} \\ &= 1.6\,\text{kpc} \left(\frac{\lambda}{300\text{nm}} \right)^{\frac{1}{2}} \left(\frac{B \sin \psi}{\text{nT}} \right)^{-\frac{3}{2}} \end{aligned} \tag{D.12}$$

---

[1] $F(x) = x \int_x^\infty K_{5/3}(z)\text{d}z$ where $K_{5/3}$ is the modified Bessel function of order $5/3$

- Lifetime:

$$\tau_r \approx 1.66 \times 10^4 \left(\frac{\gamma}{10^7}\right)^{-1} \left(\frac{B\sin\psi}{\mathrm{nT}}\right)^{-2} \mathrm{yr} \tag{D.13}$$

- Characteristic frequency:

$$\nu_c = 4.2 \times 10^{15} \left(\frac{\gamma}{10^7}\right)^2 \left(\frac{B\sin\psi}{\mathrm{nT}}\right) \mathrm{Hz} \tag{D.14}$$

- Spectral power:

$$w(\nu) = 2.34 \times 10^{-34} \left(\frac{B\sin\psi}{\mathrm{nT}}\right) \mathrm{WHz}^{-1} F\left(\frac{\nu}{\nu_c}\right) \tag{D.15}$$

D.1.2 Emission of an ensemble of electrons

The radiation we observe is emitted by many electrons, and many electrons of different energies contribute to emission at a given frequency from a fixed region of a source. The total emissivity at a given frequency is therefore given by the convolution integral of the number density of electrons with energy in the interval $E \to E + dE$, written as $n(E)dE$, and the spectral power w of one electron of that energy:

$$J(\nu) = \int_0^\infty w\left(\frac{\nu}{\nu_c(E)}\right) n(E)dE \tag{D.16}$$

As a reminder that the emissivity depends on the electron energy through the critical frequency, ν_c is written as function of energy, $\nu_c(E)$. It is also assumed here that the distribution $n(E)dE$ is isotropic in pitch angle ψ (ν_c also depends on ψ through Eqn. D.8).

Simplest case: electrons in a homogeneous, uniform field

The emissivity is proportional to the function $F(x)$. Its sharply-peaked shape corresponds roughly to a δ-function. A convolution of any function with a δ-function leaves the shape of that function unchanged, in this case the electron distribution in energy $n(E)$. Since the observed spectra have a power-law shape $I(\nu) \propto \nu^\alpha$, the electron distribution is also assumed to be of this form, *i. e.*, $n(E)dE = \kappa E^{-p}dE$. This argument can be taken further to derive the relation between the exponents α and p: as all of the electron's energy is assumed to be radiated at one frequency, the emissivity (Eqn. D.16) can then be expressed

$$J(\nu)d\nu = \left(-\frac{dE}{dt}\right) n(E)dE$$

Upon writing all terms on the right in terms of ν and the gyrofrequency ν_g, the following dependences of $J(\nu)$ are obtained:

$$J(\nu) \propto \kappa(B\sin\psi)^{(p+1)/2}\nu^{-(p-1)/2} \tag{D.17}$$

If one chooses to introduce a certain pitch angle distribution function, the dependence on the field B remains unchanged. The integration over pitch angles then only leads to different proportionality factors.

The emissivity power-law exponent α is related to the electron energy exponent p by

$$\alpha = -\frac{p-1}{2}. \tag{D.18}$$

For radio jets and lobes, exponents $\alpha \approx -0.5 \rightarrow 1$ are common, so that the "average" value for p will be about 2.5. This is a test for theories of particle acceleration in extragalactic radio sources (which will be considered in Sect. 1.1.3). A valid theory has to predict (or rather, *post*dict) the correct exponent for the accelerated electron spectrum. Although the shape of the emission spectrum from an electron distribution is the same as the energy spectrum of the electrons, its exponent is different. This arises from the quadratic dependence of the emitted frequency on the electron energy (Eqn. 4.1).

The convolution integral can also be evaluated exactly for each pitch angle. This only gives slightly different proportionality factors but no different functional form.

More realistic scenarios

In general, the magnetic field in an extragalactic radio source is not expected to be homogeneous on large scales. In addition, any line of sight is integrating the volume emissivity along the line of sight to give a surface emissivity. If there is turbulence, the field may be randomly "tangled", stretched and/or compressed, so that a distribution of pitch angles is observed in a telescope beam. Also, a single electron may be scattered during its trajectory so that all pitch angles are equally likely during its life-time (*cf.* Sect. D.1.1). The emissivity is therefore often averaged over an isotropic pitch angle distribution. The shape of the integrated spectrum remains the same, however. Isotropy is commonly assumed, but it should be stressed that this is an assumption.

