The Long And Winding Road: Lyman-α Radiative Transfer And Galactic Outflows

Andrew S. Chung

München 2016

The Long And Winding Road: Lyman-α Radiative Transfer And Galactic Outflows

Andrew S. Chung

Dissertation
an der Fakultät für Physik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Andrew S. Chung
aus Leicester

München, den 24. Marz, 2016

Erstgutachter: Prof. Dr. Simon D.M. White
Zweitgutachter: Prof. Dr. Gerhard Börner
Tag der mündlichen Prüfung: 23. Mai, 2016

Inhaltsverzeichnis

Abbildungsverzeichnis

Tabellenverzeichnis

Zusammenfassung

Das kosmologische Standardmodell stellt eine äußerst erfolgreiche Beschreibung des sichtbaren Universums dar. Um jedoch die Vorhersagen des Standardmodells für die Verteilung von Galaxien in Einklang mit den Beobachtungen zu bringen, ist die Annahme von Feedback durch galaktische Outflows erforderlich. Daher ist das Verständnis und Modellieren dieser Outflows von großem Interesse.

In der vorliegenden Arbeit präsentieren wir zwei neue numerische Lyα-Strahlungstransfer-Codes, die wir in einer zweigleisigen Strategie verwenden, um galaktischer Outflows, die von stellarer Feedback angetrieben werden, zu studieren.

Im ersten Ansatz analysieren wir das systematische Versagen der fundamentalen „Schalenmodelle“ von Outflows, welche nicht in der Lage sind die sogenannten Blue Bumps in kürzlich beobachteten Spektren zu reproduzieren. Danach präsentieren und testen wir eine einfache Erweiterung des Standard-Schalenmodells, wodurch sich die Blue Bumps auf natürliche Weise durch Fermi-artige Beschleunigung der Lyα-Photonen ergeben. Auf diese Weise können wir einige der Unzulänglichkeiten korrigieren, die sich im Schalenmodell zu zeigen beginnen, und können die Konsistenz mit den Beobachtungen zumindest vorübergehend wahren.

Trotzdem, unter Berücksichtigung zukünftiger Beobachtungen, kann davon ausgegangen werden, dass das einfache Schalenmodell aufgegeben werden muss. Daher verfolgen wir auch einen alternativen Ansatz, indem wir versuchen den hydrodynamischen Prozess der Entstehung von Galaxien im kosmologischen Kontext zu simulieren. Als nächsten Schritt berechnen wir den numerischen Strahlungstransfer um die beobachtbaren Eigenschaften der entstehenden Galaxien zu bestimmen. Auf diese Weise können wir die galaktischen Outflow-Modelle mit den Beobachtungen entsprechender Umgebungen vergleichen. Wir stellen fest, dass die Kombination dreier Observablen (Lyα-Emission, Absorption und spektraler Verlauf) starke Einschränkungen an die derzeitigen Outflow-Modelle stellen und somit letztendlich die Entwicklung besserer Modelle nahelegt.

Abstract

The standard model of cosmology has been extremely successful in explaining observations of the universe. However, bringing the standard model predictions for the distribution of galaxies into agreement with observations relies critically on the invocation of feedback processes to regulate galaxy formation via galactic outflows. As such, there is intense interest in understanding and modelling the underlying physical processes involved in these outflows.

In this thesis we present two new numerical Lyα radiative transfer codes, and apply them in a two-pronged approach to understanding galactic outflows driven by stellar feedback.

In our first approach we identify the first systematic failing of basic 'shell model' outflows – an inability to produce the 'blue bumps' seen in the spectra of recent observations. We then present, and test with numerical radiative transfer simulations, a minor extension to the standard shell model paradigm which leads naturally to the production of blue bumps via Fermi-like acceleration of Lyα photons. In this way we paper over the cracks that were starting to show in shell models, and allow them to remain consistent with observations for the present time.

However, we also cast one eye to the future, where we expect that at some point we will be forced to abandon such simple shell models. We therefore pursue an approach whereby we attempt to numerically simulate the hydrodynamical process of galaxy formation in a cosmological context. As a further step we then perform numerical radiative transfer to derive the observable properties of the formed galaxy. In this way we are able to test galactic outflow models against observations of similar systems. We find that a combination of the three observables which we simulate (Lyα emission, absorption, and spectral shape) provides a strong constraint on current outflow models, and ultimately motivates the development of better models.

Chapter 1

Introduction

1.1 Origins

This is the way the universe begins
This is the way the universe begins
This is the way the universe begins
Not with a whimper but a bang[1]

The current standard model of Big Bang cosmology (ΛCDM, see Frenk & White 2012 for a recent review) states that at the very earliest times the universe was extremely hot and dense; that this phase was followed by a period of rapid expansion known as 'cosmic inflation', where the universe rapidly expanded in size by at least ~ 60 e-foldings; and that since that time the universe has continued to expand and cool, albeit at a vastly reduced rate.

During inflation all the matter in the universe is thought to have come into being. In addition to familiar baryonic matter, this includes the dark matter content of the universe.

Recent experiments such as those carried out by the WMAP(Bennett et al., 1997) and Planck (Tauber et al., 2010) spacecraft have used the cosmic microwave background (CMB), the 'afterglow' of the big bang, to perform a census of the mass-energy content of the universe (see Figure 1.1). The results show that today dark energy, to which we attribute the observed current accelerating expansion of the universe (Riess et al., 1998; Perlmutter et al., 1999), dominates. Previous eras were the radiation-dominated era, which was followed by the matter-dominated era, finally giving way to the dark energy-dominated era that we live in.

Of the $\sim$31% of the universe which is matter, only $\sim 1/6$ is composed of baryons (i.e. 4.9% of the universe). The other $\sim 5/6$ of the matter (i.e. 26.3% of the universe) is dark matter. This mysterious, elusive type of matter was first proposed to explain observations of galaxy clusters (Zwicky, F., 1933) and galaxy rotation curves (Rubin et al., 1978). The

[1] with apologies to T.S.Eliot

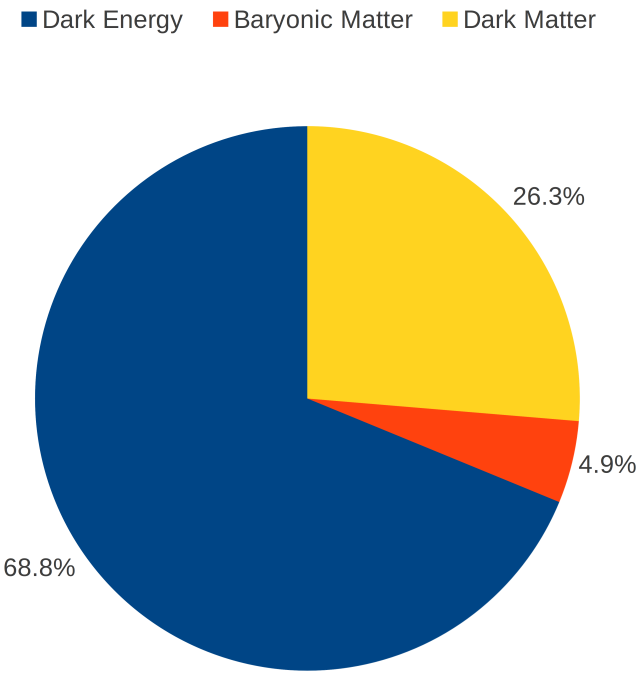

Figure 1.1: The current-day composition of the universe as measured by Planck utilising TT, TE, EE + lowP + lensing (Planck Collaboration et al., 2015)

notion that a form of matter exists which is not only invisible and weakly interacting with baryonic matter, but is also required in such quantities as to dwarf the baryonic matter that we see, at first seemed radical. However, despite the absence of a direct detection[2], evidence for the existence of dark matter has been mounting (structure formation, gravitational lensing by galaxy clusters, CMB anisotropies). This, in addition to the fact that no competing theory (i.e. an alternative to General Relativity) has enjoyed the success of dark matter models[3], has led to widespread acceptance of the dark matter hypothesis.

Testament to this acceptance can be found in the fact that dark matter has become a cornerstone of the standard model of cosmology. Indeed, the 'CDM' in ΛCDM refers to Cold Dark Matter, with 'Cold' specifying that in this model the dark matter is non-relativistic. Hot (relativistic) dark matter has effectively been ruled out as a viable model from studies of large scale structure (White et al., 1983, 1984), whereas warm (near-relativistic) dark matter remains a viable candidate (Lovell et al., 2012; Frenk & White, 2012).

Figure 1.2 shows a schematic of the timeline of the universe per ΛCDM, with several important epochs marked. We now briefly summarise some of these important features.

[2]Peebles (1998) states that 'Our reliance on hypothetical dark matter is an embarrassment; a laboratory detection would be exceedingly welcome'

[3]though the author acknowledges that this may result, at least in part, from the sociological conditions of the modern scientific endeavour.

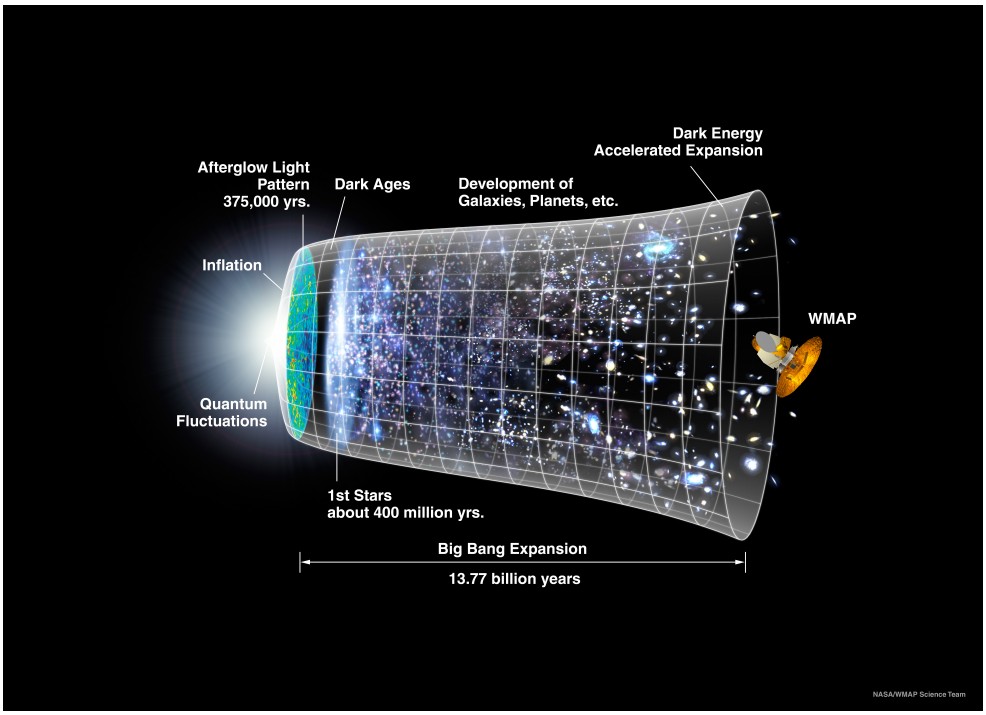

Figure 1.2: Timeline of the universe from NASA/WMAP Science Team

Firstly, from CMB measurements of curvature it appears that our observable universe is flat to the percent-level. This spatial flatness is commonly explained as a natural consequence of inflation.

The epoch marked as 'Afterglow Light Pattern' refers to the time at which two related events occurred: recombination and the advent of the last scattering surface. Up until this point the universe was so hot that the protons and electrons had too much energy to be bound into atoms. The fact that the universe was filled with free electrons rendered it opaque to radiation due to Thomson scattering of photons by the free electrons. Once the universe had expanded and cooled enough (i.e. to the point where there were few photons with energy higher than the ionisation energy of hydrogen) it became possible to form atomic hydrogen. After the formation of atoms the photons are able to free-stream through the universe. This phase transition from ionised and optically thick to neutral and optically thin is named 'recombination'. It is this free-streaming radiation that we observe as the CMB, and the last scattering surface, located at $z \sim 1100$, represents the location of CMB photons' last scatterings before the universe became optically thin.

The period between recombination and the formation of the first stars is referred to as the 'Dark Ages' due to the absence of radiation sources (and ignoring CMB photons). The dark ages are brought to an end when gravitational attraction amplifies the primordial density fluctuations to the point where the first stars begin to form via gravitational collapse.

Once stars, galaxies, and active galactic nuclei (AGN) form they produce ionising ra-

diation which is able to begin to reionise the universe. Hence, the universe underwent another phase transition from neutral to ionised in a process known as reionisation. It is important to note that reionisation does not result in the universe reverting to an optically thick state. By the time of reionisation the universe had expanded sufficiently that the density of free electrons was low enough that the universe remained optically thin. Observations and theory show that reionisation started around $z \approx 20$ and finished by $z \approx 6$. Nevertheless the details of reionisation remain uncertain. The dominant source of ionising photons is still unclear (Population III stars, galaxies, AGN) as are the exact beginning and end-points, and the spatial time evolution of the reionisation process (Barkana & Loeb, 2001; Ciardi & Ferrara, 2005). As we will see in §1.4, the study of Lyman-α (henceforth Lyα) radiative transfer through the circumgalactic medium (CGM) surrounding galaxies can help better constrain models of reionisation.

The final epoch in Figure 1.2 marks the start of the present era of dark energy accelerated expansion. In ΛCDM dark energy is a cosmological constant, Λ.

One way to study in detail the complex non-linear evolution of the universe is via simulations. In the following sections key components of cosmological simulations are discussed, while the final section of this chapter details how the work in this thesis fits into the described framework and extends our knowledge of galaxy formation and evolution.

1.2 N-body Simulations

In a way the fact that dark matter dominates the matter budget of the universe has turned out to be a blessing for cosmologists. Dark matter, at least as usually formulated, has both negligible self-interaction and negligible interaction with baryons except via the gravitational force (in both cases). This makes studying and simulating dark matter a much more tractable problem than simulating baryonic matter, which is often 'messy' (see §1.3) and requires the inclusion of many more physical processes.

Since dark matter dominates the mass budget, and only interacts with baryons gravitationally, we can use dark-matter-only simulations to investigate the large-scale structure of the universe. Whilst not true in detail due to baryonic physics, it is true that in general the baryons will trace the underlying dark matter distribution due to gravitational interaction. Since dark matter only self-interacts gravitationally the computational complexity of a naïve dark matter simulation would be $\mathcal{O}(n^2)$. However, by making simplifying approximations (Springel, 2005) one can use a tree-based algorithm to achieve the $\mathcal{O}(n \log n)$ performance typical of divide-and-conquer algorithms. The scale of cosmological N-body dark matter simulations has grown as a result of both algorithmic improvements and the explosion in available computing power. The size of the simulations can be crudely[4] characterised by the dark matter particle count which has kept pace with Moore's Law (Moore, 1965), as shown in Figure 1.3. Increasing the particle count for a given simulation box size allows a higher mass resolution, and consequently allows resolution of smaller structures.

[4]simply counting particles ignores other important factors such as the softening length and timestep size.

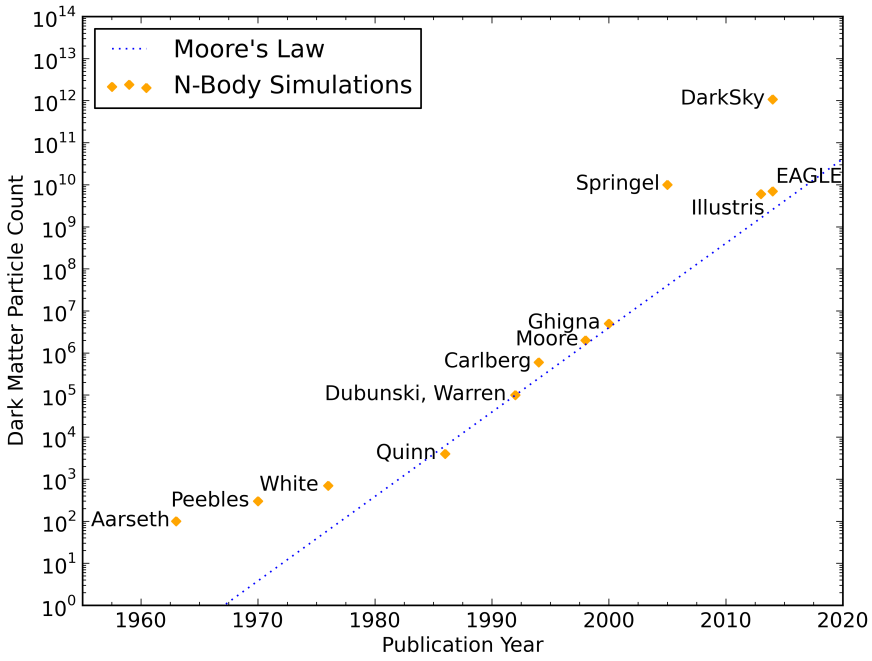

Figure 1.3: The increase in particle count of various N-body dark matter simulations over time. Plot inspired by 'Simulating Clusters of Galaxies: A Brief History of "N" and Overmerging' by Ben Moore `https://ned.ipac.caltech.edu/level5/Moore/Moore_contents.html`, extended with post-2000 data points. The dotted blue line shows the trend expected from Moore's Law, assuming a doubling timescale of 18 months and normalised to the (Dubunski, Warren) data point as in the aforementioned link.

It is widely predicted that Moore's Law will soon (within the next ten years) begin to fail, although such predictions have repeatedly been made over the 50 years that Moore's Law has so-far remained valid. Nevertheless, if this prediction should come to pass then it will likely have serious ramifications not only for the future of N-body simulations but for the whole field of scientific computation.

