

University of Southampton Research Repository
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING

Physics and Astronomy

Modelling Accretion Disk Winds in Quasars

by

Nicholas Higginbottom

Thesis for the degree of Doctor of Philosophy

May 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING

Physics and Astronomy

Doctor of Philosophy

MODELLING ACCRETION DISK WINDS IN QUASARS

by **Nicholas Higginbottom**

Outflows are ubiquitous in active galactic nuclei (AGN). They can take the form of either dramatic radio jets, which extend vast distances into the inter-galactic medium, or of much smaller scale winds - whose existence can be inferred only indirectly via their influence on the observed spectra.

There is good evidence to show that winds are likely to arise from the accretion disks thought to form the central engine of all AGN; they should therefore be observable in all such systems. The clearest observational signature comes in the form of the broad, blue-shifted absorption features seen in the spectra of broad absorption line quasars (BALQSOs) and the aim of this work is to investigate how the geometry and physical parameters of disk winds affect their absorption spectra.

We first discuss the changes made to an existing Monte Carlo radiative transfer code, PYTHON, in order to extend its capabilities to include modelling of AGN. These changes include the implementation of an approximate ionization scheme which takes account of arbitrary illuminating spectral energy distributions (SEDs), and the inclusion of heating and cooling effects likely to be important in the presence of high energy photons.

Next, we describe the second stage of the project which was to gain insight into the general properties of a wind exhibiting broad absorption line (BAL) features. We did this by performing radiative transfer and ionization calculations on a simple kinematic wind representation using PYTHON. We show that BAL features can be produced for plausible AGN parameters.

Finally, we present calculations carried out on a more complex wind geometry, generated from a hydrodynamic simulation of a line driven disk wind. This calculation does not produce BAL features at all, and in fact the wind is too highly ionised to permit efficient line driving. This result is significant because it illustrates the importance of detailed radiative transfer in hydrodynamic modelling, and suggests future work.

Contents

Declaration of Authorship	xv
Acknowledgements	xvii
Abbreviations	xix
1 Introduction	1
1.1 Active Galactic Nuclei	1
1.2 AGN Taxonomy	2
1.2.1 Seyferts	2
1.2.2 Radio Galaxies	3
1.2.3 Quasars	4
1.2.4 Broad Absorption Line Quasars (BALQSOs)	5
1.2.5 Blazars	6
1.2.6 Other related exotica	7
1.3 Unification	7
1.3.1 The physical properties of the broad line region	10
1.3.1.1 Density and temperature	10
1.3.1.2 Ionization mechanism	10
1.3.1.3 Size	11
1.3.1.4 Kinematics	11
1.3.1.5 Geometry	12
1.3.2 Cloud or photoionization models	12
1.3.3 Disk wind models	15
1.4 This thesis	17
2 Accretion	19
2.1 Accretion power	19
2.2 Accretion disks	23
2.2.1 Formation of accretion disks	23
2.2.2 α -disks	24
2.2.3 Departures from a pure accretion disk spectrum	30
2.3 Outflows	31
2.3.1 Jets	31
2.3.2 Disk winds	32
2.3.3 Feedback	34
3 Accretion disk winds in AGN	37

3.1	Driving mechanisms	37
3.1.1	Pressure driven or thermal disk winds	38
3.1.2	Radiation-pressure or line-driven winds	39
3.1.3	Magnetocentrifugal and magnetically confined winds.	42
3.2	Disk wind models	43
3.2.1	Radiation-driven winds	43
3.2.1.1	Murray and Chiang - a line-driven disk wind model	43
3.2.1.2	Proga, Stone & Kallman - a hydrodynamical, line-driven disk wind model	44
3.2.1.3	Risaliti and Elvis - a non-hydrodynamical line-driven disk wind model	47
3.2.2	Magneto-centrifugal winds	48
3.2.2.1	Emmering, Blandford and Shlosman - a model of a magneto-centrifugally accelerated disk wind	49
3.2.2.2	de Kool and Begelman - a magnetically confined, radiation-driven disk wind	50
3.2.2.3	Everett - a hybrid line-driven / magneto-centrifugal disk wind	51
4	Monte Carlo Radiative Transfer	53
4.1	Radiative transfer - the need for Monte Carlo	53
4.2	Overview of the method	55
4.3	Radiating sources	57
4.3.1	Photon packet production	57
4.3.2	Photon source 1: the central source	59
4.3.3	Photon source 2: the accretion disk	60
4.3.4	Photon source 3: the wind	60
4.3.4.1	Free-free emission	61
4.3.4.2	Free-bound emission	61
4.3.4.3	Line emission	62
4.3.4.4	Wind photon weights and directions	65
4.4	Radiative transfer in the wind	66
4.4.1	Tracking the path of photon packets	67
4.4.2	Scattering processes	68
4.4.2.1	Electron Scattering	68
4.4.2.2	Line opacity and the Sobolev approximation	69
4.4.3	Continuum Opacity	70
4.4.3.1	Bound-free opacity	71
4.4.3.2	Free-free opacity	71
4.4.4	Attenuation and scattering	71
4.4.5	Radiation estimators	72
4.5	Heating and cooling balance	73
4.5.1	Heating processes	74
4.5.2	Cooling processes	75
4.6	Ionization equilibrium calculation (dilute blackbody approximation)	76
4.7	Synthetic spectra	77
5	Radiative transfer and ionization calculations in AGN	79

5.1	Introduction	79
5.2	X-ray source	79
5.3	Ionization balance	81
5.3.1	Derivation of the approximate ionization scheme	81
5.3.2	Correction factor for recombination to excited levels	85
5.3.3	Approximate ionization calculations for an arbitrary SED	87
5.3.3.1	Modelling the mean intensity	88
5.3.3.2	The variable temperature scheme	91
5.3.4	Additional ionization and recombination processes	92
5.3.4.1	Dielectronic recombination	93
5.3.4.2	Auger Ionization	93
5.3.4.3	Charge Exchange	94
5.3.4.4	Three-body recombination	95
5.3.4.5	Supra-thermal electron ionization	95
5.3.4.6	Unresolved transition array (UTA) ionization	95
5.3.5	Validation of ionization scheme	96
5.4	Heating and cooling mechanisms	103
5.4.1	Compton heating and cooling	103
5.4.1.1	Compton heating	108
5.4.1.2	Induced Compton heating	110
5.4.1.3	Compton cooling	110
5.4.2	Dielectronic recombination cooling	112
5.4.3	Validation of new heating and cooling mechanisms	112
5.5	Treatment of excited levels	115
5.6	Atomic data	116
6	A Benchmark Disk Wind Model For BALQSOs	119
6.1	Introduction	119
6.2	The Kinematic Disk Wind Model	120
6.3	A Benchmark Disk Wind Model For BALQSOs	123
6.3.1	Basic Requirements: BALnicity and Ionization State	123
6.3.2	Black Hole Mass, Accretion Rate and Luminosity	124
6.3.3	Outflow Geometry	124
6.3.4	Wind-Launching Region	126
6.3.5	Wind Mass-Loss Rate	127
6.3.6	Velocity Law Parameters	128
6.3.7	X-ray Spectrum and Luminosity	129
6.3.8	Size and Resolution of the Numerical Grid	129
6.4	Results	130
6.4.1	The Physical State of the Outflow	130
6.4.2	Observing the Model: Synthetic Spectra	132
6.4.2.1	Sightlines looking into the wind cone	132
6.4.2.2	Sightlines looking above the wind cone	136
6.4.2.3	Sightlines looking through the base of the wind cone	137
6.5	Discussion	138
6.5.1	The Sensitivity of BAL Features to X-ray Luminosity	138
6.5.2	The Weakness of Line Emission Produced by the Outflow	143

6.5.3	Ionization Stratification and Reverberation Mapping	147
6.6	Implications	148
6.6.1	Outlook	149
6.7	Summary	150
7	RT simulations for a hydrodynamically generated wind model	153
7.1	Introduction	153
7.2	Method	155
7.2.1	The PK04 disk wind geometry	155
7.2.2	Illuminating spectrum	157
7.3	Results	160
7.3.1	Ionization state	160
7.3.2	Synthetic spectra	162
7.3.3	Temperature distribution	165
7.3.4	Ionization parameter	167
7.4	Discussion	167
7.4.1	Why is the wind hotter and more ionized than predicted by PK04?	168
7.4.2	The propagation of ionizing photons through the outflow: implications for line-driving	172
7.4.3	Implications for hydrodynamic simulations of line-driven winds	176
7.5	Conclusions	177
8	Summary and suggestions for future work	179
8.1	Summary	179
8.2	Future work	181
8.2.1	Extending the capabilities of PYTHON	182
8.2.1.1	Lines and levels	182
8.2.1.2	Improved disk models	183
8.2.1.3	Other photoionization and recombination processes	183
8.2.2	Kinematic models	184
8.2.3	Hydrodynamic models	185
References		187

List of Figures

1.1	The CIV line profile from the prototypical BALQSO - PHL 5200 reproduced from Junkkarinen et al. (1983).	5
1.2	A proposed QSO geometry (Urry and Padovani 1995).	9
1.3	A disk-wind based QSO geometry (Elvis 2000).	16
2.1	Collapse onto an accretion disk.	24
2.2	The run of temperature with radius (in units of gravitational radii) for accretion disks of different scales.	27
2.3	Accretion disk spectra for four different accretion disks.	28
2.4	The cumulative total and ionizing ($h\nu > 13.6\text{eV}$) luminosity of a quasar scale accretion disk as a function of radius, normalised to total luminosity.	29
2.5	Cartoon illustrating the Formation of a P Cygni line profile in an outflow.	33
3.1	The variation of the line-driving force multiplier $\mathcal{M}(t)$ with the ionization parameter ξ for a range of optical depth parameters t - reproduced from Stevens and Kallman (1990).	41
3.2	Schematic of line driven disk wind proposed by Murray et al. (1995).	44
3.3	Wind parameters for a hydrodynamic disk wind model reproduced from Proga et al. (2000).	46
3.4	Wind parameters for hydrodynamic disk wind model reproduced from Proga and Kallman (2004).	46
3.5	Wind streamlines for a line-driven disk wind as calculated by Risaliti and Elvis (2010).	47
3.6	Cartoon of proposed geometry (left panel) and simulated flow (right panel) from Nomura et al. (2013).	48
3.7	Schematic of magnetically driven disk wind proposed by Emmering et al. (1992).	49
3.8	Illustration of a clumpy, magnetically confined wind reproduced from de Kool and Begelman (1995).	50
3.9	Schematic of the hybrid line-driven /magneto-centrifugal disk wind proposed by Everett (2005).	51
3.10	Schematic of stratified disk wind proposed by Gallagher and Everett (2007).	52
4.1	An overview of the execution cycle of a PYTHON simulation.	56
4.2	A Planck function with $T_{eff} = 10^5\text{K}$ and associated CDF.	58
4.3	An overview of the radiation transfer scheme.	68
5.1	Collection of type 1 quasar spectral energy distributions (Richards et al. 2006).	80
5.2	The ground state recombination correction factors for Carbon.	87

5.3	Three examples of modelled cell spectra.	90
5.4	The ground state recombination correction factors for carbon, with dielectronic recombination included as an addition recombination process.	94
5.5	The flux incident on the inner edge of the thin shell model used to compare the ionization state from PYTHON and CLOUDY. This is for ionization parameter $U = 1$	97
5.6	Relative abundances of H, He, C, N and O ionization stages for PYTHON (solid lines) and CLOUDY (symbols) with a broken power law spectrum.	98
5.7	Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols).	99
5.8	Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections.	99
5.9	Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange disabled.	100
5.10	Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange and Auger ionization disabled.	101
5.11	Relative abundances of H, He, C, N and O ionization stages for PYTHON (solid lines) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY.	102
5.12	Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange and Auger ionization disabled. Dielectronic recombination is turned off in PYTHON.	103
5.13	Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe I-IX).	104
5.14	Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe X-XVIII).	105
5.15	Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe IXX-XXVII).	106
5.16	Fractional energy transfer per Compton scattering.	108
5.17	Temperature and heating/cooling rate for CLOUDY and PYTHON.	113
5.18	Heating rates for CLOUDY (symbols) and PYTHON (solid lines).	114
5.19	Cooling rates for CLOUDY (symbols) and PYTHON (solid lines).	114
6.1	A sketch illustrating the main features of our kinematic disk wind model.	120
6.2	Dependence of mass loss rate per unit area on the disk on radius for three values of λ	122
6.3	Cartoon illustrating the three broad classes of sightline discussed in the text.	125
6.4	Several input and simulated parameters of the benchmark disk-wind mode	133

6.5	Mean specific intensity in three cells	134
6.6	Simulated spectra for four sightlines	135
6.7	Electron scattering optical depth through the wind towards the origin.	138
6.8	C iv line detail for sightlines between 70° and 84°	139
6.9	Emergent α_{OX} values for benchmark model	141
6.10	Spectra at 75° and 80° sightlines for the benchmark model with a range of 2-10 keV luminosities	142
6.11	Input spectrum from the disk component in the benchmark model, together with the contribution from the central source for two different values of L_X	144
6.12	Spectrum detail around the C iv line feature for the benchmark model and others with a range of Xray luminosity and \dot{M}_{wind}	146
7.1	The density and velocity structure of the PK04 model.	154
7.2	The PK04 model plotted on logarithmic axes	157
7.3	The input spectrum used in the RT treatment of the PK04 hydro model. .	158
7.4	The mean charge of carbon in the model.	161
7.5	The mean charge of iron in the model (c.f. Sim et al. 2010b).	162
7.6	Simulated spectra for the RT treatment of the PK04 hydro model	164
7.7	Electron temperature calculated for the PK04 model.	166
7.8	Ionization parameter for the PK04 model.	168
7.9	Ionization parameter in PK04 geometry for a range of simulation parameters.	170
7.10	Illustration of the source of ionizing photons in RT treatment of PK04 geometry	173

List of Tables

1.1	A classification scheme for AGN after Urry and Padovani (1995).	8
6.1	Wind geometry parameters used in the benchmark model.	130
7.1	The key parameters adopted in the radiative transfer simulation.	160

Declaration of Authorship

I, **Nicholas Higginbottom**, declare that the thesis entitled *Modelling Accretion Disk Winds in Quasars* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as: ([Higginbottom et al. 2013, 2014](#))

Signed:.....

Date:.....

Acknowledgements

My first thank you must go to my long suffering supervisor, Christian Knigge for offering me the chance to get involved in this fascinating topic, and then for his tireless help and advice over the past few years. Thanks also to Knox Long for being so generous with his huge knowledge, and also for giving me the chance to study at the Space Telescope Science Institute for an unforgettable couple of months at the start of the project. Additionally, I'd like to thank my other two main collaborators in this work, Stuart Sim and Daniel Proga for their invaluable assistance, and to Daniel for putting me up in Las Vegas for a fantastic week. I cannot miss out my colleague James Matthews, who has graciously spent much of the first 18 months of his own PhD trying to teach me how to use version control. Don't give up on me James - I'll get there! Thanks to the students and faculty of the Physics department of the University of Southampton for collectively making the department a friendly and inspirational place to work and learn and thanks to the Science and Technology Facilities Council for allowing me to eat and fly to exciting places all round the world for the last 42 months. Finally, the biggest thank you must go to Ellie Payne, my wonderful wife who has permitted me to quit my job and become a student again. Thanks for all your support, and I promise I'll try to get a job somewhere sunny so you can quit *your* job and lie on a deckchair for a few years....

Abbreviations

AGN	Active galactic nuclei
ADAF	Advection dominated accretion flow
BAL	Broad absorption line
BALR	Broad absorption line region
BALQSO	Broad absorption line quasar
BEL	Broad emission line
BHXB	Black hole X-ray binary
BI	BALnicity index
BLR	Broad line region
BLRG	Broad line radio galaxy
CDF	Cumulative distribution function
CV	Cataclysmic variable
DR	Dielectronic recombination
FSRG	Flat spectrum radio galaxy
GPE	Gravitational potential energy
HERG	High excitation radio galaxy
ICM	Intra-cluster medium
IGM	Inter-galactic medium
IR	Infra-red
ISCO	Innermost stable circular orbit
KE	Kinetic energy
LERG	Low excitation radio galaxy
LINER	Low ionization nuclear emission-line region

LTE	Local thermodynamic equilibrium
NEL	Narrow emission line
NLR	Narrow line region
NLRG	Narrow line radio galaxy
OVV	Optically violent variable galaxy
RT	Radiative transfer
SED	Spectral energy distribution
SMBH	Super massive black hole
SSRG	Steep spectrum radio galaxy
ULIRG	Ultra-luminous infra-red galaxy
UTA	Unresolved transition array
UV	Ultraviolet
WHIM	Warm, highly ionized medium
YSO	Young stellar object

Chapter 1

Introduction

1.1 Active Galactic Nuclei

Active galactic nuclei (AGN) are extremely luminous, spatially compact sources of continuum and line radiation found at the centres of galaxies. The first examples were seen in the early 20th century as unusually bright cores in some spiral galaxies. These sources were first classified by [Seyfert \(1943\)](#) and thus became known as Seyfert Galaxies. Radio pioneers such as Karl Jansky and Grote Reber first noted the existence of extra-terrestrial radio emission ([Jansky 1933](#); [Reber 1958](#)); and radio surveys conducted in the 1950s and 60s (e.g. [Shakeshaft et al. 1955](#); [Edge et al. 1959](#); [Bennett 1962](#)) showed that at least part of this emission came from discrete, unresolved sources. Following lunar occultation work ([Hazard et al. 1963](#)), 3C 273 became the first of these radio sources to be identified with a faint, bluish star-like object. It was around this time that these objects were named quasi-stellar radio sources or quasars ([Chiu 1964](#)). Once spectral lines were identified in 3C 273, its redshift could be measured ([Schmidt 1963](#)). Assuming a cosmological origin for this redshift placed 3C 273 and other quasars at even larger distances than the Seyferts, which in turn implied that they had unprecedented luminosities. Soon after, optical surveys found many more bluish star-like objects that also had very large redshifts, but no radio emission ([Sandage 1965](#)). Originally these were known as QSOs (quasi stellar objects), but now the terms QSO and quasar are used more or less interchangeably. Here, both types are referred to as quasars.

In subsequent years, a whole “zoo” of extra-galactic objects were found, exhibiting a range of subtly different observational features. Almost all share a non-stellar, high-luminosity continuum source, and so they all came to be grouped under the AGN banner, even if the surrounding galaxy could not be detected. In this chapter, we will first discuss how AGN are classified by their observational features and then introduce models which seek to unify the different types under a single geometry.

1.2 AGN Taxonomy

The next section briefly reviews the main observational features of the different classes of AGN. Much of the information presented in this section has been taken from the excellent review of AGN taxonomy in [Peterson \(1997\)](#).

1.2.1 Seyferts

The first distinction to be made in AGN is between Seyferts and quasars. This is made on the basis of luminosity, with Seyferts being the lower luminosity counterpart to quasars. A simple criterion of $M_B > -21.5 + 5 \log h_0$ for Seyferts was originally suggested by [Schmidt and Green \(1983\)](#), and most Seyferts have a luminosity of order $10^{11}L_\odot$, concentrated in the centre of (generally spiral) galaxies. All Seyferts exhibit narrow emission lines (NELs) with widths of a few hundred km s^{-1} . These narrow lines are from ions in a wide range of ionization states. Critically, forbidden lines are also seen which implies that the electron density of whatever gas is producing the narrow lines must be below about 10^6 cm^{-3} .

An estimate of the actual density of the narrow line gas can be made by measuring the ratio of the strengths of forbidden line doublets like $[\text{O II}] \lambda\lambda 3729/3726$ and $[\text{S II}] \lambda\lambda 6716/6731$. Such lines are sensitive to the electron density, since the upper levels are populated primarily via collisional processes. At low densities, every collisional excitation will be quickly followed by radiative de-excitation, and so the ratio of the two lines in the doublet is just the ratio of the collisional excitation rates to the two upper levels. This ratio is given by the statistical weights of the upper levels. At high densities,

collisional de-excitation will dominate over radiative de-excitation, and a Boltzmann distribution between the two levels will be set up. Therefore the populations of the upper states will now be in the same ratio as the statistical weights. This gives a different ratio to the relative intensities of the doublet. The change from the low density limit to the high density limit takes place over about two orders of magnitude in density, and therefore an estimate of the density can be made depending on the exact ratio. Using this technique, the density of the narrow line region (NLR) has been shown to be in the range $10^3 - 10^6 \text{ cm}^{-3}$.

A second set of emission lines is seen in some Seyferts, but not others. This forms the basis of a subdivision into two classes. In type 2 Seyferts, the narrow lines are the only emission lines present however in the type 1 Seyferts, a set of *broad* emission lines (BELs) is also seen. Many of the same transitions are seen in both narrow and broad lines, and, as the name suggests, the widths of the broad lines are greater than those of the narrow lines, on the order of a few thousand km s^{-1} . There are no broad forbidden lines, implying a higher density for the emitting gas. Broad semi-forbidden lines *are* seen, however, which suggests that the electron density in the broad line region (BLR) is of order 10^9 cm^{-3} .

Both types of Seyferts also show a stellar component to their spectra, along with a strong, featureless continuum associated with the object at the very centre of the system. This continuum is much fainter in type 2 Seyferts. Over time, it has become clear that Seyferts do not always split perfectly into type 1 or type 2, and so a range of intermediate classes exist based upon exactly *which* emission lines are seen in the spectrum.

1.2.2 Radio Galaxies

Radio galaxies are almost always elliptical galaxies and exhibit strong radio emission, often in the form of one or two jets. The emission mechanism appears to be synchrotron radiation, and so there must be strong magnetic fields present, along with a population of hot electrons. Based on their optical spectrum and line emission, radio galaxies can themselves be split into three classes. First, narrow line radio galaxies (NLRG) have optical/ultraviolet (UV) spectra very similar to Seyfert 2 galaxies. Second, the much

rarer broad line radio galaxies (BLRG) are more analogous to Seyfert 1s. There is also a third class, which exhibit almost no line emission; these are referred to as low excitation radio galaxies (LERG). The NLRG and BLRG are both high excitation radio galaxies (HERG).

Radio galaxies can also be classified by their morphology. Faranoff-Riley type I galaxies (FRI) are brighter towards the centre of the galaxy, while Faranoff-Riley type II (FRII) galaxies show powerful jets that persist out to large radii before they impact the IGM and form radio hotspots ([Fanaroff and Riley 1974](#)). There is also a difference in the shape of the radio spectrum of the two types. The core-dominated FRI sources have a flat radio spectrum, while the jet-dominated FRII sources have a steep radio spectrum. FRI galaxies tend to be LERGs, and so appear to be generally lower luminosity systems.

1.2.3 Quasars

As the more luminous counterparts of Seyfert galaxies, quasars share many of the same observational features and are similarly categorised. Type 1 quasars have both BELs and NELs, with only narrow forbidden lines and strong continuum emission from the central source. If quasars are indeed the same type of object as Seyferts, one would expect to see a class of quasars with only narrow lines - i.e. type 2 quasars. Some NLRGs fit into this class, with their optical/UV emission showing evidence for heavy dust obscuration. Searches for radio quiet sources with type 2 line spectra have taken some time to yield results, but there are now a significant number of type 2 quasar candidates (e.g [Zakamska et al. 2003](#); [Alexandroff et al. 2013](#)). These sources also show the signatures of dust obscuration, and the reduced luminosity is the reason they have been harder to detect than their type 1 counterparts.

The strength of the radio emission from a quasar is also used as a classifier, with sources described as simply “radio loud” or “radio quiet”. The most commonly used ‘R’ definition of radio-loudness defines a quasar as radio loud if the rest-frame ratio of the radio (5 GHz) to optical (4400Å) flux density is greater than 10. Using this definition, about 15-20% of quasars are radio loud ([Kellermann et al. 1989](#)). The radio emission of radio loud quasars tends to exhibit a flat spectral index gradually steepening at shorter wavelengths.

1.2.4 Broad Absorption Line Quasars (BALQSOs)

A subset of quasars, the BALQSOs, are distinguished by having broad, deep absorption features in their spectra. These broad absorption lines (BALs) appear in moderately to highly ionized species, with considerable overlap with the species giving rise to the broad *emission* features seen in type 1 quasars and Seyferts. The absorption features are nearly always blue-shifted from the rest frequency of the emission line and often exhibit a P Cygni type profile. An example is shown in Figure 1.1. This line profile, taken from [Junkkarinen et al. \(1983\)](#), shows the C IV $\lambda 1550$ line in the prototypical BALQSO, PHL 5200. On the redshifted side of line centre, we see emission that is well modelled by a logarithmic line profile; however, on the blue shifted side, emission abruptly switches to absorption. The simplest model for this type of line shape is absorption in outflowing material (see section 2.3.2), and the width of the absorption feature in many BALQSOs suggests outflow velocities of up to 10% of the speed of light.




Figure 1.1: The CIV line profile from the prototypical BALQSO - PHL 5200 reproduced from [Junkkarinen et al. \(1983\)](#).

BALQSOs can be further subdivided into objects showing absorption only from highly ionised species (HiBALs) and those showing absorption also from less ionised species

(LoBALs). Most BALQSOs are radio quiet, although more recent studies suggest that about 20% of BALQSOs do produce radio emission ([Shankar et al. 2008](#)), albeit at a much lower luminosity than one would expect from a radio-loud quasar.

BALQSOs are most commonly classified via the so called “BALnicity index” (BI) defined by ([Weymann et al. 1991](#))

$$BI = - \int_{25000}^{3000} [1 - f(V)/0.9] CdV, \quad (1.1)$$

where $f(V)$ is the normalised flux as a function of velocity with respect to the line centre. Normalisation is carried out by calculating a continuum between the Si IV and C IV emission features. C is set to 0 initially, but once the quantity in brackets has been less than zero for 2000 km s⁻¹, it is set to 1. Whenever the quantity in brackets is greater than zero, C is reset to 0. Thus, only absorption that is continuous for at least 2000 km s⁻¹ counts towards the BI. This quantity is normally calculated for the C IV absorption feature, and a quasar is defined as a BALQSO if the BI for C IV is greater than zero.

Observationally, around 10-15% of quasars can be classified as BALQSOs (e.g [Hazard et al. 1984](#); [Trump et al. 2006](#); [Gibson et al. 2009](#)), but after correction for selection effects, this fraction increases to perhaps 20% ([Knigge et al. 2008](#), but also see [Allen et al. 2011](#)). BALs represent the most convincing evidence for large scale disk winds in AGN, and explaining these spectral features is one of the main motivations for this project.

1.2.5 Blazars

Blazars are a class of radio loud quasars that show continuum emission running from the radio all the way to γ -rays ([Sambruna et al. 1996](#)). This is a wider range than that observed in normal quasars, particularly at the high-energy end. They also show rapid variability at the highest energies, which implies variability in the whole continuum ([Ulrich et al. 1997](#)). Their optical emission exhibits a high degree of polarisation and is also variable on timescales of less than a day ([Mead et al. 1990](#)). All blazars show strong

radio emission with a fairly flat spectrum similar to that seen from FR I radio galaxies. BL Lac objects are a subclass of blazars that exhibit a featureless optical spectrum, whilst flat spectrum radio quasars (FSRQs) show strong, broad emission lines like type 1 quasars. Optically violent variable (OVV) quasars are another small subclass that exhibit particularly extreme variations.

1.2.6 Other related exotica

There are a range of other extra-galactic objects that may or may not be AGN in the conventional sense. For example, the ultra-luminous infra-red galaxies (ULIRGs) are a population of luminous galaxies that seem to emit their energy primarily in the infra-red. It is possible that these objects represent extreme type 2 Seyferts/quasars, with the optical emission even more heavily obscured by dust than a normal type 2 object.

Another class of object which shares some properties with AGN are the low-ionization nuclear emission-line regions (LINERs). These objects show rather low nuclear luminosities, and so do not fully fit the criteria for a Seyfert. Whilst they *do* show narrow emission lines (like Seyfert 2s), these lines are due to ions in relatively low ionization states. LINERs are very common, existing in perhaps 50% of all spiral galaxies, and so may represent the end state of quasar/Seyfert evolution.

Finally there is a class of radio quasars showing a steep spectral shape in the radio - like the jet-dominated FR IIs. These are labelled the steep-spectrum radio quasars (SSRQs). Table 1.1 summarises the types of AGN in a simple scheme after [Urry and Padovani \(1995\)](#).

1.3 Unification

Science seeks to find patterns in nature and, faced with such a dizzying array of different AGN, astronomers were keen from the time of the first discoveries to find some underlying model that could explain this observational diversity. The first hard evidence that all

	Type 2 Narrow Emission Lines (200-1000 km s ⁻¹)	Type 1 Broad Emission Lines (500-10000 km s ⁻¹)	Unusual
Radio Quiet	Seyfert 2 NLRG IR Quasar	Seyfert 1 Quasar (QSO)	BALQSO
Radio Loud	NLRG - FR I/FR II	Quasar BLRG SSRQ FSRQ	BL Lac Objects

Table 1.1: A classification scheme for AGN after [Urry and Padovani \(1995\)](#).

AGN could be different aspects of the same phenomenon came when [Rowan-Robinson \(1977\)](#) showed that quasars and Seyferts lie on the same luminosity function. Quasars simply lie at the rarer, higher luminosity end of the distribution. He also suggested that the difference in luminosity between Seyfert 1s and Seyfert 2s could be explained by dust extinction. He took the strong IR emission in Seyferts as evidence of dust surrounding the bright core, and suggested that Seyfert 2s simply have more dusty core than Seyfert 1s.

As we will see in the next chapter, the most likely physical mechanism to produce the very large luminosities seen in AGN is a super-massive black hole (SMBH) fed by an accretion disk. If this is true of all AGN, with the difference in luminosities simply being the mass of the SMBH and/or the rate of accretion, then the first broad brush unification has been made. If one then further postulates that radio activity is not a question of geometry, but is rather an evolutionary phase that AGN go through, then the radio properties can be set aside.

This leaves the line spectra. Whilst almost all AGN show NELs, only a subset show BELs, and fewer still show BALs. It is suggested that the narrow lines are produced in a structure called the narrow line region (NLR) ([Oke and Sargent 1968; Weedman 1977](#)), which is common to almost all classes of AGN. It appears to have a relatively large spatial extent and has even been directly imaged in nearby Seyferts (e.g. [Macchetto et al. 1994](#)).

Given the differences in the physical properties of the emitting gas, the broad lines are thought to arise in a separate structure - the “broad line region” (BLR) ([Oke and Sargent](#)

1968; Weedman 1977). The absence of these features in type 2 AGN is put down to an orientation effect with the BLR obscured from some sightlines (e.g. Lawrence and Elvis 1982; Antonucci 1993; Urry and Padovani 1995).

Figure 1.2 shows a schematic of this unification scheme based on Urry and Padovani (1995). The left-hand panel shows the location of the proposed elements of the model for a radio loud quasar, whilst the right panel shows how the object would be classified if viewed from different sightlines. In the right-hand panel, radio quiet quasars have been included by simply assuming the radio jet can be on or off.




Figure 1.2: A proposed QSO geometry (Urry and Padovani 1995).

The main elements of the model are the SMBH and associated accretion disk in the centre, both surrounded by an equatorial ring of dense, cold material - the torus. The exact nature of the torus is unknown, but its role in the unified theory is to obscure the direct view of the accretion disk from high inclination sightlines. This angle dependant obscuration provides the separation between type 1 and type 2 AGN, with the former corresponding to systems viewed at low inclination angles, so seeing the accretion disk directly. If the torus is geometrically thick enough, it can also obscure the view of the BLR, which is found within a parsec of the central region (Lawrence and Elvis 1982).

If one were to look at the object from a polar direction, the jet would be pointed straight towards the observer, and relativistic beaming would thus make the source appear to

be more luminous. Also, any fluctuations in the emission from the jet would appear to happen almost simultaneously, even if the mechanism causing the fluctuation could only propagate along the jet at the speed of light. These characteristics are typical for the blazar class of quasars.

The BLR itself is pictured in Figure 1.2 as a system of clouds. However, this is just one potential model, and the true nature of the BLR is one of the most important outstanding pieces in unified models of AGN. Models can be broadly grouped into cloud models and disk wind models although there is significant overlap between the two. However, before discussing these models, it is worth summarising the physical properties of the BLR that can be inferred directly from observations.

1.3.1 The physical properties of the broad line region

1.3.1.1 Density and temperature

As described earlier, the absence of forbidden BELs, coupled with the existence of semi-forbidden BELs suggests a characteristic hydrogen number density of order 10^9 cm^{-3} for the line-forming region. The fraction of the sky which is covered by the BLR as seen from the central source (the covering factor) can also be estimated. Very few quasars show Ly α absorption, which implies a covering factor of no more than 10% (assuming the BLR gas is optically thick in Ly α) ([MacAlpine 2003](#)). Estimating a characteristic temperature is more difficult, but the existence of Fe II lines in most AGN spectra implies a temperature of $\lesssim 35000 \text{ K}$ for at least some parts of this region. Higher temperatures would collisionally ionize this ion ([Osterbrock and Ferland 2006](#)).

1.3.1.2 Ionization mechanism

There is strong observational evidence that the BLR is photoionized. Recombination lines respond very quickly to changes in the continuum flux, which implies that the ionization state of the BLR is set by photoionization ([Osterbrock and Ferland 2006](#)). In addition, the equivalent widths of recombination lines in objects with a wide range of different luminosities are proportional to the luminosity, which implies a direct relationship

between the luminosity of the source and the ionisation state of the BLR. More evidence that the BLR is photoionized comes from the existence of fairly highly ionized species, such as O VI. If the temperature derived from the Fe II lines is representative, then only photoionization can explain the existence of such highly ionized ions. Of course, the existence of both low ionization and high ionization species in the same structure immediately implies some kind of ionization stratification.

1.3.1.3 Size

The size of the BLR can be determined by measuring the time-lag between changes in the continuum and changes in the emission lines (“reverberation mapping”). This technique works because photoionization drives the ionization state of the wind. The typical size scale of the BLR turns out to be of the order from a few light days up to a few hundred light days, depending on the mass of the SMBH (e.g. [Peterson et al. 2004](#); [Kaspi et al. 2007](#)). Reverberation mapping can also provide some insight into the ionization structure of the BLR: low ionization lines exhibit longer lags, thus appear to originate further away from the central source than lines from more highly ionized species.

1.3.1.4 Kinematics

The width of the broad emission lines cannot be thermal if the temperature of the gas is only around 10^4 K, so bulk motion of the gas must be responsible. The velocities implied by the width of the lines, coupled with the distances from reverberation mapping, suggests that the gas is moving under the gravitational influence of the SMBH. However, whether this is an orbit or an inflow / outflow is not immediately apparent from the line profiles.

BALs, first seen in quasars in the late 60s, were immediately associated with outflowing material around the SMBH - perhaps in the form of expanding shells of material, ejected in repeated outburst events ([Burbidge et al. 1966](#); [Lynds 1967](#)). The width of the absorption features is of the same order as that of the emission lines ([Lee and Turnshek 1995](#)) and very similar ionic species are seen, which suggests a simple model in which

BALs and BELs arise in a single structure. There are of course alternative explanations, the strongest alternative being that the structure producing BALs is simply present in some quasars but not others, possibly indicating an evolutionary sequence (e.g. [Surdej and Hutzemeekers 1987](#); [Fabian 1999](#); [Becker et al. 2000](#)).

1.3.1.5 Geometry

As mentioned in Section [1.2.4](#), the P Cygni profile of the BALs implies that they are produced in an expanding medium. The absorption features are very deep, but not completely black. This could mean either that part of the background source is visible “around the side” of the absorbing medium; that the optical depth is not high enough for complete extinction; or that scattering is filling in the trough. Polarisation studies have shown significant polarisation in the bottom of the absorption troughs, suggesting that scattering is responsible (e.g. [Cohen et al. 1995](#)). Furthermore, it has been shown that whatever material is responsible for the BALs has a covering factor less than 100% ([Arav et al. 1999a](#)), although from the point of view of the observer, the continuum source is completely covered. This implies a preferential viewing direction for BALQSOs, where the sightline to the continuum source intercepts the absorbing gas.

So, to summarise, the BLR appears to be a photoionized region with temperatures of order a few tens of thousands of K. It is relatively compact, and consists of material moving at virial speeds, suggesting that it is in the gravitational influence of the SMBH. It seems that the material producing emission lines is always present. However, whether it is the same material which produces BALs in some objects, or whether another component is required, is not known. In the next sections, we will introduce the two main classes of physical model for the BLR that have been suggested to date: photoionization models, and disk wind models.

1.3.2 Cloud or photoionization models

“Form follows function” is perhaps the best way to describe the development of the cloud models of the BLR. Since the BLR appears to be photoionized, it made sense from the

earliest investigations to try to tie down its structure by carrying out photoionization calculations. If one assumes that all of the BLR sees the same continuum spectral energy distribution (SED) as we observe from the Earth, the ionization state of a parcel of gas in the BLR depends on the luminosity of the continuum source, the distance of the parcel from the source, and the density of the parcel. These factors are usually combined into an “ionization parameter” - one common definition is given by

$$\xi \equiv \frac{L_x}{n(H)R_x^2}, \quad (1.2)$$

where L_x is the X-ray luminosity of the source, $n(H)$ is the number density of hydrogen atoms, and R_x is the distance of the cloud under investigation from the source. An alternative, again commonly used, definition is given by

$$U \equiv \frac{Q(H)}{4\pi r_0^2 n(H)c} = \frac{\Phi(H)}{n(H)c}, \quad (1.3)$$

where $Q(H)$ is the number of hydrogen ionizing photons emitted by the central source per second, r_0 is the distance from the source to the illuminated face of the cloud and $\Phi(H)$ is the surface flux of ionizing photons. Since U depends on the number of photons, whilst ξ depends on the integrated source luminosity, the relationship between the two depends on the source’s SED. In this thesis, we usually use U , however both are common in the literature.

Photoionization codes such as CLOUDY ([Ferland et al. 2013](#)) are used to predict line strengths using the ionization parameter as an input, and so the physical conditions in the BLR can be inferred by comparing predictions with observations. Usually, a single set of parameters cannot perfectly match observations, and so a “complete” model generally involves a large number of “clouds”, usually with the same density, but distributed in distance from the central source. A lower limit on the number of clouds can be inferred from the smoothness of the observed line profiles (e.g. [Atwood et al. 1982](#); [Arav et al. 1997](#); [Dietrich et al. 1999](#)) if the internal velocity distribution for each cloud is assumed to be thermal. The upper limit on the number of clouds is much harder to obtain, since it is degenerate in quantities like the size of the clouds and the covering factor. As the

number of clouds increases, the clouds would merge into a continuous flow (Dietrich et al. 1999).

Given the line intensity, one can estimate the number of ions in a cloud, and hence the size of the clouds (given the density), the optical depth through a cloud, and the volume filling factor of the BLR (the proportion of the volume of the BLR which is taken up by the clouds). This type of model has come to be known as the cloud or photoionization model of the BLR.

A further development of this model came from the observation that very different objects seemed to have very similar broad emission lines. This implies a preferential location for the BLR. However, it turns out that if there is a system of clouds, with a range of densities, distributed in space around the continuum source then those clouds which are in the correct ionization state for a given ion to exist would dominate the emission spectrum. This is the locally optimised cloud (LOC) model (Baldwin et al. 1995).

In such cloud models, the distribution of clouds within the volume of the BLR can be used to explain other observations. For example, as already mentioned, the lack of Ly α absorption in quasars implies a covering factor for the BLR of about 10%. However, given the energy radiated in the lines, the BLR would need to intercept *more* than 10% of the energy emitted by the central source (Maiolino et al. 2001; Gaskell et al. 2007). This apparent contradiction can be explained if the BLR clouds form a flattened geometry, so in most cases we see the central source through a hole in the BLR whilst in other directions, the BLR has an almost 100% covering fraction.

Despite the flexibility of the cloud model, there are obvious questions regarding its physical viability. For example the clouds need to be confined in some way to prevent them from simply evaporating. This confinement could be provided by some kind of intra-cloud medium¹ (Krolik et al. 1981) although magnetic confinement has also been suggested (Rees 1987).

¹A confining medium needs to move with the BLR to avoid the clouds being destroyed due to Rayleigh-Taylor instabilities (Mathews and Ferland 1987). The temperature and composition of this medium are uncertain since it must be dense enough to confine the clouds, yet it is not directly observed.