A second point has to be addressed which has been ignored so far. All calculations above assume that the power-law distribution of electrons extends to infinite energies. Whatever the acceleration process, there will in fact be a maximum attainable electron energy (considered in more detail in Sect. 1.1.3). At frequencies corresponding to electron energies (through ν_c, Eqn. D.8) above the maximum electron energy, there will be no significant emission and the spectrum exhibits a high-frequency cutoff.

If the particles are accelerated in one region of the source and then travel within an extended source, for example, down a jet or into a radio lobe, the shape of the spectrum varies with position within, because high-energy particles lose a larger fraction of their energy per unit time than low-energy particles. This variation occurs even for a homogeneous field. The change of cutoff frequency over the source could be traced exactly only with a telescope which resolves the smallest emitting regions. A telescope beam with a typical spatial resolution of at least a few parsec will integrate the spectra over a large part of the source. The superposition of spectra with different cutoff frequencies results in a power-law spectrum with a break [34]. At low frequencies, the spectrum remains unchanged. Above the break energy, it can be shown that the electron spectrum becomes one power steeper (*i. e.,* , of the form $n(E)\mathrm{d}E \propto E^{-p-1}$) so that the emission spectrum is changed to $I(\nu) \propto \nu^{\alpha-\frac{1}{2}}$ using Eqn. D.18. Above the cutoff frequency, the emissivity drops to zero. Spectra of this kind have been observed in a number of hot spots of extragalactic radio sources, all of which show optical synchrotron emission [66, 68].

D.2 Cosmological distances

For reference, we summarise cosmological formulae here. In order to convert the received flux F from an astronomical object at redshift z into the emitted luminosity L, the *luminosity distance* d_L to a source is defined as

$$d_L{}^2 = \frac{L}{4\pi F}. \tag{D.19}$$

The luminosity distance is related to the comoving coordinate distance ϖ by

$$d_L = \varpi(1 + z). \tag{D.20}$$

In order to convert an angular size into a physical size (in the limit of small angles), the angular size is multiplied by the proper or angular diameter distance d_θ, given by

$$d_\theta = \varpi/(1 + z). \tag{D.21}$$

In a general cosmology with non-zero cosmological constant Λ, expression for the ϖ as function of z is an integral that can only be evaluated numerically [43]. For small redshifts, this expression for the proper distance can be approximated to first order in z as

$$d_L = \frac{cz}{H_0}\left(1 + \frac{1}{2}(1 - q_0)z\right), \tag{D.22}$$

where the deceleration parameter q_0 is given in terms of the matter density parameter Ω_0 and the energy density parameter associated with the cosmological constant Ω_Λ by

$$q_0 = \frac{1}{2}\Omega_0 - \Omega_\Lambda. \tag{D.23}$$

We shall use $H_0 = 70h_{70}\,\mathrm{km\ s^{-1}\,Mpc^{-1}}$ [28] and $\Omega_0 = 0.3, \Omega_\Lambda = 0/7$ as indicated by the most recent results [39, *e. g.*]. For 3C 273's redshift of $z = 0.158$ [97], we obtain

$$d_L = 760\,\mathrm{Mpc}, d_\theta = 565\,\mathrm{Mpc} \tag{D.24}$$

so that $1''$ corresponds to $2.7h_{70}^{-1}\,\mathrm{kpc}$.

Bibliography

[1] Abraham, Z. & Romero, G. E. 1999, A&A, 344, 61 5.1

[2] Baade, W. 1956, ApJ, 123, 550+ 1.1.2

[3] Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, A&A, 61, 99 2.2.1

[4] Bahcall, J. N., Kirhakos, S., Schneider, D. P., et al. 1995, ApJ, 452, L91 2.6, 2.2.3, 2.2.3, 3.2

[5] Begelman, M. C., Blandford, R. D., & Rees, M. J. 1984, Rev. Mod. Phys., 56, 255 1.1.2

[6] Bell, A. R. 1978, MNRAS, 182, 147 3, 5.3

[7] —. 1978, MNRAS, 182, 443 3, 5.3

[8] Biretta, J. A. et al., eds. 2000, WFPC2 Instrument Handbook, Version 5.0 (Baltimore: STScI) 2.2.2, A.1, A.2, A.2.2, A.3.3, B.2.1, B.2.2