Figure 1.4 shows three snapshots from the Millennium dark matter simulation (Springel et al., 2005), displayed in time-order from left to right. As shown the progression of time sees the universe evolve from its initial almost homogeneous state to a halo and filamentary structure (although it remains homogeneous on large, i.e. $\gtrsim$ 100Mpc, scales). Initial conditions and cosmological parameters are given by measurements of CMB fluctuations from the WMAP and Planck experiments. Over time gravitation evolves the density field so that the initial seed overdensities are amplified, coalescing mass at overdensities by pulling in, and depleting, material from less dense regions. The (mass-)rich get richer and the (mass-)poor get poorer.

The dark matter haloes also grow by way of mergers. That is, collapsed dark matter

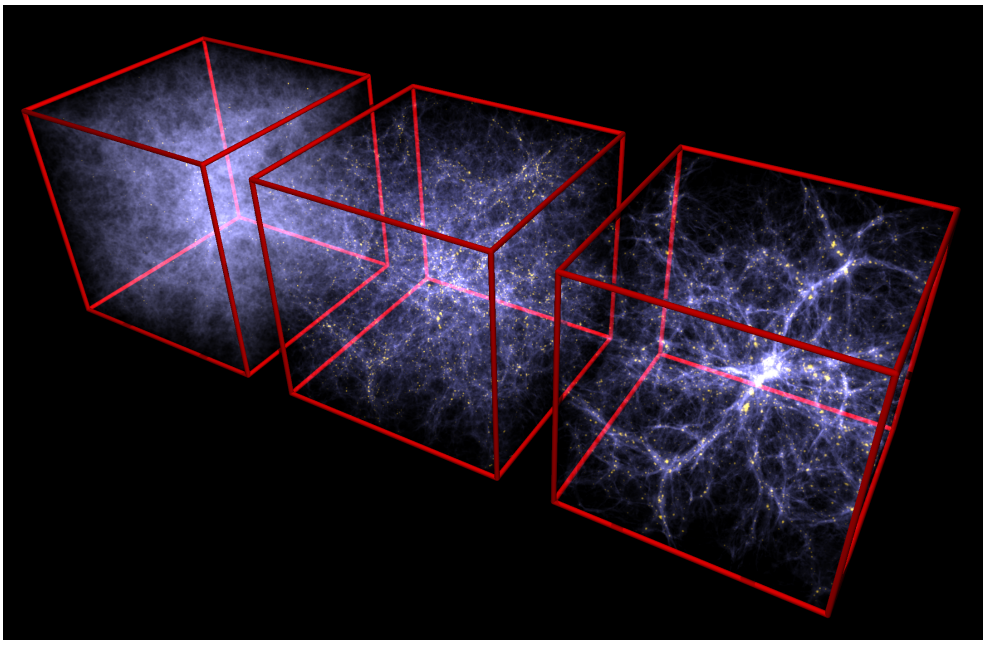

Figure 1.4: Three snapshots from the Millennium simulation, showing the evolution of the large scale structure and distribution of dark matter in the simulation box. In this figure the same simulation volume is shown in time order, from left to right. Figure by Volker Springel: `http://wwwmpa.mpa-garching.mpg.de/galform/data_vis/index.shtml`

haloes can merge with other nearby collapsed haloes to build up progressively more massive haloes. This bottom-up emergence of structure is known as hierarchical structure formation (Peebles & Dicke, 1968; Peebles, 1970, 1973, 1974, 1980), and can be represented pictorially as in Figure 1.5 (from Lacey & Cole 1993).

Ultimately for simulations to have scientific value they must make testable predictions. In this regard dark matter simulations encounter a rather fundamental problem – dark matter is not directly observable. In order to translate the dark matter simulations to something that can be observed requires the use of an observable tracer. Galaxies can serve this valuable function, providing an abundant set of luminous, observable objects which follow the underlying dark matter distribution.

Unfortunately it cannot be assumed that galaxies are unbiased tracers of the overall matter field (Kaiser, 1984). Indeed assuming a constant mass-to-light ratio to map the dark matter distribution to observable galaxies results in overpredicting the number of high and low luminosity galaxies when compared to observations.

Thus, ultimately we are forced to concede (unsurprisingly) that baryonic physical processes are important to the formation and evolution of observable structure in the universe, and must therefore be addressed head-on.

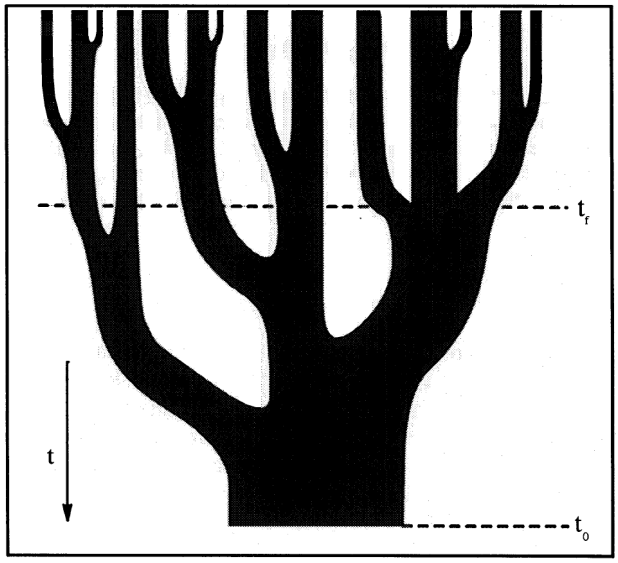

Figure 1.5: Figure from Lacey & Cole (1993). A pictorial representation of a merger tree, where time increases from the top to the bottom of the diagram. The thickness of the branches denotes the mass of the individual haloes. A horizontal slice through the tree gives the halo mass distribution at that time. The horizontal lines marked t_0 and t_f denote the present time and halo formation time respectively. Halo formation time is defined as the first time that a parent halo exceeds half of the final halo mass.

1.3 Hydrodynamical Simulations and Feedback Effects

I got a bad feeling about this
— Han Solo

As we have seen we must ultimately confront the task of simulating baryons. Essentially this means treating the gas as a fluid and evolving some initial conditions forward in time by numerically solving the corresponding fluid dynamics equations.

The field of hydrodynamical simulation is a vast, highly developed field in its own right. Within an astrophysical context there are three commonly used numerical approaches: grid-based hydrodynamics (e.g. Fryxell et al. 2000, Teyssier 2002), smoothed-particle hydrodynamics (e.g. Springel 2005, Wadsley et al. 2004), and moving-mesh codes (e.g. Springel 2011). Each approach has advantages and disadvantages, and the choice of which to use is dictated by the problem domain, researchers' prior experience, and access to code. Since we use smoothed-particle hydrodynamics (SPH) in this work we restrict our discussion to this approach.

In order to study galaxy formation and evolution we need to start with realistic initial conditions. These are obtained by first running a dark matter-only simulation, identifying a dark matter halo of interest, and tracking the particles which the halo is comprised of, or

come close to that halo, backwards through time. These particles are then replaced with dark matter and gas particles, and then again evolved forward in time with hydrodynamical codes. Via this 'zoom-in' or 're-simulation' process we are able to run hydrodynamical simulations of the formation and evolution of a single halo within its proper cosmological context without the computational expense of running a hydrodynamical simulation of the entire volume of our initial N-body simulation.

Although this already greatly reduces the scope of the problem (we simulate the particles relevant to an individual halo rather than, say, the entire universe), it is still far beyond current knowledge and computational resources to fully simulate a galaxy. To illustrate this consider a fundamental component of galaxy evolution, namely star formation. We currently do not have a complete understanding of star formation, and performing a hydrodynamical simulation of the lifecycle of even a single star is itself beyond our capabilities (even before getting into the thorny issue of simulating supernovae). It is therefore currently inconceivable to simulate the billions of stars contained in a galaxy.

Of course, the level of detail required in a simulation is related to the purpose for which the simulation is run and what we hope to learn from the simulation. When simulating rigid-body dynamics we often do not need to simulate anything below the discrete 'object' level. We routinely model fluid flows around aircraft without modelling the individual molecules comprising the air flow yet aircraft, as a rule, do not fall out of the sky. In principle, given enough resources, we could continue to simulate such systems down to an ever more microscopic level. However, given that we know how such microscopic systems behave on a macro-scale, such simulation would be an extraneous waste of resources. Sub-grid[5] prescriptions then, which statistically account for processes on a scale unresolved by simulations, are both useful and necessary.

When simulating a galaxy we make use of several sub-grid prescriptions. These models are not only below the spatial resolution but below the temporal resolution. As previously alluded to, we need to treat star formation with a sub-grid prescription whereby we typically set some criteria under which we expect star formation to occur (density threshold, convergent flows, temperature threshold), and then form entire 'simple stellar populations' rather than individual stars. The details of the populations come from population synthesis models given an initial mass function (IMF). From the mass of stars formed (given by the hydrodynamical simulations), the assumed IMF, and the population synthesis model we can calculate the number of supernovae occurring within a given stellar population particle and simulation timestep. Accounting for the effects of supernovae also requires sub-grid prescriptions. In particular we require an effective equation of state for the unresolved interstellar medium (ISM), a model for the diffusion of the metal-enriched supernovae ejecta, and a model for the injection of supernova energy into the ISM.

Much of this thesis (specifically chapters 4 and 5) concerns testing the modelling of galactic outflows, also known as galactic superwinds. These large-scale bulk outflows are driven by supernovae and blow gas out of galaxies, thereby inhibiting future star formation

[5]Though the term makes more sense for grid-based codes, this terminology is also used in reference to SPH simulations.

(this is an oversimplification - reducing star formation at early times can increase the amount of cold gas available for star formation at later times). Nevertheless the key concept is that the production of stars leads to supernovae, which then ultimately suppresses further star formation by driving superwinds and heating the surrounding gas. It is for this reason that this mechanism is known as stellar *feedback*. Such feedback, along with AGN feedback which dominates for higher mass galaxies (supernova feedback becomes less efficient above L_*, corresponding to a halo mass of $\sim 10^{13} M_\odot$), is key to suppressing star formation in order to reconcile (Benson et al. 2003) the galaxy luminosity function predicted by N-body simulations with the observed galaxy luminosity function (this discrepancy is outlined in § 1.2). As such galactic superwinds are a key component of ΛCDM, and understanding them is a prerequisite to understanding the consequences of the ΛCDM model.

In addition to regulating star formation, galactic outflows are important in that they provide an efficient means to transport galactic products, which would otherwise be trapped in galaxies, into the wider universe. Specifically galactic outflows are required for the metal enrichment of the CGM and intergalactic medium (IGM), to spread dust into the IGM, and to do the same for the intracluster and intragroup media. Taken together all of the above reasons present a compelling case for studying galactic outflows.

While this high-level overview of how to simulate a galaxy is conceptually straightforward, the uncomfortable truth is that the details of numerical hydrodynamical simulation and sub-grid models are extremely difficult to handle. This is demonstrated very clearly by the Aquila comparison project (Scannapieco et al., 2012).

The Aquila comparison project sought to investigate the effect of numerical techniques and various feedback models on the simulation of galaxy formation and evolution. To this end, 13 different gas hydrodynamics codes were run on the same halo initial conditions. Each code used its own (at the time) current-best prescriptions for gas cooling, star formation, and feedback. At z=0 the results of the runs showed significant variation between the galaxies formed with the various codes. Wide variation is seen across a variety of metrics such as galaxy morphology, stellar mass, size of the galaxy, star formation history, and the properties of the gaseous component of the galaxy. Furthermore, none of the tested codes produced realistic galaxies as measured by rotation curves and concentration. On top of this, numerical convergence was noted to be 'not particularly good' for any of the codes, and there were indications that the choice of numerical technique (SPH/AMR/Moving-mesh) was also affecting the results. However, despite these numerical issues, most of the code-to-code variation was attributed to the different feedback prescriptions adopted by each model[6].

In conclusion, our ability to reliably simulate galaxies with hydrodynamical codes is limited, with the largest differences arising from the chosen sub-grid prescriptions. It is clear then that testing and improving the sub-grid physics in our models is of paramount importance.

[6]Note however that the runs labelled G3 and AREPO in Scannapieco et al. (2012) differ only in the fact that one code is SPH and the other is Moving-mesh. That is, the sub-grid physics are the same between codes. Yet, the final stellar mass of the galaxy produced by the AREPO run is almost twice that produced in G3.

1.4 Radiative Transfer

Assume that we are able to perfectly simulate the dark matter, stellar, and gas components of a galaxy, along with their dynamics. This still does not answer the critical question: 'what does it *look* like?' This question is important because it is only by comparing the simulated observables to observational data that we can verify or falsify our models. In order to answer this question we need to address the means by which will we observe, both in reality and virtually, the galaxy. Namely, we need to simulate radiative transfer.

Due to the fact that different types of radiation propagate very differently through a medium, simulations of radiative transfer are usually limited to a single wavelength or small range of wavelengths. Therefore, the first question we need to answer is which wavelength we want to virtually observe a simulated galaxy in.

For the work described in Chapters 4 and 5 we choose to observe simulated galaxies in Lyα . This choice is motivated by several factors, but before we discuss them we need to briefly discuss the nature of Lyα radiation (for a full treatment see Harrington 1973; Neufeld 1990; Ahn et al. 2001, 2002)

The Lyman series is a series of spectral lines corresponding to transitions of electrons to and from the ground state of neutral hydrogen (HI). The transitions have ordinal labels which follow the Greek alphabet, such that the transition from the n=2 state to the n=1 state is Lyα , the transition from the n=3 state to the n=1 state is Ly-β, the transition from the n=4 to the n=1 state is Ly-γ etc. All of the Lyman series transitions correspond to the emission or absorption of UV photons.

Lyα photons have a rest frame wavelength of 1215.67Å. A Lyα photon propagates through an optically thick HI medium via resonant scattering. In this process a photon entering the medium is absorbed by an HI atom which excites the electron. After some time the electron decays back to the ground state, re-emitting a Lyα photon, and the process repeats until the photon either exits the medium or is destroyed by dust.

Since Lyα undergoes resonant scattering, the radiative transfer is sensitive to the dynamics of the gas through which the photons are propagating. This is because the resonant nature of Lyα photon scattering means that the incoming photon must be at the resonant wavelength/frequency in the *rest frame* of the absorbing atom in order to be scattered. Therefore any bulk (HI) gas motion will alter the scattering cross-section presented to the incident photon by the gas. Additionally, even the thermal motions of the HI atoms produce redshifts and blueshifts large enough to affect the Lyα scattering cross section. The thermal motion of the scattering atom also imparts a frequency shift to the emitted photon as seen in the gas frame. In summary, Lyα radiation is sensitive to the density, velocity, and temperature of the HI through which it propagates.

To illustrate this we examine some simple test cases. First we consider a monochromatic Lyα source at the centre of a spherically symmetric, static, dust-free but extremely optically thick cloud of HI gas. Figure 1.6 shows the resulting spectrum that exits the gas cloud. Note that for highly optically thick media, as is the case here, the emerging spectrum has a characteristic double-peaked profile. This is because the line-centre optical depth of the gas is so high that photons close to line centre are effectively trapped and barely

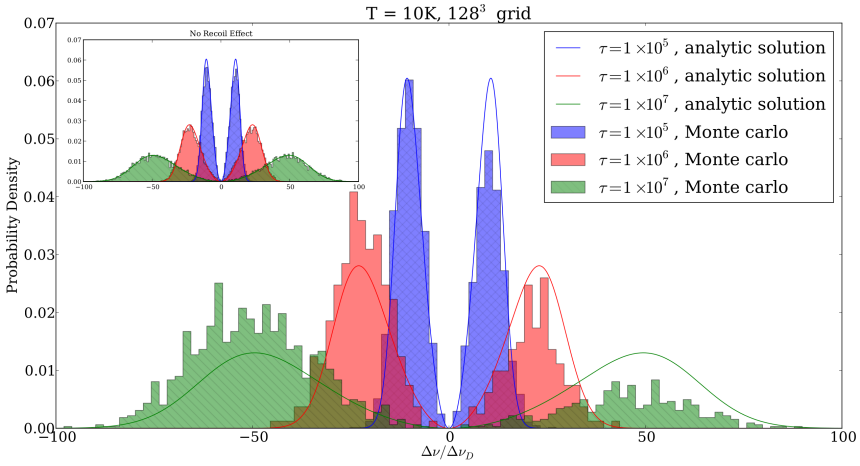

Figure 1.6: The spectrum for a monochromatic Lyα source embedded at the centre of a static sphere of HI, as measured outside the sphere. The gas has a temperature of 10K and a line-centre optical depth, τ, from the source to the edge of sphere. The main panel shows the analytic (solid lines) and numerical (filled histograms) solutions. The discrepancy is caused by the fact that the analytic solution does not account for the recoil effect. With the recoil effect artificially turned off in the numerical simulation (inset panel) the agreement with the analytic solution is excellent. Figure from Chung et al. (in prep).

propagate through physical space. It is only when a photon scatters in such a way that it attains a large frequency shift in the gas frame that it can traverse a large distance, due to the decline in the Lyα scattering cross-section of HI due to this frequency shift. It is the frequency dependence of the HI scattering cross-section which means that photons are preferentially allowed to escape the further they are from line-centre. The fact that the spectrum declines on the outer edge of the peaks is a reflection of the fact that whilst photons escape more easily when they are far from line centre, there are simply fewer of these photons which have been scattered so far out of resonance.

Figure 1.7 shows the outcoming spectrum for a similar setup, but here the gas sphere is either contracting (red data) or expanding (blue data) in a Hubble-like fashion i.e. $v_{radial} \propto r$, where v_{radial} is the radial velocity of the gas and r is the distance from the centre. If the sphere is contracting [expanding] then due to the velocity gradient the gas sees Lyα photons moving outwards as blueshifted [redshifted]. Therefore, any incident photons which initially lie redward [blueward] of line centre are blueshifted [redshifted] into resonance, once again making it hard for them to escape. Conversely, if an incident photon is initially blueward [redward] of line centre it is blueshifted [redshifted] even further out of resonance, facilitating the escape of the photon from the gas. This results in the redward [blueward] peak being depleted or, as Figure 1.7 shows, even completely erased if τ is high enough.

