UNIVERSITY OF SOUTHAMPTON

Modelling Radiative Transfer in High-Energy Astrophysical Plasmas

Ross Stefan Collins

Submitted for the degree of Doctor of Philosophy

SCHOOL OF ENGINEERING SCIENCES
FACULTY OF ENGINEERING

May 11, 2005

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING

SCHOOL OF ENGINEERING SCIENCES

Doctor of Philosophy

Modelling Radiative Transfer in
High-Energy Astrophysical Plasmas

by Ross Stefan Collins

In this thesis I describe the development of a three-dimensional radiative transfer model
that is capable of explaining observations of time-dependent broadband spectra of high-
energy astrophysical objects, such as X-ray binaries and microquasars, across the electro-
magnetic spectrum from radio waves to gamma-rays. Physical processes included in the
model are synchrotron radiation, Compton scattering, bremsstrahlung radiation, Coulomb
scattering and the time-dependent evolution of the emitting electron population's trans-
relativistic energy distribution. The model can recreate the geometry of any emitting
region, making it simple to apply to any astrophysical object. To that effect an initial at-
tempt has been made at applying the code to model the accretion disc coronae of Galactic
X-ray binaries for the purpose of explaining the cause of spectral state transitions. This
thesis concludes with a section detailing a preliminary study into the potential application
of the model to another problem in high-energy astrophysics – that of the time-dependent
emission from relativistic jets.

Contents

Preface

In this thesis I describe the design and development of a new radiative transfer model of the time-dependent emission from high-energy astrophysical plasmas. I begin in Chapter 1, by providing a background to the study of high-energy astrophysics and some of the important astrophysical objects to which the model may be applied. This is followed by a description of the fundamental physical interaction processes between photon and electron populations within a plasma in Chapter 2. In this chapter particular care has been taken to ensure the appropriate formulae are derived such that they are applicable over a trans-relativistic range of electron energies and the complete observable electromagnetic spectrum. Also presented in this chapter is the partial differential equation known as the kinetic equation, which is used to model the evolution of the electron energy distribution due to energy exchanges both within itself and with the interacting photons. Formulae are derived in this chapter for each of the energy exchange terms in the kinetic equation, which model this energy exchange for each interaction process included in this code, consisting of synchrotron radiation, Compton scattering, bremsstrahlung radiation and Coulomb scattering.

Previous models of radiative transfer have taken a Monte Carlo approach whereby *individual* photons and electrons and their interactions are considered. In Chapter 3, I present a computational model for the three-dimensional radiative transfer of photon *distributions*. Also detailed in this chapter is the numerical method used to calculate these photon distributions and the evolution of the electron distributions from the formulae presented in the previous chapter. This chapter ends with a demonstration of the capabilities of the completed model. In Chapter 4 the high-energy astrophysics radiative transfer (hereafter HEART) code is applied to modelling the X-ray and Gamma-ray spectrum emitted by accretion disc coronae. Finally, a potential application of the HEART code to the study of the time variability of relativistic jets is explained in Chapter 5.

Chapter 1

Background

1.1 High-energy astrophysics

The study of astrophysics essentially concerns the collection of photons from extrasolar sources, and measuring how the collected number of photons varies with energy, time and spatial position (as projected on the sky). The number of photons collected is often referred to as the source's *intensity* (alternatively *flux*, or *luminosity*, as will be explained in section 2.2), and the resulting intensity versus energy plot is called a *spectrum*, the intensity versus time plot is commonly known as a *light curve*, and an intensity versus spatial position plot is an *image*.

All that we know about extrasolar objects comes from the photons that they emit, and as such it is important to understand how these photons were created, and transferred to our astronomical observatory. Many astrophysical objects of interest (and most high-energy objects) appear to us as unresolved points of light, called *discrete sources*, and so our only source of knowledge about such objects comes from their spectra, and light curves. Even without imaging information we can determine the geometrical and compositional properties of an astrophysical object through a combination of spectrum and light curve observations by comparing this data with that produced by physical models.

The radiation spectrum of an astrophysical source is determined by the distribution of the emitting object's electrons (or also in some cases positrons). The electron distribution in turn traces out the distribution of matter within the astrophysical object of which its geometry, spatial density variation, temperature, and chemical composition, are all fundamental constituents. With light curve observations we can place further constraints upon the geometry on the basis of information travel time arguments. Quite simply, no astrophysical object should show variations on timescales shorter than the time it takes for light to travel across the source. Furthermore, observations of the time variability of the entire spectrum of an astrophysical source provides insights into physical models.

Many astrophysical objects are comprised of essentially very hot, tenuous gases, in a plasma state. As hydrogen is by far the most common element in the universe, most astrophysical plasmas are almost pure hydrogen, and hence are just a collection of 'free' electrons and protons, and, if not fully ionised, hydrogen atoms, with the relative number of any ions/atoms of heavier elements often being negligible. One important exception in this study is that of relativistic jets, which could possibly be formed of an electron-positron plasma.

On the assumption that the radiation (photons) and matter (electrons) in a given astrophysical object are in thermal equilibrium, a *blackbody* radiation spectrum will be emitted with the peak photon energy in the distribution corresponding to the electron temperature (see section 2.6). For most photons to be emitted in the visual waveband of light the plasma would need to have a temperature of approximately 10^3 K, or roughly that of stellar atmospheres. For a significant number of photons to be emitted at higher energies, such as the X-ray waveband, large, and relatively dense (on astrophysical scales) plasmas with temperatures of the order of 10^6 K would be required. Such temperatures can be achieved in the ultra-dense cores of stars, but tenuous astrophysical plasmas of such temperatures were not believed to exist until the first X-ray observations of the sky were undertaken. These observations found numerous discrete X-ray sources, and spawned the study of *high-energy astrophysics*.

High-energy astrophysical objects, are those that emit a significant amount of their radiation at energies greater than that of optical light. As is explained in Chapter 2, we
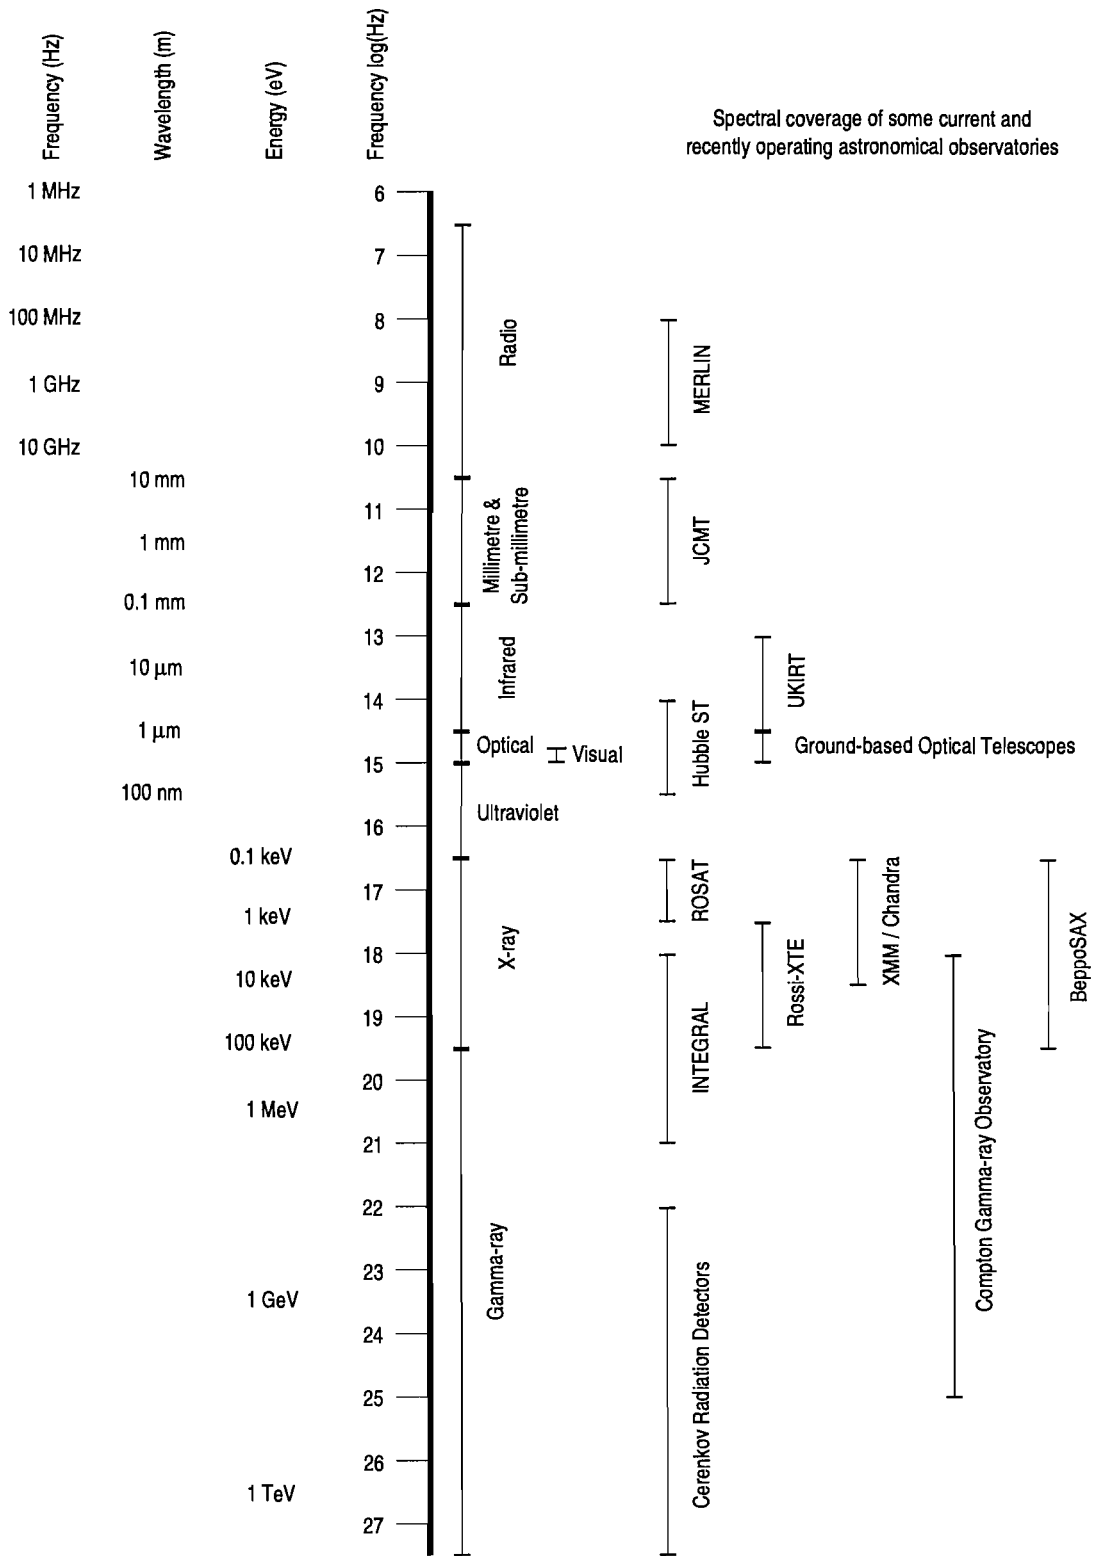
Figure 1.1: A nomogram showing the regions of the electromagnetic spectrum accessible to past and present astronomical observatories.

now know of several physical processes (both thermal and non-thermal) that can produce large quantities of high-energy photons without requiring such extremely high temperature plasmas. The nomogram in Figure 1.1 provides a useful reference to the various units used in this thesis to describe the energy ranges of photons that may be observed by current detectors.

1.2   X-ray binaries

The early X-ray observations of the sky revealed the locations of the strongest X-ray sources to be concentrated in the plane of our Galaxy. Such a distribution indicates that these objects are probably *Galactic* in origin rather than *extragalactic*. X-ray instrumentation at the time had poor energy resolution, so the primary source of information about these discrete, Galactic, X-ray sources came from their light curves. A couple of these sources were observed to pulsate with very precise, short periods – behaviour that can only be produced by neutron stars (a compact stellar object). Observations of one such object (Centaurus X-3) revealed that the pulsation period varied slightly in time, in a manner consistent with the Doppler shifting effect of a two day circular orbit. This interpretation was confirmed by observations of an eclipse in the light curve every two days – the result of the companion star passing in front of our line of sight to the X-ray source (the neutron star). So just with light curve observations it could be seen that some of these Galactic X-ray sources were binary star systems.

Further observations of eclipses in light curves revealed that most of the known X-ray sources that are distributed in the disc of our Galaxy were close binaries, consisting of a compact stellar object (a neutron star or black hole) and an ordinary, *main-sequence*, companion star. Observations of the optical spectra of these Galactic disc X-ray binaries revealed absorption line profiles of the type created by the atmospheres of very massive main-sequence stars. These *O* and *B* type stars have masses several tens that of our own sun (and that of their neutron star companions), and have relatively short lifetimes, so are typically found in the discs of spiral galaxies where they were first formed. Such X-ray binary star systems became known as high-mass X-ray binaries (HMXBs).

So what created the X-rays? Massive stars such as those found in HMXBs have very strong stellar winds that constantly eject significant amounts of stellar material from the star's atmosphere. The closely orbiting neutron star will sweep through this wind and gravitationally collect a lot of this material. Some of the material will directly impact upon its surface; but most of this plasma will have sufficient angular momentum to enter orbit around the neutron star, forming an *accretion disc*. Viscosity in the disc creates heat as gravity causes the plasma to spiral in towards the neutron star. The gravitational field in the vicinity of the neutron star is so strong that the innermost radius of the accretion disc becomes hot enough to emit blackbody radiation at X-ray energies. With newer X-ray observatories it became possible to measure the X-ray spectrum of HMXBs, which true to the theory revealed a blackbody 'hump' – the signature of the accretion disc.

The remaining Galactic X-ray sources, located in the Galactic bulge at the centre of the Galaxy, showed no signs of 'binarity' in their light curves. Observations of the optical spectra of these sources revealed emission lines in contrast to the absorption lines of ordinary star spectra, and a significantly 'bluer' continuum. Such optical spectra though are very similar to those of *cataclysmic variables* (CVs), which at the time were known
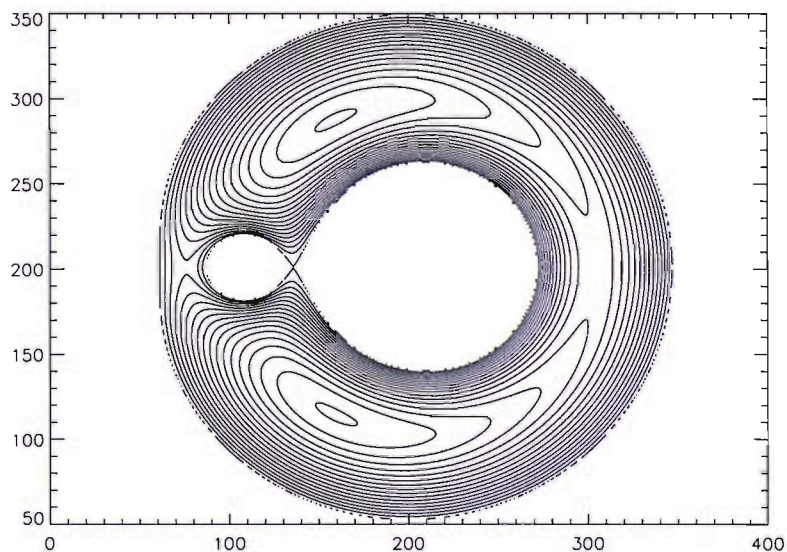

Figure 1.2: A force potential contour plot around two orbiting gravitational masses illustrating the Roche lobe and Lagrangian point, L1, defined by the crossing of the innermost contour line.

to be interacting binary systems. In a CV system stellar material is accreted by a white dwarf star from an ordinary star companion in a process called *Roche lobe overflow*. In the rotating frame of two orbiting masses there exists a point between the two masses, called the *inner Lagrangian point*, L1, where the gravitational, centrifugal, and Coriolis forces are all in balance (see Figure 1.2). If the ordinary companion star expands out to fill its Roche lobe, material at the L1 point will flow in towards the orbiting white dwarf star, and will form an accretion disc, due to its angular momentum. This accretion disc can be much brighter (and hotter) than the *donor* star's atmosphere, so it will dominate the optical spectrum of the source. The similarity between the optical spectra of CVs and the remaining Galactic X-ray sources led to the belief that perhaps these systems also contained accretion discs created by Roche lobe filling, dim, low-mass stars. Their strong X-ray emission could then be explained by the *accretor* being a neutron star or black hole, like with HMXBs, rather than the less dense white dwarfs found in CVs (see Figures 1.3 & 1.4). Further support for this theory comes from the knowledge that main-sequence stars found in the Galactic bulge are generally very old, and must therefore be low-mass stars.

The problem of a lack of evidence of binarity in the light curves observed in these Galactic bulge sources remained. This would be expected from a binary star system that was at such a great inclination that no orbital modulation effects, such as eclipses, would be observed in their light curves. However, it is unlikely that all of the sources are at such great inclinations. The solution to this problem lies in the geometry of the system. Firstly, the low-mass donor star is smaller in radius than the donor star in high-mass systems, and so would eclipse the X-ray source over a smaller range of inclinations. Secondly, X-rays are emitted from the inner radius of the accretion disc, so that at shallow inclinations the outer parts of the accretion disc block our view of the X-rays, and so no X-ray source is observed. The combination of these two effects results in the range of inclinations where an eclipsing X-ray source would be observed being much smaller than that for high-mass systems (see Figure 8.3 of Charles & Seward 1995), and so very few eclipsing sources are expected to be observed. Furthermore, with only a small portion of the donor star blocking the X-ray source, the eclipse duration would be very short, and hence difficult to detect by insensitive instruments that require long time integrations. With long duration
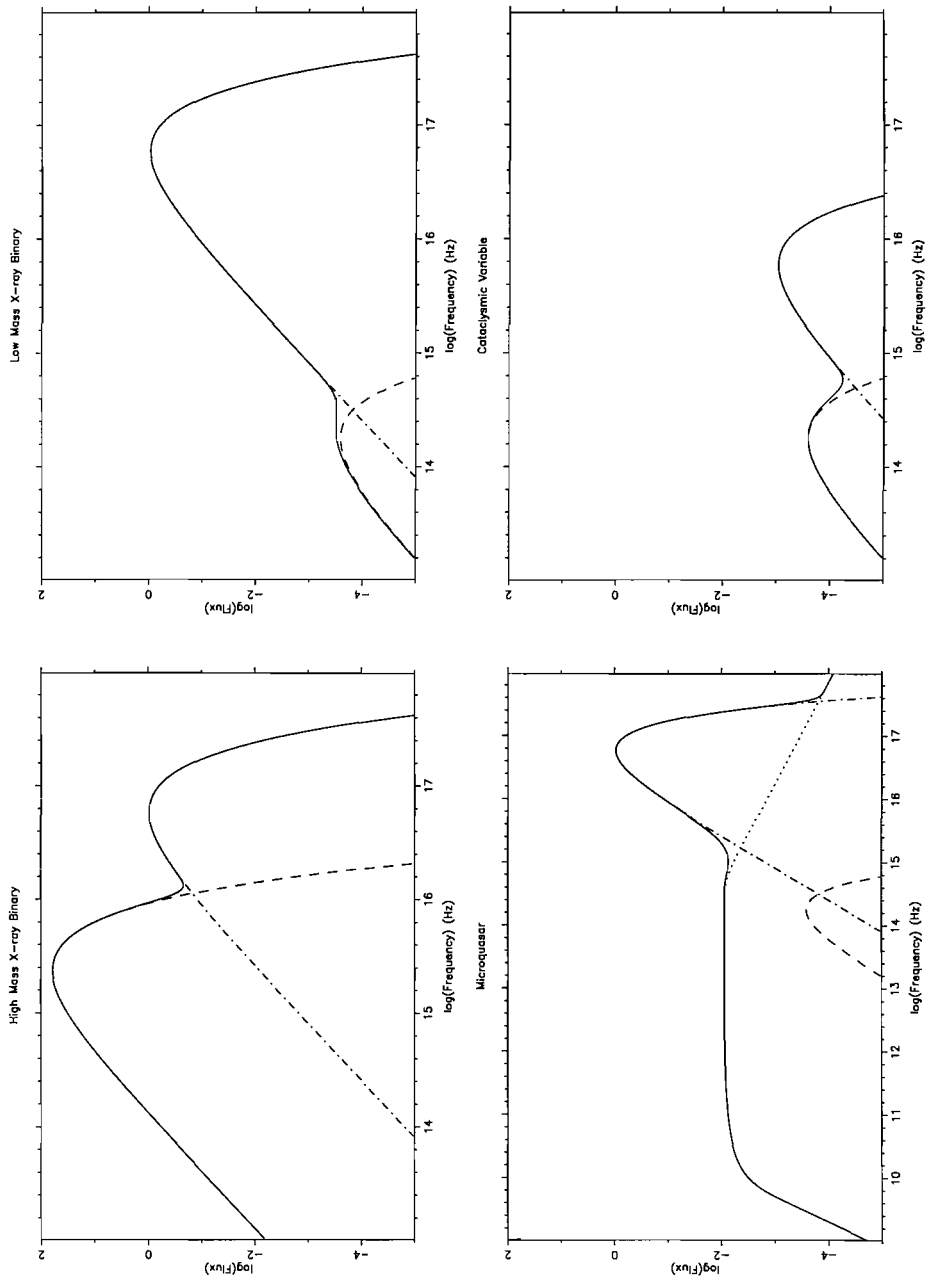
Figure 1.3: A schematic representation of canonical broadband spectra of each class of X-ray binary. The blackbody spectrum of the donor star (dashed line) in LMXB systems is obscured by the bright accretion disc (dot-dash line), and is only just visible at infrared wavelengths. However, the massive donor stars in HMXB systems are comparable in brightness to their accretion discs. Microquasar spectra reveal the broad flat spectrum of a relativistic jet (dotted line), that extends from radio wavelengths to the infrared.

Figure 1.4: The theoretical model of (a - left panel) HMXBs and (b - right panel) LMXBs. A compact object accretes a disc of material from its companion star, through the donor's stellar wind in HMXB systems, and by Roche lobe overflow in the LMXB systems. Images created using the BinSim software of Hynes (2001).

observations of these low-mass systems with more modern instrumentation eventually these short eclipses were discovered, confirming the binary nature of Galactic bulge sources, which then became known as low-mass X-ray binaries (LMXBs). Recently, spectroscopic observations of LMXBs in the infrared have identified many of the systems' donor stars, all of which show the signatures of typical low-mass main-sequence star atmospheres. Table 1.1 lists some of the observed properties from a few of the most important XRB systems.

1.2.1 Microquasars

As well as being strong sources of X-rays, many XRBs were also observed to be strong radio-sources, and surprisingly one XRB, SS 433, revealed a *resolved* image at radio-wavelengths, showing extended emission, despite being discrete at other wavelengths (Spencer 1979). More observations of extended radio emission around XRBs were made with the use of radio interferometers, which combine observations from radio telescopes around the globe to obtain radio images with far superior resolution over other wavelengths. Images of these *radio-loud* XRBs show elongated radio emission, such as that shown in Figure 1.5 of 1E1740.7-2942. This elongated emission, only visible at radio wavelengths, is remarkably similar to radio images of the radio-loud variants of quasars and other active galactic nuclei (AGN). In these extragalactic sources the radio emission

Table 1.1: Properties of some important X-ray binaries.

Identifier	Type	Accretor	Donor Star	Distance (kpc)
Scorpius X-1[1]	LMXB	1.4 $M_\odot$ n.s.?	< 0.9 $M_\odot$ M IV ?	2.8 ± 0.3
Centaurus X-3[2]	HMXB	1.2 $M_\odot$ ± 0.2 n.s.	20 $M_\odot$ O6-7 II-III	~ 8
Cygnus X-1[3]	HMXB	10 ± 5 $M_\odot$ b.h.	~ 30 $M_\odot$ O9.7	2 ± 0.3
GRS 1915+105[4]	LMXB	14 ± 4 $M_\odot$ b.h.	1.2 ± 2 $M_\odot$ K-M III	12 ± 1

Identifier	Radio-Loud	Resolved Jet	Period (d)	Inclination
Scorpius X-1	Yes	Yes	0.787	44° ± 6°
Centaurus X-3	No	No	2.087	70° ± 2°
Cygnus X-1	Yes	Yes	5.6	40° ± 15°
GRS 1915+105	Yes	Yes	33.5	70° ± 2°

[1]LMXB prototype and brightest persistent X-ray source

[2]HMXB prototype

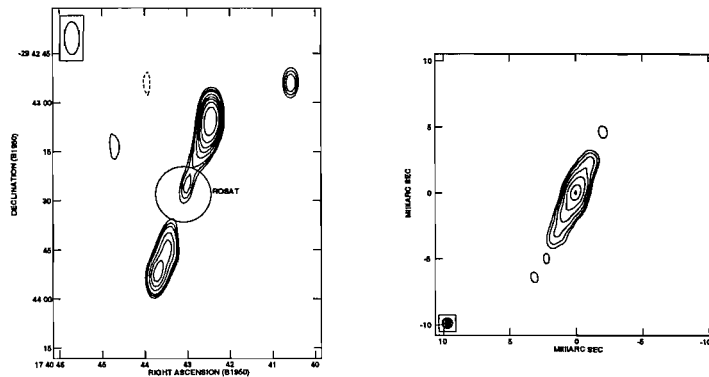[3]Black hole binary prototype

[4]Microquasar prototype

Figure 1.5: (a - left panel) The extended radio emission around the X-ray binary 1E1740.7-2942, now known to be a steady, compact, relativistic jet. Taken from Mirabel & Rodríguez (1999). (b - right panel) The steady, compact jet of GRS 1915+105. Taken from Dhawan *et al.* (2000).
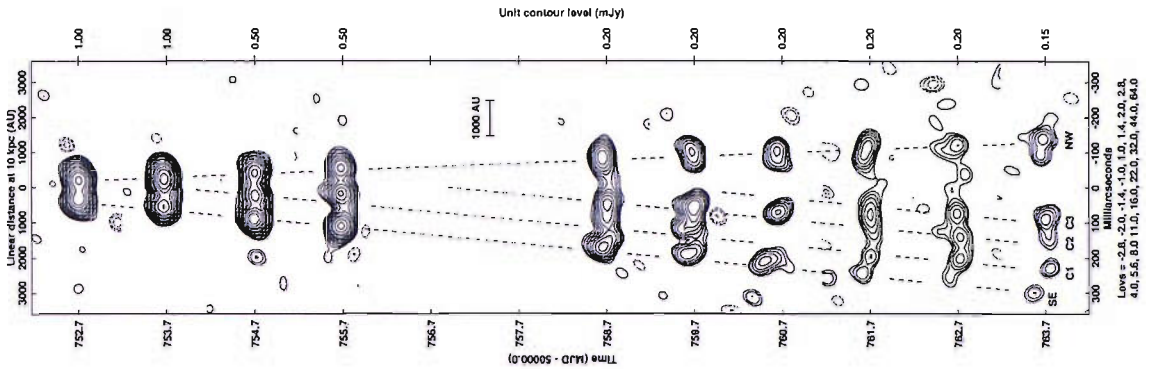
Figure 1.6: MERLIN radio observations of the superluminal ejecta from microquasar GRS 1915+105. Taken from Fender (2003).

is attributed to synchrotron radiation from relativistic jets.

In 1994, one relatively newly discovered LMXB, GRS 1915+105, was observed by Mirabel & Rodríguez (1994) using the Very Large Array (VLA), a 27 antenna array in New Mexico. The resulting series of radio images of this source (similar to the observations by Fender *et al.* (1999) shown in Figure 1.6), revealed it to be ejecting blobs of plasma in two opposite directions at an apparently *superluminal* velocity! Such behaviour had only previously been observed in the jets produced by AGN. This apparent superluminal velocity was explained by Rees (1966) as an illusion created by a relativistic jet that is inclined towards the observer. In actual fact the plasma blobs observed in the GRS 1915+105 jet were travelling at 0.92 c. The similarity between these superluminal jet sources and the highly luminous AGN known as quasars, lead to these XRBs being dubbed *microquasars*.

Compact, relativistic jets consist of a highly energetic electron (and possibly also positron) plasma emitting synchrotron radiation (see section 2.7), mainly at radio wavelengths, with a signature 'flat' spectrum (as is explained in section 5.2). Observations of the radio-loud XRB spectra have confirmed that this signature jet spectrum is present in all of the resolved compact jet sources (the microquasars), and many that still remain unresolved. Furthermore, we have so far observed superluminal ejecta from five microquasars (Fender 2003). Out of the few hundred XRBs currently known in our galaxy, about 10% are radio-loud, including both the high-mass and low-mass systems (Hjellming & Han 1995).
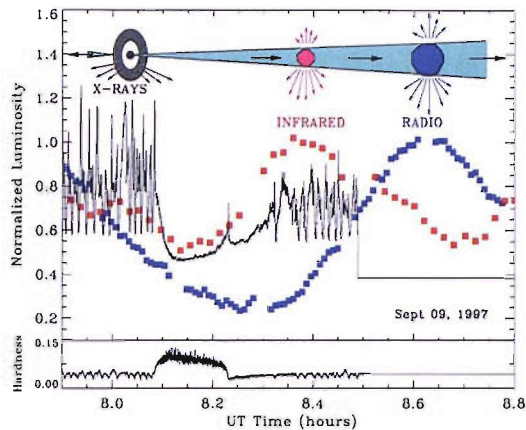
Figure 1.7: Time variability in the emission from relativistic jets, as explained by an expanding plasma blob emitting synchrotron radiation. Taken from Mirabel & Rodríguez (1999).

Microquasars show time variability in their emission at all observable frequencies. The variable X-ray emission is believed to be caused by instabilities in the accretion disc, as the disc is the strongest source of X-rays. Occasionally a dip in the X-ray emission will be observed followed by flaring of the infrared and subsequently the radio emission (see Figure 1.7). Such behaviour has been interpreted as the ejection of the inner region of the accretion disc into the relativistic jet stream. From the gradual rise of the flares it is evident that this ejection event doesn't occur instantaneously. Throughout the rise time of the flare the compact, innermost region of the conical jet becomes filled with an additional supply of fresh energetic electrons causing the jet's signature flat spectrum to rise in flux. This brightening of the compact jet (imaged in Figure 1.5) creates the infrared flare and shortly thereafter flares at progressively lower frequencies within the frequency range that reveals a flat spectrum in quiescence. When the supply of newly ejected material diminishes the higher frequency emission declines, whilst the ejecta propagate beyond the inner region of the jet. The expansion of this plasma results in flares at progressively lower radio frequencies, with the ejecta becoming visible in radio images as blobs of plasma moving at relativistic velocities (such as in Figure 1.6).

The relativistic nature of microquasar jets allows their emission to encompass such a broad range of the electromagnetic spectrum. Such high bulk velocities allow the freshly ejected material to span a very large range of radii (many orders of magnitude) during the

ejection event's short time period. When the ejection of material ceases this great velocity causes the range of radii to then rapidly decrease. The ejecta still propagate outwards at the same velocity, only now with their emission concentrated in a smaller frequency range. Many models of the emission from the ejecta at this stage of the ejection event have been developed (e.g. van der Laan 1966; Hjellming & Johnston 1988). In Chapter 5 we develop the first time-dependent model of the first stage of the ejection event where the ejecta expand over radii encompassing five orders of magnitude forming the 'compact' jet.

1.2.2 Black hole binaries

The relativistic jets from AGN are powered by supermassive black holes, leading to the suggestion that the compact object in microquasars is a black hole. Although evidence for the existence of black holes – objects that have an *event horizon* instead of a physical surface – cannot be conclusively obtained from studies of electromagnetic emission (Abramowicz *et al.* 2002), there is now very strong evidence for the existence of very compact objects with masses much greater than the Rhoades-Ruffini (1974) mass limit for neutron stars ($\sim$ 3 solar masses). For example Greiner *et al.* (2001) have presented strong evidence in favour of the existence of a $\sim$ 14 solar mass compact object in the GRS 1915+105 system. As there is not much support for the existence of theorised *solid* compact objects denser than neutron stars, any compact object with a mass greater than the mass limit for a neutron star is assumed to be a black hole.

The spectra of candidate black hole binary systems have, over time, been seen to vary between several 'states'. The two most significant states are that of the *high-soft state*, and the *low-hard state* (see Figure 1.8). The spectrum of the high-soft state is dominated by the accretion disc's blackbody in the *soft* (low-energy) X-rays, but also reveals a power-law tail extending into gamma-ray energies. This high-energy tail is believed to be created by inverse Compton scattering from a *corona* of hot electrons surrounding the accretion disc. In the low-hard state, the blackbody emission is cooler and dimmer, and is dwarfed by the *hard* (high-energy) X-ray emission that is much brighter than in the high-soft state. This hard X-ray emission in the low-hard state is a power-law in shape with a cut-off around 100 keV, but its source is unclear.
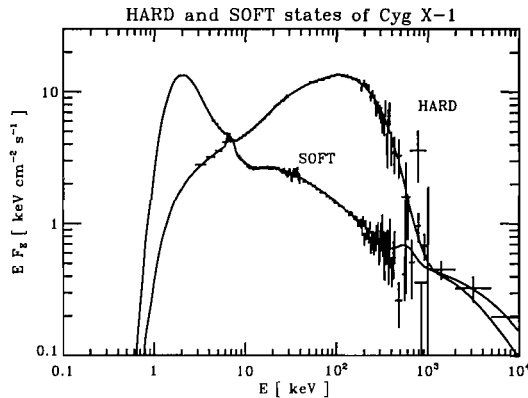
Figure 1.8: Spectral states of the black hole binary system Cygnus X-1. Taken from
Poutanen & Coppi (1998).

Whatever the source of the hard X-ray emission, it is clear that when a black hole
binary system changes from the high-soft state to the low-hard state the temperature
of the blackbody emission decreases. Accretion disc models always produce the hottest
material towards the innermost radius of the disc where the gravitational potential is
greatest, and hence this decrease in blackbody temperature suggests that the inner radius
of the accretion disc moves out to greater radii when the spectral state shifts to the low-
hard state. This in turn suggests the material from the inner region of the accretion
disc has disappeared, but it is not clear where to. It is known that many XRBs show
radio emission in the low-hard state when they do not in the high-soft state, indicating
that a jet is only present when the system is in the low-hard state (see e.g. Fender
2001). This provides a possible explanation for the removal of the inner disc material,
and can potentially explain the strengthened hard X-ray emission in this state. However,
it is also possible that the accretion disc corona is heated when the innermost material is
removed from the accretion disc, which would also provide an explanation for the observed
spectrum.

Markoff *et al.* (2001, 2003) have discovered that the optically thin synchrotron emission
tail from relativistic jet spectra provides a good fit to the observed hard X-ray emission
of two XRBs in the low-hard state. Furthermore, this same jet spectrum fits the observed
broad band spectrum of the XRBs in the low-hard state from radio up to optical fre-
quencies (see Figure 1.9). Such a model provides a good reason for correlations in the
time variability of the radio and the X-ray emission in this spectral state. Poutanen &
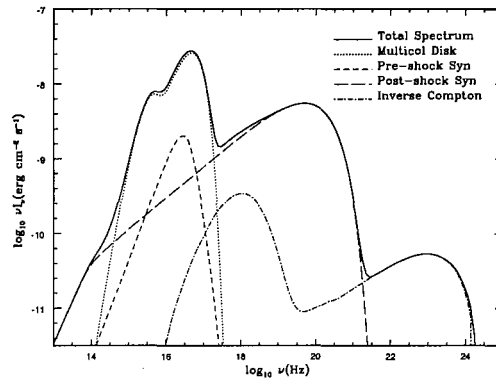
Figure 1.9: A model of the low-hard spectral state of the black hole binary system GX339-4 where the hard X-ray emission is modelled by the synchrotron emission from the relativistic jet. Taken from Markoff *et al.* (2003).

Zdziarski (2002) counter this interpretation claiming the hard X-ray emission is fit better with a model of blackbody accretion disc emission Comptonised by thermal electrons in the corona. One important difference between the two models is that the corona model explains the gamma-ray emission as a continuation of the hard X-ray power-law tail, whereas the jet model suggests the gamma-ray emission is a separate feature to the hard X-ray emission caused by inverse Compton scattering of the jet's synchrotron emission. Future gamma-ray observatories, such as INTEGRAL, may finally end the debate by providing higher resolution gamma-ray spectra.

Chapter 2

Emission and scattering processes

2.1   Introduction

The work in this thesis is primarily concerned with the problem of determining the energy distribution, or the *spectrum*, of the photons emitted by high-energy astrophysical plasmas. In this chapter we describe the fundamental physical processes that lead to the emission of photons from such a plasma. By modelling the photon distributions that we observe we can gain an understanding of the geometry and other properties of the plasmas in which they were created.

Astrophysical plasmas are, in general, ionised hydrogen gas, as hydrogen is by far the most abundant element in the universe. High-energy plasmas can be considered as simply a collection of individual electrons and protons (e–p plasmas) or in some cases electrons and positrons (e$^-$–e$^+$ plasmas, or pair plasmas). Although only the light, leptonic, particles (i.e. electrons and positrons) emit photons the proton population can influence the emission process. However, the work in this thesis is limited to covering just the electron-electron and electron-photon interaction processes; electron-proton, electron-positron, and positron-positron interactions are all ignored.

In the first half of this chapter I lay down the basic concepts and generic quantities

used to model the physics of emission processes. The second half then presents analytical equations and detailed descriptions for each individual emission or scattering mechanism that is considered in this work.

2.2   Photon distributions

The photon distribution is typically defined using a quantity known as the *specific intensity*, I_ν, which is the amount of energy, dE, in a frequency band, dν, that passes through an element of area, dA, in the time interval, dt, within a solid angle, dΩ,

$$dE = I_\nu \, dt \, dA \, d\nu \, d\Omega \tag{2.1}$$

(see Figure 2.1). The subscript, ν, denotes that the quantity, I_ν, is measured *per frequency band*, as opposed to the bolometric intensity, I, which measures the total intensity over all frequencies.

The concept of a specific intensity is represented graphically in Figure 2.1. It is a useful quantity for calculating radiative transfer internally through a source as it is independent of the emitting surface's properties. However, these properties become important in determining the emission that escapes, thus for astronomical observations the quantity
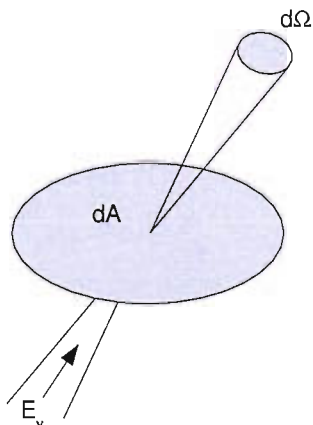
Figure 2.1: Graphical representation of the definition of the specific intensity as the amount of radiation, E_ν, that passes through an area element, dA, in the range of directions within solid angle element, dΩ, per unit of time.

measured is typically the *flux*, F_ν, which is related to the specific intensity by,

$$F_\nu = \int I_\nu \, d\Omega. \tag{2.2}$$

Hence, flux is a measure of the total photon intensity over all of the directions in which the radiation can escape from a given area. The other commonly used radiation term in astronomy, *luminosity*, generally refers to the total bolometric flux emitted by an astronomical object integrated over its entire surface. However, for observation purposes, astronomical measures of the radiation energy are generally given *per area* as the amount of radiation we detect is always a fraction of that which reaches out to our distance. No astronomical source can confine its emission to an area smaller than a detector (see Figure 2.2).