[9] Blandford, R. D. & Rees, M. J. 1974, MNRAS, 169, 395 1.1, 1.1.2, 6

[10] Blundell, K. M. & Rawlings, S. 2000, AJ, 119, 1111 1.1.4

[11] Böker, T., ed. 2000, NICMOS Instrument Handbook, Version 4.0 (Baltimore: STScI) 2.1.3, 2.2.4

[12] Bushhouse, H., ed. 1995, Synphot Users' Guide, 3rd edn. (Baltimore: Space Telescope Science Institute), `http://ra.stsci.edu/documents/SyG_95/SG_35.html` 2.2.2

[13] Camenzind, M. 1993, in Lecture Notes in Physics, Vol. 421, Jets in Extragalactic Radio Sources, ed. H.-J. Röser & K. Meisenheimer (Berlin Heidelberg New York: Springer-Verlag), 109–122 1.2

[14] Celotti, A., Ghisellini, G., & Chiaberge, M. 2001, MNRAS, 321, L1 1.1.5

[15] Chartas, G., Worrall, D. M., Birkinshaw, M., et al. 2000, ApJ, 542, 655 1.1.5

[16] Conway, R. G. & Davis, R. J. 1994, A&A, 284, 724 4.1, 10, 5.1

[17] Conway, R. G., Garrington, S. T., Perley, R. A., & Biretta, J. A. 1993, A&A, 267, 347 5, 2.2.3, 3.2

[18] Cox, C. 1994, The WFPC2 scales and alignment, Report SOB-94-10-21 (STScI) B.2.3

[19] Cox, C., Ritchie, C., Bergeron, E., MacKenty, J., & Noll, K. 1997, NICMOS Distortion Correction, IRS OSG-CAL-97-007 (STScI)   3.1.3

[20] Curtis, H. D. 1918, Publications of Lick Observatory, 13, 31   1

[21] Dermer, C. D. 1995, ApJ, 446, L63   5.1

[22] Dickinson, M., ed. 1999, NICMOS Data Handbook, Version 4.0 (Baltimore: STScI)   2.1.3

[23] Elder, F. R., Langmuir, R. V., & Pollock, H. C. 1948, Phys. Rev. Lett., 74, 52   D.1.1

[24] European Southern Observatory. 1998, ESO-MIDAS Users' Guide, MIDAS release 98NOV (Munich: ESO)   A.3.3, B.2.2

[25] Fermi, E. 1954, ApJ, 119, 1+   3

[26] Ferrari, A. 1998, ARA&A, 36, 539   1.2

[27] Flatters, C. & Conway, R. G. 1985, Nat, 314, 425+   2.6, 2.2.3

[28] Freedman, W. L., Madore, B. F., Gibson, B. K., Ferrarese, L., Kelson, D. D., Sakai, S., Mould, J. R., Kennicutt, R. C., Ford, H. C., Graham, J. A., Huchra, J. P., Hughes, S. M. G., Illingworth, G. D., Macri, L. M., & Stetson, P. B. 2001, ApJ, 553, 47   D.2

[29] Greenstein, J. L. & Schmidt, M. 1964, ApJ, 140, 1+   1.1.1, 2, 1.1.2, 1.1.4

[30] Guthrie, B. N. G. & Napier, W. M. 1975, MNRAS, 172, 85   4.4

[31] Hardee, P. E. & Norman, M. L. 1989, ApJ, 342, 680   2.2.3

[32] Hazard, C., Mackey, M. B., & Shimmins, A. J. 1963, Nat, 197, 1037   1.1.1, 2

[33] HDF collaboration. 1997, Hubble Deep Field sky histograms, `http://www.stsci.edu/ftp/science/hdf/pipeline/Sky_histograms.gif`   A.2.2

[34] Heavens, A. F. & Meisenheimer, K. 1987, MNRAS, 225, 335   1.1.2, 6, 3.2, 3.5.3, 4.1, 4.3, 6.1, 6.2, D.1.2