As described in the following section, the bulk of this thesis deals with testing galactic outflow models. For this application, where we want to probe the gas dynamics of the

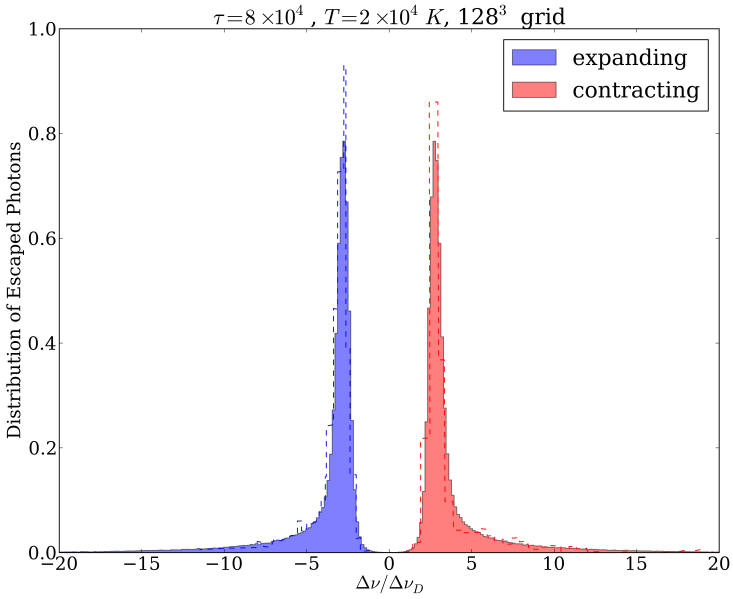

Figure 1.7: The spectrum for a monochromatic Lyα source embedded at the centre of a contracting (red) or expanding (blue) sphere of HI, as measured outside the sphere. This figure shows the numerical solutions of 2 independent Lyα radiative transfer codes: Tasitsiomi 2006 (dashed lines) and SLAF (solid, filled), the Lyα radiative transfer code developed during the course of this doctorate. Figure from Chung et al. (in prep).

CGM, the fact that Lyα is sensitive to the gas properties makes it a compelling line to observe and simulate.

Since Lyα photons interact with neutral hydrogen we also need to know the ionisation state of the gas in our simulations. Therefore a code to handle the radiative transfer of ionising radiation is also required. For this purpose we use the multi-frequency ionising UV radiative transfer code, CRASH (Ciardi et al., 2001; Maselli et al., 2003, 2009; Pierleoni et al., 2009; Graziani et al., 2013).

Finally, we note that in a cosmological context Lyα emitters have been used as a probe of reionisation (Jeeson-Daniel et al. , 2012; Dijkstra & Wyithe, 2010). In order to correctly infer the ionisation state of the IGM from observations it is necessary to know the intrinsic signal from the emitters, before it is modulated by the IGM. This provides further motivation for understanding the radiative transfer of Lyα from the galaxy, through the CGM, and into the IGM.

1.5 Putting It All Together

Put 'em together and what have you got?
bibbidi-bobbidi bibbidi-bobbidi bibbidi-bobbidi-boo
—'The Magic Song' from Cinderella [1950]

In this chapter we have outlined all of the key concepts which are needed to understand this thesis. We now explicitly define the questions we seek to address with this work.

We have seen in § 1.2 the necessity of galactic outflows to ΛCDM, at present the dominant cosmological paradigm. We have also seen that unfortunately the details of the generation of these galactic outflows are laden with uncertainties (§ 1.3). And finally, we have seen (§ 1.4) that the Lyα line has the potential to give us a window into detailed gas dynamics and properties. Combining all of the above leads naturally to the idea of applying Lyα radiative transfer to hydrodynamical models of galaxy formation and evolution in a cosmological context. In short then, the aim of this thesis is to use Lyα radiative transfer to expand our understanding of feedback, galactic outflows, and the models we use to describe them.

To this end we take a direct approach in Chapters 4 and 5 whereby we start with hydrodynamical 'zoom-in' simulations of the kind described in § 1.3. We then perform two radiative transfer post-processing steps, to account for the effects of ionising radiation and the propagation of Lyα emitted either from within the galaxy (seen in emission) or from a background source (seen in absorption). In this way we are able to produce virtual observations which can be directly compared to existing observations.

In a sense the choice of epoch at which we simulate the radiative transfer is driven by the available observations. There is little sense in simulating the observational characteristics of a galaxy for which we have no observed counterpart. Based on the available observational data we choose to simulate our galaxies at $z \sim 2-3$.

However, while it is true that the availability of observational data somewhat dictates the simulations we perform, it should not be taken as an indication that this particular epoch is randomly chosen. The fact that observations exist at $z \sim 2-3$ is no mere coincidence. $z \sim 2-3$ is the epoch of galaxy formation and the median galaxy at this time is young and star-forming. This is important since hot, young stars are strong emitters of UV radiation, which is absorbed by HI and ultimately re-emitted as the Lyα we hope to observe. Rest-frame UV and Lyα radiation at $z \sim 2-3$ is redshifted into the IR/optical part of the electromagnetic spectrum, so from a pragmatic point of view we are able to observe Lyα at $z \sim 2-3$ from large ground-based observatories. In summary galaxies at $z \sim 2-3$ exist at a cosmologically interesting epoch, and crucially have properties that allow them to be observed in Lyα .

Before we embark on such an ambitious endeavour (and we see from the uncertainties described in § 1.3 that indeed it is ambitious), we first take a step back and look at the problem from the opposite end of the complexity spectrum. In the modelling of Lyα emitting galaxies, extremely simplified 'shell-models' have been surprisingly successful in

producing the observed phenomenology (Verhamme et al., 2008; Hashimoto et al., 2015). In these models a galactic outflow is represented as a thin parameterised spherical shell of outflowing gas. It is clear that this is an unrealistic model, and yet its simplicity combined with its success in reproducing the Lyα line shape has led to widespread adoption of the model. In Chapter 3 we investigate the first observational limitations of the shell-model, and explore a simple extension which allows it to explain the observed blue bumps which pose a problem for the basic shell-model.

In conclusion this thesis presents two complementary approaches to modelling, and furthering our understanding of galactic outflows via simulations of Lyα radiative transfer.

Chapter 2

Code Details

In order to undertake the work outlined at the end of the previous chapter, and described in detail in the following two chapters, it was necessary to first develop some new codes. Here we briefly describe the nature of the two codes, and show the results of some of the verification tests we have applied to them.

2.1 Lyα Radiative Transfer (SLAF)

SLAF (Super Lyman Alpha Fighter) is a new 3D monte-carlo Lyα radiative transfer code similar to (Zheng & Miralda-Escudé, 2002; Cantalupo et al., 2005; Dijkstra et al., 2006a; Tasitsiomi, 2006; Verhamme et al., 2006; Laursen et al., 2009; Yajima et al., 2012). In order to strike a balance between efficiency, extensibility, and maintainability C++ was chosen as the implementation language, and an object-oriented design paradigm was used.

The papers cited above describe in detail how codes of this nature work. However, here we give some information specific to our code and briefly outline the critical steps involved in Lyα radiative transfer. As discussed in § 1.4 we need to use our Lyα radiative transfer code in concert with CRASH, an ionising radiative transfer code. Therefore SLAF was designed to fit into the existing CRASH pipeline, adopting the input file format of CRASH. As per CRASH the simulation is performed on a fixed grid, with each cell having a temperature, hydrogen density, ionisation fraction, and bulk velocity.

To perform the Lyα radiative transfer each source is allocated a number of Lyα photon packets, proportional to the source luminosity. Each photon packet is then tracked from the source until it exits the simulation box. The first step is to choose a random direction in which to emit the packet from the source. To propagate the photon packet from scattering to scattering an optical depth, τ, that the photon will freely stream through is randomly chosen according to the probability distribution $P(\tau) \propto e^{-\tau}$. The photon packet is then propagated in a straight line through cells, calculating the optical depth 'used up' in traversing each cell based on the frequency of the photon (shifted to the gas frame) and the temperature and HI gas density of the cell. When the optical depth of the photon packet is exhausted the photon is scattered. As the photon packet traverses each cell the

dust optical depth, τ_{dust}, is calculated based on the HI density of the cell and a specified dust-to-gas ratio appropriate to the physical scenario being simulated[1]. Thus at any point during the propagation of a photon packet the cumulative τ_{dust} is known, and whenever the photon packet intensity is required the dust attenuation factor can be applied such that $I(s) = I(0)e^{-\tau_{dust}(s)}$, where $I(0)$ is the initial photon packet intensity, $I(s)$ is the intensity evaluated at the point where the photon has traversed a path, s, and similarly $\tau_{dust}(s)$ is the dust optical depth experienced traversing s. Note that $\tau_{dust}(s)$, and consequently $I(s)$, is path-dependent since at all points on the path evaluation of τ_{dust} depends on the local dust density as given by the local gas density and the adopted gas-to-dust ratio.

For each scattering event the thermal motion of the atom to be scattered off needs to be generated. This is done by picking the thermal velocity from a non-trivial distribution, details of which are given in Zheng & Miralda-Escudé (2002). In the current version of SLAF scattering is isotropic. The direction of the scattering together with the thermal motion of the atom determines the gas-frame frequency shift of the scattered photon packet. With the new direction and gas-frame frequency the loop begins again and the photon packet is propagated to the next scattering event.

To generate surface brightness images we follow the method of Yusef-Zadeh et al. (1984). We place virtual detectors on each face of the simulation box. For each scattering event and for each of the six detector directions we create a secondary photon packet having the energy that a photon would have if it had scattered in the direction of the associated detector. This photon packet is then propagated in the direction of the detector, again being attenuated as in the normal photon propagation. Thus, when this secondary photon packet reaches the detector it has the energy the original photon packet would have had *if* it had scattered towards the detector, weighted by the fraction of the scattered photon packet energy that would reach the detector. This value is added to the cumulative total for the pixel that the secondary photon packet reaches (i.e. the pixel that the scattering event is projected onto). In this way at the end of the radiative transfer simulation each detector contains a relative brightness map. In order to convert these relative brightness maps into absolute value maps we sum the total energy of all photons which exit the simulation cube and assign it proportionally over the maps.

In order to verify our implementation we compare the SLAF output with the test case in Dijkstra et al. (2006a), for which there is an approximate analytic solution (Dijkstra et al., 2006a, Equation 9) which we repeat here:

$$J(x) = \frac{\sqrt{\pi}}{\sqrt{24}a\tau_0}\left\{\frac{x^2}{1+\cosh\left[\sqrt{2\pi^3/27}(|x^3|/a\tau_0)\right]}\right\} \tag{2.1}$$

[1]SLAF allows the relation between dust density and the absorption cross-section to be specified, along with the dust-to-gas ratio. In practice it is enough for our work in Chapters 4 and 5 that these two parameters are simultaneously chosen to satisfy the τ_{dust}-N_H relation given in Garel et al. (2012): $\tau_{dust}(\lambda) = (A_\lambda/A_V)_{Z_\odot}(Z/Z_\odot)^{1.35}\left(N_H/2.1\times10^{21}\right)(1+z)^{-\frac{1}{2}}$, where $(A_\lambda/A_V)_{Z_\odot}$ is the solar metallicity extinction curve from Mathis et al. (1983), Z is the gas metallicity, z is the redshift, and N_H is the HI column density.

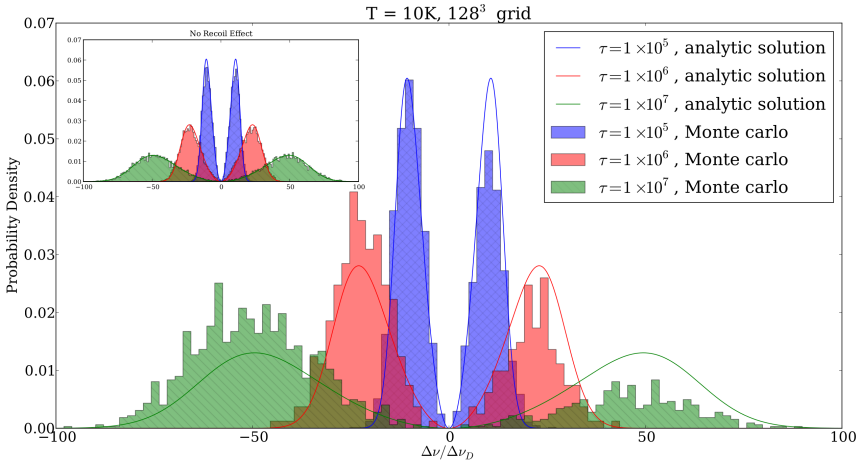

Figure 2.1: Same as Figure 1.6. Comparison of the spectrum calculated by SLAF (filled histograms) with a known analytic approximation (solid lines). The apparent discrepancy is due to the omission of the 'recoil effect' in the analytic approximation. The inset plot shows the output of our code with the 'recoil effect' disabled. Figure from Chung et al. (in prep).

where $a = A_{21}/4\pi\Delta\nu_D$, A_{21} is the Einstein A-coefficient for the Lyα transition, τ_0 is the line-centre optical depth from the centre to the edge of the sphere, $x \equiv (\nu - \nu_0)/\Delta\nu_D$, $\Delta\nu_D = v_{th}\nu_0/c$, $v_{th} = \sqrt{2k_bT/m_p}$, k_b is the Boltzmann constant, T is the gas temperature, and m_p is the mass of a proton.

In brief this test puts a single, monochromatic Lyα source at the centre of a homogeneous, static, spherically symmetric distribution of HI at 10K. In each instance the uniform gas density is chosen so that the line-centre optical depth from the source to the edge of the gas sphere is $\tau_0 = \{10^5, 10^6, 10^7\}$.

The outcoming spectrum from SLAF is compared to the analytic solution in Figure 2.1. SLAF is shown to be in excellent agreement with the analytic solution with the exception of a systematic amplification of the red peak and suppression of the blue peak compared to the analytic solution. We understand this is to be due to the 'recoil effect' which is ignored in the analytic approximation for simplicity. Figure 2.1 (inset) shows the output of SLAF with the 'recoil effect' disabled, and demonstrates almost perfect agreement with the analytic solution.

To test that SLAF can also correctly handle dynamic scenarios we make a comparison to the results presented in Tasitsiomi (2006) (Figure 3, right panel). Since we want to make a direct comparison we set up an identical scenario comprising of a single central Lyα source embedded in a spherically symmetric HI cloud. The HI cloud has a column density of 2×10^{18} cm^{-2} from the centre to the edge of the cloud, and a Hubble-like velocity gradient where the gas velocity scales proportional to radius up to a maximum inflow/outflow velocity of 200kms^{-1} at the edge of the cloud.

Figure 2.2 shows our results overlaid on the results from Tasitsiomi (2006). The results

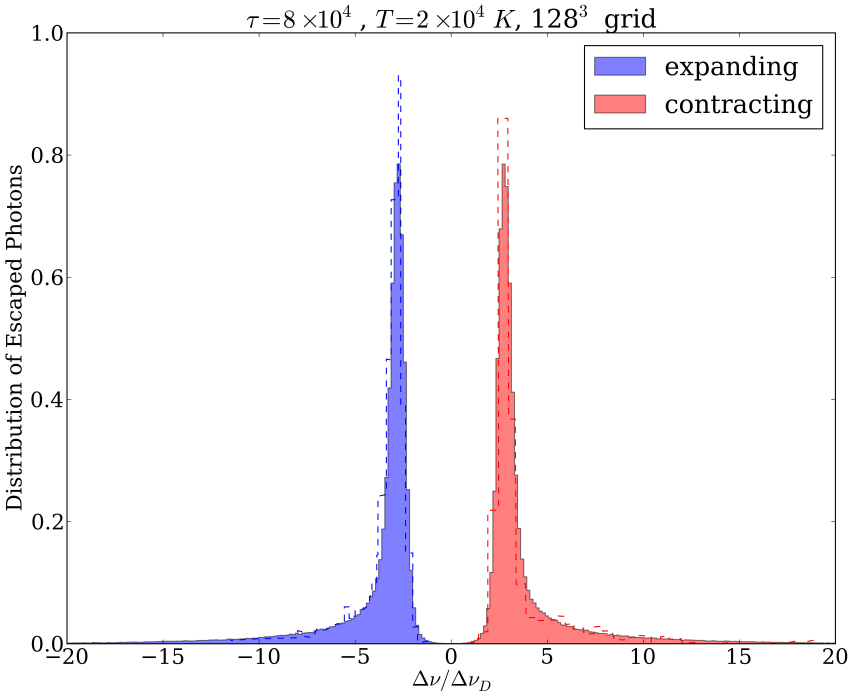

Figure 2.2: Comparison of the spectrum calculated by SLAF (histograms) with the monte carlo code (dashed lines) from Tasitsiomi (2006). Blue histograms/lines show a homogeneous expanding HI gas sphere, red histograms/lines show a homogeneous contracting gas sphere.

in blue(red) are for the expanding(contracting) cloud case. The solid histograms show our SLAF results and the dashed-line histograms show the Tasitsiomi (2006) results. There is excellent agreement between the two codes.

2.2 Absorption along Sightlines (LAF)

LAF (Lyman Alpha Fighter) is the second new code. It is used to calculate Lyα absorption line profiles for sightlines through 3D volumes of arbitrary gas distributions and associated velocity fields. Absorption equivalent widths for these sightlines can then be derived from the absorption line profiles.