In addition, it is not clear how to explain the BALs in this geometry although one possibility that has been extensively investigated is that the flattened distribution mentioned above would have some sightlines passing through a cloud, and it is possible that one would then see absorption. This geometry ties in with the observational constraint that the material producing the BALs must only partially cover the central source in order to give the observed ratio of emission to absorption (Junkkarinen et al. 1983). Another option is to turn to an evolutionary argument; for example that BALQSOs represent an early phase of quasar development (e.g Gregg et al. 2006).

1.3.3 Disk wind models

One observational characteristic of the BLR that cannot be directly explained by the simple cloud model is the origin of the width of the BELs.

Bulk motion of the line emitting material (whether clouds or something else) is clearly required, but the model does not tell us whether this is orbital, infall or outflow. It has been suggested for a long time that some kind of acceleration mechanism could be driving the clouds outwards² (e.g Blumenthal and Mathews 1975; Shields 1977), but this still leaves the question of how the clouds could survive a fairly long acceleration phase without breaking up. This was addressed by models in which the clouds are formed as condensations in an outflowing wind driven from the surface of the accretion disk (Beltrametti 1981; Weymann et al. 1982). The clouds are confined by the surrounding lower density wind, eventually evaporating as they move outwards, to be replaced by fresh condensations closer in.

As mentioned earlier, the simplest explanation for the line shape of BALs (as opposed to BELs) is that they are formed in outflowing gas. Fairly early on, it was suggested that the AGN itself could provide the driving force for such an outflow (e.g Rees 1970), although the mechanism was not (and is still not) a settled question. Non-spherically symmetric outflows seemed to be required from photon conservation arguments, and such

²Radiation pressure from the inner accretion disk was most often invoked as the force providing this acceleration in this early work.

a geometry also provides a natural explanation for the BALQSO fraction ([Junkkarinen 1983](#))³.

Around the same time, it was shown that accretion disks in AGN could produce strong winds, driven either radiatively by scattering of UV photons ([Shlosman et al. 1985](#)), or centrifugally by magnetic fields ([Blandford and Payne 1982](#)). These “disk winds” suggest a natural unifying structure, and theoretical studies (e.g [Murray et al. 1995](#)) showed that such winds could indeed produce both BELs and BALs if the conditions were correct. One of the most influential disk wind unification models was described by [Elvis \(2000\)](#). His proposed geometry for the inner parts of a QSO is shown in Figure 1.3.




Figure 1.3: A disk-wind based QSO geometry ([Elvis 2000](#)).

In this particular model, the wind consists of two phases: a warm, highly ionized medium (WHIM) in pressure balance with denser, cooler “clouds”, which give rise to the BELs. An observer looking from a polar direction would see both the unobscured continuum source and line radiation from the BEL clouds. From a sightline which intersects the opening angle of the wind, one would see absorption from the same clouds, and the object

³Spherically symmetric models can also work in some cases ([Drew and Boksenberg 1984](#)).

would appear as a BALQSO. Narrower absorption features, identified with the narrow absorption lines (NALs) seen in some quasars, would be seen from equatorial directions, since the velocity gradient of gas along the sightline to the source of UV photons is small.

This model is essentially free of assumptions about the mechanism which accelerates the wind⁴, but many other models exist in which the structure of the wind is calculated through direct modelling of a given driving mechanism. These models are discussed in Chapter 3, where the main driving mechanisms are also described.

1.4 This thesis

The aim of this thesis is to test how well the disk wind model can explain the observational features of BALQSOs. Chapter 2 continues the introductory portion of the thesis, with a brief summary of the physics of the accretion disk. This is key to the understanding of disk winds, which are described in more detail in Chapter 3. We investigated how different wind geometries affected the absorption spectrum through numerical simulations carried out using the radiative transfer and ionization code PYTHON, first described in Long and Knigge (2002). The code as it stood at the start of this work is introduced in Chapter 4.

PYTHON was originally written with the simulation of cataclysmic variable (CV) systems in mind, and significant development effort was needed to allow it to properly model winds around AGN. The new physics included in the code, and a few of the validation exercises carried out, are discussed in Chapter 5.

With the upgraded version of PYTHON in place, the next step in the project was to take a smooth, kinematic representation of a disk wind and investigate whether it could produce BAL features. Chapter 6 describes a “benchmark” model that manages to produce BALs for a physically reasonable set of AGN disk wind parameters. We then present the simulated spectra obtained from this geometry, and discuss what can be learnt from the successes and failures of benchmark model.

⁴although radiation pressure is invoked to bend the wind outwards, forming the funnel like shape

In Chapter 7, a radiative transfer and ionization calculation for a more complex line driven disk wind model is presented. First described by [Proga and Kallman \(2004\)](#), this hydrodynamic wind model represents a “realistic” geometry, which we hoped could generate more complex spectral features than the benchmark model. In fact we find that the original, necessarily simplified, radiative transfer treatment led to an underestimation of the ionization state of the wind. Since the driving mechanism for the outflow is critically dependant on the ionization state, this means that the calculated outflow geometry itself may not be completely self consistent. This is an important result given the influential nature of the model.

Finally, conclusions and suggestions for future work are presented in Chapter 8.

Chapter 2

Accretion

2.1 Accretion power

Once the true luminosity of AGN had been established, the question of what mechanism could produce such huge energy outputs quickly followed. For example, the prototypical quasar 3C 273 has a bolometric luminosity of around 6×10^{46} ergs s $^{-1}$ (Türler et al. 1999). This is around 1.6×10^{13} L $_{\odot}$, or more than a hundred times the total luminosity of all the stars in the Milky Way. The maximum amount of energy that can be extracted from matter is given by $E = mc^2$. We can thus define the efficiency, η , of an energy conversion process as the fraction of this theoretical maximum conversion rate that it can sustain. If we define \dot{M} as the rate at which matter is supplied, then the luminosity of a source powered by a process with efficiency η is

$$L = \eta \dot{M} c^2. \quad (2.1)$$

Fusion, the power source of stars, is one of the most efficient processes for converting mass into energy, with a typical efficiency (for H \rightarrow He) of 0.7% (Frank et al. 2002). If fusion were the power source of 3C 273, this would require the conversion of H to He at a rate of $\dot{M} \sim 150$ M $_{\odot}$ yr $^{-1}$. At first sight, this may not seem unreasonable. However, when one includes information about the *size* of the continuum source in quasars, things become more challenging.

The bright continuum source at the heart of AGN has never been resolved even in nearby Seyferts. Furthermore, many AGN (including 3C 273) are observed to be variable on timescales of days, some even shorter. These observational clues suggest that the continuum source is at most a few light days in size. It is hard to conceive of a stable configuration of a sufficient number of stars to produce 100 times the luminosity of a typical galaxy in a volume about the size of the solar system. There is, however, a more efficient way of extracting energy from matter - the conversion of gravitational potential energy (GPE) into radiation as that matter is accreted onto a gravitating body.

The energy, ΔE , liberated via accretion of a parcel of material of mass ΔM falling from infinity onto the surface of a body of mass M and radius R is given by

$$\Delta E = \frac{GM\Delta M}{R}. \quad (2.2)$$

Taking the time derivative of both sides of Equation 2.2 allows us to write an expression for the luminosity, L , of an accretion powered source

$$L = \frac{GM\dot{M}}{2R}, \quad (2.3)$$

where \dot{M} is the accretion rate. The factor of 2 is due to the Virial theorem, which states that a maximum of 50% of the initial GPE is available for conversion into radiation. The efficiency of the accretion process is therefore

$$\eta = \frac{L}{\dot{M}c^2} \quad (2.4)$$

$$= \frac{M}{2R} \frac{G}{c^2}. \quad (2.5)$$

$$= \frac{M}{2R} \frac{G}{c^2}. \quad (2.6)$$

Thus the accretion efficiency is proportional to the compactness, (M/R) , of the accreting object. This parameter is maximised for a black hole¹, and whilst there is no hard surface, we can usefully define its characteristic size as the radius of the “innermost

¹Actually, neutron stars can be somewhat more efficient, since more energy can be released as the accreting material impacts the surface of the neutron star.

stable circular orbit”, R_{ISCO} . Inside this radius, material plunges directly over the event horizon, carrying any remaining energy with it.

For a non-rotating (Schwarzschild) black hole, $R_{ISCO} = 6GM/c^2$, which gives a maximum efficiency of about 8% (more detailed calculations yield the more commonly used value of 6%). The ISCO reaches a minimum value of $R_{ISCO} = GM/c^2$ for prograde orbits around a maximally spinning (Kerr) black hole. In this case, a relativistic treatment is necessary and yields a maximum efficiency of about 32% (Dutau and Biermann 2005). This is more than an order of magnitude more efficient than fusion and reduces the required accretion rate for 3C 273 to around $10 M_\odot \text{ yr}^{-1}$. In addition, as we will see in section 2.2.2 below, the size of the region that would be expected to emit the bulk of the luminosity is very small, consistent with observations.

At first sight, it may appear that there is no limit to the luminosity of an accretion-powered source. Furthermore, since the efficiency appears to be independent of the mass of the accreting object, it seems that there is no requirement to have a particularly massive black hole to power even a quasar (although clearly any object which accretes $10 M_\odot \text{ yr}^{-1}$ over significant timescales will *become* very massive). There is, however, another consideration. In order to accrete material onto a luminous object, the inward force of gravity must overcome the outward radiation pressure on the in-falling material.

Let us assume that this material is ionized hydrogen, but that the free protons and electrons still strongly interact via the Coulomb force. The gravitational force is dominated by the force on the protons, and the inward force on a single proton is given by

$$F_{grav} = \frac{GMm_p}{r^2}. \quad (2.7)$$

Assuming the central object radiates isotropically with luminosity L , the radiation flux S at a distance r from the source is

$$S = \frac{L}{4\pi r^2}. \quad (2.8)$$

In a fully ionised plasma, the radiation will interact with electrons via Thomson scattering. Since the Thomson cross section (σ_T) is over a million times larger than the

comparable proton cross section, we can neglect the force exerted by the radiation on the protons. The force acting on a single electron is given by the rate of change of momentum of the photons interacting with it, which is just the energy transferred from the photons to the electron, divided by the speed of light, c . The energy transferred is equal to the flux multiplied by the cross section and so the outward force acting on a single electron is given by

$$F_{rad} = \frac{L\sigma_T}{4\pi r^2 c}. \quad (2.9)$$

In a neutral plasma, the number of protons and electrons is the same. The so-called “Eddington luminosity”, L_{edd} , is thus defined as the luminosity at which those inward and outward forces balance, so that accretion would be halted. Equating 2.7 and 2.9 and rearranging, we obtain

$$L_{edd} = \frac{4\pi GMm_p c}{\sigma_T} \simeq 1.3 \times 10^{38} \frac{M}{M_\odot} \text{ ergs s}^{-1}. \quad (2.10)$$

Thus the maximum luminosity of any object is a function of its mass, and so in order to achieve the observed luminosity of 3C 273, we require a central mass of at least $4 \times 10^8 M_\odot$. This relation is often used to place lower limits on the mass of the SMBH in AGN. Equation 2.10 can be rewritten in terms of the mass accretion rate, assuming $L = \eta m c^2$ to give an Eddington accretion rate

$$\dot{M}_{edd} = \frac{4\pi GMm_p}{\eta c\sigma_T} \simeq 3.7 \times 10^{-8} \frac{M}{M_\odot} \text{ M}_\odot \text{ yr}^{-1}, \quad (2.11)$$

for $\eta = 0.06$. Thus a $10^8 M_\odot$ SMBH has an $L_{edd} \simeq 1.3 \times 10^{46}$ ergs s $^{-1}$ and $\dot{M}_{edd} \simeq 3.7 M_\odot \text{ yr}^{-1}$.

In the above derivation, the Eddington luminosity relates to a spherically accreting and isotropically radiating body. One can therefore envision geometries where “super Eddington” accretion can take place. This could happen, for example, if the material is accreted only from certain directions, whilst the radiation is still released isotropically or vice versa. Thus the accretion flow would not see all of the radiation. The Eddington luminosity is, however, still used extensively when discussing the luminosity of sources. This is partly because it is a convenient benchmark, and partly because it is unlikely

that any source can exceed it by more than about an order of magnitude whatever its geometry.

Another interesting aspect of this treatment is to consider what happens if the in-falling material is not fully ionized. In this case, there may be many spectral lines present, due to bound-bound transitions which collectively have a cross section far larger than σ_T . Thus sub-Eddington sources can still expel strong outflows by radiation pressure alone. This is an important mechanism for the driving of disk winds, and is discussed in more detail in section 3.1.2.

2.2 Accretion disks

The preceding section demonstrated that the most likely power source for an AGN is accretion, the release of GPE from material falling towards a massive central object. However, this is not the complete answer, since the GPE will initially be converted into kinetic energy (KE) and material falling into a black hole will simply carry all of its KE with it over the event horizon. There must therefore be some mechanism by which the energy is converted into radiation. It turns out that the kinematics of collapse almost inevitably produces a structure which provides this mechanism.

2.2.1 Formation of accretion disks

Whenever material collapses under gravity, the result is likely to be a disk-like structure. This is a natural consequence of the fact that, in general, any collection of particles will always have some net angular momentum. If we take an initially spherical cloud of particles, they will begin to move along radial paths towards the centre of mass. However, since angular momentum must be conserved, as the particles get closer to the centre, their rotational velocity will increase. Particles will eventually reach the point when the radial gravitational attraction is balanced by the centrifugal force due to their rotational velocity, and so their paths will become vertical rather than radial. They will thus fall towards the plane defined by the angular momentum vector. This is illustrated in figure 2.1.




Figure 2.1: Collapse onto an accretion disk.

As particles pass through this plane, the net force will switch direction, and so they will start to oscillate about the plane, dissipating energy via interactions with other particles until the vertical component of motion is damped out, and a disk is formed. This process of energy dissipation will continue in the disk, and the orbits of particles in the disk will gradually circularise. This is because for any given value of angular momentum, a circular orbit represents the lowest energy state. We will therefore end up with a rotationally supported system of particles in a plane - the accretion disk.

There are now two important pieces of physics that must be understood. Firstly, if we are to tap the kinetic energy of the orbiting material, there must be some mechanism by which that energy is radiated away as the material moves deeper into the gravitational well. Secondly, angular momentum must be lost by material in order for it to move inwards through the structure of the disk in order to convert further GPE into KE and hence into radiation.

2.2.2 α -disks

An accretion disk is a differentially rotating system. That is, if one takes material at a given radius, neighbouring material slightly closer to the central object will orbit more quickly, while material further away will orbit more slowly. The interaction of systems of particles moving at different velocities is dealt with by the theory of viscosity, where the dynamic or shear viscosity expresses the resistance of a material to shearing forces.

The viscous force per unit area (τ) of two slabs of material separated by a distance ds , moving with a relative shearing velocity gradient dv/ds is given by

$$\tau = \rho \nu \frac{dv}{ds}, \quad (2.12)$$

where ν is defined as the kinematic viscosity of the material (measured in stokes, or $\text{cm}^2 \text{ s}^{-1}$), and ρ is the density.

In an accretion disk, this shearing force will act to transport angular momentum outwards, allowing material to slow down and move inwards. Of course this will mean that some material ends up with increased angular momentum which will have to leave the system via some other mechanism, perhaps magnetically confined outflows ([Blandford and Payne 1982](#)). In binary systems, tides can also transfer angular momentum back to the orbit. If the angular momentum cannot be shed, then the disk will continuously increase in size.

The exact source of the viscosity in an accretion disk is not fully understood, although it is clear that magnetic fields are probably important ([Balbus and Hawley 1991](#)). It turns out that one can produce a reasonable model of an accretion disk by assuming the presence of some kind of viscosity and parameterising it. The most successful model of accretion disks is the so called α -disk model, first described by [Shakura and Sunyaev \(1973\)](#). Here, α is a scaling parameter used to describe the strength of viscosity ν in terms of other disk parameters. In this model, hydrodynamic turbulence is assumed to be the source of the viscosity, i.e. the chaotic motion of small parcels of gas is assumed to transport angular momentum between neighbouring rings of the disk. In this case, dimensional analysis suggests that the kinematic viscosity can be written as

$$\nu \sim \lambda \tilde{v}, \quad (2.13)$$

where λ is the spatial scale of turbulent motion in the gas, and \tilde{v} the typical velocity of the turbulent eddies. It is reasonable to assume that the size of the largest turbulent eddies in the disk are of order the local scale height of the disk, H . Furthermore, the velocity of the eddies cannot be faster than the local sound speed, c_s . If it were, shocks would

form, and the eddies would dissipate via shock heating and radiation. The kinematic viscosity can therefore be written as

$$\nu = \alpha c_s H, \quad (2.14)$$

where α is a dimensionless parameter that combines the uncertainty in the actual scale size of eddies in the disk and in the velocity at which they travel. Based on the arguments above, it is expected that $\alpha \lesssim 1$. If one further assumes that the only source of heating for the disk is viscous dissipation, and that any radiation generated leaves the disk locally (rather than heating neighbouring parts of the disk), then the temperature of the disk can be calculated as a function of radius, yielding (Wade 1984)

$$T(r) = \left[\frac{3GM\dot{M}}{8\pi\sigma r^3} \left\{ 1 - \left(\frac{R_{in}}{r} \right)^{1/2} \right\} \right]^{1/4}. \quad (2.15)$$

Here, R_{in} is the innermost edge of the accretion disk where viscous forces are assumed to drop to zero. The shape of this relationship is shown in Figure 2.2 for four very different systems, all of which are likely to contain accretion disks. These systems are intended to be generic illustrations rather than to simulate any given astronomical object. For ease of comparison, the same ‘‘Eddington ratio’’, $\dot{M}/\dot{M}_{edd} = 0.25$, has been adopted for all four systems. They are

- A quasar: $M = 10^9 M_\odot$ and $\dot{M} = 10 M_\odot \text{ yr}^{-1}$
- A Seyfert galaxy: $M = 10^7 M_\odot$ and $\dot{M} = 0.1 M_\odot \text{ yr}^{-1}$
- An X-ray binary (XRB): $M = 10 M_\odot$ and $\dot{M} = 10^{-7} M_\odot \text{ yr}^{-1}$
- A cataclysmic variable (CV): $M = 1 M_\odot$ and $\dot{M} = 10^{-8} M_\odot \text{ yr}^{-1}$

For the quasar, Seyfert and XRB, the innermost disk radius, R_{in} is set to 6 gravitational radii, R_G , where $R_G = GM/c^2$. The outermost disk radius in these systems is set to $R_{out} = 10^4 R_{in}$. For the CV, $R_{in} = 10^9 \text{ cm}$ and $R_{out} = 100 R_{in}$.

Several interesting features are visible in Figure 2.2. Firstly we see that at large radii, the temperature scales as $T \propto r^{-3/4}$. Secondly, and much more importantly, we see that




Figure 2.2: The run of temperature with radius (in units of gravitational radii) for accretion disks of different scales.

the peak temperature of the disk around a black hole goes down with increasing black hole mass. In fact, for an accretion rate set to a fixed fraction of the black hole mass, $T_{max} \propto M^{-1/4}$. Thus, whilst an accretion disk surrounding a $10 M_\odot$ black hole would be a strong source of X-rays, this is not the case for a $10^9 M_\odot$ black hole in a quasar. Note, however, that there can be other sources of X-ray radiation in a quasar.

Finally, looking at the temperature distribution in the accretion disk of a CV, we see that the hard surface of the white dwarf at the centre of the disk means that, despite the scaling law above, the peak temperature in a CV disk is about the same as that in a quasar disk. This somewhat counter-intuitive result provided part of the motivation for this project, because some of the observational features of CVs have been successfully explained in terms of disk winds (e.g. [Drew and Verbunt 1985](#); [Long and Knigge 2002](#); [Noebauer et al. 2010](#)) and, given the apparent similarity between the disks, it is reasonable to think that these models and techniques might also help to explain the similar features seen in some AGN.

If one makes the assumption that an accretion disk will radiate like an ensemble of

blackbodies, with each annulus radiating at the local temperature given by equation 2.15, then one can produce a simulated spectrum by simply integrating over all radii. Figure 2.3 shows this for the four systems previously described.




Figure 2.3: Accretion disk spectra for four different accretion disks.

We immediately see that the spectra all take the form of stretched blackbodies. At the high frequency end, the spectrum drops off exponentially - this is the Wien tail of the hottest part of the disk. At low frequencies, $L_\nu \propto \nu^2$ - this is the sum of the Rayleigh-Jeans tails of all of the black bodies. In between these two limits is the “characteristic disk spectrum”, $L_\nu \propto \nu^{1/3}$. Figure 2.3 clearly shows the X-ray emission expected from the XRB, and also the subtle difference between the two types of AGN. The location of the high and low frequency break points in the disk spectrum depends on the values of R_{in} and R_{out} respectively, and the range of frequencies over which the characteristic disk spectrum applies depends on the ratio R_{out}/R_{in} . For the quasar, Seyfert and XRB, we set $R_{out}/R_{in} = 10^4$, so the characteristic disk spectrum applies over about three orders of magnitude (since $T \propto r^{-3/4}$). In contrast, because the CV disk exists over a much smaller range of radii, it hardly shows the characteristic spectrum at all.

Finally, we can calculate the luminosity of the accretion disk as a function of radius. Figure 2.4 shows the cumulative luminosity as a function of radius for the quasar disk described above. The luminosity is scaled to the total, and the figure shows clearly that 90% of the luminosity is radiated within $150 R_G$ of the centre. This is approximately 2×10^{16} cm or about 8 light days. This is the right size to explain the timescale of variability of quasars, as described earlier. The second line shows the cumulative ionizing luminosity (the integrated flux emitted at frequencies such that $h\nu > 13.6\text{eV}$) for the same disk, also normalised to the total luminosity. We see that the radial distribution of ionizing flux is even more concentrated towards the centre, with essentially all ionizing flux emitted inside about $50R_G$.




Figure 2.4: The cumulative total and ionizing ($h\nu > 13.6\text{eV}$) luminosity of a quasar scale accretion disk as a function of radius, normalised to total luminosity.

This treatment of the accretion disk only holds if the disk is radiatively efficient. This means that the disk is in local thermal equilibrium, with energy dissipated locally. This is expected to be true if the accretion rate is between about 1% and 100% of the Eddington rate. There are therefore two cases that require further consideration. Firstly, if the accretion rate is very low, the density of the disk will be low, and the disk will be optically thin. It will therefore not radiate efficiently, and the thermal energy of the gas

will be transported inwards with the gas (advection). In the case of a black hole, this energy will be lost when the material passes over the event horizon. This is known as an advection dominated accretion flow - ADAF ([Ichimaru 1977](#); [Narayan and Yi 1994](#)).

If the accretion rate is above the Eddington rate, the solution is not as well known. For moderately super Eddington accretion rates, modifications to the α -disk model yield a similar “slim disk” solution ([Abramowicz et al. 1988](#)). However, for even higher accretion rates, the disk becomes unable to radiate away all of the locally generated heat, and it will begin to “puff up” to form a thick disk, known as a “Polish doughnut”. This structure can sustain luminosities up to 100 times the Eddington limit ([Abramowicz et al. 1980](#)). Both types of super-Eddington disk are radiatively inefficient, with significant energy being advected into the black hole.

2.2.3 Departures from a pure accretion disk spectrum

In many accreting objects, radiation is observed at energies far in excess of those that would be expected from a system of blackbodies, extending to X-rays and even γ -rays. This radiation is best modelled by a power-law SED, although the process by which these photons are produced is not fully understood. A popular model suggests that the X-ray emission comes from some kind of disk “corona”, i.e. a hot atmosphere which increases the energy of photons originating from the disk below via a process of Compton up-scattering.

Many AGN also exhibit a fairly featureless, soft X-ray continuum, which is well described by a blackbody with radiation temperature 0.1-0.2 keV ([Gierliński and Done 2006](#)). This temperature is higher than one would expect from the accretion disk in such systems. This so-called “soft excess” lies above the extrapolation of the power-law X-ray source seen in harder X-rays, and its origin is also a topic of active research. The similarity of the radiation temperature of the soft excess radiation across a range of objects of different masses suggests that it is not Comptonised emission from the accretion disk ([Gierliński and Done 2004](#)). Alternative explanations include radiation from an optically *thin* inner part of the accretion disk ([Done et al. 2012](#)), and reflection of X-ray photons from the photoionized surface of the accretion disk ([Ross and Fabian 2005](#); [Crummy et al. 2006](#)).

2.3 Outflows

Perhaps surprisingly, wherever one sees or infers an accretion disk, one also sees strong evidence of *outflowing* material. This can take the form of the dramatic radio jets seen in radio galaxies, AGN, CVs and black hole X-ray binaries (BHXB) or the disk winds thought to exist in the same range of objects. The luminosity associated with these outflows is significant, and no discussion of accretion would be complete without mentioning some of their key features.

2.3.1 Jets

Radio jets have been seen in a wide range of astrophysical settings and seem to be intimately linked to accretion. This relationship has been seen most clearly in the BHXB case, where they appear to be produced as the accretion process switches from a high rate of accretion to a lower rate. When the accretion rate is high (the so called high/soft state), emission is dominated by an accretion disk extending down to very small radii, whereas at the lower accretion rate (the low/hard state), the disk seems to be truncated, and a jet is launched ([Fender et al. 2004](#)). In the low/hard state, accretion appears to proceed via the relatively inefficient ADAF mode.

In AGN, strong radio emission is seen in about 5-10% of sources ([Peterson 1997](#)), which suggests that there may be a similar evolutionary track for accretion in these sources ([Jester 2005](#)). However, since the mass of the central black hole is at least 5 orders of magnitude higher than in a BHXB, the state transition would take place over thousands of years (state changes in BHXBs take place over timescales of order a year), and so, observationally, AGN appear to be in a fixed state.

It is clear that there is a strong magnetic field associated with radio jets, since the emission from the extended jets is predominantly synchrotron radiation, where electrons lose energy by spiralling around magnetic field lines. It is partly for this reason that the leading candidate of a launching and collimation mechanism for radio jets relies upon magnetic field lines (e.g [Heyvaerts and Norman 1989](#)).

2.3.2 Disk winds

Less obvious observationally, but perhaps more important energetically, are the disk winds that are expected (and observed) to be launched from accretion disks. A discussion of the launching mechanisms and different models of disk winds is presented in Chapter 3.

The clearest observational signature of an outflow is the P Cygni line profile, first seen in the stellar system of the same name. In the case of P Cygni itself, a luminous blue variable star, the outflow is a stellar wind driven from the surface of the star. However, the same feature is seen in CVs and BALs (see Figure 1.1 for an example), and here a disk wind is the proposed outflow.

The P Cygni feature in BALs is due to resonance line scattering of (probably disk) photons, and Figure 2.5 shows a cartoon illustrating how the classic profile is produced. The small red circle represents a continuum source, and the larger circle represents an extended cloud of outflowing gas. Photons coming directly to the observer pass through the blue segment of the outflow. If the outflow is not fully ionized, photons can scatter via interactions with line transitions in the gas. Each line occurs at a fixed frequency, but since the gas is accelerating away from the continuum source, the photon frequency is Doppler shifted into the rest frame of the gas. For photons coming towards the observer, the resonant frequency of the line is increasingly blue shifted, and so a broad, blue-shifted absorption feature will appear in the continuum, starting at the rest frame frequency of the line and extending to a frequency corresponding to the maximum outflow velocity.

Photons are scattered out of the observer's line of sight in the blue segment of the outflow. However, interactions in the rest of the outflow can scatter photons *into* the observer's sightline. Photons leaving the far side of the continuum source will encounter gas moving away from the observer, and so will produce a red-shifted resonance line feature (e.g. position A in Figure 2.5). Photons scattering from parts of the wind moving towards the observer (e.g. position C) will similarly produce a blue-shifted resonance line. The sum of all such scattering events into the observer's line of sight will produce a broadened emission line, centred at the rest frequency of the transition. This broad resonance line




Figure 2.5: Cartoon illustrating the Formation of a P Cygni line profile in an outflow.

will be superposed on the absorbed continuum to produce the characteristic P Cygni line shape illustrated at the bottom of Figure 2.5.

This basic argument is for a smooth, spherically symmetric outflow, however more complex profiles will result if the flow is clumpy or non-isotropic. Collisionally excited (thermal) line emission in the outflow or background source can also modify the line profile.

2.3.3 Feedback

Both types of outflow mentioned in the preceding sections have the capability of interacting with the local environment surrounding the accreting object (and indeed far distant environments, in the case of the radio jets). This is of great interest, particularly in the case of AGN, where the relatively tiny SMBH at the central of a galaxy appears to have a significant impact on the evolution of the entire galaxy surrounding it. This effect is observed, for example, in the so called $M - \sigma$ relationship (e.g. [Tremaine et al. 2002](#)) between the velocity dispersion of stars in the bulge of a galaxy and the mass of the black hole at its centre. Also, the overall mass contained in galactic bulges appears correlated with the mass of the central SMBH, $M_{\text{BULGE}} \simeq 10^3 M_{\text{SMBH}}$ (e.g. [Magorrian et al. 1998](#)).

There appear to be two distinct modes by which AGN outflows interact with the host galaxy. These are the radiative or “quasar mode” of feedback which is likely associated with disk winds, and the kinetic or “radio mode” of feedback, probably associated with radio jets ([Fabian 2012](#)). The latter mode is better studied, at least in part because radio jets are observationally easier to investigate than disk winds. Indeed, some jets carve out huge bubbles in the inter-galactic medium (IGM) around galaxies, which are directly observable as cool regions in the hot IGM.

Quasar mode feedback appears to be primarily momentum-driven, with an outflow from the central accreting object simply pushing gas out of the bulge of the galaxy and thus halting star formation. It is thought that this is the initial mechanism by which star formation is quenched, since it would occur early on in the evolution of a quasar, when the accretion rate would be high, and winds would be strong. This process would thus move galaxies from the actively star-forming “blue cloud” to the more passively evolving

“red sequence” (Granato et al. 2001; Schawinski et al. 2007). In contrast, radio mode feedback is thought to be a kind of “maintenance” feedback, constantly heating the intra-cluster medium (ICM). This inhibits new star formation from taking place via cooling flows in massive red galaxies in the centres of clusters and groups (McNamara and Nulsen 2007).

Chapter 3

Accretion disk winds in AGN

Evidence for powerful disk winds has been found in a wide range of disk accreting astrophysical systems, including CVs (e.g. [Cordova and Mason 1982](#)), X-ray binaries (e.g. [Ponti et al. 2012](#)) and AGN (e.g. [Tombesi et al. 2010](#); [Weymann et al. 1981](#)). Indeed, it seems that such outflows are a fundamental and universal aspect of accretion physics. It is widely assumed that these outflows arise from the surface of accretion disks, but the mechanisms by which they are launched and accelerated are not fully understood. In this chapter, we will briefly review the leading three basic scenarios for disk wind driving mechanisms and also introduce some of the specific models for AGN disk winds.

3.1 Driving mechanisms

In order to launch and accelerate a wind from an accretion disk, some kind of driving force must act upon the gas in order to overcome the inward pull of gravity. There are three likely sources for this driving force: gas/thermal pressure; radiation pressure and centrifugal forces mediated by magnetic fields. An introduction to each of these three driving mechanisms is given in the following sections.

3.1.1 Pressure driven or thermal disk winds

Pressure-driven winds are produced when the thermal speed of gas in the disk atmosphere exceeds the local escape velocity. This requires very high temperatures. For example, at a distance of 10^{18} cm from a $10^8 M_\odot$ SMBH, the escape velocity is about 1000 km s^{-1} . Hydrogen gas with this mean thermal speed has a temperature of about $3 \times 10^8 \text{ K}$, much higher than any temperature expected in a standard α – disk (see Figure 2.2). However, a simple geometrical argument provides a plausible way of achieving such high temperatures in the outer parts of the disk.

In a disk in hydrostatic equilibrium, the vertical component of the gravitational field of the central object must be balanced by the pressure of the gas in the disk. This requirement gives an expression for the scale height, H , of the disk as a function of radius, R (Kenyon and Hartmann 1987)

$$\frac{H}{R} = \left(\frac{v_s^2 R}{GM_*} \right)^{1/2}, \quad (3.1)$$

where v_s is the local sound speed in the disk, and M_* is the mass of the central object. Since $v_s^2 \propto T(r)$, and, in the standard α – disk model, $T(r) \propto r^{-3/4}$, we see that $H/R \propto r^{1/8}$, so the disk gets thicker with radius. In fact, the geometry can be more complex than this simple picture. The accretion rate through the disk can affect the structure of the innermost parts of the disk. If the accretion rate is more than a small fraction of the Eddington rate, then the inner parts of the disk will be radiation dominated, rather than pressure dominated and the inner disk will be thicker than expected from Equation 3.1. This is expected to be the case in AGN and particularly quasars (given their implied higher accretion rate).

With the proviso that a geometrically thick inner disk could modify the details, the flared shape of the outer disk means that UV and X-ray radiation produced in the inner disk can illuminate it. Since the outer parts of the disk are characterised by a relatively low density, the two-body cooling processes which control the temperature in a cool plasma (collisionally excited line emission and radiative recombination) are inefficient. Instead, heating and cooling will be dominated by Compton effects, and the temperature will

tend to the Compton temperature, around 10^7 K. In our example system, this would be sufficient to drive a thermal wind at a distance of about 10^{19} cm from the centre. Thus, if such winds do exist, they must be produced at relatively large distances from the SMBH, past the radius at which the escape velocity is equal to the isothermal sound speed at the Compton temperature (Begelman et al. 1983). This is the Compton radius, r_C given by

$$r_C = \frac{9.8 \times 10^9}{T_{C8}} \frac{M_{SMBH}}{M_\odot} \text{ cm}, \quad (3.2)$$

where T_{C8} is the Compton temperature in units of 10^8 K (Netzer 2006).

Begelman et al. (1983) carried out a theoretical investigation of the dynamics of thermal winds, and concluded that significant mass loss was possible. Furthermore, they suggested that this type of wind could provide the hot confining gas required in cloud models of the BLR (Krolik et al. 1981). It is less clear whether thermal winds could produce BALs in the UV, since the velocities of these outflows are too low, and the temperature is too high. They have, however, been invoked as a possible source of the so called “warm absorbers” that are thought to produce the relatively narrow X-ray absorption features observed in some AGN (Krolik and Kriss 2001).

3.1.2 Radiation-pressure or line-driven winds

We have already encountered radiation pressure in the calculation of the Eddington limit in Chapter 2. In that simplified treatment, Thompson scattering was taken to be the dominant interaction mechanism between matter and radiation. However, as mentioned there, interactions with *bound* electrons via resonance lines provide a far more efficient way of transferring momentum from photons to a partially ionized or neutral gas. Since the clearest indicator of outflowing gas, the P Cygni line profile, is formed through resonance line scattering, it is likely that line driving provides at least part of the accelerating force in any wind exhibiting such features.

The definitive paper on the subject of line driving is by Castor, Abbott, and Klein (1975), who investigated the effect of scattering and absorption by lines in the atmosphere of Of

stars. They couch their equations in the form of a “force multiplier”, $\mathcal{M}(t)$, which is the factor by which the line-driving force exceeds the force due to electron scattering alone. The radiation force, f_{rad} , is given by

$$f_{rad} = \frac{\sigma_e F}{c} \mathcal{M}(t), \quad (3.3)$$

where F is the total flux and σ_e is the mass scattering coefficient of free electrons. \mathcal{M} is a function of t (an optical depth variable scaled to remove the effect of line strength), which in turn depends on the velocity profile of the gas (both thermal and bulk) and the details of the lines doing the driving. The value of t is calculated in a summation carried out over all lines in resonance over the frequency range of the calculation. $\mathcal{M}(t)$ has a maximum value of around 2000 in ideal conditions (Gayley 1995), placing a rough lower limit on the luminosity required for efficient line-driving, $L \gtrsim 0.0005L_{edd}$. Although their initial treatment was designed for stellar winds, it has also been shown that line driving could produce winds from accretion disks (e.g Shlosman et al. 1985; Proga et al. 2000).

The force multiplier treatment of line driving was extended to disk winds by Stevens and Kallman (1990) who computed the line-driving force multiplier as a function of ionization parameter, ξ^1 , for X-ray binaries. This treatment included the ionizing effect of X-rays produced from the inner regions of the accretion disks in such systems, and they found that the multiplier drops off dramatically as ξ exceeds a value of around 100. This is because the number of lines decreases dramatically as gas becomes more ionized. Figure 3.1, reproduced from the paper, shows the response of \mathcal{M} to ξ . This data has since been used in the simulation of AGN scale disk winds (see section 3.2.1.2), which are also expected to be exposed to a strong X-ray source.

It turns out that it is quite challenging to find a simple wind geometry in which ξ remains sufficiently low to allow significant line driving. As a quick illustration, let us postulate that the BLR in quasars is associated with a wind, as in the model suggested by Elvis (2000). Considering a typical quasar X-ray luminosity of around $L_x = 10^{44}$ ergs s⁻¹, a

¹see equation 1.2 for a definition




Figure 3.1: The variation of the line-driving force multiplier $\mathcal{M}(t)$ with the ionization parameter ξ for a range of optical depth parameters t - reproduced from [Stevens and Kallman \(1990\)](#).

distance to the BLR region of $R_x \simeq 10^{16}$ cm, and a density of 10^{10} cm $^{-3}$, we obtain

$$\xi = \frac{L_x}{n_H R_x^2} = \frac{10^{44}}{10^{10} \times 10^{16} \times 10^{16}} = 100; \quad (3.4)$$

we are already at the limit of line driving. To make matters worse, we have neglected the contribution of the UV bright parts of the accretion disk to the ionizing flux in our estimate.

This back-of-the-envelope calculation is a good illustration of the main difficulty faced by line-driven disk wind models. However, there are good reasons to further investigate this driving mechanism. Firstly, as previously mentioned, the presence of P Cygni features in the UV spectra of BALQSOs demonstrates that there *must* be some degree of line driving. Secondly, some BALQSOs exhibit a spectral feature colourfully known as the “ghost of Ly α ” ([Arav et al. 1995](#); [Arav 1996](#)). Observationally, this is a slight hump seen in the blue shifted absorption trough of BALs, centred at around -5900 km s $^{-1}$. The proposed mechanism for its formation is explicitly linked to line driving. More specifically, if the outflow producing the BALs is line driven, and the BEL region is inside the BAL region, then the strong Ly α BEL will contribute to the acceleration of the outflow. Photons

from this line will interact with Nv in the part of the wind where the velocity relative to the BEL region is -5900 km s^{-1} , the separation in velocity space between the Ly α and N v $\lambda\lambda 1240$ lines. This interaction will give a little extra acceleration to this layer of the wind, causing a local reduction in optical depth and hence a “hump” in the absorption line profile. Since other BALs are formed in the same part of the wind, the optical depth in those transitions will also be reduced, causing slightly less absorption at a wavelength shifted from the peak by the same velocity. This feature may already have been seen in several BALQSOs (Arav et al. 1995; North et al. 2006, but also see Cottis et al. 2010).

In order to overcome the problems of over-ionization, most models of line-driven disk winds incorporate some kind of “shielding gas” to attenuate the ionizing radiation before it encounters the wind. Some of the most successful line-driven wind models are described later in this chapter, where we will see that there is a plausible physical explanation for the existence of this attenuating shield.

3.1.3 Magnetocentrifugal and magnetically confined winds.

Accretion disks are likely to contain magnetic fields, since infalling material will retain any field it had prior to joining the disk. Magnetic fields are also clearly an important component of the radio jets seen in many AGN. It is therefore natural to ask whether these fields could also be involved in driving the observed disk winds.

In the model proposed by Blandford and Payne (1982), gas is initially lifted from the disk surface by thermal expansion, but it is confined by poloidal magnetic field lines which are frozen into the disk and rotate with it. If the field lines make an angle of less than 60° to the disk, the centrifugal force drives the hot gas out along the field lines. The gas therefore gains rotational velocity, and this “bead-on-a-wire” motion provides a mechanism for the removal of angular momentum from the disk.

A critical point in the flow is reached at the Alfvèn radius, where the magnetic pressure equals the gas pressure. At this point, the magnetic field stops controlling the flow of the gas, and the gas starts to flow radially. However, the field lines are still frozen into the gas, and so the field lines “wind up”, causing the field to become more toroidal and

stronger. This mechanism is a possible source for the collimating fields of radio jets (e.g. [Heyvaerts and Norman 1989](#)).