Photon distributions in this thesis will always be given in terms of frequency. However, other measures of photon energy are commonly utilised in astronomy for the various different regimes of the electromagnetic spectrum (e.g. see Figure 1.1). Lower energy astronomy typically uses wavelength, λ ($= c/\nu$), whereas high-energy astronomy uses photon energy, $h\nu$, in units of eV.
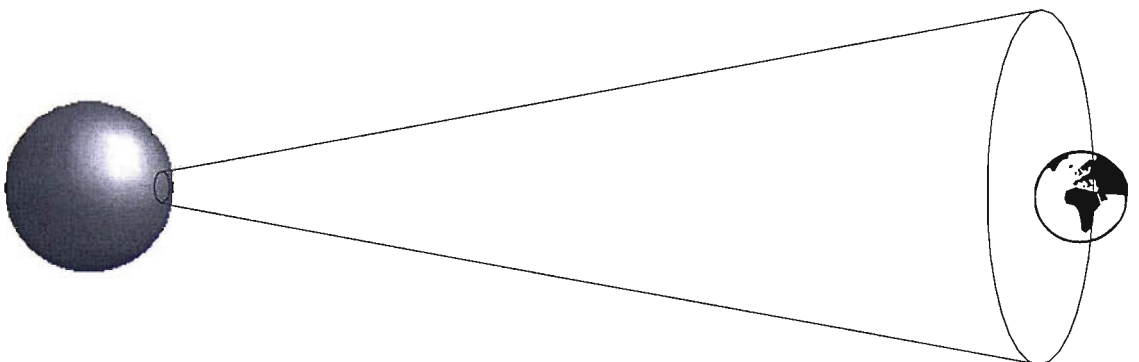
Figure 2.2: Only a small fraction of the amount of radiation emitted by an astronomical source can be detected, and therefore most measurements of this quantity are given *per area*. Note that this diagram is not to scale.

2.3   Electron distributions

As electromagnetic emission from astrophysical objects is created by the object's electron populations, it is important to know something about the physical processes that create the kinds of electron energy distributions that are commonly found in astrophysical sources.

Throughout this thesis the term, *electron distribution*, will be used to refer to the distribution of energies within a population of electrons. Dimensionless, relativistic units are used throughout as these are best suited for modelling the very large energy ranges found in the electron distributions of high-energy astrophysical objects. Therefore, we use the Lorentz factor, γ, to measure the total energy of an electron as a fraction of its rest mass energy, $m_e c^2$. An electron can't have less energy than its rest mass energy, and so γ must always be greater than unity. Electron velocities are given by, β, the velocity as a fraction of the speed of light, and for electron momentum we use $p = \sqrt{\gamma^2 - 1}$, where $\gamma = 1/\sqrt{1 - \beta^2}$. The electron distribution can be given in terms of any one of these three units, i.e. n_γ, n_β, or n_p. Typically we are interested in the number of electrons *per unit volume* with energies between γ and $\mathrm{d}\gamma$, and so express the distribution as n_γ. We will use, N_γ, to denote the total number of electrons (rather than electron number density) with energies between γ and $\mathrm{d}\gamma$.

2.3.1   Thermal distribution

Over time, and without external influence, the distribution of velocities within a population of electrons will settle down, after many collisions, to an equilibrium. This is the thermal equilibrium, with the peak velocity determined by the population's temperature. The distribution of energies is given by the Maxwell-Boltzmann equation, which in the relativistic form is

$$n_\gamma = \frac{n_e}{\theta_e K_2(1/\theta_e)} \gamma p \, e^{-\gamma/\theta_e} \tag{2.3}$$

(Synge 1957), where K_2 is the modified Bessel function of order 2, $\theta_e = kT_e/m_e c^2$ is the dimensionless, relativistic temperature of the electron distribution, and k is the Boltzmann constant. Integrating over all energies gives a total number density of n_e [m^{-3}].

2.3.2   Non-thermal distributions

To produce the high-energy radiation observed from high-energy astrophysical sources without requiring plasmas of very high temperature and density relies on the existence of electrons with ultra-relativistic energies emitting radiation through non-thermal processes. One way to produce electrons with such high energies is the Fermi acceleration process that occurs when a relativistic shock wave passes through a plasma (Fermi 1949 and see e.g. Bell 1978). Electrons that pass through the shock-front are accelerated to relativistic velocities, and multiple crossings will accelerate some electrons to even higher relativistic velocities.

The electron distribution produced by acceleration in relativistic shocks is simply given by a power-law distribution,

$$n_\gamma = n_0 \gamma^{-q}, \tag{2.4}$$

extending from γ_{min} to γ_{max}, with normalisation, $n_0 = n_e(q-1)$ $[\text{m}^{-3}]$ for $\gamma_{max} \gg \gamma_{min}$. As the minimum energy of an electron is its rest mass energy, $\gamma_{min} \geq 1$. Monte Carlo simulations of particle acceleration by shocks are used to calculate the highest electron energy in the distribution, γ_{max}, and the power-law index, q, which is found to be about 2 in the case of mildly-relativistic shocks and increasing to approximately 2.25 as the shock speed reaches the ultra-relativistic limit (e.g. Lemoine & Pelletier 2003, Achterberg *et al.* 2001, Bell 1978). The maximum energy obtained by an electron in this Fermi acceleration process depends upon the number of times it passes through the shock-front. Observations have shown that the relativistic jets of AGN can contain electrons with energies up to $\gamma_{max} = 10^7$ (Dermer & Atoyan 2002), but typically $\gamma_{max} \lesssim 10^4$ is assumed.

2.4   The principle of detailed balance

The interactions between photons and electrons within a plasma are governed by the principle of detailed balance (as discussed in e.g. Rybicki & Lightman 1979). When an electron emits a photon it will lose a certain amount of energy, $h\nu$, depending on the frequency, ν, of the emitted photon, where h is Planck's constant. Conversely, for every
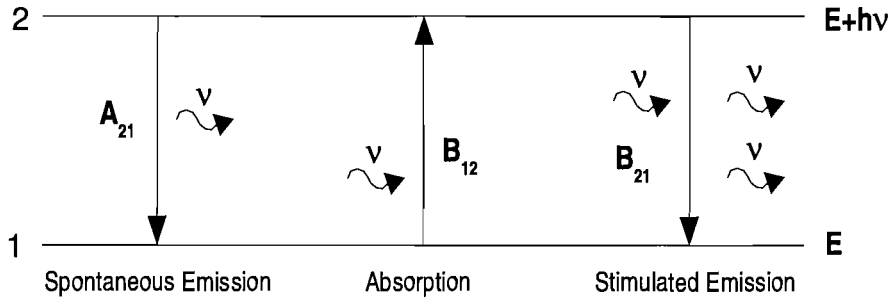
Figure 2.3: The principle of detailed balance. In this diagram, electrons can exist in two discrete energy levels, with electrons in level 2 having more energy than electrons in level 1. Electrons can move between these two levels by emitting or absorbing the energy difference of the gap between the levels in the form of a photon of frequency ν. These transitions can occur due to spontaneous emission, with a rate of A_{21}, absorption, at a rate of $B_{12}I_\nu$, or stimulated emission at a rate of $B_{21}I_\nu$.

emission process there exists an inverse process and so an electron may gain energy by absorbing a photon. Also induced emission can occur whereby a photon of frequency ν may stimulate an electron with energy $h\nu$ into emitting a second photon of the same frequency. The principle of detailed balance states that for a system in complete thermodynamic equilibrium any energy exchange process must be balanced by the rate of its inverse process.

This balance of energy exchange is illustrated in Figure 2.3 and defined by an equation of interaction rates:

$$n_1 B_{12} I_\nu = n_2 A_{21} + n_2 B_{21} I_\nu. \tag{2.5}$$

The Einstein co-efficients, A and B, define the rate of energy exchange per unit time per electron (and per specific intensity for the B co-efficient) between the two energy levels given in the subscript, and n gives the population number of the denoted energy level. Each co-efficient depends only upon the physical mechanism behind the emission process, and may be derived by determining the rate of interaction between the particles that are involved in the particular mechanism that leads to the emission of a photon. Hence,

$$\frac{dn_p}{dt} = n_e A_{21} = n_e \, n_{ip} \, c \, \sigma, \tag{2.6}$$

the production rate of photons (number n_p) is proportional to the number density of electrons, n_e, multiplied by the number density of the particles with which they are in-

teracting, n_{ip}, the speed of the interacting particle, c, and the cross-section of interaction, σ, for this particular mechanism. The interacting particles maybe electrons in the case of e^--e^- interactions (e.g. bremsstrahlung), positrons for e^--e^+ interactions (e.g. particle annihilation), or photons for $e^--\gamma$ interactions (e.g. Compton scattering).

An important quantity used in calculating emission rates is the single-particle *emissivity*, j_ν [W Hz^{-1}], which is simply the emission power per electron in the frequency band ν to $\nu+\mathrm{d}\nu$. Following equation 2.6, the emissivity is defined as

$$j_\nu = \frac{\mathrm{d}E}{\mathrm{d}t\mathrm{d}\nu} = h\nu\, n_{\mathrm{ip}}\, c\, \sigma_\nu, \tag{2.7}$$

where $\sigma_\nu = \mathrm{d}\sigma/\mathrm{d}\nu$. Integrating the emissivity over the electron distribution, n_γ, and assuming isotropy, leads us to the emission co-efficient, ϵ_ν [W m^{-3} Hz^{-1} st^{-1}],

$$\epsilon_\nu = \frac{\mathrm{d}E}{\mathrm{d}t\mathrm{d}V\mathrm{d}\Omega\mathrm{d}\nu} = \frac{1}{4\pi}\int j_\nu n_\gamma \mathrm{d}\gamma. \tag{2.8}$$

To determine whether a photon interacts with an electron it is useful to introduce the concept of a mean free-path, λ_{mfp} [m], the mean distance that a photon will travel before interacting, which depends upon the interaction cross-section as,

$$\lambda_{\mathrm{mfp}} = \frac{1}{n_{\mathrm{e}}\sigma}. \tag{2.9}$$

From this we may define an absorption co-efficient, κ_ν [m^{-1}], which expresses the number of interactions (i.e. absorptions) that we would expect to see per unit length,

$$\kappa_\nu = \int n_\gamma\, \sigma_\nu\, \mathrm{d}\gamma. \tag{2.10}$$

Integrating the absorption co-efficient over a given path-length, l, gives the number of interactions that occur while the photons travel through that length, a dimensionless quantity known as the optical depth,

$$\tau_\nu = \int \kappa_\nu\, \mathrm{d}l. \tag{2.11}$$

In a homogeneous plasma the optical depth can be thought of as the path-length divided by the mean free-path,

$$\tau = \frac{l}{\lambda_{\mathrm{mfp}}}. \tag{2.12}$$

For photon-electron interactions in the Thomson limit, where the energy of the photon is significantly less than that of the electron ($h\nu \ll \gamma m_e c^2$), the interaction cross-section is a constant known as the Thomson cross-section, σ_T. In this limit we may define a Thomson optical depth for a homogeneous plasma,

$$\tau_T = l\, n_e\, \sigma_T. \tag{2.13}$$

2.5   The radiative transfer equation

Radiative transfer is concerned with calculating the path taken by photons as they travel through a material. The nature of this path can significantly affect the photon energy and time distributions. As such, from observations of these photon distributions the matter distribution can be determined. From the principle of detailed balance we arrive at the radiative transfer equation, which describes the change in the specific intensity of light passing through a material with an emission coefficient, ϵ_ν, and absorption coefficient, k_ν, over a path-length, $\mathrm{d}l$,

$$\frac{\mathrm{d}I_\nu}{\mathrm{d}l} = -k_\nu I_\nu + \epsilon_\nu. \tag{2.14}$$

Over a distance, $\mathrm{d}l$, the specific intensity increases due to photon emission by an amount, $\epsilon_\nu\, \mathrm{d}l$, and decreases due to absorption by an amount proportional to its magnitude, $k_\nu\, \mathrm{d}l\, I_\nu$ or $\mathrm{d}\tau_\nu I_\nu$. The radiative transfer equation may be re-expressed in terms of the optical depth rather than the physical path-length, separating the terms that are properties of the matter from those of the radiation:

$$\frac{\mathrm{d}I_\nu}{\mathrm{d}\tau_\nu} = -I_\nu + S_\nu, \tag{2.15}$$

where S_ν, the source function, is the ratio of the emission coefficient to the absorption coefficient. The benefit of using the optical depth is made clear by solving the radiative transfer equation. Under the assumption that over the path-length, $\mathrm{d}l$, the material is homogeneous, we find the specific intensity as a function of optical depth through the absorbing medium to be

$$I_\nu(\tau_\nu) = S_\nu(1 - e^{-\tau_\nu}). \tag{2.16}$$

It is interesting to consider the emission in the case of optically thin materials, $\tau_\nu \ll 1$, and optically thick materials, $\tau_\nu \gg 1$. For emission from optically thick materials we find that the emission is given solely by the source function, with no dependence upon the depth of the material,

$$I_\nu(\tau_\nu) = S_\nu. \tag{2.17}$$

In the case of optically thin materials the emission strength is directly proportional to the optical depth, and hence the physical path-length, l, (in the homogeneous case),

$$I_\nu(\tau_\nu) = S_\nu \tau_\nu = \epsilon_\nu l. \tag{2.18}$$

Note that in this case the specific intensity is independent of the absorption coefficient, so no absorption occurs in this limit.

2.6 Thermal emission

An electron distribution that is in thermal equilibrium with itself (see section 2.3.1), and is sufficiently optically thick, will produce a photon distribution that is also in thermal equilibrium with the electron population, called a 'blackbody' spectrum. Despite it being physically impossible for such an equilibrium to exist – otherwise the entire universe would be comprised of isothermal material – sufficiently opaque material will emit radiation with a spectrum that approximates a blackbody. In such a material many interactions will occur along a photon's escape path-length, and hence this optically thick material will produce a photon distribution with a temperature equal to that of the electron distribution. The blackbody radiation spectral shape is given by the Planck formula,

$$I_\nu = \frac{2h\nu^3}{c^2} \frac{1}{\exp(h\nu/kT_{\mathrm{BB}}) - 1}, \tag{2.19}$$

where T_{BB} is the blackbody temperature, and h is the Planck constant. The strength of the radiation and the frequency at which it peaks increases with T_{BB}. The frequency at which the radiation strength peaks is given by Wien's law,

$$\nu_{\mathrm{peak}} = \frac{2.8k}{h} T_{\mathrm{BB}}. \tag{2.20}$$

Integrating equation 2.19 over all frequencies gives the relation between the strength of the emission and the temperature to be,

$$I = \frac{\sigma}{\pi} T_{\mathrm{BB}}^4,\tag{2.21}$$

where σ is the Stefan-Boltzmann constant. This emphasises how strongly the intensity of thermal radiation emitted by an object depends on its temperature.

2.7   Synchrotron radiation

Electrons with relativistic energies ($\beta \approx 1$) in the presence of a magnetic field will emit photons, due to the synchrotron radiation process named after the particle accelerator in which this effect was first discovered. The direction of motion of charged particles is perturbed by magnetic fields causing the particles to spiral around the lines of force. As a result of the particles moving with speeds close to light speed, the acceleration on the particles due to the magnetic field sends a wave along the particle's electric field lines, resulting in the emission of electromagnetic radiation. The frequency distribution of this radiation emissivity can be derived (see e.g. Rybicki & Lightman 1979, Longair 1994) for a single ultra-relativistic electron (with $\gamma \gg 1$) to give,

$$j_\nu = \frac{\sqrt{3}e^3 B \sin(\alpha)}{4\pi\epsilon_0 m_{\mathrm{e}}c} F_{\mathrm{S}}(x),\tag{2.22}$$

where B is the magnetic field strength, α is the pitch angle of the electron's motion with respect to the magnetic field force lines, e is the charge of an electron, ϵ_0 is the electric permittivity of free space, and m_{e} is the rest mass of an electron.

$$F_{\mathrm{S}}(x) = x \int_x^\infty K_{5/3}(z)\mathrm{d}z,\tag{2.23}$$

where $K_{5/3}(z)$ is the modified Bessel function of the second kind of order 5/3, and $x = \nu/\nu_{\mathrm{c}}$ is the frequency as a fraction of the critical frequency,

$$\nu_{\mathrm{c}} = \frac{3\gamma^2}{2\beta}\nu_{\mathrm{g}}\sin(\alpha),\tag{2.24}$$

where

$$\nu_{\mathrm{g}} = \frac{eB}{2\pi m_{\mathrm{e}}},\tag{2.25}$$
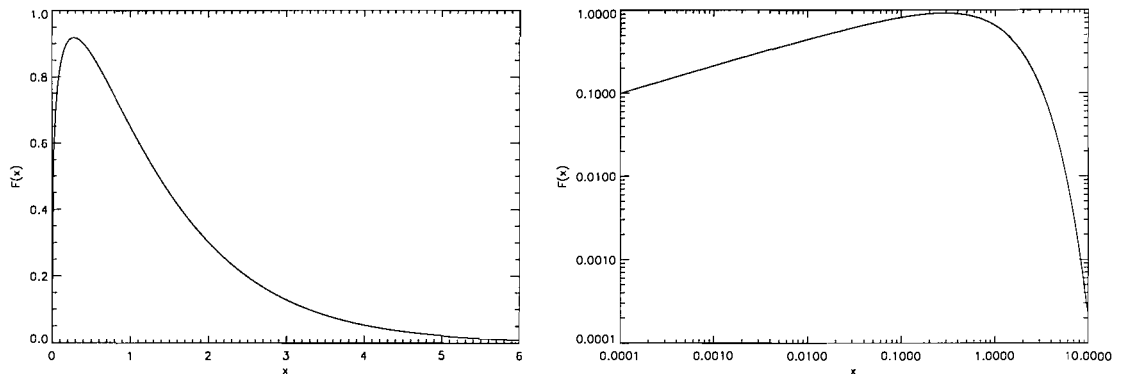
Figure 2.4: The synchrotron radiation emissivity of a single ultra-relativistic electron, in linear space (left panel), and in log space (right panel). Note that the most of the radiation is emitted at approximately the critical frequency.

is the gyro or Larmor frequency of an electron in a magnetic field of strength, B. Hence, in the ultra-relativistic limit ($\beta = 1$),

$$j_\nu = \frac{e^2 \nu}{\sqrt{3}\epsilon_0 c\gamma^2} \int_x^\infty K_{5/3}(z)\mathrm{d}z. \qquad (2.26)$$

The frequency distribution of a single electron's synchrotron emissivity is shown in Figure 2.4. The functional form of this emissivity at the frequency limits are,

$$F_S(x) = \begin{cases} \frac{4\pi}{\sqrt{3}\,\Gamma(1/3)}\left(\frac{x}{2}\right)^{\frac{1}{3}} & ; \quad x \ll 1, \\ \sqrt{\frac{\pi}{2}x}\,\mathrm{e}^{-x} & ; \quad x \gg 1. \end{cases} \qquad (2.27)$$

The analytical solution to equation 2.26 for the total emissivity of an average ultra-relativistic electron within an isotropic distribution of pitch angles was first derived by Crusius & Schlickeiser (1986), and given in the following form by Ghisellini *et al.* (1988):

$$j_\nu(\gamma) = \frac{\sqrt{3}e^2}{\epsilon_0 c}\nu_g x^2\left\{K_{4/3}(x)K_{1/3}(x) - \frac{3}{5}x\left[K_{4/3}(x)^2 - K_{1/3}(x)^2\right]\right\} \qquad (2.28)$$

where in this form

$$x = \frac{\nu}{3\gamma^2\nu_g}. \qquad (2.29)$$

This is a useful result as typical astrophysical plasmas emit only weakly polarised synchrotron radiation indicating minimal homogeneity to the magnetic field, and hence the assumption of an isotropic distribution of pitch angles is valid.

Petrosian (1981) has derived the same quantity for the case of mildly-relativistic electrons, which Ghisellini *et al.* (1988) have integrated over an isotropic pitch angle distribution to give

$$
\begin{aligned}
j_\nu(\gamma) &= \frac{\pi e^2 \nu_{\mathrm g}}{12\epsilon_0 c}\left\{1 + 2(\gamma^2 - 1)\left[1 + \gamma\ln\left(\frac{\sqrt{\gamma^2 - 1}}{1 + \gamma}\right)\right]\right\}^{-\frac{1}{2}} \\
&\quad \times \exp\left\{\frac{2\nu}{\nu_{\mathrm g}}\left[1 + \gamma\ln\left(\frac{\sqrt{\gamma^2 - 1}}{1 + \gamma}\right)\right]\right\}.
\end{aligned}
\tag{2.30}
$$

This equation was derived in the limit, $\nu \gg \gamma^2 \nu_{\mathrm g}$, but is still reasonably accurate for $\nu \geq \nu_{\mathrm g}$.

From the single electron emissivity we can calculate the emission and absorption coefficients for a distribution of electrons, which form the synchrotron emitting material, using the definitions given by equations 2.8 and 2.10. The absorption cross-section for synchrotron radiation in the Thomson limit was derived by Le Roux (1961) giving

$$
\sigma_{\mathrm S}(\nu, \gamma) = \frac{1}{8\pi m_{\mathrm e}\nu^2}\frac{1}{\gamma p}\frac{\partial}{\partial\gamma}\left[\gamma p\, j_\nu(\gamma)\right].
\tag{2.31}
$$

Integrating equation 2.10 by parts with this cross-section gives a form of the synchrotron absorption coefficient where the electron distribution is differentiated rather than the emissivity function (see Ghisellini & Svensson 1991),

$$
k_\nu = -\frac{1}{8\pi m_{\mathrm e}\nu^2}\int_{\gamma_{\min}}^{\gamma_{\max}} \gamma p\, j_\nu(\gamma)\frac{\partial}{\partial\gamma}\left(\frac{n_\gamma}{\gamma p}\right)\mathrm d\gamma.
\tag{2.32}
$$

Synchrotron emission from high-energy astrophysical sources sources typically originates from electrons with energies distributed according to a power-law, equation 2.4, which arises from electron acceleration in shocks (see section 2.3.2 for a detailed discussion). Longair (1994) presents a solution to the emission and absorption coefficient integrals for such a distribution extending from $\gamma_{\min} = 1$ to $\gamma_{\max} = \infty$, using the ultra-relativistic form of the synchrotron emissivity formula (equation 2.28) to give

$$
\begin{aligned}
\epsilon_\nu &= \frac{\sqrt{3\pi}\,e^3}{64\pi^3\epsilon_0 m_{\mathrm e}c\,(q+1)}\left(\frac{3e}{2\pi\,m_{\mathrm e}}\right)^{(q-1)/2} n_0\, B^{(q+1)/2}\,\nu^{-(q-1)/2} \\
&\quad \times \frac{\Gamma\left[(3q + 19)/12\right]\Gamma\left[(3q - 1)/12\right]\Gamma\left[(q + 5)/4\right]}{\Gamma\left[(q + 7)/4\right]}
\end{aligned}
\tag{2.33}
$$

and

$$k_\nu = \frac{\sqrt{3\pi}\,e^3}{64\pi^2\epsilon_0 m_e^2 c}\left(\frac{3e}{2\pi\,m_e}\right)^{q/2} n_0\,B^{(q+2)/2}\,\nu^{-(q+4)/2}$$
$$\times\,\frac{\Gamma\left[(3q+22)/12\right]\Gamma\left[(3q+2)/12\right]\Gamma\left[(q+6)/4\right]}{\Gamma\left[(q+8)/4\right]} \qquad (2.34)$$

respectively. For these energy limits, $n_0 = (q-1)n_e$, where n_e is the electron number density integrated over all energies.

Using the definitions given in section 2.5 under the assumption of a homogeneous emitting region the source function and optical depth of such a plasma are,

$$S_\nu = 1.00\times10^{-36}\,c_s(q)\,B^{-0.5}\,\nu^{2.5}\;[\mathrm{Wm^{-2}Hz^{-1}sr^{-1}}] \qquad (2.35)$$

and

$$\tau_\nu = 9.075\times10^3\,c_\tau(q)\,n_e\,l\,B^{(q+2)/2}\,\nu^{-(q+4)/2} \qquad (2.36)$$

respectively, where l is the path-length to our line of sight. In the case of electron energisation by mildly-relativistic shocks, $q = 2$, $c_s(2) = 0.492$, $c_\tau(2) = 1.02\times10^{11}$, and in the ultra-relativistic case, $q = 2.25$, $c_s(2.25) = 0.422$, $c_\tau(2.25) = 2.798\times10^{12}$.
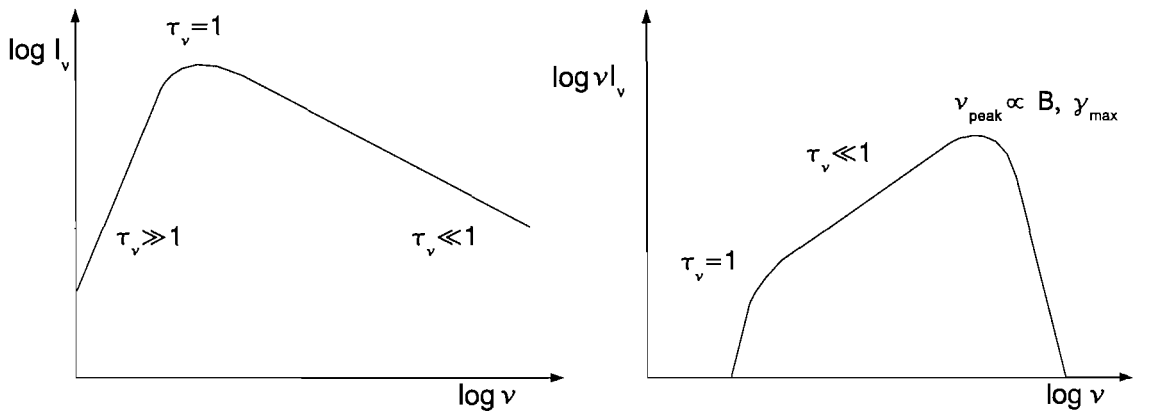
Figure 2.5: (left panel) The partially self-absorbed synchrotron emission spectrum for a power-law distribution of electrons. The emission peaks where the optical depth is unity, and decreases with a slope of 2.5 towards lower frequencies where the material is optically thick, and with a slope of $(1-q)/2$ towards higher frequencies to which the material is optically thin. (right panel) The spectral energy distribution (see text) peaks at a frequency dependent upon the highest energy electron in the emitting population and the strength of the magnetic field.

Hence, in the optically thick limit the synchrotron radiation spectrum from a power-law distribution of electrons will vary as $\nu^{2.5}$, and as $\nu^{(1-q)/2}$ in the optically thin limit. Figure 2.5 shows both the synchrotron radiation spectrum, and spectral energy distribution (SED). An SED plot, obtained by multiplying the specific intensity (or flux) by the frequency, illustrates the distribution of energy over the spectrum, rather than the distribution of photon number. Such plots are useful when comparing emission at very high frequencies (e.g. the gamma-ray regime) to that at lower frequencies, because at such high frequencies very few photons are produced, but with a net energy comparable to that emitted at lower frequencies. The frequency of the peak synchrotron radiation emission in an SED plot is determined by the critical frequency of the highest energy electron in the distribution.

Electrons lose energy when they emit synchrotron radiation, and over time this can have a significant effect upon the electron energy distribution. To determine this energy loss for a single electron we must integrate the emissivity over all emission frequencies,

$$-\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{1}{m_{\mathrm{e}}c^2} \int_0^\infty j_\nu(\gamma)\mathrm{d}\nu, \tag{2.37}$$

where the $m_{\mathrm{e}}c^2$ factor converts to relativistic units. Using equation 2.22 and integrating over an isotropic pitch-angle distribution (see e.g. Rybicki & Lightman 1979), we find

$$-\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{4\sigma_{\mathrm{T}}}{3m_{\mathrm{e}}c}U_B\gamma^2, \tag{2.38}$$

where σ_{T} is the Thomson cross-section for an electron, and

$$U_B = \frac{B^2}{2\mu_0}, \tag{2.39}$$

is the magnetic field energy density. Hence, the energy loss is proportional to γ^2, implying that the highest energy electrons lose energy more rapidly than the lowest energy electrons. This leads to the *synchrotron aging* phenomenon (Rees 1967), where radiation losses cause the electron distribution to develop a steep high-energy cut-off that shifts to lower energies with time.

2.8   Inverse Compton scattering

As well as emitting and absorbing photons, electrons will also scatter photons both in direction and to greater (or lower) energies. When a low-energy photon crosses the path of a high-energy electron it will be *up-scattered* to a higher energy while the electron loses this amount of energy in the process known as inverse Compton scattering. This process is very important in high-energy astrophysics as it provides a method of creating very high energy photons without the requirement of very high temperature plasmas.

Inverse Compton (IC) scattering requires the presence of very high energy electrons and hence is a relativistic phenomenon. Therefore a calculation of the energy gain of a scattered photon must be performed in the rest frame of the electron and then transformed into the observer's frame to account for relativistic effects. A full treatment of the derivation of the Compton scattering cross-section for all photon and electron energies is very complex, as it requires a detailed treatment of scattering angles, and how these are affected by relativistic aberrations, as well as for the potential of energy exchanges within the rest-frame. Such a derivation is presented in Jones (1968). However, most authors choose to simplify the calculation by limiting the energy regime to that of ultra-relativistic electrons ($\gamma \gg 1$) or else confine the calculation to the Thomson limit ($h\nu_0 \ll \gamma m_e c^2$).

In the Thomson limit a non-relativistic Thomson scattering event occurs in the rest frame of the electron, in which case there is no energy exchange to consider, just a scatter in angle. A Lorentz transformation will result in this scatter in angle in the rest-frame becoming an energy exchange in the observer's frame, the magnitude of which depends on the energy of the relativistic electron. In the rest frame of the electron, the incoming photon's frequency is transformed to

$$\nu_0' = \gamma\nu_0[1 - \beta\cos(\theta_0)], \tag{2.40}$$

following the definitions of Figure 2.6. All quantities in the rest frame of the electron are denoted with a prime, and those in the observer's frame are unprimed. As we are in the Thomson limit $\nu_0' = \nu_1'$. Therefore, on transforming back to the observer's frame after the
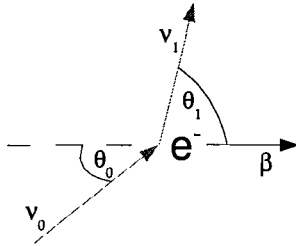
Figure 2.6: A graphical illustration of the symbols used in the derivation of the inverse Compton scattering process. A photon of frequency ν_0 approaches an electron at an angle θ_0 to the relativistic electron's velocity vector, with the velocity, β, given as a fraction of the speed of light. The photon scatters off the electron into the direction given by θ_1 and its frequency increases to ν_1.

scattering action we find

$$\nu_1 \;=\; \gamma\nu_1'[1 + \beta\cos(\theta_1)] \tag{2.41}$$

$$=\; \nu_0\gamma[1 - \beta\cos(\theta_0)]\gamma[1 + \beta\cos(\theta_1)]. \tag{2.42}$$

Hence, for $\theta_0 \approx 0$, the photon gains a factor of γ in energy upon transforming to the rest frame, and except for values of θ_1 near π it will gain another factor of γ when the energy is transformed back into the observer's frame. IC scattering increases a photon's energy by anything up to a factor of γ^2, which can be quite significant when ultra-relativistic electrons are present.

To determine how the scattered photons are distributed in frequency we not only need to know the seed photon distribution, but also the incidence angle of every photon with respect to every electron. To simplify this problem we assume that both the electrons and the photons have isotropic velocity vectors. As an inverse Compton scattering event is relativistic, the effect of relativistic aberration, where angles are transformed between frames, becomes very important. A Lorentz transformation of a scattering angle obeys,

$$\mu' = \frac{\mu - \beta}{1 - \beta\mu}, \tag{2.43}$$

where $\mu = \cos\theta$, and the prime denotes a quantity in the rest-frame of the electron. The consequence of this transformation for different Lorentz factors is shown in Figure 2.7. Relativistic aberration will, for scattering off ultra-relativistic electrons, have the effect of transforming the incidence angle of all of the incoming photons towards the electron in its
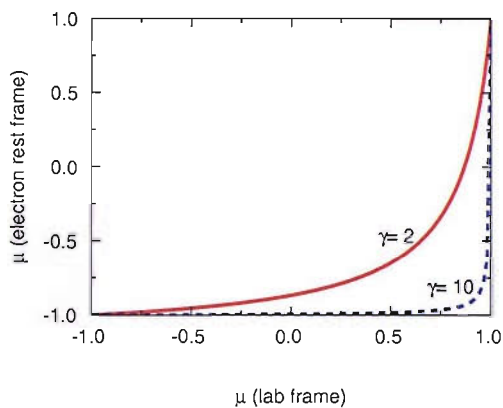
Figure 2.7: A plot of the relativistic aberration formula illustrating the effect of a frame transformation of the scattering angle, θ, where $\mu = \cos\theta$, for two Lorentz values. Relativistic aberration transforms the scattering angle in the observer's frame to always be in the direction of the electron's motion for high Lorentz values, regardless of the scattering angle in the electron's rest frame.

direction of motion in the rest-frame (see Figure 2.8). Subsequently if the photons are all scattered isotropically they will appear in the observer's frame as having been scattered along the direction of motion of the electron. As the electron distribution is assumed to be isotropic the scattered radiation will also be isotropic. However, in the case of relativistic jets there is a significant bulk motion of the electrons in one direction, and hence in this case the radiation will be scattered preferentially along the jet's direction of motion.

Figure 2.8: An illustration of the effect of relativistic aberration upon the incoming photon distribution in an ultra-relativistic inverse Compton scattering event. In the observer's frame, an isotropic distribution of incoming photons will all be scattered into the direction of the electron's motion.

Using equation 2.40 and considering that relativistic aberration at ultra-relativistic energies ensures that for an isotropic photon distribution $\theta_0 = \pi$ or 0, we find 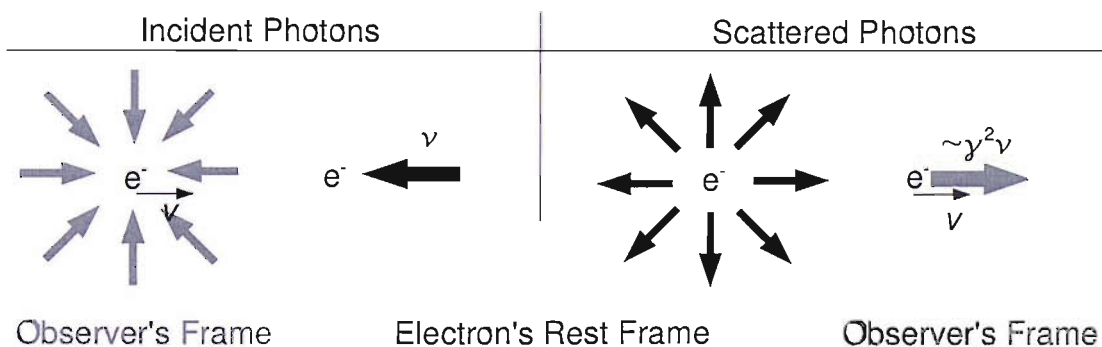that for $\beta \approx 1$, the incoming photon frequency is transformed to $\nu'_0 = 2\gamma\nu_0$ or $\nu_0/2\gamma$ in the electron's rest frame. Hence,

$$\frac{1}{2\gamma^2} < \frac{\nu_1}{\nu_0} < 4\gamma^2. \tag{2.44}$$

A full calculation of the scattered photon distribution for isotropic seed photons scattering off ultra-relativistic electrons in the Thomson limit (e.g. Blumenthal & Gould 1970), gives the following photon redistribution function

$$xF_{\rm C}(x) = 3\sigma_{\rm T} \, x \, (2x\ln(x) + x + 1 - 2x^2), \tag{2.45}$$

$$x = \frac{\nu_1}{4\gamma^2\nu_0}, \tag{2.46}$$

which is related to the frequency dependent interaction cross-section as $xF_{\rm C}(x) = \nu_0\sigma_\nu$. Therefore, using equation 2.7, bearing in mind that the interacting particles are the seed photons (hence $n_{\rm ip} = n_{\rm p}$) and that the number density of photons is given by

$$n_{\rm p} = \frac{4\pi}{hc} \int_{\nu_{0,\min}}^{\nu_{0,\max}} I_{\nu_0} \frac{{\rm d}\nu_0}{\nu_0} \tag{2.47}$$

we may calculate the emissivity of IC scattering to be

$$j_{\nu_1} = 4\pi \int_{\nu_{0,\min}}^{\nu_{0,\max}} I_{\nu_0} xF_{\rm C}(x) \frac{{\rm d}\nu_0}{\nu_0}. \tag{2.48}$$

Thus we may calculate the emission coefficient for IC scattering by a distribution of electrons, n_γ, using equation 2.8, to find

$$\epsilon_{\nu_1} = \int_{\nu_{0,\min}}^{\nu_{0,\max}} \int_{\gamma_{\min}}^{\gamma_{\max}} n_\gamma \, xF_{\rm C}(x){\rm d}\gamma \, I_{\nu_0} \frac{{\rm d}\nu_0}{\nu_0}, \tag{2.49}$$

where $\gamma_{\min} > \sqrt{\nu_1/4\nu_0}$, from equation 2.44, and $\gamma_{\max} < m_{\rm e}c^2/h\nu_0$ due to the Thomson limit, which is not imposed upon the frequency limits. The specific intensity of IC scattered radiation produced by this material depends upon the path-length that the radiation takes through the material, which in the optically thin limit we assume is the line of sight path-length, l, and hence $I_\nu = l\epsilon_\nu$.