LAF calculates the absorption line profile by casting rays through the simulation box from one face to the opposite face, such that each ray is perpendicular to the source face and corresponds to a sightline of interest. Each ray starts with an associated flat spectrum, represented by an array of intensity values for a set of frequency bins (the width of the bins is variable, depending on a parameter which specifies the resolution of the spectrum to be calculated). As the ray traverses the simulation box, the physical properties of each cell through which it passes are used to calculate the optical depth of the cell to the incident ray as a function of frequency. From this the absorption line profile for the cell, *in the cell*

frame, is calculated. This line profile is then redshifted or blueshifted based on the bulk velocity of the gas in the cell, and applied to the spectrum of the incident ray (which we define to be in the lab frame). By doing this for all cells through which the ray passes the cumulative absorption line profile is calculated.

In order to test LAF we construct a simple test case. We create a sphere of HI gas with radius $r = 25.6$kpc, uniform HI number density $n = 1 \times 10^{-10}\text{cm}^{-3}$, and a temperature of 5000K on a 512^3 grid.

We then use LAF to calculate the absorption line profiles, and hence absorption equivalent widths, for lines of sight through the simulation volume. For such a simple geometry the equivalent width (EW) of the absorption lines along a line of sight at impact parameter, b, is given by the following analytic expression (derived and expanded in Appendix B):

$$\text{EW}_\nu(b) = \int_{\nu_{min}}^{\nu_{max}} [1 - I(\nu, b)] d\nu \tag{2.2}$$

where

$$I(\nu, b) = e^{-2n\sqrt{R^2 - b^2}\sigma(\nu)} \tag{2.3}$$

and $\sigma(\nu)$ is the cross-section at frequency ν, n is the HI density, and R is the radius of the gas sphere.

Figure 2.3 shows a comparison between the LAF output and the analytic solution. The output of LAF is in excellent agreement with the analytic solution, with a small deviation arising only from the fact that the sphere has been quantised onto a grid. This can be seen in the fact that the agreement gets worse as the impact parameter approaches the radius of the sphere. This can be understood by realizing that as a sightline moves closer to the edge of the sphere the segment intersecting the gas sphere decreases. Therefore since the cell size is fixed the relative error due to the discretisation onto a grid grows as the traversed optical depth decreases with higher impact parameter.

The static sphere test is useful for testing that our code gives correct results for a variety of densities, but is limited to static gas distributions. In order to test that LAF is handling non-static gas correctly we perform the following simple test. A slab of HI is set up with a steep (fixed) velocity gradient across it. This allows us to take advantage of the fact that our simulation is performed on a discrete grid. We shoot a ray through the gas, in the direction of the velocity gradient. As long as the velocity difference from cell to adjacent cell is large compared to the Doppler width, the absorption line profiles resulting from each cell should not overlap when superimposed. Therefore the expected result is a series of repeated absorption line profiles, with the line centres spaced apart equally at a frequency shift corresponding to the velocity shift from cell to cell. Figure 2.4 shows the result from LAF. The red vertical lines show the expected shift in the line profile for a velocity shift of 10km/s from cell to cell. The blue line shows the line profile output from LAF, in essentially perfect agreement with the expected velocity shift.

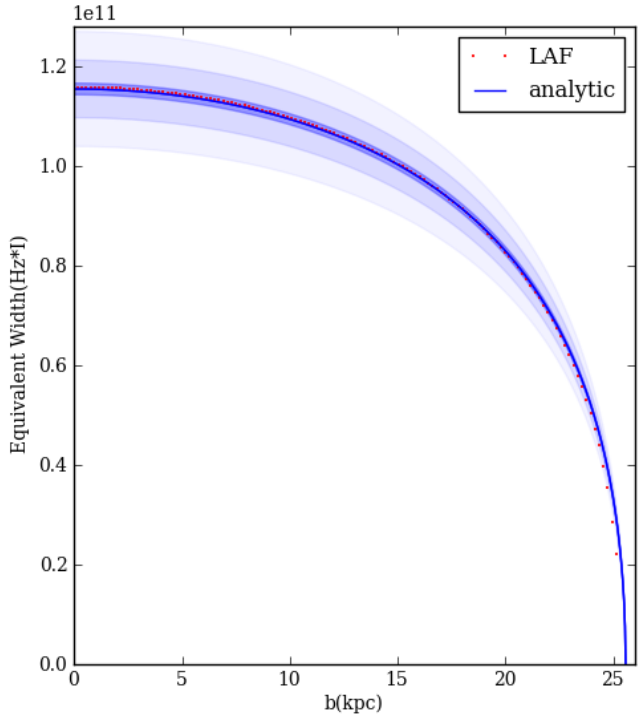

Figure 2.3: Equivalent width vs impact parameter for a 25.6kpc static, homogeneous HI sphere. The blue line is the analytic solution while the shaded regions show 1%, 5% and 10% errors. The red dotted line shows the output from LAF to be in excellent agreement with the analytic solution.

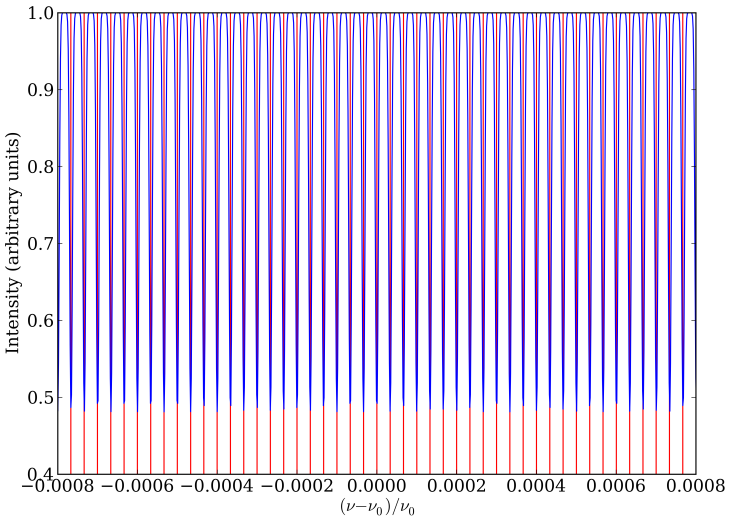

Figure 2.4: The resulting spectrum from the velocity gradient test. The blue line shows the output from LAF, whereas the vertical red lines show the expected positions of the velocity-shifted line centres.

Chapter 3

Fermi-like Acceleration of Lyman-α

This chapter is adapted from (Chung et al., 2016) by Andrew S. Chung, Mark Dijkstra, Benedetta Ciardi, and Max Gronke

We explore the impact of Fermi-like acceleration of Lyα photons across shock fronts on the observed Lyα spectral line shape. We first confirm the result of Neufeld & McKee (1988) that this mechanism gives rise to extended blue wings which may have been observed in some radio galaxies. Our Monte-Carlo radiative transfer calculations further show that in a minor modification of the shell-model, in which we add an additional static shell of hydrogen, this process can naturally explain the small blue bumps observed in a subset of Lyα emitting galaxies, which have been difficult to explain with conventional shell-models. Blue bumps can be produced with an additional column density of static hydrogen as small as $N_{\rm HI}^{\rm static} \ll N_{\rm HI}^{\rm shell}$, and typically occur at roughly the outflow velocity of the shell. In our model, the spectra of so-called 'blue-bump objects' might reflect an evolutionary stage in which the outflows regulating the escape of Lyα photons are still engulfed within a static interstellar medium.

3.1 Background

Observations indicate that the escape of Lyα photons is facilitated enormously by the presence of outflowing interstellar gas (e.g. Kunth et al., 1998). Scattering off these outflows effectively Doppler shifts Lyα photons out of resonance into the red wing of the absorption line profile. In models of this process the outflow is often represented with a geometrically thin shell of gas (e.g. Verhamme et al., 2006; Schaerer et al., 2011). In spite of its simplicity, this so-called 'shell-model' has been very successful at reproducing observed spectra line shapes (Verhamme et al., 2008; Hashimoto et al., 2015).

However, the shell-model has recently been shown to have difficulties, especially in reproducing the strength of 'blue bumps' in a subset of observed spectra (Kulas et al., 2012; Chonis et al., 2013; Hashimoto et al., 2015). Adams et al. (2009) studied what may

be viewed as an extreme example of blueshifted Lyα emission in which the Lyα line of spatially extended Lyα emission around a radio galaxy as a whole is blueshifted relative to the observed 21-cm absorption, with emission extending beyond 1500 km s^{-1} into the blue wing of the line profile.

Blueshifted Lyα radiation arises naturally when the photons scatter through optically thick inflowing gas (Zheng & Miralda-Escudé, 2002; Dijkstra et al., 2006a,b). Indeed, Adams et al. (2009) show that their data can be reproduced if more than $10^{12}M_{\odot}$ of cold inflowing gas is present. This large mass of cold neutral gas inside a massive dark matter halo ($M \gtrsim 10^{13}M_{\odot}$) though, is at odds with expectations from theory, which predicts that the gas should be predominantly accreted in the hot-mode (e.g. Kereš et al. 2005; Figure 7 of Dekel & Birnboim 2006).

An alternative process which gives rise to blueshifted Lyα emission is described by Neufeld & McKee (1988; henceforth NM88), who show that scattering across a shock front can give rise to such a blueshift. NM88 presented analytical calculations, and therefore were forced to adopt simplifying assumptions, namely that the shock-crossing acceleration process is modelled by Lyα photons bouncing between two partially transparent, frequency-preserving mirrors. In reality, however, the frequency of a Lyα photon is not preserved at each 'reflection'. Instead, the frequency of the photon diffuses through frequency-space as the photon is scattered by the material on either side of the shock front, in turn changing the optical depth of the gas to the photon. In addition, they assume that photons cross the shock front isotropically. Finally, the analysis presented by NM88 was restricted to a simplified geometry of two adjacent semi-infinite slabs.

In spite of an increasing prevalence of Lyα radiative transfer Monte-Carlo codes (Zheng & Miralda-Escudé, 2002; Cantalupo et al., 2005; Dijkstra et al., 2006a; Tasitsiomi, 2006; Verhamme et al., 2006; Laursen et al., 2009; Yajima et al., 2012), they have so far not been used to investigate this Fermi-like acceleration mechanism (while acknowledging the inaccuracy of doing so, we henceforth refer to this as 'Fermi acceleration' for brevity). The goals of this chapter are two-fold: (*i*) numerically study Fermi acceleration of Lyα photons across a shock front without the simplifying assumptions that were required in the analytical treatment, and (*ii*) investigate whether this mechanism can help to reproduce observed blue bumps in spectra that have been difficult to explain with conventional shell-models. The outline of this chapter is as follows: § 3.2 describes the basic acceleration mechanism, § 3.3.1 details the simulations which we use to investigate the physical mechanism, § 3.3.2 presents the output of our numerical experiments, § 3.4 discusses the applications of our results, and in § 3.5 we summarise our findings.

3.2 Blueshifting Mechanism

Lyα photons are resonantly scattered by HI, making the Lyα radiative transfer complex. The propagation of Lyα photons is affected by both bulk gas motion and, in detail, microscopic motion of individual hydrogen atoms. The result is that Lyα photons undergo a random-walk like motion in both physical and frequency-space (see Dijkstra, 2014, for a

detailed review).

As Lyα photons diffuse away from the resonance frequency, their mean free path - and hence their escape probability - increases. As a result, the spectrum of Lyα photons emerging from static optically thick media consists of two peaks, which are distributed symmetrically around the resonance frequency (as frequency diffusion can occur to lower and higher energies with equal probability[1]). When the scattering medium is contracting [expanding] however, the frequency diffusion preferentially occurs towards higher [lower] frequencies. As a result, the spectrum of Lyα photons emerging from contracting optically thick media is blueshifted. An alternative way to see this is that the converging flow of contracting gas is doing work on the Lyα photons as they scatter outwards, which increases their mean energy (Zheng & Miralda-Escudé, 2002; Dijkstra et al., 2006a,b).

The Fermi acceleration mechanism described in NM88 invokes similar converging flows of gas. The mechanism is illustrated with a simplified geometry in Figure 3.1, which shows two semi-infinite adjacent slabs of neutral gas. The slab on the left is moving to the right with velocity $v_{\rm s}$ into a stationary slab. If $v_{\rm s}$ exceeds the sound speed in the left slab, then the two slabs are separated by a shock front. When a Lyα source is in the vicinity of this shock front we expect some Lyα photons to diffuse through space and cross the front. When a photon traverses the shock front, the Doppler boost will impart a blueshift in the local gas frame. Partially coherent scattering off atoms with thermal velocity v_{th} mostly preserves the blueshift of this photon: for a photon with frequency $x_{\rm s} \sim v_{\rm s}/v_{\rm th}$ each scattering event pushes back the photon to the line centre by an average amount $-1/|x_{\rm s}|$ (Osterbrock, 1962). When the photon scatters back across the shock front before this 'restoring force' has returned the photon to line center, the blueshift of the photon increases with each shell crossing. This is what NM88 referred to as 'Fermi acceleration'.

For Fermi acceleration of a Lyα photon to occur we require it to scatter after crossing the shock front. The optical depth to a Lyα photon that is blueshifted to v_b through a slab of gas with column density $N_{\rm HI}$ is $\tau(v_{\rm b}) \sim 0.6(N_{\rm HI}/10^{19}\,{\rm cm}^{-2})(v_b/200\,{\rm km\,s}^{-1})^{-2}$. The condition $\tau(v_b) \geq 1$ thus translates to $N \geq 1.6\times10^{19}(v_b/200{\rm km\,s}^{-1})^2\ {\rm cm}^{-2}$. Inversely, a photon can be Fermi accelerated to a maximum blueshift of $v_b \sim 155(N_{\rm HI}/10^{19}\,{\rm cm}^{-2})$ km s^{-1}, *in the frame of the scattering medium.* After scattering, the Lyα photon can get an additional Doppler shift $\sim v_{\rm s}$ depending on whether the scattering medium is moving or not.

3.3 Fermi Acceleration in Monte-Carlo Simulations

In this Section we describe the setup and results of the numerical simulations which we use to study the Fermi acceleration mechanism briefly summarised in § 3.2.

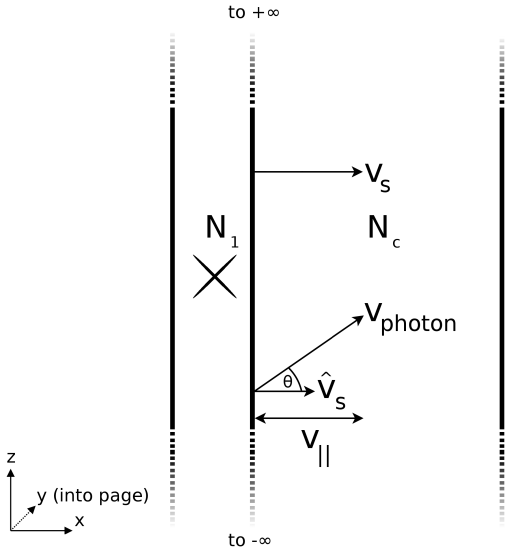

Figure 3.1: Schematic showing the setup used for studying the basic acceleration mechanism (see § 3.2 and § 3.3.1). The Lyα source is denoted by the cross in the left slab. The left slab has HI column density N_1, and the right slab has HI column density N_c. The gas in the left slab moves with velocity v_s, photons cross the shock front with velocity v_{photon}, and the angle between v_s and v_{photon} is labelled θ. $v_{\parallel}$ is then the projection of v_{photon} onto the unit direction vector of the shock front, $\hat{v}_s$.

3.3.1 Simulation Setup

We use modified versions of the SLAF 3D Monte-Carlo Lyα radiative transfer code (Chung et al., in prep), as well as the code described in Dijkstra et al. (2006a). The vanilla version of SLAF is a grid-based code, able to handle arbitrary gas and source distributions within a finite volume of 3D space. For this work we have modified it, removing the dependence on an underlying grid-structure, so that the code can handle the specific setup described in NM88.

We choose to focus on the geometry discussed in NM88 to facilitate a straightforward comparison. Explicitly, there are two distinct regions of gas, represented by slabs of infinite width (y) and height (z). The slab on the left side represents outflowing gas, and has a bulk velocity of $v_{\rm s} = 400$ km s^{-1} from left to right. The HI column density across the outflowing slab is $N_1 = 1 \times 10^{20}{\rm cm}^{-2}$. The slab on the right side represents static gas with an HI column density $N_c = 10^{21}$ cm$^{-2} \gg N_1$. A shock front exists at the interface of the two slabs. Lyα photons are emitted at line centre in the rest frame of the outflowing slab on the left, which represents Lyα emission powered by star formation triggered by the passing shockwave (see NM88). We ignore dust throughout the analysis in § 3.3.2, which simplifies the interpretation of our results and does not affect our main results. We discuss

[1]For low gas temperatures ($T = 10$ K) the energy deposited in the recoiling scattering atom becomes important, and the red peak is enhanced.

the impact of this assumption separately when we discuss applications of our results in § 3.4.1 and § 3.4.2. Figure 3.1 shows the described experimental setup.

Figure 3.2 was produced from a run tracking 5×10^4 crossing events, while Figures 3.3 and 3.4 were produced from runs of 5×10^4 photons. In all cases the results were checked for convergence, and found to already be converged with much lower photon counts.

3.3.2 Results

3.3.2.1 Shock Crossing Statistics

Table 3.1 shows the fraction of photons escaping the simulation volume, f_n, as a function of the total number of shock crossings, n. The second and third columns indicate the f_n found in our simulations and the analytic estimate, respectively. The following paragraphs explain how we obtained these analytic estimates.

Photons emitted in the centre of the outflowing slab are equally likely to leave this slab on the left or right hand side. We therefore expect $f_0 = 1/2$ of all photons not to cross the shock front, which is in agreement with the simulated fraction.