Magnetocentrifugal winds do not require any kind of additional driving mechanism once the plasma has been launched. They are therefore good solutions in systems which either lack the luminosity for other types of driving (e.g. YSOs) or are over-ionized ([Proga 2007](#)). However, the magnetic field of the disk may also “confine” a wind predominantly driven by another mechanism, thereby maintaining a high density and low ionization parameter ([de Kool and Begelman 1995; Proga 2000](#)). This overcomes one of the main problems with line-driven winds - the tendency of the strong ionizing flux of the central source to overionize the wind and thus remove its driving mechanism.

3.2 Disk wind models

This section briefly reviews the geometries that have been calculated or postulated for a few of the most influential disk wind models. Wherever possible a figure is provided to illustrate the proposed geometry. This is not intended to be an exhaustive list; in particular, it is mostly restricted to models that attempt at least in part to explain the origin of BALs.

3.2.1 Radiation-driven winds

The observational clues of P Cygni line profiles in BALQSOs, and the ghost of Ly α offer strong evidence that line driving is occurring at some level in disk winds. The question is whether it can be the sole driving mechanism. In this section some of the models that aim to explain disk winds via line driving are introduced. A recurring theme is the requirement for some kind of shield to protect the outflow from overionization.

3.2.1.1 Murray and Chiang - a line-driven disk wind model

In the first of a series of influential papers, [Murray et al. \(1995\)](#) compute a disk wind model arising from an accretion disk surrounding a $10^8 M_\odot$ SMBH accreting at a rate of

$1.8M_{\odot} \text{ yr}^{-1}$ (about 50% of the Eddington luminosity for $\eta = 0.06$). They find a solution which produces a line-driven disk wind arising at radii outside about 10^{16} cm which carries away about 10% of the accreting gas. The opening angle of the wind is around 5° , which gives a covering factor of around 10% for the outflow as seen by the central source of UV radiation.

In order for line driving to be efficient, the ionization parameter must be $\xi \lesssim 100$ (see Figure 3.1). However, the X-ray source at the centre tends to overionize the wind. The solution proposed by Murray and Chiang takes the form of so-called “hitchhiking gas”, which is drawn along with the line-driven gas due to pressure gradients. It rapidly becomes overionized and falls back, but shields the outflowing wind in the process. Their figure illustrating the geometry is reproduced below as Figure 3.2.




Figure 3.2: Schematic of line driven disk wind proposed by [Murray et al. \(1995\)](#).

[Murray et al. \(1995\)](#) also demonstrate that BELs can be produced from the dense base of the wind, while BALs are formed in the outer parts of the wind where it has been accelerated to large outflow velocities. This model was fleshed out in several follow up papers. For example, [Murray and Chiang \(1995\)](#) suggest that the material at the innermost edge of the wind could be the source of the warm absorbers seen in X-rays, whilst [Murray and Chiang \(1997\)](#) demonstrate that the characteristic single-peaked line profile shapes of BELs are a natural consequence of line formation in a disk wind.

3.2.1.2 Proga, Stone & Kallman - a hydrodynamical, line-driven disk wind model

Using the ZEUS hydrodynamic code, with some modifications to include the effect of line-driving, [Proga, Stone, and Kallman \(2000\)](#), hereafter referred to as PSK) produced the

first of a series of hydrodynamic models of line-driven disk winds. In these simulations, the ionization state of the gas in the model was set by the X-ray luminosity of a centrally located source, L_X , taking account of continuum attenuation between the source and each point in the wind. This gives an ionization parameter for each part of the wind, which is then translated into a line driving force multiplier using the parametrisations from [Stevens and Kallman \(1990\)](#) (Figure 3.1).

In their model the central source was also a source of UV photons, with luminosity L_{UV} , and these photons together with those produced from the central parts of the disk interact with the ionized gas, via resonance lines, to drive a disk wind. The temperature of the gas was also computed from the ionization parameter, and then the equations of motion were solved. The solution is reproduced here in Figure 3.3. In this initial model, the mass of the central object is assumed to be $10^8 M_\odot$, and the accretion rate is taken to be $\dot{M}_{acc} = 1.8 M_\odot \text{ yr}^{-1}$ (50% of the Eddington rate). This is used to calculate the total disk luminosity L_D , and the luminosity of the central object is then set equal to L_D and split equally between L_{UV} and L_X .

PSK report that a strong outflow forms in this simulation, carrying a mass of $\dot{M}_{wind} = 0.5 M_\odot \text{ yr}^{-1}$ and starting at a radius of about 10^{16} cm . The flow quickly accelerates to the velocities required to produce BAL profiles in line with observations. The covering factor of this wind is around 0.2, centred at angles around 75° . However, the flow is unsteady, producing dense knots on the timescale of a few years in the rest frame of the quasar. Importantly, they also find a region inside the launching radius where a wind is initially lifted from the disk, but is overionized by the central X-ray source, and thus falls back to the disk. This failed wind provides shielding for the successful wind, doing the job of the hitch-hiking gas in [Murray et al. \(1995\)](#).

This model is further tested in [Proga and Kallman \(2004\)](#), in which only disk photons are counted as contributors to the line-driving force, and the central X-ray source is only attenuated by electron scattering (although it is assumed to have a lower luminosity of $L_X = 0.1 L_D$). Even under these less favourable conditions, a wind is still launched; its density and velocity fields are reproduced in Figure 3.4. Synthetic profiles for the




Figure 3.3: Wind parameters for a hydrodynamic disk wind model reproduced from [Proga et al. \(2000\)](#).

C IV absorption line were also presented and shown to be consistent with the shapes expected for BAL features.




Figure 3.4: Wind parameters for hydrodynamic disk wind model reproduced from [Proga and Kallman \(2004\)](#).

3.2.1.3 Risaliti and Elvis - a non-hydrodynamical line-driven disk wind model

A related wind model was presented by [Risaliti and Elvis \(2010\)](#). They consider a SMBH mass of $2 \times 10^8 M_\odot$ accreting at 50% of the Eddington rate and model it both analytically and with a code called QWIND. This code calculates the radiation pressure on the wind based upon a simplified estimate of the local ionization parameter. They do not solve the equations of hydrodynamics for the wind, rather they concentrate on more approximate calculations which allows a wider range of different parameters to be investigated.




Figure 3.5: Wind streamlines for a line-driven disk wind as calculated by [Risaliti and Elvis \(2010\)](#).

They find that there is an inner region where the wind becomes overionized and falls back to the disk surface, forming a shield similar to the hitchhiking gas postulated by [Murray et al. \(1995\)](#) and the failed wind found by [Proga et al. \(2000\)](#). At large radii, the UV flux from the disk is too low to accelerate the wind to escape velocity, and so once more the wind fails. However, in between these two regions, at around 200 Schwarzschild radii ($\sim 6 \times 10^{15}$ cm) there is a thin region where the ionization state is low enough, and the UV flux high enough, to allow efficient line driving. This drives an outflow at about 20° above the disk plane, with a covering fraction of a few percent. One of their wind models is reproduced here as Figure 3.5.

Investigating the effect of the Eddington ratio of the central objects, they find that the wind moves outwards and accelerates as L/L_{Edd} increases towards 0.5, but then decreases

as the wind becomes more ionized, and line-driving becomes less efficient. For ratios less than about 0.3, no wind is launched.

Very similar calculations were carried out by [Nomura et al. \(2013\)](#). However, they used an improved ionization scheme and a more realistic mass outflow rate. A cartoon illustrating their calculation is reproduced in the left hand panel of Figure 3.6, along with one of their wind solutions for a $10^8 M_\odot$ SMBH accreting at 50% of the Eddington rate.




Figure 3.6: Cartoon of proposed geometry (left panel) and simulated flow (right panel) from [Nomura et al. \(2013\)](#).

They find that BAL features would be expected for a range of SMBH masses between $10^7 - 10^{8.5} M_\odot$ and for a range of Eddington ratios from 0.3 to 0.9. For larger SMBH masses they do not find a wind, because the cooler disk produces an insufficient UV flux to drive it.

3.2.2 Magneto-centrifugal winds

Most models of magnetically driven winds do not suggest that the magnetic effects are the sole accelerating mechanism. Some use the magnetic field lines to constrict the outflowing gas, thereby increasing its density and reducing the ionization state to the point where line driving becomes possible. Others include the “bead on a wire” mechanism of magneto-centrifugal acceleration as both an additional driving mechanism and a way of extracting angular momentum from the disk. In this section, some of the models which explicitly rely on magnetic fields are reviewed.

3.2.2.1 Emmering, Blandford and Shlosman - a model of a magneto-centrifugally accelerated disk wind

Emmering et al. (1992) set out to explain BELs in terms of dense molecular clouds flung out from the outer parts of the accretion disk by magnetic stresses. These clouds quickly become exposed to the ionizing flux from the central parts of the disk and thus become photoionized. The clouds remain dense enough to maintain the ionization state seen in BLR regions, because they are confined by the magnetic field lines. A cartoon illustrating the model is reproduced in Figure 3.7.




Figure 3.7: Schematic of magnetically driven disk wind proposed by Emmering et al. (1992).

The clouds are launched from the outer part of the disk, with a typical launch distance of about a parsec for an AGN with ionizing luminosity of 10^{46} ergs s^{-1} . At this distance, the bulk velocity of the clouds would not produce the observed line widths, and so line broadening due to high temperature electrons in an intra-cloud medium is invoked. They are able to show that the typical cloud sizes, electron densities, filling factors and velocities are comparable with observations. They also suggest that, with some modifications, BALs could be explained by their model.

3.2.2.2 de Kool and Begelman - a magnetically confined, radiation-driven disk wind

A simple disk wind model in which magnetic fields are included as a factor is presented by [de Kool and Begelman \(1995\)](#). They investigate whether BAL features can be produced in a clumpy wind arising in the outlying parts of a disk. Using similar arguments to [Emmering et al. \(1992\)](#), they suggest that such an outflow could first be confined by magnetic field lines threading the disk and wind, and then compressed by radiation pressure from the central source. An illustration of the wind geometry from their paper is reproduced in Figure 3.8.




Figure 3.8: Illustration of a clumpy, magnetically confined wind reproduced from [de Kool and Begelman \(1995\)](#).

The initial launching mechanism is suggested to be magnetic buoyancy or interactions between cool molecular clouds in the disk, which gives some clouds a small vertical velocity. As the clouds leave the disk, they are exposed to the radiation from the central source and are heated up. The magnetic field prevents sideways expansion, and the plasma is “squeezed” up between the field lines, becoming a filamentary outflowing wind. The wind is then accelerated outwards via UV line driving.

This model suggests that the wind is launched at about 10^{18} cm from the central source, about 100 times further out than the line-driven disk wind of [Murray et al. \(1995\)](#). At such a large distance, there is no issue in maintaining the ionization state required to

show UV BALs, and no shielding gas is required. They find that a mass-loss rate of $\sim 1 M_{\odot} \text{ yr}^{-1}$ is required for a $10^8 M_{\odot}$ SMBH to produce the right covering factor to explain the BAL fraction.

3.2.2.3 Everett - a hybrid line-driven / magneto-centrifugal disk wind

[Everett \(2005\)](#) explores how line driving can be combined with a magneto-centrifugal model of a disk wind from a disk around a $10^8 M_{\odot}$ SMBH. He chooses an Eddington ratio of $L/L_{edd} = 0.01$ for his fiducial model and splits the wind into two components to simplify the calculations. At radii close to the assumed launch radius of $3 \times 10^{16} \text{ cm}$, the wind becomes over-ionized, but the magneto-centrifugal effect means that the wind is still able to accelerate outwards. This inner wind then provides a shield for an outer wind, which is in a lower ionization state and can therefore be accelerated via line driving to very high velocities. A schematic of this model is reproduced here as Figure 3.9.

When the Eddington ratio of the central source is increased, the wind tends to become overionized, even with the shield, and line driving declines. For lower values of L/L_{edd} , the UV flux is insufficient for line driving.




Figure 3.9: Schematic of the hybrid line-driven / magneto-centrifugal disk wind proposed by [Everett \(2005\)](#).

This model was extended somewhat by [Gallagher and Everett \(2007\)](#) with the incorporation of a third wind component from the dusty torus. The driving mechanism for this last component is suggested to be UV radiation pressure on dust grains. This dusty wind is highly equatorial and would be expected to produce the strong absorption seen in the rarely observed type 2 Quasars. This extended geometry is reproduced in Figure 3.10.




Figure 3.10: Schematic of stratified disk wind proposed by [Gallagher and Everett \(2007\)](#).

Chapter 4

Monte Carlo Radiative Transfer

This chapter introduces our photoionization and radiative transfer code, PYTHON, first described in [Long and Knigge \(2002\)](#). Significant modifications to the code have been made in the process of this project, but the basic approach is unchanged, and this is described in detail here. In addition, all of the physics included in the original version of PYTHON are still used in the new version, and so these are also described here. The main exception to this is the ionization calculation, where a new approach has been introduced which improves the treatment of AGN spectral energy distributions (SEDs). This updated scheme is discussed in the next chapter, but the original scheme is described here for completeness.

4.1 Radiative transfer - the need for Monte Carlo

Radiative transfer (RT) calculations treat the passage of a ray of given specific intensity through a medium. As the ray progresses, the specific intensity can increase if the medium is radiating energy, and it can decrease due to opacity. These two processes are summarised in the RT equation

$$\frac{dI_\nu}{ds} = -\kappa_\nu I_\nu + j_\nu, \quad (4.1)$$

where dI_ν/ds is the rate of change of specific intensity I_ν with distance travelled along a path segment of length ds , κ_ν is the absorption coefficient and j_ν is the emission coefficient. In order to solve the RT equation, it can be recast in terms of another variable, the “optical depth”

$$\tau_\nu = \kappa_\nu ds. \quad (4.2)$$

Dividing both sides of Equation 4.1 by κ_ν the RT equation can be rewritten as

$$\frac{dI_\nu}{d\tau_\nu} = -I_\nu + S_\nu, \quad (4.3)$$

where S_ν is the “source function” for the medium, defined as

$$S_\nu = \frac{j_\nu}{\kappa_\nu}. \quad (4.4)$$

The RT equation can now be solved if the intensity and source function are expressed as functions of the optical depth. The formal solution of the RT equation is

$$I_\nu(\tau_\nu) = I_\nu(0)e^{-\tau_\nu} + \int_0^{\tau_\nu} e^{-(\tau_\nu - \tau'_\nu)} S_\nu(\tau'_\nu) d\tau'_\nu, \quad (4.5)$$

where the first term on the right hand side refers to the attenuation of the ray, and the second term is the contribution of emission in the medium along the path the ray has taken. If the geometry is sufficiently simple, then this equation can be used to compute an emergent spectrum. However, if one wishes to calculate RT through a geometry with little or no symmetry or with a clumpy distribution of opacity and emission, then the problem becomes analytically intractable. In such cases the equation of RT needs to be solved numerically.

The Monte Carlo method is one such numerical method. The fundamental idea is to simulate RT through a distribution of material by tracing the path of individual photons through a computer representation of the medium under consideration. As the photons progress through the model, they can be absorbed by the medium and more can be generated through radiative processes. When they finally exit the medium, they can be added up to produce an emergent spectrum. The method is named after the casino

on the French Riviera because it is largely probabilistic, starting with the generation of photons. These are produced randomly in such a way as to properly represent the photon distribution that would be seen in the real object. There are many variants of the Monte Carlo method, and in the rest of this chapter we will describe how the method is implemented in our RT code, PYTHON.

4.2 Overview of the method

An RT simulation in PYTHON begins with the construction of a computational grid of “cells”. Physical parameters such as density, temperature and velocity are associated with each cell. These are either calculated from a kinematic prescription for a wind, or input directly from an externally defined wind model. The initial ionization state of the wind is set by assuming each cell is illuminated by a dilute blackbody. The mean intensity in the cell is therefore

$$J_\nu = WB_\nu(T_r), \quad (4.6)$$

where the radiation temperature, T_r , is set to the initial temperature of the cell, and the dilution factor, W , is *initially* assumed to be purely geometric, i.e.

$$W = 0.5 \left((1 - \sqrt{\frac{1 - R_*^2}{r^2}}) \right), \quad (4.7)$$

where R_* is the radius of the central star and r is the distance of the cell from the star.

A system of radiating sources is then defined, which can include an accretion disk and a central object. Photons are produced from these sources and then tracked through the wind. They are used to compute the heating in grid cells via the various heating processes discussed in section 4.5.1, and therefore an improved estimate of the cell temperature. They are also used to compute the frequency-averaged mean intensity, J , and mean frequency, $\bar{\nu}$, for the cell.

In the original version of the code described here, these parameters are in turn used to calculate improved estimates of T_r and W . This results in a more accurate model of J_ν , based upon the radiation which actually reaches the cell and thereby taking some account

of absorption and emission between the source and the cell. This improved model of J_ν , along with the new temperature, is then used to calculate an updated ionization state. At this point, a new population of photons is generated, and the process is repeated in an iterative fashion until the temperature of each cell converges.

Once the wind temperature and ionization state have converged, synthetic spectra may be computed. To this end, photons are once more produced and propagated though the converged wind model. These photons are captured as they leave the model and used to produce synthetic spectra at user-defined sightlines. The photons are typically produced over a reduced frequency range to improve signal-to-noise, and other variance reduction techniques are used to further improve the spectrum. This execution cycle is summarised below in Figure 4.1.




Figure 4.1: An overview of the execution cycle of a PYTHON simulation.

4.3 Radiating sources

In PYTHON, emission processes are dealt with separately from absorption processes. The former are treated as sources of photon packets, generated at the start of a calculation cycle, whilst the latter are calculated as the photon packets pass through the wind. The wind itself can generate photon packets. However, once converged, the wind is assumed to be in thermal and radiative equilibrium and so it is not a *net* source of energy, rather it reprocesses absorbed packets. This section deals with the different sources of photons, the physics of emission and the details of how the various mechanisms are implemented.

4.3.1 Photon packet production

The photon packet, defined by its frequency ν and its weight w , is the basic modelling unit in PYTHON. In the most simple implementation, all photon packets have the same weight. This weight is equal to the total luminosity of all radiating sources integrated over the frequency range required, divided by the number of photon packets, N_{phot} , i.e.

$$w = \frac{\sum_{sources} \int_{\nu_{min}}^{\nu_{max}} L_{\nu}}{N_{phot}}. \quad (4.8)$$

Thus w has units of luminosity, ergs s⁻¹.

The frequency of a packet is chosen randomly in such a way as to produce the correct SED. This is achieved by converting the required SED into a cumulative distribution function, P_{ν} ,

$$P_{\nu} = \frac{\int_{\nu_{min}}^{\nu} L_{\nu}(\nu') d\nu'}{\int_{\nu_{min}}^{\nu_{max}} L_{\nu}(\nu') d\nu'}. \quad (4.9)$$

The cumulative distribution function (CDF) monotonically increases from 0 at ν_{min} to 1 at ν_{max} . Figure 4.2 illustrates the form of the CDF for a Planck function with $T_{eff} = 10^5$ K.

In order to select a photon frequency, one therefore calculates a random number from a uniform distribution between 0 and 1 and inverts the CDF to find the frequency corresponding to this probability. In practice this approach has a few limitations, so various techniques are employed to improve results.




Figure 4.2: A Planck function with $T_{eff} = 10^5$ K and associated CDF.

First, it is often desirable to ensure that a sufficient number of photons are produced at certain frequencies, for example around photoionization edges. If these frequencies lie on the Wien tail of a blackbody distribution, it is rather unlikely that many such photons would be generated, since the vast bulk of the luminosity is produced at lower frequencies. We see this in Figure 4.2, where the CDF at 2×10^{16} Hz (near the photoionization edge of the important ion C IV) is 0.987. Thus if we sample directly from the CDF, only 13 photons in every 1000 would have a frequency above this value. However, the precise shape of the Wien tail is critical in accurately modelling the photoionization state of a plasma. It is therefore important to produce a sufficient number of energy packets from this part of the spectrum to ensure that this population of high energy photons has a fair chance of propagating to all parts of the wind.

In order to cope with such cases, PYTHON implements a strategy called “importance sampling”. Here, the frequency range over which photons are generated is subdivided into bands. The weight of photon packets is permitted to differ between bands, which allows one to produce a large number of low-weight photons to represent parts of the SED far away from the peak luminosity that are still important to the calculation. Following changes to the ionization scheme described in Chapter 5, this “banding” of photon generation has become much more important, since accurate modelling of the spectrum in a given cell requires good sampling across the frequency range. The relatively

more simple dilute blackbody model of the spectrum in a cell (used in the original version of PYTHON described here; Equation 4.6) is more forgiving of sparse sampling. This is because all that is required is a good estimate of the radiation temperature and frequency-averaged mean intensity (see section 4.6).

The second limitation is a computational one: it is not always possible to construct an algebraic expression for the CDF. For example, in the case of a power-law SED, the CDF is a simple function, and one can easily invert it. However, in the case of a blackbody, it is not possible solve Equation 4.9 analytically. Photon packets can also be produced by line emission, and the SED for this process is more complex still. In such cases, the CDF is constructed numerically on a fixed grid of P_ν . A random number is then used to select a bin, and the exact frequency chosen for the photon is calculated by interpolating between the two bin limits. This interpolation can take account of the known form of the SED - e.g. if we are in the Rayleigh Jeans or Wien limits of a blackbody spectrum, or failing that, linear interpolation can be used.

Having described the basic mechanism of photon generation, the next sections describe the three main sources of photons in the original version of PYTHON.

4.3.2 Photon source 1: the central source

Prior to the start of the current project, most models contained a central star, defined by its temperature T_* , its radius, R_* and its mass M_* . The radiation from the star can either be modelled as a simple blackbody or using externally defined stellar atmosphere models. The Eddington approximation of limb darkening is also accounted for, in the CDF for the photon packet's direction of emission. One may also define an additional source of photons, the “boundary layer”, which is of use in modelling CVs where high frequency photons can be produced in a transition region between the accretion disk and the surface of the star (Popham and Narayan 1995). In an AGN model, the star is replaced by a source of X-rays (see section 5.2).

4.3.3 Photon source 2: the accretion disk

The accretion disk is modelled as an α -disk, described in section 2.2.2. This allows calculation of the luminosity of an annulus of the disk at radius r of width Δr

$$L_{\text{annulus}} \simeq 2\pi r \Delta r \sigma T^4. \quad (4.10)$$

The disk is radially sub-divided into equal-luminosity annuli, and a random number is used to select from which annulus a photon will be emitted. The temperature of each annulus is assumed to be constant, and the frequency of the photon generated is normally chosen from a blackbody distribution or an imported stellar atmosphere model. In the latter case, the local surface gravity is calculated taking into account the central object only. The actual location of the photon packet emission within an annulus is randomised along with the direction of emission, which takes account of foreshortening and limb darkening.

4.3.4 Photon source 3: the wind

The wind is also a source of radiation, generating photon packets via free-free (bremsstrahlung), free-bound (recombination) and thermal, collisionally excited, bound-bound (line) emission. In a converged model, the wind is treated as being in radiative equilibrium, so it is not a *net* source of photons, but rather a reprocessing medium. Thus the luminosity of a converged wind cell summed over all three emission processes is equal to the radiative heating rate in that cell, less any non-radiative cooling due to adiabatic expansion. Before convergence, this balance does not hold. Hot cells can cool down by emitting more energy than they absorb, and cool cells may not emit all the energy they absorb. In the first ionization cycle, before a heating rate for each cell is known, wind emission is calculated using an initial estimate of the wind temperature. The details of the emission processes as implemented in PYTHON are introduced in the rest of this section.

4.3.4.1 Free-free emission

The free-free emissivity per unit volume is given by

$$4\pi j_{\nu}^{ff} = n_e n_i \frac{2^5 \pi}{3} \left(\frac{2\pi}{3} \right)^{1/2} \frac{Z^2 e^6}{m^2 c^3} \left(\frac{m}{kT_e} \right)^{1/2} g^{ff} e^{-h\nu/kT_e} \quad (4.11)$$

for an ion j of density n_i and charge Z . When generating photons from the wind, this expression is integrated over the frequency range of the simulation and summed over all ionic species present. In practice, in typical conditions, the free-free emissivity will always be dominated by interactions of electrons with H II, He II and He III ions, so the summation is only carried out over these three species. The gaunt factor, g^{ff} , is set to unity. So, in a given band,

$$L^{ff} = V \sum_{ions} \int_{\nu_{min}}^{\nu_{max}} 4\pi j_{\nu}^{ff} d\nu. \quad (4.12)$$

4.3.4.2 Free-bound emission

The free-bound emissivity due to an ion in the ground state of ionization state $i + 1$, recombining to level l of ionization state i is calculated via the Milne relation from the photoionization cross section $\sigma_{\nu(i,l)}^{bf}$,

$$4\pi j_{\nu(i,l)}^{fb} = n_e n_{i+1,l} h\nu \left(\frac{2\pi m k}{h^2} \right)^{-3/2} \frac{8\pi}{c^2} \frac{g_{i,l}}{g_e g_{i+1,l}} T_e^{-3/2} \nu^2 \sigma_{i,l}^{bf}(\nu) e^{-(\nu - \nu_t)/kT_e}. \quad (4.13)$$

Thus the total luminosity due to free-bound emission is

$$L^{fb} = V \sum_{ions} \sum_{levels} \int_{\nu_{min}}^{\nu_{max}} 4\pi j_{\nu(i,l)}^{fb} d\nu, \quad (4.14)$$

where the summation is carried out over all ions and also all levels. The level populations for excited levels are calculated using

$$n_{i,l} = n_i \frac{W g_l e^{-E_k/kT_r}}{z_i(T_r)}, \quad (4.15)$$

where z_i is the partition function for the ion calculated from

$$z_i = \sum_{levels} W g_l e^{-E_l/kT_r}. \quad (4.16)$$

The statistical weight of level l is g_l , and E_l is the excitation energy. If $W = 1$, and $T_r = T_e$, the requirements for local thermodynamic equilibrium (LTE), then the equation reduces to the normal Boltzmann equation. This represents a first-order approximation to a non-LTE level population, where the levels are radiatively excited. The radiative weight W and the radiation temperature T_r are the parameters calculated for the black-body model of the mean intensity in each cell (Equation 4.6).

4.3.4.3 Line emission

The final emission process is thermal, collisionally excited, line emission. To understand this process, it is worth a quick diversion into the microphysics of line interactions. When a photon interacts with an ion via a line interaction, the ion in question is excited to a higher energy state, and the photon is destroyed. This process can be described as “line opacity”. The excited ion can relax in one of two ways. Firstly, it can radiatively de-excite, thereby re-emitting essentially the same photon; the photon is said to have scattered. Alternatively, the excited ion can be collisionally de-excited, and the photon is lost - this process is true absorption.

An ion can also be excited by interactions with the electron pool - this is collisional excitation, and the excited ion can once again relax collisionally (in which case the energy is returned to the electron pool) or radiatively, in which case a photon is generated. This latter process is what we would call collisionally excited line emission, and this is what we want to treat as line emission from the wind.

The source function for line emission, taking account of both scattering and true absorption is

$$S_\nu = (1 - q)J_\nu + qB_\nu(T_e), \quad (4.17)$$

where q is the fraction of excited ions which relax collisionally rather than radiatively. The first term therefore represents the scattered photons from the incoming radiation and the second term is the thermally generated photons. PYTHON treats line emission in the “two-level atom” approximation. The source function for line transitions between upper level u and lower level l can be written as

$$S = \frac{j^{line}}{\kappa^{line}} = \frac{n_u A_{ul}}{n_l B_{lu} - n_u B_{ul}}, \quad (4.18)$$

where A_{ul} is the Einstein coefficient for spontaneous emission, B_{lu} is the coefficient for induced absorption, and B_{ul} is the coefficient for induced emission. The number densities in the two levels (n_u and n_l) are related by the rate equation,

$$n_l(c_{lu} + B_{lu}J_\nu) = n_u(c_{ul} + A_{ul} + B_{ul}J_\nu) \quad (4.19)$$

where c_{lu} and c_{ul} are the collisional rates for excitation and de-excitation respectively. Combining these expressions, and comparing the result with Equation 4.17 allows an expression for q to be written down,

$$q = \frac{c_{ul} (1 - e^{-h\nu/kT_e})}{A_{ul} + c_{ul} (1 - e^{-h\nu/kT_e})}. \quad (4.20)$$

From Equation 4.17, the line emission coefficient is given by

$$j^{line} = q\kappa^{line}B_\nu(T_e). \quad (4.21)$$

The total line absorption coefficient is given by

$$\kappa^{line} = \frac{h\nu}{4\pi}(n_l B_{lu} - n_u B_{ul}) = \frac{c^2}{8\pi\nu^2} \left(\frac{g_u}{g_l} n_l - n_u \right) A_{ul}, \quad (4.22)$$

and thus we can write an expression for the line emission coefficient

$$j^{line} = q \frac{h\nu}{4\pi} \left(\frac{g_u}{g_l} n_l - n_u \right) A_{ul} \left(\frac{1}{e^{h\nu/kT_e} - 1} \right). \quad (4.23)$$

The total line luminosity for a cell of volume V is given by

$$L^{lines} = V \sum_{lines} 4\pi j^{line}. \quad (4.24)$$

However, not all photons produced by a particular line transition in a cell would be expected to escape that cell, due to interactions with the same line in neighbouring atoms. This is called “line trapping”, and it turns out that Equation 4.23 requires only slight modification to take account of this.

The chance of a line photon escaping the region in which it is generated can be described in terms of the “escape probability” formalism. The probability that such a photon will escape in a given direction \hat{n} is given by the escape probability,

$$P_{esc} = \frac{1 - e^{-\tau(\hat{n})}}{\tau(\hat{n})}, \quad (4.25)$$

where $\tau(\hat{n})$ is the optical depth of the cell in the direction \hat{n} .

Therefore, the proportion of photons which escape the region in which they are generated is just the angle-averaged escape probability $\langle P_{esc} \rangle$. This means that a fraction $1 - \langle P_{esc} \rangle$ will be reabsorbed locally. As already discussed, such interactions result in an excited ion which can de-excite radiatively or collisionally. In the latter case, the photon energy returns to the electron pool, thereby reducing the expected line luminosity. The fraction of excited ions that collisionally de-excite is just the probability of collisional de-excitation, q . We define a new “absorption fraction”, Q , which is the fraction of photons that are absorbed by a line interaction, taking account of line trapping:

$$Q = \frac{q}{q + \langle P_{esc} \rangle (1 - q)}. \quad (4.26)$$

We can then re-write Equation 4.23 to take account of this by simply replacing q with Q ,

$$j_{esc}^{line} = Q \frac{h\nu}{4\pi} \left(\frac{g_u}{g_l} n_l - n_u \right) A_{ul} \left(\frac{1}{e^{h\nu/kT_e} - 1} \right). \quad (4.27)$$

The actual line luminosity of the cell is then

$$L^{lines} = V \sum_{lines} \langle P_{esc} \rangle 4\pi j_{esc}^{line}. \quad (4.28)$$

This is the equation for line emission used in the simulations presented in this thesis.

Equation 4.27 contains n_u , the number density of ions in the excited state. This is effectively a correction term, which takes account of stimulated emission. The ratio between number densities in the two levels is calculated in the on-the-spot approximation¹ using Equation 4.19, assuming that the mean intensity can be modelled as a dilute blackbody (Equation 4.6). This gives

$$\frac{n_u}{n_l} = \frac{c_{lu} + B_{lu}WB_\nu(T_r)}{c_{ul} + A_{ul} + B_{ul}WB_\nu(T_r)}. \quad (4.29)$$

The collisional de-excitation rate coefficient, c_{ul} , is calculated from the collision strength, which is in turn calculated using the g-bar, or van Regemorter (1962) approximation. The collisional excitation coefficient c_{lu} is calculated from c_{ul} using a detailed balance argument. The population in the lower level n_l is calculated using Equation 4.15 if that level is the ground state. If both levels are excited, then the population of each is calculated with respect to the ground state separately.

4.3.4.4 Wind photon weights and directions

The total wind luminosity is the sum of the contributions from all three mechanisms,

$$L^{wind} = L^{lines} + L^{fb} + L^{ff}, \quad (4.30)$$

and is added to the disk and central source luminosities to give the total luminosity. The weight of each photon packet is then calculated, with the number of photons generated from each source set to be in proportion to the source luminosity. The number of photons produced by each wind emission mechanism is split in the same way.

¹The on-the-spot approximation assumes all locally emitted radiation is reabsorbed close to where it is produced, meaning that the local mean intensity is dominated by the radiation arriving from remote sources.

The frequency of free-free and free-bound photon packets is decided via a numerically calculated CDF, just as for blackbody photons. For line photon packets, the frequency is chosen by selecting which transition emits the photon, based on a CDF calculated as a sum over the frequency-ordered lines

$$P(\nu_i) = \frac{\sum_{i'=0,i} L_{\nu_{i'}}}{L_{\text{lines}}}. \quad (4.31)$$

Free-free and free-bound wind photons are emitted isotropically, but line photons are emitted with a directional CDF based upon the angle-dependent escape probability, given by Equation 4.25.

Once the photon packets have been launched, they are tracked through the wind structure. In the next section, we will briefly discuss the sources of opacity in the wind; the method by which the path of a photon is determined; and how the estimators of the radiation field in each cell are computed. These are used to determine the ionization state of the wind.

4.4 Radiative transfer in the wind

The equation of radiative transfer (Equation 4.1) includes both emission and absorption processes in the calculation of the rate of change of intensity as a beam of radiation passes through a medium. In PYTHON they are treated separately. Emission takes the form of the production of photon packets by the wind and was discussed in section 4.3.4.

Absorption is formally split into scattering and true absorption, although it is worth noting that this is an artificial distinction. Physically, scattering is just absorption followed quickly by re-emission of an “identical” photon. It is convenient to treat this as the *same* photon, thereby making the distinction between scattering and true absorption (where the photon is “destroyed”). Scattering processes therefore do not change the weight of a photon packet, but they may cause the packet to change direction. True absorption processes cause a reduction in weight of the photon packet, with the balance of the energy going into heating the wind.

PYTHON therefore splits absorption processes into scattering and true absorption processes. Electron scattering and line scattering are treated as pure scattering processes, causing the photon packet to change direction without loss of weight². In contrast, free-free and photoionization processes are treated as true absorption processes, which can remove energy from a photon packet along its entire path through the wind, but do not change its direction.

4.4.1 Tracking the path of photon packets

The path a photon packet takes through the wind is a series of straight path segments, along which its weight can be reduced by true absorption. These path segments are linked by scattering events. Thus the procedure for tracking a photon packet involves searching for locations at which the packet will scatter. This is a probabilistic process, based upon the cumulative scattering optical depth along the path segment.

When a photon packet enters a grid cell containing plasma, a randomised optical depth, \mathcal{T} , is generated

$$\mathcal{T} = -\ln(1 - p), \quad (4.32)$$

where p is a random number with a uniform distribution between 0 and 1. As the packet moves through the the cell, the cumulative scattering optical depth, τ_{scat} due to electron and line scattering along its path is calculated. If/when this equals \mathcal{T} , the packet scatters. If the packet reaches the boundary of a cell, the accumulated value of τ_{scat} is retained, along with the \mathcal{T} , and the photon continues into the next cell. This process continues until the photon eventually scatters or leaves the wind. Once a new scattering location is found, the total *absorption* opacity along the path segment from the last scattering location is calculated, and used to reduce the weight of the packet - this lost energy heats the cells that the packet has passed through. The procedure continues until the photon packet leaves the wind, strikes an absorbing surface, or has its weight reduced to zero by continuum opacity. The procedure is summarised in Figure 4.3 below.

²In fact, when scattering from a line, the weight of the photon packet *can* be reduced since, as already discussed, a proportion of the excited ions will de-excite via collisional processes rather than radiative ones. This is taken account of in the code - see Section 4.4.4.




Figure 4.3: An overview of the radiation transfer scheme.

The following sections briefly describe the implementation of the various scattering and absorption mechanisms, starting with electron scattering.

4.4.2 Scattering processes

4.4.2.1 Electron Scattering

The electron scattering absorption coefficient in the Thompson limit depends only on the electron density, n_e , and the Thompson scattering cross section, σ_T

$$\kappa^{es} = \sigma_T n_e. \quad (4.33)$$

The electron scattering optical depth along a path of length ds is then simply

$$\tau^{es} = \kappa^{es} ds, \quad (4.34)$$

which simply increases monotonically along the path of a photon packet through a cell. If electron scattering were the only scattering process, then discovering the point where the packet scatters would be a simple matter. However, in the situations we are interested in, line opacity is important.

4.4.2.2 Line opacity and the Sobolev approximation

Line opacity in a wind requires special treatment, since it is not a continuous process. PYTHON treats line opacity using the Sobolev approximation, in which each line is treated separately, with resonance events taking place only on physically distinct “Sobolev surfaces”. These surfaces are where, through Doppler shifting of the photon frequency into the local frame, a photon comes into resonance with a given line. Therefore, the process by which a photon is propagated through the wind involves calculating the locations of these surfaces in the wind. A photon packet of frequency ν will enter resonance with a line with rest frequency ν_{lu} at a location in the wind where the local velocity v is such that

$$\nu = \nu_{lu} \left(1 - \frac{v}{c}\right). \quad (4.35)$$

At this point, the optical depth associated with the line is calculated via the equation

$$\delta\tau = \frac{\pi e^2}{mc} f_{lu} \left(n_l - n_u \frac{g_l}{g_u} \right) \frac{c}{\nu_{lu} |dv/ds|}, \quad (4.36)$$

where $|dv/ds|$ is the velocity gradient through the Sobolev surface. The level populations n_u and n_l are calculated using Equation 4.29.

The procedure for computing the path of a photon through a cell is therefore one of calculating the distance to the next Sobolev surface at which the photon may undergo line scattering. The scattering optical depth due to electron scattering along this path is calculated. If the total scattering optical depth exceeds \mathcal{T} before the packet reaches

the next Sobolev surface, than the packet scatters via electron scattering. Otherwise, the scattering optical depth associated with the line is calculated and added to the total scattering optical depth. If this combined electron and line scattering opacity exceeds τ_{scat} , then the photon scatters from that line. If not the process is repeated until the photon does scatter by a different line, via electron scattering, or is absorbed completely by the wind due to its weight being reduced to zero.

Application of the Sobolev approximation requires that the macroscopic velocity gradient ($|dv/ds|$) along the photon's flight path is much greater than the thermal velocity of the gas divided by the typical length scale over which interactions occur (v_{th}/l). If this requirement is met the Sobolev surfaces can be correctly assumed to be physically distinct. A second requirement is that the physical conditions of the gas (temperature, ionization state etc.) do not change on the scale of the Sobolev length, l_s , defined by

$$l_s = \frac{v_{th}}{dv/ds}. \quad (4.37)$$

Physically, the Sobolev length represents the depth of a Sobolev surface. This interpretation comes about because the thermal broadening of the line profile in velocity space translates into a distance in this formalism. In our case, these requirements are usually well satisfied.

4.4.3 Continuum Opacity

Once the location where a photon packet scatters has been calculated, the continuum opacity along the last path segment is used to take account of attenuation through reduction in the weight of the photon packet. The energy lost contributes to the heating of the cell. Bound-free and free-free processes are both treated in PYTHON as sources of continuum opacity; the details of how they are implemented in PYTHON is discussed in the next two sections.

4.4.3.1 Bound-free opacity

The bound-free absorption coefficient is given by

$$\kappa_{\nu}^{bf} = n_{i,l} \sigma_{\nu(i,l)}^{bf}, \quad (4.38)$$

where $n_{i,l}$ is the number of density of level l of ion i , and $\sigma_{\nu(i,l)}^{bf}$ is the photoionization cross section of that ion/level. The total bound-free opacity to a photon packet of frequency ν is therefore simply a summation over all ions in the wind cell. Excited levels are included in the calculation, and the requisite level populations are again calculated using Equation 4.15, as for the free-bound calculation.