A trans-relativistic solution to the problem of Compton scattering in the Thomson limit has been derived by Ensslin & Kaiser (2000), giving a photon redistribution function of

$$xF_{\rm C}(x) \quad = \quad \frac{3\sigma_{\rm T}}{8p^5}\left\{ -\frac{|1-\rho|}{4p}\left[1 + \rho(10 + 8p^2 + 4p^4) + \rho^2\right] + \right.$$

$$\rho(1+\rho)\left[\frac{3+3p^2+p^4}{\gamma}-\frac{3+2p^2}{2p}\left(2\ln(p+\gamma)-|\ln(\rho)|\right)\right]\Bigg\},\quad(2.50)$$

where $\rho=\nu_1/\nu_0$, and p is the relativistic momentum of the electron as defined in section 2.3. The range of frequencies over which a photon may be scattered by an electron of a given energy is now

$$\left|\ln\left(\frac{\nu_1}{\nu_0}\right)\right|<2\ln(p+\gamma),\qquad(2.51)$$

which agrees with equation 2.44 in the ultra-relativistic limit. Figure 2.9 plots this redistribution function, which at ultra-relativistic energies ($\gamma\gg1$) is identical to the function produced by equation 2.45. The dispersion of frequencies over which a photon may scatter is very narrow at low energies and becomes broader at high energies. At mildly-relativistic energies a photon is most likely to scatter without a change in frequency, whereas at ultra-relativistic energies the photon is most likely to increase in frequency by a factor of $2\gamma^2$ after scattering. Also it is clear from the second set of plots in Figure 2.9 that for all energies the probability of a *down-scatter* to lower frequencies decays exponentially, but is almost as probable as an up-scatter from mildly-relativistic electrons. Hence, this method does consider the effects of traditional Compton scattering as well as the astrophysically more important inverse Compton scatter process.
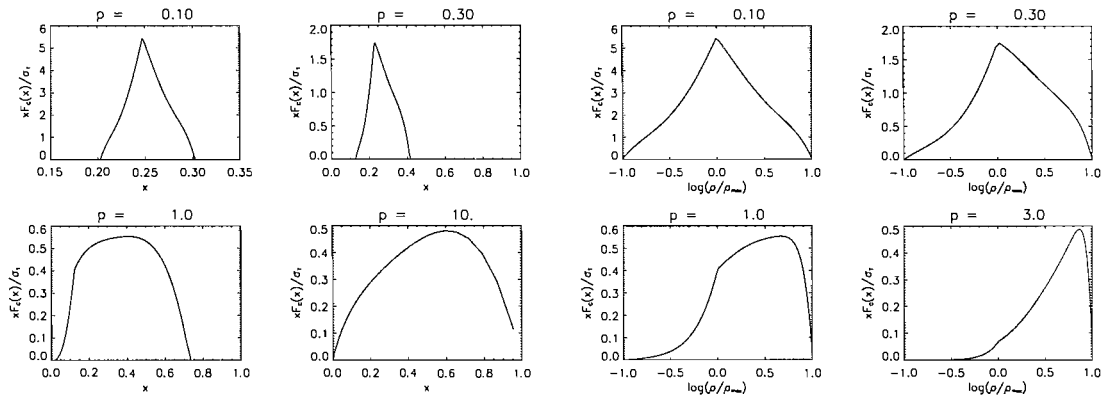
Figure 2.9: The photon re-distribution function for Compton scattering by an electron with relativistic momenta $p=0.1,0.3,1$ and 10 in the Thomson limit. The four plots on the left show this function with respect to x as defined by equation 2.46, with the lowest energy plot rescaled along the x-axis to illustrate the profile of the distribution function at low energies. The four plots on the right are scaled along the x-axis by the logarithm of the frequency increase ν_1/ν_0 as a fraction of the highest possible frequency gain $\rho_{\mathrm{max}}=(p+\gamma)^2$.

Beyond the Thomson limit the energy of the photon becomes comparable to that of the electron. Therefore, an energy exchange between the photon and the electron will occur in the rest-frame and the frequency of the photon increases as

$$\nu_1 = \frac{\nu_0}{1 + \frac{h\nu_0}{m_e c^2}(1 - \cos\theta)}. \tag{2.52}$$

Also a quantum mechanical phenomenon known as the Klein-Nishina effect becomes relevant, where the scattering cross-section reduces from the classical Thomson cross-section as the photon frequency increases. The Klein-Nishina cross-section in the high-frequency limit is

$$\sigma_{KN} = \sigma_T \frac{3 m_e c^2}{8 h\nu}\left[\ln\left(\frac{2h\nu}{m_e c^2}\right) + 0.5\right], \tag{2.53}$$

which tends to zero as the seed photon frequency tends to infinity. Therefore, at high seed photon frequencies the overall number of photons scattered by any given electron decreases.

A relatively simple photon redistribution function for Compton scattering by ultrarelativistic electrons in the Klein-Nishina regime was derived by Jones (1968) as

$$x F_C(x) = 3\sigma_T \, x \left\{ 2q\ln(q) + (1-q)\left[(1+2q) + \frac{\left(4\gamma q \frac{h\nu_0}{m_e c^2}\right)^2}{2\left(1 + 4\gamma q \frac{h\nu_0}{m_e c^2}\right)}\right]\right\}, \tag{2.54}$$

$$q = \frac{x}{1 - \frac{h\nu_1}{\gamma m_e c^2}}. \tag{2.55}$$

In the Thomson limit this formula reduces to that of equation 2.45.

A photon redistribution function that applies to all electron energies and photon frequencies without being constrained to the Thomson limit may be found in Jones (1968). Although it is too lengthy to include here, plots of this function in the Klein-Nishina regime for a range of electron energies are shown in Figure 2.10a. This figure clearly shows that overall scattering is reduced and that the maximum frequency gain or loss of a scattered photon is no longer given by equation 2.51.

Obviously, the Compton scattered photon distribution is strongly dependent upon the distribution of electrons and seed photons. One important astrophysical occurrence of Compton scattering is in synchrotron radiation emitting plasmas. If the optical depth
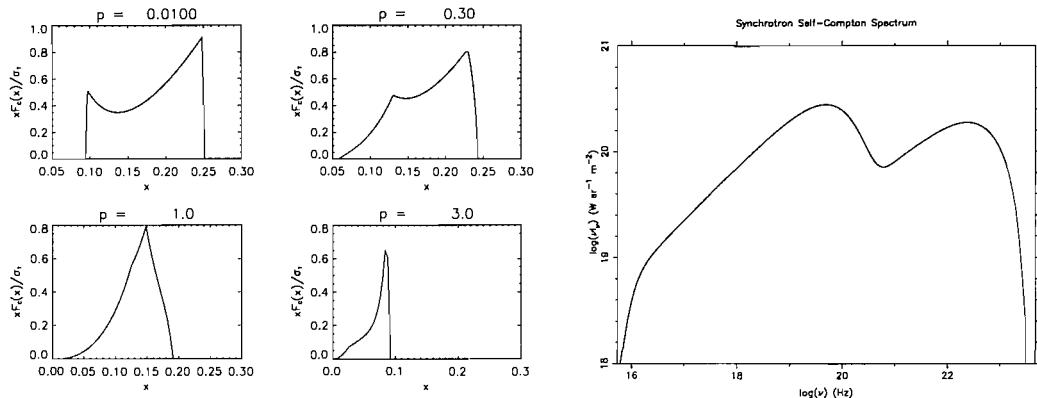
Figure 2.10: (a - left panel) The re-distribution function including Klein-Nishina effects for the Compton scattering of a 10^{20} Hz seed photon by an electron with relativistic momenta $p = 0.01, 0.3, 1$ and 3. Note that as well as significant alterations to the profile the overall magnitudes are smaller for larger p. (b - right panel) The spectral energy distribution of synchrotron self-Compton emission.

in these plasmas is sufficiently high, then a process known as synchrotron self-Compton (SSC) emission will become important. Here the same population of relativistic electrons will both emit synchrotron photons and inverse Compton scatter these photons to higher energies. We can determine the resulting spectrum by calculating the seed photon distribution, I_{ν_0}, using the synchrotron radiation formulae, and then using this distribution to calculate the scattered distribution from equation 2.49. An example spectrum for scattering by a non-thermal power-law distribution of electrons is shown in Figure 2.10b, and is similar to typical blazar spectra (see e.g. Chiaberge & Ghisellini 1999), where the emission from a relativistic jet is inclined directly towards the observer.

As for synchrotron radiation we can calculate the energy loss of a Compton scattering event by integrating the single-electron emissivity over all frequencies (see equation 2.37). For the simple case of an ultra-relativistic electron in the Thomson limit we find the energy loss rate to be

$$-\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{4\sigma_{\mathrm{T}}}{3m_{\mathrm{e}}c}U_{\mathrm{rad}}\gamma^2, \tag{2.56}$$

where

$$U_{\mathrm{rad}} = \frac{4\pi}{c}\int_{\nu_{0,\mathrm{min}}}^{\nu_{0,\mathrm{max}}} I_{\nu_0}\mathrm{d}\nu_0, \tag{2.57}$$

is the energy density of the seed photons, and $\nu_{0,\mathrm{max}}$ is limited by the Thomson scattering

approximation, and thus is a function of γ. It is interesting to compare the synchrotron radiation loss rate to the IC loss rate, as they are both functionally dependent on γ^2, so their ratio is simply given by

$$\frac{\dot{\gamma}_{\rm IC}}{\dot{\gamma}_{\rm SR}} = \frac{U_{\rm rad}}{U_B}. \tag{2.58}$$

Thus for IC scattering to be the dominant process over synchrotron radiation the radiation field energy density must be greater than that of the magnetic field.

2.9   Bremsstrahlung

Bremsstrahlung radiation is the emission of photons due to an electron (or positron) being accelerated by a close encounter with another charged particle (be it another electron/positron or an ion). Hence, in determining the emission rate of photons due to the bremsstrahlung process the interaction cross-section needs to be derived. As with the other emission processes described in this chapter, a full derivation of the emissivity that is valid for both mildly-relativistic and ultra-relativistic electrons is extremely complicated. Most authors choose to derive the bremsstrahlung cross-section in the limiting cases, and hence thermal bremsstrahlung is treated separately to relativistic thermal bremsstrahlung or non-thermal bremsstrahlung. Also the interaction cross-section depends upon whether the interacting particle is an electron, positron, or ion.

In the ultra-relativistic limit the bremsstrahlung cross-section is the same for all possible interactions, and is given by Zdziarski, Coppi & Lamb (1990) as

$$\sigma_{\nu,\rm ur}(\gamma) = \frac{3\alpha_{\rm f}\sigma_{\rm T}}{2\pi\nu}\xi\left(\xi + \frac{1}{\xi} - \frac{2}{3}\right)\left(\ln\left(\frac{2\gamma^2\xi}{1-\xi}\right) - \frac{1}{2}\right), \tag{2.59}$$

where $\alpha_{\rm f}$ is the fine-structure constant, and $\xi = 1 - h\nu/\gamma m_{\rm e}c^2$ gives the emitting electron's energy as a fraction of its initial value. For the interaction of mildly relativistic electrons (or positrons) such as in a thermal distribution, Gould (1980) derived an interaction cross-section of

$$\sigma_{\nu,\rm mr}(\gamma) = \frac{\alpha_{\rm f}\sigma_{\rm T}}{10\pi\nu}\left\{\xi\left[20 - 6\frac{1+\xi^4}{(1+\xi^2)^2}\right] + \right.$$
$$\left.\ln\left(\frac{1+\xi}{1-\xi}\right)\left[6(1+\xi^2) - \frac{7(1-\xi^4)^2 + 3(1-\xi^2)^4}{2(1+\xi^2)^3}\right]\right\}. \tag{2.60}$$

Further interaction cross-sections for mildly-relativistic electron-ion encounters are presented in Svensson (1982), but are not included here as only electron plasmas are considered in this work.

Using the techniques of section 2.4, we can deduce that the emissivity of a mildly relativistic electron interacting with other mildly relativistic electrons is

$$j_{\nu,\mathrm{mr}}(\gamma) = h\nu \; c \; n_{\mathrm{e,mr}} \; \sigma_{\nu,\mathrm{mr}}(\gamma). \tag{2.61}$$

We consider the emissivity of ultra-relativistic electrons to be the sum of their interactions with the mildly-relativistic electrons and together with their interactions with themselves,

$$j_{\nu,\mathrm{ur}}(\gamma) = h\nu \; c \; (n_{\mathrm{e,mr}} + 2n_{\mathrm{e,ur}})\sigma_{\nu,\mathrm{ur}}(\gamma). \tag{2.62}$$

As shown by Zdziarski, Coppi & Lamb (1990), the emission of ultra-relativistic electrons interacting with themselves is twice as large as ultra-relativistic electrons interacting with mildly-relativistic electrons. Hence the total emission co-efficient for the plasma is

$$\epsilon_\nu = \frac{h\nu c}{4\pi\sigma_{\mathrm{T}} l} \left[\tau_{\mathrm{T,mr}} \int_1^{\gamma_{\mathrm{ur}}} n_\gamma \sigma_{\nu,\mathrm{mr}} \mathrm{d}\gamma + (\tau_{\mathrm{T,mr}} + 2\tau_{\mathrm{T,ur}}) \int_{\gamma_{\mathrm{ur}}}^{\gamma_{\mathrm{max}}} n_\gamma \sigma_{\nu,\mathrm{ur}} \mathrm{d}\gamma \right], \tag{2.63}$$

where γ_{ur} is the arbitrary division energy between ultra-relativistic and mildly-relativistic electrons, and γ_{max} is the highest energy electron in the distribution. Correspondingly we can derive the absorption co-efficient from equation 2.10 to be

$$\begin{aligned}
\kappa_\nu \;=\; & \frac{hc}{8\pi m_{\mathrm{e}}\sigma_{\mathrm{T}} l\nu} \left[\tau_{\mathrm{T,mr}} \int_1^{\gamma_{\mathrm{ur}}} \frac{n_\gamma}{\gamma p} \frac{\mathrm{d}}{\mathrm{d}\gamma}(\gamma p \sigma_{\nu,\mathrm{mr}}) \mathrm{d}\gamma \right. \\
& \left. + (\tau_{\mathrm{T,mr}} + 2\tau_{\mathrm{T,ur}}) \int_{\gamma_{\mathrm{ur}}}^{\gamma_{\mathrm{max}}} \frac{n_\gamma}{\gamma p} \frac{\mathrm{d}}{\mathrm{d}\gamma}(\gamma p \sigma_{\nu,\mathrm{ur}}) \mathrm{d}\gamma \right].
\end{aligned} \tag{2.64}$$

Finally, the energy loss rate for electrons is

$$-\dot{\gamma}(\gamma) = \begin{cases} \frac{h}{\sigma_{\mathrm{T}} l m_{\mathrm{e}} c} \left[\tau_{\mathrm{T,mr}} \int \nu\sigma_{\nu,\mathrm{mr}} \mathrm{d}\nu \right] & ; \quad \gamma < 2, \\ \frac{h}{\sigma_{\mathrm{T}} l m_{\mathrm{e}} c} \left[(\tau_{\mathrm{T,mr}} + 2\tau_{\mathrm{T,ur}}) \int \nu\sigma_{\nu,\mathrm{ur}} \mathrm{d}\nu \right] & ; \quad \gamma \geq 2. \end{cases} \tag{2.65}$$

The bremsstrahlung spectrum is calculated in the same manner as for synchrotron radiation, by using the radiative transfer equation (see equation 2.16). Bremsstrahlung radiation has a characteristically flat spectrum, with an absorption cut-off where $\tau_\nu > 1$, as shown in Figure 2.11. The shape of the bremsstrahlung spectrum depends only weakly
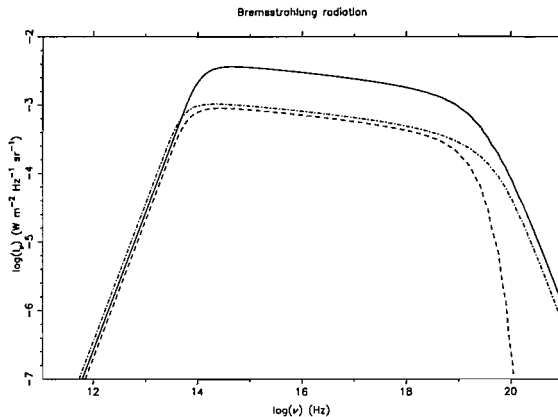
Figure 2.11: A characteristic bremsstrahlung spectrum illustrating the radiation from a hybrid (solid line), thermal (dashed line) and non-thermal (dot-dashed line) electron distribution.

on the electron distribution parameters and hence is almost identical for both thermal distributions and non-thermal distributions of equal average energies and densities. However, its normalisation strongly depends on the electron density.

2.10 The kinetic equation

As a consequence of emitting, absorbing and scattering radiation, the electron energy distribution will evolve with time. We model this evolution with the kinetic equation, a partial differential equation, first derived for ultra-relativistic distributions by McCray (1969). Ghisellini *et al.* (1988) present the kinetic equation without ultra-relativistic assumptions as

$$\frac{\partial n_\gamma}{\partial t} = \frac{\partial}{\partial \gamma}\left[\Lambda(\gamma)n_\gamma + H(\gamma)\gamma p\frac{\partial}{\partial \gamma}\left(\frac{n_\gamma}{\gamma p}\right)\right] + Q(\gamma) - \frac{n_\gamma}{T(\gamma)}, \tag{2.66}$$

where $\Lambda(\gamma)$ [s^{-1}] is the total cooling rate due to free energy losses. Essentially this is the sum of all single-electron emissivities integrated over all frequencies. Hence, for radiative cooling only,

$$\Lambda(\gamma) = -\dot{\gamma}_{\mathrm{S}} + -\dot{\gamma}_{\mathrm{C}} + -\dot{\gamma}_{\mathrm{B}}, \tag{2.67}$$

where the synchrotron cooling rate, $-\dot{\gamma}_{\mathrm{S}}$, is given by equation 2.37, and the bremsstrahlung cooling rate, $-\dot{\gamma}_{\mathrm{B}}$, by equation 2.65. The cooling rate due to Compton scattering is

calculated following Blumenthal & Gould (1970) as

$$-\dot{\gamma}_C = \frac{4\pi}{m_e c^2} \int \int \frac{\nu_1 - \nu_0}{\nu_1} I_{\nu_0} x F_C(x) \frac{d\nu_0}{\nu_0} d\nu_1, \qquad (2.68)$$

where we have used the full emissivity that includes both the gain of scattered photons and the loss of seed photons in calculating the energy exchanges. Hence, the possibility of energy gains are included.

The factor $H(\gamma)$ [s^{-1}] contains the effects of induced losses and gains upon the electron distribution, essentially the absorption of photons by synchrotron and bremsstrahlung processes,

$$H(\gamma) = \frac{1}{2m_e^2 c^2} \int \frac{I_\nu}{\nu^2} [j_{\nu_S}(\gamma) + j_{\nu_B}(\gamma)] d\nu. \qquad (2.69)$$

The remaining factors $Q(\gamma)$ [m^3 s^{-1}] and $T(\gamma)$ [s] are source and sink functions respectively, where possible sources would be the injection of non-thermal electrons (e.g. equation 3.10), and possible sinks are electron escape where $T(\gamma) = t_{\mathrm{esc}}$, or the adiabatic expansion of the plasma in volume, V, where $T(\gamma) = V/\dot{V}$.

Radiative cooling will cause an initially power-law distributed non-thermal electron distribution to 'age' whereby the highest energy electrons cool faster than the lower energy electrons. This phenomenon (first studied by Rees 1967 in the case of synchrotron radiation cooling) leads to a non-thermal distribution of low-energy electrons with a high-energy cut-off that moves to lower energies with time. An example of this radiative cooling process is shown in Figure 2.12a. In the absence of any heating this aging process will continue at an ever decreasing rate. However, both synchrotron and bremsstrahlung radiation processes have inverse processes that will heat low-energy electrons by the absorption of low-energy photons. For certain parameters an equilibrium can be achieved between the heating rate and the cooling rate causing the electron distribution to *thermalise* into a Maxwell-Boltzmann type distribution. Ghisellini *et al.* (1988) described this mechanism in the case of synchrotron emitting plasmas as the 'synchrotron boiler'. Compton scattering will also heat low-energy electrons, but typically the strong cooling of inverse Compton scattering dominates, preventing any thermalisation.

When a non-thermal power-law distribution following $Q(\gamma) = Q_0 \gamma^{-q}$ is continually injected into the plasma a superposition of power-law distributions of many different ages

occurs, leading to the electron distribution having a power-law with an index of $-q - 1$ that is steeper than the injected distribution. If electrons are continually injected, an equilibrium is more readily achieved if some electrons escape from the system. Hence, the concept of an escape time-scale, t_{esc}, is usually employed, with the physical justification that an electron plasma is not necessarily a closed system and that electrons may escape. Placing this escape time-scale into the kinetic equation results in the normalisation of the electron distribution decreasing over time as we have chosen to make the time-scale independent of energy.

A more physically correct model would replace the injection and escape terms with a calculation of the effect of shock-waves or magnetic re-connection events heating the existing electron distribution by adding extra terms to the kinetic equation co-efficients. Some early attempts to model the creation of a power-law distribution of non-thermal electrons from an initially cool low-energy distribution using the kinetic equation have already been undertaken (Li, Kusunose & Liang 1996), but are beyond the scope of this current work.
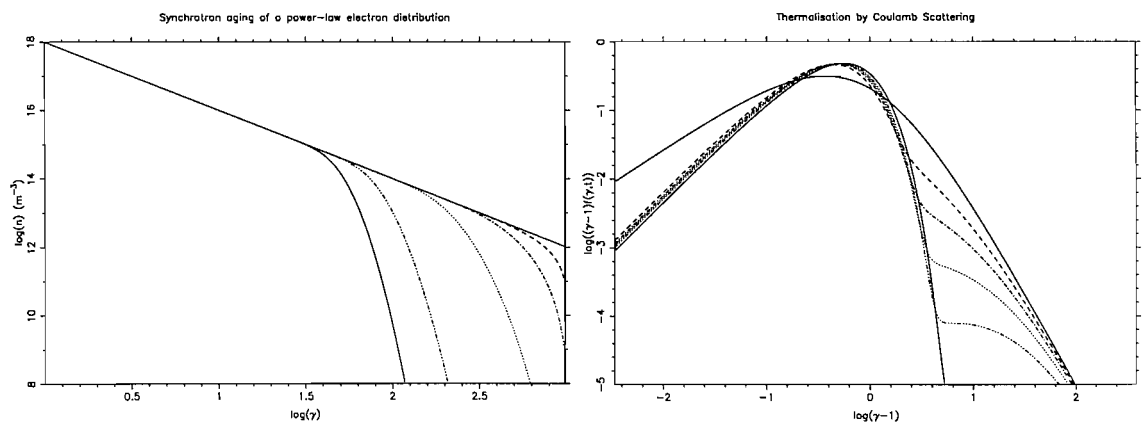
Figure 2.12: (a - left panel) Synchrotron aging whereby the radiative cooling of an initially power-law distributed non-thermal electron distribution leads to a high-energy cut-off in the electron distribution that evolves with time. (b - right panel) Thermalisation due to the Coulomb scattering process. An initial power-law distribution of electrons evolves towards a Maxwellian distribution. The distribution at times $t = 0$, 0.125, 0.25, 0.5, 1.0, and 3.0 $\ln \Lambda \, t_c$ are plotted, where t_c is the Coulomb timescale, given in equation 2.75.

2.11   Coulomb scattering

When considering the evolution of electron distributions it is also important to consider the non-radiative energy exchange process known as Coulomb scattering. Coulomb scattering describes a process similar to bremsstrahlung where an energy exchange occurs when two electrons interact. However, in the case of Coulomb scattering we are dealing with the net electric field in the plasma caused by the whole distribution of electrons transferring energy to individual electrons. This process has the effect of thermalising the distribution.

Our treatment of Coulomb scattering follows that of Nayakshin & Melia (1998), which was the first to model the thermalisation of an arbitrary electron distribution by this process. Their method was to define a kinetic equation of the form,

$$\frac{\partial n_\gamma}{\partial t} = -\frac{\partial}{\partial \gamma}\left[a(\gamma)n_\gamma\right] + \frac{\partial^2}{\partial \gamma^2}\left[D(\gamma)n_\gamma\right], \tag{2.70}$$

to calculate the redistribution of energy within the electron distribution. Here the $a(\gamma) = <\dot{\gamma}>$ co-efficient gives the mean of the energy exchange, and the $D(\gamma) = \mathrm{d}/\mathrm{d}t <(\Delta\gamma)^2>$ co-efficient gives the dispersion of the energy exchange.

The co-efficients are derived to be

$$a(\gamma) = -\frac{3\sigma_T c \ln \Lambda}{4p\gamma}\int \frac{n_{\gamma'}}{p'\gamma'}(\gamma - \gamma')\chi(\gamma,\gamma')\mathrm{d}\gamma', \tag{2.71}$$

where $\ln \Lambda$ is the Coulomb logarithm, which is a weak function of energy that is assumed to be constant with a value of 20, and

$$\begin{aligned}\chi(\gamma,\gamma') &= \int_{\gamma^-}^{\gamma^+} \frac{x^2 \mathrm{d}x}{\sqrt{(x+1)(x-1)^3}} \\ &= \left[\sqrt{x^2-1} - \sqrt{\frac{x+1}{x-1}} + 2\ln\left(\sqrt{\frac{x-1}{2}} + \sqrt{\frac{x+1}{2}}\right)\right]_{\gamma^-}^{\gamma^+}\end{aligned} \tag{2.72}$$

and $\gamma^\pm = \gamma\gamma'(1 \pm \beta\beta')$. The dispersion is derived as

$$D(\gamma) = \frac{3\sigma_T c \ln \Lambda}{4p\gamma}\int \frac{n_{\gamma'}}{p'\gamma'}\left[\zeta(\gamma,\gamma') - \frac{1}{2}(\gamma-\gamma')^2\chi(\gamma,\gamma')\right]\mathrm{d}\gamma', \tag{2.73}$$

where we have used the corrected version given by Blasi (2000), and

$$\zeta(\gamma,\gamma') = \int_{\gamma^-}^{\gamma^+} \frac{x^2}{\sqrt{x^2-1}}\left[\frac{(\gamma+\gamma')^2}{2(1+x)} - 1\right]\mathrm{d}x$$

$$= \frac{1}{2} \left[\frac{(\gamma + \gamma')^2 \sqrt{x^2 - 1}}{(1 + x)\sqrt{(1 + x)^2 - 2x - 2}} + (\gamma + \gamma')^2 \sqrt{x^2 - 1} - x\sqrt{x^2 - 1} \right.$$

$$\left. - \left[1 + (\gamma + \gamma')^2 \right] \ln \left(x + \sqrt{x^2 - 1} \right) \right]_{\gamma^-}^{\gamma^+}. \tag{2.74}$$

The Coulomb scattering process will thermalise a distribution of electrons in the characteristic timescale of

$$t_{\mathrm{c}} = (n_{\mathrm{e}} c \sigma_{\mathrm{T}} \ln \Lambda)^{-1} \tag{2.75}$$

with a temperature that depends on the average energy of the initial non-thermal distribution. This thermalisation process is illustrated in Figure 2.12b, where a power-law distribution of electrons evolves towards a Maxwellian thermal distribution. The slight shift in the energy of the distribution's thermal peak is due to a numerical artefact and is generally suppressed when radiative cooling effects are included.

Chapter 3

The HEART Model

3.1 Introduction

Attempts to model the spectra of HEA sources have in the past relied upon the equations derived in the previous section to calculate I_ν for a given single electron distribution, n_γ, that does not vary with time or space, and thus determine the relevant system parameters. More advanced methods consider the effect of energy losses upon the electron distribution, and for a given set of parameters, attempt to determine the *steady-state* solution for n_γ, and thus I_ν, by injecting high-energy electrons to balance the losses due to cooling. This is achieved by finding solutions to the partial differential equation known as the *kinetic equation*, which is described in section 2.10, and one such popular code for doing this is EQPAIR (Coppi 1992, 1999). The major shortcoming of this method is that it cannot calculate time-dependent phenomena to re-create light curve observations. Furthermore, it assumes the source is a homogeneous object where escaping photons encounter the same electron distribution, n_γ, throughout their escape path.

In order to counter these problems the radiative transfer within the source must be calculated. In the past this has only been achieved with the use of Monte Carlo simulations, first developed by Pozdnyakov, Sobol' & Sunyaev (1977), which determine the path taken by individual photons through the emitting object. By restricting the focus to individual

photons, the effect of energy losses upon the electron distribution becomes harder to cal-
culate, and such methods can only deal with limited photon energy ranges. Furthermore,
Monte Carlo simulations of radiative transfer are incapable of modelling changes to the
electron distribution on timescales comparable to the photon crossing time.

The development of the High-Energy Astrophysics Radiative Transfer (HEART) model
described in this section, was motivated by the problem of calculating reliable time-
dependent spectra over a wide range of photon energies. Using the kinetic equation
method as a starting point, the problem of radiative transfer was solved by dividing
the emitting object into multiple zones, which we call *cells*, each containing an individual
electron distribution. Radiative transfer can then be simulated by transferring the photon
distributions, each calculated independently for every cell, to adjacent cells on time-steps
determined by the light crossing time over each cell width. Hence, inhomogeneities may be
modelled by giving cells differing initial electron distributions. When cells emit radiation,
the corresponding energy loss incurred by the electron distribution and the additional
heating from Coulomb scattering may be calculated using the kinetic equation.

The HEART code is also the first model to freely evolve arbitrary electron distributions
over a truly trans-relativistic distribution of energies whilst including accurate modelling
of the photon distributions created by these trans-relativistic electrons due to Compton
scattering, bremsstrahlung, and synchrotron emission. Both transitional behaviour and
equilibrium solutions for both distributions may be found, the latter by balancing an
injection rate of electrons with the radiative cooling and free escape of electrons. In the
following sections I describe the numerical design of the HEART model in full detail.

3.2   Conceptual design

The HEART code tackles the problem of calculating the emission of radiation from high-
energy astrophysical plasmas by constructing a three-dimensional model of the emitting
plasma divided into homogeneous regions. These discrete cells, containing unique electron
distributions, are modelled as tessellating cubes with each cube face attached to no more
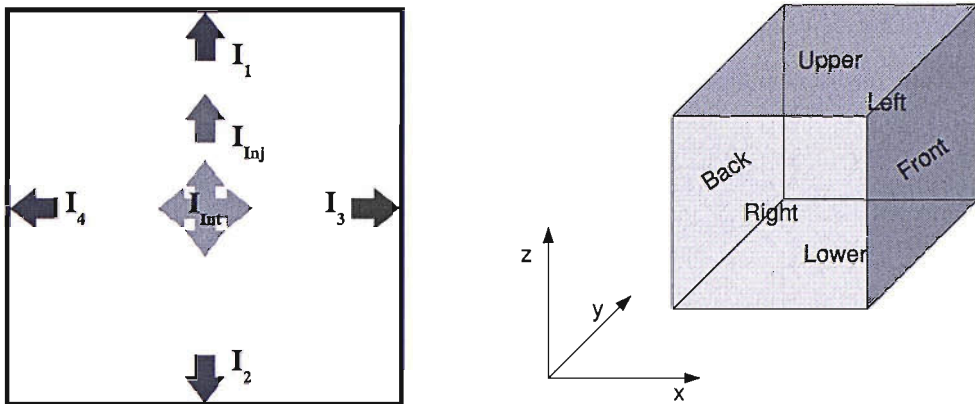
Figure 3.1: (a - left panel) An illustration of the photon distributions within a two-dimensional HEART cell. Photon distributions that enter the cell from adjacent cells are transferred to the opposing face of the cell. The distribution at each cell face is numbered as I_1, I_2 etc. The cell's intrinsic emission, I_{Int}, e.g. synchrotron emission, is added to the distribution at each face. Some simulations, such as the one discussed in Chapter 4, require external radiation to be injected into the cell and this is denoted by I_{Inj}. (b - right panel) The notation used to describe the faces of a three-dimensional cell, with respect to the principle axes.

than one adjacent cell. Each cell emits an intrinsic distribution of photons in all directions due to various emission processes. To simplify the problem of transferring this radiation between the cells we only consider six directions of motion, those perpendicular to the six faces of the cubic cells. Therefore, assuming the intrinsic emission is emitted isotropically, the radiation is distributed evenly to each cell face and transferred to the adjacent cells. If there are no adjacent cells at a given face then that radiation escapes. To model the radiative transfer process accurately it is important to ensure that a given photon distribution maintains its original direction of motion. Thus, as well as a unique electron distribution, each cell contains six separate photon distributions, one for each direction. Figure 3.1 defines our notation for the cell faces and photon distributions present in each cell.

Photon distributions are measured by the specific intensity, as defined in section 2.2. This quantity is simple to transfer between cells as it is independent of direction, surface area, and time. Numerically, the quantity of transferred emission at each cell face is equal to the cell's intrinsic emission. Hence, using the definitions in Figure 3.1, $I_1 = I_1 + I_{Int}$,

$I_2 = I_2 + I_{\text{Int}}$, and so on. All transferred quantities are equal. However, when measuring the total quantity of radiation in a cell, I_{cell}, we use an *average* specific intensity and hence

$$I_{\text{cell}} = I_{\text{Int}} + \frac{I_1 + I_2 + I_3 + I_4 + I_5 + I_6}{6}, \tag{3.1}$$

prior to the transfer of the intrinsic emission to each cell face.

The total surface specific intensity observed from the modelled plasma in a given direction is also defined by the average specific intensity of all of the photon distributions escaping from cell faces perpendicular to that direction. The flux is found by integrating this specific intensity over all solid angles, defined by

$$d\Omega = \frac{dA}{D^2}, \tag{3.2}$$

where D is the distance to the observer. Hence, the total surface flux is simply given by the specific intensity multiplied by the surface area of each emitting cell face and divided by the distance to the observer squared.

Using cubic cells ensures that the radiation transfer time across the entire width of a cell is equal in all directions. This timescale determines the logical time step duration, Δt, for our simulation

$$\Delta t = \frac{\Delta w}{c}, \tag{3.3}$$

where Δw is the cell width, and c is the speed of light. Therefore, on each time step incoming radiation will pass through the entire width of the cell. Intrinsic radiation is also modelled as having passed through the entire width of the cell in one time step.

Dividing a plasma of size w into multiple cells allows processes on timescales that are shorter than the light crossing time of the entire plasma, $t_{\text{cross}} = w/c$, to be modelled. However, it is important to realise that the simulation will undergo a transitional period of flux increase between $t = 0$ and t_{cross}, because photons from the deepest cells will not yet have escaped from the plasma. This phenomenon is best illustrated in the one-dimensional case described in Figure 3.2.

As a result of the simulation time step being inversely proportional to the number of cells in one dimension, R_{cells}, the computing time for a given three-dimensional simulation
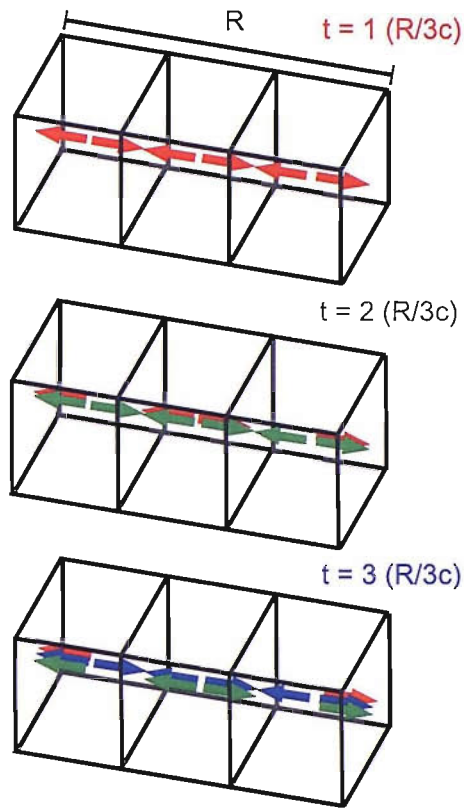
Figure 3.2: At the start of a radiative transfer simulation only the radiation from the outer cells has had time to escape, so the total emitted flux after the first time-step is just $1/N$ of equilibrium flux, where N is the number of cells in this one-dimensional example. After N time steps, or the light crossing time, the total escaping flux will reach an equilibrium in the absence of any other time-dependent effects.

increases as R_{cells}^4. Therefore, spatial resolution is severely limited. Each emission process provides separate constraints upon the spatial resolution to ensure that cells are sufficiently small to model any inhomogeneities, and these limiting factors are described in the relevant sections.

The following sections are split into two main groups. Firstly, section 3.3 describes the computational algorithms that handle the entire ensemble of cells. Then, described in section 3.4, are the numerical methods used by the HEART code to calculate changes to the electron and photon distributions, through emission processes, within each individual cell.

3.3   Computational method for multiple cells

3.3.1   Construction of the three-dimensional model

The HEART code is designed to model the three-dimensional geometry of any astrophysical object. This is achieved by dividing the object into cubes, which we call cells. At present, the code can model any three-dimensional, solid shape with the spatial resolution only limited by the available memory of the user's computer. All shapes are defined by supplying the Cartesian equation that defines the shape's outer boundary, together with the spatial resolution required, given by the number of cells in one dimension, R_{cells}. The code has not been extended to include hollow shapes where inner boundaries would also have to be specified, as this would complicate the radiation transfer process. In the case of hollow shapes, it is better to construct the solid equivalent, and then declare the cells within the inner boundary to be devoid of electrons. This will allow the radiative transfer to still be modelled accurately.

Cells are created in the form of a *linked list* of cell objects, where each cell object is accessed by following a link from the previously created cell. However, this data structure is only one-dimensional and, to model radiative transfer between cells, links to every adjacent cell must be created. Each cell's position co-ordinate is given by the distance of the cell centre to the origin. If R_{cells} is an odd number, then the origin lies in the centre of the first cell, otherwise the origin lies at the centre of the first cell's lower face (see Figure 3.1b for face definitions). Cells are created such that all cell centre co-ordinates lie within the region defined by the outer boundary equation.

To describe how the code constructs the three-dimensionally linked list of cells in the shape described by the boundary equation, I shall use the analogy of a tree (also see Figure 3.3). The trunk is formed by starting at the origin and moving along the positive z-axis one cell's width at a time until the boundary is reached. Then, returning to the origin, the trunk of cells is completed in the same manner for the negative z-axis. Returning to the trunk cell at the uppermost boundary, a branch of cells are then created along the positive x-axis until the boundary is reached, and then returning to the trunk cell the branch is

completed along the negative x-axis. Similarly, for each branch cell, starting at the cell at the positive x-axis limit and working down to the negative x-axis limit, a line of twig cells along the y-axis is formed, first in the positive y-axis direction and then along the negative y-axis. When the line of twig cells attached to the last branch cell is complete the code returns to the next trunk cell, and fills the remaining volume of the shape in the same manner.

Links between cells are formed on a reciprocal basis. When a new cell is created it is given a link to the cell in the previous position, and that cell is given a link to the newly created cell. Creating a linked list in this manner does not automatically ensure that a cell is linked to every neighbour. This will only be true for the trunk cells. The branch cells will be missing links to cells adjacent to their upper and lower faces, and the twig cells will be missing those links as well as links to cells adjacent to their front and back faces. Therefore, when a branch is being formed, reciprocal links are always formed to cells adjacent to the upper faces of the branch cells. As we start at the uppermost trunk cell and work downwards this will ensure that every branch cell is fully linked to its adjacent cells. The cell adjacent to the upper face is located by following the links back
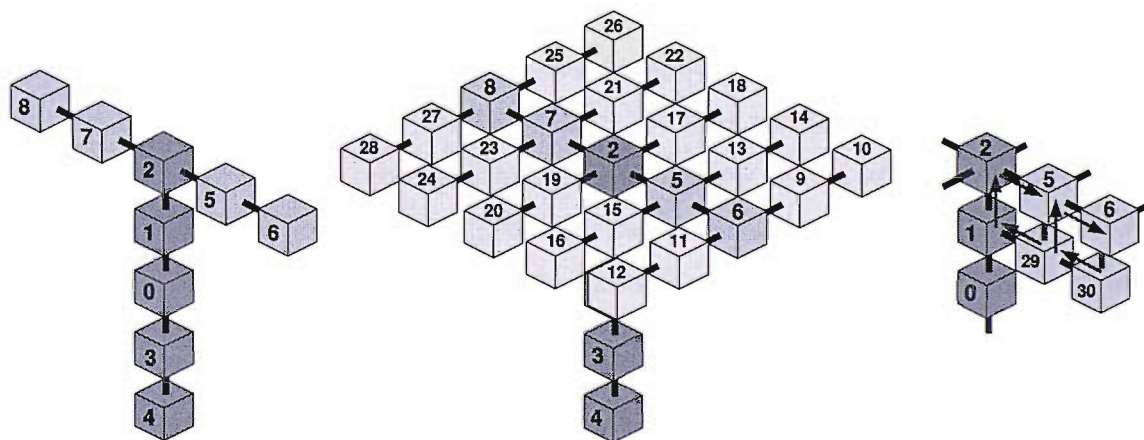
Figure 3.3: An illustration showing how a cube shape would be created by our algorithm, from individual cells. The numbers denote the order in the which the cells are created, and the cube has only be partially constructed for clarity. The black lines denote the initial links between cells. The diagram on the far right illustrates the algorithm to link cells from adjacent branches. Black lines denote existing reciprocal links, grey ones denote reciprocal links formed by this algorithm, and the grey arrows show the path taken by the algorithm to locate the desired cell.
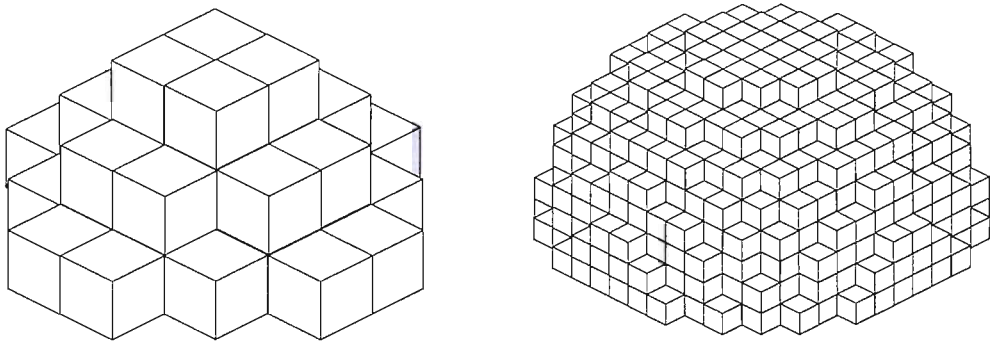
Figure 3.4: A hemisphere geometry modelled at differing resolutions, $R_{cells} = 6, 16$. The total number of cells are 68 and 1088 respectively.

along the branch by one cell, up by one cell (a link that will exist for trunk cells and all subsequent branch cells), and then along the upper branch by one cell (see Figure 3.3). This general technique is also applied to the twig cells, where reciprocal links are formed to cells attached to their upper faces, and cells attached to their back faces.