[35] Heinz, S. & Begelman, M. C. 1997, ApJ, 490, 653+   5.1, 5.1

[36] Hippelein, H. & Meisenheimer, K. 1991, local software   A.3.3

[37] Holtzman, J. A., Hester, J. J., Casertano, S., Trauger, J. T., Watson, A. M., Ballester, G. E., Burrows, C. J., Clarke, J. T., Crisp, D., Evans, R. W., Gallagher, J. S., I., Griffiths, R. E., Hoessel, J. G., Matthews, L. D., Mould, J. R., Scowen, P. A., Stapelfeldt, K. R., & Westphal, J. A. 1995, PASP, 107, 156   B.3

[38] HST ETC. 1998, HST on-line exposure time calculator, `http://www.stsci.edu/ftp/instrument_news/WFPC2/wfpc2_site.html`   A.3.3

[39] Hu, W., Fukugita, M., Zaldarriaga, M., & Tegmark, M. 2001, ApJ, 549, 669   D.2

[40] Hughes, P. A. & Miller, L. 1991, in Beams and jets in astrophysics, ed. P. A. Hughes, Cambridge Astrophysics Series No. 19 (Cambridge University Press), 1   5.1

[41] Jester, S. 1999, An optical spectral index map of the jet in 3C 273, Progress report, Universität Heidelberg 2.2.2, 3.1.2, 3.5.2

[42] Jester, S., Röser, H.-J., Meisenheimer, K., Perley, R., & Conway, R. G. 2001, A&A, 373, 447 3.3.1, 3.5.2, 4.4, 5.2.2, B.3

[43] Kayser, R., Helbig, P., & Schramm, T. 1997, A&A, 318, 680 D.2

[44] Krist, J. 1999, Long-Term Trends in the NICMOS Camera 2 Obscuration Pattern and Aberrations, NICMOS ISR-99-011 (Baltimore: STScI) 2.2.5

[45] —. 1999, `Tiny Tim` frequently asked questions, `http://www.stsci.edu/software/tinytim/tinytim_faq.html` B.2.2

[46] —. 1999, `Tiny Tim` PSF modelling software, `http://www.stsci.edu/software/tinytim/tinytim.html` 3.1.1

[47] Krist, J. & Burrows, C. 1994, Large Angle Scattering in WFPC2 and Horizontal "Smearing" Correction, Instrument Science Report WFPC2 94-01 (Baltimore: STScI) A.3.2

[48] Laing, R. A. 1981, ApJ, 248, 87 3.3.1

[49] Lallo, M. 1998, Observation Logs Documentation, `http://www.stsci.edu/ftp/instrument_news/Observatory/obslog/OL_1.html` B.2.1, B.2.3

[50] —. 2001, The SI Aperture File, `http://www.stsci.edu/instruments/observatory/siaf.html` B.2.3

[51] Leahy, J. P. 1991, in Beams and jets in astrophysics, ed. P. A. Hughes, Cambridge Astrophysics Series No. 19 (Cambridge University Press), 100–186 3.3.1, 4.3.1

[52] Lelièvre, G., Nieto, J. L., Horville, D., Renard, L., & Servan, B. 1984, A&A, 138, 49 2.6, 2.2.3, 2.2.3

[53] Lesch, H. & Birk, G. T. 1998, ApJ, 499, 167+ 4.4

[54] Liénard, A. M. 1898, 'Eclairage électrique, 14, 16 D.1.1

[55] Longair, M. 1994, High-energy astrophysics, 2^{nd} edn., Vol. 2 (Cambridge: CUP), p. 276 5.1

[56] Longair, M. S. 1994, High energy astrophysics. Vol.2: Stars, the galaxy and the interstellar medium (Cambridge: Cambridge University Press, 2nd ed.) 4.3, D.1, D.1, D.1.1

[57] Manolakou, K., Anastasiadis, A., & Vlahos, L. 1999, A&A, 345, 653 4.4

[58] Marshall, H. L., Harris, D. E., Grimes, J. P., et al. 2001, ApJ, 549, L167 1.1.5, 5.2.2, 5.3, 5.6, 5.3, 5.7

[59] McLeod, B. A. 1997, in The 1997 HST Calibration Workshop with a new generation of instruments, ed. S. Casertano, R. Jedrzejewski, C. D. Keyes, & M. Stevens (Baltimore: STScI), 281+ 2.2.4

[60] Meisenheimer, K. 1983, PhD thesis, Universität Heidelberg   3

[61] Meisenheimer, K. 1990, Optische Synchrotronstrahlung aus extragalaktischen Radio-quellen, Habilitationsschrift, Universität Heidelberg   3.3.1, 4.1