4.4.3.2 Free-free opacity

The free-free absorption coefficient is calculated from the free-free emission given in Equation 4.11 via Kirchoff's law,

$$\frac{j_{\nu}^{ff}}{\kappa_{\nu}^{ff}} = B_{\nu}(T). \quad (4.39)$$

This states that material at temperature T in (local) thermodynamic equilibrium with its surroundings must re-radiate absorbed energy with a frequency distribution which maintains a blackbody radiation field. We can therefore write

$$\kappa_{\nu}^{ff} = n_e n_i \frac{4}{3} \left(\frac{2\pi}{3} \right)^{1/2} \frac{Z^2 e^6}{m^2 h c} \left(\frac{m}{k T_e} \right)^{1/2} g^{ff} \nu^{-3} \left(1 - e^{(-h\nu/k T_e)} \right). \quad (4.40)$$

4.4.4 Attenuation and scattering

The photon packet passes through the cell until the total optical depth to scattering along its path is equal to the random optical depth, \mathcal{T} , generated at the start of this path segment. The total absorption optical depth along the path segment is then simply the sum of both continuum absorption coefficients multiplied by the path length. This

summation yields the *absorption* optical depth,

$$\tau_{abs} = \kappa_{\nu}^{ff} ds + \kappa_{\nu}^{bf} ds. \quad (4.41)$$

The photon weight just before it scatters, w_{abs} , is given by

$$w_{abs} = w_{in} e^{-\tau_{abs}}. \quad (4.42)$$

If the packet's path ends on a line resonance, then the new direction of the photon can be chosen from an isotropic distribution or via the escape probability formalism given in Equation 4.25. As discussed earlier, line scattering is not a purely elastic process, since there is a chance that the excited ion will de-excite via a collisional process rather than a radiative one. This means that the photon packet should lose a proportion of its energy during the scattering process. We have already calculated the proportion of photons that are absorbed (as opposed to scattered) in a line interaction. This is given by q (Equation 4.20) or Q (Equation 4.26), if we want to take account of line trapping. Thus the energy lost from the photon packet is given by Qw_{abs} , and so the new photon packet weight w_{out} is given by

$$w_{out} = (1 - Q) w_{abs}. \quad (4.43)$$

The alternative to line scattering is electron scattering. In this case, the photon packet is scattered in a direction chosen from an isotropic distribution, and its weight is not reduced. A new random optical depth, \mathcal{T} , is now calculated, and the next path length calculated. This procedure is repeated until the photon leaves the wind.

4.4.5 Radiation estimators

Several parts of the method, most notably the ionization scheme, require a model of the specific intensity of radiation in the cell. As mentioned earlier, in the original version of PYTHON this is assumed to take the form of a dilute blackbody, defined by a dilution factor, W , and radiation temperature, T_r . In order to calculate these parameters, two

“radiation estimators” are logged during the photon transport phase of the simulation. These estimators are the frequency-averaged mean intensity, J , and the mean frequency of photon packets, $\bar{\nu}$:

$$J = \frac{1}{4\pi V} \sum_{paths} w_{ave} ds, \quad (4.44)$$

$$\bar{\nu} = \frac{\sum_{paths} w_{ave} \nu ds}{\sum_{paths} w_{ave} ds}, \quad (4.45)$$

where w_{ave} is the average packet weight along each path segment,

$$w_{ave} = (w_{abs} - w_{in}) / \tau_{abs}. \quad (4.46)$$

W and T_r are then calculated from these radiation estimators by comparison with the values expected from a blackbody ([Mazzali and Lucy 1993](#)):

$$T_r = \frac{h\bar{\nu}}{3.832 \times k} \quad (4.47)$$

$$W = \frac{\pi J}{\sigma T_r^4}. \quad (4.48)$$

4.5 Heating and cooling balance

Once a population of photon packets has passed through the wind, PYTHON calculates an updated estimate of the cell electron temperature. This is done by assuming that cells should be in thermal equilibrium, with the heating rate equal to the cooling rate. The heating rate is calculated from the weight loss of photon packets passing through the cell. It is therefore a fixed quantity for a given ionization cycle. The cooling rate is expressed as a function of temperature, and so the new temperature for the cell is calculated by finding the temperature, T , where

$$\sum_{processes} H = \sum_{processes} C(T). \quad (4.49)$$

This approach would directly yield an answer for the cell temperature if the heating rate was temperature independent, and the cooling rate was simply a function of temperature. However, both the heating and cooling rates are complex functions of not only the temperature, but also the ionization state of the gas. This dependency means that an iterative process is required, where the temperature is updated after each ionization cycle and used to calculate an updated ionization state, giving a different heating rate in the next cycle.

In order to guard against instabilities, the temperature of a cell is only allowed to change by a fixed fraction of its current value after each ionization cycle. As the temperature converges, and the cooling rate approaches the heating rate, this fraction is reduced. A simulation is considered converged once most cells have reached a fixed temperature.

The details of the heating and cooling mechanisms are discussed in the next two sections.

4.5.1 Heating processes

Heating in a cell occurs when a photon packet loses weight due to one of the processes described in the previous section. In this section, a few of the details associated with calculating the heating rates are discussed.

Electron scattering events are treated as elastic, so the only continuum opacities that remove energy from a photon packet are free-free and bound-free. All of the energy a photon packet loses to free-free opacity is transferred directly to the electron pool, so the total free-free heating rate is given by

$$H^{ff} = w_{in} (1 - e^{-\tau_{abs}}) \frac{\tau^{ff}}{\tau_{abs}}. \quad (4.50)$$

In contrast, if a photon is absorbed via bound-free opacity, a proportion of the photon's energy is lost in overcoming the binding energy of the electron. The balance ends up as the kinetic energy of the liberated electron, i.e. it represents the heating of the electron pool. Therefore the heating rate due to bound-free absorption is given by

$$H^{bf} = w_{in} (1 - e^{-\tau_{abs}}) \frac{\tau^{bf}}{\tau_{abs}} \left(1 - \frac{\nu_t}{\nu}\right), \quad (4.51)$$

where ν_t is the ionization threshold frequency.

The photon packet can also lose energy via collisional de-excitation in line interactions.

The total heating rate due to this process is given by

$$H^{lines} = \sum_{lines} w_{abs} - w_{out} \quad (4.52)$$

$$= \sum_{lines} w_{abs} - (Q)w_{abs} \quad (4.53)$$

$$= \sum_{lines} (1 - Q)w_{abs}, \quad (4.54)$$

where w_{abs} is the weight of the photon packet following true absorption, prior to the scattering event.

4.5.2 Cooling processes

The wind is assumed to be in thermal equilibrium, and PYTHON searches for a temperature which balances heating and cooling processes after each cycle of photon packet propagation. Three of the cooling processes, free-free, free-bound and lines have already been described in the context of photon emission from the wind. The total cooling effect due to each of these three processes is calculated by integrating the frequency-dependent luminosities associated with each process for a range of temperatures, thereby constructing temperature-dependent cooling rates. For example, the temperature-dependent cooling rate due to the free-free process is calculated as

$$C^{ff}(T) = \int_0^{\infty} L^{ff}(\nu, T) d\nu. \quad (4.55)$$

In an outflowing wind with diverging streamlines, the density of the gas drops as it moves outwards. Since, in the sort of wind we typically model, the bulk velocity of the wind is much higher than the thermal velocity of the gas within it, this will cause adiabatic (expansion) cooling. This mechanism is included in the code as a non-radiative cooling mechanism and is given by

$$C^{ad} = P \frac{dV}{dt}, \quad (4.56)$$

where P is the pressure and dV/dt is the rate of change in the volume of a co-moving parcel of gas in a given cell. The pressure is assumed to be dominated by free electrons, so, for a temperature T , we have $P = n_e kT$. In our model, each cell has a fixed volume, but the effective expansion rate can be written as $V(\nabla \cdot v)$, where the divergence of the velocity, $\nabla \cdot v$, of the cell is calculated by summing the partial velocity gradients along the three axes at the centre of the cell. Thus we have

$$C^{ad}(T) = n_e k_B T V(\nabla \cdot v). \quad (4.57)$$

4.6 Ionization equilibrium calculation (dilute blackbody approximation)

The original version of PYTHON used a highly simplified but computationally efficient approach to calculating the ionization state of a wind cell. This was based on the modified Saha approach first suggested by [Mazzali and Lucy \(1993\)](#) and can be summarised by

$$\frac{n_{i+1} n_e}{n_i} = W [\zeta + W (1 - \zeta)] \left(\frac{T_e}{T_r} \right)^{1/2} \left(\frac{n_{i+1} n_e}{n_i} \right)_{T_r}^*. \quad (4.58)$$

This equation is derived in section [5.3.1](#), where the derivation is used to introduce the more flexible scheme now implemented in PYTHON. The left hand side of the equation is the ratio of the number density of two adjacent ionization stages of an element, multiplied by the electron number density. The starred term on the right is the same ratio, calculated in local thermodynamic equilibrium (LTE) via the Saha equation. W and T_r are the parameters of the dilute blackbody used to model the mean specific intensity in the cell (see Equation [4.6](#)). Finally, ζ is the ratio of recombinations going directly to the ground state relative to the number of recombinations going to all states.

Since this equation only links pairs of ionization levels, the calculation of the complete ionization state proceeds via an iterative process. First, an initial estimate of n_e is made. The number density of each neutral ion is set to 1, and then the population of each subsequent state is calculated relative to this. The number densities of ions of different elements are then rescaled to take account of the relative abundances, and then

the electron density n_e is recalculated. This new value of n_e is then used for a second calculation, and the process continues until n_e converges.

At this point, we have a new estimate of T_e and a new ionization state for the wind. This is the end of an ionization cycle. If the wind has converged in T_e , then the ionization and temperature structure of the wind is frozen, and synthetic spectra may be calculated. If more ionization cycles are needed to produce convergence, then a new population of photon packets are generated and another ionization cycle begins.

4.7 Synthetic spectra

In PYTHON, synthetic spectra can be generated at a series of user defined sightlines. The calculation proceeds via a sequence of spectral cycles, which are similar to ionization cycles, but with two main differences. Firstly, the frequency range over which photons are produced is normally much reduced. In the ionization cycles, it is important to produce enough photons to ensure that all applicable heating and cooling mechanisms for the wind can be properly modelled. It is also important that the range of frequencies is large enough to allow a good estimate of the mean intensity in the cell to be calculated. As we will see later, this is even more critical in the case of the modified version of PYTHON presented in the next chapter. However, since the thermal and ionization state of the wind is fixed for the spectral cycles, it is not necessary to carry out these calculations.

Secondly, whenever a photon packet scatters in the wind, the optical depth from the scattering location to the edge of the wind along all requested sightlines is calculated in addition to the normal scattering calculation. Normally, the photon packet would scatter in a random direction. However, in this case, the photon is split into as many parts as there are sightlines, each weighted according to the chance it would have gone in that direction (thus preserving the total luminosity). The remains of the original packet proceed as normal, but each of the “extracted” packets is added to the spectrum in the given sightline, attenuated by the optical depth in that direction.

The final spectrum is calculated by summing the weights of the extracted photon packets for each sightline in a set of frequency bins. If the extraction technique is not used, then a

spectrum at a given sightline is made by summing photon packets leaving the simulation in a small range of angles around that sightline. Clearly in this case many, if not most, of the photon packets will not happen to leave at the right angles, and so will be wasted. This illustrates the efficiency of the extraction algorithm.

The specific luminosity in a given frequency bin is the sum of all photon packet weights, divided by the width of the bin,

$$L_\nu = \sum_{\nu-\delta\nu/2}^{\nu+\delta\nu/2} w/\delta\nu. \quad (4.59)$$

The observed monochromatic flux for any given observer distance d is then simply calculated by dividing by the surface area of a sphere,

$$F_\nu = \frac{L_\nu}{4\pi d^2}. \quad (4.60)$$

Reducing the frequency range over which photon packets are produced in the spectral cycle increases the number of photons in each frequency bin, and so improves the signal-to-noise ratio of the resultant spectra. The extraction process does the same thing, since every photon contributes at least once to each spectrum.

Chapter 5

Radiative transfer and ionization calculations in AGN

5.1 Introduction

Chapter 4 introduced the original version of PYTHON, our radiative transfer and ionization code. It was written to enable the modelling of disk winds in CVs, and the physics included in the code reflect this. This project, however, seeks to understand disk winds in AGN, where different physical processes can become important. In this chapter, we will discuss the additional physics we have sought to model in PYTHON to allow meaningful simulations of AGN, and present the results of some test calculations.

5.2 X-ray source

Figure 5.1 shows a collection of quasar spectral energy distributions. This shows that a type 1 quasar spectrum consists of several parts. The optical/UV part of the spectrum is generally believed to arise from an optically thick accretion disk and is therefore already modelled reasonably appropriately in PYTHON. However, the X-ray part of the spectrum arises from a different source, and given the lack of consensus as to the physical origin of this component, we do not try and model it in a self-consistent manner. Instead, we




Figure 5.1: Collection of type 1 quasar spectral energy distributions (Richards et al. 2006).

include it as a separate radiation source, with an SED parameterised by a photon power law index, α_x , and the integrated luminosity from 2-10 keV.

Despite the fact that the detailed geometry of the X-ray source is not known, it is still important to choose a physically motivated location for it. A popular explanation for the production of AGN X-rays is Compton up-scattering of disk photons by a hot corona. This picture suggests that the X-ray source should be close to the location of high frequency seed photons, i.e. the inner part of the accretion disk. This is backed up by recent reverberation measurements, which show that the X-ray source is located only about $4R_G$ above the accretion disk (Emmanoulopoulos et al. 2014). Other possibilities

are the base of a radio jet, or perhaps a spherically symmetrical accretion flow at the centre of the accretion disk ([Meyer et al. 2000](#)).

Based upon these considerations, we choose a simple source geometry - a centrally located spherical source with radius set equal to the innermost radius of the accretion disk. It radiates isotropically, with no limb-darkening.

5.3 Ionization balance

Disk winds are generally thought to be photoionized, and therefore accurate calculation of the ionization state of the wind requires accurate modelling of the local mean intensity. The ionization calculation in the original version of PYTHON makes the assumption that the SED is well approximated by a dilute blackbody. However, with the addition of the X-ray source described above, this will no longer be the case. Additionally, absorption in the wind can strongly modify the SED seen in a cell, and the original scheme could only treat this very approximately. Following [Sim et al. \(2008\)](#), we have therefore extended the original scheme to work for an arbitrary SED. It is instructive to examine the physical basis of the original scheme before describing our upgraded method.

5.3.1 Derivation of the approximate ionization scheme

The approximate ionization scheme used in PYTHON originally is based upon calculating LTE abundances using the Saha equation and then applying correction factors to take account of the true photoionization and recombination rates. It has already been summarised in Chapter 4 (Equation 4.58).

The Saha equation links the relative number of ions in two successive stages of ionization and is derived by noting that, in LTE, collisional processes dominate, and so free electrons and ions are in statistical equilibrium. The ratio of the number densities in any two states is then simply set by statistical and quantum mechanical arguments. A version for two

successive ions, both in the ground state, is

$$\frac{n_{i+1,0}^* n_e^*}{n_{i,0}^*} = \frac{g_e g_{i+1,0}}{g_{i,0}} \left(\frac{2\pi m_e k T}{h^3} \right)^{3/2} e^{-\frac{\chi_{i,0,0}}{kT}}. \quad (5.1)$$

In this equation, $n_{i,0}^*$ and $n_{i+1,0}^*$ are the number densities of ions in ionization state i and $i + 1$, whilst n_e^* is the number density of electrons. The star in the superscript reminds us that these densities are calculated in LTE, and the zero in the subscript shows that the equation refers only to the ground state of the ions in question. In the following derivation, we replace the ratio on the left hand side of the equation with the symbol $\Phi_{i,0,0}^*(T)$. The first term on the right hand side is the ratio of the ground state multiplicities of the two ions, and the electron multiplicity. T is the temperature, and $\chi_{i,0,0}$ is the ionization potential.

In a typical nebula it is unlikely that LTE will hold, because the densities are too low for collisional processes to dominate. However, if the overall ionization state is in equilibrium, the number of ionizations from ion i per unit time must equal the number of recombinations back into i . Thus we can write

$$n_i \sum \gamma_{i \rightarrow i+n} = n_{i+1} n_e \sum \alpha_{i+1 \rightarrow i}, \quad (5.2)$$

where the left hand side is a summation over all ionizing processes with rate coefficients γ , and the right hand side is a summation over all recombination processes with rate coefficients α . Note that ionizations do not necessarily only increase the charge of an ion by one, hence the $i + n$ in the subscript of the ionization rate coefficient, where $n = 1$ or more. In LTE, we can apply the principle of “detailed balance” - that is the rate of each process is the same as that of its inverse. Choosing photoionization and radiative recombination, we can therefore write

$$n_{i,0}^* \gamma_{i,0 \rightarrow i+1,0}^{p*} = n_{i+1,0}^* n_e^* \alpha_{i+1,0 \rightarrow i,0}^{r*} \quad (5.3)$$

where $\gamma_{i,0 \rightarrow i+1,0}^{p*}$ is the photoionisation rate coefficient, and $\alpha_{i+1,0 \rightarrow i,0}^{r*}$ is the radiative recombination rate coefficient, both linking the ground states of successive ionization stages. These rates are starred to remind us that they refer to LTE. The recombination

rate coefficient is for a Boltzmann distribution of electron velocities at temperature T , and the photoionization rate coefficient is for a Planck distribution of photon energies at the same temperature. Rearranging we see that

$$\frac{n_{i+1,0}^* n_e^*}{n_{i,0}^*} = \frac{\gamma_{i,0 \rightarrow i+1,0}^{p*}}{\alpha_{i+1,0 \rightarrow i,0}^{r*}} = \Phi_{0,0,i}^*(T). \quad (5.4)$$

Even if we are *not* in LTE, we can make the same simplification to Equation 5.2 as long as photoionization and radiative recombination are the *dominant* processes. This is reasonable in a low density radiation dominated plasma. We can therefore write

$$\frac{n_{i+1,0} n_e}{n_{i,0}} = \frac{\gamma_{i,0 \rightarrow i+1,0}^p}{\alpha_{i+1,0 \rightarrow i,0}^r}, \quad (5.5)$$

where we now have the densities we are interested in in terms of the actual photoionization and recombination rates. Multiplying top and bottom of Equation 5.5 by the LTE rate coefficients, we can write

$$\begin{aligned} \frac{n_{i+1,0} n_e}{n_{i,0}} &= \frac{\gamma_{i \rightarrow i+1}^p}{\alpha_{i+1 \rightarrow i}^r} \frac{\gamma_{i,0 \rightarrow i+1,0}^{p*}}{\alpha_{i+1,0 \rightarrow i,0}^{r*}} \frac{\alpha_{i+1,0 \rightarrow i,0}^{r*}}{\gamma_{i,0 \rightarrow i+1,0}^{p*}} \\ &= \Phi_{0,0,i}^*(T) \frac{\gamma_{i \rightarrow i+1}^p}{\gamma_{i,0 \rightarrow i+1,0}^{p*}} \frac{\alpha_{i+1,0 \rightarrow i,0}^{r*}}{\alpha_{i+1,0 \rightarrow i,0}^r}. \end{aligned} \quad (5.6)$$

The first term on the right hand side of Equation 5.6 is the LTE abundance ratio from the Saha equation, the second term is the actual photoionization rate coefficient divided by the LTE rate coefficient, and the last term is the LTE radiative recombination rate coefficient divided by the actual coefficient. These last two are our “correction” factors. Let us deal with them in turn.

To work out the photoionization correction factor, we start by writing an equation for the photoionization rate γ^p . It is simply the number of photons of sufficient energy to cause an ionisation multiplied by the ionisation cross section. The energy flux E_ν in the frequency range $d\nu$ is

$$E_\nu = 4\pi J_\nu d\nu \quad (5.7)$$

where J_ν is the mean specific intensity. Therefore, the energy absorbed by a given ion

is the flux multiplied by the cross section of the ion $\sigma_{i,0}^p(\nu)$. Therefore E_ν^{abs} , the energy absorbed by the ion per unit time is

$$E_\nu^{abs} = 4\pi J_\nu \sigma_{i,0}^p(\nu) d\nu. \quad (5.8)$$

The photoionization rate coefficient is the number of ionizations per unit time, and so we must divide by the energy of a photon, and integrate over all ionizing photons to obtain an expression for $\gamma_{i,0 \rightarrow i+1,0}^p$, viz

$$\gamma_{i,0 \rightarrow i+1,0}^p = 4\pi \int_{\nu_0}^{\infty} \frac{J_\nu \sigma_{i,0}^p(\nu)}{h\nu} d\nu. \quad (5.9)$$

In LTE, the mean specific intensity is, by definition, a blackbody, so $J_\nu = B_\nu(T)$, yielding

$$\gamma_{i,0 \rightarrow i+1,0}^{p*} = 4\pi \int_{\nu_0}^{\infty} \frac{B_\nu(T) \sigma_{i,0}^p(\nu)}{h\nu} d\nu. \quad (5.10)$$

Thus, the photoionization rate correction factor, $S_{i,0,0}$, is

$$S_{i,0,0} = \frac{\gamma_{i,0 \rightarrow i+1,0}^p}{\gamma_{i \rightarrow i+1,0}^{p*}} = \frac{\int_{\nu_0}^{\infty} J_\nu \sigma_{i,0}^p(\nu) \nu^{-1} d\nu}{\int_{\nu_0}^{\infty} B_\nu(T) \sigma_{i,0}^p(\nu) \nu^{-1} d\nu} \quad (5.11)$$

The recombination rate correction factor is more simple. To first order the recombination rate coefficient is inversely proportional to the square root of the temperature of electrons. Furthermore, even if we are not in LTE, it is a reasonable approximation to say that the electrons still follow a Boltzmann distribution. This is because the typical timescale of electron-electron interactions is much shorter than the typical timescale of photoionizations. Thus any electrons produced by photoionization will very quickly interact with the electron pool and reach equilibrium. This means that the only difference between the LTE recombination rate and the true recombination rate will be due to differences in electron temperature. If electrons in our plasma have temperature T_e , and T is the temperature at which we calculate the LTE abundance ratio, the recombination rate correction factor is

$$\frac{\alpha_{i+1,0 \rightarrow i,0}^{r*}}{\alpha_{i+1,0 \rightarrow i,0}^r} = \sqrt{\frac{T_e}{T}} \quad (5.12)$$

Replacing the correction factors in Equation 5.6 with equations 5.11 and 5.12, we obtain the fundamental equation for this approximate ionization method

$$\frac{n_{i+1,0}n_e}{n_{i,0}} = \Phi_{0,0,i}^*(T) S_{i,0,0} \sqrt{\frac{T_e}{T}}. \quad (5.13)$$

In the simple case where the mean intensity can be approximated by a dilute blackbody with radiation temperature T_r , we can write $J_\nu = WB_\nu(T_r)$. If we calculate the Saha abundances at temperature T_r , the approximate ionization equation simplifies to

$$\begin{aligned} \frac{n_{i+1,0}n_e}{n_{i,0}} &= \Phi_{0,0,i}^*(T_r) \frac{\int_{\nu_0}^{\infty} WB_\nu(T_r) \sigma_{i,0}^p(\nu) \nu^{-1} d\nu}{\int_{\nu_0}^{\infty} B_\nu(T_r) \sigma_{i,0}^p(\nu) \nu^{-1} d\nu} \sqrt{\frac{T_e}{T_r}} \\ &= W \sqrt{\frac{T_e}{T_r}} \Phi_{0,0,i}^*(T_r), \end{aligned} \quad (5.14)$$

which is the basic form of Equation 4.58, the scheme that was implemented in the version of PYTHON prior to this project. The extra factor of $[\zeta + W(1 + \zeta)]$ is a further correction which takes account of recombinations to levels other than the ground state and photoionizations from excited levels. This factor is discussed in the next section.

5.3.2 Correction factor for recombination to excited levels

In the derivation above, we have assumed that all photoionizations are from the ground state, and all recombinations are to the ground state. The first assumption is very commonly used in RT calculations of low density gas clouds and is known as the “nebular approximation”. The argument is that the typical timescale for radiative de-excitation of an ion in an excited state to the ground state is much less than the typical time between photoionization events.

The second assumption is not as sound, however, since excited states are available to all electrons during recombination, and so the cross section of the ion to electrons is larger than one would expect from the ground state alone. Fortunately it is easy to come up with a correction factor. We start by rewriting Equation 5.2 to incorporate

recombinations to all excited levels l of ionization state i

$$\frac{n_{i+1,0}n_e}{n_{i,0}} = \frac{\gamma_{i,0 \rightarrow i+1,0}^p}{\sum_l \alpha_{i+1,0 \rightarrow i,l}^r}. \quad (5.15)$$

Multiplying top and bottom by the recombination rate to the ground-state only and rearranging gives

$$\frac{n_{i+1,0}n_e}{n_{i,0}} = \frac{\gamma_{i,0 \rightarrow i+1,0}^p}{\alpha_{i+1,0 \rightarrow i,0}^r} \frac{\alpha_{i+1,0 \rightarrow i,0}^r}{\sum_l \alpha_{i+1,0 \rightarrow i,l}^r}. \quad (5.16)$$

The second term is now a correction factor which is the ratio of recombinations going to the ground-state only to all recombinations. We represent this correction factor with the symbol $\zeta(T)$. The rest of the derivation of the approximate ionization equation proceeds as in the previous section, apart from a factor of $\zeta(T_e)$, where ζ is calculated at the electron temperature T_e , a suitable choice for recombination processes. Thus we obtain

$$\frac{n_{i+1,0}n_e}{n_{i,0}} = W\zeta(T) \sqrt{\frac{T_e}{T_r}} \Phi_{0,0,i}^*(T_r). \quad (5.17)$$

This represents the first stage of modifications to Equation 5.14. However, following [Mazzali and Lucy \(1993\)](#), PYTHON includes further correction factors which take account of photoionizations from excited levels. However, these arguments are based upon the assumption that the mean intensity is well modelled by a dilute blackbody. Since this is not the case in our AGN models, we do not use these extra correction factors in the new ionization scheme described below. Therefore the derivations of these extra factors are not included here.

The recombination rate coefficients used to calculate $\zeta(T)$ are generated in one of two ways. Ideally, we use tabulated values of the total recombination rate and the ground state resolved rate from the CHIANTI database ([Dere et al. 1997](#)). If these are not available, we calculate the rate coefficients from the photoionization rate coefficients via the Milne relation. Figure 5.2 shows the values of $\zeta(T)$ calculated for the ionic states of carbon.




Figure 5.2: The ground state recombination correction factors for Carbon.

5.3.3 Approximate ionization calculations for an arbitrary SED

The approximate ionization scheme derived above is implemented in PYTHON in the dilute blackbody form, which is computationally very fast and accurate in the limit where the sources of radiation in the model are thermal. For the AGN simulations proposed in this project, there are strongly non-thermal sources of ionizing radiation, and early testing revealed that this particular implementation did not give an accurate representation of the ionization state in this case.

Fortunately, there is a simple way to increase the flexibility of the scheme to take account of any SED. Instead of assuming $J_\nu = WB_\nu$ in Equation 5.11, the photoionization rate correction factor can be evaluated by direct numerical integration of the SEDs in the numerator and denominator. The denominator is simply an integral of the photoionization cross-section multiplied by a Planck function with the same temperature as that used to calculate the LTE abundances from the Saha equation. The numerator is an integral of the cross section multiplied by the mean intensity actually seen in the cell. This calculation therefore requires a model of the mean intensity. This modelling is the most complex element in the new ionization scheme and is described next.

5.3.3.1 Modelling the mean intensity

In the original version of PYTHON, the mean intensity is assumed to be a dilute blackbody, but the new implementation discussed above requires a more flexible model of the mean intensity in a cell as a function of frequency. We first split the full frequency range up into a series of bands. In each band, j , the mean intensity can be represented by a power law of the form

$$J_{\nu,j} = K_{pl,j} \nu^{\alpha_{pl,j}}, \quad (5.18)$$

or an exponential

$$J_{\nu,j} = K_{exp,j} e^{(-h\nu/kT_{exp,j})}. \quad (5.19)$$

The parameters defining the two models are the two normalisation factors, $K_{pl,j}$ and $K_{exp,j}$, the power law spectral index, $\alpha_{pl,j}$, and the exponential temperature, $T_{exp,j}$. The exponential representation is parameterised in a way that is meant to ensure that the temperature has physical meaning, i.e. in the case of a thermal source of photons it can be compared to the known radiation temperature to check for consistency.

The SED parameters are calculated using three estimators of the radiation field in the cell, which are logged during the propagation of photon packets. These three estimators are: the mean frequency of photon packets, $\bar{\nu}_j$, the frequency integrated mean intensity, J_j , and the standard deviation of the frequency, $\sigma_{\nu,j}$. These estimators are collected in each of the bands used to model the spectrum, each covering a frequency band $\nu_{min} \rightarrow \nu_{max}$. J_j and $\bar{\nu}_j$ are just “band limited” versions of the estimators collected in the original version of PYTHON, but all three are given below for completeness:

$$J_j = \frac{1}{4\pi V} \sum w_{ave} ds, \quad (5.20)$$

$$\bar{\nu}_j = \frac{\sum \nu w_{ave} ds}{\sum w_{ave} ds}, \quad (5.21)$$

$$\sigma_{\nu,j} = \sqrt{\frac{\sum \nu^2 w_{ave} ds}{\sum w_{ave} ds} - \bar{\nu}^2}, \quad (5.22)$$

where w_{ave} is the mean weight of a photon packet with frequency between the band limits, over a path segment of length ds . In the rest of this section, the subscript j is dropped for clarity, but all calculations are assumed to be carried out using band-limited values. Similarly, integrals are always over the range of frequencies in the band.

Once the estimators have been calculated following a flight of photons, they are used to compute values for the spectral model parameters. We first calculate the mean frequency from the mean intensity

$$\bar{\nu}_{mod} = \frac{\int \nu J_\nu d\nu}{\int J_\nu d\nu}. \quad (5.23)$$

We can then substitute the two different spectral model expressions for J_ν , which allows Equation 5.23 to be solved for two estimates of the mean frequency. For the power law representation we obtain

$$\bar{\nu}_{mod} = \frac{(\alpha_{pl} + 1)(\nu_{max}^{\alpha_{pl}+2} - \nu_{min}^{\alpha_{pl}+2})}{(\alpha_{pl} + 2)(\nu_{max}^{\alpha_{pl}+1} - \nu_{min}^{\alpha_{pl}+1})}, \quad (5.24)$$

and for the exponential, setting $x = -h/kT_{exp}$, we find

$$\bar{\nu}_{mod} = \frac{1}{x} \frac{(x\nu_{max} - 1)e^{x\nu_{max}} - (x\nu_{min} - 1)e^{x\nu_{min}}}{e^{x\nu_{max}} - e^{x\nu_{min}}}. \quad (5.25)$$

We then find a value for α_{pl} and T_{exp} by numerically solving the equation $\bar{\nu} - \bar{\nu}_{mod} = 0$. It is then straightforward to use the measured frequency-integrated mean intensity to calculate the normalisation factors K_{pl} and K_{exp} via

$$J = \int_{\nu_{min}}^{\nu_{max}} J_\nu d\nu. \quad (5.26)$$

At this point, we have two potential models for the mean intensity in the cell. In order to choose between them, we calculate the standard deviation of the frequency for each model,

$$\sigma_{\nu,mod} = \sqrt{\frac{\int \nu^2 J_\nu d\nu}{\int J_\nu d\nu} - \bar{\nu}_{mod}^2} \quad (5.27)$$

and compare this to the estimator. We choose the model which yields the best match. Once we have a series of models for J_ν , the numerator in the photoionization rate correction factor (Equation 5.11) can be evaluated as a sum of integrals over each band

as

$$S_{i,0,0} = \frac{\sum_j \int_{\nu_{min}}^{\nu_{max}} J_{\nu,j} \sigma_{i,0}^p(\nu) \nu^{-1} d\nu}{\int_{\nu_0}^{\infty} B_{\nu}(T) \sigma_{i,0}^p(\nu) \nu^{-1} d\nu}. \quad (5.28)$$

Figure 5.3 shows J_{ν} for three cells in a wind illuminated by an accretion disk plus an X-ray source with a power law SED. The model is described fully in Chapter 6, but here we focus on just three cells to illustrate the spectral models.




Figure 5.3: Three examples of modelled cell spectra.

The top panel is for a cell exposed directly to the illuminating SED, and we see that the model reproduces the mean intensity very well. The step between the end of the exponential model chosen for the Wien tail of the disk and the power law model at higher frequencies highlights the fact that we do not enforce continuity across band boundaries.

In the middle panel, the cell is in the bulk of the wind, and bound-free absorption has caused a reduction in the mean intensity between 3×10^{16} and 10^{18} Hz. Once more, the model follows J_ν very well. The cell represented by the final panel is deep within the wind, and the bound-free absorption is consequently much greater. No photon packets at all have survived between 3×10^{16} and 3×10^{18} Hz, and it is situations like this which potentially cause the biggest problems when trying to model J_ν . Our approach is to change ν_{min} and ν_{max} for any band which does not see photons across the whole band to reflect the range of frequencies that it *does* see. We then fit a model to this reduced band, and only carry out the integrals for photoionization rate coefficients over that reduced frequency range. The model fitted to the SED in the bottom panel is truncated at the edges of the absorption feature and fits the SED fairly well. Clearly, this scheme has limitations. For example, if there are multiple deep absorption features in a given band, they will be hard to model. However, bound-free is the main absorption process in the winds we model and this tends to give very broad features like those in Figure 5.3.

5.3.3.2 The variable temperature scheme

The model of mean intensity allows us to calculate the numerator of the photoionization rate correction factor $S_{i,0,0}$ (Equation 5.11) more accurately than with the dilute blackbody formulation. However, we still need to calculate the denominator, which is the integral of a blackbody multiplied by the frequency-dependent photoionization cross section for the lower ion. The temperature of the blackbody is the same as that at which the Saha equation is evaluated, and this can lead to computational issues.

In order to overcome these issues, we note that there is actually nothing in the approximate ionization scheme which tells us at which temperature to calculate the Saha equation. Previous choices have been made for reasons of computational efficiency. In the original scheme, evaluating the Saha equation at the radiation temperature of the dilute blackbody removes the requirement to carry out an integral for the denominator - a significant computational saving. Therefore, in this case the choice is obvious.

However, in the new formulation, we always have to calculate the integral in the denominator, so we have a free hand in selecting the temperature at which to calculate the Saha

equation. One choice might be to evaluate it at T_e , since this would remove the need to correct for the recombination rate (the square root term in Equation 5.14). However, with a significant population of X-ray photons contributing to J_ν , we would expect non-zero populations of highly ionized ions. If the temperature at which the LTE abundances are calculated is too low, the photoionization threshold of the more highly ionized ions will be significantly past the peak of the blackbody spectrum at that temperature. Thus, the integral in the denominator can become very small - quickly approaching machine accuracy. At this point, the scheme becomes unstable.

Equation 5.4 suggests a good choice of temperature. It shows that the photoionization rate divided by the recombination rate (both calculated in LTE) is equal to $n_{i+1}n_e/n_i$. The integral in the denominator of equation 5.11 is also the LTE photoionization rate. Equation 5.4 shows that this must be non-zero if $n_{i+1}n_e/n_i$ is non-zero. We can therefore calculate a reasonable temperature for the Saha equation and hence the temperature of the blackbody in the denominator of Equation 5.11 by numerically solving Equation 5.1 to find the temperature where $n_{i+1}n_e/n_i = 1$.

5.3.4 Additional ionization and recombination processes

Whilst photoionization and radiative recombination are the dominant processes which set the ionization state of a diffuse, radiation dominated plasma such as is expected around an AGN, there are other processes that can have an influence. As it stands, the new ionization scheme described above can incorporate other processes which link adjacent ionization states of a given ion. Ionization processes would simply become additive terms in the numerator of the photoionization rate correction factor (Equation 5.11), whilst recombination processes would appear in the denominator of the ζ term which would become

$$\zeta(T) = \frac{\alpha_{i+1,0 \rightarrow i,0}^r}{\sum_l \alpha_{i+1,0 \rightarrow i,l}^r + \alpha^x + \alpha^y + \dots} \quad (5.29)$$

where the $\alpha^{x,y}$ terms represent the recombination rate coefficients for additional processes. The rest of this section describes some of the additional ionization and recombination processes that could be considered for inclusion in PYTHON, starting with dielectronic recombination (DR), which *has* already been included via this method.

5.3.4.1 Dielectronic recombination

In this process, a free electron is captured by an ion, but instead of releasing a photon immediately, excess energy carried by the electron goes into exciting one of the existing electrons. This differs from normal radiative recombination in that it is a resonant process. The cross-section of the ion to electrons is strongly peaked at energies corresponding to the energy difference between different electron orbitals in the ion. Once in the excited state, the ion can decay in one of two ways. First, it can undergo “auto-ionization”, where one of the electrons is re-emitted. Thus the ion ends up in the same state as before the event. Clearly this outcome will not affect the ionization state of the gas.

The second outcome is that the ion relaxes via the emission of a photon, and the ion remains in the less ionized state. This second process is of interest in the ionization balance. As a proof of concept, we have implemented this process into the ionization scheme as an extra additive term in the denominator of $\zeta(T)$ using DR rate coefficients from the CHIANTI database. Figure 5.4 shows how $\zeta(T)$ is modified by this additional process. The solid lines are the new values, with the original values also plotted as dashed lines. We see that the recombination rate for C VII is unchanged. This is simply because there are no electrons in the bare ion to take part in the process. DR is only important for C VI and C V above a temperature of about 10⁶ K. The recombination rate for the lithium-like C IV is most affected by DR.

5.3.4.2 Auger Ionization

Auger Ionization is the inverse process of DR, i.e. it corresponds to a high energy photon causing an inner-shell electron to be ejected from the ion. This leaves the ion




Figure 5.4: The ground state recombination correction factors for carbon, with dielectronic recombination included as an additional recombination process. The dashed lines are the original correction factors shown in Figure 5.2.

in an excited state, and there is a second process of relaxation where an outer electron de-excites into the gap left by the electron initially removed. This causes emission of an energetic photon, or possibly an auto ionization event whereby an outer electron is ejected leaving the ion two stages more ionized. This is a significantly more complex process to implement in our scheme, since it couples non-adjacent ionization states. We have therefore not implemented it as yet. In section 5.3.5, we will consider its importance in the ionization balance of carbon, and in section 8.2.1.3, we will discuss a strategy for possible future implementation.

5.3.4.3 Charge Exchange

In the process of charge exchange, an electron is passed between two interacting ions. Clearly, if the two ions belong to same species, this leaves the net ionization state unchanged. However, this process can also occur between different species. An important example is OI/II and HI/II. This is a particularly favourable interaction, because the ionization potential of both neutral species is around 13.6 eV. Again, this is a difficult

process to implement in the current scheme, since it “cross-connects” ions of different elements. Nevertheless, the scheme outlined in section 8.2.1.3 would also work for this process.

5.3.4.4 Three-body recombination

Three-body recombination is a recombination process where two electrons interact near an ion, leaving one to recombine and the other to carry away the excess energy. This process is only important in relatively high density plasmas, since the rate scales as n_e^2 . We have therefore not attempted to include it in the current project.

5.3.4.5 Supra-thermal electron ionization

In a plasma that is illuminated by a flux of very high energy photons, photoionization can release very high energy electrons. These will thermalise fairly quickly, however, during this process they can take part in collisional ionization of other ions in the plasma. This may well be an important process, particularly in locations where photoionization absorption heavily attenuates the far UV/soft X-ray flux. In order to properly include this process, we would need to include a model of the electron energy distribution in each cell. This is beyond the scope of the work presented here.

5.3.4.6 Unresolved transition array (UTA) ionization

UTA ionization takes place when one of the inner shell lines of iron with an open M shell (Fe I - Fe XVI) is excited by a soft X-ray photon. An $n=2$ electron is excited to the $n=3$ level and, upon relaxation, there is sufficient energy for autoionization to take place. These lines form a broad, unresolved feature in the soft X-ray spectra of AGN (Behar et al. 2001). Implementation of this mechanism requires a multi-level model of the iron atom, which is a task beyond this project.

5.3.5 Validation of ionization scheme

In order to test our new ionization scheme, several validation exercises were carried out in which we compared the ionization state of a simple spherical shell model calculated with PYTHON to that predicted by the widely used photoionization code CLOUDY (Ferland et al. 2013). The geometry used was a shell of depth 1 cm and radius 10^{11} cm. This ensures that the plasma is optically thin, thereby allowing testing of the ionization scheme in isolation. The hydrogen density was set to $n(H) = 10^7$ cm $^{-3}$. For the CLOUDY runs, the illuminating spectrum was set using the command

`Table power law spectral index -0.9 low=0.01, hi=1470,`

which produces a broken power law spectrum with spectral index $\alpha = -0.9$ between 0.01 Rydbergs and 1470 Rydbergs. It is characterised by $\alpha = 5/2$ for low frequencies and $\alpha = -2.0$ at high frequencies. This spectrum is not intended to replicate any astrophysical object, but is convenient to use because it can be defined identically in both codes. The overall luminosity is set in CLOUDY via the `ionization parameter` command. CLOUDY uses the U definition of ionization parameter (Equation 1.3). Thus the luminosity is set indirectly by requiring a given ionization parameter at the inner edge of the cloud. This spectrum is illustrated in Figure 5.5. The noise in some parts of the PYTHON spectrum is due to the limited number of photon packets produced at the low flux end of each photon generation band.