Using these techniques we can create any three-dimensional shape using cubic cells at any desired spatial resolution, with links between every adjacent cell (see examples in Figure 3.4). However, we do assume that the maximum length of cells in the z-direction occurs at x = 0, y = 0, and that the maximum length of cells in the x-direction on a given x-y plane occurs at y = 0. For example, if the height of the object is 10 cells in the z-direction at (0,0) and is 20 cells at (5,0) the code will fail to model the shape. In this simple example, the problem may be solved simply by shifting the origin, but for shapes with a concave surface, like a bowl, it can be difficult to find a suitable point for the origin. Furthermore, some rays leaving the surface of such shapes will intersect with another part of the shape. Therefore, when a shape has a concave surface it should be treated like a hollow shape, by creating *empty cells*, allowing the radiative transfer to be accurately modelled.

3.3.2 Radiation transfer

Storing the cells inside a linked list data structure simplifies the task of transferring photons between the cells. As illustrated in Figure 3.1a, every cell contains a separate photon

distribution for each face of the cell. Following the definition of the simulation time step as the time required for a photon distribution to traverse the entire length of a cell, the final process at the end of every simulated time step is to transfer these photon distributions to the corresponding faces of the adjacent cells. So, for example, upper face photon distributions are transferred to the upper face of the cell above. As the entire photon distribution is transferred in every case it is important to ensure that the algorithm that performs the radiation transfer only moves photon distributions to empty spaces. Therefore, the first step of this process is to remove the photon distributions that escape from the emitting object. All photons distributions from faces not attached to other cells, i.e. faces with *null* links, are summed and multiplied by the surface area to give the total surface flux. This quantity is then divided by $4\pi D^2$, where D is the distance to the observer, to give the observed flux.

With the surface flux photon distributions removed, the radiation transfer algorithm traverses the linked list of cells until a cell with a null link is found. The algorithm then transfers photons along a line, starting with the distribution that fills the space created for the null link face, as described by Figure 3.5. Faces in other dimensions of that cell with null links are then detected, and radiation transfer along those lines is then performed. When all null linked faces have had cell photon distributions transferred to them, then the algorithm moves on to the next cell in the list. This continues until all cells have been covered.

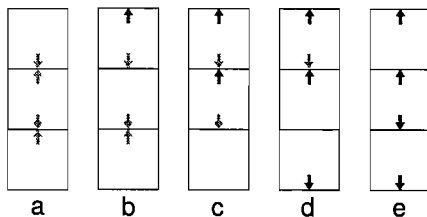**Figure 3.5:** A one-dimensional illustration of the radiative transfer algorithm applied to three cells. Black arrows denote photon distributions that have been transferred, grey arrows those that have not yet been transferred. On one time step every photon distribution travels the length of an individual cell. Each step taken by the radiative transfer algorithm, over a single time step, is shown here, and labelled alphabetically.

3.4   Numerical method

3.4.1   Distribution discretisation

Both the photon and electron distributions will encompass a large range of energies, over several orders of magnitude, so both distributions are divided into logarithmic grids. In our code the number of divisions, or bins, in each grid are flexible, however the total computation time is very sensitive to this number and hence we strive to minimise the number of bins required. An accurate solution to the partial differential equation that describes the evolution of the electron distribution can be achieved with just 100 bins (see section 3.4.6). Therefore, we divide the electron distribution into 128 bins ($R_\gamma = 128$), and as this also provides sufficient resolution for the photon distribution over most of the typically required frequency ranges we also split this distribution into 128 bins ($R_\nu = 128$).

The logarithmic grids are calculated using the following formula to determine the value of each abscissa,

$$x_i = x_{\min} \left(\frac{x_{\max}}{x_{\min}} \right)^{\frac{i}{N}} \tag{3.4}$$

where i is the grid index, which obeys $0 \leq i \leq N$, and N is the number of bins (there are $N + 1$ abscissas). Bin widths are defined as,

$$\begin{aligned} \Delta x_i &= x_{i+1} - x_i \\ &= x_i \left[\left(\frac{x_{\max}}{x_{\min}} \right)^{\frac{1}{N}} - 1 \right], \end{aligned} \tag{3.5}$$

where $0 \leq i \leq N - 1$. The grid value for each bin is taken as the bin's mid-point in logarithmic space,

$$\begin{aligned} x_{i+\frac{1}{2}} &= 10^{\frac{1}{2} \log(x_{i+1}) + \frac{1}{2} \log(x_i)} \\ &= \sqrt{x_{i+1} x_i}, \end{aligned} \tag{3.6}$$

where $0 \leq i \leq N - 1$. The logarithmic mid-point is favoured other the mid-point in linear space, as it is more accurate at integrating the power-law distributions that the photon and electron distributions typically follow.

The photon distribution is calculated over a frequency grid, where $\nu_i = x_{i+\frac{1}{2}}$, and

$\Delta \nu_i = \Delta x_i$, for $x_{\max} = \nu_{\max}$ and $x_{\min} = \nu_{\min}$. The electron distribution is calculated over a relativistic energy grid for simulations where the electrons are distributed in significant numbers over a large range of energies. Hence, $\gamma_i = x_{i+\frac{1}{2}}$, and $\Delta \gamma_i = \Delta x_i$, for $x_{\max} = \gamma_{\max}$ and $x_{\min} = \gamma_{\min}$. In simulations where extra resolution is required at low energies the electron distribution is calculated over a relativistic momentum grid, p_i, with constant spacing, and converted back to relativistic energy values using the definitions of section 2.3. This effectively provides a relativistic energy grid that, in log-space, has smaller intervals at lower energies.

Other than the numerical solution to the kinetic equation, all numerical integration and differentiation in this model is calculated using simple first-order methods. For example, integration is discretised as

$$\int_{\gamma_{\min}}^{\gamma_{\max}} n_\gamma \mathrm{d}\gamma \Rightarrow \sum_{i=0}^{N-1} n_i \Delta \gamma_i. \tag{3.7}$$

By the nature of our grid designs, this numerical integration is effectively the composite trapezium method. Differentiation is less precise, because unlike $\Delta \gamma_i$, Δn_i is an unknown quantity and must be approximated by

$$\Delta n_i = \frac{n_{i+1} - n_{i-1}}{2}, \tag{3.8}$$

which is obviously undefined at the boundaries. In this case we assume that the gradient at each boundary bin is equal to the gradient at the adjacent bin.

3.4.2   Electron distributions

For simulations where we do not model the evolution of the electron distribution, the code can create an initial electron distribution for each cell that is either a thermal distribution, a non-thermal distribution (for example a power-law) or a hybrid of the two distributions. Using the equations of section 2.3 we initialise thermal electron distributions with

$$n_i = \begin{cases} \frac{n_e}{\theta_e K_2(1/\theta_e)} \gamma_i \sqrt{\gamma_i^2 - 1} \, \exp\left(-\frac{\gamma_i}{\theta_e}\right) & ; \quad p_e \geq 0.15, \\ \sqrt{\frac{2}{\pi}} \frac{n_e}{\theta_e^{3/2}} \gamma_i \sqrt{\gamma_i^2 - 1} \, \exp\left(\frac{1-\gamma_i}{\theta_e}\right) & ; \quad p_e < 0.15, \end{cases} \tag{3.9}$$

where the parameters are the electron number density, n_e [m^{-3}], and the electron temperature, $p_e = \sqrt{(\theta_e + 1)^2 - 1}$. For low temperature distributions the normalisation term

rapidly reaches numbers that are too high for a computer to calculate, whilst the exponential term decreases by the same factor. Therefore, we employ the approximation for the Bessel function in the $x \gg 1$ limit (Press *et al.* 1992) to calculate low temperature, $p_e < 0.15$, distributions. Hence, for distributions where $p_e < 0.15$, the total number density integrated over all energies differs, very slightly, from n_e.

In simulations where we model the evolution of the electron distribution no initial distribution is set, and a non-thermal distribution is continually injected. Due to the effect of Coulomb scattering over a sufficient length of time such a distribution will reach a thermal equilibrium. As non-thermal electrons are continually injected the resulting equilibrium will be a hybrid distribution. The injection of non-thermal power-law distributed electrons is described by

$$Q_i = Q_0 \gamma_i^{-q}, \tag{3.10}$$

$$n_i = n_i + Q_i \Delta t, \tag{3.11}$$

where Q_0 [m^{-3} s^{-1}] determines the injection rate, q gives the power-law index, and Δt is the simulation time-step.

3.4.3   Synchrotron radiation

The synchrotron emission photon distribution is calculated by solving the emission and absorption coefficients by numerical integration of the electron distribution and synchrotron emissivity function. Each cell is considered to be a homogeneous region and so the photon distribution may be calculated by using the homogeneous solution to the radiative transfer equation.

To calculate the synchrotron emissivity we divide the electron distribution into ultra-relativistic and mildly-relativistic energies with the divide arbitrarily chosen to be $\gamma = 2$. For the ultra-relativistic, $\gamma \geq 2$, electrons we adopt the tangled magnetic field solution of Ghisellini *et al.* (1988) given in equation 2.28. The numerous Bessel functions make the calculation of this function computationally intensive. Therefore, we pre-calculate a 'look-up table' for the synchrotron emissivity at the start of the simulation. To ensure a

high resolution for this look-up table we only store the values between $x_{min} < x < x_{max}$, divided into 10^4 bins of constant width in logarithmic space. For $x > x_{max}$ and $x < x_{min}$, we calculate the emissivity in real-time using the series expansion of the function in the $x \gg 1$ and $x \ll 1$ limits respectively. Hence, we set the values of x_{max} and x_{min} to the point where the limiting functions are approximately 99% accurate, giving $x_{min} = 0.001$ and $x_{max} = 10.0$. The synchrotron emissivity of ultra-relativistic electrons, given by equation 2.28, in the $x \ll 1$ limit is

$$j(x) = 5.27 \times 10^{-25} B x^{\frac{1}{3}} \; [\text{W Hz}^{-1}] \tag{3.12}$$

and in the $x \gg 1$ limit,

$$j(x) = 3.57 \times 10^{-26} B x e^{-2x} \; [\text{W Hz}^{-1}]. \tag{3.13}$$

Synchrotron emission from mildly-relativistic, $\gamma < 2$, electrons is poorly understood, and hence there is less agreement upon the best equation to model the emissivity. Equation 2.28 was derived in the ultra-relativistic limit and for $\gamma \lesssim 2$ fails to produce the correct emissivity. Following Ghisellini et al. (1988) we adopt equation 2.30, which provides a good description of the synchrotron emissivity, at frequencies where $\nu \gg \nu_g$, from $\gamma < 2$ electrons. However, this function fails to accurately model the emission at frequencies $\nu \lesssim \nu_g$ from $\gamma < 2$ electrons, and therefore we assume that for $\nu \leq \nu_g$ the emissivity from $\gamma < 2$ electrons is zero.

The emission coefficient is simply calculated by numerically integrating equation 2.8. There are various ways of calculating the absorption coefficient (Ghisellini & Svensson 1991) and at present our model uses equation 2.32. We rewrite this equation as

$$k_\nu = \frac{1}{8\pi m_e \nu^2} \int_{\gamma_{min}}^{\gamma_{max}} j_\nu(\gamma) \left[n_\gamma \left(\frac{\gamma^2 + p^2}{\gamma p^2} \right) - \frac{dn_\gamma}{d\gamma} \right] d\gamma, \tag{3.14}$$

and then numerically differentiate the electron distribution and integrate the equation in the manner described in section 3.4.1. The synchrotron radiation photon distribution emitted by the cell can then be calculated using equation 2.16, with the path-length equal to the full width of the cell. This photon distribution is equally distributed to all six faces of the cell. The resulting spectrum for an arbitrary power-law distribution of electrons contained within in a single cell is shown in Figure 3.6.
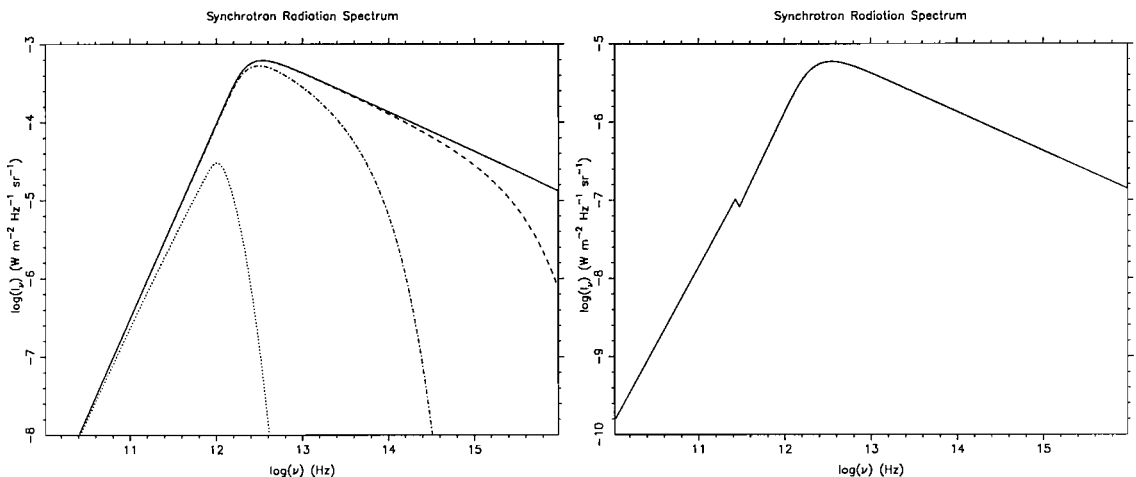
Figure 3.6: (a - left panel) The spectrum of self-absorbed synchrotron radiation from a power-law electron distribution with $q = 2$ and differing maximum energies, as calculated by the HEART code for frequencies where $\nu \gg \nu_g$. The solid line shows the case where the critical frequency, ν_c, of the highest energy electron is much greater than the frequency where the optical depth is unity. In this case the HEART code accurately reproduces the correct power-law slopes of 2.5 in the optically thick region of the spectrum and -0.5 in the optically thin region. The other lines show the effect of lowering the highest electron energy and hence decreasing the critical frequency; $\nu_c = 10^{18}$ Hz (solid line), 10^{16} Hz (dashed line), 10^{14} Hz (dash-dotted line), and 10^{12} Hz (dotted line). (b - right panel) In this instance the gyro frequency has been increased to $\nu_g = 5 \times 10^{11}$ Hz illustrating the inaccuracy of the HEART code below this frequency.

The peak of the radiation distribution occurs at the frequency, $\nu_{\tau=1}$, for which the electron distribution has an optical depth of unity. The electron distribution is optically thin to higher frequencies, where it emits synchrotron radiation with a power-law spectrum of index $(1 - q)/2$. At lower frequencies, to which the electron distribution is optically thick, more radiation is absorbed producing a power-law spectrum of index 2.5. The optically thin region of the radiation spectrum eventually ends with an exponential cut-off when the critical frequency, ν_c, of the highest energy electron is reached. The spectrum is accurately reproduced by the HEART code for frequencies much greater than the the gyro frequency, ν_g. Synchrotron emission at frequencies where $\nu \approx \nu_g$ is poorly modelled by the HEART code, as shown in Figure 3.6b. At frequencies just above the gyro frequency the slope flattens to a value of 2.3, and below the gyro frequency the spectrum is accurately modelled with a slope of 2.0 as is appropriate for optically thick radiation in this regime.

Assuming that the emissivity from mildly relativistic electrons below the gyro frequency is zero creates a kink in the spectrum at the gyro frequency. Therefore, the HEART code is only accurate when $\nu_g \ll \nu_{\tau=1}$.

We also model the effect of synchrotron absorption upon any incoming radiation that enters the cell. The incoming photon distribution for each direction is absorbed according to,

$$I_\nu = I_\nu e^{-\tau_\nu}. \tag{3.15}$$

3.4.4   Compton scattering

Numerical calculations of Compton scattering, being quite computationally expensive, were traditionally simplified by using a delta function to describe the photon redistribution at low energies. This agrees with the result of a derivation of Thomson-limited Compton scattering that oversimplifies the problem by assuming the incoming photon field is isotropic in the rest-frame rather than the observer's frame (see e.g. Rybiciki & Lightman 1979, Chiaberge & Ghisellini 1999). Furthermore, the photon redistribution function for scattering off ultra-relativistic electrons is also flawed in such derivations, giving a mean frequency gain that is too small. As mentioned in Coppi & Blandford (1990), in the case of high optical depth plasmas where multiple scatterings occur, the use of a delta function to describe low energy scatters results in a bumpy superposition of each successive scatter whereas in reality we observe a smooth power-law. Therefore, when modelling Compton scattering by mildly-relativistic electrons we must use the correct dispersion profile.

Traditionally Klein-Nishina effects have also been simplified by assuming the scattering cross-section is zero at frequencies beyond the Klein-Nishina limit ($h\nu \gtrsim \gamma m_e c^2$). Again in the case of cool plasmas with a high optical depth this assumption is flawed. Photons scattered to the limiting frequency pile-up because they cannot be scattered further resulting in an unphysical spike in the radiation spectrum. Therefore, it is important to model the Klein-Nishina effect as a gradual decay in the number of scatters when the photon energy begins to exceed the electron energy.

Unlike most other models we evolve a trans-relativistic electron distribution and thus need to treat the mildly-relativistic electrons in the same manner as the ultra-relativistic electrons. Previous models (e.g. Coppi 1992, 1999) would assume that the lower energy electrons were quickly thermalised and hence calculated the expected Compton scattered photon distribution analytically from a Maxwell-Boltzmann distribution. By using the Kompaneets equation the effect of multiple scatterings could be modelled (see Figure 3.8 for a comparison between the spectra produced by this model and the spectra produced by the HEART model). However, as we model truly arbitrary distributions that are free to evolve we cannot use analytical methods. Instead, we calculate the photon redistribution function for all electron energies and convolve it with the seed photon distribution and the electron distribution to obtain the resulting spectrum for a single scatter. Although a single cell in our model can only scatter a given photon once, multiple scatters are still possible as they are a natural result of our radiative transfer model for multi-cell simulations. Hence, the total number of scatters modelled is limited by the number of cells along the photon's escape path ($N_{\mathrm{scat,max}} \approx R_{\mathrm{cells}}$) in our model. This does not represent a problem if we ensure that the optical depth of any one cell does not exceed unity, and hence modelling multiple scatters within the path-length defined by a single cell is not required.

For the reasons stated above we have elected to model Compton scattering using the exact photon redistribution function of Jones (1968). Coppi & Blandford (1990) have presented a corrected version of this function that may be calculated by numerical integration

$$xF_{\mathrm{C}}(x) = 3\sigma_{\mathrm{T}}\, x^2 \int \frac{\mathrm{d}P}{\mathrm{d}z} dz. \tag{3.16}$$

The function $\mathrm{d}P/\mathrm{d}z$ may be found in appendix A as it is too lengthy to include here. Unfortunately this integration is quite computationally intensive and so to shorten the execution time of the simulation we load a pre-calculated table of the exact photon redistribution function. This, of course, limits our accuracy as the simulation energy grids are chosen at run-time and may not necessarily be identical to those of the pre-calculated table. Furthermore, photon redistribution functions typically only depend upon one parameter, the frequency gain $\rho = \nu_1/\nu_0$; trans-relativistic ones also depend upon the electron energy. As the exact photon redistribution function also considers Klein-Nishina effects there is a further parameter dependence upon the seed photon frequency and so a three-dimensional

table of values must be calculated. For such a table to be accurate for any arbitrarily chosen set of photon and electron energy grids it would need to be exceptionally large. Therefore, to avoid this problem we narrow the range of values for each parameter modelled with the exact formula by calculating the redistribution function look-up table at run-time using the limiting case formulae given in section 2.8.

An IDL program is used to calculate the exact photon redistribution function table, and this performs the numerical integration of the function dP/dz (given in appendix A) with the QROMB function of Press et al. (1992). This function becomes very time consuming to integrate at high γ, as it behaves like the relativistic aberration function, and also at low frequencies. Thankfully in both of these cases we reach the ultra-relativistic and Thomson limits where there are only slight changes to the dispersion profile that are adequately modelled by the limiting case formulae. The three parameters used in calculating this table are $k = \gamma h\nu_0/m_e c^2$, which gives the proximity to the Klein-Nishina regime, p, the electron's energy given in units of relativistic momenta, and $x = \rho/(p+\gamma)^2$ the frequency gain as a fraction of the maximum frequency gain. The dispersion profile is calculated in the range $0 \leq x \leq 1$ (Klein-Nishina effects only reduce the maximum scattered frequency to less than $x = 1$), and is divided into 256 linear grid points to minimise under-scatter. The other two parameters need not be so accurately modelled as the dispersion profiles change slowly, and so they are both modelled by logarithmic grids of just 64 points, with the ranges $0.08 \leq p \leq 10$ and $10^{-2} \leq k \leq 10^3$. Beyond the upper limit for k, Klein-Nishina effects are assumed to reduce the scattering cross-section to zero, and below the lower limit the Thomson limit assumption applies so we can calculate the photon redistribution function for the run-time energy grids using the formula of Ensslin & Kaiser (2000) (see equation 2.50). Similarly beyond the upper limit for p the ultra-relativistic assumption applies and we can calculate the photon redistribution function for the run-time energy grids using the formula of Jones (1968) (see equation 2.54).

The lower limit of p is thus also the lowest value of p that the HEART code can model. The decision regarding the lowest energy electron to model is in part influenced by the requirement to prevent the possibility of over-scatter, where the width of the dispersion profile becomes smaller than the energy grid's bin width. However, such over-scattering

is not a significant problem for the photon spectrum as at these energies photons are only scattered in direction without an energy gain. The more significant constraint on the choice of the lowest energy electron results from a feature of our numerical model. In a simulation where electrons only lose energy due to inverse Compton scatterings there is a minimum energy electron that can just increase the energy of a photon enough to move it into the next frequency bin. Therefore, beneath that energy electrons cannot lose energy and hence electrons would pile-up at this minimum energy. Therefore, for a given photon frequency resolution (R_ν) there is a corresponding minimum electron energy that may be modelled.

The photon re-distribution function for the run-time energy grids is calculated at the start of every simulation and stored in an $R_\nu \times R_\nu \times R_\gamma$ sized array to save calculation time during time loops. For all frequencies and $\gamma \geq 10$ we use the ultra-relativistic solution of equation 2.54, and for $\gamma < 10$ and frequencies where $k \leq 10^{-2}$ we use the Thomson limit solution of equation 2.50. Only for scatterings off electrons with $\gamma < 10$ in the Klein-Nishina regime where $k > 10^{-2}$ do we rely upon sampling the pre-calculated table of the exact photon redistribution function values. Although scatterings from low energy electrons are generally weak, when multiple scatterings occur they become significant and so it is important to model the decline in scattering due to the Klein-Nishina effect.

The scattered photon distribution is calculated from equation 2.49 by numerically integrating in the standard manner

$$I_{\nu_1} = w \int_{\nu_0,\min}^{\nu_0,\max} \int_{\gamma_{\min}}^{\gamma_{\max}} n_\gamma \, x F_C(x) \mathrm{d}\gamma \, I_{\nu_0} \frac{\mathrm{d}\nu_0}{\nu_0}, \tag{3.17}$$

with $I_{\nu_0} = I_{\mathrm{cell}}$, i.e. a combination of the intrinsic emission, injected radiation, and any radiation that passes through the cell. After scattering, the directionality of the incoming photons is lost, so their contribution to the IC emission is distributed evenly to all six faces of the cell together with the scattered intrinsic emission. Therefore, the model does account for the possibility of individual photons taking 'random walks' through the plasma and hence having effective path-lengths greater than the dimensions of the plasma.

We also account for the loss of seed photons by subtracting the integrated total contri-

bution from each seed frequency, ν_0, over all scattered frequencies, ν_1,

$$I'_{\nu_0} = I_{\nu_0} - w \int_{\nu_{1,\min}}^{\nu_{1,\max}} \int_{\gamma_{\min}}^{\gamma_{\max}} n_\gamma \, x F_C(x) \mathrm{d}\gamma \, I_{\nu_0} \frac{\mathrm{d}\nu_1}{\nu_1}. \tag{3.18}$$

Each seed photon distribution must be modified individually, hence in the above equation $I_{\nu_0} = I_{\mathrm{Int}}, I_1, I_2$ etc. To ensure that the total photon number is conserved we can perform a test upon the code, by injecting some seed photons with a blackbody distribution given by equation 2.19. Calculating the scattered radiation spectrum from a non-thermal, power-law electron distribution, with various values of the optical depth, τ_T, and from thermal electrons at various temperatures produced the spectra shown in Figure 3.7. Using equation 2.47 to calculate the number of photons from the specific intensity we find that the ratio of photons that form the scattered spectrum over the injected seed photons is equal to the optical depth τ_T, as expected, and that equation 3.18 correctly subtracts
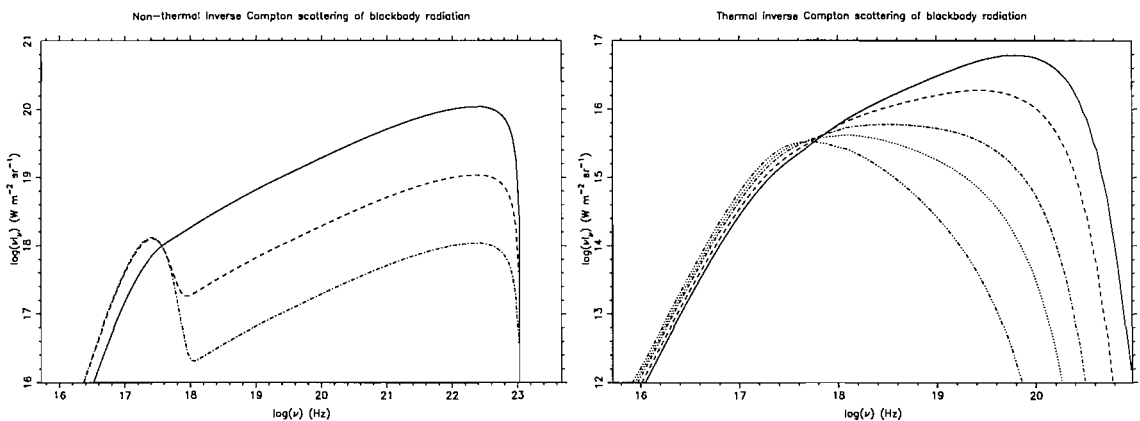

Figure 3.7: (a - left panel) The spectral energy distribution of non-thermal, inverse Compton scattered blackbody radiation for different optical depths, $\tau_T = 1.0$ (solid line), 0.1 (dashed line), and 0.01 (dot-dashed line). The optical depth determines the fraction of seed photons scattered, so for $\tau_T = 1.0$ all of the blackbody photons have been inverse Compton scattered. The slope of the inverse Compton radiation depends upon the electron distribution's power-law index as $(3 - q)/2$, therefore the spectral slope becomes negative for $q > 3$. The sharp high-frequency cut-off is due to the Klein-Nishina effect. (b - right panel) The spectral energy distribution of thermal inverse Compton multiple scattered blackbody radiation from electron distributions with an optical depth of $\tau = 5$ and different temperatures, $kT_e = 200$ keV (solid line), 150 keV (dashed line), 100 keV (dot-dashed line), 75 keV (dotted line), and 50 keV (double-dot dashed line). The power-law like scattered spectra are the result of many superpositions of scattered blackbody-like spectra.
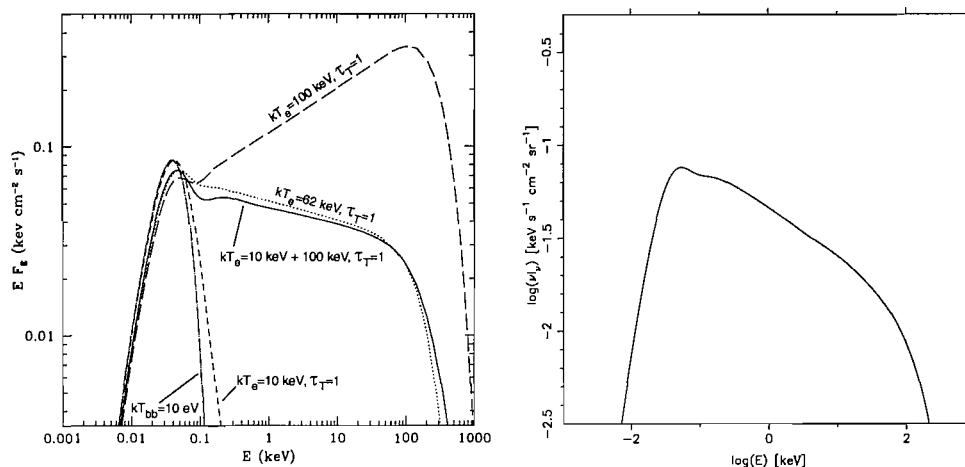
Figure 3.8: A comparison between the resulting equilibrium spectra from EQPAIR (left panel, dotted line only, taken from Coppi 1999) and HEART (right panel) for a model of $kT_{BB} = 10$ eV blackbody radiation passing through a spherical plasma of thermal electrons at a temperature $kT_e = 62$ keV with an optical depth of $\tau = 1$. This illustrates a difference in the resulting spectrum between these two alternative models for the process of multiple Compton scattering in thermal electron distributions. The HEART model is created using a spherical geometry, with a spatial resolution of $R_{cells} = 16$, and with the blackbody seed photons being injected isotropically in the centre cell.

the scattered fraction of photons from the seed photon distribution. Therefore, for an optical depth of unity, all of the seed photons are scattered. For optical depths greater than unity, equation 3.17 will scatter more photons than actually exist in the seed photon distribution, and so this equation is not physically valid for $\tau_T > 1$. In practice this is not a problem, because to accurately model radiative transfer each cell would need to be optically thin. However, it is important for the purposes of photon conservation that cells never have an optical depth greater than unity.

A comparison between the thermal Comptonisation spectra produced by HEART to that produced by the EQPAIR model (Coppi 1999) is shown in Figure 3.8. The difference between these two spectra can be explained by the HEART model not simulating the same number of high-order scatterings (i.e. multiple scatters of the same photon many times) as the EQPAIR model, resulting in a weaker flux at higher energies (also, to a lesser extent, a slightly stronger flux at lower energies). This is due to our limitation of modelling only one scattering event per cell, hence the requirement of many cells to accurately model the strength of the high-order scattering component to the spectrum. Therefore, we would

need a substantially higher resolution simulation to accurately recreate the EQPAIR model spectrum. Furthermore, it is difficult in a multi-cell model to treat the injection of seed photons into the spherical plasma in an identical way to a single cell model such as EQPAIR. The geometry of injection model will also affect the number of higher-order scatters seen in the spectrum, and so is another factor that has contributed to the differences between these two model's Comptonisation spectra. Obviously, the HEART model will recreate the greater number of higher-order scatters if the temperature or optical depth of the plasma were increased. However, the resulting overall spectral shape would differ.

We can also model the spectral energy distribution produced by synchrotron self-Compton (SSC) emission, where a power-law electron distribution produces the seed photons in the form of synchrotron radiation and then scatters this radiation in the inverse Compton process. Observations of SSC emission reduce the number of free parameters for the synchrotron emission models, because the ratio of the IC radiation luminosity to the synchrotron radiation luminosity is uniquely determined by the strength of the magnetic field.

3.4.5 Bremsstrahlung

As we aim to make no distinction between thermal electrons and non-thermal electrons it is important that we treat the bremsstrahlung emission in the same way for all energies. Following Svensson (1982) we use separate cross-sections for the high-energy and the low-energy electrons and apply correcting factors to smooth the transition between the two functions. For $\gamma < 2$ electrons we use the mildly-relativistic formula of equation 2.60 multiplied by a factor of γ and for $\gamma \geq 2$ electrons we use the ultra-relativistic formula of equation 2.59 together with a correction factor of 1.27 to ensure that the transition between the two functions is smooth. A table of cross-section values is calculated at the start of every simulation for the run-time energy grids, from which the emission and absorption coefficients can be calculated for the current electron distribution by numerically integrating equations 2.63 and 2.64 respectively.

The specific intensity of bremsstrahlung radiation emitted by each cell is calculated

in precisely the same manner as with synchrotron radiation using the radiative transfer solution assuming the cell is homogeneous, and distributing the radiation equally to all six faces of the cell. Transiting radiation is also absorbed as,

$$I_\nu = I_\nu e^{-\tau_{\nu,B}}. \tag{3.19}$$

3.4.6 Electron distribution evolution

The kinetic equation may be written in the form of a Fokker-Planck equation, the numerical solutions to which are well studied (Park & Petrosian 1996). In the Fokker-Planck form the kinetic equation is written as,

$$\frac{\partial n_\gamma}{\partial t} = \frac{\partial}{\partial \gamma}\left[B(\gamma)n_\gamma + C(\gamma)\frac{\partial n_\gamma}{\partial \gamma}\right] + Q(\gamma) - \frac{n_\gamma}{T(\gamma)}, \tag{3.20}$$

where

$$B(\gamma) = \Lambda(\gamma) - H(\gamma)\left(\frac{\gamma^2 + p^2}{\gamma p^2}\right) \tag{3.21}$$

and $C(\gamma) = H(\gamma)$. To calculate the co-efficients we numerically integrate equations 2.67 and 2.69, where specific intensity in equation 2.69 is that of all of the radiation that may be absorbed, i.e. $I_\nu = I_{\text{cell}}$ (after scattering). The Coulomb scattering co-efficients $a(\gamma)$ and $D(\gamma)$ from equation 2.70 can be converted to this form as

$$B(\gamma) = \frac{\mathrm{d}D(\gamma)/2}{\mathrm{d}\gamma} - a(\gamma) \tag{3.22}$$

$$C(\gamma) = D(\gamma)/2. \tag{3.23}$$

A table of values for the χ and ζ functions used to calculate these co-efficients are calculated at the start of every simulation for the run-time energy grid. Both co-efficients are then calculated on every time loop by numerical integration over the electron distribution in the usual manner.

When modelling the electron distribution evolution it is vital that the thermalisation process is accurately modelled with sufficient time-resolution. Therefore, so as not to place further constraints on the spatial resolution of the model, we separately re-calculate the Coulomb co-efficients and solve the Fokker-Planck equation as many times as required within a radiative transfer time-step to ensure that $\Delta t \leq t_c/20$. Note that the radiative

co-efficients are not recalculated for each sub-time-step, but are always included in the kinetic equation that is being solved.

We employ the Chang-Cooper method (Chang & Cooper 1970) to solve the Fokker-Planck equation, which discretises the equation as

$$\frac{n_i^{k+1} - n_i^k}{\Delta t} = \frac{F_{i+1/2}^{k+1} - F_{i-1/2}^{k+1}}{\Delta \gamma_i} - \frac{n_i^{k+1}}{T_i} + Q_i \tag{3.24}$$

where k is the time index, i is the energy index, and the flux, F, is given by the bracketed term in equation 3.20. The Chang-Cooper method differs from other methods in the way the flux is discretised, by using centred difference on the diffusive term (the term with $C(\gamma)$) and weighted difference on the advective term $(B(\gamma))$, giving

$$F_{i+1/2}^{k+1} = (1 - \delta_{i+1/2})B_{i+1/2}n_{i+1}^{k+1} + \delta_{i+1/2}B_{i+1/2}n_i^{k+1} + C_{i+1/2}\frac{n_{i+1}^{k+1} - n_i^{k+1}}{\Delta \gamma_{i+1/2}}, \tag{3.25}$$

$$= \frac{C_{i+1/2}}{\Delta \gamma_{i+1/2}}\left[W_{i+1/2}^+ n_{i+1}^{k+1} - W_{i+1/2}^- n_i^{k+1}\right], \tag{3.26}$$

where the weighting function is given by,

$$\delta_i = \frac{1}{w_i} - \frac{1}{\exp(w_i) - 1}, \tag{3.27}$$

and

$$w_i = \frac{B_i}{C_i}\Delta \gamma_i, \tag{3.28}$$

$$W_i^{\pm} = W_i \exp\left(\pm\frac{w_i}{2}\right), \tag{3.29}$$

$$W_i = \frac{w_i/2}{\sinh(w_i/2)}. \tag{3.30}$$

In this method the boundary condition is that there should be no flux across the energy grid boundaries, hence, for $i = 0, ..., I$,

$$F_{I+1/2}^{k+1} = F_{-1/2}^{k+1} = 0. \tag{3.31}$$

This should conserve particle number in the absence of any sinks or sources, and guarantees positive solutions. In light of the boundary condition we may re-write equation 3.24 as a tridiagonal system of linear equations

$$\begin{cases} -a_i n_{i-1}^{k+1} + b_i n_i^{k+1} - c_i n_{i+1}^{k+1} &= r_i, \\ a_0 = c_I &= 0, \end{cases} \tag{3.32}$$

where

$$a_i = \frac{\Delta t}{\Delta \gamma_i} \frac{C_{i-1/2}}{\Delta \gamma_{i-1/2}} W^-_{i-1/2}, \qquad (3.33)$$

$$c_i = \frac{\Delta t}{\Delta \gamma_i} \frac{C_{i+1/2}}{\Delta \gamma_{i+1/2}} W^+_{i+1/2}, \qquad (3.34)$$

$$b_i = 1 + \frac{\Delta t}{\Delta \gamma_i} \left[\frac{C_{i-1/2}}{\Delta \gamma_{i-1/2}} W^+_{i-1/2} + \frac{C_{i+1/2}}{\Delta \gamma_{i+1/2}} W^-_{i+1/2} \right], \qquad (3.35)$$

$$r_i = \Delta t Q_i + n^k_i. \qquad (3.36)$$

Note that for the purpose of clearly illustrating the numerical method all indices used here define grid abscissas rather than bin centres as used in the code. To solve the tridiagonal system of linear equations we use the TRIDAG routine of Press et al. (1992). Park & Petrosian (1996) demonstrate that the solution is very accurate for logarithmic energy grids containing just 100 bins, as long as $\Delta \gamma_i / \gamma_i \ll 1$, which can be calculated using equation 3.5. We have tested our implementation of this method for non-thermal electron distributions emitting self-absorbed synchrotron emission, and find that the electron distribution always remains positive with the total number of electrons being conserved.