[62] Meisenheimer, K. 1996, in Lecture Notes in Physics, Vol. 471, Jets from Stars and Galactic Nuclei, ed. W. Kundt (Berlin Heidelberg New York: Springer Verlag), 57+   1.1.4

[63] Meisenheimer, K. & Heavens, A. F. 1986, Nat, 323, 419   1.1.2, 3.2, 3.6, 4, 6.1

[64] Meisenheimer, K., Neumann, M., & Röser, H.-J. 1996, in Lecture Notes in Physics, Vol. 471, Jets from Stars and Galactic Nuclei., ed. W. Kundt (Berlin Heidelberg New York: Springer Verlag), 230+   5, 3.5.2, 3.5.3, 4.3.1, 10, 4.8, 4.4, 4.9, 4.11, 4.4, 5.2.2, 5.3, 6.1

[65] Meisenheimer, K. & Röser, H. J. 1993, in Astronomy and Astrophysics, Extension and Supplement to Volume 2, Subvolume a: Instruments, Methods, Solar System, ed. H. H. Voigt, Landolt-Börnstein (Springer-Verlag, Heidelberg), 17–33   3

[66] Meisenheimer, K., Röser, H.-J., Hiltner, P. R., et al. 1989, A&A, 219, 63   1.1.2, D.1.2

[67] Meisenheimer, K., Röser, H.-J., & Schlötelburg, M. 1996, A&A, 307, 61+   3.5.3

[68] Meisenheimer, K., Yates, M. G., & Röser, H.-J. 1997, A&A, 325, 57   1.1.2, 1.1.4, 3.3.2, 4.1, 4.3, 10, 4.4, 5.1, 6.1, D.1.2

[69] Mirabel, I. F., Rodriguez, L. F., Cordier, B., Paul, J., & Lebrun, F. 1992, Nat, 358, 215   1

[70] Myers, S. T. & Spangler, S. R. 1985, ApJ, 291, 52   4.3.1

[71] Neumann, M. 1995, PhD thesis, Universität Heidelberg   1.1.5, 3.5.3, 4.3.1, 1

[72] Neumann, M., Meisenheimer, K., Roeser, H. ., & Fink, H. H. 1997, A&A, 318, 383   1.1.5

[73] Neumann, M., Meisenheimer, K., & Röser, H.-J. 1997, A&A, 326, 69   5, 3.2, 3, 3.13, 3.5.2, 5.2.1, 5.2.1

[74] NICMOS team. 2001, NICMOS calibration anomalies, `http://www.stsci.edu/NICMOS/si.pl?nav=calibration:anomalies&sel=objtype:anomaly` 2.2.4, 2.2.4

[75] —. 2001, NICMOS Synthetic Dark Request Form, `http://www.stsci.edu/cgi-bin/NICMOS/si.pl?nav=calibration:cal_tools&sel=id:261` 2.2.4

[76] Oke, J. B. & Schmidt, M. 1963, AJ, 68, 288+   2

[77] Owen, F. N., Hardee, P. E., & Cornwell, T. J. 1989, ApJ, 340, 698   3.3.1, 5.1

[78] Pacholczyk, A. G. 1970, Radio astrophysics. Nonthermal processes in galactic and extragalactic sources (Series of Books in Astronomy and Astrophysics, San Francisco: Freeman, 1970)   4.3, D.1, 13

[79] Padman, R., Lasenby, A. N., & Green, D. A. 1991, in Beams and jets in astrophysics, ed. P. A. Hughes, Cambridge Astrophysics Series No. 19 (Cambridge University Press), 484–564 1

[80] Perley, R. A. 1999, in ASP Conf. Ser. 180: Synthesis Imaging in Radio Astronomy II, 275+ 2.1.1, 2.2.1

[81] Perley, R. A., Röser, H.-J., & Meisenheimer, K. 1997, A&A, 328, 12 1.1.4

[82] Perlman, E. S., Biretta, J. A., Sparks, W. B., Macchetto, F. D., & Leahy, J. P. 2001, ApJ, 551, 206 3.5.2, 6, 6.1

[83] Perlman, E. S., Biretta, J. A., Zhou, F., Sparks, W. B., & Macchetto, F. D. 1999, AJ, 117, 2185 3.3.1, 6.3

[84] Perlman, E. S., Marshall, H. L., & Biretta, J. A. 2001, in Presented at meeting on ”Mass Outflow in Active Galactic Nuclei”, Washington DC, March 8-10 2001., available at `astro-ph/0106183`, 6183+ 1.1.5, 6, 5.5, 6.1