Photons that enter the static slab for the first time are Doppler boosted by an average amount $v_b = 2v_s/3$ (see § 3.3.2.2). The probability of these photons being transmitted all the way through the slab is given by $T = 4/(3\tau_i)$, where τ_i denotes the optical depth of the entire static slab to photons that enter at frequency x_i (in the frame of the slab, see Neufeld 1990; note that this transmission probability only applies when the optical depth of the slab to incoming photons is $\gg 1$). Substituting numbers yields $T \approx 0.051 v_{b7}^2 / N_{\rm HI}^{20}$, where $N_{\rm HI}^{20}$ denotes the HI column density of the slab in units of 10^{20} cm^{-2}, and v_{b7} denotes v_b in units of 100 km s^{-1} (notation adopted from NM88). We then get $f_1 = (1 - f_0) \times T \approx 0.018$, which compares reasonably to the fraction $f_1 = 0.023$ we found in our simulation (exact agreement is not expected as photons cross the slab over a range of frequencies).

Photons that cross the front twice are back in the low column density outflowing slab, and appear Doppler shifted by an average amount $v_b \sim 4v_s/3$ in the slab frame (see § 3.3.2.2). The slab will appear optically thin to these photons, and we estimate the transmission probability from $T \sim \exp(-2\tau_i) + 0.5[1 - \exp(-2\tau_i)]$. Here, $\tau_i = N_1 \sigma_\alpha(v_b)$ denotes the optical depth through the outflowing shell for the incoming photons, where $\sigma_\alpha(v_b)$ is the Lyα absorption cross-section at v_b. The factor of 2 in the exponent accounts for the fact that the photons enter the slab under an angle θ with probability $P(\cos\theta) \propto \cos\theta$ (see § 3.3.2.2). The transmission probability is thus the sum of the probability (*i*) that photons are transmitted directly through the slab (the $\exp(-2\tau_i)$ term), and (*ii*) that photons that do scatter eventually escape the slab without crossing the shock again. For the latter, we assume that photons are equally likely to escape on the left and right hand side after scattering, which seems reasonable given that the slab appears optically thin on average. Substituting numbers we obtain $f_2 = (1 - f_0 - f_1)T \sim 0.28$, which agrees with the simulation result.

With $T = 4/(3\tau_i)$ and $v_b \sim 2v_s$, we estimate $f_3 = (1 - f_0 - f_1 - f_2) \times T \sim 0.068$, 48.5% larger than what we obtain from the simulation. We attribute this discrepancy to

Table 3.1: Fraction of photons undergoing n successive shock crossings before exiting the simulation (middle column). The right column refers to the same fraction calculated with an analytic method.

total # shock crossings (n)	fraction (f_n)	analytic estimate
0	0.50820	0.5
1	0.02299	0.018
2	0.26215	0.28
3	0.04578	0.068
4	0.12525	...
5	0.01283	...
6	0.01944	...
7	0.00149	...
8	0.00168	...
9	0.00008	...
10	0.00011	...

the approximation $T = 4/(3\tau_i)$ breaking down at larger shifts from line centre.

The previous analysis allows us to understand quantitatively the simulation results. More photons escape after undergoing an even number of shock crossings than odd. This is simply a reflection of the fact that in our experimental setup we have specified that $N_c \gg N_1$.

As we pointed out in § 3.1, it is theoretically possible that a Lyα photon that has been Fermi accelerated into the line wing can subsequently diffuse back into the core of the line prior to escaping or crossing the shock front again. When this happens, any memory of previous Fermi acceleration is erased. However, our numerical simulations indicate that in practice this almost never occurs, and therefore we can identify odd [even] numbered contiguous shock crossings with photons which exit on the right [left] side of the simulation.

3.3.2.2 Velocity Shift vs Shock Crossing Number

We denote the angle at which a photon crosses the shock front with θ (see Figure 3.1). For an isotropic distribution of shock crossing directions we have $P(\mu) = 1$, where $\mu \equiv \cos\theta$. The average Doppler shift experienced by a photon as it crosses the shock front is $v_b = v_s \int_0^1 d\mu\, \mu P(\mu) = v_s/2$. After l crossings the photon experiences an average blueshift of $lv_s/2$ (NM88).

With our code we are able to track individual photons as they cross the shock front, which allows us to directly measure $P(\mu)$. Figure 3.2 shows the number count, $N(\mu)$, for photons that cross the l^{th} time for various l. $P(\mu)$ only differs from $N(\mu)$ by a normalisation factor. Here, l refers to the current (rather than final) contiguous crossing count of a particular photon as it crosses the shock front. A photon which finally exits the simulation after n crossings is represented n times (once for each crossing). Figure 3.2 shows clearly

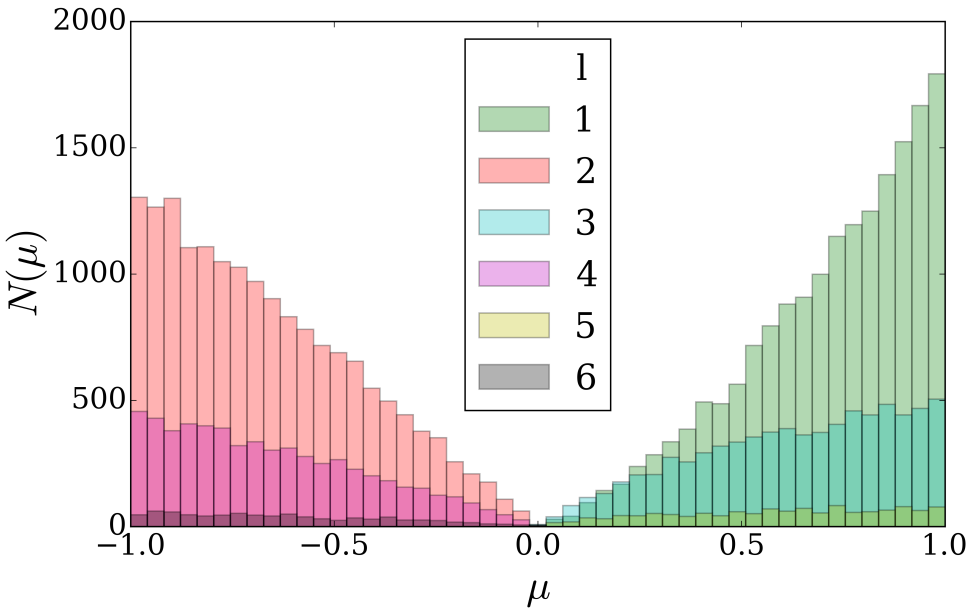

Figure 3.2: Distribution of the projection of the shock-crossing angle for various values of l. For each value of l the plotted distribution includes all photons that make an l^{th} crossing regardless of what their final crossing count, n, is.

that the distribution of photon crossing projections is not isotropic[2]. Instead, we find that $P(\mu) \propto \mu$. This distribution has been found in previous analyses (e.g. Ahn et al., 2001; Garavito-Camargo et al., 2014), and its origin is discussed in Appendix A.1. For this distribution we expect $v_s \int_0^1 d\mu\, \mu P(\mu) = 2v_s/3$ (see also Figure A.1), and hence that a photon experiences a Doppler boost $2lv_s/3$ after l shock crossings[3].

Figure 3.3 shows the velocity shift v_b of exiting photons as a function of the total number of shell crossings n, in the frame of the gas into which the photon is crossing (i.e. for odd [even] numbered shock crossings the velocity shift is given in the static [outflowing] frame). Results from our Monte-Carlo simulations are represented by the data points. The red dashed line shows the analytic result under the assumption that photons cross the shock front isotropically (as in Neufeld & McKee, 1988), while the green solid line shows the analytic result for photons crossing the shock front according to $P(\mu) \propto \mu$. For $n \leq 3$ the simulation results follow the $\frac{2lv_s}{3}$-relation (where here n is a good proxy for l), after which it approaches the isotropic shock-crossing case. The reason for this transition is that as the photons get blueshifted further into the wing of the line profile their mean free path

[2]In Figure 3.1 $\mu > 0$ [$\mu < 0$] for photons that cross the shock front from left-to-right [right-to-left]. In Figure 3.2 only photons on an odd [even] shock crossing contribute to $N(\mu)$ for positive [negative] values of μ.

[3]Note that if a photon crosses the shock front at an angle θ, then the total hydrogen column density to the edge of the slab is $N_{\rm HI}/\mu$, where $N_{\rm HI} = N_1$ [$N_{\rm HI} = N_c$] for the slab on the left [right] in Figure 3.1. The angle-averaged column to the edge of the slab is given by $\langle N_{\rm HI} \rangle = \int_0^1 d\mu\, P(\mu) N_{\rm HI}/\mu = 2N_{\rm HI}$, as used in § 3.2.1.

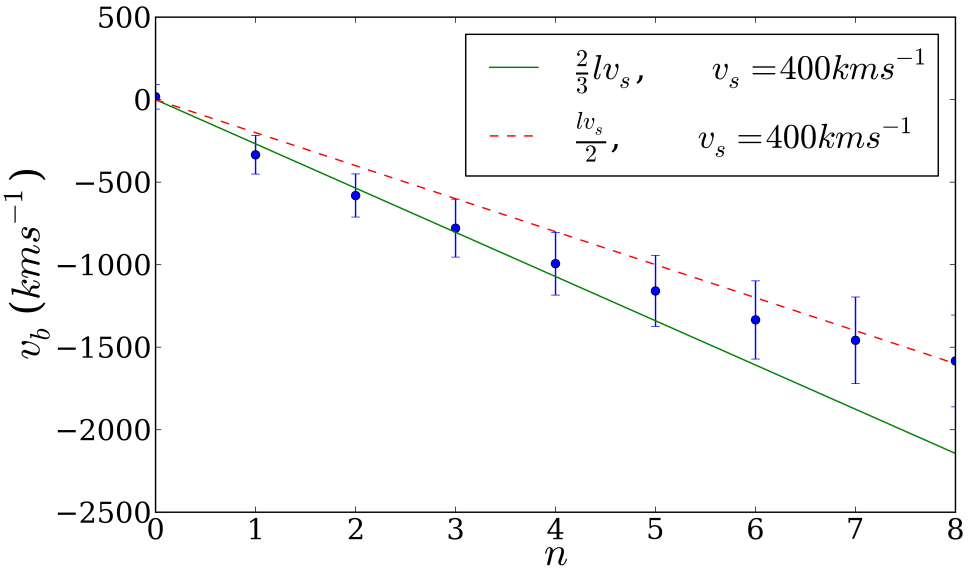

Figure 3.3: Mean velocity offset for photons exiting with n shock crossings as found in our simulations (points), as measured in the frame of the gas. The error bars show the standard deviation within the distribution. The dashed red line shows the $\frac{1}{2}lv_s$-relation which applies to isotropic shock crossings, whereas the green line refers to the $\frac{2}{3}lv_s$-relation which applies to anisotropic shock-crossings of the form $P(\mu) \propto \mu$ (see text).

increases. As we discuss in detail in Appendix A.1, we expect a transition to isotropic shock crossing when the mean free path becomes comparable to the thickness of the slab.

Finally, Figure 3.4 shows the spectrum of photons emerging from the slab, as measured in the lab frame. The result can be easily understood in terms of our previous analysis. More specifically, the total spectrum (black histogram), can be explained as follows.

- The peak at $v_b \sim 300$ km s^{-1} is composed primarily of the $f_0 = 50\%$ of all photons that did not cross the shock front. These photons diffused outward of the left slab, and would have emerged with a characteristic double peaked emission line profile centered around the resonance frequency (e.g. Adams, 1972; Harrington, 1973; Neufeld, 1990). However, when we Doppler boost these photons back into the lab frame this double peak is diluted by the fact that the photons escaping the slab do so at various angles. The Doppler boost is dependent on the exit angle, with an average Doppler boost of $\langle\mu\rangle v_s \sim 270$ km s^{-1}. The grey crosshatched histogram shows the spectrum of these photons.

- The small fraction of photons that escape after a single shock crossing, $f_1 \sim 2\%$, is indicated by the purple hatch-filled histogram and escapes with a mean blueshift of $\sim 2v_s/3 \sim 300$ km s^{-1} (also see Figure 3.3, these photons escape from the static shell, and no additional Doppler boost into the lab frame is required).

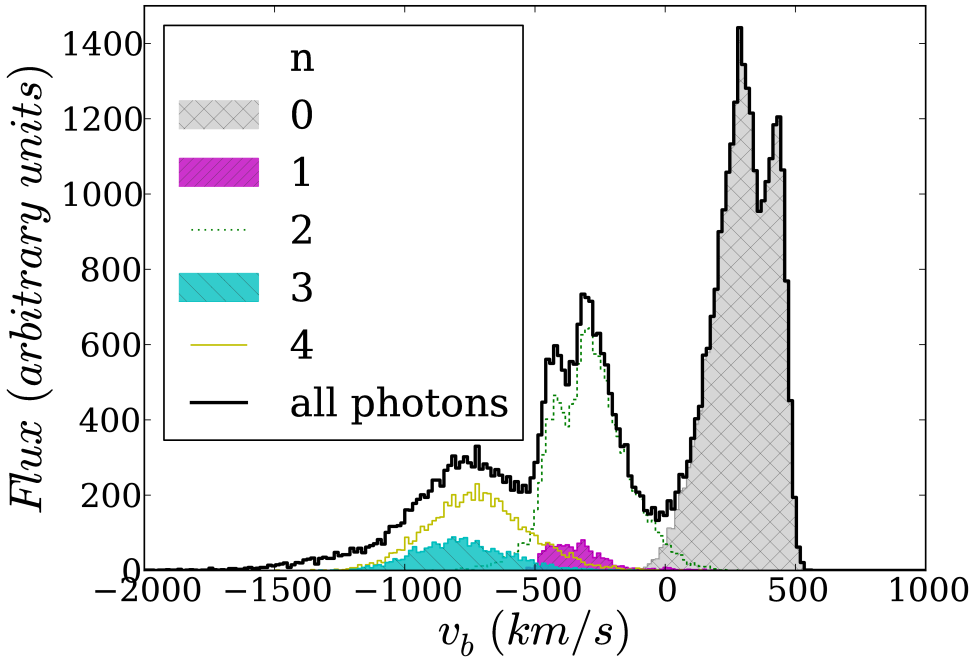

Figure 3.4: Spectrum of photons exiting the slabs with a single source, lab frame. n is the number of shock crossings that the photon undergoes prior to exiting the simulation (see § 3.3.2.1).

- The photons that escape after two shock crossings (green dotted line histogram) account for $f_2 \sim 26\%$ of the total. Figure 3.3 shows that these photons have accumulated a blueshift of $\sim 4v_s/3 \sim 530$ km s^{-1} in the frame of the outflowing gas. A Doppler boost back into the lab-frame transforms these photons back to a blueshift of $\sim 2v_s/3 \sim 300$ km s^{-1}.

A similar reasoning applies to photons that escape after a larger number of crossings, and shows that it is easy to understand the shape of the total spectrum: the redshifted peak (grey crosshatched histogram) at $v_b \sim 2v_s/3$ is composed of photons that did not cross the shock front. The first peak blueward ($v_b < 0$) of the Lyα resonance is at $v_b \sim -2v_s/3$ and consists of photons that crossed the shock front $1-2$ times. Similarly, the second blueward peak at $v_b \sim -4v_s/3$ consists of photons that crossed the shock 3-4 times, etc. The prevalence of these peaks depends on a number of factors, including v_s, N_1, N_c, and also the distribution of Lyα sources relative to the HI gas.

Finally we note that our results here are in excellent agreement with the analytic estimate of the maximum blueshift detailed in § 3.2. Setting $N = 10^{21}$ cm^{-2} in the analytic formula gives $v_b \approx 1600$ km s^{-1}, which lies in the blue tail of Figure 3.4.

3.4 Applications

We now explore, by means of our Monte-Carlo codes and assuming some simple models, whether the Fermi acceleration mechanism can provide a viable explanation for some observed systems.

3.4.1 Extended Blue Wings in Radio Galaxies

The analysis by NM88 was motivated by observations of 3C326.1, a radio source with a Lyα spectral line profile with emission extending far into the blue wing of the line profile (Djorgovski 1988, Strauss et al. 1987). NM88 modelled 3C326.1 with a shock front that was propagating into a collection of dense clumps. The shock passage triggers star formation and Lyα emission inside the post-shock gas. The subsequent radiative processes in each clump proceed just as in the slab models we discussed previously. The details of the emerging spectrum depend on the distribution and covering factor of the clumps, but for simplicity and to facilitate a direct comparison to NM88 we follow their prescription and assume that the emerging spectrum is a superposition of the spectra of individual clumps. NM88 further show that dust does not affect the emerging spectra, provided that the gas-to-dust ratio $\xi_{\rm dust} \lesssim 0.016$, where $\xi_{\rm dust}$ denotes the dust-to-gas ratio relative to the local interstellar value. They argue that this limit is acceptable because the dust grains would be destroyed by the passage of the shock front. We finally assume all clumps - and therefore the radiative transfer processes inside them - to be identical, then we can simply adopt the spectrum from Figure 3.4. In reality, from our findings in § 3.3.2.2, we expect that a distribution of clump properties would give rise to a superposition of spectra with peaks in different locations, and therefore that the emission in the blue wing of the line profile would have less prominent features[4].

Putting these reservations aside, we observe that qualitatively our simulated spectrum shows similarities to that predicted by NM88. We should note though that in Figure 3.4 we omit the direct sources at $v_b = 0$, which NM88 put in by hand as a delta function at $v_b = 0$.