When there are no heating effects the kinetic equation will pile-up electrons in the lowest energy bin due to the no flux boundary condition. This physical inaccuracy may eventually be reflected in the radiation spectrum. Therefore, after solving the Fokker-Planck equation, if the number of electrons in the lowest energy bin becomes significantly larger than that in the adjacent bin, then it is reset to be equal to the amount in the adjacent bin. In effect this is a sink process, removing electrons from the distribution that have cooled below the lowest energy modelled. However, this is only ever important in simulations where radiative losses completely dominate Coulomb scattering and thus no thermal component can be maintained in the electron distribution.

It is important to check that energy is conserved in its transfer between the photon and electron distributions. The total energy available for transfer in the electron distribution is the kinetic energy, $(\gamma - 1)m_e c^2$, integrated over the distribution and volume. Integrating the specific intensity over frequency and volume, $(4\pi/c)I_\nu dV d\nu$, will give the total photon energy. The total energy loss from the electron distribution after having been modified by the kinetic equation should equal the total energy in the photon distribution. We find this is true when the only source of energy transfer is through self-absorbed synchrotron

radiation and the time step is smaller than the cooling timescale. Obviously, the accuracy of the numerical method decreases with larger time steps, which can result in a significant departure from energy conservation.

3.5   Simulation overview, parameters and constraints

The exact ordering of the physical processes on each time step of the simulation is not arbitrary, and here we shall explain the logic behind our chosen order of events. With the three-dimensional configuration of cells created, and all of the electron distributions initialised, the simulation can start performing the sequence of events for each time loop. First, the intrinsic photon distribution is calculated from the synchrotron and/or bremsstrahlung emission of the current electron distribution and the transiting radiation is absorbed by the appropriate amount. Next, the total photon distribution within the cell, I_{cell}, is then determined and Compton scattered with appropriate adjustments to each of the seed distributions to account for lost photons. The intrinsic plus scattered radiation is then distributed equally into each transiting photon distribution in all 4π steradian directions.

Whenever photons are created, absorbed, or scattered there is a corresponding energy change to the electron distribution, and so for each of the above processes the kinetic equation co-efficients are calculated. The thermalisation effect of Coulomb scattering within the electron distribution is then calculated by solving the kinetic equation multiple times with time-steps of $\Delta t \leq t_{\mathrm{c}}/20$ until the time of a radiative transfer time-step has been completed. Once the changes to the photon distributions and the electron distribution for the time-step have been calculated, the surface flux is removed from the outer cells before the transiting photon distributions are transferred to the next cell. The simulation continues until a specified end-time, t_{max}, is reached, or until the photon distribution has reached an equilibrium (i.e. the rate of change of the entire distribution with time becomes negligibly small).

The model parameters depend upon the type of simulation that is being performed, but a complete list, for any given geometry, is given in Table 3.1. Presently the code

Settings		
R_{cells}		The spatial resolution of the model, or the number of cells along the full span of the model across the x-axis. N_{cells} is the total number of cells in the simulation and hence depends upon both R_{cells} and the chosen geometry.
R_ν		Photon distribution energy resolution, or the number of frequency bins.
R_γ		Electron distribution energy resolution, or the number of energy bins.
ν_{min} & ν_{max}	[Hz]	The lowest and highest photon frequencies to be modelled.
t_{max}	[s]	The maximum duration of the simulation (if no equilibrium is first achieved).
General parameters		
D	[m]	The distance between the plasma and the observer.
w	[m]	Total width of the plasma along the x-axis.
Models with the injection of non-thermal electrons		
γ_{min}		The lowest electron energy injected (normally unity).
γ_{max}		The highest electron energy to be injected.
Q_0	$[\text{m}^{-3}\ \text{s}^{-1}]$	Injection rate of non-thermal electrons.
q		Power-law index for non-thermal electrons.
t_{esc}	[s]	The electron escape time-scale for injection models.
t_{inj}	[s]	The duration of electron injection (normally infinity).
Magnetic models		
B	[T]	Magnetic field strength.
Models with non-evolving electron distributions		
n_e	$[\text{m}^{-3}]$	The number density of thermal electrons.
kT_e	[keV]	The thermal electron distribution temperature.

Table 3.1: HEART model parameters. All time quantities are typically scaled to t_{cross}.

initialises all cells to the same values, but can be easily adapted to model spatial variations of any of the parameters. To accurately model the electron distribution evolution it is important to ensure the photon distribution frequency range encompasses every frequency with significant photon emission.

Limitations on the energy, and spatial (and hence time) resolution place constraints upon the model. For most models the constraints on the energy resolution of the electron and photon distributions should be well met by the default settings. However, the spatial resolution will be freely varied to minimise the computation time required by a simulation. Greater spatial resolution improves the ability of a radiative transfer model to accurately model the varying path-lengths of light through a given object's geometry. However, accurate modelling of plasma geometries is not likely to be a major requirement of our simulations. The main advantage of greater spatial resolution for the HEART code lies in its ability to model inhomogeneous plasmas. Therefore, the required spatial resolution is dependent upon the significance of variations in the electron distributions within the plasma.

An initially homogeneous plasma will always become inhomogeneous through the emission of radiation, which cools the electron distribution. Therefore the first spatial resolution constraint is that $\Delta t \lesssim t_{cool}(\gamma_{max})$, where

$$t_{cool}(\gamma_{max}) = \frac{\gamma_{max}}{\Lambda(\gamma_{max})} \qquad (3.37)$$

and $\Lambda(\gamma_{max})$ is given by equation 2.67. It is of even greater importance to ensure this constraint is met for the purpose of having time steps sufficiently small enough to ensure that the kinetic equation conserves energy. The second constraint is that the optical depth to inverse Compton scattering must be less than unity, $\tau_T < 1$. Finally, if any inhomogeneity is imposed on the initial electron distributions such as making the normalisation radius dependent, $n_e = n_e(r)$, then this will add further constraints upon the spatial resolution.

3.6   Illustrations of HEART capabilities

3.6.1   The radiative transfer approach to thermal Compton scattering

To show how the HEART code models the process of multiple Compton scatters upon a photon distribution we present here a simple one-dimensional configuration of cells, each containing an identical thermal distribution of electrons of fixed temperature, as illustrated in Figure 3.9. For clarity, energy losses upon these thermal electron distributions are not modelled. A seed photon distribution, I_{seed}, is injected into the first cell where it undergoes first order scattering, whereby a fraction of the total number of photons, given by the optical depth of the cell, $\Delta\tau_{\text{T}}$, gain energy due to a single scattering event. The photon distribution is transferred to the next cell, now consisting of the remaining unscattered seed photons and the fraction (one sixth) of the photons that were scattered along the direction of propagation of the seed photons. In the second cell the same fraction of photons will be scattered, some of which may be the photons that were scattered in the previous cell and hence a second order of scattering is modelled. When the radiation finally escapes from the outermost cell it will have undergone up to N orders of scattering, where $N = R_{\text{cells}}$, or the total number of cells in the case of this one-dimensional example.

Figure 3.10a illustrates this evolution process for a blackbody seed photon distribution, with temperature, $kT_{\text{BB}} = 100$ eV, where each cell contains a thermal distribution of electrons with optical depth, $\Delta\tau_{\text{T}} = 0.1$, at a temperature of $kT_{\text{e}} = 150$ keV, for a total optical depth of $\tau_{\text{T}} = 5$ (hence, $R_{\text{cells}} = 50$). This figures reveals the photon distribution as it emerges from each individual cell and illustrates how the superposition of multiple orders of scattering leads to a power-law like scattered distribution (e.g. Coppi 1992).
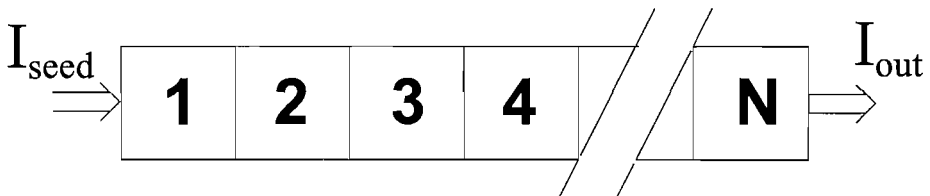
Figure 3.9:   A one-dimensional configuration of cells to illustrate the process of multiple Compton scatters in the HEART model.

This figure also illustrates the eventual loss of seed photons, and also of low-energy, singly scattered photons as the optical depth increases. Furthermore, the photon distribution does not reach an equilibrium immediately after emerging from the final cell. This is because photons may be backscattered multiple times before finally escaping – effectively increasing the path-length. Hence, the last spectrum shown in Figure 3.10a is the equilibrium emerging photon spectrum as found after two light-crossing time-scales, $t = 2t_{\rm cross}$ (as defined in section 3.2).

It is vital that we sub-divide the modelled plasma into optically thin cells, with $\Delta\tau_{\rm T} \leq$ 0.1. To illustrate this consider a plasma of unit optical depth modelled by a single cell and as modelled by ten optically-thin cells. In the former case every seed photon will be scattered, but only once, hence only first order scattering is modelled. In the later
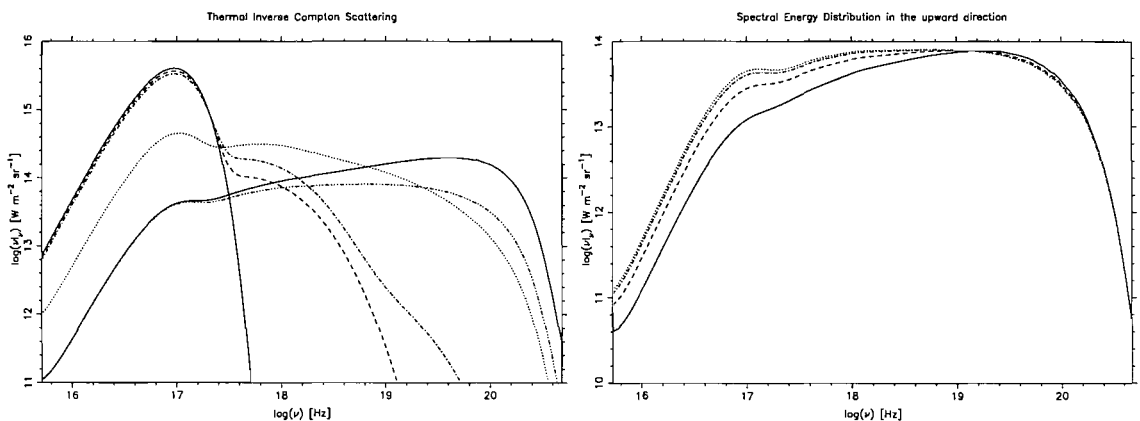

Figure 3.10: (a - left panel) Following the evolution of a photon spectrum starting from an injected blackbody seed distribution (solid line) as it passes through successive cells, undergoing inverse Compton scattering in each. The plasma is divided into 50 cells, and is modelled as consisting of a thermal distribution of electrons at a temperature of $kT_{\rm e} = 150$ keV with a total optical depth of $\tau_{\rm T} = 5$. The solid line gives the initial blackbody seed distribution, dashed line the first order of scattering as the radiation leaves the first cell, dot-dashed line the second order after passing through the second cell, dotted line the photon distribution at the half-way point, triple-dot-dashed line the emerging photon distribution as it leaves the final cell. Finally the second solid line gives the equilibrium photon distribution as found emerging from the cell after $t = 2t_{\rm cross}$. (b - right panel) The emerging photon spectrum from a $\tau_{\rm T} = 5$ plasma as modelled with $\Delta\tau_{\rm T} = 0.5$ (solid line), 0.25 (dashed line), 0.1 (dot-dashed line), and 0.05 (dotted line).

case every seed photon will not necessarily be scattered, and some photons may undergo multiple scatterings (up to ten in this case) before escaping the plasma, resulting in a very different spectrum. This effect is illustrated in Figure 3.10b for the more common scenario of a plasma with optical depth $\tau_T = 5$ and temperature $kT_e = 150$ keV being modelled with individual cell optical depths varying between $0.05 \leq \Delta\tau_T \leq 0.5$.

3.6.2   The evolution of a hybrid electron distribution

Another unique feature of the HEART code is its ability to model the evolution of a truly arbitrary electron distribution over a trans-relativistic range of energies as it comes into equilibrium with the pervading photon distribution. When simulating the evolution of an electron distribution we begin by continually injecting a non-thermal, power-law electron distribution. Coulomb scattering within the distribution will lead to the lower-energy electrons thermalising, and hence a hybrid distribution emerges consisting of thermal electrons at low-energies and non-thermal power-law distributed electrons at high-energies. Thus, as the thermal distribution is a product of the non-thermal distribution we cannot directly choose its temperature and density. Rather we alter the conditions to which the non-thermal distribution is subject in such a way that we obtain the desired thermal distribution. Therefore, in this section we demonstrate how the HEART model parameters affect the final equilibrium temperature and density of the evolved thermal component.

The temperature of the thermal component depends both upon the balance between the average energy of the injected non-thermal electrons and the strength of radiative losses. Therefore, we begin by demonstrating the effect of the parameters Q_0, q, t_{esc}, and γ_{min} (see Table 3.1 for definitions) upon the equilibrium temperature of the thermal distribution in the absence of any radiative losses. Figure 3.11 shows example equilibrium hybrid electron distributions for a sample of parameter values after they have thermalised through Coulomb scattering. Note that these are simple, single cell models.

Increasing the minimum energy of the injected electron distribution has the greatest effect upon the evolved thermal distribution, increasing the temperature from $kT_e = 130$ keV to $kT_e = 225$ keV with a $\gamma_{min} = 1.2$, and $kT_e = 525$ keV with a $\gamma_{min} = 2.0$ (Fig-
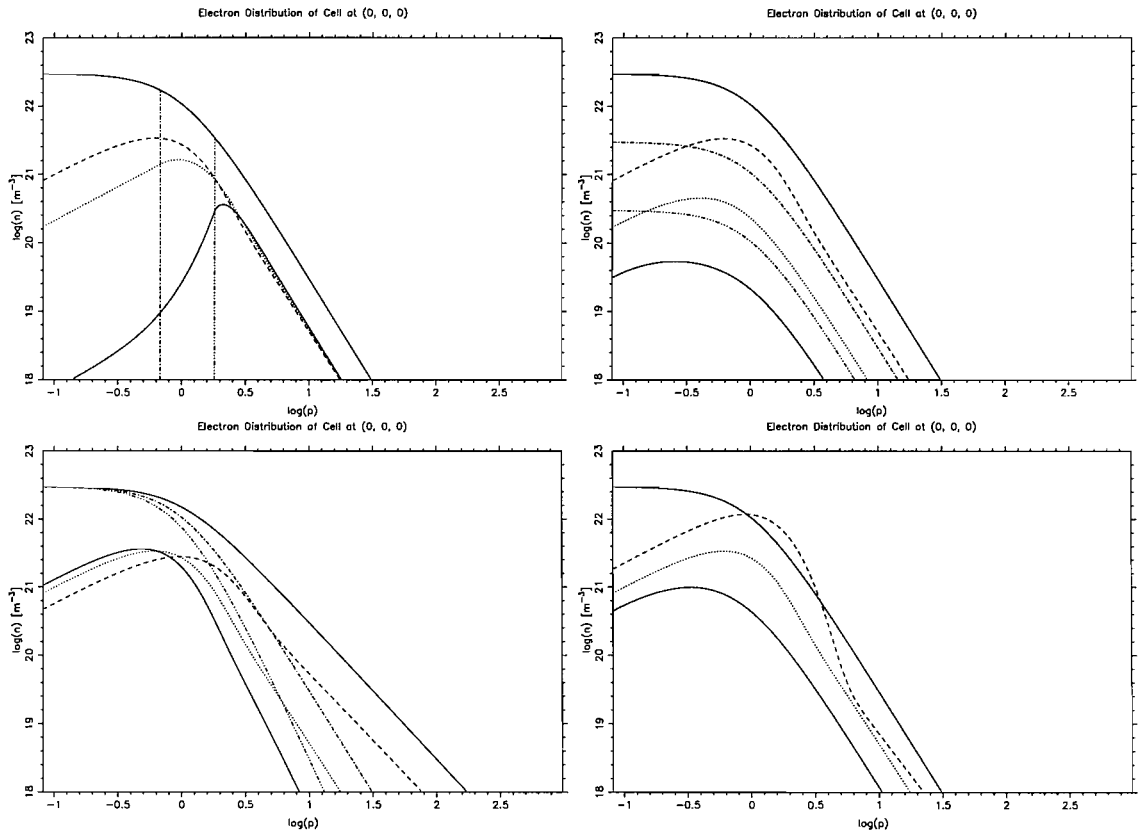
Figure 3.11: Evolving thermal electron distributions without radiative losses from an injected non-thermal power-law distribution with default parameters of $Q_0 = 3.0\times10^{22}$ m^{-3} t_{cross}^{-1}, $q = 3.0$, $t_{\text{esc}} = 0.2$ t_{cross}, $\tau_{\text{T}} = 1.0$, $t_{\text{cross}} = 3 \times 10^{-3}$ s, $\gamma_{\text{min}} = 1.0$, and $\gamma_{\text{max}} = 10^3$. With these default parameters the thermal distribution evolves to a temperature of $kT_e = 130$ keV at equilibrium. Each plot shows the initial distribution (i.e. the injected distribution integrated over the time-step) for each parameter as well as the evolved equilibrium distribution in each case. (a - upper-left panel) Increasing the minimum injected energy from $\gamma_{\text{min}} = 1.0$ (solid line) to 1.2 (dot-dashed line), and 2.0 (triple-dot-dashed line), produces thermal components with temperatures of $kT_e = 130$ (dashed line), 225 (dotted line) and 525 (solid line) keV, respectively. (b - upper-right panel) Decreasing the injected density from $Q_0 = 3.0 \times 10^{22}$ (solid line) to 3.0×10^{21} (dot-dashed line), and 3.0×10^{20} (triple-dot-dashed line) m^{-3} t_{cross}^{-1}, produces thermal components with temperatures of $kT_e = 130$ (dashed line), 70 (dotted line) and 30 (solid line) keV, respectively. (c - lower-left panel) Increasing the injected power-law index from $q = 2.0$ (solid line) to 3.0 (dot-dashed line), and 4.0 (triple-dot-dashed line), produces thermal components with temperatures of $kT_e = 225$ (dashed line), 130 (dotted line) and 90 (solid line) keV, respectively. (d - lower-right panel) Decreasing the escape-time scale from $t_{\text{esc}} = 1.0$ (solid line) to 0.2 (dot-dashed line), and 0.04 (triple-dot-dashed line) t_{cross}, produces thermal components with temperatures of $kT_e = 210$ (dashed line), 130 (dotted line) and 50 (solid line) keV, respectively.

ure 3.11a). Increasing the density of the injected electrons, Q_0 (see Figure 3.11b), not only increases the overall electron density, but also increases both the relative density of the thermal component to the non-thermal component (due to more efficient Coulomb thermalisation) and the temperature of the thermal component. Varying the power-law index of the injected non-thermal electrons (Figure 3.11c) has a significant impact upon the average energy of the injected electrons and so higher indices lead to cooler thermal components. Finally, the escape time-scale, $t_{\rm esc}$, has a significant effect upon the evolved thermal component (see Figure 3.11d). For slow escape time-scales, $t_{\rm esc} > 0.2\ t_{\rm cross}$, thermalisation is efficient and so a hotter thermal component is evolved with a greater relative density over the non-thermal component. With fast escape time-scales, $t_{\rm esc} < 0.2\ t_{\rm cross}$, thermalisation is suppressed, and so only a cool thermal component is evolved that is relatively weak compared to the non-thermal component.

Next, in Figure 3.12, we illustrate the effects of Compton cooling upon the thermalisation process by injecting a blackbody seed photon distribution into the electron distribution. Increasing the temperature of the seed photon distribution reduces the temperature of the evolved thermal electron distribution (Figure 3.12a), and reducing the intensity of the seed photons (for a given temperature) weakens the relative strength of the radiative losses and so the temperature of the evolved thermal electron distribution increases (Figure 3.12b). In models with substantial radiative losses the injected electron distribution parameters only weakly affect the outcome of the evolved thermal electron distribution's temperature. For example, in Figure 3.12c, we vary the escape time-scale through the same range as in Figure 3.11d, yet the thermal component always evolves to approximately the same temperature. However, as is clearly seen in Figure 3.12c the escape time-scale still has a strong affect upon the relative strength of the thermal component with respect to the non-thermal component of the evolved hybrid distribution. In the extreme case of strong radiative losses as shown in Figure 3.12d, the thermal component cools to $kT_{\rm e} \ll 10$ keV and even increasing the minimum injected energy, $\gamma_{\rm min}$, fails to overcome the radiative cooling.
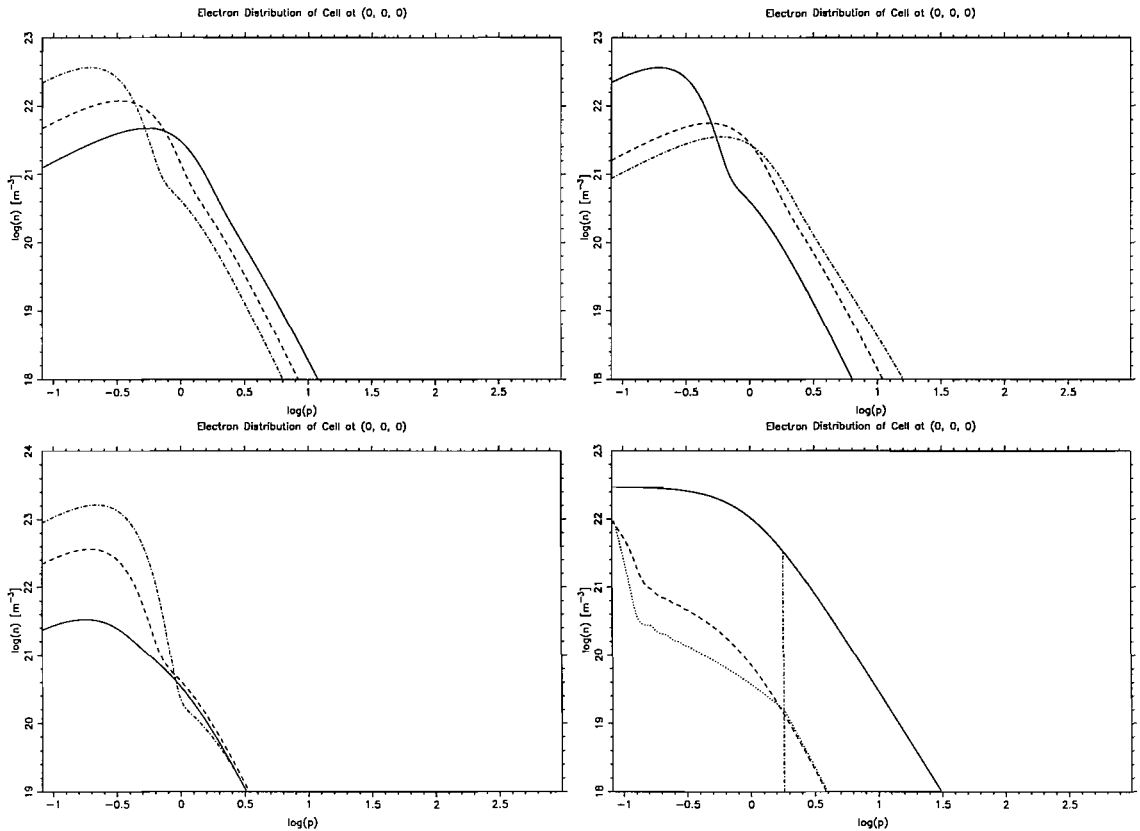
Figure 3.12: Evolving thermal electron distributions subject to radiative losses from Compton cooling. The injected electron distribution parameters are as the default parameters in Figure 3.11, in addition a blackbody photon distribution is injected into the cell that interacts with the electron distribution by Compton scattering. (a - upper-left panel) Evolved thermal components for blackbody seed photons of kT_{BB} = 0.2 (solid line), 0.3 (dashed line), and 0.4 (dot-dashed line) keV, resulting in thermal distributions of temperature kT_e = 100, 50 and 20 keV, respectively. (b - upper-right panel) Evolved thermal components for blackbody seed photons of kT_{BB} = 0.4 keV, but reduced in intensity by a factor of 1.0 (solid line), 10.0 (dashed line), and 100.0 (dot-dashed line), resulting in thermal distributions of temperature kT_e = 20, 90 and 130 keV, respectively. (c - lower-left panel) Injecting kT_{BB} = 0.4 keV blackbody seed photons and varying the escape time-scale from t_{esc} = 0.04 (solid line), 0.2 (dashed line) and 1.0 (dot-dashed line) t_{cross}. (d - lower-right panel) In the extreme case of the injected electron distribution (solid line) being irradiated by kT_{BB} = 0.4 keV photons, increased in intensity by a factor of 6 the evolved thermal component cools to $kT_e \ll 10$ keV (dashed line). Even increasing the minimum injected energy to γ_{min} = 2.0 (dot-dashed line), fails to increase the temperature of the thermal component (dotted line).

Chapter 4

The HEART Model - Application to Accretion Disc Coronae

4.1 Introduction

As discussed in section 1.2.2 the origin of the observed X-ray spectrum from black hole binaries in their two states is strongly debated. Clearly in both states there is a blackbody spectrum emitted by the system's accretion disc which explains the soft X-ray spectrum. However, the origin of the observed power-law spectra at higher energies is less well understood. Most models suggest that the high-soft spectrum is best described by emission from a Comptonising accretion disc corona. Some models suggest that this is also true for the low-hard spectrum whereas others favour a synchrotron jet emission explanation for the high-energy spectrum origin. Our approach to solving this problem is to improve upon the steady-state accretion disc corona models of the past by applying the HEART code to model the time varying nature of such a corona, with the ultimate aim of explaining the cause of state transitions as well as individual state spectra.

The best corona model fits to the two spectral states of the black hole binary Cygnus X-1 were made by McConnell *et al.* (2002) using the EQPAIR code. Their model fits are illustrated in Figure 4.2. Clearly, their fits show a decrease in the temperature of the
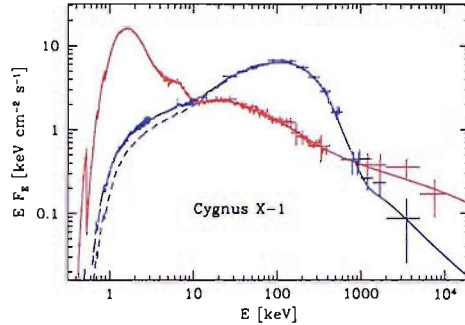
Figure 4.1: Observations of the X-ray and gamma-ray spectrum of Cygnus X-1 in both the low-hard and high-soft states. Note that the nature of the non-thermal power-law slope at gamma-ray energies is quite poorly constrained, and that in the low-hard state there is not a clear blackbody component. Taken from McConnell *et al.* (2002).
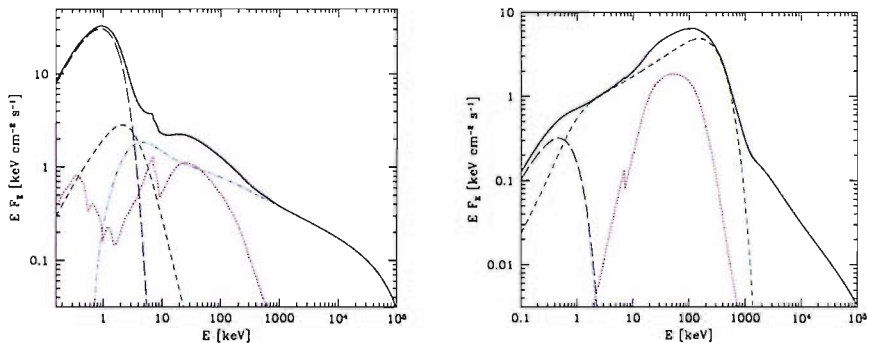
Figure 4.2: Model fits to the spectra of Cygnus X-1 shown in Figure 4.1, in (a - left panel) the high-soft state and (b - right panel) the low-hard state. In both plots the solid line denotes the total spectrum, the dashed line is the accretion disc blackbody, the short dashed line is the thermal Compton scattered component, and the dot-dashed line is the non-thermal Compton scattered component. The dotted line represents a Compton reflection component. Taken from McConnell *et al.* (2002). These spectra differ from the observations shown in Figure 4.1 as they are intrinsic, i.e. uncorrected for galactic absorption.

accretion disc blackbody spectrum between the high-soft and low-hard states. As the temperature of an accretion disc is dictated by the strength of the gravitational potential, the simplest explanation for a decrease in temperature would be the loss of the inner part of the accretion disc. This inner accretion disc material could end up in the black hole, a relativistic jet (commonly observed in the low-hard states of black hole binaries, see e.g. Fender 2001) or into the corona. Their fit suggests the corona is dominated by non-thermal electrons with a small thermal component in the high-soft state, but when the accretion disc temperature drops the corona contains a much stronger thermal component, which may be due to the injection of accretion disc material. Cygnus X-1 can remain in the low-hard state continuously for several years, in total spending 90% of its time in this state (see Zdziarski *et al.* 2002).

We will attempt to fit the spectrum of Cygnus X-1 using the HEART code to model the emission of an accretion disc corona. External blackbody radiation from the accretion disc will be injected into the plasma where it will be Compton scattered by the plasma's electron distribution. We assume that the non-thermal component of the plasma's electron distribution arises from heating due to magnetic reconnection events or shock-waves within the corona, which is modelled as an injection of power-law distributed electrons. The thermal component of the electron distribution will be modelled as the natural result of Coulomb scatterings thermalising the injected non-thermal electrons. We choose not to include the effects of bremsstrahlung processes in such a model as its radiation and energy losses are insignificant compared to the effects of Compton scattering. Also synchrotron radiation is not modelled as there is no strong evidence that a synchrotron component exists in the visible spectrum, and the heating effect of synchrotron absorption is again negligible in the relevant parameter space. By altering the strength of the injected radiation the balance in the electron distribution will change. Hence, the HEART code could be used to explain state transitions in the spectrum of accretion disc corona as a natural consequence of this process.
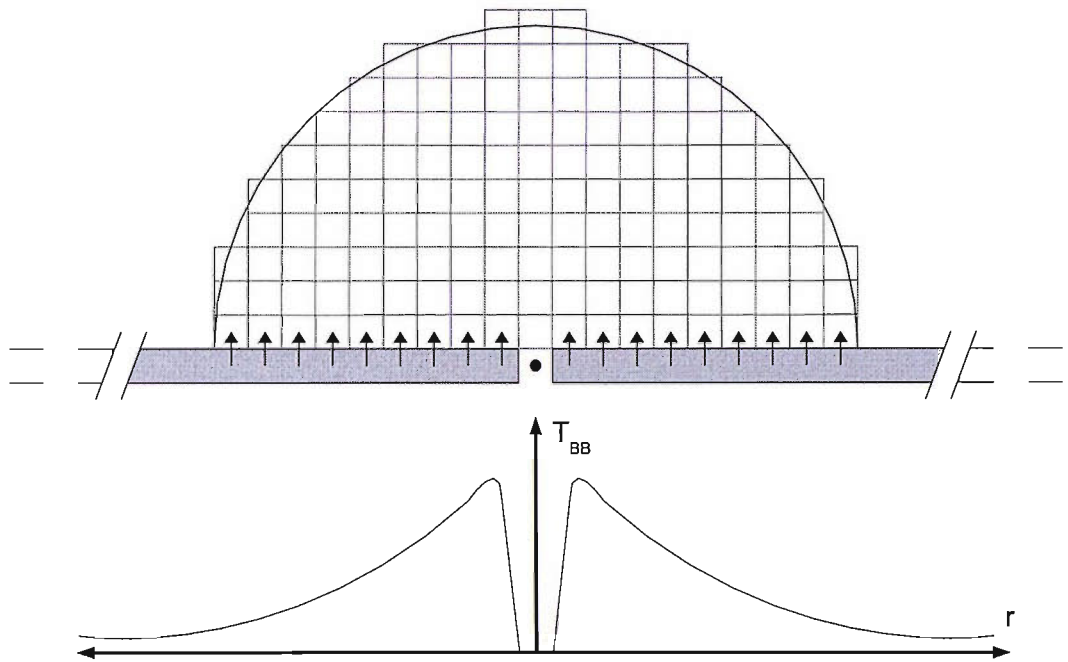
Figure 4.3: A cross-section of our accretion disc corona emission model. The corona is described by a hemisphere of energetic electrons positioned above a blackbody emitting accretion disc. Blackbody radiation, of a temperature determined by the profile plotted below the cross-section, is injected vertically into the corona's base cells.

4.2   The external corona model

Many different geometrical designs for the Comptonising corona have been put forward in attempts to explain the observed emission. However, a priori, the choice of geometry is arbitrary as there is no source of information (other than that provided by the model) to bring favour to any particular geometry. Therefore, for our initial attempts to model state changes we adopt the simple geometry of a hemisphere placed directly above the accretion disc, as illustrated in Figure 4.3. Obviously, in reality the corona will exist above and below the disc, but as they are symmetric only one side needs to the modelled.

In this geometrical configuration the injected seed photons for the Compton scattering process are unidirectional, which in turn will lead to an anisotropic distribution of scattered photons. Most of the scattered energy will be directed back down towards the accretion disc where it will be reprocessed and reflected (see Ghisellini *et al.* 1991, and Haardt &

Maraschi 1993). To simplify our model we treat the scattered radiation as being isotropic, and ignore the additional *reflection* component in the emitted spectrum. We note however that this reflection component will increase the overall strength of Compton cooling in the corona by providing an additional supply of seed photons.

4.2.1   A multi-colour accretion disc

Blackbody radiation from the surface of the accretion disc is injected into the corona of energetic electrons, where it is inverse Compton scattered. The temperature of the accretion disc varies with radius, so the injected blackbody spectrum also varies with radius. To model this injected blackbody spectrum we use the multi-colour disc model of Frank *et al.* (1992), which gives the radial dependence of the temperature as

$$T_{\mathrm{BB}}(r) = \frac{T_{\max}}{0.488} \left(\frac{3r_{\mathrm{s}}}{r}\right)^{\frac{3}{4}} \left(1 - \sqrt{\frac{3r_{\mathrm{s}}}{r}}\right)^{\frac{1}{4}}. \tag{4.1}$$

The $3r_{\mathrm{s}}$ term is the innermost radius of an accretion disc, given by the last stable orbit, around a non-rotating black hole, where r_{s} is the Schwarzschild radius. We only model the accretion disc radiation that is emitted perpendicularly to the disc. Hence, no radiation is injected into the corona within the innermost disc radius. The temperature profile is plotted with respect to the accretion disc in Figure 4.3, revealing that the maximum temperature is slightly offset from the inner radius of $r = 3\ r_{\mathrm{s}}$. To calculate the injected emission spectrum from the accretion disc we use the formula for blackbody radiation given by equation 2.19.

Accurate modelling of the complete temperature profile of the accretion disc would place another constraint upon the spatial resolution of the HEART code if for each cell just a single blackbody spectrum is injected that corresponds to the cell's centre radius. This constraint may be removed by sampling more than one radius per cell, and thus inject an integrated blackbody spectrum of several temperatures. The injected radiation spectrum is calculated prior to the main time loop of the simulation, and so the integration of the multi-temperature blackbody spectra should not affect the simulation computation time significantly.

The integration of the full multi-colour disc spectrum over the entire base of the corona is achieved by taking temperature samples at $0.1\,r_s$ intervals in both the x and y directions of the lower cell faces until the entire face is covered. Blackbody spectra of each of these sampled temperatures are then summed and divided by the total number of samples to give the integrated multi-temperature spectrum for the cell. This spectrum is injected into the cell on each time step of the simulation. Tests of this integration method reveal that it is sufficiently accurate to ensure that all models of any spatial resolution always produce identical total multi-colour disc spectra.

Sampling spectra over the face of a cell was chosen over integrating the radial temperature profile because calculating the area of each radial segment covered by each cell face would be quite complicated as every cell is aligned differently with respect to the radial axis of the disc. However, we also model the contribution to the observed flux of the outer region of the disc that lies beyond the corona (even though this radiation is not used as seed photons for the corona). For this we integrate over the radial profile to give a specific intensity of radiation from the disc of

$$
\begin{aligned}
I_{\mathrm{disc}}(\nu) &= \frac{\int_{r_{\min}}^{r_{\max}} I_{\mathrm{BB}}(\nu, T_{\mathrm{BB}}(r)) 2\pi r \mathrm{d}r}{\int_{r_{\min}}^{r_{\max}} 2\pi r \mathrm{d}r}, \\
&= 2 \frac{\int_{r_{\min}}^{r_{\max}} r I_{\mathrm{BB}}(\nu, T_{\mathrm{BB}}(r))}{(r_{\max}^2 - r_{\min}^2)} \mathrm{d}r
\end{aligned}
\tag{4.2}
$$

where $I_{\mathrm{BB}}(\nu, T)$ is the blackbody spectrum given by equation 2.19, $T_{\mathrm{BB}}(r)$ is the multi-colour disc temperature relation given by equation 4.1, $r_{\min} = r_{\mathrm{c}} = w/2$, the radius of the corona, and $r_{\max} = r_{\mathrm{out}}$, the outer radius of the accretion disc. Both radii are scaled to Schwarzschild radii. The integration is performed by defining a logarithmic grid of 100 radial points between $r_{\min}$ and $r_{\max}$, and integrating the function as described in section 3.4.1 for numerical integration over the energy grids. The disc spectrum, which is added to the total flux from the corona, is given by

$$
F_{\mathrm{disc}}(\nu) = \frac{2\pi \cos i [r_{\max}^2 - r_{\min}^2]}{D^2} I_{\mathrm{disc}}(\nu),
\tag{4.3}
$$

where i is the inclination of the disc to the observer's line-of-sight, and D is the distance to the observer.
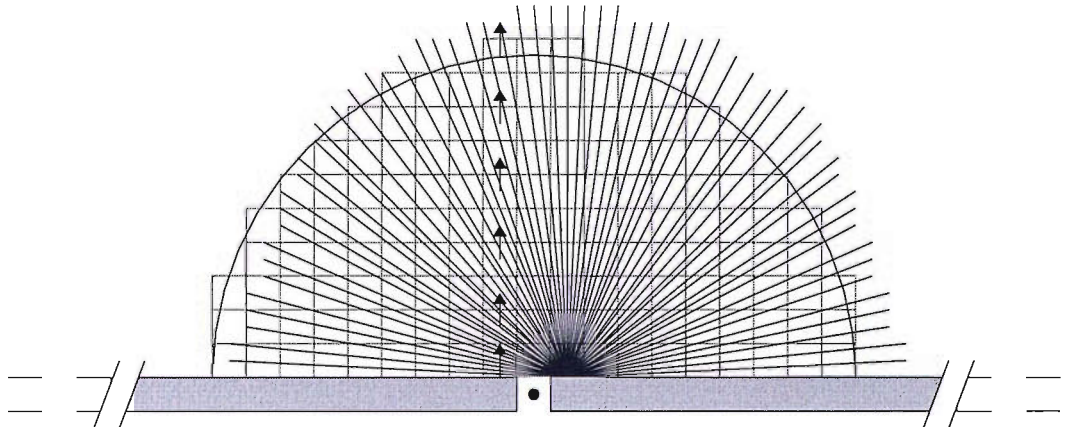
Figure 4.4: An illustration comparing the true propagation of unscattered photons through a corona to the method used in our model. Radiation is injected into the base of the corona from the accretion disc material directly beneath. In our model, these injected photons will, in the absence of scattering, only propagate vertically through the corona, as illustrated by the vertical arrows through the first column of cells to the left-hand side of the black hole. In reality, as illustrated for one cell to the right-hand side of the black hole by the diverging lines, these photons will be emitted in every direction, and thus propagate through every cell of the corona. We can account for the greater strength of radiation, but not for the geometric effect.