In PYTHON a new mode was written in order to precisely replicate both the spherical shell geometry and the broken power law SED. We use the photon packets passing through a cell to compute the ionization parameter via

$$U = \frac{1}{cVn(H)} \sum_{h\nu \geq 13.6eV} \frac{w_{ave}ds}{h\nu}, \quad (5.30)$$

where w_{ave} is the average photon packet weight along a path length ds in a cell of volume V . The summation gives a number of photons, which, when divided by the cell volume, gives a photon density. For these runs we fixed the plasma temperature in both codes at 10 000 K, i.e. we only test ionization, not thermal equilibrium.




Figure 5.5: The flux incident on the inner edge of the thin shell model used to compare the ionization state from PYTHON and CLOUDY. This is for ionization parameter $U = 1$.

Figure 5.6 shows the predicted relative abundances for ionic species of hydrogen, helium, carbon, nitrogen and oxygen. We see that there is excellent agreement for hydrogen and helium, but there are differences for the other three elements. It is instructive to investigate the reason for these differences, since this informs future work to extend the number of ionization and recombination processes implemented in PYTHON.

We will take carbon as an example, and begin by plotting the ratio of the number density of adjacent states, rather than simply the relative abundances. This directly allows us to understand which pair of states are to blame for any discrepancy. This plot is presented in Figure 5.7. It is immediately apparent that there is an offset between the two codes for the ratio of $n(\text{C II})/n(\text{C I})$, and this is down to a difference in the photoionization cross section data used in the two codes. CLOUDY uses cross sections taken from Verner et al. (1996a) (hereafter referred to as VFKY), whereas PYTHON uses cross sections taken from TOPBASE (Cunto et al. 1993) for H, He, C, N, O and Fe. The TOPBASE data includes data for inner shell photoionization edges and is easily extrapolated to very high frequencies, where even small cross sections can be important. However, there is a slight discrepancy between the two data sets for the exact threshold frequency. PYTHON also includes a limited subset of the VFKY data, and Figure 5.8 shows the ratio of number




Figure 5.6: Relative abundances of H, He, C, N and O ionization stages for PYTHON (solid lines) and CLOUDY (symbols) with a broken power law spectrum.




Figure 5.7: Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols).

densities of adjacent states for carbon using this data.




Figure 5.8: Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections.

The $C\text{ II}/C\text{ I}$ ratio is now in good agreement, but we see that PYTHON still underestimates the $C\text{ IV}/C\text{ III}$ ratio. This shows that the $C\text{ III}$ photoionization rate is underestimated. In order to investigate what mechanism is responsible for the remaining differences, we have turned off a few mechanisms in CLOUDY. We begin by turning off charge exchange. The

results of this simulation are shown in Figure 5.9. Note that the CLOUDY simulations do not converge for $\log(U) < -5$ with charge exchange disabled.




Figure 5.9: Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange disabled.

Clearly charge exchange was a significant process, leaving only the C V/C IV ratio in error. This could be caused by too much C V or too little C IV. Auger effects are particularly important for lithium-like ions like C IV, since this ion is the one which has the smallest difference between the ionization potential and the energy required to eject an inner shell electron. We therefore turn off Auger ionization in CLOUDY. The results for carbon are shown in Figure 5.10, and the hunch that Auger ionization was to blame is confirmed.

Figure 5.11 shows the ion fractions for H, He, C, N and O for VFKY photoionization cross sections for the PYTHON predictions, and with charge exchange and Auger ionization disabled for CLOUDY. We now see excellent agreement. Inclusion of these two ionization mechanisms has not proved possible in the timescale of the research program described here. Looking back at Figure 5.7, however, we see that for ionization parameters above around 10^{-3} , we obtain good agreement with CLOUDY even without them. We are always in this regime and we therefore note the important ionization mechanisms for




Figure 5.10: Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange and Auger ionization disabled.

future implementation, but our ionization calculations are acceptable for use in the type of radiation fields seen around AGN.

As a final test, we have switched off the dielectronic recombination effect in PYTHON to observe how important it is. Figure 5.12 shows the ratio of number densities of adjacent carbon ions. Comparing this with Figure 5.4, we see that the ratios between the upper three ion pairs are unchanged, which is as expected because dielectronic recombination is not important for these ions at $T = 10^4$ K. The largest effect is for C IV ion, and this is reflected in the significant change in the ratio of that ion to C III. The excess of C IV is due to the lower recombination rate when dielectronic recombination is neglected.

Finally, Figures 5.13, 5.14 and 5.15 show a comparison between CLOUDY and PYTHON for iron. In general, even for this very complex element, agreement is reasonable. In particular PYTHON agrees with CLOUDY for the value of U at which each species is dominant to within an order of magnitude in most cases. We do very well once we get to lithium-, helium- and hydrogen-like iron ions. Whilst iron is a very important element in the X-ray spectrum of AGN, it is much less critical to get abundances of iron ions exactly right for UV spectra. As long as the heating and cooling effects due to iron are about right, our treatment should be good enough for the current project.




Figure 5.11: Relative abundances of H, He, C, N and O ionization stages for PYTHON (solid lines) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY.




Figure 5.12: Ratio of number density of adjacent ionization states of carbon for PYTHON (solid lines) and CLOUDY (symbols). PYTHON calculation uses VFKY photoionization cross sections and CLOUDY calculation has charge exchange and Auger ionization disabled. Dielectronic recombination is turned off in PYTHON.

5.4 Heating and cooling mechanisms

In the original version of PYTHON, photons passing through a cell are able to lose energy through continuum absorption from bound-free and free-free processes. They can also lose energy when undergoing line scattering. This lost energy goes into heating the cell, which is then able to cool through free-free, free-bound and line emission, as well as adiabatic cooling. For the kind of SED seen in AGN, there are additional heating and cooling processes that set the plasma temperature. It is important to include them if we wish to obtain correct converged temperatures in such cases. This section describes the additional processes included in the new version of PYTHON.

5.4.1 Compton heating and cooling

At the start of this project, all scattering of photons by electrons in PYTHON was dealt with in the Thompson limit. This approximation is entirely suitable for situations where the typical photon energy is less than the rest energy of an electron, $h\nu \ll m_e c^2$, i.e. $\nu \ll 1.23 \times 10^{20}$ Hz. However, with our inclusion of X-rays into the illuminating spectra, we now reach photon energies up to 50 keV, representing a frequency of 1.2×10^{19} Hz,




Figure 5.13: Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe I-IX).




Figure 5.14: Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe x-XVIII).




Figure 5.15: Relative abundances of Fe ionization stages for PYTHON (solid line) and CLOUDY (symbols) with a broken power law spectrum. Charge exchange and Auger ionization disabled in CLOUDY (Fe IXX-XXVII).

which is approaching this limit. Unlike Thompson scattering, Compton scattering is not a purely elastic process, and so removes energy from the photons.

Treatment of an individual Compton scattering event requires consideration of the direction of the photon and electron before and after the scattering process, taking account of both energy and momentum conservation. The energy transfer between the electrons and photons depends on the energy of each, and the change of direction during the interaction. However, since we are dealing with a photon packet representing a large number of photons, interacting with a cell containing electrons with a range of velocities, we require a statistical treatment.

If one considers a population of electrons with velocities that can be represented by a Boltzmann distribution, then the mean energy transfer can be calculated by an integration over all possible angles. For a photon of energy $\epsilon = h\nu$, the mean fractional energy transfer per scattering event is given (in the non-relativistic limit) by ([Rybicki and Lightman 1979](#))

$$\frac{\Delta\epsilon}{\epsilon} = \frac{4kT - \epsilon}{mc^2}. \quad (5.31)$$

A plot of the mean fractional energy change of a photon per interaction is shown in Figure [5.16](#).

$\Delta\epsilon/\epsilon$ is negative when $h\nu > 4kT$, which represents a transfer of energy from the photons to the electrons - this is Compton heating of the gas. This case is plotted in the top panel of Figure [5.16](#) and occurs when high energy photons pass through a relatively cool gas.

When low frequency photons pass through a relatively hot gas, the opposite happens, and we have the inverse Compton effect - the energy of photons is increased by collisions with hot electrons. This case is shown in the lower panel of Figure [5.16](#). This is the process thought to be responsible for the X-ray power law tail seen in the SED of AGN, where seed photons from the accretion disk are Compton up-scattered by an atmosphere of hot electrons.




Figure 5.16: Fractional energy transfer per Compton scattering.

In the new version of PYTHON, we include Compton heating as a form of opacity which removes weight from photon packets during the photon propagation part of a calculation. It acts as an additional heating mechanism alongside free-free and bound-free (see Section 4.4.4). Compton cooling is implemented as an additional cooling term in the heating/cooling calculations (see Section 4.5). The next three subsections describe how Compton heating and cooling are implemented.

5.4.1.1 Compton heating

The “heating” component of Equation 5.31 is given by

$$\left(\frac{\Delta\epsilon}{\epsilon} \right)_{HEATING} = \frac{h\nu}{mc^2}, \quad (5.32)$$

which represents the mean energy loss from photon to electron in a single interaction. In PYTHON wind heating is calculated during the photon packet transport phase of a

simulation. The photon packets lose weight to continuum processes along their paths through the wind, and that lost weight represents a heating input to the wind cells. In order to calculate how much energy is lost from a photon packet due to this process, we need to compute the number of interactions that take place along a path segment of length ds . The differential cross section of an electron to a photon of frequency ν is given by the Klein-Nishina formula. Since we are carrying out an angle-averaged calculation, we use the angle-averaged Klein-Nishina formula from [Rybicki and Lightman \(1979\)](#)

$$\sigma_c(\nu) = \sigma_T \cdot \frac{3}{4} \left[\frac{1+x}{x^3} \left\{ \frac{2x(1+x)}{1+2x} - \ln(1+2x) \right\} + \frac{1}{2x} \ln(1+2x) + \frac{1+3x}{(1+2x)^2} \right] \text{ cm}^2, \quad (5.33)$$

where $x = h\nu/mc^2$. Therefore, a photon packet of frequency ν will scatter $n_e \sigma_C(\nu)$ times per cm travelled, and since each interaction will change its energy by a factor given by Equation 5.32, we can write an absorption coefficient for this process:

$$\kappa_\nu^{Comp} = n_e \sigma_c(\nu) \frac{h\nu}{mc^2} \text{ cm}^{-1}. \quad (5.34)$$

Equation 4.41 becomes

$$\tau_{abs} = \kappa_\nu^{ff} ds + \kappa_\nu^{bf} ds + \kappa_\nu^{Comp} ds, \quad (5.35)$$

and the weight of a photon packet passing through a given cell is further reduced due to Compton “opacity”. The lost weight is an additional heating input into the cell,

$$H^{Comp} = w_{in} (1 - e^{-\tau_{abs}}) \frac{\tau^{Comp}}{\tau_{abs}}, \quad (5.36)$$

where

$$\tau^{Comp} = \kappa_\nu^{Comp} ds. \quad (5.37)$$

Note that this implementation is not strictly physical; in reality, the energy lost from a photon in such an interaction results in a change in its frequency, whereas we represent this energy loss as a reduction in the weight of the photon packet, at fixed frequency. Since we are primarily interested in ions with ionization energies up to a few hundred eV,

and lines in the UV, this approximation will not significantly affect results. This is because the fractional frequency shift for relevant photon packets would be much less than 1% per interaction, and we are normally operating in the optically thin regime. Nonetheless, we note that the spectrum at frequencies where $h\nu$ approaches mc^2 ($\nu \gtrsim 10^{19}$ Hz) should be treated with caution.

5.4.1.2 Induced Compton heating

In situations where

$$J_\nu \gtrsim \frac{2h\nu^3}{c^2}, \quad (5.38)$$

induced (or stimulated) Compton heating can also take place. This heating rate is identical to the Compton heating rate, but with an additional multiplicative term ([Cloudy & associates 2011](#)): the photon occupation number, $\eta_i(\nu)$, given by

$$\eta_i(\nu) = \frac{J_\nu}{\frac{2h\nu^3}{c^2}}. \quad (5.39)$$

We can therefore write an induced Compton absorption coefficient

$$\kappa_\nu^{ind Comp} = \eta_i(\nu) n_e \sigma_c(\nu) \frac{h\nu}{mc^2}. \quad (5.40)$$

We use the spectral model from the previous ionization cycle to estimate J_ν when calculating this heating mechanism. The absorption coefficient is included (multiplied by ds) as a further additive term in Equation 5.35, and the induced Compton heating input is calculated as for Compton heating. Again, the packet weight is reduced rather than the frequency and so the same caveat applies here as in Compton heating.

5.4.1.3 Compton cooling

When $4kT > h\nu$, the fractional energy change per interaction is positive, and so energy passes from the electrons to the photons. This process is inverse Compton scattering which cools the plasma. We therefore include it along with the other cooling terms that are calculated after each cycle of photon propagation (see section 4.5.2).

Each interaction results in a fractional energy loss from the plasma of

$$\left(\frac{\Delta\epsilon}{\epsilon}\right)_{COOLING} = \frac{4kT}{mc^2}. \quad (5.41)$$

Taking a ray of intensity I_ν , the energy per solid angle, per unit frequency interval per second intercepting an electron in the gas is equal to $I_\nu\sigma_c(\nu)$. Therefore the total energy lost by this ray is equal to $I_\nu\sigma_c(\nu) \times 4kT/mc^2$.

The total energy loss by all rays passing through a given cell, with electron number density n_e and volume V is therefore given by

$$C^{Comp} = Vn_e \frac{4kT}{mc^2} \int_{\Omega} \int_{\nu} I_\nu \sigma_c(\nu) d\nu d\Omega. \quad (5.42)$$

Since this process is only important for low frequency photons, the relevant cross section is the Thompson cross section, σ_T . The double integral then reduces to 4π times the frequency-averaged mean intensity J in the cell, which we already log as part of the spectral modelling. Thus, the Compton cooling rate (C^{Comp}) is given by

$$C^{Comp} = 16\pi V \sigma_T J \frac{kT_e}{m_e c^2}. \quad (5.43)$$

As with the other cooling mechanisms, this is a function of the electron temperature T_e , and so can be included in the heating and cooling balance calculations.

PYTHON generates photon packets to balance free-free, recombination and line cooling. However, in the case of Compton cooling, it is not so clear how to proceed. Compton cooling does not *generate* photons, rather it increases the frequency of photons passing through the hot plasma - Compton up-scattering. As implemented, the cooling mechanism is calculated after the photon packets have been propagated, so we cannot change the frequency retrospectively. We therefore neglect Compton cooling as a source of wind radiation, but note the approximation. In cases where Compton cooling is an important process, we will likely underestimate the emergent flux at high frequencies.

5.4.2 Dielectronic recombination cooling

Since we have included dielectronic recombination in the ionization scheme, it should also be included as a cooling mechanism for consistency. For radiative recombinations, the recombination rate is calculated from the photoionization cross sections via the Milne relation. The analogous calculation for dielectronic recombination cooling would be to carry out an integration of the Auger ionization cross sections, since this is the inverse process. We have not yet carried out this work, and so we include only a rough estimate of the dielectronic recombination cooling rate.

The temperature-dependent dielectronic recombination rate coefficient multiplied by the electron number density and the number density of the recombining ion gives the number of recombinations per unit volume. As an estimate of the energy removed from the cell by each recombination, we use the mean electron energy. The total dielectronic recombination rate in a cell is therefore given by

$$C^{DR}(T) = V \sum_{ions=i} n_e n_i 2k T_e \alpha_i^{DR}(T). \quad (5.44)$$

This approximate treatment allows us to understand the relative importance of this cooling process, although the current implementation is not fully self-consistent.

Free-free, free-bound and line cooling processes in the wind all produce photons packets to balance the cooling rate and so maintain the total luminosity of the the system. Since we do not have the frequency-resolved emissivity (which would come from the Auger ionization Milne relation), we do not have the data required to generate photon packets of the correct frequency. We therefore again neglect this process, but note that if dielectronic recombination is a significant cooling mechanism, we could be underestimating the total luminosity of the wind.

5.4.3 Validation of new heating and cooling mechanisms

As with the ionization calculations, we turn to CLOUDY as a well tested comparison for our calculations. We use the same geometry and illuminating spectrum as in Section

5.3.5, but this time the temperature is solved for as well. We already know that for a given temperature, PYTHON produces a slightly different ionization state from CLOUDY, so we limit this effect by disabling Auger ionization, charge exchange and three-body recombination in the CLOUDY calculations.

Figure 5.17 shows the results of the comparison, with electron temperature plotted against ionization parameter on the left hand y-axis. The volumetric heating and cooling rates for the two codes are plotted on the right hand axis - since both successfully converge, the heating and cooling curves overlap each other. The agreement is excel-




Figure 5.17: Temperature and heating/cooling rate for CLOUDY and PYTHON.

lent, particularly at the extremes, and we can clearly see that both codes converge to a similar Compton temperature for $2 \lesssim \log(U) \lesssim 7$. The slight difference in the Compton temperature is mainly due to differences between the cross sections used in CLOUDY and PYTHON. The departure from the Compton temperature at very high values of U arises because (with the simple SED we are using here), the photon occupation number at low frequencies exceeds that of a blackbody. Thus induced Compton heating becomes important. The heating rate is proportional to J^2 - hence the change in gradient of the heating curve.

Figure 5.18 shows the heating rate for both codes split by physical mechanism. At low values of U , the dominant heating mechanism is photoionization, and we agree very

well with CLOUDY (although note that we have turned off two important ionization mechanisms in CLOUDY). The disagreement is slightly worse between $0 \lesssim \log(U) \lesssim 2$, and the slightly higher heating rate in PYTHON is the reason for the shape of the T vs $\log(U)$ curve being different around this point. This is where iron ionization heating (including un-modelled UTA heating) becomes dominant.




Figure 5.18: Heating rates for CLOUDY (symbols) and PYTHON (solid lines).

Figure 5.19 shows the cooling rates split up by mechanism. Once more we see very good agreement between the two codes.




Figure 5.19: Cooling rates for CLOUDY (symbols) and PYTHON (solid lines).

5.5 Treatment of excited levels

Much effort has been expended in this project to allow accurate calculation of the ionization balance of gas illuminated by a non-thermal SED. This has in large part been successful, as evidenced by the promising agreement between ionization states predicted in PYTHON and CLOUDY. However, this is not the complete picture, since ions can exist in excited states. The excitation level of ions is particularly important in the calculation of line emission, but it can also affect the detailed ionization state through ionization from excited levels. There has not been sufficient time to address excitation of ions for non-thermal SEDs in this project, but in this section we will discuss the approximations that have been made.

There are three main areas in which the calculation of level populations affects a RT and photoionization calculation. First, in the ionization calculations, the Saha equation only links two individual levels of an ion. We have chosen to calculate the relative populations of the ground state of two adjacent ions. One can generalise the calculation to include all levels of the two adjacent ions by replacing the ground state degeneracies ($g_{i+1,0}$ and $g_{i,0}$ in Equation 5.1) with the partition functions for the two ions. In the original version of PYTHON, the partition functions are calculated using a first order non-LTE correction, given by Equation 4.16. This equation is based upon the assumption that the levels are radiatively excited, and the specific intensity can be modelled as a dilute blackbody. Initial testing showed that this approximation was not appropriate for the non-thermal SEDs we were modelling, and in fact worsened our agreement with CLOUDY. The correct way to deal with this would be to solve for level populations using photo-excitation rates based on the actual specific intensity in a cell and also collisional excitation rates based on the electron density and velocity distribution. This is a complex task, beyond the scope of the project. In situations where we are modelling the specific intensity with exponential and power law models rather than a dilute blackbody, we therefore assume that all ions exist only in the ground state. Given the good agreement between PYTHON and CLOUDY, this appears to be reasonable over a large range of ionization parameters.

In the original version of PYTHON, the partition function is also used in the calculation of

bound-free opacity, where photoionization from excited levels is included. If we assume all ions are only in the ground state, this can clearly affect the total level of bound-free opacity and heating. Free-bound cooling would be similarly affected, since it is calculated via the Milne relation, and excited levels are also included in this calculation. However, we can be reasonably sure that if the calculation of the ionization *state* is not affected adversely by the assumption of all ions only being in the ground state, then the effect on the heating and cooling rates will be similarly negligible. This is borne out once more by the good agreement with CLOUDY for the heating and cooling rates.

Finally, level populations are central to the calculation of line opacities and emission. In the original version of PYTHON, level populations were calculated using a two-level atom model (see Section 4.3.4.3 and specifically Equation 4.29 for details). This takes account of stimulated excitation and de-excitation through the normal assumption that the mean intensity is a dilute blackbody. This formulation has not been changed in the new version of PYTHON, and so the mean intensity will not be accurately modelled. It is perhaps surprising therefore, that the overall line cooling rate appears to be correctly calculated, at least for the geometry and illuminating spectrum used to create figure 5.19. In this case, the modelled dilute blackbody has a very small radiative weight, and so the populations in the two level atom are set by collisional excitation and spontaneous radiative de-excitation. Presumably, the line populations in CLOUDY are also set collisionally.

It seems therefore that use of the less accurate dilute blackbody model for mean intensity in the two level atom approximation does not cause significant error, at least to the overall line cooling rate. It is less clear at the time of writing how it affects the modelling of which lines actually emit the energy. The way forward is clear - use the same model for mean intensity as used in the ionization calculations. This is an improvement that will take place in the next phase of development of PYTHON.

5.6 Atomic data

For the calculations described in this work, we include atomic data for H, He, C, N, O, Ne, Na, Mg, Al, Si, S, Ar, Ca and Fe. Elemental abundances are taken from [Verner](#)

et al. (1994), and the basic atomic data, such as ionization potential and ground state multiplicities, are taken from Verner et al. (1996b). For H, He, C, N, O and Fe, we use TOPBASE (Cunto et al. 1993) data for energy levels and photoionization cross sections from both ground state and excited levels. For other elements, we use the analytic approximations for ground state photoionization cross sections given in Verner et al. (1996a). For bound-bound transitions, we use the line list given in Verner et al. (1996b), which contains nearly 6000 ground-state-connected lines. The levels associated with these lines are computed from the line list information using a method similar to that described by Lucy (1999).

We adopt dielectronic recombination rate coefficients and total radiative rate coefficients from the CHIANTI database version 7.0 (Dere et al. 1997; Landi et al. 2012). Ground-state recombination rates are taken from Badnell (2006) where available and otherwise computed from the photoionization data via the Milne relation. Finally, we use Gaunt factors for free-free interactions given by Sutherland (1998).

Chapter 6

A Benchmark Disk Wind Model For BALQSOs

6.1 Introduction

In Section 1.3.3 we introduced a disk wind model for the BLR in AGN which suggests a unification scheme linking type 1 quasars and BALQSOs via wind emission and absorption. In order to assess the feasibility of such a model, it is clearly important to determine the physical properties – i.e. the geometry, kinematics and energetics – of these outflows. However, while there have been many previous efforts in this direction (see Chapter 3) there has not really been a convergence towards a unique physical description. Our aim here is to test if a simple, physically motivated disk wind model can give rise to spectra that resemble those of BALQSOs when the line of sight towards the central engine lies within the wind cone.

In Section 6.2, we describe the family of kinematic disk wind models we use to describe AGN/QSO outflows. In Section 6.3, we discuss observational and theoretical constraints that restrict the relevant parameter space of the model. These considerations allow us to define a benchmark disk wind model as a reasonable starting point for investigating the impact of disk winds on the spectra of AGN/QSO.




Figure 6.1: A sketch illustrating the main features of our kinematic disk wind model.

In Section 6.4, we present the results from simulations of this benchmark model, including synthetic spectra. We then go on to discuss the implications of what we see.

6.2 The Kinematic Disk Wind Model

Since the driving mechanism, geometry and dynamics of AGN/QSO disk winds are all highly uncertain, we use a flexible, purely kinematic model to describe the outflow. This allows us to describe a wide range of plausible disk winds within a single, simple framework. The specific prescription we use was developed by [Shlosman and Vitello \(1993, SV93 hereafter\)](#) to model the accretion disk winds observed in CVs, i.e. accreting white dwarf binary systems.

The geometry of our outflow model is illustrated in Figure 6.1. A biconical wind is taken to emanate between radii r_{min} and r_{max} in the accretion disk, with the wind boundaries

making angles of θ_{min} and θ_{max} with the axis of symmetry. At each radius, r_0 , within this range, the wind leaves the disk with a poloidal (non-rotational) velocity vector oriented at an angle θ to the axis of symmetry, with θ given by

$$\theta = \theta_{min} + (\theta_{max} - \theta_{min}) x^\gamma, \quad (6.1)$$

where

$$x = \frac{r_0 - r_{min}}{r_{max} - r_{min}}. \quad (6.2)$$

The parameter γ can be used to vary the concentration of poloidal streamlines towards either the inner or outer regions of the wind, but throughout this work we fix it to 1, so the poloidal streamlines are equally spaced in radius.

The poloidal velocity, v_l , along a streamline in our model is given by

$$v_l = v_0 + [v_\infty(r_0) - v_0] \frac{(l/R_v)^\alpha}{(l/R_v)^\alpha + 1}, \quad (6.3)$$

where l is distance measured along a poloidal streamline. This power-law velocity profile was adopted by SV93 in order to give a continuous variation in the derivative of the velocity and a realistic spread of Doppler-shifted frequencies in the outer portion of the wind ($l > R_v$). We have similar requirements. The initial poloidal velocity of the wind, v_0 is (somewhat arbitrarily) set to 6 km s^{-1} for all streamlines, comparable to the sound speed in the disk photosphere. The wind then speeds up on a characteristic scale length R_v , defined as the position along the poloidal streamline at which the wind reaches half its terminal velocity, v_∞ . The terminal poloidal velocity along each streamline is set to a fixed multiple of the escape velocity at the streamline foot point,

$$v_\infty = f v_{esc}, \quad (6.4)$$

so the innermost streamlines reach the highest velocities. The power-law index α controls the shape of the velocity law: as α increases, the acceleration is increasingly concentrated around $l = R_v$ along each streamline. For large α , the initial poloidal velocity stays low near the disk and then increases quickly through R_v to values near v_∞ .




Figure 6.2: Dependence of mass loss rate per unit area on the disk on radius for three values of λ . \dot{M} and $\theta(r_0)$ are set as in the benchmark model defined in Section 6.3

The wind is assumed to initially share the Keplerian rotation profile of the accretion disk, i.e. $v_{\phi,0}(r_0) = v_K(r_0)$ at the base of the outflow. As the wind rises above the disk and expands, we assume that specific angular momentum is conserved, so that the rotational velocities decline linearly with increasing cylindrical distance from the launch point.

The SV93 prescription also allows us to control the mass loading of the outflow as a function of radius. More specifically, the local mass-loss rate per unit surface area on the disk (\dot{m}_{wind}) is given by

$$\dot{m}_{wind} \propto \dot{M}_{wind} r_0^\lambda \cos[\theta(r_0)] \quad (6.5)$$

where \dot{M}_{wind} is the global mass loss rate through the wind. The run of mass loss rate per unit area with r_0 is illustrated in Figure 6.2.

The combination of local mass-loss rate, launching angle and poloidal velocity is sufficient to uniquely determine the density at any point in the wind via a continuity equation (see SV93 for details). The resulting wind is a smooth, single phase outflow without clumps. There is evidence for structure in BALQSO outflows, apparent from complex line shapes (e.g. [Ganguly et al. 2006](#); [Simon and Hamann 2010](#)) and time variability (e.g. [Capellupo et al. 2011, 2012, 2013](#)), but this is something our simple model cannot address. However,

hot star winds and CV winds also exhibit variability and small-scale structure, yet much has been learned about these winds by focussing on the global, underlying smooth flow field. This is the approach followed here.

6.3 A Benchmark Disk Wind Model For BALQSOs

One of the main challenges in trying to model BALQSOs is that the physical characteristics of the disk winds that produce the BALs are highly uncertain. The kinematic model described in Section 6.2 is sufficiently flexible to describe a wide range of outflow conditions, but the price for this flexibility is a potentially huge parameter space. We have therefore used a variety of theoretical and observational constraints (and considerable trial and error) to design a benchmark disk wind model whose predicted spectrum resembles that of a BALQSO (for sightlines cutting through the outflow).

Since we have not yet carried out a systematic exploration of the relevant parameter space, we make no claim that this fiducial model is optimal, nor that it faithfully describes reality. Indeed, in Section 6.5, we will discuss some of its shortcomings. However, we believe it is nevertheless extremely useful as a benchmark. More specifically, it provides (i) a first indication of what can be realistically expected of disk wind models; (ii) a better understanding of the challenges such models face; (iii) a convenient starting point for further explorations of parameter space.

In the rest of this section, we will discuss the constraints that have informed our choice of parameters for the benchmark model, which are summarised in Table 6.1.

6.3.1 Basic Requirements: BALnicity and Ionization State

Our primary aim here is to find a set of plausible parameters for a QSO wind that yield synthetic spectra containing the features expected of a BALQSO from a range of viewing angles. As noted in Section 6.1, in the majority of BALs, the HiBALs, the absorption features are due to highly ionized species such as N V $\lambda\lambda 1240$, C IV $\lambda\lambda 1550$, and Si IV $\lambda\lambda 1400$. The C IV feature, in particular, is most often used to identify BALQSOs, so

we judge the initial success of candidate wind models by their ability to produce this feature. In practice, this is primarily a constraint on the ionization state of the wind: we require that C IV should be present – if not necessarily dominant – throughout an appreciable fraction of the outflow.

BALs are usually identified via the so-called BALnicity index (BI) (Weymann et al. 1991), which is a measure of absorption strength in the velocity range between -3000 km s^{-1} and -25000 km s^{-1} . It is therefore not sufficient to produce C IV anywhere in the outflow – in order for the model to resemble a BAL, C IV needs to be present in reasonable concentrations in regions characterized by the correct velocities.

6.3.2 Black Hole Mass, Accretion Rate and Luminosity

Our goal is to model a fairly typical, high-luminosity (BAL)QSO, so we adopt a black hole mass of $10^9 M_\odot$, along with an Eddington ratio of $\epsilon = 0.2$. For an assumed accretion efficiency of $\eta = 0.1$, this corresponds to an accretion rate of $\dot{M}_{\text{accretion}} \sim 5 M_\odot \text{yr}^{-1}$ and a bolometric disk luminosity of $L_{bol} \simeq 2.4 \times 10^{46} \text{ ergs s}^{-1}$.

6.3.3 Outflow Geometry

BALQSOs represent $\simeq 20\%$ of all QSOs (e.g. Knigge et al. 2008; Hewett and Foltz 2003, but also see Allen et al. 2011). There are two obvious ways to interpret this number: (i) either all QSOs spend $\simeq 20\%$ of their life as BALQSOs (“evolutionary unification”; e.g. Becker et al. 2000), or (ii) all QSOs contain disk winds, but the resulting BAL signatures are visible only from $\simeq 20\%$ of all possible observer orientations (“geometric unification”; e.g. Elvis 2000).¹ Here, we adopt the geometric unification picture, with the aim of testing its viability.

As illustrated in Figure 6.3, our wind geometry allows for three distinct classes of sight-line towards the central engine. Looking from the polar direction, “above the wind cone”,

¹A 20% BAL fraction could be an underestimate, since intrinsic BALQSO fractions tend to be derived from magnitude-limited samples, with no allowance for the possibility that foreshortening, limb-darkening or wind attenuation may reduce the continuum flux of a typical BALQSO relative to that of a typical QSO (Krolik and Voit 1998). For the purpose of the present paper, we ignore this potential complication.




Figure 6.3: Cartoon illustrating the three broad classes of sightline discussed in the text.

the sightline does not intersect the wind at all, and looking from the equatorial direction, the sightline passes through the base of the wind cone where the projected line of sight velocities are too low to produce BAL features. We anticipate that BAL features will result when the central engine (or, strictly speaking, the UV-bright portions of the accretion disk) are viewed from the third class of sightline - looking “into the wind cone”. We therefore expect such sightlines to subtend $\simeq 20\%$ of the sky as seen by the central engine. Note that this description of the sightlines only applies to models like our benchmark model, in which most of the emission arises from inside the wind launching radius.

If we define θ_{min} and θ_{max} as the angles made by the inner and outer boundaries of the wind with respect to the disk normal, these constraints translate into a simple relationship between θ_{min} and θ_{max} , i.e.

$$0.2 = \cos(\theta_{min}) - \cos(\theta_{max}). \quad (6.6)$$

Most geometrical unification scenarios assume that BAL-producing QSO winds flow predominantly along the equatorial plane of the disk, rather than along the polar axis (e.g. [Weymann et al. 1991](#); [Ogle et al. 1999](#); [Elvis 2000](#); however, see [Zhou et al. 2006](#) and [Brotherton et al. 2006](#) for a contrary view). In line with this general view, we have found that ionization states capable of producing C IV and other typical BAL features are more naturally produced in equatorial, rather than polar disk winds. This is because one of the main challenges faced by disk wind models is to prevent the over-ionization of the

outflow by the central X-ray source (see below). Other things being equal, equatorial winds are less ionized, partly because of limb darkening, and partly because the disk appears foreshortened from equatorial directions. In addition, in geometries like ours, where the wind arises from parts of the disk radially outside the location where the ionizing radiation is produced, an equatorial wind exhibits self-shielding, i.e. the radiation must pass through the inner parts of the wind to reach the outer parts. Therefore the ionization state can be lower than for a polar wind, where all parts of the wind can see the ionizing radiation directly.

Based on these considerations, we adopt a mostly equatorial geometry with $\theta_{min} = 70^\circ$ and $\theta_{max} = 82^\circ$ for our benchmark model. With this choice, 66% of the remaining sightlines lie “above” the wind cone, while 14% view the central source through the base of the wind. If geometric unification is correct and complete – in the sense that the emission lines in non-BAL QSOs are produced by the same outflow that produces the absorption features seen in BALQSOs (e.g. [Elvis 2000](#)) – we would ideally see different classes of QSO as we view a given system from the three distinct types of viewing angles.

6.3.4 Wind-Launching Region

There is no consensus on the location of the wind-launching region on the accretion disk. However, there are several empirical and theoretical considerations that can inform the design of our benchmark model.

First, within the geometric unification scenario, QSO disk winds are responsible not only for BAL features, but also for the broad emission lines (BELs) seen in non-BAL QSOs. We can therefore use the result of reverberation mapping studies of this broad line region (BLR) to guide us in locating the outflow. [Kaspi et al. \(2000\)](#) have carried out reverberation mapping of the C IV BLR for high luminosity quasars. A representative example is S5 0836+71, a quasar with a black hole mass of $2.6 \times 10^9 M_\odot$, for which they found a rest frame delay of 188 days. This delay implies a distance of the line-forming region from the UV continuum source of $1260 r_g$.² Since the C IV line-forming region

²Note that the UV continuum is produced at very small radii compared to the wind launch radius and so can be treated as originating from the centre of the model for these estimates

may lie substantially downstream in the outflow, this is an upper limit on the distance of the launching region from the centre.

Second, BAL features exhibit variability on timescales from days to years (e.g. [Capellupo et al. 2011, 2012, 2013](#)). If interpreted in terms of wind features crossing our line of sight, the *shortest* variability time scales – which arise in the innermost parts of the line-forming region – correspond to distances of $\sim 3 \times 10^{16}$ cm or, equivalently, $\sim 200 r_g$ from the centre (for an assumed $M_{BH} \simeq 10^9 M_\odot$).

Third, one of the most promising mechanisms for driving QSO disk winds is via radiation pressure in spectral lines (e.g. [Proga et al. 2000; Proga and Kallman 2004](#)). Empirical support for this idea comes from the presence of the “ghost of Lyman α ” in some N v BAL features (see [Arav et al. 1995; Arav 1996; North et al. 2006](#); but also see [Cottis et al. 2010](#)). In published simulations where such winds are modelled in a physically consistent manner, the typical launching radii are (a few) $\times 100 r_g$ (e.g. [Proga et al. 2000; Risaliti and Elvis 2010](#)).

Fourth and finally, it is physically reasonable to assume that the maximum velocity achieved by the outflow corresponds roughly to the fastest escape velocity in the wind launching region on the accretion disk. Observationally, $v_{max} \sim 20,000 \text{ km s}^{-1}$, which, for a $10^9 M_\odot$ black hole, corresponds to the escape velocity from $\simeq 400 r_g$.

Taken together, all of these considerations suggest that a reasonable first guess at the location of the wind-launching region on the disk is on the order of a few hundred r_g . In our benchmark model, we therefore adopt $r_{min} = 300 r_g$ and $r_{max} = 600 r_g$.

6.3.5 Wind Mass-Loss Rate

One of the most fundamental parameters of any wind model is the mass-loss rate into the outflow, \dot{M}_{wind} . In our benchmark model, we set $\dot{M}_{wind} = \dot{M}_{acc} = 5 M_\odot \text{ yr}^{-1}$. This value was arrived at mainly by trial and error, but with a conservative preference for values satisfying $\dot{M}_{wind} \lesssim \dot{M}_{acc}$. As shown explicitly in Section 6.5.1, in practice, we found that we required $\dot{M}_{wind} \simeq \dot{M}_{acc}$ in order to produce BAL features for even modest X-ray luminosities.

Strictly speaking, any model with $\dot{M}_{wind} \gtrsim \dot{M}_{acc}$ is not entirely self-consistent, since the presence of such an outflow would alter the disk's temperature structure (e.g. [Knigge 1999](#)). However, we neglect this complication, since the vast majority of the disk luminosity arises from well within our assumed wind-launching radius. Thus even though our model ignores that the accretion rate *is* higher further out in the disk, it correctly describes the innermost disk regions that produce virtually all of the luminosity.

Finally, in the absence of evidence to the contrary, we take the simplest possible prescription for the run of the mass-loss rate with radius across the wind-launching region, i.e. $\lambda = 0$ (see Equation 6.5 in Section 6.2).

6.3.6 Velocity Law Parameters

Ideally, the parameters defining the poloidal velocity law of the wind would be predicted by the relevant acceleration mechanism. However, since this is not yet possible, we have simply adopted parameters that result in reasonably BAL-like C IV features.

As noted in Section 6.3.4 above, it seems physically plausible that the terminal outflow velocity along a poloidal streamline will roughly reflect the local escape velocity at the streamline footprint. Having already adopted wind launching radii based on this idea, we therefore also adopt $v_\infty = v_{esc}$, i.e. $f = 1$ in Equation 6.4.

The other two parameters defining the poloidal velocity law are R_v , the poloidal distance at which the velocity is equal to half the terminal velocity, and α , the exponent describing how the velocity varies with poloidal distance. As discussed above, the locations of BAL- and BEL-forming regions within the wind likely extend to distances of up to about $2000 r_g$, so this is a reasonable number to adopt for R_v . For the $10^9 M_\odot$ black hole in our benchmark model, this corresponds to $R_v = 10^{18}$ cm. Finally, in the absence of other information, we adopt $\alpha = 1$. This value produces a relatively slow acceleration as a function of poloidal distance along a streamline and seems to produce reasonable BAL shapes. We will explore the effect of changing these parameters in future work.

6.3.7 X-ray Spectrum and Luminosity

The X-ray luminosity of the central source is set to $L_X = 10^{43}$ ergs s^{-1} in our benchmark model. This value was arrived at through the requirement that the model should produce BALs which constrains the X-ray luminosity, since X-rays can strongly affect the ionization state of the outflow. We set the spectral index of the X-ray spectrum to $\alpha_X = -0.9$, as suggested by [Giustini et al. \(2008\)](#).