We note that the two largest peaks immediately redward and blueward of line centre cannot be attributed to photons exiting without crossing the shock front as in NM88 (i.e. the standard double-peak profile produced by a static, optically thick medium). Instead, our analysis shows clearly that these peaks are primarily caused by photons which have undergone 0 and $\{1, 2\}$ shock crossings respectively. The smaller, secondary blueward peak is primarily composed of photons which have undergone 3 or 4 shock crossings. The primary differences between the spectra predicted by us and by NM88 are in the asymmetry of the two main peaks, and the ratio of the peaks. As we have previously mentioned, the source distribution and physical parameters of the clumps affect the outcoming spectrum,

[4]Even for fixed clump properties, the emerging Lyα spectrum depends on the placement of Lyα sources. In particular, changing only the source position leaves the peak positions unchanged, but modulates their relative amplitudes. We have explicitly verified, however, that our spectrum barely changed if we assumed a uniform distribution of Lyα sources throughout the outflowing slab.

and it is likely that tuning of either or both of these will alleviate this discrepancy. The two blue peaks in our model appear to be shifted redward compared to the model predictions from NM88. However, a comparison of the observed spectrum in NM88 and our model shows that the positions of our blue peaks are also a good fit to the data. In fact, if the amplitude of the peaks is ignored (this is dependent on the exact source distribution), the position of the peaks in our model is arguably a better fit to the data than the NM88 model.

We now consider the observed Lyα spectra of $z \sim 3.4$ radio galaxy B2 0902+34 (Adams et al., 2009). There is broad agreement in the shape of the observed and predicted spectra, with the broad wing extending well beyond ~ 1000 km s^{-1}, and some spectra displaying prominent peaks blueshifted by ~ 1000 km s^{-1}. The spectra from B2 0902+34 vary widely depending on where exactly on the galaxy the fiber is placed. It is however true to say that for the majority of fibers the most prominent peak has been observed to lie either at the systemic velocity, which was determined from the observed 21-cm absorption signature, or slightly blueward of it. This would appear to be in tension with our model here which, following NM88, predicts the most prominent peak to appear at $v_b = 0$ (not shown in Figure 3.4) as a result of direct 'blister' sources on the edge of the clump. Recalling that the amplitude ratios of the peaks is determined by the source distribution, we speculate that this tension could be relieved if the most prominent peak in the B2 0902+34 observations is instead identified as the first blueshifted peak in our model, and the absence of a pronounced peak at systemic velocity attributed to a lower prevalence of 'blister' sources in this system.

3.4.2 Blue Bump Spectra

As we mentioned in § 3.1, observed Lyα spectral line profiles can often be reproduced surprisingly well with shell-models. In these models, a central Lyα source is surrounded by a geometrically thin shell of gas. The shell-models contain two parameters describing the Lyα source: (*i*) the assumed FWHM of the Lyα line prior to scattering, and (*ii*) the strength of the Lyα emission line, which is quantified by the equivalent width (EW). The shell itself is described by four additional parameters: (*i*) the HI column density of the shell $N_{\rm HI}^{\rm shell}$, (*ii*) its outflow velocity $v_{\rm shell}$, (*iii*) its dust content $\tau_{\rm d}$, and (*iv*) its velocity dispersion b.

Some recent analyses have pointed out that shell-models have difficulties explaining blue bumps in observed spectra (e.g. Kulas et al., 2012; Chonis et al., 2013; Hashimoto et al., 2015). Nice examples can be seen e.g. in Figure 7 of Chonis et al. (2013). More recently Hashimoto et al. (2015) pointed out that the shell-models, when applied to blue-bump objects, consistently require the intrinsic Lyα line to have a FWHM which is too large (an excessively high FWHM was only required for blue-bump objects).

Because Fermi acceleration naturally gives rise to blueshifted emission, we investigate a minor modification of the shell-model in which we embed the outflowing shell within a static gas cloud. This modification can be interpreted as the situation in which the outflow is still propagating into the static interstellar medium (possibly prior to breaking out of it).

This configuration now includes a shock front as in our previous analysis. We specifically study a shell-model with parameters based loosely on those inferred by Hashimoto et al. (2015): $N_{\rm HI}^{\rm shell} = 10^{19}$ cm^{-2}, $v_{\rm shell} = 200$ km s^{-1}, and FWHM=200 km s^{-1}. We further assume $\tau_{\rm d} = 0$, EW$= \infty$ (i.e. pure Lyα emission), and $b = 12.9$ km s^{-1} (corresponding to gas at 10^4K), but note that these assumptions do not affect our results at all[5]. Finally, the key new model ingredient is the addition of a static shell of gas adjacent to the outflowing shell. This shell is characterised primarily by its column density $N_{\rm HI}^{\rm stat}$. We assume that it has no dust and that it has the same temperature as the outflowing shell.

The results of this analysis for $N_{\rm HI}^{\rm stat} = 0$ (i.e. the original shell-model; solid black histogram), $N_{\rm HI}^{\rm stat} = 10^{18}$ cm^{-2} (red dashed line), and $N_{\rm HI}^{\rm stat} = 5 \times 10^{18}$ cm^{-2} (blue dotted line) are shown in Figure 3.5, where it is clear that a small additional column of hydrogen ($N_{\rm HI}^{\rm stat} \ll N_{\rm HI}^{\rm shell}$) dramatically affects the spectrum blueward of the systemic velocity. This large change of the spectrum can be easily understood, as the outflowing shell with $N_{\rm HI}^{\rm shell} = 10^{19}$ cm^{-2} directly transmits a significant fraction of Lyα photons. However, the surrounding static shell is optically thick to photons emitted near line centre, because these photons still appear close to the centre of the line in the frame of this gas. The static shell therefore effectively reflects back photons into the outflowing shell[6]. The reflected photons appear blueshifted by $\sim v_{\rm shell}$ in the frame of the outflowing shell, where their newly acquired large blueshift makes them escape efficiently. For example, photons that are scattered by 90° after being reflected back into the outflowing shell escape with a blueshift of $\sim v_b$, which is indeed where the new peaks in the spectrum emerge. This process is depicted schematically in Figure 3.6. The analytic estimate for the maximum blueshift in § 3.2 provides a decent estimate of the maximum blueshift $v_b \sim 160{\rm km\,s^{-1}} + v_s \sim 350$ km s^{-1} (where the 160 km s^{-1} was the maximum shift for a column density of 10^{19} cm^{-2} in the frame of the scattering medium).

3.5 Discussion and Conclusions

In this chapter we have presented a detailed analysis of Fermi acceleration of Lyα photons across a shock front, a process that was first studied analytically by NM88. Our Monte-Carlo radiative transfer simulations of the slab model (Figure 3.1) confirm the basic results in NM88.

In particular, we find that each time a photon crosses a shock front that is propagating at v_s, it experiences a Doppler boost $v_b \sim (0.5 - 0.7)v_s$ in the gas frame. The precise numerical coefficient depends on the opacity of the gas to the Lyα photons. We also

[5]For completeness we have presented our numerical results which include dust in Appendix A.2. These results show clearly that dust barely affects the blue bumps. This result is not surprising: our modification only adds a small amount of additional hydrogen, and this additional gas triggers Fermi-acceleration into the (blue) wings of the Lyα line profile, where Lyα photons escape more easily.

[6]In Figure A.4 we show that this mechanism produces blue bumps for column densities as low as $N_{\rm HI}^{\rm stat} \sim 10^{15}$ cm^{-2}. This is because a static shell of gas with $N_{\rm HI}^{\rm stat} \sim 10^{15}$ cm^{-2} remains optically thick to Lyα photons emitted close to line centre, and can therefore reflect back photons into the expanding shell.

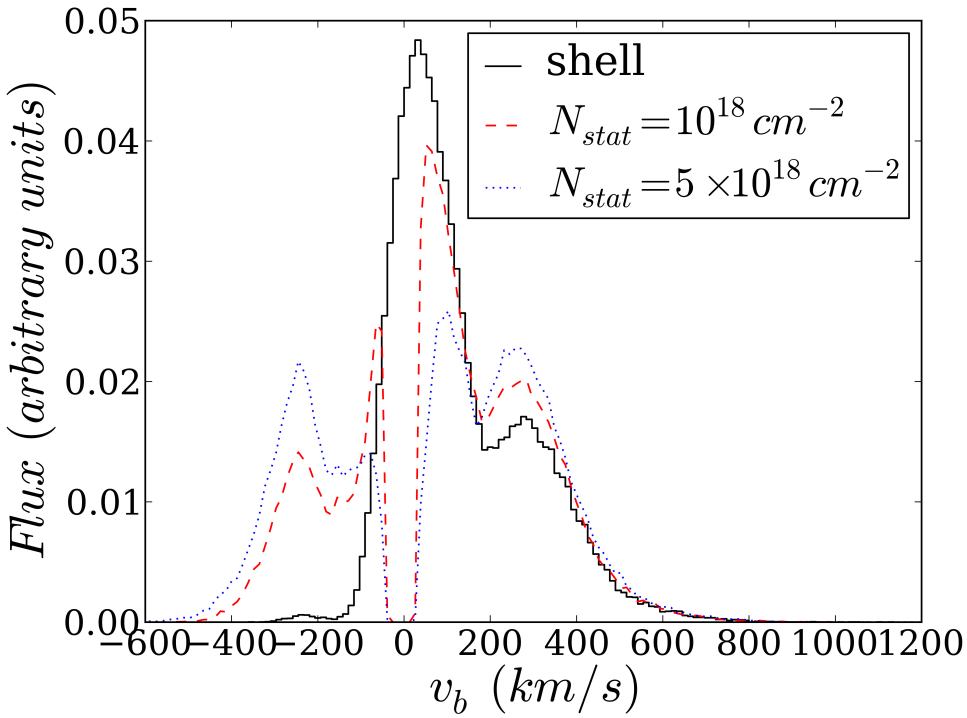

Figure 3.5: Predicted Lyα spectra emerging from a Lyα source surrounded by a shell of HI gas outflowing at 200 km s^{-1} (the shell-model, see text) embedded within a static neutral shell with a column density $N_{\rm HI}^{\rm stat}$. The solid black histogram shows the spectrum for $N_{\rm HI}^{\rm stat} = 0$ (i.e. the standard shell-model), while the red dashed and blue dotted line shows the spectrum for $N_{\rm HI}^{\rm stat} = 10^{18}$ cm^{-2} and $N_{\rm HI}^{\rm stat} = 5 \times 10^{18}$ cm^{-2}, respectively. The plot demonstrates that adding even a small amount of hydrogen ($1/10^{\rm th}$ of that in the shell) triggers the on-set of Fermi acceleration, which gives rise to a blue bump.

find the blueshift of the Lyα photons to increase in proportion with the number of shock crossings l, i.e. $v_b \propto (0.5 - 0.7) l v_s$.

We discussed how our results can help to explain extended blue wings observed in spectra of radio galaxy 3C326.1 (which was proposed previously by NM88), but we note that our line as a whole appears redshifted by $\sim$ 300 km s^{-1}. We reached the same conclusions for the radio galaxy B2 0902+34 (Adams et al., 2009). We nevertheless consider Fermi acceleration a plausible alternative to the model proposed by Adams et al. (2009), which involves the collapse of $> 10^{12} M_{\odot}$ of neutral gas.

Lastly, we have shown that Fermi acceleration naturally gives rise to blue bumps in Lyα spectra, which are difficult to reproduce with conventional shell-models (at least those with reasonable values for the FWHM of the intrinsic Lyα line, see Hashimoto et al. 2015). We presented a natural extension to shell-models in which the shell is expanding into static gas that contains a low column density of HI[7]. These models can give rise to blue bumps

[7]If we interpret the shell as dense gas that is swept up by feedback processes, then the static gas can

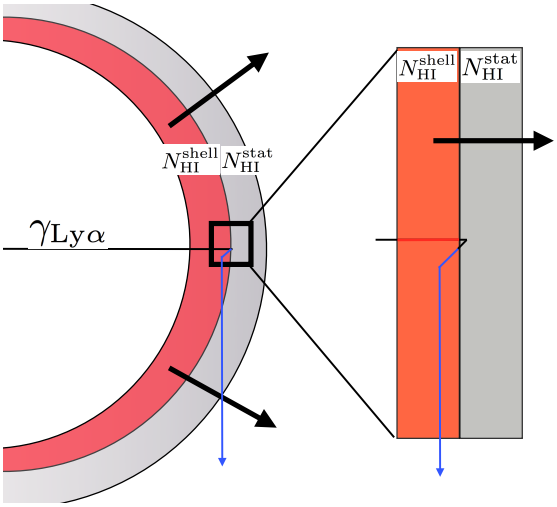

Figure 3.6: Schematic illustration of the origin of photons of the blue bump observed in Lyα spectra in the shell-model. The outflowing shell is surrounded by a static ISM.

in Lyα spectra without abandoning the simplicity of the conventional shell-model[8]. In our model, the bump consists of photons that initially streamed through the outflowing shell, but which were reflected back into this shell by the surrounding static interstellar medium. This suggests that blue bumps in Lyα spectra are associated with outflows that are still confined to the ISM of the galaxy, which may represent an earlier stage in the evolution of the galaxy.

be interpreted as gas that is being swept up.

[8]It is possible that this simply signals the break-down of the shell-model, and that radiative transfer through more complex gas geometries needs to be studied. There are numerous works studying Lyα transfer through more complex, simulated gas distributions (e.g. Tasitsiomi, 2006; Laursen et al., 2009; Barnes et al., 2011; Verhamme et al., 2012; Behrens & Braun, 2014; Lake et al., 2015; Yajima et al., 2015), though results from these calculations have not been compared to observed spectra in detail. Note however, that the sensitivity of the Lyα transfer scattering process to the distribution and kinematics of neutral gas implies that these predictions depend on sub-grid feedback prescriptions as well as the spatial resolution of the simulation (which ideally needs to be sub-pc, see Dijkstra & Kramer 2012 for a discussion). Simulations therefore have their own uncertainties, and thus provide a complementary route to addressing Lyα scattering on interstellar scales.

Chapter 4

Constraints on Galactic Outflow Models from the CGM in Lyα

This chapter is adapted from a paper to be submitted for publication by MNRAS. Authors: Andrew S. Chung, Mark Dijkstra, Benedetta Ciardi, Koki Kakiichi, Michael Aumer, and Thorsten Naab.

Galactic outflows are critical to our understanding of galaxy formation and evolution. However the details of this feedback process remain unclear. We compare Lyα observations of the circumgalactic medium (CGM) of Lyman Break Galaxies (LBGs) with mock observations of their simulated CGM. We use cosmological hydrodynamical 'zoom-in' simulations of a LBG which contains strong, momentum-driven galactic outflows. Simulation snapshots at $z = 2.2$ and $z = 2.65$ are used, corresponding to the available observational data. The simulation is post-processed with the radiative transfer code CRASH to account for the impact of ionising photons on hydrogen gas surrounding the simulated LBG. We generate mock absorption line maps for comparison with data derived from observed close galaxy-galaxy pairs. We perform calculations of Lyα photons scattering through the CGM with our newly developed Monte-Carlo code SLAF, and compare to observations of diffuse Lyα halos around LBGs. Our fiducial galactic outflow model reproduces many aspects of the currently observed characteristics of the CGM in Lyα . Our results suggest that galactic outflows affect Lyα absorption and emission around galaxies mostly at impact parameters $b < 50$ kpc, while cold accretion flows dominate at larger distances. We discuss the implications of this result, and underline the potential constraining power of CGM observations - in emission and absorption - on galactic outflow models.

4.1 Background

The processes which govern galaxy formation and evolution are one of the main open issues in modern cosmology. In the standard ΛCDM cosmological scenario strong feedback

is needed at both the low and high-mass ends of the galaxy mass function in order for models to match observations (White & Frenk, 1991). 'Feedback' typically refers to the complex processes through which star formation and accretion onto black holes deposit energy and momentum back into their surroundings. The details of feedback are not well understood (see Ciardi & Ferrara 2005 for a review on the topic). Because of its importance though, it is fundamental to study as many (potential) observational probes of this process as possible.

There is an increasing amount of data on the so-called 'circum-galactic' medium (CGM), which has been defined to be the region around galaxies out to a distance of $r \sim 300$ kpc and with a velocity offset from the galaxy's systemic redshift of up to $\Delta v \sim 300$ km s^{-1} (Rudie et al., 2012). Steidel et al. (2010) (hereafter S2010) note that the CGM provides a 'laboratory in which the effects of galaxy formation and AGN accretion (e.g. radiative and hydrodynamical feedback and its recent history) can be measured on scales that are not accessible using direct observations of galaxies'. In this chapter, we explore whether observations of the CGM in Lyα can constrain galactic outflow models, by comparing simulations of the CGM with observations.

Steidel et al. (2011) (hereafter S2011) found spatially extended Lyα emission from the CGM of Lyman Break Galaxies (LBGs) by stacking Lyα observations of 92 individual galaxies. Similar stacking analyses have revealed (fainter) Lyα halos around Lyα selected galaxies (i.e. Lyman Alpha Emitters, LAEs) (Matsuda et al., 2012; Momose et al., 2014)[1].

S2011 proposed that the diffuse halos arise from Lyα photons, produced in star-forming regions, scattering off outflowing material as they escape the galaxy. The presence of these outflows was inferred from the ubiquitous blueshifted low-ionisation absorption lines, and their interaction with Lyα photons was inferred from the redshifted Lyα emission lines (S2010). Here, star formation is both the source of the Lyα photons (produced by recombination in the HII regions around young stars) and of the stellar feedback which drives the outflowing material. S2011 provided a simple analytic model for the scattering of Lyα photons through the outflow, and showed that it is a good fit to their stacked observations. The radiative transfer of Lyα photons was treated though in an extremely approximate fashion.