4.2.2   Injected radiation transfer

Only radiation with vertical direction vectors is injected into the base cells of the hemisphere. So in the absence of scattering no emission will escape from the sides of the hemisphere, and the injected multi-colour accretion disc blackbody spectrum will escape from the upper surfaces of the hemisphere cells. Inverse Compton scattering by energetic electrons within the corona will cause the strength of the escaping blackbody spectrum at the upper surfaces to diminish, and a high-energy Compton scattered spectrum will escape from all of the surfaces. In reality the injected radiation from the accretion disc will be directed in all directions, not just perpendicular to the disc, as illustrated in Figure 4.4.

Therefore, when calculating the total photon distribution in the cell for use as the seed distribution for inverse Compton scattering, we consider the injected radiation to be

multi-directional. Hence, equation 3.1 is re-written as,

$$I_{\text{cell}} = I_{\text{Int}} + I_{\text{Inj}} + \frac{I_1 + I_2 + I_3 + I_4 + I_5 + I_6}{6}, \tag{4.4}$$

where I_{Inj} is the injected distribution which has been given the same weighting as the multi-directional intrinsic distribution. As a consequence of treating injected radiation differently to other transiting radiation it must be given a unique distribution in the cell rather than simply adding it to the upper face distributions (see Figure 3.1). When performing the radiative transfer process this special photon distribution for injected radiation is transferred along with the upward travelling photon distribution to the cell above. There it is again kept separate from the other photon distributions to ensure that unscattered seed photons are treated as encompassing a π solid angle, rather than the $4\pi/6$ solid angle that regular transiting radiation occupies. Unscattered seed photons that reach the uppermost cells of the corona are included in the total upward escaping radiation flux.

4.3   Cygnus X-1 observations

We do not aim to accurately fit the observed spectra of Cygnus X-1, as that would require us to model features such as Compton reflection of corona radiation off of the accretion disc. Instead we use both the model fits shown in Figure 4.2 and the corresponding observations shown in Figure 4.1 to obtain a characteristic spectrum that agrees with the general features of the observations. Decisions upon the relative importance of each feature of the spectrum are strongly influenced by the quality of the observations. For example, the high-energy observations place poor constraints upon the slope of the non-thermal Compton tail, and the low-energy observations of the black-body spectrum are strongly affected by galactic absorption. The model fits are also useful when deciding upon which features are most important for the HEART model to recreate, as they highlight the regions where Compton reflection has a dominant effect upon the radiation spectrum.

The fluxes and frequencies of the observed features that we are trying to model are listed in Table 4.1. For the high-soft state, we aim to recreate the flux and frequency of the modelled blackbody peak together with a point, b, observed at the transition between the blackbody and Compton component and three points (c, d, and e) of the observed

power-law tail. If possible, we also aim to re-create the low-frequency tail of the multi-colour disc radiation by fitting point a, but it is of low priority as this feature is strongly affected by galactic absorption. For the low-hard state we aim to recreate the blackbody peak as modelled in Figure 4.2 together with two points (a and b) describing the slope of the modelled thermal component, its peak, a point (c) after the peak, the inflexion where the non-thermal component becomes dominant, and a point (d) on the non-thermal tail. We also try to model the general slope of both the thermal and non-thermal com-

High-Soft state		
Feature	Frequency [Hz]	νF_ν flux [W m^{-2}]
Point a	4.8×10^{16}	1.6×10^{-11}
Blackbody peak	2.2×10^{17}	4.8×10^{-11}
Point b	1.5×10^{18}	5.8×10^{-12}
Point c	2.4×10^{18}	3.6×10^{-12}
Point d	3.7×10^{19}	1.6×10^{-12}
Point e	1.7×10^{21}	3.2×10^{-13}
Low-Hard state		
Point	Frequency [Hz]	νF_ν flux [W m^{-2}]
Blackbody peak	1.1×10^{17}	5.1×10^{-13}
Point a	1.1×10^{17}	8.2×10^{-13}
Point b	9.7×10^{17}	2.4×10^{-12}
Thermal peak	3.7×10^{19}	7.5×10^{-12}
Point c	1.2×10^{20}	2.5×10^{-12}
Thermal/Non-thermal inflexion	2.8×10^{20}	3.5×10^{-13}
Point d	8.2×10^{20}	1.4×10^{-13}
Thermal slope, +0.5		
Non-thermal slope, -0.95		

Table 4.1: Observational points of the spectrum of Cygnus X-1 as shown in Figures 4.2 and 4.1 to be fit with the HEART corona model. These points may be found overlaid upon the model spectra given in the following sections. See text for an explanation of each point's significance. Flux values are accurate to within ±0.2 W m^{-2}, and the frequencies of the peaks and inflexions are accurate to within ±0.2 Hz.

ponent, though an accurate representation of the slope of the non-thermal component is not necessary as the observational data poorly constrains the range of the possible slope values.

Slopes, peaks, and points of inflexion in the model spectrum are automatically found by the HEART code by differentiating the photon distribution to find turning points, thus both the modelled frequencies and the fluxes of these features are fit. The remaining, arbitrarily chosen, points that are noted as "Points" in Table 4.1 (as opposed to "peaks", "slopes", or "inflexions") are fit by manually requesting the code to return the photon flux at the given frequencies, hence just the flux is fit for the given frequency. All fits are, of course, limited in accuracy by the simulation's frequency resolution.

4.4 Applying the external corona model to Cygnus X-1

4.4.1 Modelling photon evolution only

Initially we have applied this multi-cell model with multi-colour disc blackbody injection to simulate the radiative transfer of Compton scattered radiation through the corona without modelling the evolution of the electron distribution. We allow the photon distribution to reach an equilibrium that corresponds to a fixed, pre-determined hybrid electron distribution, and then vary the parameters until the equilibrium photon distribution agrees with the observed spectrum. This quick, simple model should be capable of reproducing the observed steady-state spectral fits and help steer the development of future time-dependent models with a knowledge of which electron distributions reproduce the desired spectra.

Additional parameters that the corona model requires over the standard HEART model parameters of Table 3.1 are defined in Table 4.2.[1]

---

[1]Note that kT_{max} is merely the temperature value that best describes the blackbody spectrum emitted by the accretion disc. In actual fact the true temperature of the accretion disc differs by a factor between 1.6 and 1.8 (see Gierlinski & Done (2004) and Davis et al. (2005)) as it does not emit like a true blackbody. As this temperature is a free parameter the disregarded temperature correction has no affect upon our models.

kT_{max}	[keV]	The peak temperature of the accretion disc
M_{BH}	$[M_\odot]$	Black hole mass, used to define Schwarzschild radii scale.
r_{in}	$[r_{\mathrm{s}}]$	Inner radius of the accretion disc
r_{out}	$[r_{\mathrm{s}}]$	Outer radius of the accretion disc
i	[degrees]	Inclination of the accretion disc to the observer
τ_{A}		Optical depth of intrinsic absorption in seed photon source

Table 4.2: Additional parameters for the HEART model of accretion disc coronae. We also re-define the plasma width, stating it as a coronal radius, $r_{\mathrm{c}} = w/2$.

We begin by fitting the frequency blackbody peak for the high-soft state by varying kT_{max}, for a fixed $r_{\mathrm{in}} = 3\ r_{\mathrm{s}}$, the innermost stable orbit. To obtain the correct normalisation of the peak we then alter the black hole mass (which varies the area of the accretion disc region that has a temperature of T_{max}), and the distance, D, of the disc to the observer. Also the width of the corona, w, affects the normalisation of the blackbody peak, so for any alterations to w the normalisation of the blackbody peak will need to be refit. To obtain a fit for the low-frequency tail of the blackbody radiation (point a in Table 4.1) we alter r_{out} and the inclination, i, which affect the intensity of the cooler outer disc radiation that does not pass through the corona. Again the corona width in this model affects this fit, and as the observations for this point are heavily affected by galactic absorption we typically leave these parameters fixed to $r_{\mathrm{out}} = 150\ r_{\mathrm{s}}$, and $i = 50°$.

Finally to fit the blackbody peak of the low-hard state, we increase r_{in} until the peak lies at the correct frequency. However, we find that although simply truncating the inner radius of the accretion disc achieves the correct peak frequency, the resulting blackbody spectrum is significantly more intense than the equivalent spectra in the models shown in Figure 4.2, as is illustrated in Figure 4.5.[2] Therefore, we apply an additional correction

---

[2]In truth, the models of McConnell et al. (2002) fail to accurately model the soft X-ray spectrum in the low-hard state, underestimating its flux by a factor of ten (see dashed line in Figure 4.1). Therefore, although we recreate the model spectrum of McConnell et al. (2002) we note that the more intense seed photon distribution modelled by a simple truncation of the inner-disc as shown in Figure 4.5 does indeed provide a better fit to the observed spectrum (see also Di Salvo et al. 2001 for a discussion of modelling the complex soft excess displayed by Cygnus X-1 in the low-hard state). This might also lead to a corona that is, overall, cooler in equilibrium.
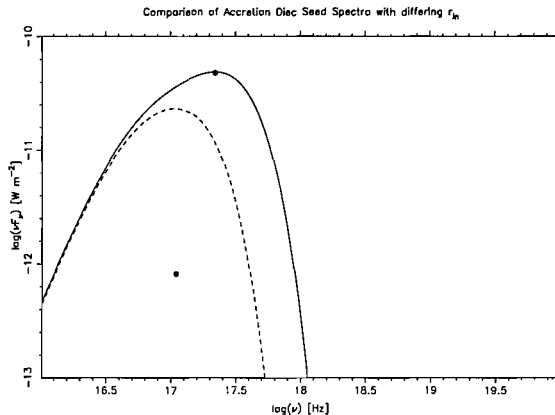
Figure 4.5: Blackbody spectra of a multi-colour accretion disc as modelled around a $5M_\odot$ black hole, with $kT_{max} = 0.34$ keV, $r_{in} = 3\ r_s$ (solid line), and $13\ r_s$ (dashed line), $r_{out} = 100\ r_s$ at an inclination of $i = 50°$. Also plotted are the positions of the high-soft and low-hard state blackbody peaks as modelled in Figure 4.2. Clearly the blackbody peak in the case of the high-soft state is matched. However, in the low-hard state simply truncating the disc to lower the frequency of the blackbody peak frequency is not sufficient as the spectrum remains too intense. Hence, the requirement of an intrinsic absorption factor, τ_A, to our models.

factor to reduce the intensity of the blackbody radiation until the normalisation is correct. Although it is possible to recreate this effect by shrinking the outer radius of the accretion disc down until the disc is a narrow ring, it is perhaps more likely that the cause of this loss of intensity is due to some additional absorption or scattering obscuring the observed blackbody. Therefore, we treat this correction factor as an absorption process, defining an intrinsic optical depth for the blackbody radiation, τ_A.

To simplify the creation of a state transition model it would be preferable to ensure that the corona remains the same size in both states. Therefore, we freely determine the corona size for one spectral state and then force the other to use the same size. It is better to fit the low-hard state first as it provides a greater constraint upon the minimum corona width ($r_{corona} > r_{in}$) and has a Comptonised spectrum that is more sensitive to the optical depth of the corona.

To fit the low-hard state we begin by just modelling the slope and normalisation of the non-thermal high-energy tail, using the parameters Q_0, and q. As the electron distributions are not evolved, the simulation initialises the non-thermal distribution with a
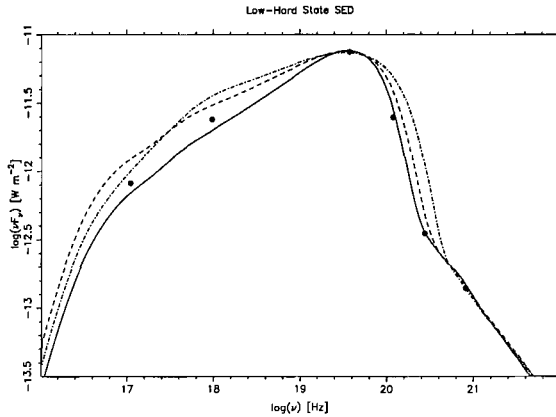
Figure 4.6: (a - left panel) Example model spectra for the low-hard state illustrating the effect of varying the coronal width, $w = 25\ r_{\rm s}$ solid line, 33 $r_{\rm s}$ dashed line, and 50 $r_{\rm s}$ dot-dashed line. For each width the unscattered blackbody component and the peak of the thermal Comptonised component have been refit. Hence each model has a different optical depth and thermal electron temperature.

normalisation of Q_0 [m^{-3}] and no further injection occurs. Next we fit the frequency of the thermal peak by altering the temperature of the thermal component of the electrons, $kT_{\rm e}$, and obtain the correction normalisation by varying $n_{\rm e}$. We refit the non-thermal high-energy tail, and the inflexion between the thermal and non-thermal component by altering the ratio of $n_{\rm e}/Q_0$ and q as necessary. To fit the low-frequency points in the spectrum we vary the width of the corona, but this also affects the frequency of the thermal peak. Therefore, for every change to the corona's width, both the blackbody peak and the thermal peak need to be refit. The effect on the low-frequency points of varying the corona width is shown in Figure 4.6.

For the high-soft state we begin by recreating the non-thermal Compton spectrum component shown in the EQPAIR model, and then attempt to fit the observations as best we can by just simply adding in a thermal component to the spectrum from a low temperature electron distribution. Due to the nature of the way we fit these models to the observations, each iteration has to be performed by hand, and so to minimise the simulation duration we have used a coarse spatial resolution of $R_{\rm cells} = 6$. The energy grids can be modelled with a high resolution though, and so we use $R_\gamma = 128$, and $R_\nu = 192$. The best fit models are shown in Figure 4.7, together with the model parameters. In the high-soft state varying the spatial resolution has no effect upon the resulting spectral model, and
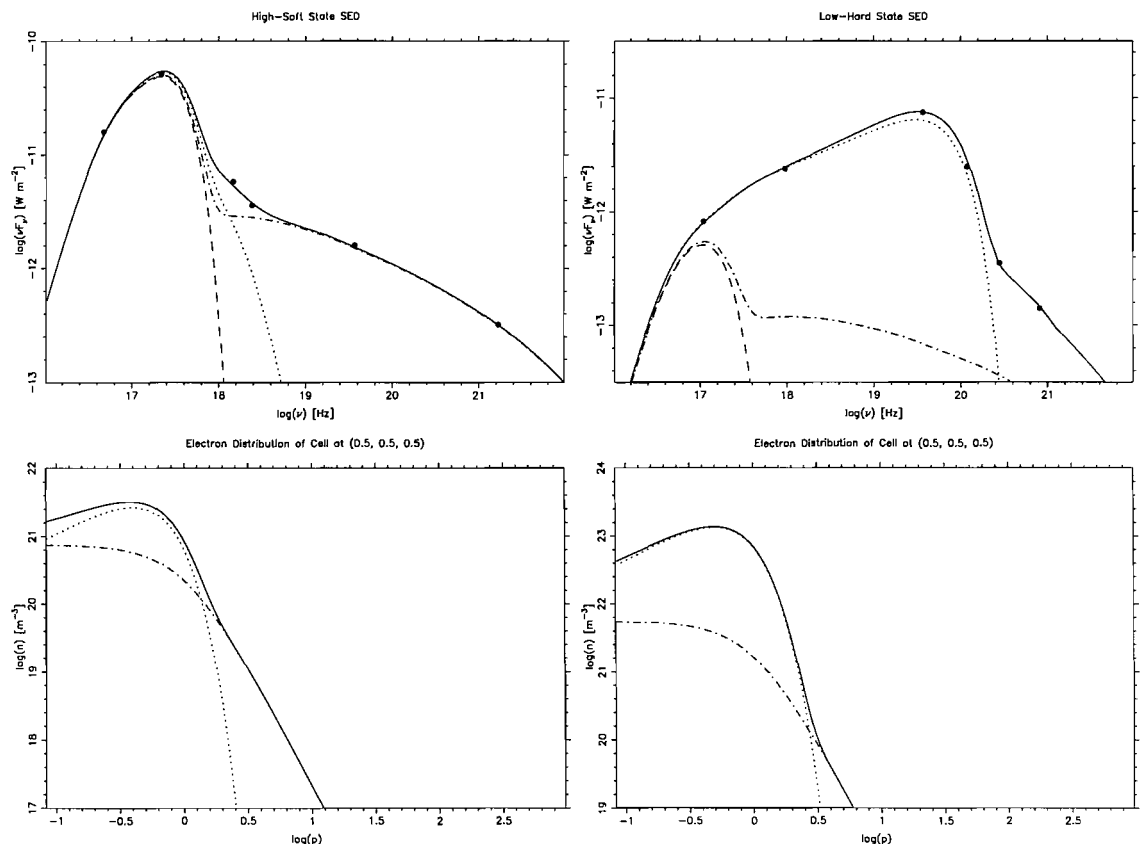
Figure 4.7: Approximate fits to the two spectral states of Cygnus X-1, (a - left panel) the high-soft state and (b - right-panel) the low-hard state, using the HEART code with fixed electron distributions. Each spectral plot shows the total spectrum (solid line), the unscattered blackbody component only (dashed line), just the non-thermal Comptonised component and blackbody (dot-dashed line), and just the thermal Comptonised component and blackbody (dotted line). Also plotted are the points to be fit, as listed in Table 4.1. Below each spectrum is a sample cell's electron distribution, which are modelled with $\gamma_{min} = 1$ and $\gamma_{max} = 1000$. In both models the system is assumed to be at $D = 2000$ pc, with $i = 50°$, $M_{BH} = 5.0 M_\odot$, and an accretion disc with $kT_{max} = 0.34$ keV, $r_{out} = 100 \, r_s$, with a $r_c = 31 \, r_s$ corona. The high-soft state is modelled with a non-thermal electron distribution of $Q_0 = 7.5 \times 10^{20} \, m^{-3}$, $q = 3.55$, and a thermal electron distribution of $n_e = 8.0 \times 10^{20} \, m^{-3}$, and $kT_e = 65$ keV, giving a total optical depth of $\tau_c = 1/30$. The low-hard state is modelled with $r_{in} = 15.0 \, r_s$, a non-thermal electron distribution of $Q_0 = 5.5 \times 10^{21} \, m^{-3}$, $q = 3.5$, and a thermal electron distribution of $n_e = 6.2 \times 10^{22} \, m^{-3}$, $kT_e = 95$ keV, giving a total optical depth along the x-axis of $\tau_c = 2$. In the low-hard state the blackbody is reduced with an intrinsic absorption factor of $\tau_A = 3.83$.
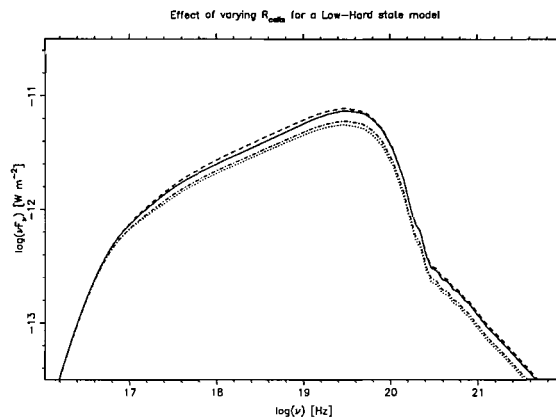
Figure 4.8: (b - right panel) The effect of increasing the spatial resolution of a low-hard spectral state model, $R_{cells} = 6$ (solid line), 8 (dashed line), 12 (dot-dashed line), 16 (dotted line). The model converges to a unique solution for $R_{cells} > 12$. However, computational time restraints restrict us to a resolution of $R_{cells} = 6$.

in the low-hard state the effect is quite small (see Figure 4.8). Note that in these fixed electron distribution models we need to inject steeper power-laws for the non-thermal electron distribution than for models with freely evolving electron distributions to model the steepening effect of energy losses.

4.4.2   A full photon and electron evolution model

Knowing that it is possible to obtain rough fits to the observed spectra with fixed electron distributions, the next step is to develop a physically correct model, where the photons and electrons are in equilibrium with *each other* not just themselves. We achieve this by continually injecting a non-thermal power-law distribution of electrons and allowing the distribution to cool through inverse Compton scatterings and thermalise through Coulomb scatterings, whilst allowing electrons to escape over a reasonable time-scale to counter the continuing injection. Varying the model parameters to make the resulting equilibrium photon spectra agree with observations of Cygnus X-1 in both states will allow us to see how the model parameters differ between the two states, and hence we can gain an insight into the physical difference between the states.

In practice, however, we find that for the geometry described in this chapter a thermal electron distribution of the required temperature to recreate the model fits shown in Figure

4.7 cannot be maintained. In a model with an external corona where all of the accretion disc radiation is injected vertically through the base, the energy losses due to inverse Compton scattering dominate over any thermalisation mechanism and so the electrons rapidly cool to temperatures of $kT_e \ll 10$ keV (see e.g. Figure 4.9). This is in agreement with the Dove *et al.* (1997) Monte Carlo models for such geometries. Under such strong radiative cooling conditions even increasing the average energy of the injected electrons by, say, increasing the minimum energy is not sufficient to balance out the dominant Compton cooling that results from injecting the required seed photon spectrum (see examples in section 3.6.2).

As the Compton scattered radiation in the high-soft state is dominated by the non-thermal component we can still produce a reasonable fit to this spectrum using the vertical injection model without trying to obtain a good fit to the thermal component. In Figure
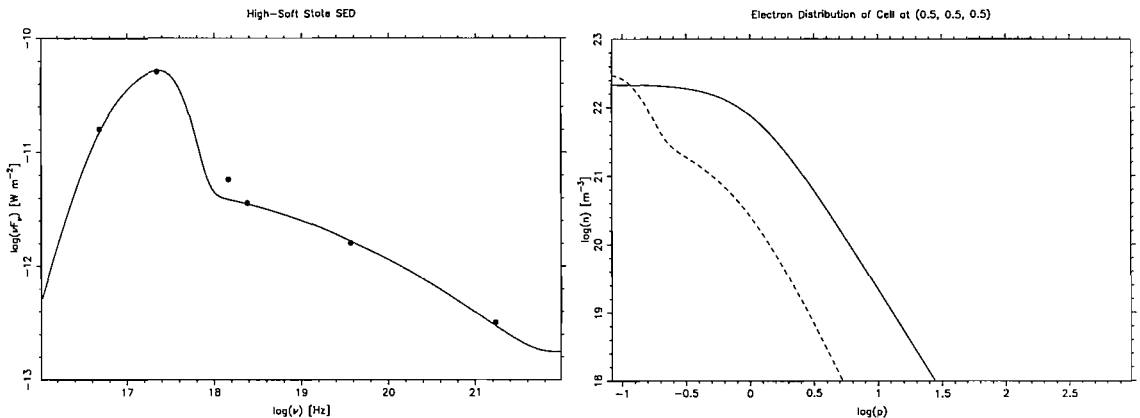

Figure 4.9: The equilibrium photon spectrum (left panel) from a freely evolved electron distribution (right panel - the solid line gives the initial injected distribution, and the dashed line gives the final equilibrium distribution) that gives the best agreement with the observed spectrum of Cygnus X-1 in the high-soft state, using the vertical injection into an external corona model. The parameters are the same as in the fixed electron distribution model, but with no initial thermal distribution and an injected non-thermal electron distribution of $Q_0 = 1.3 \times 10^{23}$ m$^{-3}t_{\mathrm{cross}}^{-1}$, $q = 3.0$. The thermalised component evolves to an equilibrium temperature of $kT_e \ll 10$ keV due to strong Compton cooling. Therefore, we use a fast escape time-scale of $t_{\mathrm{esc}} = 0.01\, t_{\mathrm{cross}}$ to suppress the relative strength of the thermal component to the non-thermal component of the electron distribution. The optical depth to Compton scattering with these parameters is $\tau_c = 0.02$.

4.9 we show the best fit photon spectrum together with an example equilibrium electron distribution from one of the cells. Note that for all regions of the corona, Compton cooling is too strong to allow any cell to maintain an electron distribution temperature of $kT_e > 10$ keV. As the low-hard state is dominated by a thermal Compton component it cannot be recreated with the external corona model.

4.5   The internal corona model

The fact that external corona models cannot maintain thermal electron distributions was also noted by Dove *et al.* (1997), who produced a Monte-Carlo simulation of thermal-only inverse Compton scattering in coronae. They propose moving the corona to a position *within* the inner edge of the accretion disc. Hence, not all the radiation from the accretion disc will pass through the corona, just those photons with vectors that intercept the corona. Thus by decreasing the intensity of radiation entering the corona, the strength of the cooling effect of inverse Compton scattering upon the electron distribution will diminish.

Recreating such a model for the HEART code would be quite complicated as it requires a determination of the three-dimensional distribution of radiation vectors impinging upon the discretely modelled spherical corona. To greatly simplify this task we assume that the multi-colour disc radiation approaches a spherical corona parallel with the $x - y$ plane from infinity (see Figure 4.10). To account for the greater inclination of the corona's surface to this radiation at increasing heights, we only inject radiation to cells within the height range $|z| < w/4$, i.e. only half of the corona's height. The radiation is injected into every cell with an open face in the x or y direction and is treated as a transiting photon distribution in the relevant direction, unlike our treatment of the seed photons in the vertical injection model as described in section 4.2.2. The injected radiation is the multi-colour disc spectrum as integrated over the entire accretion disc, with $r_{\min} = r_{\mathrm{in}}$, and $r_{\max} = r_{\mathrm{out}}$. The intensity as given by equation 4.2 is divided by six, because each injected photon distribution only enters one face of a cell and hence passes through a $4\pi/6$ solid angle. The disc spectrum as calculated in equation 4.3 is also added to the observed
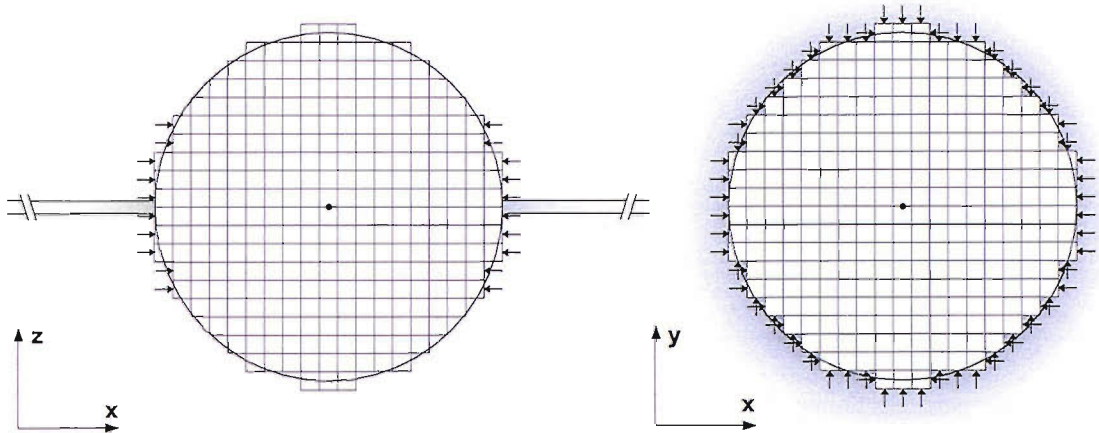
Figure 4.10: A simplified internal corona model. A spherical corona fills the location between the black hole and the accretion disc's innermost edge. Radiation is treated as coming from infinity and is injected horizontally into the corona in the x and y directions. This radiation is received over only half of the corona's height to account for a loss of intensity due to surface inclination effects. The diagram on the left gives the side view of the x-z plane, and the diagram on the right gives the top down view of the x-y plane.

flux of the model.

For this new geometry we need to refit the blackbody spectrum using the same procedure as for the external corona model. First we fit the high-soft state unscattered blackbody peak and low frequency tail with a fixed $r_{in} = 3\ r_s$. Followed by fitting the modelled low-hard state blackbody peak by truncating to the inner part of the disc to larger radii. We find that with this model we need a slightly hotter ($kT_{max} = 0.35$ keV), larger ($r_{out} = 150\ r_s$) disc inclined at a greater angle ($i = 60°$) to the fit the same observations. Note that unlike the external corona model, in this model the coronal radius has no impact upon the strength of the unscattered blackbody radiation if $r_c \ll r_{out}$. Also the inclination affects the intensity of the entire blackbody spectrum, not just the low-frequency tail. The low-hard state blackbody can be recreated by truncating the disc to $r_{in} = 13.0\ r_s$, and applying an intrinsic absorption factor of $\tau_A = 3.76$.

To fit the Comptonised spectra of each state we use a model with freely evolving electron distributions. Again, having already determined the seed photon distribution parameters,
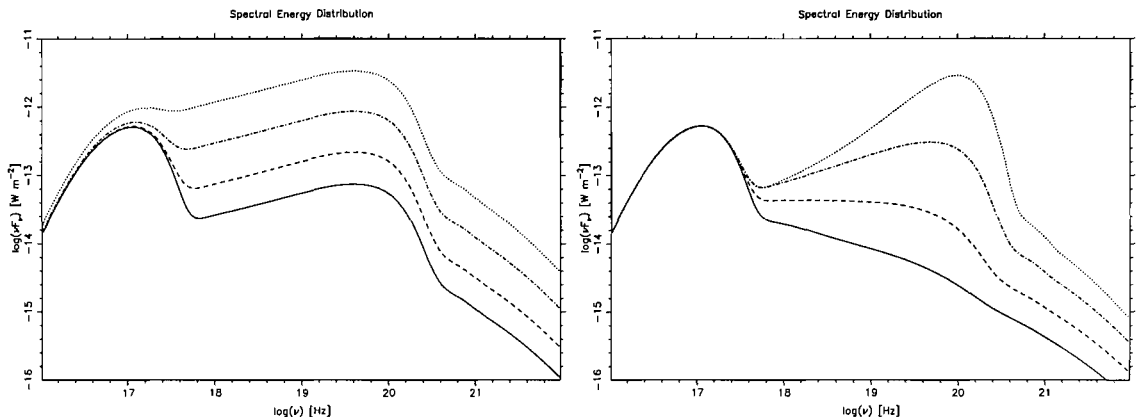
Figure 4.11: Example internal corona spectra with (a - left panel) differing corona radii $r_c = 30$ (solid line), 50 (dashed line), 100 (dot-dashed line), 200 r_s (dotted line), and (b - right panel) differing optical depths $\tau_c = 1/3$ (solid line), 2/3 (dashed line), 1.25 (dot-dashed line), 2.5 (dotted line). In this model the radius of the corona is strongly constrained by the observed relative strength of the Compton component over the unscattered disc blackbody spectrum. The optical depth affects the slope of the thermal Compton component.

we start by fitting the Comptonised radiation of the low-hard state spectrum, and then fit the high-soft state using the same corona radius. The sensitivity of the low-hard state spectrum to the corona radius is well illustrated by this model, because it is the only parameter that determines the relative normalisation of the Comptonised spectrum to the blackbody spectrum, as shown in Figure 4.11a. In fitting the low-hard state Comptonised spectrum we begin by fitting the non-thermal tail's slope by varying q. Then from the observed transition between the unscattered blackbody spectrum and the Comptonised spectrum we can deduce the size of the corona, r_c. As illustrated in Figure 4.11b, the slope of the thermal component is determined by the optical depth. The relative intensity of the thermal peak to the non-thermal tail is determined by the t_{esc} parameter. With fast escape time-scales the non-thermal component is stronger for a given optical depth as the density of the injected distribution, Q_0, must be increased. With slower escape time-scales there is more time to thermalise the injected electrons and so the relative strength of the thermal component increases. The time to reach equilibrium for the electron distribution, and ultimately the photon distribution, is also determined by the escape time-scale, and the optimum time is typically $t_{esc} \approx 0.2\ t_{cross}$.

In practice, with the internal corona model we require much larger coronae than for the

external corona model in order to recreate the observed relative strength of the Compton component. This poses a problem as when $r_c \gtrsim r_{out}$ the corona contributes a larger fraction of the observed unscattered blackbody component than the accretion disc because it has a larger surface area! This effect, visible in the spectra of the larger coronae shown in Figure 4.11a, is a consequence of the assumption that the accretion disc lies well beyond the surface of the corona becoming invalid and so we should not treat the injected seed photons as plane-parallel rays approaching the corona from infinity. Although we can compensate for the unrealistically bright blackbody component by reducing the intensity of the seed photons, this model is still geometrically unphysical. The low-hard state is best fit by a corona of radius $r_c > 1000\ r_s$, but as the model becomes ever more incredible as the corona expands beyond the outer radius of the accretion disc, we chose to limit the size of the corona to a thousand Schwarzschild radii.

The advantage of this internal corona model with horizontal injection is that it allows for the evolution of a high-temperature thermal electron distribution, which in turn produces the dominant thermal "hump" in the Comptonised spectrum. To match the observed Comptonised spectrum exactly the thermal electron distribution has to evolve to exactly the right temperature. However, in following the steps above, to match the slope of the thermal hump, its normalisation, and the slope of the non-thermal high-energy tail and its normalisation, there are no more free parameters. Hence, we have no control over the final equilibrium temperature, which is determined by the balance between the average injected electron energy and Compton cooling. It is a product of the injected electron distribution parameters, which are chosen to fit the spectral features listed above, and the photon injection model that determines the strength of Compton cooling. Therefore, if the equilibrium temperature disagrees with the observations then the model doesn't fully account for the extent of either the cooling processes or the heating processes.

The spectra and parameters of the best fit models for both states are shown in Figure 4.12. We find that in this low-hard state model Compton cooling is negligible and thus the electron distributions throughout the corona are isothermal at equilibrium, having reached an equilibrium temperature just based upon their average injection energy. Unfortunately, this equilibrium temperature for the best fit model parameters is too high, at $kT_e \approx 125$
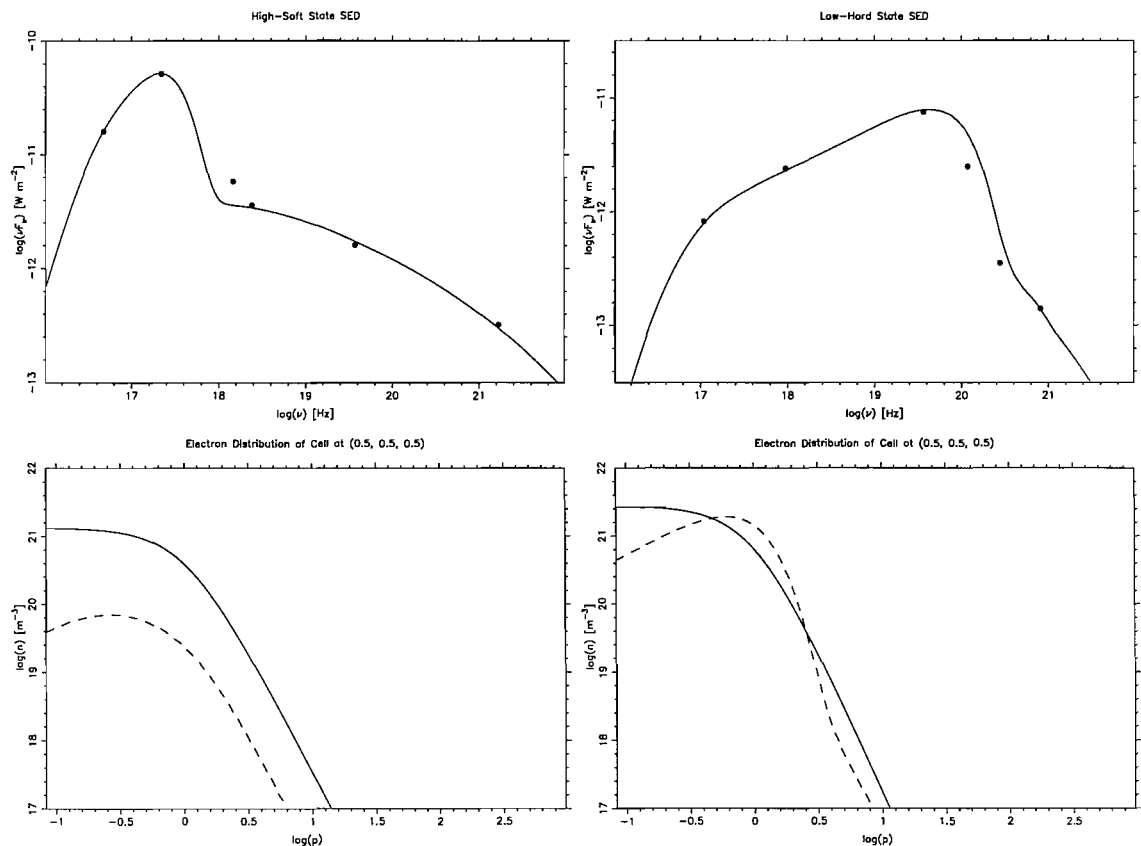
Figure 4.12: Approximate fits to the two spectral states of Cygnus X-1, (a - left panel) the high-soft state and (b - right-panel) the low-hard state, using the HEART code with the internal corona model. Below each spectrum are the freely evolved electron distributions from a sample cell that give rise to the spectra. The solid line gives the initial injected distribution which is modelled with $\gamma_{\min} = 1$ and $\gamma_{\max} = 1000$, and the dashed line gives the final equilibrium distribution. In both models the system is assumed to be at $D = 2000$ pc, with $i = 60°$, $M_{\mathrm{BH}} = 5.0 M_\odot$, and an accretion disc with $kT_{\max} = 0.35$, $r_{\mathrm{out}} = 150\ r_{\mathrm{s}}$, with a $r_{\mathrm{c}} = 1000\ r_{\mathrm{s}}$ corona. The high-soft state is modelled with an injected non-thermal electron distribution of $Q_0 = 8.0 \times 10^{21}$ m$^{-3} t_{\mathrm{cross}}^{-1}$, $q = 3.6$, with $t_{\mathrm{esc}} = 0.01\ t_{\mathrm{cross}}$, which thermalises to a temperature of $kT_{\mathrm{e}} \approx 30$ keV and has an optical depth of $\tau_{\mathrm{c}} = 0.03$ at equilibrium. The low-hard state is modelled with $r_{\mathrm{in}} = 13.0\ r_{\mathrm{s}}$, and an injected non-thermal electron distribution of $Q_0 = 1.65 \times 10^{22}$ m$^{-3} t_{\mathrm{cross}}^{-1}$, $q = 4.2$, with $t_{\mathrm{esc}} = 0.25\ t_{\mathrm{cross}}$, which thermalises to a temperature of $kT_{\mathrm{e}} \approx 125$ keV and has an optical depth of $\tau_{\mathrm{c}} = 1.25$ at equilibrium. To counter the unphysical nature of the geometry the blackbody intensity in the high-soft state is reduced with an intrinsic absorption factor of $\tau_{\mathrm{A}} = 3.07$, and with $\tau_{\mathrm{A}} = 6.31$ in the low-hard state.

keV, to achieve a good fit to the peak of the low-hard spectrum caused by thermal Compton scattering. To reduce the equilibrium temperature to $kT_e \approx 95$ keV, which, as revealed by the fixed electron distribution models, produced the best spectral fits, would require a model with greater Compton cooling.