Quasar spectra are often characterised in terms of the ratio of X-ray to bolometric (or UV) luminosity. This ratio is usually characterized by the spectral index between X-rays and optical/UV, α_{OX} , defined as

$$\alpha_{OX} = 0.3838 \log \left(\frac{L_\nu(2 \text{ keV})}{L_\nu(2500 \text{ \AA})} \right), \quad (6.7)$$

where $L_\nu(2 \text{ keV})$ and $L_\nu(2500 \text{ \AA})$ are the inferred monochromatic luminosities at 2 keV (X-ray) and 2500 Å (UV), respectively. Ignoring the angle dependence of the disk emission, we obtain a mean value of $\alpha_{OX} \sim -2.2$ for the parameters used in our benchmark model. This is considerably lower than observed in typical non-BAL QSOs ([Just et al. 2007](#)) and lies near the lower (X-ray weak) end of the range of α_{ox} seen in BALQSOs (e.g [Giustini et al. 2008; Gibson et al. 2009](#)). The X-ray weakness of the model, as well as the sensitivity of our results to the adopted X-ray luminosity, are discussed in more detail in Section 6.5.1.

6.3.8 Size and Resolution of the Numerical Grid

The outer edge of our simulation grid is set to 1×10^{19} cm, large enough to ensure that all spectral features of interest arise entirely within the computational domain. In all simulations shown in this paper, this domain is sampled by a cylindrical grid composed of 100×100 logarithmically spaced cells. We have carried out resolution tests to ensure that this is sufficient to resolve the spatially varying physical conditions and ionization state of the outflow.

Free Parameters	Value
M_{BH}	$1 \times 10^9 M_{\odot}$
\dot{M}_{acc}	$5 M_{\odot} yr^{-1} \simeq 0.2 \dot{M}_{Edd}$
α_X	-0.9
L_X	$1 \times 10^{43} \text{ ergs s}^{-1}$
$r_{disk}(\text{min}) = r_X$	$6r_g = 8.8 \times 10^{14} \text{ cm}$
$r_{disk}(\text{max})$	$3400r_g = 5 \times 10^{17} \text{ cm}$
\dot{M}_{wind}	$5 M_{\odot} yr^{-1}$
r_{min}	$300r_g = 4.4 \times 10^{16} \text{ cm}$
r_{max}	$600r_g = 8.8 \times 10^{16} \text{ cm}$
θ_{min}	70.0°
θ_{max}	82.0°
λ	0
v_{∞}	$v_{esc}(\text{f=1})$
R_v	$1 \times 10^{18} \text{ cm}$
α	1.0
Derived Parameters	Value
$L_{\nu}(2500\text{\AA})$	$6.3 \times 10^{30} \text{ ergs s}^{-1} \text{ Hz}^{-1}$
$L_{\nu}(2keV)$	$1.2 \times 10^{25} \text{ ergs s}^{-1} \text{ Hz}^{-1}$
L_{bol}	$2.4 \times 10^{46} \text{ ergs s}^{-1}$
M_{bol}	-27.4
M_u	-26.2
α_{OX}	-2.2

Table 6.1: Wind geometry parameters used in the benchmark model.

6.4 Results

Having defined a benchmark disk wind model for AGN/QSOs, we are ready to examine its predicted physical and observable characteristics. In Section 6.4.2 we will present synthetic spectra reminiscent of BALQSOs for sightlines looking into the wind cone. However, it is helpful to examine the physical and ionization state of the benchmark model before analyzing these spectra in detail.

6.4.1 The Physical State of the Outflow

The top four panels in Figure 6.4 show a selection of physical parameters of the converged wind model. Considering first the electron density (top left panel), we see that our choice of kinematic parameters has given rise to an outflow with a fairly high density of $n_H \sim 10^{10} \text{ cm}^{-3}$ at its base. However, the density declines quickly as we move outwards,

due to the expansion and acceleration of the wind. Hydrogen is fully ionized throughout the entire outflow, so $n_H \simeq n_e$ everywhere.

The electron temperature in the wind ranges from $\sim 10^3$ K near the base of the wind to more than 10^5 K (top right panel). The highest temperatures are in a thin layer near the “top” of the wind, at distance of $\simeq 10^{18}$ cm from the central engine. These regions are hot because they are directly exposed to the radiation field of the accretion disk and the X-ray source. These regions also shield the wind material “behind” them, however, and thus help to ensure more moderate temperatures in the rest of the outflow. In fact, much of the disk wind is heated to temperatures near $T_e \sim 10^4$ K, which are quite conducive to the formation of the high-ionization lines typically seen in (BAL)QSOs, such as C IV, Si IV and N V.

The middle panels of Figure 6.4 illustrate the velocity structure of our benchmark model, separated into poloidal (middle left panel) and rotational (middle right panel) components. The poloidal velocities show the relatively gradual speed-up of the outflow as a function of distance, which is due to our choice of velocity law exponent. As one would expect, low velocities are found near the base of the wind and high velocities are reached only quite far out. It is this variation in poloidal velocity which produces BAL features as photons produced by the effectively point-like central UV source are scattered out of the line of sight by progressively faster moving wind parcels.

By contrast, the highest rotational velocities are found near the disk plane, where they are effectively equal to the Keplerian velocities in the disk. They then decline linearly with increasing cylindrical distance from the rotation axis, because we assume that wind material conserves specific angular momentum. The projected rotational velocity along a sightline to the central source is zero (or nearly so), and therefore rotational velocities do not play a major part in the formation of BAL features, although they should help in producing broad *emission* lines.

The bottom left and right panels in Figure 6.4 show the ionization parameter and C IV ionization fraction throughout the outflow, respectively. Interestingly, the model generates sufficient concentrations C IV to make strong BALs even though $U > 1$ throughout

most of the wind. This contrast with the results shown for a thin-shell model in Figure 5.6, where this ion is not present for such values of U . The main reason for this difference is that the wind is self-shielding: absorption in the inner parts of the outflow (primarily due to photoionization) can strongly modify the SED seen by the outer parts of the wind. This, in turn, changes the ionization parameter at which particular ions are preferentially formed. Figure 6.5 illustrates the effect of the self-shielding on the SED seen by different parts of the outflow.

6.4.2 Observing the Model: Synthetic Spectra

Figure 6.6 shows the spectra we predict for observers located along the directions indicated in the lower right panel of Figure 6.4. These viewing angles correspond to sightlines to the central source that look above the wind cone ($i = 40^\circ$), into the wind cone ($i = 75^\circ$ and $i = 80^\circ$), and through the base of the wind cone ($i = 85^\circ$). For comparison, we also show a pure, unobscured continuum spectrum for each viewing angle; these continua were generated by re-running the model with the wind density set to a negligible value.

6.4.2.1 Sightlines looking into the wind cone

The two middle panels in Figure 6.6 show the emergent spectra for sightlines through the wind ($i = 75^\circ$ and $i = 80^\circ$). These spectra clearly contain BALs, i.e. broad, blue-shifted absorption features associated with several strong, high-ionization lines in the UV region. More formally, we calculate BIs of 11400 km s^{-1} ($i = 75^\circ$) and 9900 km s^{-1} ($i = 80^\circ$) for the C IV lines we predict for these two sightlines. Thus observed systems displaying these spectra would certainly be classified as bona-fide BALQSOs. This is one of the main results of our work.

Many of the transitions exhibiting BAL features in the spectra for both sightlines also show a redshifted emission component, forming the other part of the classic P-Cygni profile seen in such sources. This can be interpreted as the red-shifted part of a classic rotationally broadened, double-peaked emission line, which is produced primarily by




Figure 6.4: The electron density (top left), electron temperature (top right), poloidal velocity (mid left), rotational velocity (mid right), ionization parameter (bottom left) and proportion of C ions in the C IV ionization state (bottom right) for the benchmark model. Only the positive xz plane is shown, the wind is rotationally symmetrical around the z axis. Note the logarithmic scales and the difference in scales for the x and z axes. The location of cells for which ‘cell’ spectra are presented in figure 6.5 are shown in the ionization parameter plot and the black lines on the C IV plot show sightlines through the wind to the origin used to produce the spectra plotted in figure 6.6.




Figure 6.5: Mean specific intensity in three cells. The red line is a cell with C IV as the dominant Carbon ionization stage, in the base of the wind. The green line is for a cell with significant C IV in a fast moving part of the wind, and the blue line is for a cell at the top of the wind. fully exposed to the ionizing continuum. The thick lines represent the unobscured spectrum that would be seen in that region. The vertical line marks the location of the C IV photoionization edge.

scattering in the base of the wind. The blue-shifted part of the line profile is not visible, since it is superposed on and/or absorbed by the BAL feature.

Another interesting feature of the BAL spectra is that the Si IV absorption feature is narrower than the C IV feature for all sightlines. This is due to the lower ionization potential of the silicon ion, meaning that it is produced in a more limited part of the wind. The relative strengths of features is broadly in agreement with observations (Gibson et al. 2009) and demonstrates that, at least to first order, we are predicting the correct ionization state in the BAL-forming portion of the wind.

The continuum for both sightlines through the wind is suppressed relative to the unobscured SED (see Figure 6.6). This attenuation of the continuum is predominantly due to electron scattering. The optical depth to electron scattering through the wind is shown in Figure 6.7, along with the corresponding attenuation factor, $F/F_0 = e^{-\tau_{es}}$.




Figure 6.6: Simulated spectra for four sightlines. The top panel shows the 40° sightline, over the top of the wind to the brightest parts of the accretion disk, whilst the next panel shows the 75° sightline, looking through the upper parts of the wind. The next panel is for 80° looking through the lower part of the wind and finally, the bottom panel shows the 85° sightline which is almost equatorial, and views the bright central source through the very base of the wind. For each sightline, the unobscured continuum is also plotted for comparison along with a scaled continuum. The scaled continuum is the unobscured continuum scaled to equal the simulated spectrum away from line features, and therefore takes account of electron scattering. The location of some of the spectral features most relevant for BALQSOs are marked.

For example, at $i = 75^\circ$ our model gives $\tau_{es} \simeq 0.3$ and $F/F_0 \simeq 0.7$, in agreement with Figure 6.6.

Detailed C IV line profiles predicted by our model for a more finely spaced grid of sightlines are shown in Figure 6.8, along with the BI calculated for each profile. BALs (i.e. features with $BI > 0$) are observed for inclinations $73^\circ < i < 83^\circ$, which represents $\simeq 17\%$ of all possible sightlines.

None of the BALs are completely black. Examination of the history of the photon packets contributing to the flux seen at the bottom of the lines reveals that all have scattered at least once, which means that this radiation could exhibit significant polarisation. Observations of BALs do indeed find significant polarisation in the troughs (e.g [DiPompeo et al. 2011](#)).

6.4.2.2 Sightlines looking above the wind cone

We now consider the $i = 40^\circ$ sightline, which views the central engine from “above” the wind cone. In a pure geometric unification scenario, this sightline may be expected to produce a classic Type I QSO spectrum. What we observe from the simulation is a slightly enhanced continuum, along with some broad, but weak emission lines (as in [Sim et al. 2008](#)).

The continuum enhancement is mainly due to electron scattering *into* this line of sight. As we have seen in Section 6.4.2.1 and Figure 6.7, the base of the wind is marginally optically thick to electron scattering in directions along the disk plane. Thus photons scattering in this region tend to emerge preferentially along the more transparent sightlines perpendicular to this plane, and the wind essentially acts as a reflector.

More importantly, the emission lines superposed on the continuum do correspond to the typical transitions seen in Type I QSOs, but they are weaker than the observed features. For example, the C IV emission line in our model has an equivalent width of only 1.4 Å. By contrast, the equivalent width of the C IV line in a typical QSO with a continuum luminosity of $L_\lambda(1550\text{\AA}) \simeq 10^{43}$ ergs s $^{-1}$ is $\simeq 60$ Å ([Xu et al. 2008](#)). Other sightlines above the wind cone ($0^\circ \leq i \leq 70^\circ$) yield qualitatively similar spectra. Overall, the

presence of the right “sort” of emission lines is encouraging, but their weakness is an obvious shortcoming of the model. We will discuss the topic of line emission in more detail in Section 6.5.2.

6.4.2.3 Sightlines looking through the base of the wind cone

Let us finally consider the highest inclinations, which correspond to sightlines that do not lie along the wind cone, but for which the UV continuum source (i.e. the central region of the accretion disk) is viewed through the dense, low-velocity base of the outflow. It is not obvious what type of QSO one might expect to see from such sightlines, since in the standard model of QSOs, one might expect the torus to obscure such sightlines.

Figure 6.7 shows that the $\tau_{es} > 7$ for this inclination, and indeed we find that essentially all of the radiation emerging from the model in this direction has been scattered several times within the wind before ultimately escaping along this vector. This is in line with results presented by [Sim et al. \(2010a\)](#), who found that, for Compton-thick winds, AGN spectra at high inclination are dominated by scattered radiation. Recent observational work (e.g. [Treister et al. 2009](#)) has demonstrated that there is a population of AGN in the local universe where the X-ray source is completely obscured, and the X-ray spectrum is dominated by reflection. These high inclination sightlines could represent these so called “reflection dominated Compton thick AGN”.

The shape of the predicted emission features is dominated by the rotational velocity field in the wind. Taking C IV as an example and examining Figure 6.4, we see that the fractional abundance peaks in a region where the rotational velocity is significantly higher than the outflow velocity. Scattering from this region (along with any thermal emission) will produce the double-peaked line profile that is characteristic of line formation in rotating media (e.g. [Smak 1981; Marsh and Horne 1988](#)). The blue part of this profile is then likely suppressed, since it is superposed on and/or absorbed by blue-shifted absorption in the base of the wind.




Figure 6.7: Electron scattering optical depth through the wind towards the origin. The heavy line is the expected attenuation.

6.5 Discussion

The results shown in the previous section confirm that a simple disk wind model can produce the characteristic BAL features seen in about $\simeq 20\%$ of QSOs. However, they also reveal some significant shortcomings of the model. In this section, we will consider some of the key questions raised by our results in more detail and point out promising directions for future work.

6.5.1 The Sensitivity of BAL Features to X-ray Luminosity

As noted in Section 6.3.7, our benchmark model is rather X-ray weak compared to most real (BAL)QSOs. To quantify this, Figure 6.9 shows the viewing angle dependence of α_{OX} for our benchmark model. Clearly, α_{OX} (i.e. the X-ray to optical ratio) is a strong function of inclination. Two effects are responsible for this.




Figure 6.8: C IV line detail for sightlines between 70° and 84° . Each spectrum is scaled to a linear continuum, fitted at velocity $\pm 30000 \text{ km s}^{-1}$. The BALnicity Index for the spectrum is marked on each plot.

First, the UV and optical fluxes are dominated by the accretion disk and thus decline steeply with inclination, due to foreshortening and limb-darkening. By contrast, the X-ray source is assumed to emit isotropically (although the disk can obscure the “lower” hemisphere). Thus, in the absence of any outflow, α_{OX} increases with inclination, as the X-rays become relatively more important. This is the trend marked by the solid line in Figure 6.9.

Second, sightlines to the central engine that cross the wind cone can be affected by absorption and radiative transfer in the outflow. In particular, bound-free absorption

preferentially suppresses the X-ray flux along such sightlines in our model, and hence reduces α_{OX} . The crosses in Figure 6.9 show the predicted values of α_{OX} taking both effects into account. This shows that the suppression of α_{OX} is strongest for sightlines looking directly into the wind, i.e. the same viewing angles that produce BALs.

For our benchmark model, we find that $\alpha_{OX} = -2.44$ for $i = 40^\circ$ and $\alpha_{OX} = -2.89$ for $i = 75^\circ$. If these viewing angles can be taken to represent non-BAL QSOs and BALQSOs respectively, we can compare these values to those expected from the observed scaling relation between $L_{2500\text{\AA}}$ and α_{OX} found by [Just et al. \(2007\)](#) for non-BAL QSOs,

$$\alpha_{OX} = (-0.140 \pm 0.007) \log(L_{2500\text{\AA}}) + (2.705 \pm 0.212). \quad (6.8)$$

Taking the values of $L_{2500\text{\AA}}$ for the two sightlines from our benchmark model, we obtain empirical predictions of $\alpha_{OX} = -1.66 \pm 0.43$ for $i = 40^\circ$ and $\alpha_{OX} = -1.52 \pm 0.43$ for $i = 75^\circ$. However, we still have to correct for the fact that BALQSOs are known to be X-ray weak (presumably due to absorption in the outflow). [Gibson et al. \(2009\)](#) find that the median $\Delta\alpha_{OX} = \alpha_{OX,BALQSO} - \alpha_{OX,QSO} = -0.17$ for all BALQSOs with $BI > 0$, but the distribution of $\Delta\alpha_{OX}$ is quite wide, and systems with large $BI \gtrsim 1000 \text{ km s}^{-1}$ tend to have larger $\Delta\alpha_{OX}$. If we adopt $\Delta\alpha_{OX} \simeq -0.3$ as typical for strong BALQSOs (like our benchmark model), the empirically predicted value for the $i = 75^\circ$ BALQSO sightline becomes $\alpha_{OX} \simeq -1.8 \pm 0.4$ (where the quoted uncertainty is purely that associated with the scaling relation for non-BAL QSOs).

These numbers make the X-ray weakness of the benchmark model clear: at fixed $L_{2500\text{\AA}}$, the emergent X-ray flux predicted by the model is at least 2 orders of magnitude lower than the average observed values. Equivalently, the model corresponds to a $2-3\sigma$ outlier in the α_{OX} distribution of (BAL)QSOs.

To investigate whether we can increase L_X substantially while retaining the observed BAL features for sightlines looking into the wind cone, Figure 6.10 shows the behaviour of the emergent spectra for $i = 75^\circ$ and $i = 80^\circ$ as the X-ray luminosity is increased from $L_X = 10^{43} \text{ erg s}^{-1}$ to $L_X = 10^{44} \text{ erg s}^{-1}$. We will focus on the behaviour of C IV and O VI as representative examples. For $i = 75^\circ$, the C IV BAL feature becomes




Figure 6.9: Emergent α_{OX} values for benchmark model (symbols) and without wind (solid line)

narrower with increasing L_X : it gradually disappears, starting from the blue edge. The feature is lost entirely by $L_X = 5 \times 10^{43}$ ergs s^{-1} . The effect on the O VI feature is more abrupt, with almost no effect up to a luminosity of 3.7×10^{43} ergs s^{-1} , but complete disappearance for twice that luminosity. Similarly, for $i = 80^\circ$, C IV weakens gradually in models with larger L_X , while in this case O VI shows hardly any response to L_X until we reach $L_X = 10^{44}$ ergs s^{-1} . Note that C IV is still present with a clear BAL signature at this X-ray luminosity for $i = 85^\circ$.

The strong sensitivity of the UV resonance transition to L_X in this range can be understood by considering the location of the photoionization edges for these species. Figure 6.11 shows these locations relative to the angle-averaged disk spectrum and the X-ray spectra corresponding to $L_X = 10^{43}$ erg s^{-1} and $L_X = 10^{44}$ erg s^{-1} . The ionization edge of C IV falls very close to the frequency above which the X-ray source takes over from




Figure 6.10: Spectra at 75° (top panel) and 80° sightlines for the benchmark model with a range of 2-10 keV luminosities

the disk as the dominant spectral component. In fact, for $L_X = 10^{43} \text{ erg s}^{-1}$, photo-ionization of C IV is driven primarily by disk photons, while for $L_X = 10^{44} \text{ erg s}^{-1}$, it is mainly due to photons from the X-ray source. The photo-ionization rate of higher-ionization species like O VI is dominated by the X-ray source even at $L_X = 10^{43} \text{ erg s}^{-1}$.

We conclude that moderate increases in L_X (by a factor of a few) would be feasible without any other parameter adjustments, but larger changes will begin to destroy BAL features for more and more sightlines. However, it is reasonable to anticipate that the negative effects of increasing L_X on BAL features may be mitigated by simultaneously increasing the mass-loss rate in the wind, \dot{M}_{wind} . We have therefore carried out a 2-dimensional sensitivity study, in which we vary both L_X and \dot{M}_{wind} . The results – focussing particularly on the key C IV line profile for $i = 75^\circ$ – are shown in Figure 6.12. In each panel of this figure, we also give the BI measured for the profile shown, as well as values of F/F_0 ; the ratio of the continuum at -30000 km s^{-1} to that obtained in a model without a wind.

The second column in Figure 6.12 illustrates the effect of increasing L_X while keeping the mass-loss rate at the benchmark value of $\dot{M}_{wind} = 5M_{\odot} \text{ yr}^{-1}$ (as in Figure 6.10). We see again that the C IV BAL feature persists up to at least $L_X = 10^{44} \text{ erg s}^{-1}$, although it becomes both weaker and narrower with increasing L_X . For higher mass-loss rates (third and beyond columns in Figure 6.12), the C IV profile becomes increasingly insensitive to L_X . So, as expected, increasing \dot{M}_{wind} allows for the production of strong BAL features in species like C IV even at higher irradiating L_X . Thus one way to address the X-ray weakness of the benchmark model (i.e. increase α_{ox}) is to increase both L_X and \dot{M}_{wind} .

This remedy has a side effect: any increase in \dot{M}_{wind} also implies an increase in τ_{es} , the electron-scattering optical depth through the outflow, which reduces the *observed* X-ray and optical luminosity of the BALQSO. For moderate inclinations and mass-loss rates, the emergent continuum flux will scale simply as $e^{-\tau_{es}}$, although for the highest inclinations and \dot{M}_{wind} , scattered radiation will dominate the continuum and break this scaling. In any case, our main conclusion is that, for high mass-loss rates, we are able to produce BALs over a wider range of L_X but at the expense of significant continuum suppression. If this is to provide a realistic route towards geometrical unification, it implies that the *intrinsic* luminosities of BALQSOs must be higher than suggested by their observed brightness.

There are, of course, other possibilities. For example, if the outflow is clumpy (e.g. Kirolik et al. 1981; Arav et al. 1999b), some of the X-ray radiation may pass through the wind unimpeded, thus increasing the observed L_X for sightlines inside the wind cone. Such a solution would have the added advantage that the higher density within the clumps would naturally result in a lower ionization state for a given X-ray flux. Relaxing the assumption of a smooth flow would, of course, take us back in the direction of cloud models of the BLR and BALR. We plan to investigate the pros and cons of such scenarios in future work.

6.5.2 The Weakness of Line Emission Produced by the Outflow

As mentioned in Section 6.4.2.2, our benchmark model does not produce strong emission line features, especially at low inclination models where the geometric unification scenario




Figure 6.11: Input spectrum from the disk component in the benchmark model, together with the contribution from the central source for two different values of L_X . The photoionization cross sections for C IV, O VI and Si IV are also plotted over the same frequency range.

suggests Type I QSOs should be seen. This suggests that modifications are required in some aspect of the structure or the physics that is incorporated into PYTHON, if the geometric unification scenario is correct.

In considering this shortcoming, it is interesting to compare with the disk wind modelling carried out by Murray et al. 1995 [hereafter MCGV95]. They explicitly model a line-driven outflow, and the broad characteristics of their disk wind are sufficiently similar to our benchmark model to make a comparison interesting. MCGV find that their outflow produces BAL troughs when viewed at high inclinations. However, they also find that their model produces sufficient collisionally excited emission in species like C IV to explain the BEL in both BALQSOs and non-BAL QSOs. In their model, this C IV emission arises in a layer of the wind close to the disk plane where C IV is the dominant ionization stage of carbon. They argue that this region extends from $\simeq 10^{16}$ cm to $\simeq 10^{17}$ cm and is characterized by $T_e \simeq 20,000$ K and $n_e \simeq 10^{10}$ cm $^{-3}$. It is the integrated, collisionally

excited emission from this region that dominates the C IV line flux in their model.

In the case of our benchmark model, Figure 6.4 shows that the fractional C IV abundance is high ($f_{\text{C IV}} \gtrsim 0.1$) in a region extending from $x \simeq 5 \times 10^{16}$ cm to $x \simeq 3 \times 10^{17}$ cm. However, the density in this region is considerably lower than in the C IV line-forming region considered by MCGV: $10^7 \text{ cm}^{-3} \lesssim n_e \lesssim 10^9 \text{ cm}^{-3}$. It is this lower density that almost certainly explains the difference in C IV emissivity in the models. In a later version of their model (Murray and Chiang 1998), the production of C IV is dominated by lower density ($n_e \simeq 10^8 \text{ cm}^{-3}$) material lying at larger radii $r \gtrsim 10^{18}$ cm. In these models, it is the large emitting volume that explains the strength of collisionally excited C IV feature.

These comparisons suggest that moderate changes to the benchmark model may be enough to produce BELs via collisionally excited line emission in the outflow. It is likely that such changes would go in the direction of increasing the density in some parts of the wind, either by modifying existing wind parameters or by introducing clumpiness. In fact, we can clearly see that the emission component of the C IV line in Figure 6.12 does increase with increasing \dot{M}_{wind} . However, even our highest \dot{M}_{wind} models do not yet produce enough emission at low inclination angles. Thus, while the benchmark model shows some promise in the context of geometric unification, it remains to be seen whether (and how) it can be modified to produce the strong emission lines seen in both BALQSOs and non-BAL QSOs. This will be the subject of a separate paper.

Before we leave this section, it is interesting to ask why the *scattering* of C IV photons in the outflow does not produce emission lines of sufficient strength in our model. In CV winds, for example, modelling suggests that resonance scattering is the primary mechanism for producing the observed emission line profiles (Shlosman and Vitello 1993; Knigge et al. 1995, LK02, Noebauer et al. 2010). The critical factor in understanding the relative importance of scattered photons to the line formation process is the combination of foreshortening, limb-darkening, outflow orientation and covering factor. The outflows in CVs are thought to be ejected roughly perpendicular to the accretion disk, albeit over quite a large range of opening angles. Moreover, CV winds emerge directly above the continuum emitting regions of the disk, whereas BALQSO winds cover only $\simeq 20\%$ as




Figure 6.12: Spectrum detail around the C IV line feature for the benchmark model and others with a range of X-ray luminosity and \dot{M}_{wind} . Spectra are all for $i = 75^\circ$ and are normalised to a linear continuum model fitted at $\pm 30000 \text{ km s}^{-1}$. Values of F/F_0 show the ratio of the continuum at -30000 km s^{-1} to that obtained in a model without a wind. The scales for each subplot are given on the upper right hand subplot and the baseline model is highlighted.

seen from the continuum source. Thus, in a CV, an outflow that is optically thick in C IV over some velocity interval will intercept virtually all of the continuum photons within this interval – and scatter them into other sightlines, where they emerge as apparent line emission. Conversely, our QSO wind scatters only a small fraction of the continuum photons, even in frequency intervals in which the wind is optically thick. As a result, scattered photons will not contribute as much to emission line formation.

6.5.3 Ionization Stratification and Reverberation Mapping

In our benchmark model, the ionization state of the outflow tends to increase with increasing distance along poloidal streamlines (Figure 6.4). This observation is easily understood. Far away from the central engine, in the region of the flow where the BAL features are produced, the wind geometry approximates an optically thin, spherical outflow. In such an outflow, $U \propto v(r)$, so as long as the wind accelerates, the ionization parameter must also increase outwards.

This trend may appear to contradict reverberation mapping studies of the BLR, which suggest ionization stratification in the opposite sense. For example, the C IV BLR radius at fixed UV luminosity inferred by [Kaspi et al. \(2007\)](#) is $\simeq 2 \times$ smaller than that obtained by [Kaspi et al. \(2005\)](#) for H β . This sense of ionization stratification is also supported by recent microlensing results (e.g. [Guerras et al. 2013](#)).

However, our model is not necessarily in conflict with these empirical findings. As discussed in Section 6.5.2, it is clear that the parts of the wind producing BALs (where U increases outwards) are unlikely to be co-spatial with the parts of the wind producing BELs. More specifically, only the dense base of the disk wind is likely to produce significant amounts of collisionally excited line emission. Thus, in the context of disk wind models, the reverberation mapping results are unlikely to probe the “radial” ionization structure of the outflow, but may instead probe the stratification along the base of the wind in the x direction. Here, the ionization parameter can (and does) decrease outwards in our disk wind models.

6.6 Implications

Here we consider implications of our benchmark model for AGN feedback in galaxy evolution and identifying possible driving mechanisms for the wind.

In energetic terms, effective feedback – sufficient to establish the $M_{BH} - \sigma$ relation, for example – seems to require $L_k/L_{bol} \simeq 0.005 - 0.05$ (Di Matteo et al. 2005; Hopkins and Elvis 2010). Here, $L_k = \frac{1}{2}\dot{M}_{wind}v_\infty^2$ is the kinetic luminosity of the outflow. Our benchmark model is characterized by $v_\infty \simeq 20,000 \text{ km s}^{-1}$ and $\dot{M}_{wind} = 5 \text{ M}_\odot \text{ yr}^{-1}$, which yields $L_k/L_{bol} \simeq 0.025$. Thus an outflow of this type would be capable of providing significant amounts of energy-driven feedback. The kinetic luminosity of the benchmark model is also broadly in line with the empirical scaling between L_{bol} and L_k found by King et al. (2013) for black holes across the entire mass-scale, from compact binary systems to AGN/QSO.

King (2003, 2005, 2010) has argued that QSO outflows interact with their host galaxy in a *momentum*-, rather than energy-driven fashion. In this case, the strength of the feedback an outflow can deliver depends on its momentum flux, $\dot{M}_{wind}v_\infty$. This scenario naturally produces the observed $M_{BH} - \sigma$ relation if the outflows responsible for feedback satisfy $\dot{M}_{wind}v_\infty \simeq L_{Edd}/c$; this is the single-scattering limit for momentum transfer in a radiatively-driven outflow from a QSO radiating at L_{Edd} . For our benchmark model, we find $\dot{M}_{wind}v_\infty \simeq 0.7L_{bol}/c \simeq 0.1L_{Edd}/c$. In other words, the momentum flux in our model is close to the single-scattering limit for its bolometric luminosity. Since black hole growth and cosmological feedback are thought to be dominated by phases in which $L_{bol} \simeq L_{Edd}$ (Soltan 1982; Yu and Tremaine 2002; Di Matteo et al. 2005), outflows of this type would probably also meet momentum-based feedback requirements.

The momentum flux in the outflow is also a key parameter for any disk wind driven by radiation pressure, such as the line-driven winds considered by MCGV95 and Proga et al. (2000). In particular, such winds typically satisfy the single-scattering limit for momentum transfer from the radiation field to the outflow, $\dot{M}_{wind}v_\infty \leq L_{bol}/c$. As noted above, our benchmark model also satisfies this limit. However, this is somewhat misleading: our disk wind subtends only $\simeq 20\%$ of the sky as seen from the central

source, and it intercepts an even smaller fraction of the QSOs bolometric luminosity (due to foreshortening and limb-darkening of the disk radiation field). The momentum flux in our benchmark model actually exceeds the single-scattering limit if only momentum carried by photons that actually intercept the wind is taken into account. Thus, in the context of radiatively-driven winds, either the mass-loss rate in our model must be overestimated or multiple scattering effects must be important. The latter has been suggested for quasi-spherical Wolf-Rayet star winds (Lucy and Abbott 1993; Springmann 1994; Gayley 1995) but is much more challenging for the biconical wind considered here. Alternatively, the winds of QSOs may not be driven (exclusively) by radiation pressure, with magnetic/centrifugal forces providing the most obvious alternative (e.g. Blandford and Payne 1982; Emmering et al. 1992; Proga 2003).

Our discussion of all these considerations is preliminary: the benchmark model is far from perfect, and we have not yet carried out a comprehensive survey of the available parameter space. However, the first column in Figure 6.12 shows that a reduction in the mass-loss rate from the benchmark value of $\dot{M}_{wind} = 5M_{\odot} \text{ yr}^{-1}$ to $\dot{M}_{wind} = 1M_{\odot} \text{ yr}^{-1}$ virtually destroys the BAL features, even for $L_X = 10^{43} \text{ erg s}^{-1}$. Thus, at least for our adopted wind geometry and kinematics, it seems that any outflow capable of producing BAL features is also likely to provide significant feedback and pose a challenge to line-driving.

6.6.1 Outlook

In our view, the most significant shortcomings of the benchmark model are its intrinsic X-ray weakness and its inability to produce collisionally excited BELs. Both problems might be resolved by moderate changes in wind parameters, so an important step will be to explore the parameter space of our model. We plan to adopt two complementary approaches in this: (i) systematic grid searches; (ii) targeted explorations guided by physical insight. The latter approach is important, since compromises will have to be made in the former: a detailed and exhaustive search over all relevant parameters is likely to be prohibitive in terms of computational resource requirements. As an example

of how physical insight can guide us, we note that both the X-ray and emission line problems might be resolved by increasing the density in specific wind regions.

6.7 Summary

This work represents the first step in a long-term project to shed light on the nature of accretion disk winds in QSOs. These outflows may be key to the geometric unification of AGN/QSO and might also provide the feedback required by successful galaxy evolution scenarios.

In this pilot study, we have focussed on the most obvious signature of these outflows, the broad, blue-shifted absorption lines seen in $\simeq 20\%$ of QSOs. We have constructed a kinematic disk wind model to test if it can reproduce these features. This benchmark model describes a rotating, equatorial and biconical accretion disk wind with a mass-loss rate $\dot{M}_{wind} \simeq \dot{M}_{acc}$.³ A Monte Carlo ionization and radiative transfer code, PYTHON, was used to calculate the ionization state of the outflow and predict the spectra emerging from it for a variety of viewing angles. Our main results are as follows:

- (i) Our benchmark model succeeds in producing BAL-like features for sightlines towards the central source that lie fully within the wind cone.
- (ii) Self-consistent treatment of ionization and radiative transfer is necessary to reliably predict the ionization state of the wind and the conditions under which key species like C IV can efficiently form.
- (iii) The benchmark model does not produce sufficient collisionally excited line emission to explain the BELs in QSOs. However, we argue that moderate modifications to its parameters might be sufficient to remedy this shortcoming.
- (iv) The ionization structure of the model, and its ability to produce BALs, are quite sensitive to the X-ray luminosity of the central source. If this is too high, the

³ Strictly speaking, this is not entirely self-consistent, since the presence of such an outflow would alter the disk's temperature structure. However, we neglect this complication, since the vast majority of the disk luminosity arises from well within our assumed wind-launching radius.

wind becomes overionized. In our benchmark model, L_X is arguably lower than indicated by observations. Higher values of L_X may require higher outflow columns – e.g. via higher \dot{M}_{wind} – in order to still produce BALs.

- (v) For our adopted geometry and kinematics, $\dot{M}_{wind} \gtrsim \dot{M}_{acc}$ is required in order to produce significant BAL features. The kinetic luminosity and momentum carried by such outflows is sufficient to provide significant feedback.

Chapter 7

RT simulations for a hydrodynamically generated wind model

7.1 Introduction

As discussed in Chapter 3, several mechanisms have been proposed to produce the accelerating force for disk winds in AGN, including gas/thermal pressure, magnetocentrifugal forces and radiation pressure acting on spectral lines (“line-driving”). All of these mechanisms have been shown to be capable of producing outflows in suitable conditions. However, line driving is particularly attractive, since the requisite force in this case arises naturally when (primarily UV) photons produced by the central engine are scattered by strong resonance lines. Many such lines are observed directly in the UV spectra of (BAL)QSOs, so line driving must certainly be acting in some measure.

Following the kinematic disk wind simulations presented in the last chapter, it seemed natural to carry out simulations for a more realistic physical model. In particular, a hydrodynamic model such as that presented by [Proga and Kallman \(2004, hereafter PK04\)](#) naturally forms higher density “knots”, so it is interesting to ask whether such high density regions could produce enhanced emission line strengths. Moreover, the RT




Figure 7.1: The density (colours) and poloidal velocity (arrows) structure of the PK04 model. Radial lines delineate the three zones described in the text.

treatment in the original simulation was necessarily simplified, so we were also interested in seeing whether a full radiative transfer and ionization simulation in PYTHON would give the same ionization structure as that found in PK04.

The PK04 wind model (introduced in Section 3.2.1.2) has already been subjected to two “post-processing” radiative transfer calculations of the type we are proposing here. Both of these were concerned with the impact of the predicted outflow on the X-ray spectra of AGN/QSO. First, Schurch et al. (2009) performed a 1-D simulation which

found that the wind produced observable X-ray spectral features. Second, [Sim et al. \(2010b\)](#), hereafter SP10) carried out a multi-dimensional radiative transfer simulation in which the ionization state and temperature structure of the wind were self-consistently computed. They confirmed that the wind was able to imprint a variety of characteristic features into the X-ray spectra of AGN. However, they also showed that scattering in the (failed) outflow is critically important in setting the ionization state of the wind in regions that would otherwise be shielded from the central engine.

Here, we therefore carry out a full, multi-dimensional radiative transfer calculation for the PK04 model that allows us to predict the ionization state, temperature structure and emergent UV spectra for this outflow. This extends the work done by SP10 to the longer wavelengths and lower ionization stages that are critical for the line-driving mechanism and for the formation and the formation of BALs.

7.2 Method

In this section, we will review the geometry and kinematics of the outflow computed by PK04. We also describe the form of the illuminating spectrum we adopt and discuss how it relates to that used in PK04 and SP10. We use version 77 of PYTHON for the calculations.

7.2.1 The PK04 disk wind geometry

We take as our input the hydrodynamic wind model computed by PK04 at time-step 955. This represents a time sufficiently far from the simulation start that the wind geometry no longer depends on initial conditions. In addition, at this point, the flow is only weakly time-dependent, so this time-step is broadly representative of the flow as a whole. The figures in PK04 also focused on this time-step, as did both SP10 and [Schurch et al. \(2009\)](#). The hydrodynamic calculations in PK04 were carried out using ZEUS ([Stone and Norman 1992](#)), modified to include the effect of the radiation field on the driving force and the temperature balance of the wind. The model is illustrated in Figure 7.1, with

the poloidal component of velocity field superimposed upon the density structure. PK04 noted three main regions (delineated on Figure 7.1 by radial lines):

- A: A hot, low-density, infalling region near the pole.
- B: A hot transition zone in which the wind struggles to escape.
- C: A warm, dense, equatorial flow that is shielded from the central source by a high-density ‘failed wind’ region close to the origin

The ZEUS grid is polar, logarithmic in both the r - and θ -direction, with finer discretisation close to the disk and the origin. To carry out the radiative transfer calculation, we create an identical grid and import the density and velocity structure of the ZEUS model. The density in each of our grid cells is constant and taken to be the value at the centre of each matching grid cell in the hydrodynamic model, whilst the velocity field in each cell is interpolated across each cell based upon the values at the vertices in the original model. This discretization of the velocity field lends itself more naturally to our use of the Sobolev approximation in the treatment of line scattering.

Figure 7.2 shows the same density information as Figure 7.1, but on logarithmic axes in order to highlight the dense parts of the wind closer to the disk plane. A wedge-like region at the base of the wind clearly visible - the disk atmosphere. In our code, photons emitted by the accretion disk start from the disk plane and therefore would have to make their way through this atmosphere before entering the rest of the wind. Such dense regions, in which photons scatter a very large number of times, are extremely computationally expensive in a Monte-Carlo code. We therefore remove this part of the flow from the computation, permitting photons to fly unimpeded from the disk into the lower part of the wind. This is a reasonable approximation since, to first order, the dense disk atmosphere would be expected to thermalise to the same temperature as the disk beneath it. Hence the upper surface of that region would be expected to emit photons with the same spectral distribution as our disk.

The dashed line in Figure 7.2 shows the angle at which we make this cut - $\theta = 89.5^\circ$. The temperature of the lower surface of the flow where it is primarily irradiated by photons




Figure 7.2: The density of the PK04 model, plotted on logarithmic axes. The solid lines show the boundaries between the three zones described in the text, and the dashed line at $\theta = 89.5^\circ$ is where we truncated the wind in our representation of the model.

from the disk is approximately equal to the disk below it. This is the same temperature as would be expected if the surface was in direct contact with the upper surface of an extended disk atmosphere. We are therefore confident that this simplification does not significantly affect the results obtained for the bulk of the wind.

7.2.2 Illuminating spectrum

Given the density and velocity structure of the wind, we next need to specify the parameters of the radiation sources in our model, i.e. the accretion disk and X-ray-emitting central ‘corona’. To permit comparison between our results and those of PK04 and SP10, we match the parameters used in those studies as closely as possible.