Motivated by these models, Dijkstra & Kramer (2012) performed a systematic study of Lyα transfer through phenomenological models of spherically and cylindrically symmetric, large-scale, clumpy outflows. In these models, clumps were exclusively outflowing, with a one-to-one correspondence between outflow velocity (v) and distance from the galaxy (r). Following S2011, the velocity profile was inspired by 'momentum-driven' wind models in which the outflow accelerates as $a \propto r^{\alpha}$ ($\alpha \sim 1.5$). Dijkstra & Kramer (2012) constrained the HI properties of their clumps by matching the galaxy-galaxy pair absorption line presented by S2010. While these models can simultaneously explain the presence of Lyα halos and the amount of absorption, they also predict that a non-negligible fraction of Lyα pho-

[1]Feldmeier et al. (2013) did not find Lyα halos around $z \sim 2$ LAEs (also see Smith et al., 2012), and discuss that systematic uncertainties associated with stacking could reduce the statistical significance of previously reported detections. Momose et al. (2014) used larger samples of LAEs to confirm these systematic effects, but still obtained significant detections of Lyα halos.

tons did not scatter at all. These photons should be visible as a bright point source, which is absent in the data. As Dijkstra & Kramer (2012) pointed out, this problem could be due to the absence of low-column density HI systems, and/or a consequence of the simplified velocity profile of the outflow in their models.

In this work, we take a new complementary approach, and use cosmological hydrodynamical 'zoom' simulations from Genel et al. (2012) to generate a model CGM of an LBG. These simulations contain strong galactic outflows which are also momentum-driven, and provide us with a complex CGM that may more closely reflect reality than previous models. Importantly, the simulations contain inflowing 'cold streams' (e.g. Kereš et al., 2005; Dekel et al., 2009), which can contribute significantly to the amount of absorption measured in the CGM (van de Voort et al., 2012; Goerdt et al., 2012) and possibly to the emission (e.g. Dijkstra & Loeb, 2009; Rosdahl & Blaizot, 2012). While these simulations do not have the resolution to properly resolve the feedback processes and the kinematics of the cold gas, it is important to check how they compare to the available data on the CGM.

While there are a number of previous works which use simulations to study observational signatures of the CGM, our work distinguishes itself by simultaneously considering the CGM in Lyα emission and absorption. Previous works that used simulations have focussed either on emission (e.g. Laursen et al., 2009; Faucher-Giguère et al., 2010; Barnes et al., 2011; Rosdahl & Blaizot, 2012) or absorption (e.g. Goerdt et al., 2012; Shen et al., 2013). As was shown by Dijkstra & Kramer (2012), joint constraints from Lyα in absorption and emission are much more powerful than either data set individually. Finally, in contrast to previous studies that modelled the CGM in absorption, our simulations are post-processed with an ionising photon radiative transfer code (CRASH, Ciardi et al. 2001) and account for local sources of ionising radiation, which can be more important than the overall ionising background (especially at close distances to the galaxy, see e.g. Shen et al., 2013).

This chapter is laid out as follows: in § 4.2 we describe the simulations, § 4.3 presents the output of our pipeline and compares our results to observations, § 4.4 discusses the results and our conclusions[2].

4.2 Simulations

4.2.1 Cosmological Hydrodynamical Simulations

It is not currently computationally feasible to perform full hydrodynamic cosmological simulations with sufficient resolution to resolve the detailed gas dynamics of the CGM. Therefore we use the cosmological 'zoom-in' simulations from Genel et al. (2012), which start with an N-body dark matter only simulation. A region of space is cut out around a massive dark matter halo and re-simulated, adding baryons and hydrodynamic physics

[2]Throughout this work, we used the following cosmological parameters: $\Omega_\Lambda = 0.74$, $\Omega_m = 0.26$, $\Omega_b = 0.044$, $h = 0.72$, $n = 0.95$, and $\sigma_8 = 0.77$ (REF).

using a modified version of GADGET-2 (Springel 2005, Oppenheimer & Davé 2006, and Oppenheimer & Davé 2008).

Genel et al. (2012) use a modified Oppenheimer & Davé (2008) wind model, which implements momentum-driven winds powered by stellar feedback. In brief, star-forming particles become wind particles and are stochastically kicked perpendicular to the plane of the galaxy. The strength of the kick given to a particular wind particle is given by $v_{wind} = \sigma(4 + 4.29\sqrt{f_L - 1})$, where σ is the velocity dispersion of the galaxy, and f_L is the luminosity factor stochastically chosen in the range 1.05-2. With this model v_{wind} is higher than that used in Oppenheimer & Davé (2008), with typical wind velocities of $\approx$ 400-700 km s^{-1}. This ensures that the wind particles escape the disk. The mass-loading factor, which is the wind mass loss rate divided by the star formation rate, is typically $\eta \approx 4$.

It is worth noting that this wind model was used to generate a suite of simulations which reproduce the metallicity and ionisation of the intergalactic medium (IGM), the galaxy mass-metallicity relation, the high galactic gas fraction at high redshift ($z \gtrsim 2$), and the fact that galaxies contain a low fraction of cosmic baryons ($\sim 5 - 10\%$ at $z = 0$ Fukugita & Peebles 2004). In other words, despite its simplicity, this model simultaneously reproduces several observational constraints and scaling relations.

The hydrodynamic simulation, covering a region of $\approx$ 5 Mpc comoving, has a mass resolution of $8 \times 10^5 M_\odot$ for baryonic particles, and $5 \times 10^6 M_\odot$ for dark matter particles. The gravitational softening length of the baryonic particles is $200h^{-1}$ pc comoving. Since we ultimately want to compare to S2010 and S2011, we select galaxies with a similar mass (i.e. $M_* \approx 10^{10.5} M_\odot$) and use the GADGET snapshots at $z = 2.2$ (absorption) and $z = 2.65$ (emission), which are equivalent to the mean redshifts of the observations in S2010 and S2011, respectively. Further details of the specific galaxy under consideration in this chapter (identified as $s396$) can be found in Table 1 of Genel et al. (2012). In brief, at $z = 2.2$, $s396$ resides in a $1.5 \times 10^{12} M_\odot$ dark matter halo, has a stellar mass, $M_* = 2.5 \times 10^{10} M_\odot$, a star formation rate of $14 M_\odot$ yr^{-1}, and an intrisic Lyα luminosity of 3.024×10^{44} erg s^{-1} (assuming $f_{esc}^{ion} = 0.02$ and $f_{esc}^{Ly\alpha} = 1$). The star formation rate is comparable to the median star formation rate of the 'Lyα Em' sub-sample from S2011, which is $18.6 M_\odot$ yr^{-1}.

The simulations provide, as part of their output, the sites where star formation occurs during the simulation. This is the source of the stellar feedback which drives the galactic wind. In order to calculate the photon budget for the radiative transfer post-processing we model the star forming regions with the STARBURST99 (Leitherer et al. 1999) population synthesis code. Each star-forming particle is treated as a simple stellar population (SSP), where we assume a Chabrier (2003) Initial Mass Function (IMF) to be consistent with the assumptions made for the hydrodynamical simulations. We use the instantaneous star formation mode of STARBURST99, and integrate in time. Thus for the duration of the burst we ascertain the mean ionising photon count, Lyα luminosity, and the time-averaged ionising photon spectrum of the SSP.

Finally, for comparison to our fiducial model we also run hydrodynamical simulations with the same initial conditions but with less efficient feedback. For these runs we use the Springel & Hernquist (2002) feedback prescription which does not give an explicit kick to

wind particles nor decouple them from hydrodynamics. Throughout this chapter we refer to this model as the inefficient feedback model. In Chapter 5 this model is referred to as *n71*.

4.2.2 Radiative Transfer Simulations

A smaller box of side ~ 600 comoving kpc (corresponding to ~ 328 physical kpc at z=2.65 and ~ 187 physical kpc at z=2.20) was cut out of the SPH simulation output and gridded onto a discrete, uniform grid with dimension $N_c = 256$. This ensures a high-enough resolution to give converged results in the radiative transfer calculations. The output provides the temperature and density fields of the gas as well as the location and luminosity of the stellar population sources. Once augmented with information on the ionisation state of the gas, it can be used as initial conditions for performing the ionising radiation radiative transfer of the galaxy's local sources.

To obtain the initial conditions for the ionisation state we assume the presence of a Haardt & Madau (2012) uniform UV background (UVB) and the photoionisation equilibrium between the gas and the UVB. With the initial conditions defined by the above procedure the effect of the local sources given by the hydro simulations is added with CRASH (Ciardi et al., 2001; Maselli et al., 2003, 2009; Pierleoni et al., 2009; Graziani et al., 2013), a ray-tracing Monte Carlo 3D radiative transfer code which follows the propagation of the ionising continuum and its effect on the gas it crosses.

The output of CRASH includes the temperature and ionisation state of each cell in the simulation volume. We use 5×10^4 ionising photon packets per source and have checked that the results are converged such that the ionization state of the simulation box is not significantly affected by using more photon packets. We refer the interested reader to the original papers for more details on the code CRASH.

The recombination timescale of the ionised gas is long but nevertheless some recombination should occur. Because the UVB is not explicity included in the radiative transfer calculation, the gas at high galactocentric radius which has been highly ionised by the initial UVB, could artificially recombine during the CRASH run. To cope with this we estimate which cells have a UV flux dominated by the UVB and which by local sources. This is done by summing, for each cell, the ionising flux from all sources assuming a r^{-2} falloff and comparing this total to the UVB ionising flux. In all subsequent post-processing steps we use either the initial ionisation state calculated under UVB photoionisation equilibrium for the UVB dominated cells, or the CRASH ionisation state for cells dominated by local sources.

As we are neither able to resolve nor handle computationally the interstellar medium we simply remove it from the galaxy and approximate its effect by parameterizing the Lyα and ionising continuum escape fractions ($f_{esc}^{Ly\alpha}$ and f_{esc}^{ion} respectively). This choice is also motivated by the main goal of the chapter, i.e. an investigation of the impact of outflowing/inflowing material and not of the ISM. We remove the ISM based on two criteria: a density and a radial distance threshold. That is, grid cells are tagged as belonging to the ISM and subsequently removed if their density is above a density threshold $n_{th} = 0.5$ cm^{-3}

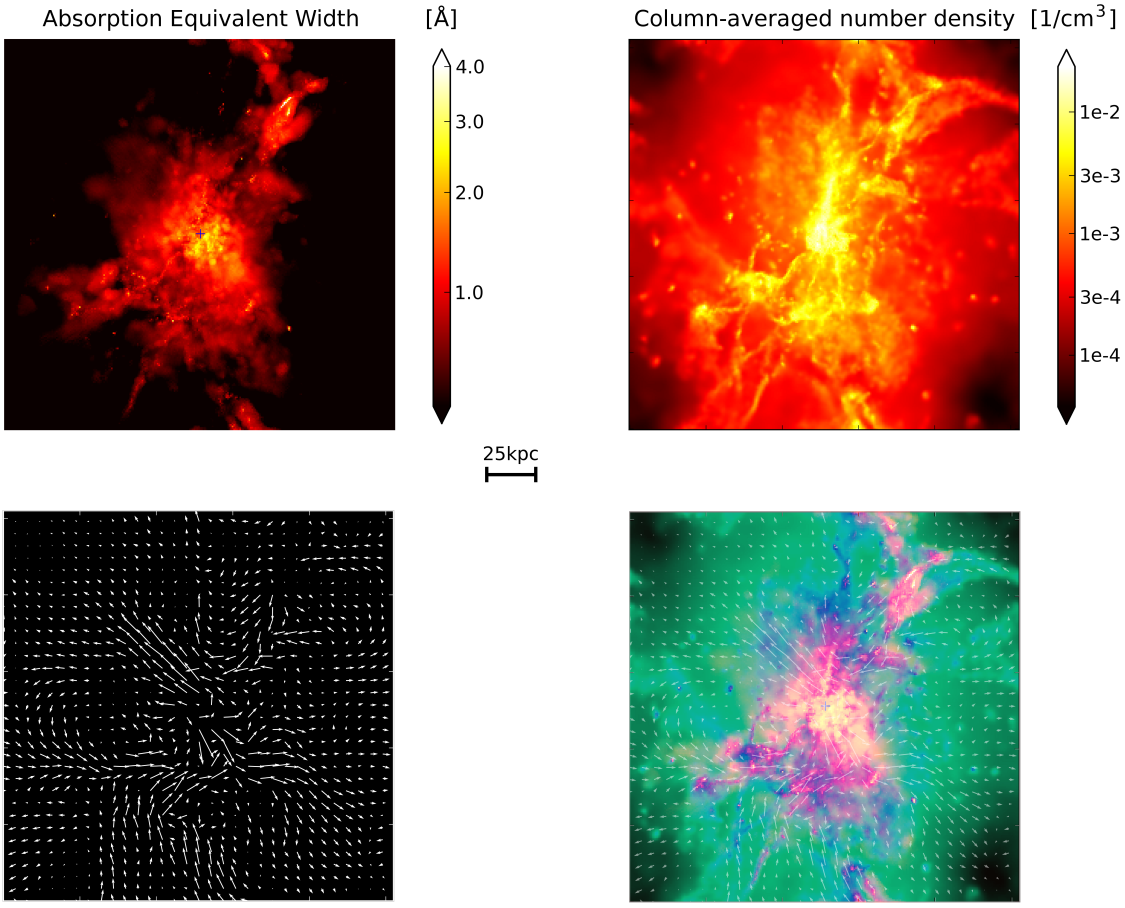

Figure 4.1: *Upper-left panel:*Synthetic equivalent width map of a ∼187 kpc (physical) simulation box at $z = 2.2$. The '+' symbol denotes the centre of the galaxy as defined by the gas centre-of-mass. *Upper-right panel:* The same simulation box but showing column-averaged HI number density. *Lower-left panel:* Depiction of the velocity field structure for the central slice. *Lower-right panel:* Composite image additively overlaying the equivalent width, density, and velocity fields of the previous 3 panels. For clarity the hue of the column-averaged density plot has been shifted to green-blue-cyan.

and they also lie within ∼10 kpc of the centre of mass of the galaxy. This galactocentric radius threshold just serves to ensure that only gas which is part of the galaxy itself is identified as ISM, and avoids removing high-density clumps in the CGM. In practice our results are insensitive to the exact radius threshold used. Likewise we tested density thresholds of $n_{th} = \{0.1, 1.0\}$ cm^{-3} and found very little variation in our results. Here we assume a value of $f_{esc}^{ion} = 0.02$ as in Gnedin et al. (2008). To perform the Lyα radiative transfer we use a total of $\sim 10^5$ photon packets and assume $f_{esc}^{Ly\alpha} = 1.0$, but as discussed later we renormalise the results to assume a different Lyα escape fraction.

The Lyα radiative transfer is performed with SLAF, a new code which we developed during the course of this work. SLAF is a Monte-Carlo Lyα radiative transfer code in the vein of many previous works (e.g. Zheng & Miralda-Escudé, 2002; Cantalupo et al., 2005; Dijkstra et al., 2006a; Tasitsiomi, 2006; Verhamme et al., 2006; Laursen et al., 2009; Yajima et al., 2012), and can be applied to arbitrary 3D gas distributions and velocity fields. All Lyα photons are injected into the CGM at line-centre. We have briefly investigated the effect of a gaussian injection line profile and find that for a gaussian with a standard deviation of 150 km s^{-1} the shape of the surface brightness profile is not altered considerably. However we leave the discussion of a detailed treatment of the injection profiles for a future paper. We also developed a code, LAF, to calculate absorption line profiles for sightlines through the simulation volume, using the ionisation state and temperature output of the CRASH radiative transfer. With these two codes we can generate emission maps (using SLAF), and absorption equivalent-width maps (using LAF). Technical details of SLAF and LAF are discussed in Chapter 2.

We include a very simple dust model in our Lyα radiative transfer calculations. Dust is treated as a grey absorber with a characteristic grain radius and dust-to-hydrogen ratio tuned to satisfy the redshift-dependent dust optical depth, τ_{dust}, relation from Garel et al. (2012). This gives the dust optical depth as a function of HI column density, $\tau_{dust}(\lambda) = (A_\lambda/A_V)_{Z_\odot} (Z/Z_\odot)^{1.35} (N_H/2.1 \times 10^{21}) (1+z)^{-\frac{1}{2}}$. As in Garel et al. (2012), $(A_\lambda/A_V)_{Z_\odot}$ is the solar metallicity extinction curve from Mathis et al. (1983), Z is the gas metallicity, and N_H is the HI column density. This prescription is applied to each cell in the simulation volume with the assumption of solar metallicity everywhere. Although this assumption is dubious, our testing has shown that our results are insensitive to this assumption.

4.3 Results

In this section we present the results we obtain by simulating the galaxy in the manner described above. We break down the results by the two quantities which can be directly compared to observations.

For reference we also present the results from the inefficient feedback model, though we leave detailed comparative discussion to § 4.4.

4.3.1 Absorption

The upper-left panel of Figure 4.1 shows a synthetic equivalent width map derived from the absorption line profiles output by LAF. The first step to create the map is to choose an orientation to view the simulation volume face-on. Then, from each cell on the far face of the box a ray is cast perpendicular to the face, towards the observer, for a total of N_c^2 rays. Each ray starts with a perfectly grey spectrum, which for simplicity we set to 1.0 such that $\{\forall\lambda : I_0(\lambda) = 1.0\}$. As it traverses the simulation volume the absorption line profile of the intervening HI is imprinted onto the spectrum. This is accomplished by Doppler shifting the existing spectrum into the gas frame of each cell in turn. For the n^{th} cell the optical depth, $\tau_n(\lambda)$, is calculated for each sampled wavelength of the spectrum, and the intensity exiting the cell derived as $I_{n+1}(\lambda) = I_n(\lambda)e^{-\tau_n(\lambda)}$, where $I_n(\lambda)$ is the intensity entering the cell. Thus when each ray reaches the observer its spectrum is known. From this the equivalent width of the total imprinted absorption line for each ray can be calculated as $EW = \int_{line}[1-\frac{I(\lambda)}{I_0(\lambda)}]d\lambda$, where $I(\lambda)$ is the final intensity reaching the observer and as noted earlier $I_0(\lambda) = 1.0$. This results in an $N_c \times N_c$ equivalent-width map.