[85] Röser, H.-J., Conway, R. G., & Meisenheimer, K. 1996, A&A, 314, 414 3.2, 5.2.1, 5.2.1

[86] Röser, H.-J. & Meisenheimer, K. 1987, ApJ, 314, 70 1.1.4

[87] —. 1991, A&A, 252, 458 1.1.2, 1.1.4, 5, 2.1, 2.1.2, 2.6, 2.2.3, 2.2.3, 3, 3.5.2, 3.13, 4.4, 5.2, 5.2.1, 5.2.2, 5.3, 5.3

[88] Röser, H.-J. & Meisenheimer, K., eds. 1999, Lecture Notes in Physics, Vol. 530, The radio galaxy Messier 87 (Proceedings of a conference held at Schloß Ringberg: Springer-Verlag Berlin Heidelberg New York.) 1.1.4

[89] Röser, H. J., Meisenheimer, K., Neumann, M., Conway, R. G., & Perley, R. A. 2000, A&A, 360, 99 1.1.5, 5, 3.5.2, 3.5.3, 5.1, 5.2.2, 5.3, 5.3, 5.3

[90] Röser, H.-J., Meisenheimer, K., Neumann, M., et al. 1997, in Reviews of Modern Astronomy, Vol. 10, 253–262 1.2

[91] Rybicki, G. B. & Lightman, A. P. 1979, Radiative processes in astrophysics (New York, Wiley-Interscience, 1979. 393 p.) D.1, D.1.1, D.1.1

[92] Ryle, M., Smith, F. G., & Elsmore, B. 1950, MNRAS, 110, 508+ 1.1

[93] Sambruna, R. M., Urry, C. M., Tavecchio, F., et al. 2001, ApJ, 549, L161 5.2.2, 5.3, 5.3

[94] Schmidt, M. 1963, Nat, 197, 1040 1.1.1

[95] Schwartz, D. A., Marshall, H. L., Lovell, J. E. J., et al. 2000, ApJ, 540, L69 1.1.5

[96] Stockton, A. 1978, ApJ, 223, 747 2

[97] Strauss, M. A., Huchra, J. P., Davis, M., Yahil, A., Fisher, K. B., & Tonry, J. 1992, ApJS, 83, 29 D.2

[98] Tavecchio, F., Maraschi, L., Sambruna, R. M., & Urry, C. M. 2000, ApJ, 544, L23 1.1.5

[99] Thiele, M. & Camenzind, M. 1998, Astronomische Gesellschaft Meeting Abstracts, Abstracts of Contributed Talks and Posters presented at the Annual Scientific Meeting of the Astronomische Gesellschaft at Heidelberg, September 14–19, 1998, poster #P12, 14, 12+ 1.2

[100] Thommes, E. 1996, PhD thesis, Universität Heidelberg 3

[101] Thompson, R. I., Rieke, M., Schneider, G., Hines, D. C., & Corbin, M. R. 1998, ApJ, 492, L95 2.1.3

[102] Thomson, R. C., Mackay, C. D., & Wright, A. E. 1993, Nat, 365, 133 2.2.3

[103] Trauger, J. T., Vaughan, A. H., Evans, R. W., & Moody, D. C. 1995, in Calibrating Hubble Space Telescope Post Servicing Mission, ed. A. Koratkar & C. Leitherer (Baltimore: STScI), 379 3.1.3, B.3

[104] van der Laan, H. & Perola, G. C. 1969, A&A, 3, 468+ 5.1, 5.1, 5.1

[105] Voit, M. et al., eds. 1998, HST Data Handbook, Version 3.1 (Baltimore: STScI) 2.2.2, 2.2.2, A.1, A.3.1, B.2.1

[106] Wardle, J. F. C., Homan, D. C., Ojha, R., & Roberts, D. H. 1998, Nat, 395, 457 3

[107] Whitmore, B. & Wiggs, M. S. 1995, Charge Transfer Traps in the WFPC2, Instrument Science Report WFPC2 95-03 (Baltimore: STScI) 2.2.2, A.3.1

[108] Wilson, A. S., Young, A. J., & Shopbell, P. L. 2000, ApJ, 544, L27 1.1.5

[109] —. 2001, ApJ, 547, 740 1.1.5

List of Figures

List of Tables

Acknowledgements