Figure 7.3: The input spectrum using the simulation, together with the 10 keV bremsstrahlung spectrum assumed in the original PK04 calculation. Also shown are the composite spectra from Laor et al. (1997) and Zheng et al. (1997) used to inform the luminosities used in PK04. The vertical lines illustrate the positions of various band boundaries discussed in the text

PK04 assumed the presence of a $10^8 M_\odot$ supermassive black hole (BH) at the centre of their accretion disk and adopted an accretion rate on the BH of $1.8 M_\odot \text{yr}^{-1}$. This rate gives a disk luminosity of $L_D = 6.5 \times 10^{45} \text{ ergs s}^{-1}$ assuming an efficiency of $\simeq 6\%$. In their hydrodynamic simulation, PK04 split the illuminating luminosity into two components. First, they defined an ‘ionizing’ luminosity (L_X) that was used to compute the temperature and ionization parameter throughout the wind. Second, they defined a ‘UV’ luminosity (L_{UV}) that was used to compute the actual line-driving force. However, in regions of the wind that are shielded from the X-ray source by optically thick material, the wind temperature was set to that of the disk below the wind, thus implicitly assuming that these parts of the wind would be heated and ionized by disk photons.

The balance between the two components was motivated by the composite spectra presented in [Laor et al. \(1997\)](#) and [Zheng et al. \(1997\)](#); the radio quiet composite is shown in Figure 7.3. The ‘ionizing luminosity’, L_X , defined by PK04 as the integrated luminosity of the central source above 13.6 eV (912 Å), was taken to be 10% of L_D , or 6.5×10^{44} ergs s $^{-1}$. The ionization parameter at any point in the flow was then computed from the local X-ray flux by assuming $1/r^2$ geometric dilution as well as $e^{-\tau_{es}}$ attenuation, where τ_{es} is the optical depth towards the origin due to electron scattering. This ionization was then used to calculate the line-driving force using the analytic formulae of [Stevens and Kallman \(1990\)](#). These formulae assume that the frequency distribution of ionizing photons takes the form of a 10 keV bremsstrahlung spectrum with no low-frequency cutoff. The calculations of the wind temperature structure in PK04 used the same characteristic temperature for the X-ray source in order to ensure self-consistency. A bremsstrahlung spectrum of this form, normalised to give the correct value of L_X is also plotted on Figure 7.3.

The UV luminosity in PK04 was defined as the luminosity in a band running from 200 to 3000 Å and was set to 90% of L_D . It should be noted that there is a degree of inconsistency here - a significant part of the ‘UV’ band is actually above the 13.6eV lower limit of the ‘ionizing’ band, and we see from Figure 7.3 that the composite spectrum carries significant ionizing luminosity that would contribute to L_X , as well as L_{UV} . However, increasing L_X to take account of this would mean that the luminosity around the peak of the bremsstrahlung spectrum would be much too high. This illustrates the difficulty of approximating an observed SED with a highly simplified model.

In our current calculation, we use a thin accretion disk with the same BH mass and accretion rate adopted by PK04, even though such a model may not produce the best match to observations (e.g [Davis and Laor 2011](#); [Done et al. 2012](#); [Slone and Netzer 2012](#)). We also include a power-law X-ray source assumed to arise in a central spherical ‘corona’. This X-ray source is defined by the same photon index, Γ , and luminosity in the 2-10 keV range, L_{2-10} , that was used by SP10. In that work, the X-ray parameters were chosen so as to give a reasonable match to observations whilst maintaining the same luminosity in the 2-10 keV range as the X-ray parameters in PK04. More specifically, we

Parameter	Value
M_{BH}	$10^8 M_{\odot}$
\dot{M}_{acc}	$1.8 M_{\odot} yr^{-1} \simeq 0.5 \dot{M}_{Edd}$
Γ_X	2.1
$L_{2-10 eV}$	$2.7 \times 10^{44} \text{ ergs s}^{-1}$
$r_{disk}(\text{min}) = r_X$	$6r_g = 8.8 \times 10^{13} \text{ cm}$
$r_{disk}(\text{max})$	$2700r_g = 4 \times 10^{16} \text{ cm}$

Table 7.1: The key parameters adopted in the radiative transfer simulation.

follow SP10 in adopting $L_{2-10} = 2.7 \times 10^{44} \text{ ergs s}^{-1}$ and $\Gamma_X = 2.1$. Moreover, following PK04, we set the radius of the X-ray source, r_X , equal to the innermost radius of the accretion disk. With these choices, our SED yields $L_X \sim 0.7L_D$ and $L_{UV} \sim 0.9L_D$ in the notation of PK04. Note the two components sum to more than unity because we count ionizing photons from the disk in both L_X and L_{UV} . Figure 7.3 shows that our SED provides a reasonable match to the composite spectrum derived from observations and also produces roughly the same luminosity at high energies as the bremsstrahlung spectrum. However, the ionizing luminosity in our model is much higher than that in the PK04 simulation, since we allow both disk and central source photons to contribute to the ionizing flux. We return to this point in Section 7.4.1, where we explicitly check the impact of this difference on the calculated ionization state of the wind. Our adopted parameters are summarised in Table 7.1.

7.3 Results

We are now ready to present the key properties of our converged wind model. We will first illustrate the ionization state of the wind and then present the simulated emergent spectra. Finally, we will also show the temperature structure and ionization parameter throughout the outflow.

7.3.1 Ionization state

We begin by considering the ionization state of carbon, since strong, blueshifted absorption associated with C IV 1550 Å is a defining feature in the UV spectra of essentially all BALQSOs. However, as illustrated in Figure 7.4, carbon is actually fully ionized




Figure 7.4: The mean charge of carbon in the model.

throughout much of the outflow. In the very dense 'failed wind' at the base of the wind, carbon is in a much lower ionization state. This is partly due to the high density in this region, but also to photoelectric absorption in the outer skin of the wind, which removes a large proportion of the soft X-ray photons responsible for ionizing carbon. In PK04, the warm outflow behind the failed wind also remained in a low ionization state. However, in our calculation, the carbon in this area is fully ionized. This key difference is due to scattered and reprocessed radiation and is discussed in more detail in section 7.4. Figure 7.4 also illustrates the sight-lines for which we produce synthetic spectra.

Figure 7.5 shows the mean charge of iron in the model and can be directly compared to Figure 3 in SP10. The computed ionization states are very similar. This similarity shows that the inclusion of an accretion disk in our simulation as a source of ionizing photons has not altered the conclusions of the earlier work with regards to species whose




Figure 7.5: The mean charge of iron in the model (c.f. [Sim et al. 2010b](#)).

ionization state is mainly set by the incident X-ray flux.

7.3.2 Synthetic spectra

Figure 7.6 shows synthetic spectra computed for the four sightlines illustrated in Figure 7.4. For comparison, the spectrum that would be seen in the absence of any absorbing material (labelled “unabsorbed”) is also shown as a faint grey line for each sightline. The 80° sightline passes through the failed wind region where carbon is not fully ionized. This is the most likely place to look for absorption from moderately ionized species. However, the simulated spectrum for this sightline (the heavy black line) is almost entirely featureless except for a small absorption feature at 255 Å due to a resonance line of Fe XXIV. The other three spectra in this plot show the proportion of the net spectrum made up from unscattered, singly scattered and multiply scattered photons. These contributions

reveal why we do not see any absorption features associated with the species present in the failed wind: essentially all of the photons escaping along this direction have scattered at least once and about 50% of them have scattered more than once. Thus the net spectrum for this sightline is produced by photons that have scattered *around* the failed wind, rather than passing through it. Fe XXIV is present in the equatorial, outflowing wind behind the failed wind region (see Figure 7.5) and scattered photons pass through regions of the wind where this ion is present. The direct (unscattered) component of the radiation field escaping along this sightline is very heavily attenuated by a combination of photoionization absorption and electron scattering.

The 60° sightline skims the top of the failed wind. Here, the component of the synthetic spectrum associated with photons that did not undergo any scattering events is only slightly less than 50% of the total. As the viewing angle moves even further up, to 55°, most of the photons escape unscattered. The Fe XXIV feature is still present at 60°, but only in the direct component. The spectrum predicted for the 55° sightline is completely featureless.

The final spectrum is computed for a sightline of 40°, and here we see a continuum enhancement at wavelengths below about 700 Å. This is a consequence of the geometry of the wind model: photon packets with wavelength less than 700 Å are predominantly produced by parts of the accretion disk inwards of the innermost radial extent of the wind (60 R_G). The spectrum is therefore made up of a direct component of photons passing unimpeded from the disk into the relatively low density polar parts of the wind (direct component), and photons reflecting from the inner face of the failed wind (scattered component). By contrast, a significant proportion of photon packets with wavelength longer than 700 Å are produced from parts of the accretion disk under the dense failed wind, and this component will tend to be absorbed or scattered back into the disk. We therefore only see a fraction of the direct component. Even when supplemented by scattered radiation, the resulting flux is less than would be seen in the wind were not present at all. The continuum enhancement at low inclination angles was also predicted in the X-ray spectra computed in SP10, where it was also attributed to reflection from the innermost illuminated face of the ‘failed wind’ region.




Figure 7.6: The simulated spectrum for three sightlines. The grey line is the flux that would be observed if there were no wind present. The heavy line shows the simulated spectrum for the converged wind model, and the remaining three lines split that spectrum into components. The blue line shows the part of the escaping flux associated with photons that have not scattered in the model, the green line shows the part associated with photons that have scattered just once, and the red line shows the part associated with multiply scattered photons.

7.3.3 Temperature distribution

The top two panels of Figure 7.7 show the electron temperature T_e computed for the hydrodynamic model on two different scales. The left hand panel matches the scale of the two lower plots, while the right hand panel is optimised to show the full range of temperatures obtained in the current calculations. The lower left panel shows the temperature structure estimated by PK04, whilst the lower right hand panel shows the temperature structure found by SP10. The clearest difference between the temperature computed in the hydrodynamic simulation and that found here is in the low density in-falling flow towards the polar direction. In this region, Compton heating and cooling are the dominant mechanisms, and so the temperature depends on the mean frequency of photons in the region. Since we include the radiation field produced by the accretion disk in our ionization and thermal equilibrium calculations of ionization, the mean frequency we find is far lower than that assumed in PK04, where only the 10 keV Bremsstrahlung spectrum was used to compute the ionization state and temperature structure. The temperature structure found by SP10 is closer to ours, since they also included an assumed disk spectrum in their calculation of Compton cooling. In any case, the exact temperature in this polar region is actually not all that important, since the density there is far too low to imprint any spectral features.

Another interesting difference is seen in the transition zone, where the PK04 calculation contains very hot ‘streamers’ of shock-heated gas. We do not model this non-radiative heating mechanism, and so our wind does show these structures. Again, this difference does not affect the simulated spectra or our wider conclusions, since this region is fully ionized in both cases.

Finally, a much more important difference is in the warm equatorial flow. This region lies behind the dense ‘failed wind’ region which is very cold all three calculations. In PK04, the ionization parameter was used to calculate the temperature in the model. The ionization parameter, in turn, was calculated using the X-ray luminosity of the central source, attenuated by electron scattering opacity and distance. In the dense equatorial flow, this estimate of the ionization parameter is very low, and the treatment used in PK04 (described in detail in [Proga et al. 2000](#)) yields unrealistically low temperatures. In




Figure 7.7: The electron temperature (T_e) of the model calculated here, compared to that from PK04 and SP10. The top two panels show the current calculation, but on different colour scales. The upper right scale is optimised for the data, whilst the top left is on an identical scale to the lower two plots. The lower left panel is the temperature from the original PK04 calculation, whilst the lower right panel shows the temperature from SP10.

such cases, they assumed that the gas would be in local thermodynamic equilibrium, and the temperature was set to be equal to the temperature of the disk below it. However, our more detailed radiative transfer treatment permits photons to scatter around the failed wind, and these scattered photons turn out to be the dominant source of ionization and heating in the equatorial flow region. We therefore obtain temperatures of a few hundred thousand Kelvin for this region, significantly hotter than in PK04, but similar to the temperatures obtained by SP10, whose calculation also allowed for scattering. We discuss the critical consequences of this finding in Section 7.4.

7.3.4 Ionization parameter

In Figure 7.8, we show the ionization parameter U , computed for each cell in the model based on the actual photon packets that have passed through the cell. The standard definition of U is

$$U = \frac{Q(H)}{4\pi r_0^2 n(H)c}, \quad (7.1)$$

where $Q(H)$ is the number of photons emitted by the source per second, r_0 is the distance from the illuminated cloud to the source, and $n(H)$ is the number density of hydrogen. Thus U is simply the ratio of the number density of ionizing photons, $n(\gamma_H)$, to the number density of hydrogen atoms. We explicitly track this information in our simulation, and so our version of U for each cell is given by

$$U = \frac{n(\gamma_H)}{n(H)}. \quad (7.2)$$

In the optically thin case, both definitions will give identical values of U . However, our version self-consistently takes account of any absorption between the source and the cell, along with additional photons which may scatter into the cell from other lines of sight.

Figure 7.8 shows immediately why we do not have any C IV in the wind. In order for a significant fraction of carbon ions to be found in this ionization stage, we require $\log(U) \sim -2$ (Higginbottom et al. 2013). There are no parts of the wind where this is true. There are also significant implications of this ionization parameter distribution on the self-consistency of the hydrodynamic model, since efficient line driving is only possible for significantly lower values of ionization parameter than we predict for the outflowing parts of the simulation. We discuss this further below.

7.4 Discussion

We are now in a position to consider some key aspects and consequences of our results in more detail. First, we will attempt isolate the cause of the discrepancy between the temperature structure and ionization state found by PK04 and those calculated above. Second, we will look at the way in which photons actually propagate through the




Figure 7.8: The ionization parameter $\log(U)$ in each cell of the model. The contours show the underlying density of the model.

PK04 model in order to address what the structure of a viable line-driven disk wind might be. Third and finally, we will discuss the implications of our study for dynamical simulations of line-driven winds.

7.4.1 Why is the wind hotter and more ionized than predicted by PK04?

There are several key differences between the calculation of the temperature and ionization state of the wind in PK04 and that described above:

- Input spectrum: PK04 used a 10 keV Bremsstrahlung spectrum to compute both the ionization parameter and the temperature of the wind; by contrast, we have

used an observationally motivated AGN spectrum that includes ionizing radiation associated with both the accretion and a central X-ray source.

- Wind radiation: in our model, photons produced by the wind itself are able to interact with (e.g. ionize) other parts of the outflow; in PK04, the wind was assumed to be transparent to its own cooling radiation.
- Scattering: in PK04, the local radiation field at each point in the wind was assumed to be dominated by *direct* radiation from the central engine, with electron scattering treated as a purely absorptive process; here we take full account of all opacity mechanisms and in particular we track the propagation of scattered photons through the flow.

To assess the relative importance of these effects, Figure 7.9 shows the distribution of the ionization parameter throughout for two additional test cases. The ionisation parameter is a useful parameter to compare here, since it largely defines the ionization state of the wind and is also used to compute the temperature and line-driving force in PK04.

The upper right panel in Figure 7.9 shows the ionization parameter as presented in Figure 1 in PK04. However, we have converted the definition of the ionization parameter adopted in PK04 (ξ) to that adopted here (U), using their approximate formula $\log(U) = \log(\xi) - 1.75$. Comparing this to Figure 7.8, we immediately see that our values of U are consistently higher than those found by PK04.

To test how much of this difference is due simply to the differences between the adopted illuminating spectra, we have rerun our simulation, with the accretion disk removed as a radiation source and with the X-ray spectrum modified to a power-law that approximates the 10 keV Bremstrahlung spectrum assumed in PK04, i.e. $F_\nu \propto \nu^0$. This X-ray spectrum is normalized so as to give the same luminosity over the 2-10 keV range as the PK04 SED, and the high-frequency cutoff of the Bremsstrahlung spectrum is approximated by truncating the power law at 10 keV. This input spectrum is used for all three of the new calculations shown in Figure 7.9.

The ionization parameter calculated for the wind in the absence of a disk radiation field is shown in the upper left panel of Figure 7.9. Based on this, we can conclude




Figure 7.9: The upper right panel shows the ionization parameter from PK04 (scaled from their definition ξ to ours via $\log(U) = \log(\xi) - 1.75$). The remaining three panels all show results from PYTHON calculations that use an approximation to the 10 keV bremsstrahlung spectrum assumed in PK04 as the only source of photons. The two panels on the left are based on the same simulation, but the upper panel shows U calculated taking into account all photons in the cell, whilst the lower panel only counts photons arriving directly from the X-ray source (i.e. unscattered photons). The lower right panel shows U from a calculation where the wind is permitted to cool via radiative processes, but the resulting photons are not actually produced and tracked in the simulation. This is analogous to the approximation made in PK04. Here, again, only unscattered photons are used to calculate U .

that the difference between the ionizing SED adopted by us and in PK04 is the main reason for the discrepancy between the estimates of U in the polar wind regions. This is unsurprising since the wind is optically thin in these directions. However, the critical region for the viability line-driving and the production of broad UV absorption lines is the warm, equatorial outflow. Here, U is uniformly very small in the PK04 data, but remains at $\log(U) \gtrsim 2$ outside the dense failed wind region even in our simulation without disk photons. This ionization parameter is too high to support line driving (Stevens and Kallman 1990).

In the lower left panel of Figure 7.9, we show U from the same no-disk calculation, but taking into account only photons that have not undergone any scattering events, i.e. only photons arriving directly from the central source or emitted by the wind. We now see a moderate decrease in U throughout the whole model except the bulk of the failed wind. This is consistent with the failed wind efficiently scattering photons into the polar regions in our standard calculation, and with regions shielded by the failed wind being affected strongly by radiation scattered around the shield. However, the ionization parameter in the warm, equatorial outflow region is still higher than in PK04 and still too high for efficient line driving.

Finally, in PK04, the wind is able to cool via radiative processes, but is assumed to be optically thin to its own radiation. We therefore carry out another no-disk, unscattered-photons-only calculation, but now also disable wind emission in the PYTHON simulation. More specifically, the wind is still allowed to cool via radiative processes, but no photons are produced that could interact with other parts of the wind. The result is shown in the lower right panel of Figure 7.9. Only in this last calculation, which most closely mimics the approximations made in PK04, does $\log(U)$ drop to a level where line driving could become efficient. It should be noted that drawing this distinction between the wind photons and scattered photons is largely artificial. Since PYTHON treats the wind as being in radiative and thermal equilibrium, and there are no non-radiative heating processes implemented (e.g. shock heating or magnetic reconnection) all photons that reach the wind have originated in the disk or central source. Some will have scattered,

whilst others have been absorbed and re-emitted. We call the latter wind photons, but what we call scattering is really just a special case of absorption and re-emission.

7.4.2 The propagation of ionizing photons through the outflow: implications for line-driving

We have seen that taking account of wind emission, multiple scattering and including the accretion disk as a source of ionizing radiation has resulted in an outflow that is too highly ionized for line driving to work efficiently as an acceleration mechanism. However, this does not mean that line driving could not produce an outflow at all. If one were to compute a new step in the hydrodynamical calculation based on our new ionization state, the line-driving forces, and hence the outflow structure, would certainly change (relative to the same time step in PK04's original simulation). The key question is whether the end result would be a total quenching of the wind or whether the flow might adjust itself to a new configuration in which line driving can be effective again. A self-consistent 3-dimensional radiation-hydrodynamic calculation would be needed to answer this question definitively, and we briefly comment on the prospects for this in Section 7.4.3 below. However, by considering the way in which ionizing radiation actually propagates through the outflow, we can already comment on the general properties of any viable line-driven disk wind model.

In Figure 7.10, we show some characteristics of the ionizing photons that reach three particular cells in the wind. One cell (1 - marked with the triangle in the upper panel) is in the highly ionized, low density part of the model and serves as a comparison to the other two cells (2 - marked with a circle, and 3 - marked with a square), both of which are in the outflowing part of the model. In PK04, cells 2 & 3 are characterized by the relatively low ionization states required for efficient line driving. In our calculations, these cells are highly ionized, so that line driving could not be effective in them. Also shown in the upper panel are the starting locations of the ionizing photons emitted by the wind that pass through cell 2. Broadly speaking, these photons turn out to originate in the outside skin of the failed wind. This region sees and absorbs a significant fraction of the ionizing radiation generated by the disk and the X-ray source. Since the wind is




Figure 7.10: Illustration of the source of ionizing photons contributing to the radiation field in three cells. The top panel shows the location of the three cells (filled symbols), along with the starting locations of wind generated photons that reach cell 2 (open circles). The underlying colours show the density of the model - note that the colour scale here is different to that used in Figure 7.2. The two straight lines in this plot correspond to two specific sightlines: the top line marks $i \simeq 26.5^\circ$, the sightline passing from the origin through cell 1; the bottom line marks $i \simeq 58^\circ$, a sightline passing through the transition region between the failed wind and the polar regions of the outflow. The middle panel shows the launch radius of the disk generated photons that reach the three cells (plot symbols match the symbols used to mark each cell in the top panel). The proportion of photons originating from the central X-ray source is also shown via the isolated symbols near the left hand side of the plot. The bottom panel shows the initial launch angle (relative to the vertical axis) of the same photons. The vertical lines mark the two angles also shown as sightlines in the top panel. In the lower two plots, the ordinate is normalised to the total number of photons reaching the cell in question.

assumed to be in radiative and thermal equilibrium, all of this absorbed energy must be locally re-emitted.

The middle panel in Figure 7.10 shows the distribution of the initial launch radii of the ionizing disk photons that pass through our test cells. The detached symbols near the left-hand-side of the plot show the corresponding fraction of photons emitted by the central X-ray source for comparison. We can see immediately that, for all cells, the majority of ionizing photons arise in the inner parts of the accretion disk. In PK04, the region above this part of the disk was outside the computational domain, and so it is empty in our calculation as well. Photons generated in this part of the disk are therefore able to reach the wind without attenuation. In the hydrodynamic calculation, this was done for reasons of computational efficiency, but it is also physically reasonable, since the temperature and ionization state of the gas above this part of the disk would likely be too high to permit efficient line driving. Nonetheless, the additional failed wind region that would likely result if this region was included in the simulation would certainly modify the spectrum incident on the rest of the wind and might result in a lower ionization state there. Extending the computational domain down to the innermost parts of the disk would therefore be an important next step in understanding line-driven winds. Useful steps towards this have already been taken in radiation-(magneto)hydrodynamic simulations of accretion disks (e.g. [Ohsuga and Mineshige 2011](#); [Jiang et al. 2013](#)).

The middle panel of Figure 7.10 also shows that the launch radius distribution of ionizing photons is similar for all three test cells, even though two of the cells (nos. 2 and 3) are located in the shadow of the failed wind region, as seen from the central engine. Thus the failed wind does not provide an effective shield from ionizing radiation. As already shown in Figure 7.6 and discussed in Section 7.3, this is because scattering allows photons to reach parts of the flow that would otherwise be shadowed. Interestingly, the radiation field in cells 2 and 3 also includes some photons generated in the outer parts of the disk; for example, about 10% of the ionizing photons reaching cell 3 come directly from the disk beneath it. This is possible because the outflow density and opacity has dropped significantly at these larger radii.

Finally, the lower panel of Figure 7.10 shows the distribution of launch angles (relative to the vertical axis) for all ionizing photon packets produced by the disk and central X-ray source that ultimately pass through the three test cells. The direct sightline from the origin to cell 1 corresponds to an angle of 26.5° , and the launch angle of most photons that reach this cell is, in fact, close to this value. However, a significant fraction of the ionizing photons in this cell were originally emitted in a more equatorial direction and only reach the cell after scattering off the failed wind. These photons produce the hump near 65° in the photon launch angle distribution for cell 1. Cells 2 and 3 (which lie at a sightline from the origin of 80°) are shielded from the direct radiation produced by the central X-ray source and inner disk by the failed wind. The ionizing photons reaching these cells instead exhibit a wide distribution of initial launching angles, peaking around 58° . This direction is marked by the lower straight line in the top panel of Figure 7.10. Photons taking this initial path interact with the hot transition region of the flow, i.e. the upper skin of the failed wind. When they enter this region, the optical depth along their trajectory is much higher than that perpendicular to it, so they are preferentially scattered both up into the polar regions and down into the parts of the flow behind the failed wind. Photons with initially more equatorial directions will initially strike the failed wind region. A significant fraction of these photons will be reflected back across the inner disk and enter the transition region on the other side of the grid. Some of these photons will then once again be scattered in the outflow region behind the failed wind.

This analysis suggests that it is the transition region (i.e. region B in Figures 7.1 and 7.2) that is mainly responsible for redirecting ionizing photons into the outflow regions behind the failed wind. This applies not only to photons emitted by the central X-ray source and inner disk, but also to photons produced by the “front” of the wind itself. Thus, even though *the failed wind is doing its job in preventing direct illumination* from reaching the outflow behind it, ionizing photons can get around it via scattering in the transition region.

How might we expect the outflow to respond? Given the over-ionization of the outflow behind the failed wind, it seems likely that the outflow in this region would also fail. This would result in a larger failed wind region - perhaps reducing the ionization state of

the material behind this extended block. However, since the failed wind is already highly optically thick in the radial direction, and since scattered and reprocessed radiation is already dominant in the shielded region, this effect is likely to be quite small.

Given the importance of the transition region in redirecting ionizing photons into otherwise shielded parts of the flow, it seems likely that it is this region that would have to be modified to allow the formation of a line-driven outflow. In the PK04 model, the geometry of this transition region allows photons that would otherwise pass over the outflow to be efficiently scattered down into it. If this transition region instead had a more curved upper surface profile, photons missing the failed wind could not be scattered so easily back down into the outflow behind this region. Is it plausible that the flow might adjust itself to form such a geometry? Possibly. In the PK04 model, line driving already only just works in the transition region, since the ionization state there is fairly high. Given the even higher ionization state we predict for this region, line driving would likely fail completely, causing the material to drop back to the disk, perhaps forming the geometry required to produce an effective shield. However, fully self-consistent calculations would be required to test this idea.

7.4.3 Implications for hydrodynamic simulations of line-driven winds

The analysis we have carried out on the PK04 geometry has clearly shown that reprocessing and scattering effects have a significant – and actually dominant – effect on the ionization state of the wind. This, in turn, must strongly affect the efficiency of the line driving mechanism itself. The clear conclusion from our results is therefore that these radiative effects must somehow be incorporated into hydrodynamic simulations.

The most obvious way to achieve this would be to carry out a full radiative transfer and ionization calculations after each time step of the hydrodynamic simulation. However, the main radiative transfer calculation presented here takes about 20 hours of processing time on 192 processors of the Southampton University Iridis supercomputer. Thus, even if such a scheme could be implemented, it is not at all clear that it is computationally feasible at this time. Nonetheless, it is important to develop at least an approximate method along these lines in order to obtain robust results from line-driven wind simulations.

It should be noted that it is not only full hydrodynamic simulations of line-driven winds that have treated radiative transfer in a simplified manner. Calculations such as those presented in [Risaliti and Elvis \(2010\)](#) and [Nomura et al. \(2013\)](#) also neglect scattered and reprocessed radiation and give rise to wind solutions similar to that presented in PK04. It is therefore likely that similar effects would be seen in these models.

7.5 Conclusions

We have presented results from comprehensive radiative transfer and photoionization calculations for the line-driven AGN disk wind predicted by the hydrodynamic simulations carried out in PK04. These simulations were computationally expensive and therefore treated the radiative transfer and ionization in a simplified manner. Here, we have focused on one snapshot from the hydrodynamic calculation and carried out a much more detailed Monte-Carlo simulation of the interaction between the AGN radiation field and the outflow.

Our main result is that the ionization state of the outflow is much higher than estimated by PK04, to the extent that line-driving would become inefficient. The over-ionized flow also no longer produces the broad UV absorption lines that are the key observational tracers of disk winds in AGN/QSOs. The main reason for this change in the predicted ionization state of the flow is that self-shielding becomes much less effective when radiative transfer effects are fully accounted for. More specifically, the failed wind region that protects the outflow from over-ionization in the simulations of PK04 can actually be “circumnavigated” by ionizing photons via scattering and reprocessing.

We conclude that hydrodynamic disk wind simulations need to take account of scattering and reprocessing in order to robustly assess the viability of line driving as an acceleration mechanism. This should ideally take the form of a self-consistent treatment, in which the radiative transfer and hydrodynamics are calculated simultaneously. However, the kind of ‘post-processing’ approach we have used here is at least useful in validating purely hydrodynamic models. For outflows driven non-radiatively, it already provides a means to self-consistently predict the observational characteristics of the flow.

Chapter 8

Summary and suggestions for future work

8.1 Summary

In this thesis, we have described the work carried out over the past three and a half years to investigate the viability of disk-winds as a physical explanation for the origin of the BALs seen in 10-20% of quasars with the long term goal of producing a unified model for the broad emission and absorption line regions in quasars.

The first section of this work was to extend the capabilities of PYTHON, our radiative transfer and ionization code to produce reliable results for AGN. The first modification was to include a source of X-ray photons. This was observationally motivated - all AGN show non-thermal X-ray emission, although the source is not fully understood. The inclusion of this extra source of ionizing photons necessitated a significant modification of the ionization scheme used by PYTHON. Originally, PYTHON used a computationally fast method for calculating the ionization state of the plasma, but this method is only accurate if the mean intensity can be adequately represented by a dilute blackbody. This is not the case if an X-ray power law is present. We therefore implemented a more flexible scheme, which models the mean intensity using power law and exponential components. This spectral model is then used to calculate the ionization state. We

compared the ionization state calculated by PYTHON using the new scheme with the output of CLOUDY for a simple geometry, and demonstrated excellent agreement for ions of interest.

The second modification was to include extra heating and cooling mechanisms required to properly calculate the temperature of a plasma illuminated by X-rays. Compton effects were the most important, and, again, comparisons with CLOUDY show that PYTHON obtains good results for simple geometries where the capabilities of the two codes overlap.

Armed with the updated version of PYTHON the first application of the code was to investigate a simple wind geometry. The aim was to see if BAL features could be produced from a disk wind at all, and if so what physical properties were required. We therefore began by constructing a smooth, biconical disk wind, with physical parameters based on observed quantities wherever possible. We demonstrated that BAL features could indeed be produced in this “benchmark” disk wind model. Moreover, we found that winds producing strong BAL features can potentially provide feedback at the levels required to influence the evolution of their host galaxies.

There are, however, two main challenges for our benchmark model. First, it is very sensitive to the X-ray luminosity of the central source, with realistic values of α_{OX} tending to overionize the wind. Increased mass loss through the wind (\dot{M}_{wind}) is able to partially recover the required ionization state, but this introduces other difficulties, such as the self-consistency of a model where the accretion rate is significantly less than the outflow rate. Second, line emission is seen only at a relatively low level in our model. This is important, since if QSO unification is to be achieved via disk winds, they must produce strong line emission for sightlines where no absorption is seen.

This initial work was promising, but we were also interested in investigating a more realistic wind model. Hydrodynamic wind models naturally produce density variations in the outflow, and one might expect denser regions to maintain lower ionization states in the presence of a given radiation field. In addition, increasing density increases the emission measure of the plasma, thereby increasing line emission. Thus, a more realistic wind could potentially solve both of the shortcomings of the simple wind at the same time.

We therefore used PYTHON to compute the ionization structure of a line-driven disk wind model first reported in [Proga and Kallman \(2004\)](#) and also produced synthetic spectra for this model. The results were very interesting, although not what we were initially hoping for. The original work used a necessarily simplified ionization and radiative transfer scheme, however the more complete RT treatment provided by PYTHON predicted that the outflow was very highly ionized. This meant that no BAL features were seen, but more importantly line driving would fail to provide the expected acceleration in that model. This is an important result, and whilst it does not suggest that line driving cannot be the driving mechanism for disk winds it does demonstrate that one needs to incorporate radiative transfer and ionization calculations into the hydrodynamic simulation.

8.2 Future work

The work done to date represents the first steps in a long-term effort to understand the promises and limits of the disk wind paradigm in explaining the properties of quasars. Future work will take place in three parallel strands.

First, whilst the physics included in the current version of PYTHON produces good results in most parts of a typical wind model, there are a few shortcomings that should be addressed in future versions. Some of the more pressing issues and likely solutions are described in Section [8.2.1](#).

Second, kinematically defined disk wind models like that introduced in Section [6.2](#) are a good way of gaining physical insight. The calculations of simple geometries tend to be relatively fast, and therefore large parameter spaces can be efficiently searched. Section [8.2.2](#) briefly describes some possible future directions.

Finally, we have demonstrated the importance of including detailed RT and ionisation calculations in radiation hydrodynamic simulations of disk winds. This is an exciting area of research, and we are well placed to make a significant contribution to it. We will discuss possible starting points in section [8.2.3](#).

8.2.1 Extending the capabilities of PYTHON

8.2.1.1 Lines and levels

One of the major shortcomings of the benchmark kinematic disk wind model was the weakness of the emission lines. In section 6.5.2, we argue that the relatively low covering factor of the wind means that resonance line scattering cannot contribute as much to the lines as in CVs. Assuming that model is at all realistic, we must therefore seek alternative mechanisms for the production of lines.

Collisionally excited line emission is implemented in PYTHON via a two-level atom model, as described in section 4.3.4, with the level populations set via the assumption of a dilute blackbody radiation field. We see from the comparisons with CLOUDY discussed in section 5.4.3 that the overall line luminosity is about right, but this is not to say that the energy comes out in the correct lines. The approximation used in PYTHON for populating the levels linked by a line is that the mean intensity can be modelled by a dilute blackbody. This is not the case for our AGN model, and this inconsistency may be a reason for the lack of thermal line emission for resonance lines.

The new ionization scheme maintains a model of the mean intensity in each cell, and, in principle, this could be used to calculate radiatively populated level populations in a more self consistent manner. In addition, this scheme could be used to calculate the partition function for each ionic species, which could improve the ionization calculations.

There is also a second, unmodelled, source of line radiation: recombination cascades. We permit recombination to excited levels via the ζ term in the ionization equation, but we do not explicitly follow the radiative cascade as the excited ion relaxes to the ground state. Such cascades are an important source of line emission (particularly for the Balmer lines of hydrogen), but including it as a process is not simple. In order to capture the process, one needs to deal with collisional and radiative excitation and de-excitation between all of the levels in an ion. This work is underway for hydrogen and helium, via the macro-atom formalism described by Lucy (2002, 2003). It is possible that this work could be extended to include *some* important ions of other elements, but a full treatment for all ionic species would prove very computationally expensive.

8.2.1.2 Improved disk models

The accretion disk model used in PYTHON for the models presented here are all based upon the radial temperature distribution from a thin disk model. This is a reasonable first approximation, but the spectra of real AGN suggest that modifications to this model should be considered.

One frequently observed departure from the thin disk spectrum is the so called “soft X-ray excess”. The excess refers to radiation below 1 keV, which lies above the extrapolation of the power law seen at higher energies. The highest frequency radiation produced by an accretion disk is set by the temperature of the innermost annulus, and, for SMBH of $M \sim 10^8 M_\odot$, this is around a few times 10^4 K. A blackbody with this temperature peaks in the EUV, and cannot explain the soft excess. [Done et al. \(2012\)](#) suggest a model for the soft excess where only the outer parts of the disk produce blackbody emission. At smaller radii, they find that Compton upscattering from the disk’s own atmosphere can reprocess EUV from the inner disk into soft X-rays, explaining the excess. A simple modification of the disk to reproduce this model has been tested in PYTHON and could be included.

Another possible shortcoming of our disk model arises when we have large outflow mass fluxes. [Knigge \(1999\)](#) discusses the possible impact of mass loss via a wind on the temperature distribution of an accretion disk. Again, such a scheme could be implemented in PYTHON.

Finally, PYTHON does not at present include general relativistic (GR) effects such as the gravitational redshift of photons due to the central SMBH. Gravitational redshift can actually produce a secondary line driving force as the photon frequency changes ([Dorodnitsyn 2003](#)). In principle these effects can be included, although more complex GR effects such as light-bending would be more challenging.

8.2.1.3 Other photoionization and recombination processes

As mentioned in section 5.3.4, there are several other ionization and recombination processes that affect astrophysical plasmas in the kind of state we want to model. These are

Auger ionization, charge exchange and possibly three-body recombination. The current ionization scheme is based upon a computationally fast approximate Saha formulation, and it is possible to fairly easily include extra processes that link adjacent states. This has already been done for dielectronic recombination. However, this procedure is harder to implement for Auger ionization, since it links non-adjacent ionic stages, or for charge exchange, which links ions of different elements. Furthermore, the calculation of rate coefficients to implement such processes actually removes the computational efficiency that was the original motivation for the scheme.

The next stage to improve the ionization state calculations of PYTHON is almost certainly to move to a full rate-equation scheme, albeit still based upon spectral models for the mean intensity in a cell. Specifically, this would involve the construction of a matrix of rate coefficients, which can link any ion to any other ion. This is then solved via a matrix inversion operation, still requiring an iteration on n_e .

8.2.2 Kinematic models

Whilst there are many competing models of disk winds, models defined using a simple parametric approach are a good way of gaining insight that can be applied more widely. This was the reasoning behind the kinematic model described in Chapter 6. A possible next stage of this research strand is to address the two shortcomings seen in that benchmark model.

In order to maintain a moderate ionization state in the presence of high X-ray luminosities, many previous studies have included a dense “shielding” region into their models (e.g. [Murray and Chiang 1995](#)). We could do the same, either as a separate element, or by finding a wind geometry which naturally produces such a shield (the dense base of a wind with slow acceleration is a likely candidate). If this takes the form of an extra, dense component, then it may also produce the enhanced line emission required to address unification issues.

The great benefit of simple kinematically defined models is the speed with which simulations proceed. Once a successful geometry is discovered, a grid of models with varying

parameters can be calculated, thereby producing an ‘atlas’ of observational signatures of different classes of wind solutions. In principle this would permit a statistical investigation linking disk wind geometries with the proportions of different classes and subclasses of AGN. This is a long term goal.

8.2.3 Hydrodynamic models

In chapter 7, we showed the results of a detailed RT calculation on a hydrodynamic disk wind model. The ionization state of the wind was shown to be too high for efficient line-driving, which called into question the self-consistency of the original calculation.

We also suggested how the wind might respond to this changed ionization state - it is distinctly possible that a new stable outflow could result, albeit with a different geometry. An exciting next step to this project would be to take the wind model with its new ionization state, and feed this back into the hydrodynamic simulation to see how it responds. Solving the problems of how to do this would pave the way for a linked RT / hydrodynamic simulation, where after each hydro step (or several steps) an RT step is calculated, which would update the ionization state of the wind. This would permit a fully integrated radiation hydrodynamic simulation of AGN disk winds to be carried out.

References

M. A. Abramowicz, M. Calvani, and L. Nobili. Thick accretion disks with super-Eddington luminosities. *ApJ*, 242:772–788, December 1980. doi: 10.1086/158512.

M. A. Abramowicz, B. Czerny, J. P. Lasota, and E. Szuszkiewicz. Slim accretion disks. *ApJ*, 332:646–658, September 1988. doi: 10.1086/166683.

R. Alexandroff, M. A. Strauss, J. E. Greene, N. L. Zakamska, N. P. Ross, W. N. Brandt, G. Liu, P. S. Smith, J. Ge, F. Hamann, A. D. Myers, P. Petitjean, D. P. Schneider, H. Yesuf, and D. G. York. Candidate type II quasars at $2 < z < 4.3$ in the Sloan Digital Sky Survey III. *MNRAS*, 435:3306–3325, November 2013. doi: 10.1093/mnras/stt1500.

J. T. Allen, P. C. Hewett, N. Maddox, G. T. Richards, and V. Belokurov. A strong redshift dependence of the broad absorption line quasar fraction. *MNRAS*, 410:860–884, January 2011. doi: 10.1111/j.1365-2966.2010.17489.x.

R. Antonucci. Unified models for active galactic nuclei and quasars. *ARA&A*, 31:473–521, 1993. doi: 10.1146/annurev.aa.31.090193.002353.

N. Arav. The “Ghost of LY alpha ” as Evidence for Radiative Acceleration in Quasars. *ApJ*, 465:617, July 1996. doi: 10.1086/177447.

N. Arav, K. T. Korista, T. A. Barlow, and Begelman. Radiative acceleration of gas in quasars. *Nature*, 376:576–578, August 1995. doi: 10.1038/376576a0.

N. Arav, T. A. Barlow, A. Laor, and R. D. Blandford. Keck high-resolution spectroscopy of MRK 335: constraints on the number of emitting clouds in the broad-line region. *MNRAS*, 288:1015–1021, July 1997.

N. Arav, R. H. Becker, S. A. Laurent-Muehleisen, M. D. Gregg, R. L. White, M. S. Brotherton, and M. de Kool. What Determines the Depth of Broad Absorption Lines? Keck HIRES Observations of BALQSO 1603+3002. *ApJ*, 524:566–571, October 1999a. doi: 10.1086/307841.