In the high-soft state, the more intense radiation leads to more significant Compton cooling and reduces the equilibrium temperature of the lower-half of the corona. Furthermore, as there is not such a strong constraint upon the relative strengths of the thermal and non-thermal components in this spectral state, we are free to vary the escape time-scale solely to achieve the correct equilibrium temperature (which affects the equilibrium temperature as shown in section 3.6.2). However, models with strong thermal components, and hence slow escape time-scales, do not produce good fits to the observations. This is because, in this model, the relative density of the thermal electron distribution and it's temperature are tied together in such a way that the best fitting combination found in the fixed electron distribution models of section 4.4.1 is impossible to achieve. Therefore, when modelling the high-soft state we again just fit the dominant non-thermal component by suppressing the thermalisation with a fast escape time-scale, which also leads to cool isothermal electrons throughout the corona, with $kT_e \approx 30$ keV.

4.6   An alternative model

Finally, we attempt to model the low-hard state of the Cygnus X-1 corona using a much simpler model. Instead of using the accretion disc as a source of seed photons for Compton scattering, we model the corona as a spherical plasma with a central core sufficiently dense to emit its own blackbody radiation. Hence, in this model, the blackbody seed photons are injected at the centre of the spherical plasma and are emitted radially outwards (see Figure 4.13). For odd values of R_{cells} the injected radiation is added to the transiting radiation at each of the six faces of the central cell. For even values of R_{cells} the injected radiation is added to the transiting radiation at each outward face of the central eight cells.
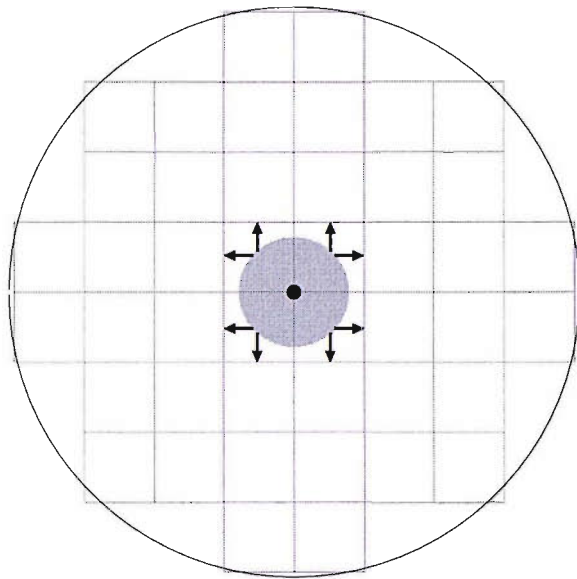
Figure 4.13: Two-dimensional cross-section of an alternative model for the corona in the low-hard state. Instead of using the accretion disc to provide the seed photons, we hypothesise the existence of a dense inner corona, spherical in shape, that injects blackbody photons into the surrounding corona. The injected seed photons are then transferred to the surrounding cells in all directions. In the above diagram the dense inner corona is shown as a grey circle, with the black hole as a black circle. The design is symmetrical, therefore the axes of this cross-section are irrelevant.

To recreate the observed low-hard state spectrum with this model we begin by obtaining the correct blackbody spectrum for the seed photons. The peak frequency of the blackbody spectrum is determined by the temperature kT_{BB}. For a given optical depth to Compton scattering, τ_c, the size of the corona, r_c, in this model, only affects the normalisation of the entire spectrum. Therefore, to recreate the peak flux of the blackbody seed photons we model the propagation of the seed photons through the corona in the absence of Compton scattering and adjust the size of the corona until we have achieved the correct peak flux for the unscattered blackbody spectrum.

Again, the slope of the non-thermal Compton scattered radiation component, or the high-energy "tail", is given by power-law index, q, of the non-thermal injected electrons. Having obtained the correct value for q, the slope of the thermal Compton scattered radiation component, the main "hump", is then governed by the optical depth of the corona, τ_c, which for a fixed coronal radius, r_c, is controlled by adjusting the density, Q_0,
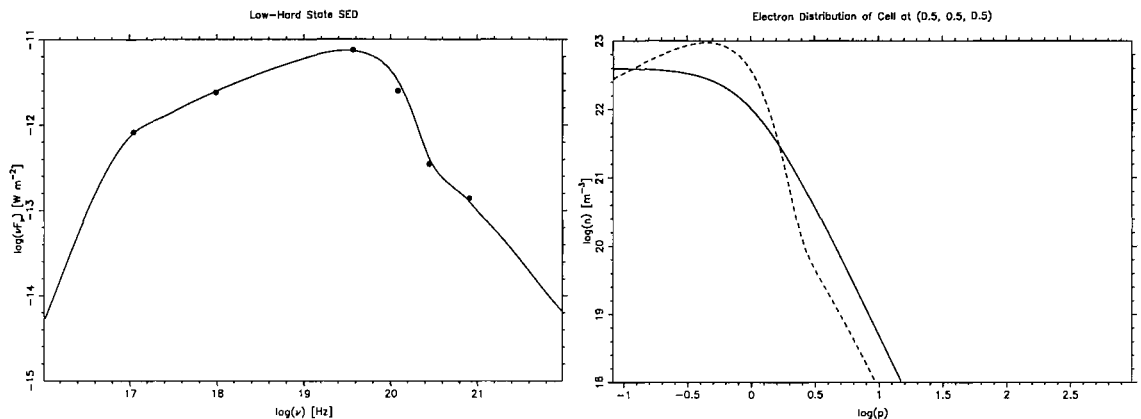
Figure 4.14: The equilibrium photon spectrum (left panel) from a freely evolved electron distribution (right panel - the solid line gives the initial injected distribution, and the dashed line gives the final equilibrium distribution) that gives the best agreement with the observed spectrum of Cygnus X-1 in the low-hard state, using the alternative model of internal injection within a corona of size $r_c = 5.0 \times 10^5$ m or 18 r_s (with a $10 M_\odot$ black hole). The seed photon distribution is a blackbody spectrum with $kT_{BB} = 0.12$ keV, and the injected electron distribution parameters are $Q_0 = 2.4 \times 10^{23}$ m$^{-3} t_{cross}^{-1}$, $q = 3.9$, and $t_{esc} = 0.45$ t_{cross}. At equilibrium the optical depth of the corona is $\tau_c = 1.29$, with thermal electron distributions ranging in temperature from $kT_e = 80$ keV in the centre, to $kT_e = 125$ keV in the outer edge of the corona.

of the injected electrons. As the flux of the unscattered seed photons decreases as the optical depth increases, it necessary at this stage to readjust the size of the corona, r_c, to obtain the correct normalisation for the spectrum, and thus also adjust Q_0 to maintain the same optical depth.

Once both the normalisation and slope of the main hump of radiation from thermal Compton scattering has been matched to the observations, the last step is to obtain the correct normalisation for the high-energy tail. As this is created by non-thermal Compton scattered radiation its normalisation is controlled by the normalisation of the non-thermal component of the hybrid electron distribution. We adjust this by altering the escape time-scale, t_{esc}, (as detailed in section 3.6.2). However, this parameter also in turn affects both the slope of the thermal Compton scattered radiation and the slope of the non-thermal Compton scattered radiation. So finally, q, and τ_c have to be slightly readjusted for the new value of t_{esc}.

The model spectrum for the parameters that best fit the observations is shown in Figure 4.14. In this model we only need a small corona, $r_c = 5 \times 10^5$ m or $r_c = 18 \ r_s$ (for a $10 M_\odot$ black hole), to recreate the observed spectrum. Compton cooling in such a small corona is significant at the centre where the thermal electron distribution reaches an equilibrium temperature of $kT_e = 80$ keV. However, at the outer edges of the corona Compton cooling is again negligible, and the electrons find equilibrium at a temperature of $kT_e = 125$ keV. So despite the centre cell electron distribution in Figure 4.14 agreeing with that of the fixed electron distribution model of Figure 4.7, the spectrum in this model is still too "hot".

To cool the equilibrium temperature of just the outer cell electrons it would be logical to make the average injected electron energy radially dependent. The parameter that is most suited to radial variations is the injected density, Q_0, as the gravitational potential of the black hole gives rise to a r^{-1} density profile. To test this out we make the injected density dependent upon the distance of the cell to the origin of our co-ordinate system (the geometric centre of the spherical corona) following $Q_0(r) = Q_0(0.866/r)$, where r is the distance from the cell centre to the origin in units of cell width. Therefore, a density of Q_0 is injected into the cells at the centre of the corona and this decreases with distance from the centre for the outer cells. However, in practice we find that making Q_0 a function of radius doesn't affect the radial dependence of the equilibrium temperature. This is because it doesn't change the average *injected* energy of the distribution, and in these models in particular, the effect of decreasing the overall density of the plasma does not lead to reduction in the average energy of the distribution at *equilibrium*. As shown in section 3.6.2, changes to the overall density of the plasma *can* still have a small effect upon the equilibrium energy for certain parameters. A larger effect upon the equilibrium energy can be achieved by making the escape time-scale a function of radius.

With this internal injection model, there are no constraints upon the size of the corona, and by balancing the intensity of the seed photons using the τ_A parameter we can freely choose to model a corona of any size. Therefore, to compare with the previous horizontal injection model, we use the internal injection model to investigate a corona of radial size $r_c = 1700 \ r_s$. Following the above procedure to find the best fitting electron injection pa-

		Plasma Parameters			Injected Seed Photon Parameters				
State	Model	D [pc]	r_{c} [r_{s}]	$kT_{\max}$ [keV]	M_{BH} [$M_{\odot}$]	r_{in} [r_{s}]	r_{out} [r_{s}]	i	τ_{A}
High-Soft	ECORONA *fixed*	2000	31	0.34	5	3	100	50°	N/A
High-Soft	ECORONA	2000	31	0.34	5	3	100	50°	N/A
High-Soft	ICORONA	2000	1000	0.35	5	3	150	60°	3.07
Low-Hard	ECORONA *fixed*	2000	31	0.34	5	15	100	50°	3.83
Low-Hard	ICORONA	2000	1000	0.35	5	13	150	60°	6.31
Low-Hard	ICORONA_IINJ	2000	18	0.12	10	N/A	N/A	N/A	N/A

		(Injected) Electron Distribution Parameters							
State	Model	$\gamma_{\min}$	$\gamma_{\max}$	Q_0 [m^{-3}]	q	t_{esc} [t_{cross}]	n_{e} [m^{-3}]	kT_{e} [keV]	τ_{c}
High-Soft	ECORONA *fixed*	1	1000	7.5×10^{20}	3.55	N/A	8.0×10^{20}	65	1/30
High-Soft	ECORONA	1	1000	$1.3 \times 10^{23}\, t_{\mathrm{cross}}^{-1}$	3.0	0.01	N/A	$\ll 10$	1/50
High-Soft	ICORONA	1	1000	$8.0 \times 10^{21}\, t_{\mathrm{cross}}^{-1}$	3.6	0.01	N/A	~ 30	1/30
Low-Hard	ECORONA *fixed*	1	1000	5.5×10^{21}	3.5	N/A	6.2×10^{22}	95	2
Low-Hard	ICORONA	1	1000	$1.65 \times 10^{22}\, t_{\mathrm{cross}}^{-1}$	4.2	0.25	N/A	~ 125	1.25
Low-Hard	ICORONA_IINJ	1	1000	$2.4 \times 10^{23}\, t_{\mathrm{cross}}^{-1}$	3.9	0.45	N/A	~ 80 to ~ 125	1.29

Table 4.3: Best fit HEART parameters in modelling Cygnus X-1.

rameters, we find that such a large corona is again isothermal, with insignificant Compton cooling and electrons reaching an equilibrium temperature of $kT_e = 140$ keV.

The parameters of all the HEART corona model spectral fits are compared in Table 4.3.

4.7   State transitions

Although the internal corona model clearly requires an unphysical geometry it is still capable of approximately recreating both of the observed states and so can be used to investigate changes to the broad X-ray spectrum during the transition period between each of the states.   By allowing the model to reach an equilibrium for a given state we can initiate a state change by spontaneously altering the injected electron and seed photon distributions and allowing the model to reach a new equilibrium. Starting with the high-soft state equilibrium we truncate the inner part of the accretion disc, dropping the temperature and intensity of the injected blackbody spectrum, and correspondingly increase the density of the injected electron distribution by altering the parameter values
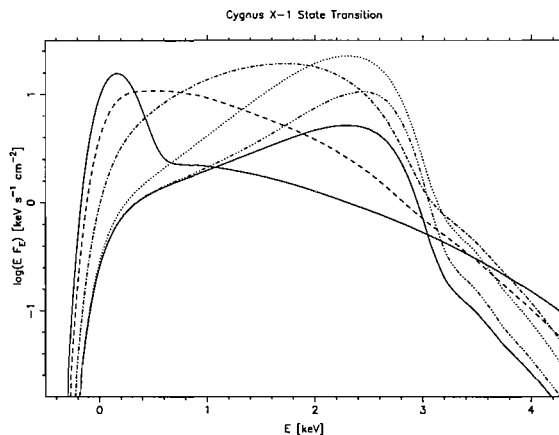
Figure 4.15: A state transition as simulated by the HEART code with an internal corona model, showing the evolution from a high-soft state to the low-hard state. The spectra include the effect of galactic absorption. The solid line denotes both of the equilibrium states, and the times plotted are $t = 0$ (solid line), 4/6 (dashed line), 8/6 (dot-dashed line), 14/6 (dotted line), 20/6 (triple-dot-dashed line), 6 (solid line) t_{cross} where $t = 0$ is defined as the instant when the injection parameters change and $t_{cross} = 0.1$ s.

Figure 4.16: Observations of Cygnus X-1 in various states from Zdziarski *et al.* (2002).

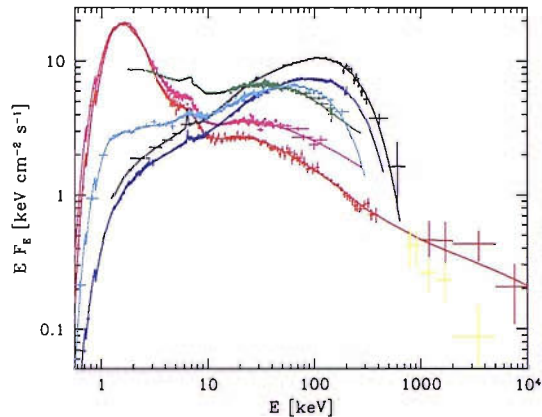to those determined in the previous section.

We allow the simulation to reach the high-soft equilibrium, and then instantaneously change the injection parameters to those that produce the low-hard spectrum. After six light crossing times the photon distribution reaches the low-hard state equilibrium. The time evolving spectra between the states are shown in Figure 4.15, and comparison observations are shown in Figure 4.16. To more readily allow comparisons between the model and the observations we alter the spectra to account for galactic absorption with an optical depth given by

$$\tau_{\text{GA}} = 2.0 \times 10^{-26} \, N_{\text{H}} \, \left(\frac{h\nu}{1\text{keV}} \right)^{-8/3} , \qquad (4.5)$$

following Longair (1994), where the hydrogen column density $N_{\text{H}} = 6.0 \times 10^{-25}$ m^{-2} for Cygnus X-1 (following McConnell *et al.* 2002). It is clear that this simple model of an instantaneous change does not entirely agree with the observations. Either the parameter change should be more gradual, or say, there should be a delay between the alteration of the seed photon injection parameters and the subsequent changes to the electron injection parameters. Alternatively it is possible that this model cannot reproduce the observed variations over time, but it cannot be ruled out yet.

One of most important changes between the electron distribution parameters in the two states is the escape time-scale, t_{esc}, as it determines the relative optical depths of the thermal and non-thermal components of the evolved electron distribution. In the low-hard state we clearly have a thermal component with a greater optical depth than

in the high-soft state. It is difficult to provide a physical explanation or why the escape time-scale should change in this way, or at all. However, as the physical basis for having an escape time-scale is fairly weak anyway we see no need to explain the cause of this parameter varying. As mentioned previously, an ideal model would avoid the need for an escape time-scale by creating a non-thermal component from the existing distribution rather than by an injection mechanism.

4.8   Conclusions

Much more work clearly needs to be done in modelling the state transitions of X-ray binary systems. Here we have presented three potential models for explaining the nature of the X-ray spectrum, none of which are entirely successful. The external corona model cannot sustain the high-temperature thermal distribution of electrons required to explain the low-hard state spectrum in terms of thermal Comptonisation, as previously shown by Dove *et al.* (1997). Maintaining high-temperature thermal electron distributions is the greatest challenge to thermal Comptonisation models. Having hybrid distributions containing non-thermal components helps as it allows the Comptonised spectra to be recreated with lower electron temperatures (see Coppi 1999). However, energy losses due to Comptonisation can prevent even these lower temperature distributions from being maintained. Gierlinski *et al.* (1999) solve this problem by providing an extra source of heating to the electron distribution. They propose that the magnetic reconnection events responsible for creating the non-thermal component also create hot 50 MeV ions that heat the electrons through Coulomb scattering.

Our solution to the problem of maintaining a high-temperature thermal distribution was to lessen the intensity of the seed photons and hence minimise the Compton cooling through our internal corona geometry. We have found that this internal corona model requires extremely long scattering path lengths to recreate the Comptonised component of the spectrum. Other research, for example Kazanas, Hua, & Titarchuk (1997), has suggested that the time variability of the X-ray spectrum also can only be explained for sources that span 10^4 r_s, and similarly we expect accretion disc coronae to be of substantial

size, $\sim 10^3 \ r_s$, (Church 2001) to explain the eclipse dipping behaviour. However, such large coronae would not be compatible with the geometrical model used in our internal corona model. Furthermore, we find that in our internal injection models, which do not have geometrical size constraints, Compton cooling is weak leading to the evolution of thermal electron distributions with temperatures much greater than those observed.

As an alternative, long scattering path lengths could be modelled by replacing the internal spherical corona geometry with that of a conical jet that in the vertical direction extends to radii much greater than $10^4 \ r_s$. Again the accretion disc seed photons could be injected into the jet in the same manner as for the internal corona model, but in this case there would be no geometrical limit to the size of the corona as it would never obscure the accretion disc. However, a conical geometry for a Comptonising corona is not a gravitationally stable configuration and so the plasma would have to be *dynamical*, as we would expect from a relativistic jet. At present the HEART code cannot model dynamical plasmas. On this topic, there has been much interest in explaining the X-ray spectrum in the low-hard state in terms of a jet rather than a traditional Comptonising corona in recent years. There is evidence of the existence of a jet in many X-ray binary systems in this state, and the model of Markoff *et al.* (2001, 2003, 2004) can explain the low-hard state spectrum purely as synchrotron emission from the jet with the high-energy tail a result of synchrotron self-Compton emission.

In this chapter we have studied models whereby the low-hard state, high-energy X-ray spectrum is produced by thermal Compton scattering of blackbody seed photons. The shape of the spectrum, and specifically the peak of the thermal "hump", can only be recreated by plasmas containing a thermal electron distribution with a certain specific temperature. Therefore, before such models can be conclusively accepted as a plausible means of explaining the origin of the high-energy X-rays it is vital that they can prove that the existence of a thermal electron distribution of that specific temperature in the plasma is physically justifiable. At present, I believe the models detailed here do not provide a sufficiently complete description of all the relevant heating and cooling processes, to say this with any certainty. Most importantly I find the continual injection and escape model to explain the presence of the non-thermal component of the electron distribution

unsatisfactory. Producing the high-energy power-law component of the electron distribution through modelling the acceleration process within the kinetic equation would both allow the model to maintain a closed electron population and provide a more physically accurate description of energy exchange within the electron population. Furthermore, the composition of the plasma also affects the temperature of the electron distribution. For example, if hot ions are present, which are not modelled here, they will also provide a means of heating the electron distribution.

Finally, in our models of accretion disc coronae we have thus far not included the effects of bremsstrahlung radiation. Bremsstrahlung emission both provides an extra source of seed photons, and an additional cooling effect upon the electron distribution. In practice, this means models that include the effects of bremsstrahlung require lower optical depths to recreate the same slope to the thermal hump in the spectrum, and also lead to slightly cooler thermal components within the hybrid electron distribution.

Thus far we have barely touched upon the capabilities of the HEART code to model observations of time variability in the emission from accretion disc coronae or other high-energy plasmas. One possible application of the HEART code would be to model the effect of magnetic reconnection events creating local flares in some part of the corona, and to determine how such an event would affect time-dependent emission from the corona. Such a model would make full use of the three-dimensional capabilities of the HEART code.

Chapter 5

Emission Models for Relativistic Jets

5.1 Introduction

Future development could allow the HEART model to form the basis of a detailed model for time-dependent emission from relativistic jets. In this section we provide a motivation for the construction of such a model by revealing the flaws inherent in currently used jet emission models. In section 5.2 we describe the famous 'flat' spectrum model of relativistic jets, followed by our development of a time-dependent version of this simple model in section 5.3. The failure of this type of model in explaining the behaviour of non-steady-state spectra is described in section 5.4.

5.2 The Blandford-Königl steady-state jet model

The presence of a relativistic jet in an unresolved astrophysical source is recognised by its signature 'flat' spectrum. This spectral diagnostic was the result of the first emission model for a relativistic jet, calculated by Blandford & Königl (1979, hereafter BK79). At this time relativistic jets from active galactic nuclei (AGN) had been resolved in radio

images and it was known that these jets were narrow, conical streams of plasma, ejected from the AGN at relativistic velocities. BK79 proposed a model where the dominant emission process in relativistic jets is synchrotron radiation. As explained in section 2.7, synchrotron emission occurs when relativistic electrons are in the presence of a magnetic field; conditions that we believe exist in relativistic jets.

Relativistic jets from AGN are less dense than their surrounding medium, which causes internal shock waves to propagate through the jet as it penetrates into its surroundings. These shock waves energise the electrons in the jet plasma to ultra-relativistic energies via the process described in section 2.3.2. The resulting electron distribution follows a power-law given by equation 2.4, with an index of $q \approx 2$ for acceleration by mildly-relativistic shocks, and $q \approx 2.25$ in the ultra-relativistic limit.

It is believed that collimated formation of relativistic jets is guided by a twisted magnetic field that extends throughout the length of the jet. This magnetic field provides the second ingredient required for the emission of synchrotron radiation. For simplicity this magnetic field is assumed to be tangled on small scales, and hence anisotropies in the electron pitch angles may be ignored. This assumption is justified by observations showing that the jet emission is only weakly polarised, indicating minimal homogeneity in the magnetic field. The specific intensity of synchrotron emission from a power-law distribution of electrons in a tangled magnetic field was calculated in section 2.7.

These ingredients provide the three physical parameters of a synchrotron emitting jet plasma model – the electron number density, n, magnetic field strength, B, and jet half-width, w, (see Figure 5.1). As the jet plasma propagates outwards it expands; decreasing the electron density and weakening the magnetic field strength. Therefore, the jet plasma parameters all vary with radius in a way which may be expressed as a power-law dependence

$$n(r) \propto r^{\alpha_n}, \ B(r) \propto r^{\alpha_B} \ \text{and} \ w(r) \propto r^{\alpha_w}.$$

Hence,

$$n(r) = n_* \left(\frac{r}{r_*} \right)^{\alpha_n}, \ B(r) = B_* \left(\frac{r}{r_*} \right)^{\alpha_B} \ \text{and} \ w(r) = w_* \left(\frac{r}{r_*} \right)^{\alpha_w},$$

where r_* is an arbitrary radius at which we normalise the jet parameters as $n_* = n(r = r_*)$,
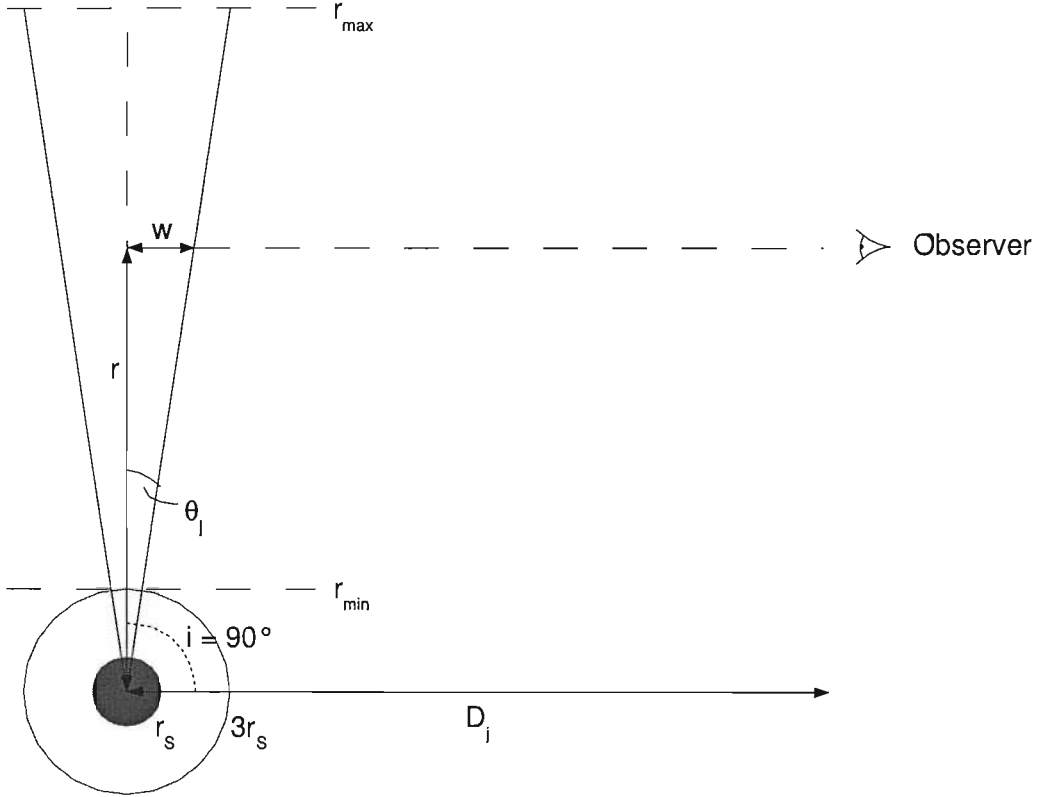
Figure 5.1: A schematic diagram of the jet emission model's geometry. We model the jet as a cone, of fixed opening angle, θ_j, inclined at 90° to the observer's line of sight. The emitting region exists between $r_{min} < r < r_{max}$, where r_{min} is assumed to be the radius of the last stable orbit around a non-rotating black hole, which is 3 r_s, or three times the radius of the black hole's event horizon (the Schwarzschild radius, r_s).

$B_* = B(r = r_*)$, and $w_* = w(r = r_*)$. Typically in a jet model we use the base of the jet as this radius, hence $r_* = r_{\text{base}}$. A measure of each parameter that describes its radial variation independent of r_* may also be defined in the following manner, e.g.

$$w_0 = \frac{w_*}{r_*^{\alpha_w}}. \tag{5.1}$$

We can now see how the source function and optical depth of the jet plasma vary with radius, from equations 2.35 and 2.36 respectively,

$$S_\nu(r) = 1.00 \times 10^{-36} \, c_s(q) \, \frac{b_*}{w_*} \left(\frac{r}{r_*} \right)^{-(\alpha_B/2)} \nu^{2.5} \; [\text{Wm}^{-2}\text{Hz}^{-1}\text{sr}^{-1}] \tag{5.2}$$

and

$$\tau_\nu(r) = 9.075 \times 10^3 \, c_\tau(q) \, a_* \left(\frac{r}{r_*}\right)^{\alpha_n + \alpha_B \frac{q+2}{2} + \alpha_w} \nu^{-\frac{q+4}{2}}, \qquad (5.3)$$

where we have defined two parameters that determine the normalisation of the optical depth and flux,

$$a_* = n_* B_*^{\frac{q+2}{2}} w_* \qquad (5.4)$$

and

$$b_* = B_*^{-0.5} w_* \qquad (5.5)$$

respectively. In this model, we have assumed that the jet's inclination to our line of sight is 90°, and that a photon's escape path-length, l, is on average equal to the jet's half-width, w.

To determine the total jet spectrum we must calculate the flux, and hence integrate the specific intensity over all solid angles. The range of solid angles, $d\Omega$, over which an an individual radial jet element, dr, emits is given by

$$d\Omega = \frac{2\pi w(r) dr}{D_j^2}, \qquad (5.6)$$

where D_j is the observer's distance from the jet. Assuming that the jet plasma is homogeneous in the lateral direction within each of these radial segments, we may calculate the total flux emitted by a relativistic jet by placing equation 5.2 into the specific intensity equation 2.16, and integrate with respect to $d\Omega$, to find

$$F_\nu = \frac{2\pi 10^{-36} c_s(q)}{D_j^2} b_{\mathrm{base}} \, \nu^{2.5} \int_{r_{\mathrm{base}}}^{r_{\mathrm{top}}} \left(\frac{r}{r_{\mathrm{base}}}\right)^{\alpha_w - (\alpha_B/2)} \left[1 - e^{-\tau_\nu(r)}\right] dr \; [\mathrm{W m^{-2} Hz^{-1}}]. \qquad (5.7)$$

In the BK79 model the jet expands adiabatically, with constant opening angle, resulting in a conical shape where the jet's width increases in direct proportion to its length, so $\alpha_w = 1$. This expansion is restricted to the two dimensions perpendicular to the jet's propagation direction, and any energy losses due to the adiabatic expansion are compensated for by an unspecified re-energisation process. Two-dimensional volume expansion will reduce the electron density as the inverse square of the radial position, hence $\alpha_n = -2$, and will reduce the magnetic field energy density in the same proportion, so the magnetic field strength decreases as the square-root of that rate, and hence $\alpha_B = -1$. These values

represent the 'equipartition case', where there is an equipartition of energy between the magnetic field and the electrons at all times.

Only for these values will a jet emission model produce a flat spectrum, see Figure 5.2a. If energy losses occur in the jet plasma as it propagates outwards, be it adiabatic, radiative, or otherwise, then the magnitude of these power-law indices will increase leading to an inverted spectrum. The flat spectrum is formed as a result of the partially self-absorbed synchrotron emission spectrum shifting to lower frequencies as the optical depth decreases along the jet's radius. A decrease in the optical depth will also lead to a decrease in peak flux, but for the values of α_n, α_B, and α_w used in the BK79 model, the flux normalisation increases with the exact radial dependence required to compensate for this loss of flux. Hence, the flux at the peak frequency remains constant whilst the peak frequency decreases, giving a broad total spectrum with a power-law index of zero.
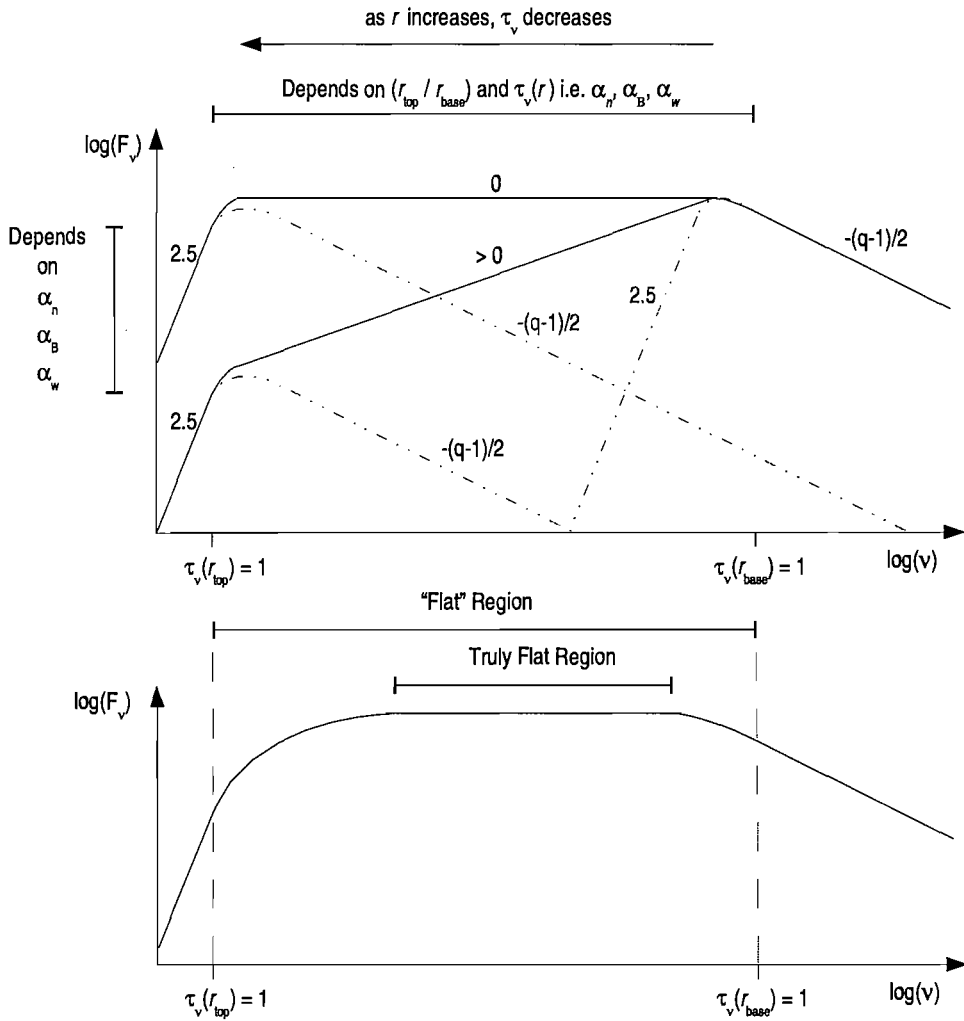
This effect can be seen by an approximate analytical solution to equation 5.7. Let us assume the jet extends out beyond the maximum observable radius for all frequencies of interest. This radius, $r_{\mathrm{max\,ob}}(\nu)$, is defined as the radius where the magnetic field strength is too weak to produce synchrotron radiation of frequency, ν. If the critical frequency (see equation 2.24) of the most energetic electrons falls below the observing frequency, ν, then the maximum radius from which we receive radiation of frequency ν is reached. The energy of the most energetic electrons, γ_{max}, does not vary with radius in the BK79 model (a non-lossy jet) and therefore the critical frequency as a function of radius only depends upon the magnetic field strength, $\nu_{\mathrm{c}}(r) \propto B(r)$. Hence, $r_{\mathrm{max\,ob}} \propto \nu^{1/\alpha_B}$.

The flux integration may now be approximated by separating the optically thin and optically thick regions of the jet at the boundary given by $\tau_\nu(r) = 1$. From equation 5.3 it can be seen that this radius, $r_{\mathrm{th}} \propto \nu^{\alpha_{\mathrm{th}}}$, where $\alpha_{\mathrm{th}} = (q+4)/(2\alpha_n + (q+2)\alpha_B + 2\alpha_w)$. Therefore, assuming $q = 2$, and employing the optical depth limits given in section 2.5,

$$F_\nu \quad \propto \quad \int_{r_{\mathrm{th}}(\nu)}^{r_{\mathrm{max\,ob}}(\nu)} r^{\alpha_n + 1.5\alpha_B + 2\alpha_w}\, \nu^{-0.5}\, \mathrm{d}r + \int_0^{r_{\mathrm{th}}(\nu)} r^{\alpha_w - 0.5\alpha_B}\, \nu^{2.5}\, \mathrm{d}r \qquad (5.8)$$

$$\propto \quad \nu^0,$$

for the BK79 power-law index values only.

as r increases, τ_ν decreases

Depends on (r_{top}/r_{base}) and $\tau_\nu(r)$ i.e. α_n, α_B, α_w

$\log(F_\nu)$

Depends on α_n α_B α_w

2.5

0

2.5

>0

2.5

$-(q-1)/2$

$-(q-1)/2$

2.5

$-(q-1)/2$

$\tau_\nu(r_{top}) = 1$ $\tau_\nu(r_{base}) = 1$ $\log(\nu)$

"Flat" Region

Truly Flat Region

$\log(F_\nu)$

$\tau_\nu(r_{top}) = 1$ $\tau_\nu(r_{base}) = 1$ $\log(\nu)$

ν scale is determined by normalisation of τ_ν, hence a_{base} for given α_n, α_B, α_w

F_ν scale is determined by S_ν and τ_ν normalisation, hence both a_{base} and b_{base} for given α_n, α_B, α_w

Figure 5.2: (a - top plot) Steady-state spectra from jet emission models. The decrease in optical depth with radius along the jet shifts the self-absorbed synchrotron spectrum (see Figure 2.5) to lower frequencies. If the jet extends to a radius, r_{top}, that is many orders of magnitude greater than its initial radius, r_{base}, the total spectrum will be very broad, consisting of a convolution of synchrotron spectra from differing radii. In the equipartition case (i.e. a BK79-type jet) the radial dependencies of the plasma parameters cause a flat spectrum to be produced − the peak flux remains constant as the frequency of the peak decreases. In a lossy jet an inverted spectrum is formed. (b - bottom plot) The flat region of the spectrum is not truly flat throughout the frequency space between $\tau_\nu(r_{top}) = 1$ and $\tau_\nu(r_{base}) = 1$. This makes it harder to distinguish between a mildly inverted spectrum and short flat spectrum, when there are only a few data points.

The spectra of finite length relativistic jets may be calculated by numerically integrating equation 5.7. As illustrated in Figure 5.2a, a jet of finite length only has a flat spectrum over a finite frequency range. It encompasses every frequency for which the jet material, over the length of the emission region, is optically thick at the smallest radius and has become optically thin at the largest radius. We may estimate the size of this flat region by assuming it extends over the frequency range given by $\tau_\nu(r_{\mathrm{top}}) = 1$ for ν_{min} and $\tau_\nu(r_{\mathrm{base}}) = 1$ for ν_{max}. Hence,

$$\frac{\nu_{\mathrm{max}}}{\nu_{\mathrm{min}}} = \left(\frac{r_{\mathrm{base}}}{r_{\mathrm{top}}}\right)^{1/\alpha_{\mathrm{th}}}, \tag{5.9}$$

where $\alpha_{\mathrm{th}} = -1$ for a BK79-type jet. These values can then be scaled to the desired frequency regime by a_{base}, with the flux normalisation determined by b_{base}. This frequency range for the flat region of the spectrum is not accurate, because the jet material becomes optically thin when $\tau_\nu(r_{\mathrm{top}}) \ll 1$ rather than when $\tau_\nu(r_{\mathrm{top}}) < 1$. Therefore, the flat region is in fact slightly shorter, as illustrated in Figure 5.2b.

5.3   A time-dependent emission model for equipartition jets

In Collins *et al.* (2003) we developed the basic BK79 model of equipartition (or non-lossy) jets into a time-dependent emission model, with the aim of explaining the time variability of jet spectra. This can simply be achieved by making either or both of the free physical parameters of the model vary with time. However, it would be impossible for the parameter values to vary along the entire length of the jet instantaneously because a relativistic jet extends over such a great distance that such behaviour would require information to travel faster than the speed of light. Therefore, we change the parameter values of the plasma that is being *injected* at the base of the jet with time, and allow this plasma to propagate outwards at the jet's bulk velocity, v_{j}.

Either an increase in the value of a_{base} or b_{base} with time will create a flux increase. Increasing b_{base} has the advantage that the entire flat spectrum will raise in flux, and hence all frequencies will increase by the same amount, at the same time. However, a_{base} would have to remain constant, which would require a increase in the injected electron density

of the appropriate rate to exactly counter the decrease in the magnetic field strength (by the definitions giving in equations 5.4 and 5.5). From a physical viewpoint this seems to be an unlikely scenario, so we choose to vary the value of a_{base}, which may be interpreted as an increase in the injected electron density.