N. Arav, K. T. Korista, M. de Kool, V. T. Junkkarinen, and M. C. Begelman. Hubble Space Telescope Observations of the Broad Absorption Line Quasar PG 0946+301. *ApJ*, 516:27–46, May 1999b. doi: 10.1086/307073.

B. Atwood, J. A. Baldwin, and R. F. Carswell. Echelle spectrometry of the Seyfert galaxies NGC 3783 and Markarian 509. *ApJ*, 257:559–569, June 1982. doi: 10.1086/160011.

N. R. Badnell. Radiative Recombination Data for Modeling Dynamic Finite-Density Plasmas. *ApJS*, 167:334–342, December 2006. doi: 10.1086/508465.

S. A. Balbus and J. F. Hawley. A powerful local shear instability in weakly magnetized disks. I - Linear analysis. II - Nonlinear evolution. *ApJ*, 376:214–233, July 1991. doi: 10.1086/170270.

J. Baldwin, G. Ferland, K. Korista, and D. Verner. Locally Optimally Emitting Clouds and the Origin of Quasar Emission Lines. *ApJ*, 455:L119, December 1995. doi: 10.1086/309827.

R. H. Becker, R. L. White, M. D. Gregg, M. S. Brotherton, S. A. Laurent-Muehleisen, and N. Arav. Properties of Radio-selected Broad Absorption Line Quasars from the First Bright Quasar Survey. *ApJ*, 538:72–82, July 2000. doi: 10.1086/309099.

M. C. Begelman, C. F. McKee, and G. A. Shields. Compton heated winds and coronae above accretion disks. I Dynamics. *ApJ*, 271:70–88, August 1983. doi: 10.1086/161178.

E. Behar, M. Sako, and S. M. Kahn. Soft X-Ray Absorption by Fe^{0+} to Fe^{15+} in Active Galactic Nuclei. *ApJ*, 563:497–504, December 2001. doi: 10.1086/323966.

M. Beltrametti. Thermal instabilities in radiatively driven winds - Application to emission line clouds of quasars and active galactic nuclei. *ApJ*, 250:18–30, November 1981. doi: 10.1086/159344.

A. S. Bennett. The revised 3C catalogue of radio sources. *MmRAS*, 68:163, 1962.

R. D. Blandford and D. G. Payne. Hydromagnetic flows from accretion discs and the production of radio jets. *MNRAS*, 199:883–903, June 1982.

G. R. Blumenthal and W. G. Mathews. Theoretical emission line profiles in QSOs and Seyfert galaxies. *ApJ*, 198:517–526, June 1975. doi: 10.1086/153628.

M. S. Brotherton, C. De Breuck, and J. J. Schaefer. Spectropolarimetry of PKS 0040-005 and the orientation of broad absorption line quasars. *MNRAS*, 372:L58–L62, October 2006. doi: 10.1111/j.1745-3933.2006.00226.x.

E. M. Burbidge, C. R. Lynds, and G. R. Burbidge. ON the Measurement and Interpretation of Absorption Features in the Spectrum of the Quasi-Stellar Object 3c 191. *ApJ*, 144:447, April 1966. doi: 10.1086/148629.

D. M. Capellupo, F. Hamann, J. C. Shields, P. Rodríguez Hidalgo, and T. A. Barlow. Variability in quasar broad absorption line outflows - I. Trends in the short-term versus long-term data. *MNRAS*, 413:908–920, May 2011. doi: 10.1111/j.1365-2966.2010.18185.x.

D. M. Capellupo, F. Hamann, J. C. Shields, P. Rodríguez Hidalgo, and T. A. Barlow. Variability in quasar broad absorption line outflows - II. Multi-epoch monitoring of Si IV and C IV broad absorption line variability. *MNRAS*, 422:3249–3267, June 2012. doi: 10.1111/j.1365-2966.2012.20846.x.

D. M. Capellupo, F. Hamann, J. C. Shields, J. P. Halpern, and T. A. Barlow. Variability in quasar broad absorption line outflows - III. What happens on the shortest timescales? *MNRAS*, 429:1872–1886, March 2013. doi: 10.1093/mnras/sts427.

J. I. Castor, D. C. Abbott, and R. I. Klein. Radiation-driven winds in Of stars. *ApJ*, 195:157–174, January 1975. doi: 10.1086/153315.

H.-Y. Chiu. Gravitational collapse. *Physics Today*, 17(5):21–34, 1964. doi: <http://dx.doi.org/10.1063/1.3051610>. URL <http://scitation.aip.org/content/aip/magazine/physicstoday/article/17/5/10.1063/1.3051610>.

Cloudy & associates. *Hazy, a brief introduction to CLOUDY C10*, 2011.

M. H. Cohen, P. M. Ogle, H. D. Tran, R. C. Vermeulen, J. S. Miller, R. W. Goodrich, and A. R. Martel. Spectropolarimetry of Two Broad Absorption Line Quasars with the W. M. Keck Telescope. *ApJ*, 448:L77, August 1995. doi: 10.1086/309602.

F. A. Cordova and K. O. Mason. High-velocity winds from a dwarf nova during outburst. *ApJ*, 260:716–721, September 1982. doi: 10.1086/160291.

C. E. Cottis, M. R. Goad, C. Knigge, and S. Scaringi. Searching for the signature of radiative line driving: on the absence of Ly α -NV line-locking features in a large sample of BALQSOs. *MNRAS*, 406:2094–2112, August 2010. doi: 10.1111/j.1365-2966.2010.16839.x.

J. Crummy, A. C. Fabian, L. Gallo, and R. R. Ross. An explanation for the soft X-ray excess in active galactic nuclei. *MNRAS*, 365:1067–1081, February 2006. doi: 10.1111/j.1365-2966.2005.09844.x.

W. Cunto, C. Mendoza, F. Ochsenbein, and C. J. Zeippen. Topbase at the CDS. *A&A*, 275:L5, August 1993.

S. W. Davis and A. Laor. The Radiative Efficiency of Accretion Flows in Individual Active Galactic Nuclei. *ApJ*, 728:98, February 2011. doi: 10.1088/0004-637X/728/2/98.

M. de Kool and M. C. Begelman. Radiation Pressure–driven Magnetic Disk Winds in Broad Absorption Line Quasi-stellar Objects. *ApJ*, 455:448, December 1995. doi: 10.1086/176594.

K. P. Dere, E. Landi, H. E. Mason, B. C. Monsignori Fossi, and P. R. Young. CHIANTI - an atomic database for emission lines. *A&AS*, 125:149–173, October 1997. doi: 10.1051/aas:1997368.

T. Di Matteo, V. Springel, and L. Hernquist. Energy input from quasars regulates the growth and activity of black holes and their host galaxies. *Nature*, 433:604–607, February 2005. doi: 10.1038/nature03335.

M. Dietrich, S. J. Wagner, T. J.-L. Courvoisier, H. Bock, and P. North. Structure of the broad-line region of 3C 273. *A&A*, 351:31–42, November 1999.

M. A. DiPompeo, M. S. Brotherton, and C. De Breuck. Very Large Telescope Spectropolarimetry of Broad Absorption Line QSOs. *ApJS*, 193:9, March 2011. doi: 10.1088/0067-0049/193/1/9.

C. Done, S. W. Davis, C. Jin, O. Blaes, and M. Ward. Intrinsic disc emission and the soft X-ray excess in active galactic nuclei. *MNRAS*, 420:1848–1860, March 2012. doi: 10.1111/j.1365-2966.2011.19779.x.

A. V. Dorodnitsyn. Line-driven winds in the presence of strong gravitational fields. *MNRAS*, 339:569–576, February 2003. doi: 10.1046/j.1365-8711.2003.06205.x.

J. Drew and F. Verbunt. Investigation of a wind model for cataclysmic variable ultraviolet resonance line emission. *MNRAS*, 213:191–213, March 1985.

J. E. Drew and A. Boksenberg. Optical spectroscopy of two broad absorption line QSOs and implications for spherical-symmetric absorbing wind models. *MNRAS*, 211:813–831, December 1984.

I. Dutan and P. L. Biermann. The Efficiency of Using Accretion Power of Kerr Black Holes. In M. M. Shapiro, T. Stanev, and J. P. Wefel, editors, *Neutrinos and Explosive Events in the Universe*, page 175, January 2005.

D. O. Edge, J. R. Shakeshaft, W. B. McAdam, J. E. Baldwin, and S. Archer. A survey of radio sources at a frequency of 159 Mc/s. *MmRAS*, 68:37–60, 1959.

M. Elvis. A Structure for Quasars. *ApJ*, 545:63–76, December 2000. doi: 10.1086/317778.

D. Emmanoulopoulos, I. E. Papadakis, M. Dovčiak, and I. M. McHardy. General relativistic modelling of the negative reverberation X-ray time delays in AGN. *MNRAS*, 439:3931–3950, March 2014. doi: 10.1093/mnras/stu249.

R. T. Emmering, R. D. Blandford, and I. Shlosman. Magnetic acceleration of broad emission-line clouds in active galactic nuclei. *ApJ*, 385:460–477, February 1992. doi: 10.1086/170955.

J. E. Everett. Radiative Transfer and Acceleration in Magnetocentrifugal Winds. *ApJ*, 631:689–706, October 2005. doi: 10.1086/432678.

A. C. Fabian. The obscured growth of massive black holes. *MNRAS*, 308:L39–L43, October 1999. doi: 10.1046/j.1365-8711.1999.03017.x.

A. C. Fabian. Observational Evidence of Active Galactic Nuclei Feedback. *ARA&A*, 50: 455–489, September 2012. doi: 10.1146/annurev-astro-081811-125521.

B. L. Fanaroff and J. M. Riley. The morphology of extragalactic radio sources of high and low luminosity. *MNRAS*, 167:31P–36P, May 1974.

R. P. Fender, T. M. Belloni, and E. Gallo. Towards a unified model for black hole X-ray binary jets. *MNRAS*, 355:1105–1118, December 2004. doi: 10.1111/j.1365-2966.2004.08384.x.

G. J. Ferland, R. L. Porter, P. A. M. van Hoof, R. J. R. Williams, N. P. Abel, M. L. Lykins, G. Shaw, W. J. Henney, and P. C. Stancil. The 2013 Release of Cloudy. *Rev. Mexicana Astron. Astrofis.*, 49:137–163, April 2013.

J. Frank, A. King, and D. J. Raine. *Accretion Power in Astrophysics: Third Edition*. Cambridge University Press, January 2002.

S. C. Gallagher and J. E. Everett. Stratified Quasar Winds: Integrating X-ray and Infrared Views of Broad Absorption-line Quasars. In L. C. Ho and J.-W. Wang, editors, *The Central Engine of Active Galactic Nuclei*, volume 373 of *Astronomical Society of the Pacific Conference Series*, page 305, October 2007.

R. Ganguly, K. R. Sembach, T. M. Tripp, B. D. Savage, and B. P. Wakker. High-Resolution Absorption Spectroscopy of Multiphase, High-Metallicity Gas Associated with the Luminous Quasar HE 0226-4110. *ApJ*, 645:868–889, July 2006. doi: 10.1086/504395.

C. M. Gaskell, E. S. Klimek, and L. S. Nazarova. NGC 5548: The AGN Energy Budget Problem and the Geometry of the Broad-Line Region and Torus. *ArXiv e-prints*, November 2007.

K. G. Gayley. An Improved Line-Strength Parameterization in Hot-Star Winds. *ApJ*, 454:410, November 1995. doi: 10.1086/176492.

R. R. Gibson, L. Jiang, W. N. Brandt, P. B. Hall, Y. Shen, J. Wu, S. F. Anderson, D. P. Schneider, D. Vanden Berk, S. C. Gallagher, X. Fan, and D. G. York. A Catalog of Broad Absorption Line Quasars in Sloan Digital Sky Survey Data Release 5. *ApJ*, 692:758–777, February 2009. doi: 10.1088/0004-637X/692/1/758.

M. Gierliński and C. Done. Is the soft excess in active galactic nuclei real? *MNRAS*, 349:L7–L11, March 2004. doi: 10.1111/j.1365-2966.2004.07687.x.

M. Gierliński and C. Done. Energy-dependent variability and the origin of the soft X-ray excess in active galactic nuclei. *MNRAS*, 371:L16–L20, September 2006. doi: 10.1111/j.1745-3933.2006.00199.x.

M. Giustini, M. Cappi, and C. Vignali. On the absorption of X-ray bright broad absorption line quasars. *A&A*, 491:425–434, November 2008. doi: 10.1051/0004-6361:200810363.

G. L. Granato, L. Silva, P. Monaco, P. Panuzzo, P. Salucci, G. De Zotti, and L. Danese. Joint formation of QSOs and spheroids: QSOs as clocks of star formation in spheroids. *MNRAS*, 324:757–768, June 2001. doi: 10.1046/j.1365-8711.2001.04369.x.

M. D. Gregg, R. H. Becker, and W. de Vries. FR II Broad Absorption Line Quasars and the Life Cycle of Quasars. *ApJ*, 641:210–216, April 2006. doi: 10.1086/500381.

E. Guerras, E. Mediavilla, J. Jimenez-Vicente, C. S. Kochanek, J. A. Muñoz, E. Falco, and V. Motta. Microlensing of Quasar Broad Emission Lines: Constraints on Broad Line Region Size. *ApJ*, 764:160, February 2013. doi: 10.1088/0004-637X/764/2/160.

C. Hazard, M. B. Mackey, and A. J. Shimmins. Investigation of the Radio Source 3C 273 By The Method of Lunar Occultations. *Nature*, 197:1037–1039, March 1963. doi: 10.1038/1971037a0.

C. Hazard, D. C. Morton, R. Terlevich, and R. McMahon. Nine new quasi-stellar objects with broad absorption lines. *ApJ*, 282:33–52, July 1984. doi: 10.1086/162174.

P. C. Hewett and C. B. Foltz. The Frequency and Radio Properties of Broad Absorption Line Quasars. *AJ*, 125:1784–1794, April 2003. doi: 10.1086/368392.

J. Heyvaerts and C. Norman. The collimation of magnetized winds. *ApJ*, 347:1055–1081, December 1989. doi: 10.1086/168195.

N. Higginbottom, C. Knigge, K. S. Long, S. A. Sim, and J. H. Matthews. A simple disc wind model for broad absorption line quasars. *MNRAS*, 436:1390–1407, December 2013. doi: 10.1093/mnras/stt1658.

N. Higginbottom, D. Proga, C. Knigge, K. S. Long, J. H. Matthews, and S. A. Sim. Line-driven Disk Winds in Active Galactic Nuclei: The Critical Importance of Ionization and Radiative Transfer. *ArXiv e-prints*, February 2014.

P. F. Hopkins and M. Elvis. Quasar feedback: more bang for your buck. *MNRAS*, 401: 7–14, January 2010. doi: 10.1111/j.1365-2966.2009.15643.x.

S. Ichimaru. Bimodal behavior of accretion disks - Theory and application to Cygnus X-1 transitions. *ApJ*, 214:840–855, June 1977. doi: 10.1086/155314.

K. G. Jansky. Radio Waves from Outside the Solar System. *Nature*, 132:66, July 1933. doi: 10.1038/132066a0.

S. Jester. A Simple Test for the Existence of Two Accretion Modes in Active Galactic Nuclei. *ApJ*, 625:667–679, June 2005. doi: 10.1086/429812.

Y.-F. Jiang, J. M. Stone, and S. W. Davis. On the Thermal Stability of Radiation-dominated Accretion Disks. *ApJ*, 778:65, November 2013. doi: 10.1088/0004-637X/778/1/65.

V. T. Junkkarinen. A non-spherically symmetric model for absorption regions near quasars. *ApJ*, 265:73–84, February 1983. doi: 10.1086/160654.

V. T. Junkkarinen, E. M. Burbidge, and H. E. Smith. Spectrophotometry of the broad absorption-line QSO PHL 5200. *ApJ*, 265:51–72, February 1983. doi: 10.1086/160653.

D. W. Just, W. N. Brandt, O. Shemmer, A. T. Steffen, D. P. Schneider, G. Chartas, and G. P. Garmire. The X-Ray Properties of the Most Luminous Quasars from the Sloan Digital Sky Survey. *ApJ*, 665:1004–1022, August 2007. doi: 10.1086/519990.

S. Kaspi, P. S. Smith, H. Netzer, D. Maoz, B. T. Jannuzi, and U. Giveon. Reverberation Measurements for 17 Quasars and the Size-Mass-Luminosity Relations in Active Galactic Nuclei. *ApJ*, 533:631–649, April 2000. doi: 10.1086/308704.

S. Kaspi, D. Maoz, H. Netzer, B. M. Peterson, M. Vestergaard, and B. T. Jannuzi. The Relationship between Luminosity and Broad-Line Region Size in Active Galactic Nuclei. *ApJ*, 629:61–71, August 2005. doi: 10.1086/431275.

S. Kaspi, W. N. Brandt, D. Maoz, H. Netzer, D. P. Schneider, and O. Shemmer. Reverberation Mapping of High-Luminosity Quasars: First Results. *ApJ*, 659:997–1007, April 2007. doi: 10.1086/512094.

K. I. Kellermann, R. Sramek, M. Schmidt, D. B. Shaffer, and R. Green. VLA observations of objects in the Palomar Bright Quasar Survey. *AJ*, 98:1195–1207, October 1989. doi: 10.1086/115207.

S. J. Kenyon and L. Hartmann. Spectral energy distributions of T Tauri stars - Disk flaring and limits on accretion. *ApJ*, 323:714–733, December 1987. doi: 10.1086/165866.

A. King. Black Holes, Galaxy Formation, and the $M_{BH}-\sigma$ Relation. *ApJ*, 596:L27–L29, October 2003. doi: 10.1086/379143.

A. King. The AGN-Starburst Connection, Galactic Superwinds, and $M_{BH}-\sigma$. *ApJ*, 635:L121–L123, December 2005. doi: 10.1086/499430.

A. L. King, J. M. Miller, J. Raymond, A. C. Fabian, C. S. Reynolds, K. Gültekin, E. M. Cackett, S. W. Allen, D. Proga, and T. R. Kallman. Regulation of Black Hole Winds and Jets across the Mass Scale. *ApJ*, 762:103, January 2013. doi: 10.1088/0004-637X/762/2/103.

A. R. King. Black hole outflows. *MNRAS*, 402:1516–1522, March 2010. doi: 10.1111/j.1365-2966.2009.16013.x.

C. Knigge. The effective temperature distribution of steady-state, mass-losing accretion discs. *MNRAS*, 309:409–420, October 1999. doi: 10.1046/j.1365-8711.1999.02839.x.

C. Knigge, J. A. Woods, and J. E. Drew. The application of Monte Carlo methods to the synthesis of spectral line profiles arising from accretion disc winds. *MNRAS*, 273:225–248, March 1995.

C. Knigge, S. Scaringi, M. R. Goad, and C. E. Cottis. The intrinsic fraction of broad-absorption line quasars. *MNRAS*, 386:1426–1435, May 2008. doi: 10.1111/j.1365-2966.2008.13081.x.

J. H. Krolik and G. A. Kriss. Warm Absorbers in Active Galactic Nuclei: A Multitemperature Wind. *ApJ*, 561:684–690, November 2001. doi: 10.1086/323442.

J. H. Krolik and G. M. Voit. What Is the True Covering Factor of Absorbing Matter in BALQSOs? *ApJ*, 497:L5, April 1998. doi: 10.1086/311274.

J. H. Krolik, C. F. McKee, and C. B. Tarter. Two-phase models of quasar emission line regions. *ApJ*, 249:422–442, October 1981. doi: 10.1086/159303.

E. Landi, G. Del Zanna, P. R. Young, K. P. Dere, and H. E. Mason. CHIANTI—An Atomic Database for Emission Lines. XII. Version 7 of the Database. *ApJ*, 744:99, January 2012. doi: 10.1088/0004-637X/744/2/99.

A. Laor, F. Fiore, M. Elvis, B. J. Wilkes, and J. C. McDowell. The Soft X-Ray Properties of a Complete Sample of Optically Selected Quasars. II. Final Results. *ApJ*, 477:93, March 1997. doi: 10.1086/303696.

A. Lawrence and M. Elvis. Obscuration and the various kinds of Seyfert galaxies. *ApJ*, 256:410–426, May 1982. doi: 10.1086/159918.

L. W. Lee and D. A. Turnshek. On Correlations between the Broad Absorption Lines and Adjacent Broad Emission Lines in QSO Spectra. *ApJ*, 453:L61, November 1995. doi: 10.1086/309753.

K. S. Long and C. Knigge. Modeling the Spectral Signatures of Accretion Disk Winds: A New Monte Carlo Approach. *ApJ*, 579:725–740, November 2002. doi: 10.1086/342879.

L. B. Lucy. Improved Monte Carlo techniques for the spectral synthesis of supernovae. *A&A*, 345:211–220, May 1999.

L. B. Lucy. Monte Carlo transition probabilities. *A&A*, 384:725–735, March 2002. doi: 10.1051/0004-6361:20011756.

L. B. Lucy. Monte Carlo transition probabilities. II. *A&A*, 403:261–275, May 2003. doi: 10.1051/0004-6361:20030357.

L. B. Lucy and D. C. Abbott. Multiline transfer and the dynamics of Wolf-Rayet winds. *ApJ*, 405:738–746, March 1993. doi: 10.1086/172402.

C. R. Lynds. A Quasi-Stellar Source with a Rapidly Expanding Envelope. *ApJ*, 147:396, January 1967. doi: 10.1086/149021.

G. M. MacAlpine. Ionizing Continuum and Emission-Line Energetics of Quasars and Seyfert Galaxies. In M. Reyes-Ruiz and E. Vázquez-Semadeni, editors, *Revista Mexicana de Astronomia y Astrofisica Conference Series*, volume 18 of *Revista Mexicana de Astronomia y Astrofisica Conference Series*, pages 63–68, September 2003.

F. Macchetto, A. Capetti, W. B. Sparks, D. J. Axon, and A. Boksenberg. HST/FOC imaging of the narrow-line region of NGC 1068. *ApJ*, 435:L15–L18, November 1994. doi: 10.1086/187583.

J. Magorrian, S. Tremaine, D. Richstone, R. Bender, G. Bower, A. Dressler, S. M. Faber, K. Gebhardt, R. Green, C. Grillmair, J. Kormendy, and T. Lauer. The Demography of Massive Dark Objects in Galaxy Centers. *AJ*, 115:2285–2305, June 1998. doi: 10.1086/300353.

R. Maiolino, M. Salvati, A. Marconi, and R. R. J. Antonucci. The Ly-edge paradox and the need for obscured QSOs. *A&A*, 375:25–29, August 2001. doi: 10.1051/0004-6361:20010808.

T. R. Marsh and K. Horne. Images of accretion discs. II - Doppler tomography. *MNRAS*, 235:269–286, November 1988.

W. G. Mathews and G. J. Ferland. What heats the hot phase in active nuclei? *ApJ*, 323:456–467, December 1987. doi: 10.1086/165843.

P. A. Mazzali and L. B. Lucy. The application of Monte Carlo methods to the synthesis of early-time supernovae spectra. *A&A*, 279:447–456, November 1993.

B. R. McNamara and P. E. J. Nulsen. Heating Hot Atmospheres with Active Galactic Nuclei. *ARA&A*, 45:117–175, September 2007. doi: 10.1146/annurev.astro.45.051806.110625.

A. R. G. Mead, K. R. Ballard, P. W. J. L. Brand, J. H. Hough, C. Brindle, and J. A. Bailey. Optical and infrared polarimetry and photometry of blazars. *A&AS*, 83:183–204, April 1990.

F. Meyer, B. F. Liu, and E. Meyer-Hofmeister. Evaporation: The change from accretion via a thin disk to a coronal flow. *A&A*, 361:175–188, September 2000.

N. Murray and J. Chiang. Active Galactic Nuclei Disk Winds, Absorption Lines, and Warm Absorbers. *ApJ*, 454:L105, December 1995. doi: 10.1086/309775.

N. Murray and J. Chiang. Disk Winds and Disk Emission Lines. *ApJ*, 474:91, January 1997. doi: 10.1086/303443.

N. Murray and J. Chiang. Photoionization of Disk Winds. *ApJ*, 494:125, February 1998. doi: 10.1086/305183.

N. Murray, J. Chiang, S. A. Grossman, and G. M. Voit. Accretion Disk Winds from Active Galactic Nuclei. *ApJ*, 451:498, October 1995. doi: 10.1086/176238.

R. Narayan and I. Yi. Advection-dominated accretion: A self-similar solution. *ApJ*, 428:L13–L16, June 1994. doi: 10.1086/187381.

H. Netzer. A Thermal Wind Model for the X-Ray Outflow in GRO J1655-40. *ApJ*, 652:L117–L120, December 2006. doi: 10.1086/510067.

U. M. Noebauer, K. S. Long, S. A. Sim, and C. Knigge. The Geometry and Ionization Structure of the Wind in the Eclipsing Nova-like Variables RW Tri and UX UMa. *ApJ*, 719:1932–1945, August 2010. doi: 10.1088/0004-637X/719/2/1932.

M. Nomura, K. Ohsuga, K. Wada, H. Susa, and T. Misawa. Modeling Line-Driven Disk Wind for Broad Absorption Lines of Quasars. *PASJ*, 65:40, April 2013.

M. North, C. Knigge, and M. Goad. A new sample of broad absorption-line quasars exhibiting the ghost of Lyman α . *MNRAS*, 365:1057–1066, February 2006. doi: 10.1111/j.1365-2966.2005.09828.x.

P. M. Ogle, M. H. Cohen, J. S. Miller, H. D. Tran, R. W. Goodrich, and A. R. Martel. Polarization of Broad Absorption Line QSOS. I. A Spectropolarimetric Atlas. *ApJS*, 125:1–34, November 1999. doi: 10.1086/313272.

K. Ohsuga and S. Mineshige. Global Structure of Three Distinct Accretion Flows and Outflows around Black Holes from Two-dimensional Radiation-magnetohydrodynamic Simulations. *ApJ*, 736:2, July 2011. doi: 10.1088/0004-637X/736/1/2.

J. B. Oke and W. L. W. Sargent. The Nucleus of the Seyfert Galaxy NGC 4151. *ApJ*, 151:807, March 1968. doi: 10.1086/149486.

D. E. Osterbrock and G. J. Ferland. *Astrophysics of gaseous nebulae and active galactic nuclei*. University Science Books, 2006.

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

B. M. Peterson, L. Ferrarese, K. M. Gilbert, S. Kaspi, M. A. Malkan, D. Maoz, D. Merritt, H. Netzer, C. A. Onken, R. W. Pogge, M. Vestergaard, and A. Wandel. Central Masses and Broad-Line Region Sizes of Active Galactic Nuclei. II. A Homogeneous Analysis of a Large Reverberation-Mapping Database. *ApJ*, 613:682–699, October 2004. doi: 10.1086/423269.

G. Ponti, R. P. Fender, M. C. Begelman, R. J. H. Dunn, J. Neilsen, and M. Coriat. Ubiquitous equatorial accretion disc winds in black hole soft states. *MNRAS*, 422:L11, May 2012. doi: 10.1111/j.1745-3933.2012.01224.x.

R. Popham and R. Narayan. Accretion disk boundary layers in cataclysmic variables. 1: Optically thick boundary layers. *ApJ*, 442:337–357, March 1995. doi: 10.1086/175444.

D. Proga. Winds from Accretion Disks Driven by Radiation and Magnetocentrifugal Force. *ApJ*, 538:684–690, August 2000. doi: 10.1086/309154.

D. Proga. Numerical Simulations of Mass Outflows Driven from Accretion Disks by Radiation and Magnetic Forces. *ApJ*, 585:406–417, March 2003. doi: 10.1086/345897.

D. Proga. Theory of Winds in AGNs. In L. C. Ho and J.-W. Wang, editors, *The Central Engine of Active Galactic Nuclei*, volume 373 of *Astronomical Society of the Pacific Conference Series*, page 267, October 2007.

D. Proga and T. R. Kallman. Dynamics of Line-driven Disk Winds in Active Galactic Nuclei. II. Effects of Disk Radiation. *ApJ*, 616:688–695, December 2004. doi: 10.1086/425117.

D. Proga, J. M. Stone, and T. R. Kallman. Dynamics of Line-driven Disk Winds in Active Galactic Nuclei. *ApJ*, 543:686–696, November 2000. doi: 10.1086/317154.

Grote Reber. Early radio astronomy at wheaton illinois. *Proceedings of the IRE*, 46(1):15–23, 1958.

M. J. Rees. On Multiple Absorption Redshifts in Quasi-Stellar Objects. *ApJ*, 160:L29, April 1970. doi: 10.1086/180517.

M. J. Rees. Magnetic confinement of broad-line clouds in active galactic nuclei. *MNRAS*, 228:47P–50P, October 1987.

G. T. Richards, M. Lacy, L. J. Storrie-Lombardi, P. B. Hall, S. C. Gallagher, D. C. Hines, X. Fan, C. Papovich, D. E. Vanden Berk, G. B. Trammell, D. P. Schneider, M. Vestergaard, D. G. York, S. Jester, S. F. Anderson, T. Budavári, and A. S. Szalay. Spectral Energy Distributions and Multiwavelength Selection of Type 1 Quasars. *ApJS*, 166:470–497, October 2006. doi: 10.1086/506525.

G. Risaliti and M. Elvis. A non-hydrodynamical model for acceleration of line-driven winds in active galactic nuclei. *A&A*, 516:A89, June 2010. doi: 10.1051/0004-6361/200912579.

R. R. Ross and A. C. Fabian. A comprehensive range of X-ray ionized-reflection models. *MNRAS*, 358:211–216, March 2005. doi: 10.1111/j.1365-2966.2005.08797.x.

M. Rowan-Robinson. On the unity of activity in galaxies. *ApJ*, 213:635–647, May 1977. doi: 10.1086/155195.

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

R. M. Sambruna, L. Maraschi, and C. M. Urry. On the Spectral Energy Distributions of Blazars. *ApJ*, 463:444, June 1996. doi: 10.1086/177260.

A. Sandage. The Existence of a Major New Constituent of the Universe: the Quasistellar Galaxies. *ApJ*, 141:1560, May 1965. doi: 10.1086/148245.

K. Schawinski, D. Thomas, M. Sarzi, C. Maraston, S. Kaviraj, S.-J. Joo, S. K. Yi, and J. Silk. Observational evidence for AGN feedback in early-type galaxies. *MNRAS*, 382:1415–1431, December 2007. doi: 10.1111/j.1365-2966.2007.12487.x.

M. Schmidt. 3C 273 : A Star-Like Object with Large Red-Shift. *Nature*, 197:1040, March 1963. doi: 10.1038/1971040a0.

M. Schmidt and R. F. Green. Quasar evolution derived from the Palomar bright quasar survey and other complete quasar surveys. *ApJ*, 269:352–374, June 1983. doi: 10.1086/161048.

N. J. Schurch, C. Done, and D. Proga. The Impact of Accretion Disk Winds on the X-Ray Spectra of Active Galactic Nuclei. II. Xscort + Hydrodynamic Simulations. *ApJ*, 694:1–11, March 2009. doi: 10.1088/0004-637X/694/1/1.

C. K. Seyfert. Nuclear Emission in Spiral Nebulae. *ApJ*, 97:28, January 1943. doi: 10.1086/144488.

J. R. Shakeshaft, M. Ryle, J. E. Baldwin, B. Elsmore, and J. H. Thomson. A survey of radio sources between declinations -38° and $+83^\circ$. *MmRAS*, 67:106, 1955.

N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. *A&A*, 24:337–355, 1973.

F. Shankar, X. Dai, and G. R. Sivakoff. Dependence of the Broad Absorption Line Quasar Fraction on Radio Luminosity. *ApJ*, 687:859–868, November 2008. doi: 10.1086/591488.

G. A. Shields. The origin of the broad line emission from Seyfert galaxies. *Astrophys. Lett.*, 18:119–123, 1977.

I. Shlosman and P. Vitello. Winds from accretion disks - Ultraviolet line formation in cataclysmic variables. *ApJ*, 409:372–386, May 1993. doi: 10.1086/172670.

I. Shlosman, P. A. Vitello, and G. Shaviv. Active galactic nuclei - Internal dynamics and formation of emission clouds. *ApJ*, 294:96–105, July 1985. doi: 10.1086/163278.

S. A. Sim, K. S. Long, L. Miller, and T. J. Turner. Multidimensional modelling of X-ray spectra for AGN accretion disc outflows. *MNRAS*, 388:611–624, August 2008. doi: 10.1111/j.1365-2966.2008.13466.x.

S. A. Sim, L. Miller, K. S. Long, T. J. Turner, and J. N. Reeves. Multidimensional modelling of X-ray spectra for AGN accretion disc outflows - II. *MNRAS*, 404:1369–1384, May 2010a. doi: 10.1111/j.1365-2966.2010.16396.x.

S. A. Sim, D. Proga, L. Miller, K. S. Long, and T. J. Turner. Multidimensional modelling of X-ray spectra for AGN accretion disc outflows - III. Application to a hydrodynamical simulation. *MNRAS*, 408:1396–1408, November 2010b. doi: 10.1111/j.1365-2966.2010.17215.x.

L. E. Simon and F. Hamann. The origins of a rich absorption line complex in a quasar at redshift 3.45. *MNRAS*, 409:269–283, November 2010. doi: 10.1111/j.1365-2966.2010.17306.x.

O. Slone and H. Netzer. The effects of disc winds on the spectrum and black hole growth rate of active galactic nuclei. *MNRAS*, 426:656–664, October 2012. doi: 10.1111/j.1365-2966.2012.21699.x.

J. Smak. On the Emission Lines from Rotating Gaseous Disks. *Acta Astron.*, 31:395, 1981.

A. Soltan. Masses of quasars. *MNRAS*, 200:115–122, July 1982.

U. Springmann. Multiple resonance line scattering and the 'momentum problem' in Wolf-Rayet star winds. *A&A*, 289:505–523, September 1994.

I. R. Stevens and T. R. Kallman. X-ray illuminated stellar winds - Ionization effects in the radiative driving of stellar winds in massive X-ray binary systems. *ApJ*, 365:321–331, December 1990. doi: 10.1086/169486.

J. M. Stone and M. L. Norman. ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests. *ApJS*, 80:753–790, June 1992. doi: 10.1086/191680.

J. Surdej and D. Hutsemekers. Geometry of the mass-outflows around broad absorption line QSOs and formation of the complex Ly-alpha + N V line profile. *A&A*, 177:42–50, May 1987.

R. S. Sutherland. Accurate free-free Gaunt factors for astrophysical plasmas. *MNRAS*, 300:321–330, October 1998. doi: 10.1046/j.1365-8711.1998.01687.x.

F. Tombesi, M. Cappi, J. N. Reeves, G. G. C. Palumbo, T. Yaqoob, V. Braito, and M. Dadina. Evidence for ultra-fast outflows in radio-quiet AGNs. I. Detection and statistical incidence of Fe K-shell absorption lines. *A&A*, 521:A57, October 2010. doi: 10.1051/0004-6361/200913440.

E. Treister, C. M. Urry, and S. Virani. The Space Density of Compton-Thick Active Galactic Nucleus and the X-Ray Background. *ApJ*, 696:110–120, May 2009. doi: 10.1088/0004-637X/696/1/110.

S. Tremaine, K. Gebhardt, R. Bender, G. Bower, A. Dressler, S. M. Faber, A. V. Filippenko, R. Green, C. Grillmair, L. C. Ho, J. Kormendy, T. R. Lauer, J. Magorrian, J. Pinkney, and D. Richstone. The Slope of the Black Hole Mass versus Velocity Dispersion Correlation. *ApJ*, 574:740–753, August 2002. doi: 10.1086/341002.

J. R. Trump, P. B. Hall, T. A. Reichard, G. T. Richards, D. P. Schneider, D. E. Vanden Berk, G. R. Knapp, S. F. Anderson, X. Fan, J. Brinkman, S. J. Kleinman, and A. Nitta. A Catalog of Broad Absorption Line Quasars from the Sloan Digital Sky Survey Third Data Release. *ApJS*, 165:1–18, July 2006. doi: 10.1086/503834.

M. Türler, S. Paltani, T. J.-L. Courvoisier, M. F. Aller, H. D. Aller, A. Blecha, P. Bouchet, M. Lainela, I. M. McHardy, E. I. Robson, J. A. Stevens, H. Teräsranta, M. Tornikoski, M.-H. Ulrich, E. B. Waltman, W. Wamsteker, and M. C. H. Wright. 30 years of multi-wavelength observations of 3C 273. *A&AS*, 134:89–101, January 1999. doi: 10.1051/aas:1999125.

M.-H. Ulrich, L. Maraschi, and C. M. Urry. Variability of Active Galactic Nuclei. *ARA&A*, 35:445–502, 1997. doi: 10.1146/annurev.astro.35.1.445.

C. M. Urry and P. Padovani. Unified Schemes for Radio-Loud Active Galactic Nuclei. *PASP*, 107:803, September 1995. doi: 10.1086/133630.

H. van Regemorter. Rate of Collisional Excitation in Stellar Atmospheres. *ApJ*, 136:906, November 1962. doi: 10.1086/147445.

D. A. Verner, P. D. Barthel, and D. Tytler. Atomic data for absorption lines from the ground level at wavelengths greater than 228Å. *A&AS*, 108:287–340, December 1994.

D. A. Verner, G. J. Ferland, K. T. Korista, and D. G. Yakovlev. Atomic Data for Astrophysics. II. New Analytic FITS for Photoionization Cross Sections of Atoms and Ions. *ApJ*, 465:487, July 1996a. doi: 10.1086/177435.

D. A. Verner, E. M. Verner, and G. J. Ferland. Atomic Data for Permitted Resonance Lines of Atoms and Ions from H to Si, and S, Ar, Ca, and Fe. *Atomic Data and Nuclear Data Tables*, 64:1, 1996b. doi: 10.1006/adnd.1996.0018.

R. A. Wade. A double grid of accretion disc model spectra for cataclysmic variable stars. *MNRAS*, 208:381–398, May 1984.

D. W. Weedman. Seyfert galaxies. *ARA&A*, 15:69–95, 1977. doi: 10.1146/annurev.aa.15.090177.000441.

R. J. Weymann, R. F. Carswell, and M. G. Smith. Absorption lines in the spectra of quasistellar objects. *ARA&A*, 19:41–76, 1981. doi: 10.1146/annurev.aa.19.090181.000353.

R. J. Weymann, J. S. Scott, A. V. R. Schiano, and W. A. Christiansen. A thermal wind model for the broad emission line region of quasars. *ApJ*, 262:497–510, November 1982. doi: 10.1086/160443.

R. J. Weymann, S. L. Morris, C. B. Foltz, and P. C. Hewett. Comparisons of the emission-line and continuum properties of broad absorption line and normal quasi-stellar objects. *ApJ*, 373:23–53, May 1991. doi: 10.1086/170020.

Y. Xu, W.-H. Bian, Q.-R. Yuan, and K.-L. Huang. The origin and evolution of CIV Baldwin effect in QSOs from the Sloan Digital Sky Survey. *MNRAS*, 389:1703–1708, October 2008. doi: 10.1111/j.1365-2966.2008.13545.x.

Q. Yu and S. Tremaine. Observational constraints on growth of massive black holes. *MNRAS*, 335:965–976, October 2002. doi: 10.1046/j.1365-8711.2002.05532.x.

N. L. Zakamska, M. A. Strauss, J. H. Krolik, M. J. Collinge, P. B. Hall, L. Hao, T. M. Heckman, Ž. Ivezić, G. T. Richards, D. J. Schlegel, D. P. Schneider, I. Strateva, D. E. Vanden Berk, S. F. Anderson, and J. Brinkmann. Candidate Type II Quasars from the Sloan Digital Sky Survey. I. Selection and Optical Properties of a Sample at $0.3 < Z < 0.83$. *AJ*, 126:2125–2144, November 2003. doi: 10.1086/378610.

W. Zheng, G. A. Kriss, R. C. Telfer, J. P. Grimes, and A. F. Davidsen. A Composite HST Spectrum of Quasars. *ApJ*, 475:469, February 1997. doi: 10.1086/303560.

H. Zhou, T. Wang, H. Wang, J. Wang, W. Yuan, and Y. Lu. Polar Outflows in Six Broad Absorption Line Quasars. *ApJ*, 639:716–723, March 2006. doi: 10.1086/499768.