By allowing $a_{base} = a_{base}(t)$, we introduce a time dependence to the observed flux function (equation 5.7). However, if the injected electron density at the base of the jet varies with time, then it will also vary with radius, as there is a lag in the time for the injected electrons to reach greater radii along the jet. Therefore, we introduce a transformed variable for the time, t,

$$t'(t,r) = t - \frac{r - r_{base}}{v_j}.\tag{5.10}$$

Hence equation 5.7 becomes

$$F_\nu(t) = \frac{2\pi 10^{-36} c_s(q)}{D_j^2}\, b_{base}\, \nu^{\frac{5}{2}} \int_{r_{base}}^{r_{top}(t)} \left(\frac{r}{r_{base}}\right)^{\frac{3}{2}} \left[1 - e^{-\tau_\nu(t,r)}\right]\, dr\ \ [\text{Wm}^{-2}\text{Hz}^{-1}],\tag{5.11}$$

with

$$\tau_\nu(t,r) = 9.075 \times 10^3\, c_\tau(q)\, a_{base}(t') \left(\frac{r}{r_{base}}\right)^{-\frac{q+4}{2}} \nu^{-\frac{q+4}{2}},\tag{5.12}$$

where $r_{top}(t) = r_{base} + v_j t$, and we have used the equipartition values for the radial power-law indices. We have also assumed a constant bulk velocity of the jet material, v_j.

As we have no prejudice towards the functional dependence of the injected electron density with time we have investigated three scenarios. Although each scenario describes an increase in injected electron density from zero, in reality we expect an increase from some quiescent electron density. The resulting spectral behaviour of the model is the same, but in the latter case our results would just be superimposed upon a pre-existing steady flat spectrum. This underlying flat spectrum would be more extended then that which can vary in the observed timescale. The three scenarios are, with t expressed in units of t_{max} (the time of the maximum injected electron density):

(i) A simple spontaneous increase in the injected electron density, followed by a spontaneous decrease

$$a_{base}(t) = \begin{cases} a_{max} & ;\ \ 0 \le t \le 1, \\ 0.0 & ;\ \ t > 1. \end{cases}\tag{5.13}$$

This model will effectively produce the steady-state model spectra of Figure 5.2 at, and only at, $t = 1$.

(ii) A linear increase in the injected electron density, followed by a linear decrease

$$a_{\text{base}}(t) = a_{\text{max}} \left(1 - |t - 1|\right). \tag{5.14}$$

(iii) A Gaussian electron density injection function

$$a_{\text{base}}(t) = a_{\text{max}} \, \mathrm{e}^{-25(t-1)^2}. \tag{5.15}$$

The time varying flux was calculated by performing a Romberg numerical integration (e.g. Press et $al.$ 1992) of equation 5.11 over the time period between $t = 0$ and $t = 2\,t_{\text{max}}$, for just two frequencies, ν_{low} and ν_{high}. For each scenario, the simulation was performed for two values of a_{max}, calculated such that ν_{high} always describes the light curve of a frequency that lies within the flat region of the steady-state spectrum. However, the higher value of a_{max}, a_{high}, shifts the spectrum such that ν_{low} describes the light curve of a frequency that lies within the optically thick cut-off region, and for a_{low}, ν_{low} also lies within the flat region. The steady-state spectra for both values of a_{max} are shown in Figure 5.3. The resulting light curves for each scenario are shown in Figures 5.4, 5.5 and 5.6, respectively.
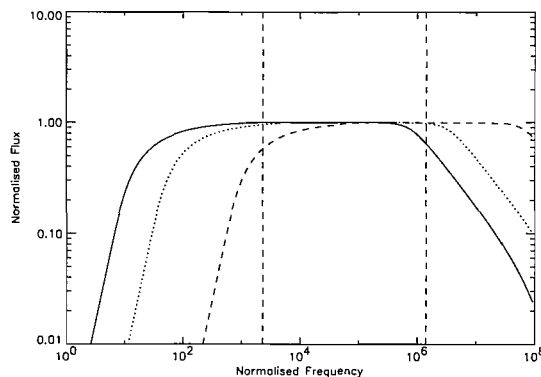

Figure 5.3: Normalised steady-spectra illustrating the positions of ν_{high} and ν_{low} (denoted by the vertical dashed lines) relative to the flat region for two values of a_{base}, a_{high} (dashed line) and a_{low} (dotted line), as used in the following light curve simulations.
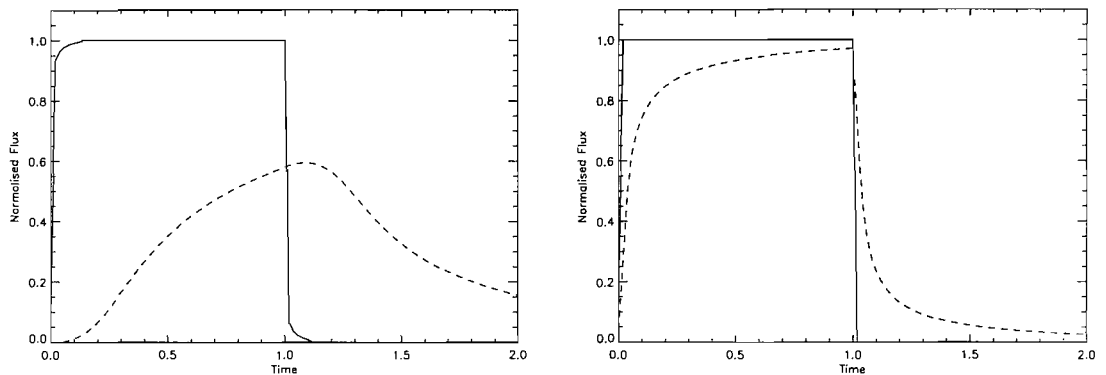
Figure 5.4: Model light curves for ν_{high} (solid line) and ν_{low} (dashed line), over a period of $2\,t_{\text{max}}$, for (a - left panel) $a_{\text{max}} = a_{\text{high}}$ and (b - right panel) $a_{\text{max}} = a_{\text{low}}$, defined in Figure 5.3, with a constant value of injected electron density.
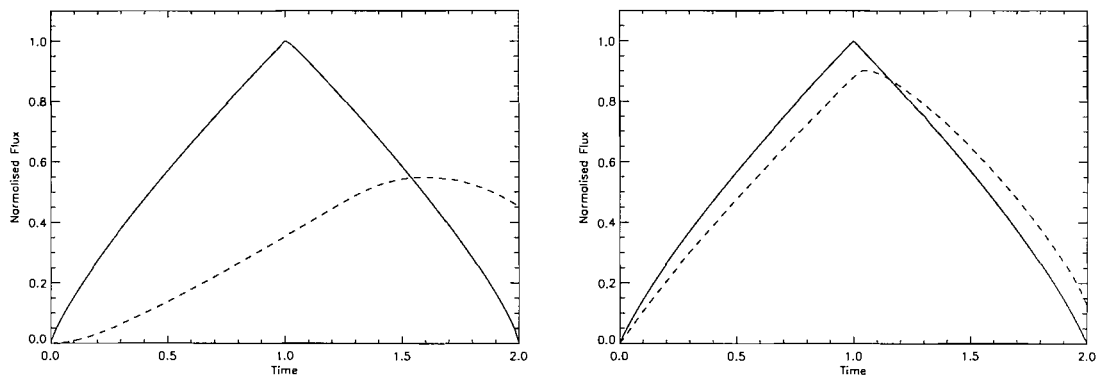

Figure 5.5: Model light curves for ν_{high} (solid line) and ν_{low} (dashed line), over a period of $2\,t_{\text{max}}$, for (a - left panel) $a_{\text{max}} = a_{\text{high}}$ and (b - right panel) $a_{\text{max}} = a_{\text{low}}$, defined in Figure 5.3, with a linear electron density injection profile.

Figure 5.6: Model light curves for ν_{high} (solid line) and ν_{low} (dashed line), over a period of $2\,t_{\text{max}}$, for (a - left panel) $a_{\text{max}} = a_{\text{high}}$ and (b - right panel) $a_{\text{max}} = a_{\text{low}}$, defined in Figure 5.3, with a Gaussian electron density injection profile.
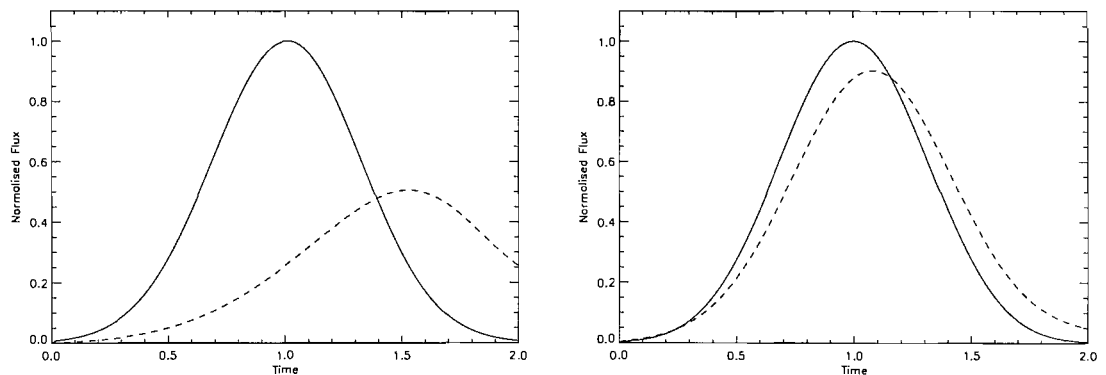
As the increase in electron density propagates along the jet, a flat spectrum, which initially forms at the high-frequency end, extends towards lower frequencies. The light curves peak at the time when the flat spectrum has extended to include the frequency of observation. Hence, the light curve of the higher frequency emission always peaks before that of the lower frequency emission because it originates from smaller radii within the jet. The profile of the light curve in Figure 5.4b clearly demonstrates how a flat spectrum is created as a function of time. The flat region of the spectrum is formed very rapidly at the high-frequency end, but its progression towards lower frequencies slows down over time. When the electron injection is switched off, the high-frequency end of the flat spectrum very rapidly disappears but the lower frequency region decays more gradually.

For higher values of a_{max}, the flat spectral region is shifted towards higher frequencies, so a given frequency peaks at a later time. Furthermore, for the lower frequencies, such as ν_{low} in Figure 5.4, the flat spectral region never reaches the frequency of observation, and so a slow flux increase, followed by a slow decay is observed. Conversely, as the observation frequency approaches the optically thin cut-off frequency, the observed light curve profile will tend towards the injection function profile. This may also be described by the optically thin radius for the frequency of observation, ν, approaching the base of the emission region, $r_{th}(\nu) \to r_{base}$.

The linear and Gaussian injection functions both display light curve profiles, shown in Figures 5.5 and 5.6, that tend towards the injection profile as $r_{th}(\nu) \to r_{base}$. The time-lag observed between the flux peaks of the two frequencies can be reduced by decreasing a_{max}, such that both frequencies lie closer to the, rapidly formed, high-frequency end of the flat spectrum. This also has the effect of decreasing the ratio between the two frequencies' peak fluxes, until both frequencies lie in the flat region where it becomes unity (assuming the emission region is sufficiently large). For sufficiently small values of a_{max} the peak flux ratio will begin to decrease again, as the higher frequency emission becomes optically thin.

This time variability model can be used to determine the physical properties of relativistic jets from observations of their light curves. We would expect such observations to be best modelled by the gradual increase and decrease of the injected electron density

given by the linear and Gaussian functions. Any such injection function would predict approximately the same time-lags between flux peaks, peak flux ratios, and flux normalisations for a given value of a_{max} only with differing light curve profiles. Therefore, having measured the rise time of an observed flare to give t_{max} we can then calculate light curve models to find the value of a_{max} that agrees with the observed time-lags (or peak flux ratios) between a given set of frequencies. From the flux normalisation the corresponding value of b_{base} may be found, and hence from these two parameters B_{base} and n_{max} may be extracted. In the following section we employ this technique to analyse observations of flaring in the emission from the relativistic jet source GRS 1915+105.

5.4   Modelling flares in the jet emission from microquasar GRS 1915+105

5.4.1   Observations

On 20th May 1999, Fender & Pooley (2000) observed the microquasar GRS 1915+105 simultaneously at two wavelengths: at 2.2-μm with the IRCAM3 instrument on the United Kingdom Infrared Telescope (UKIRT), and at 1.3-mm with the SCUBA instrument on the James Clerk Maxwell Telescope (JCMT). The resulting light curves recorded over a two hour period are shown, overlaid, in Figure 5.7. The light curves show large-amplitude quasi-periodic oscillations, each lasting ~ 1000 s, believed to be ejection events in the jet (see section 1.2.1). Detailed views of the first and fourth peaks are shown in Figure 5.8.

The infrared data from UKIRT have been dereddened with an assumed infrared K-band extinction of $A_K = 3.3$ mag. However, there is an uncertainty in the infrared flux values of at least 40% (Fender et $al.$ 1997) as the precise value for the absorption correction is not known. Furthermore, we expect that the infrared flux will have some unknown background contribution due to the emission from regions of the GRS 1915+105 system other than the jet. The error values on individual data points are of the order of 5% (Fender, private communication).
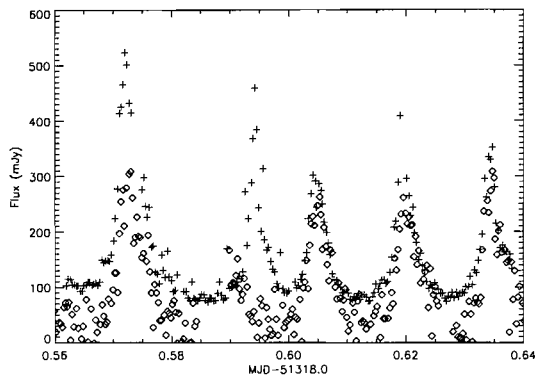
Figure 5.7: Infrared 2.2-μm data (crosses) from the IRCAM3 instrument on UKIRT, and millimetre 1.3-mm data (diamonds) from the SCUBA instrument on JCMT. Taken from Fender & Pooley (2000).

The millimetre data show flux peaks that occur almost simultaneously with the infrared peaks. The time-lag between the millimetre and infrared peaks is less than 50 s, though the sampling rate is not sufficient to provide a more exact determination. Importantly, the data reveal a ratio between the infrared flux and the millimetre flux that is very close to unity. It is unlikely that the flux from an astrophysical object at two such different frequencies covering three orders of magnitude would be so close unless the spectrum is that of a relativistic jet, which has a very small or zero spectral index (see section 5.2). This information agrees with the belief that such relatively low frequency emission from microquasars is synchrotron radiation from the system's relativistic jet (see section 1.2.1). Therefore, we believe that our time-dependent model of relativistic jet emission should be capable of modelling this data, despite only having knowledge of the emission from just two frequencies.

5.4.2   Modelling the flares

Figure 5.8 clearly shows light curve profiles of the flares that differ from our time-dependent jet emission model. The observations show a rapid rise in flux followed by a slow decay. This may be indicative of the time-dependent behaviour of the injected electron density. However, we believe that the basic parameters − time-lags, flux ratios and normalisations − are independent of the light curve profile and should still be accurately modelled by the
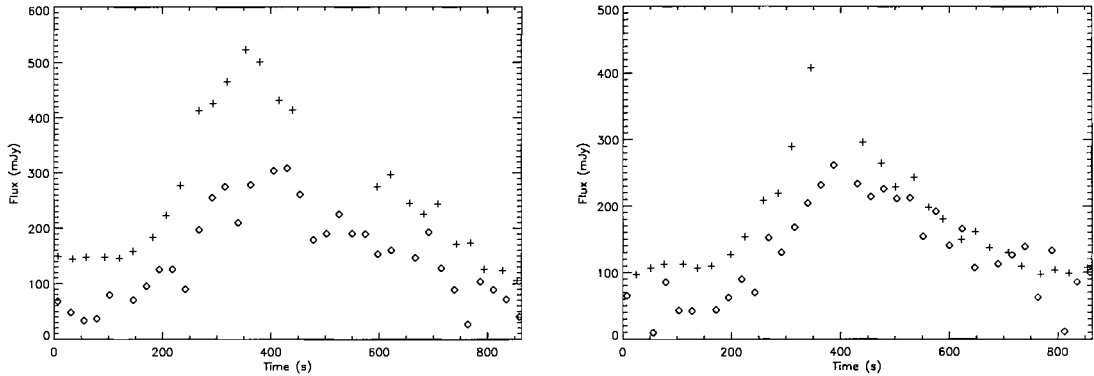
Figure 5.8: Detailed view of (a - left panel) first flux peak (b - right panel) fourth flux peak in the data shown in Figure 5.7.

Gaussian variant of our model.

Non-unity flux ratios are a product of our time-dependent jet emission model when the flat region of the spectrum does not extend fully to cover both frequencies of observation. Hence, one frequency will have a lower flux as no region of the jet has an optical depth of unity to that frequency. We note here that such close to unity flux ratios may also be explained by a mildly inverted spectrum from a lossy jet that extends over a greater range of radii.

To model these flares we assume that the electron acceleration is caused by ultra-relativistic shocks, and hence $q = 2.25$. In practice changing this value to the mildly-relativistic shock value results in no significant difference to the model. We assume a constant bulk velocity for the jet material, using the same value as Kaiser *et al.* (2000), of $v_j = 0.6c$. Our model implicitly assumes that we view the jet at an inclination of $i = 90°$, which is in reasonable agreement with observed inclinations of $i \approx 70°$ (Mirabel & Rodríguez 1994). The opening half-angle of the jet, θ_j, is observed to be less than 4° (Fender *et al.* 1999), so we take $w_0 = 0.05$ (see equation 5.1). For the distance to GRS 1915+105, D_j, we assume 11 kpc (Fender *et al.* 1999).

We assume that the radius of the initial shock acceleration, r_{base}, is where the relativistic electrons are first injected and thus is limited to the size of the last stable orbit for a non-rotating black hole, which is approximately three times the Schwarzschild radius, r_s.

w_0	D_j	v_j	r_{base}	t_{max}
0.05	11 kpc	0.6 c	1.0×10^5 m	350 s

Table 5.1: Fixed parameters of the jet model, determined by observations of GRS 1915+105.

Taking the lower mass limit for the GRS 1915+105 black hole of 10 $M_\odot$ (Greiner *et al.* 2001), gives the lowest plausible value for $r_{\text{base}} = 3r_s = 1.0 \times 10^5$ m.

The value of t_{max} is assumed to equal the rise time of the observed infrared flux peaks, which is ~ 350 s. The longer decay period in the observed light curve profiles is similar to model profiles shown in Figure 5.4 for those frequencies to which the jet material is always optically thick. Hence, this would suggest the flat spectrum extends to higher frequencies, giving a delay between t_{max} and the time of peak infrared flux. However, the short time delays between the infrared and millimetre peak fluxes can only be created for values of a_{max} that place the infrared frequency within the flat spectral region. Therefore, the infrared flux rise time should be approximately equal to t_{max}.

The observationally determined fixed model parameters are summarised in Table 5.1. Our model neglects any relativistic effects on the emission spectra, such as Doppler boosting and time dilation, as these will be small for the adopted parameters (see section 5.4.3). Using these model parameters we performed several simulations of the predicted light curves for the infrared (1.4×10^{14} Hz) and the millimetre (2.3×10^{11} Hz) frequencies with different values of the a_{max} parameter. This resulted in a relationship, for the fixed parameters, between the time-lag and the ratio of the peak fluxes, shown in Figure 5.9a. The grey rectangle denotes the region of parameter space that is consistent with the observation data, and shows that the model is inconsistent with the observations. However, we have not considered the uncertainty in the infrared flux values due to the imprecise extinction measure in this waveband (see section 5.4.1). Furthermore, we would expect the donor star in the GRS 1915+105 system, along with other 'background' sources, to contribute to the quiescent infrared emission. Therefore, if we believe the model to be correct for these values of the fixed parameters, we can use this relationship to determine the jet's contribution to the infrared emission. Following this premise we find that

synchrotron emission from the jet accounts for approximately 70% of the observed peak infrared emission. The infrared flux between flares is approximately 150 mJy, or 30% of the peak emission, and therefore it seems likely that this represents a non-varying background component. Subtracting this component from the data will make the model agree with the observations.

Either the observed time-lag or the ratio between the flux peaks may be used in the modelling process to determine the value of a_{max} for the system. Since the flux ratio is relatively dubious due to the uncertainty in the infrared flux, we choose to use the time-lag to determine a_{max} with greater confidence. It is also preferential to use the time-lag as this is a quantity specific to a time-dependent model, whereas the flux ratio is also an observational result of the steady-state model.

From the simulations we can see that the optical depth parameter values that agree with the observed time-lag of 25 ± 25 s are given by $\log(a_{max}) = 31.5 \pm 1.5$ (in SI units). The flux normalisation parameter, b_{base}, may then be determined by fitting the model's
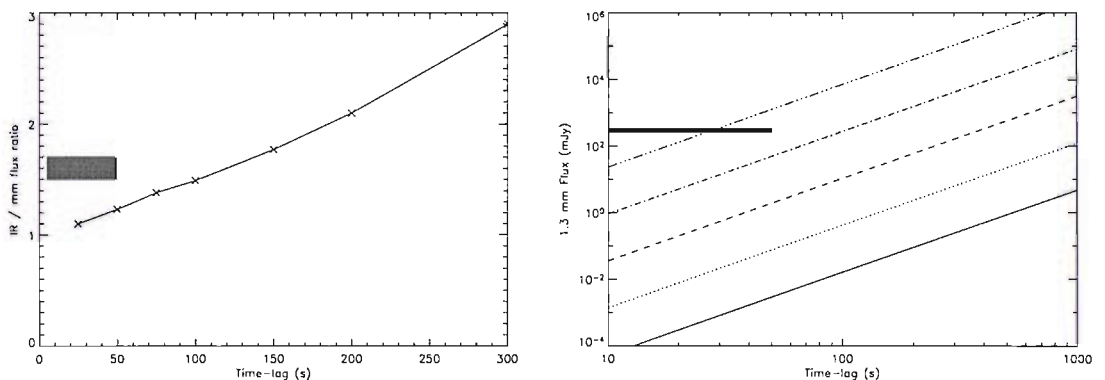

Figure 5.9: (a - left panel) The predicted time-lag and flux ratios between the flux peaks at 1.4×10^{14} Hz and at 2.3×10^{11} Hz for $3 \times 10^{31} \leq a_{max} \leq 10^{37}$. The grey rectangle shows the region where the predicted relationship would agree with the data. (b - right panel) The predicted flux for 1.3-mm emission as a function of the time-lag between the infrared and millimetre flux peaks. Each line represents a different maximum injected electron density, $n_{max} = 10^{16}$ (bottom, solid line), $10^{22}, 10^{28}, 10^{34}$, and 10^{40} (top, triple dot-dashed line) m^{-3}. The time-lags are solely determined by the values of a_{max}, and the grey rectangle denotes the values that are consistent with the observations.

light curve to the more reliable millimetre peak flux, which is consistently ~ 300 mJy. Our model, with a_{max} fixed at its most likely value, produced millimetre peaks of the observed flux level with $b_{\mathrm{base}} = 5.7 \times 10^6$ (in SI units), corresponding to a magnetic field strength of $B_{\mathrm{base}} = 7.8 \times 10^{-7}$ T. Hence, to produce the observed time-lags we need peak injected electron densities to reach $\log(n_{\mathrm{max}}) = 40.8 \pm 1.5 \log(\mathrm{m}^{-3})$. Clearly such high electron densities are physically unjustifiable. By performing a series of simulations with varying a_{max} for different fixed peak injected electron density values, we produced the plot shown in Figure 5.9b. This plot clearly illustrates the inescapable requirement of our model for such high electron density values.

Finally, to maintain the emission of the highest frequency radiation at the outer edge of the jet, the electron energy distribution must extend to a Lorentz factor of $\gamma_{\mathrm{max}} = 10^7$ for our fitted value of the magnetic field strength. Although this value is high there is some evidence, from extragalactic jet observations (Dermer & Atoyan 2002), that electrons in jets can be accelerated to such energies. However, this high value of γ_{max} is a direct consequence of the low fitted value of the magnetic field strength, which would be higher in a model that does not require extremely high electron densities.

5.4.3 Conclusions

From measurements of the time-lag observed between the peak infrared flux and millimetre flux, together with the reliable millimetre peak flux, we have ascertained that our time-dependent jet emission model is incompatible with the observations presented here for the relativistic jet source GRS 1915+105. Our model requires unphysical electron densities to be present within the jet to produce the observed flux for the given time-lags. In order for the model to recreate the observed time-lags with justifiable electron density values it must either predict a flux $\sim 10^4$ times larger for the fitted value of n_{max}, or else predict the observed time-lag with a value of a_{max} that is $\sim 10^9$ times larger.

We are quite confident that the fixed parameters of the model are reasonably accurate. In order to make the model fit the observations, either the distance to the jet, D_{j}, or the jet opening angle, given by w_0, must be decreased by many orders of magnitude, which

is unreasonable. Similarly, due to physical constraints, t_{max} and r_{base} cannot be altered sufficiently to significantly affect the model results. As a consequence of not modelling relativistic effects, any value for the bulk velocity, v_j, that is similar to the speed of light will produce approximately the same result.

However, the bulk gas flow within the jet is relativistic. Therefore, if relativistic effects can cause the observed time-lag to be considerably shorter than the time-lag in the rest frame of the jet material, then the observed flux will be predicted from a lower electron density. Although time dilation has the opposite effect, relativistic Doppler shifts can in principle cause this effect. However, in the case of the GRS 1915+105 jet, the inclination of the jet axis to our line of sight is well constrained to $\sim 70°$, and for such inclinations the maximum effect on the observed flux from relativistic Doppler shift and Doppler beaming is only 20%. Furthermore, the simplification in using the equivalent line of sight optical path-length for a jet at an inclination of 90° does not significantly effect the results. Only for substantially smaller viewing angles and highly relativistic jet velocities do relativistic effects lead to a significant shortening of observed time-lags compared to the jet rest frame.

Therefore, the model presented here will fit the observations only if the relativistic jet in GRS 1915+105 is inclined further towards us, at a smaller angle than the $\sim 70°$ that was measured by Mirabel & Rodríguez (1994). This is possible as there have been numerous examples recently of the precession of relativistic jets in other microquasar systems (see e.g. Romero et al. 2002), and there have not been enough measurements of GRS 1915+105 to date to rule out this possibility. However, until such a precession of the jet can be confirmed by other means we must assume that the equipartition time-dependent jet emission model is flawed.

The time variability model would be more convincing had it inherently reproduced the observed light curve profile, and this might provide a clue to the physical source of the time variability. It is possible that this profile is a representation of the time variability in the injected electron density. Although a function that produced the observed light curve profile could be constructed, it would not alter the results of the model fit. Alternatively, the slow decay phase may be an effect of the radiative transfer process in the optically thick jet, which has not been modelled here, but could be included by applying the HEART

model to relativistic jets.

It is possible that the fundamental equipartition description of relativistic jet emission is flawed. Such models require continuous, ubiquitous re-energisation of the electron distribution throughout the jet, and it is not clear how this could be achieved. Even the steady-state model of BK79 is incapable of producing such small flux ratios with such large fluxes. However, modifications to the time-dependent emission model to allow for the different radial dependences of the parameters to model lossy jets also failed to produce a viable model.

We have so far assumed that a near unity flux ratio between the two frequencies of observation implies a flat or mildly inverted spectrum. However, without observations at other frequencies we cannot be certain that this is true. One possibility is that the higher frequency emission is caused by the inverse Compton scattering of the lower frequency synchrotron emission. Such synchrotron self-Compton emission would produce the correlated time variability behaviour, but would not necessarily produce such small flux ratios.

In conclusion, either the observed infrared and millimetre flares in the emission from GRS 1915+105 do not originate from the system's relativistic jet or an alternative model is required to explain relativistic jet emission from microquasars. Observations of the jet spectrum and time variability at different frequencies will help resolve the issue of whether the spectrum is flat, and whether we are observing optical depth effects. Observing a relationship between the flux ratio and the time-lag would confirm that the flux ratio is due to an optical depth effect rather than due to a mildly inverted spectrum.

This result is true for any model that attempts to explain such correlated time variability in the emission between two different wavelengths (longer than or within the infrared band) from GRS 1915+105 as being due to a variation in the optical depth of a relativistic jet that obeys the physical model of BK79 and radiates predominantly synchrotron radiation with a flat spectrum. This physical time variation can be any propagating alteration in the optical depth along the length of the jet, be it due to a change in the density of the injected material or due to a shock wave propagating through the jet. The speed of

this propagation must be close to the speed of light to explain such a short delay between the flux peaking at each of the two wavelengths. This is because the peak flux at each wavelength is associated with a characteristic radius, which for such a large difference in scale between wavelengths leads to a large difference in scale between characteristic radii. Decreasing the speed of propagation places even greater demands upon the density of the jet material, and as it close to the speed of light in this model it cannot be increase significantly.

Acknowledgements

Firstly, I wish to acknowledge the helpful support of my two supervisors, Christian Kaiser, and Simon Cox, in guiding me through the maze of astronomy, physics, mathematics, computer science, bureaucracy and scientific writing required during my PhD research.

I would like to thank Sergei Nayakshin for taking the time to help me implement his Coulomb scattering equations. Rob Fender for allowing me to use his data in Chapter 5. Some of the information provided in the background chapter of this thesis was sourced from Charles & Seward (1995), and reviews by Mirabel & Rodríguez (1999) and Fender (2003). I am also grateful to Rob Hynes for his efforts in creating the BinSim program (Hynes 2001) that I have used to create the X-ray binary system visualisation diagrams.

On a personal note I'd like to thank my parents for making it possible for me to reach this stage of my education, and for providing me with a boat to live on for the final months of my research! Also my house-mates, Rich, Katie, and Kirsty for their superb company over these past few years. My friends and colleagues, Carolyn, Ed, Will, Nick, Jose, Andy, Mike, Remon, Colin, Katherine, Vanessa, Liz, Phil, Anna, Ralph and Jim to name but a few for the inspiration they have provided me, stimulating nights down the Crown, the multitude of squash matches and trips to the beach. And last, but not least, my ever changing office mates for providing me with such a great working environment, Julie, Simon, Diana, Steven, Will, and Ed.

Finally, I should thank the department of Electronics and Computer Science for funding my PhD through an EPSRC grant, and providing me with the Toshiba Satellite Pro laptop that allowed me to work in the Astronomy group and has somehow survived these three, computationally intensive, years!

Ross – November 2004

Appendix A

Formulae for exact Compton

scattering

To calculate the exact Compton scattering function, $xF_C(x)$, that includes both quantum mechanical (Klein-Nishina) and relativistic effects and is valid for all frequencies and energies we follow the procedure of Coppi & Blandford (1990) who use the corrected fomulae dervied by Jones (1968).

$$xF_C(x) = 3\sigma_T \, x^2 \int_{z_-}^{z_+} \frac{dP}{dz} dz. \tag{A.1}$$

$$
\frac{dP}{dz} = \frac{1}{(1-\beta z)\sqrt{\beta^2 + \epsilon^2 + 2\beta\epsilon z}} \times \left\{ 2y_0 k - ak^2 + \frac{1 + y_0^2 - 2ay_0 k + a^2 k^2}{\sqrt{a^2 - b^2}} \right.
$$
$$
\left. + \frac{1}{k^2(1-\beta z)} \left[k^2 + \frac{ak^2(2b-a)}{(a-b)\sqrt{a^2-b^2}} + \frac{a(1-y_0)^2 + 2kb^2(1-y_0) - ba^2 k^2}{(a^2 - b^2)^{3/2}} \right] \right\}, \tag{A.2}
$$

where $\epsilon = h\nu_0/\gamma m_e c^2$, $k = \gamma m_e c^2/h\nu_1$, $a = 1 - \beta z - (1-y_0)/k$, $b = \delta/k$, and

$$y_0 = \frac{(\epsilon + \beta z)(\rho + \epsilon\rho - 1 + \beta z)}{\rho(\beta^2 + \epsilon^2 + 2\beta\epsilon z)}, \tag{A.3}$$

$$\delta = \frac{\beta\sqrt{(1-z^2)[\rho^2\beta^2 + 2\rho\epsilon(1-\rho)(1-\beta z) - (\rho - 1 + \beta z)^2]}}{\rho(\beta^2 + \epsilon^2 + 2\beta\epsilon z)}, \tag{A.4}$$

with $\rho = \nu_1/\nu_0$. The integration limits are

$$z_\pm = \max\left[\pm 1, \frac{1 - \rho\left(d \mp \sqrt{d^2 - 1/\gamma^2}\right)}{\beta}\right],$$ (A.5)

where $d = 1 + \epsilon - \epsilon\rho$.

Bibliography

Abramowicz M.A., Kluźniak W., and Lasota J.-P., 2002, *A&A*, **396**, L31

Achterberg A., Gallant Y.A., Kirk J.G., and Guthmann A.W., 2001, *MNRAS*, **328**, 393

Bell A.R., 1978, *MNRAS*, **443**, 182

Blandford R.D., and Königl A., 1979, *ApJ*, **232**, 34

Blasi P., 2000, *ApJ*, **532**, L9

Blumenthal G.R., and Gould R.J., 1970, *Rev. Mod. Phys.*, **42**, 237

Chang J.S., and Cooper G., 1970, *J. Comp. Phys.*, **6**, 1

Charles P.A., and Seward F.D., 1995, *Exploring the X-ray Universe*, CUP

Chiaberge M., and Ghisellini G., 1999, *MNRAS*, **306**, 551

Church M.J., 2001, *AdSpR*, **28**, 323

Collins R.S., Kaiser C.R., and Cox S.J., 2003, *MNRAS*, **338**, 331

Coppi P.S., and Blandford R.D., 1990, *MNRAS*, **245**, 453

Coppi P.S., 1992, *MNRAS*, **258**, 657

Coppi P.S., 1999, in *High Energy Processes in Accreting Black Holes*, eds. J. Poutanen
    and R. Svensson, ASP Conf. Series, astro-ph(9903158)

Crusius A., and Schlickeiser R, 1986, *A&A*, **164**, L16

Davis S.W., Blaes O.M., Hubeny I., Turner N.J., 2005, *ApJ*, **621**, 372

Dermer C.D., Atoyan A.M., 2002, *ApJ*, **568**, L81

Dhawan V., Mirabel I.F., and Rodríguez L.F., 2000, *ApJ*, **543**, 373

Di Salvo T., Done C., Zycki P.T., Burderi L., and Robba N.R., 2001, *ApJ*, **547**, 1024

Dove J.B., Wilms J., Maisack M., and Begelman M.C., 1997, *ApJ*, **487**, 759

Fender R.P., Pooley G.G., Brocksopp C., and Newell S.J., 1997, *MNRAS*, **290**, L65

Fender R.P., Garrington S.T., McKay D.J., Muxlow T.W.B., Pooley G.G., Spencer R.E.,

Stirling A.M., and Waltman E.B., 1999, *MNRAS*, **304**, 865

Fender R.P., and Pooley G.G., 2000, *MNRAS*, **318**, L1

Fender R.P., 2001, *MNRAS*, **322**, 31

Fender R.P., 2003, in *Compact Stellar X-Ray Sources*, eds. W.H.G. Lewin and M. van der Klis, CUP, (astro-ph/0303339)

Fermi E., 1949, *Phys. Rev. 2nd Series*, **75**, 1169

Frank J., King A.R., and Raine D.J., 1992, *Accretion Power in Astrophysics*, Second Edition. CUP

Ghisellini G., Guilbert P.W., and Svensson R., 1988, *ApJ*, **334**, L5

Ghisellini G., George I.M., Fabian A.C., and Done C., 1991, *MNRAS*, **248**, 14

Ghisellini G., and Svensson R., 1991, *MNRAS*, **252**, 313

Gierlinski M., and Done C., 2004, *MNRAS*, **347**, 885

Gould R.J., 1980, *ApJ*, **238**, 1026

Greiner J., Cuby J.G., and McCaughrean M.J., 2001, *Nature*, **414**, 522

Haardt F., and Maraschi L., 1993, *ApJ*, **413**, 507

Hjellming R.M., and Johnston K.J., 1988, *ApJ*, **328**, 600

Hjellming R.M., and Han X., 1995, *X-ray Binaries*, eds. W.H.G. Lewin and M. van der Klis, CUP

Hynes R.I., 2001, in *The Physics of Cataclysmic Variables and Related Objects*, eds. B.T. Gaensicke, K. Beuermann, and K. Reinsch, ASP Conf. Series, (astro-ph/0110058)

Jones F.C., 1968, *Phys. Rev.*, **167**, 1159

Kaiser C.R., Sunyaev R., and Spruit H.C., 2000, *A&A*, **356**, 975

Kazanas, D., Hua X.-M., and Titarchuk L., 1997, *ApJ*, **480**, 735

Lemoine M., and Pelletier G., 2003, *ApJ*, **589**, 73

Le Roux R., 1961, *Ann. Astrophys.*, **24**, 71

Li H., Kusunose M., and Liang E.P., 1996, *ApJ*, **460**, L29

Longair M.S., 1994, *High-Energy Astrophysics*, Volume 2. CUP

Markoff S., Falcke H., and Fender R.P., 2001, *A&A*, **372**, L25

Markoff S., Nowak M., Corbel S., Fender R.P., and Falcke H., 2003, *A&A*, **397**, 645

Markoff S., and Nowak M.A., 2004, *ApJ*, **609**, 972

McConnell M.L., et al., 2002, *ApJ*, **572**, 984

McCray R., 1969, *ApJ*, **156**, 329

Mirabel I.F., and Rodríguez L.F., 1994, *Nature*, **371**, 46

Mirabel I.F., and Rodríguez L.F., 1999, *ARA&A*, **37**, 409

Nayakshin S., and Melia F., *ApJS*, **114**, 269

Park B.T., & Petrosian V., 1996, *ApJS*, **103**, 255

Petrosian V., 1981, *ApJ*, **251**, 727

Poutanen J., and Coppi P., 1998, *Phys.Scripta*, **T77**, 57

Poutanen J., and Zdziarski A.A., 2002, in *New Views on Microquasars*,
    eds. Ph. Durouchoux, Y. Fuchs, and J. Rodriguez,
    Center for Space Physics: Kolkata, (astro-ph/0209186)

Pozdnyakov L.A., Sobol' I.M. and Sunyaev R.A., 1977, *Soviet Astr.*, **21**, 708

Press W.H., Teukolsky S.A., Vetterling W.T., and Flannery B.P., 1992,
    Numerical Recipes. Second Edition. CUP

Rees M.J., 1966, *Nature*, **211**, 468

Rees M.J., 1967, *MNRAS*, **136**, 279

Rhoades C.E. Jr., Ruffini R., 1974, *Phys. Rev. Lett.*, **32**, 324

Romero G.E., Kaufman Bernadó M.M., and Mirabel I.F., 2002, *A&A*, **393**, L61

Rybicki G.B., and Lightman A.P., 1979, *Radiative Processes in Astrophysics*,
    Wiley-Interscience

Spencer R.E., 1979, *Nature*, **282**, 483

Svensson R., 1982, *ApJ*, **258**, 335

Synge J.L., 1957, *The Relativistic Gas*, Amsterdam: North-Holland

van der Laan H., 1966, *Nature*, **211**, 1131

Zdziarski A.A., Coppi P.S., and Lamb D.Q., 1990, *ApJ*, **357**, 149

Zdziarski A.A., Poutanen J., Paciesas W.S., and Wen L., 2002, *ApJ*, **578**, 357