Our simulation box is 187kpc (physical) on a side at redshift 2.2. At this scale the IGM intervening between the edge of the box and the observer could have an impact on the calculation of EW, and should in principle be taken into account. As per Laursen et al. (2011) this would primarily affect the spectrum blueward of line-centre. However, since the IGM at z=2.2 is mostly ionised, we do not expect significant interaction with the Lyα line. To explicitly test that this is the case we simulated a 2x larger box, centred on the same point, and found good agreement with the results presented here.

The equivalent width maps cannot be directly compared to observational data due to the fact that in reality not every line of sight has a sufficiently bright background galaxy or quasar along it. In other words, the relative rarity of galaxy-galaxy pairs and galaxy-quasar pairs prohibits the creation of such a map observationally. Nevertheless it is instructive to see what could be revealed in absorption given limitless background sources, especially in light of the other panels of Figure 4.1. The upper-right panel of Figure 4.1 shows the HI density averaged along the line of sight of each pixel in the map, the lower-left panel shows the projection of the velocity field of the central slice, and the lower-right panel shows a composite of all the other panels. We can see that generally speaking the morphology of the absorption equivalent width map follows that of the underlying HI distribution. It is clear that there is a correspondence of structures in the absorption map to the structures, such as the heavy central concentration of HI and the dense filaments, in the density map. Examination of the velocity field (lower panels) shows material inflowing along the filaments, and being blown out in a biconical outflow. Although the absorption map and velocity fields are somewhat messy, the lower-right panel shows how the velocity field (and hence outflow) affects the absorption map. Unsurprisingly, since the absorption depends on both the density and velocity field, larger velocities tend to coincide with larger absorption equivalent widths. Figure 4.1, in particular the lower-right panel, demonstrates the complex interplay between these 3 quantities.

Figure 4.2 is a plot of absorption equivalent width vs impact parameter (b), i.e. an

absorption equivalent width radial profile generated from the absorption map in Figure 4.1 (upper-left). It was created by sampling points in the map radially from the centre of the galaxy (denoted by a '+' in Figure 4.1). Each pixel from the absorption map in Figure 4.1 is plotted in Figure 4.2 as a black point. The solid red line shows the equivalent width profile obtained by azimuthal averaging of the map. The blue points and error bars are the data points from the galaxy-galaxy pairs in S2010. Figure 4.2 shows that for this viewing angle our simulation shows good agreement with the S2010 observations from $\sim$ 40kpc outwards, mildly underpredicts absorption compared to the data point at $b = 31$kpc and increasingly underpredicts the absorption towards the centre of the galaxy.

Figure 4.3 shows the absorption equivalent width profile as viewed from all three axis-aligned orientations as green, red, and blue solid lines. The red line in Figure 4.3 corresponds to the data in Figure 4.2. The mean of the three orientations is shown as a black solid line. We note that the mean equivalent width profile from our simulations is in quite good agreement with the observations. Similarly to the single orientation presented in Figure 4.2 the simulation is within the observational error bars for the outer 2 observational points, $\sim 12\%$ below at $b \sim 31$kpc and $\sim 43\%$ below at $b \sim 0$. The lower absorption equivalent width predicted by our simulations in the inner $\sim$ 40kpc of the galaxy is indicative of the simulated HI gas density being too low, the simulated velocity dispersion being too low, or a combination of both. We return to discuss this topic later in § 4.4, taking into account the emission results presented in § 4.3.2.

For comparison the inefficient feedback model is also shown in Figure 4.3 with dashed lines. Since this model lacks strong stellar feedback the galaxy has formed more stars and thus has a stellar ionising photon luminosity around 4x that of the galaxy formed with the fiducial model. In order to make a meaningful comparison then, we rescale the luminosity to match that of the galaxy formed under the strong feedback model.

Figure 4.3 shows that the inefficient feedback model also does a good job of reproducing the observations in the outer CGM. However it shows markedly worse agreement with observations in the inner $\sim$ 40kpc, consistently underpredicting the absorption equivalent width.

4.3.2 Emission

The observed Lyα luminosity of the simulated galaxy is 3.023×10^{44} erg s^{-1}, under the assumption that the Lyα escape fraction is 100%. This is extremely close to the intrinsic luminosity (see § 4.2). In the simulations of the Lyα emission of our galaxy, we ignore the component of the signal arising from recombination radiation in the CGM. We justify this by noting that a calculation of the recombination rate, and the resulting Lyα emission shows that the contribution from recombination within the simulated volume is less than 10% of the total Lyα emission.

Figure 4.4 shows the Lyα surface brightness profile of the galaxy as viewed from a single side. The data in red assume $f_{esc}^{Ly\alpha} = 0.144$, i.e. 14% of the Ly$\alpha$ photons generated in the galaxy diffuse through the ISM and escape into the CGM without attenuation. This is motivated by S2011, which quotes the Lyα escape fraction as 14.4% for the galaxy sub-

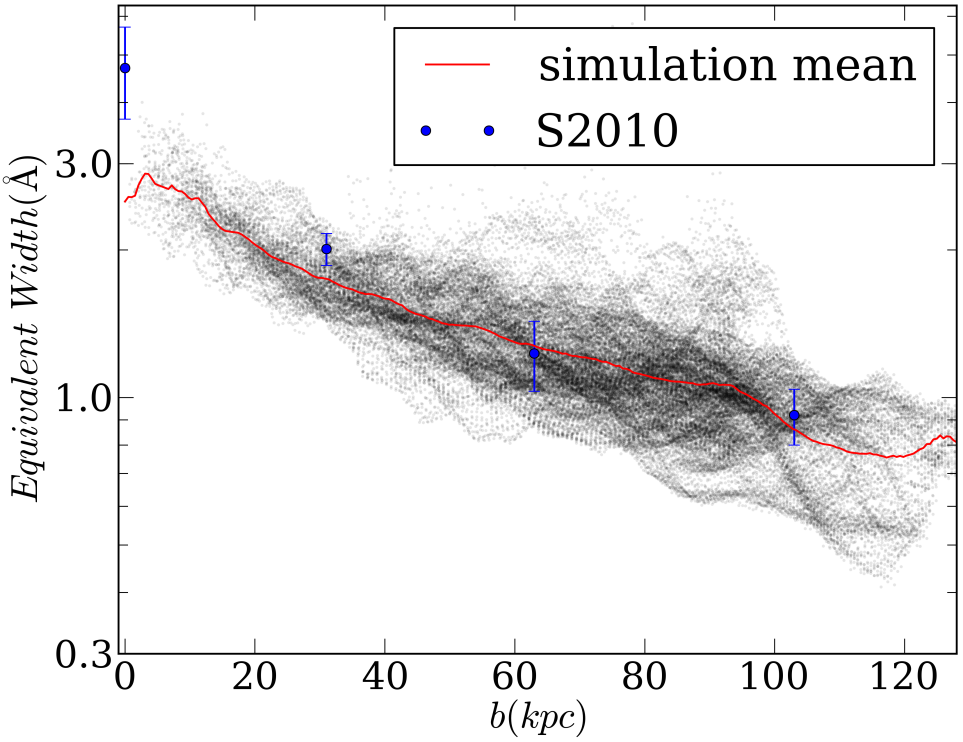

Figure 4.2: Equivalent width profile derived from the upper-left panel of Figure 4.1. Each pixel is plotted in the EW-b plane as a black point, with the mean profile shown by the solid red line. Observations from S2010 are shown in blue.

sample we compare to. The data in blue show the upper limit of $f_{esc}^{Ly\alpha} = 1.0$. Changing $f_{esc}^{Ly\alpha}$ simply modulates the total energy of the Lyα photons injected into the CGM. Since we use enough photon packets to sample the radiative transfer we similarly modulate the energy assigned to each photon packet to vary $f_{esc}^{Ly\alpha}$, which has the result of shifting the normalisation of the points in Figures 4.4. Here, solid lines represent the mean radial profile (which is the observational quantity plotted in S2011) of the simulated surface brightness images, dashed lines represent median radial profiles, and the shaded regions show the lower and upper quartiles. The simulated surface brightness images are degraded to 1" FWHM resolution prior to creating these profiles, in order to match the resolution of the S2011 data, represented by the green dashed line.

Interestingly, at times the mean profile rises above the median profile, and even above the upper quartile of the distribution at some impact parameters, which is indicative of the fact that at these impact parameters the distribution is not gaussian but skewed, dragging the mean up. This is caused by a small amount of substructure in the surface brightness images - small, bright star-forming clumps. The impact parameter at which these clumps reside depends on projection effects and thus on the observation angle. In this work we only simulate one galaxy, but given the essentially random projection of substructure we expect that a fairer comparison which stacked many simulated galaxies

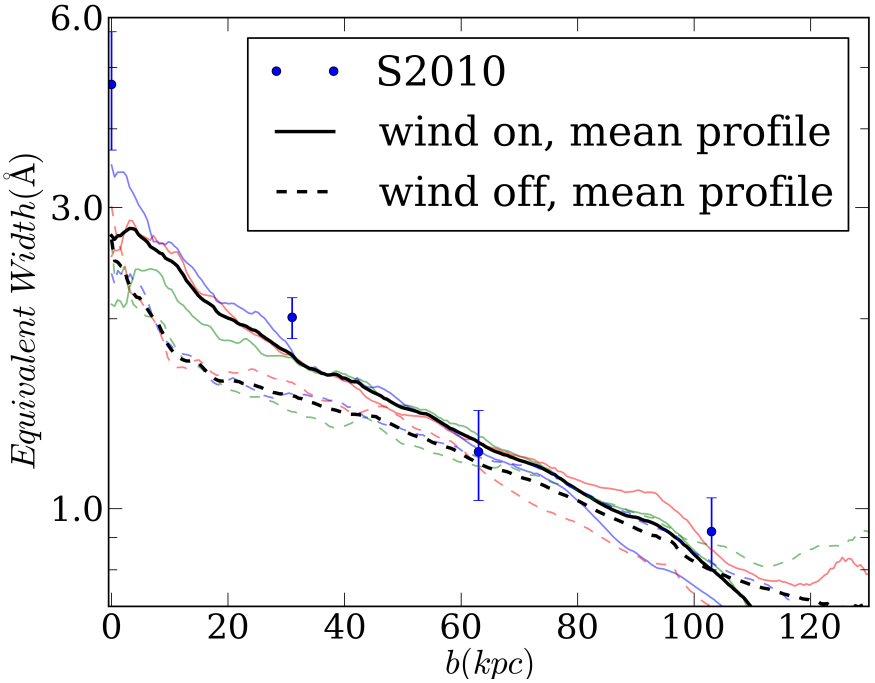

Figure 4.3: Comparison of equivalent width profiles from simulations with (solid lines) and without (dashed lines) strong stellar feedback. Green, red, and blue lines are profiles as viewed along the x, y, and z axes respectively. Black lines are averaged over the three primary axes. Observations from S2010 are shown as blue points and error bars.

would have the mean close to the median. This is because stacking more galaxies would reduce the weight given to outlying substructure: for each profile with substructure at a given impact parameter, many more profiles would not have substructure at that same impact parameter. This would significantly diminish the effect of these small clumps on the mean. We argue then that in the case of Figure 4.4 where there is little substructure it is perhaps better to compare the median profile to the results from S2011, which are themselves a stack of 52 galaxies. Considering the median profile in Figure 4.4 it is clear that our simulation is a good fit to observations from $\sim$ 40kpc out to where the observations end at 80kpc. Below $\sim$ 40kpc the surface brightness profile given by our simulation starts to rise above the observed profile - our simulated surface brightness profile is too peaked. The model proposed in Dijkstra & Kramer (2012) also exhibits a similar rise away from the observations in this inner region.

Figure 4.5 shows the average surface brightness profile obtained by stacking the six surface brightness maps corresponding to viewing the simulated galaxy from the six sides of the simulation volume. When stacking several viewing angles the shape of the profile remains very similar as there is suprisingly little variation in the profile when the galaxy is viewed from different orientations. The main difference between the profiles seen from

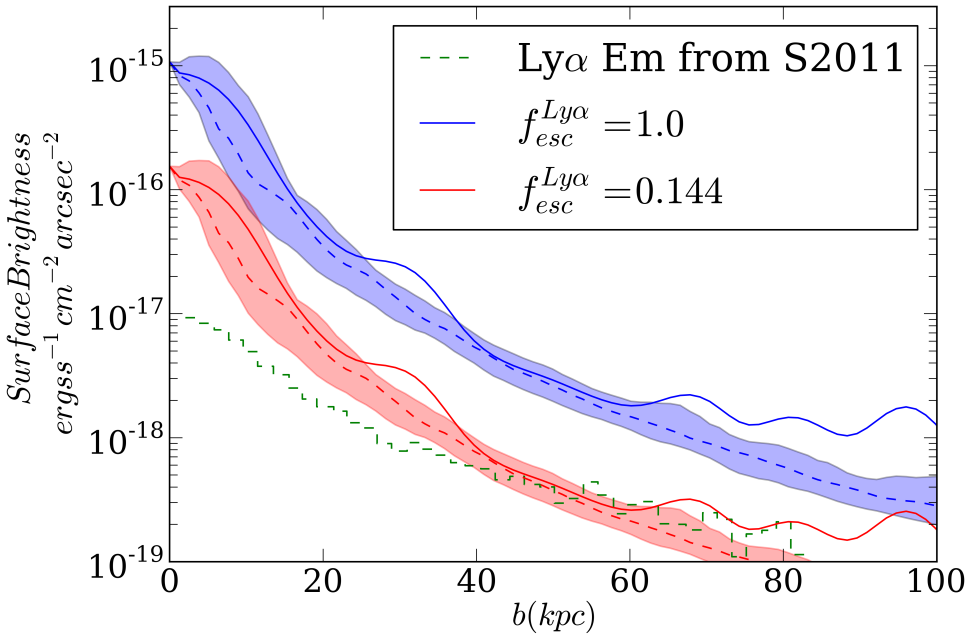

Figure 4.4: Lyα surface brightness as a function of impact parameter for our simulated galaxy at $z = 2.65$. The blue lines are for $f_{esc}^{Ly\alpha} = 1.0$, while the red lines assume $f_{esc}^{Ly\alpha} = 0.144$. In both cases the solid line is the azimuthal mean, the dashed line is the azimuthal median, and the shaded regions show the interquartile range. The dashed green line shows the 'Lyα Em' sub-sample from S2011.

different orientations is in the location and size of the bumps and peaks. The large bump at 80kpc comes from a bright peak visible along a single axis.

Figure 4.5 also shows the inefficient feedback model in red. As was done for the absorption results, the stellar luminosity of the inefficient feedback model has been scaled down to match that of the fiducial model. The medians of the two models are almost indistinguishable, thus the median surface brightness profile from the inefficient feedback model is also compatible with the data from $\sim$ 35kpc to 80kpc. However, when we compare the mean and median profiles of the inefficient feedback model we see a much larger discrepancy than between the mean and median profiles of the fiducial model. Whereas the fiducial model has good agreement between the mean and median for the vast majority of the profile, the mean of the inefficient feedback model is significantly above the median from $\sim$ 35kpc outwards. As discussed earlier the discrepancy between mean and median arises when there is substructure present. In this case there is much more substructure in the inefficient feedback model than the fiducial model. From this single realisation of a galaxy it is not possible to tell if the inefficient feedback model generically predicts an increased presence of substructure, but if this is a generic property of this model then we argue that here the median is not a good indicator of what we can expect from mean-stacking many simulated galaxies. This apparent contradiction to what has been advocated earlier is due

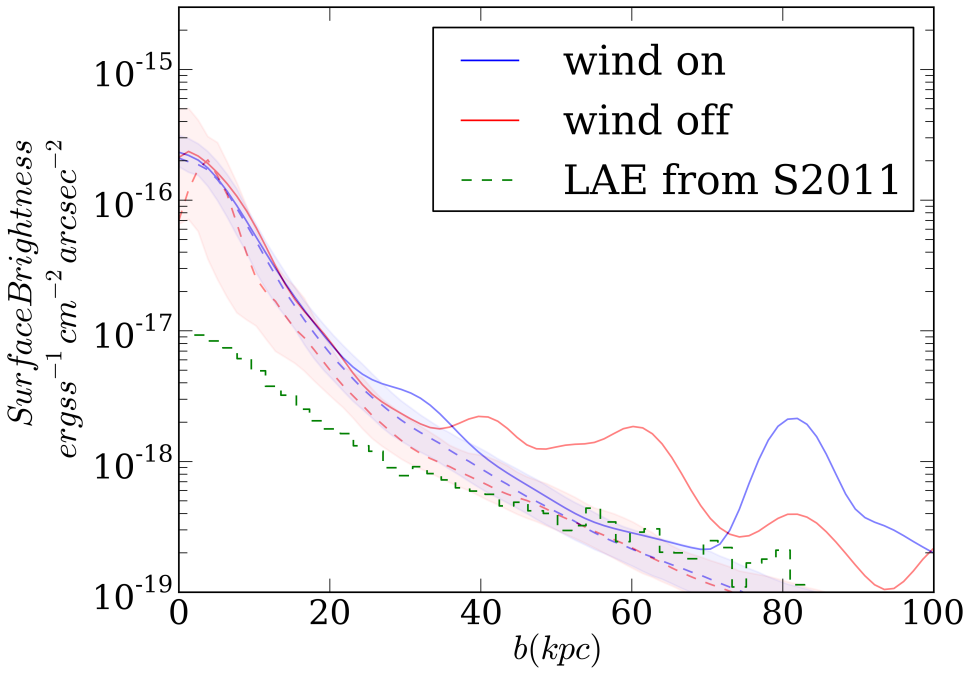

Figure 4.5: Six orientations of the simulated galaxy stacked to give the average surface brightness as a function of impact parameter. The blue points show the results for our wind model whereas the red points show the results of the galaxy simulated without strong feedback. In both cases $f_{esc}^{Ly\alpha} = 0.14$ as per S2011. In both cases the solid line is the azimuthal mean, the dashed line is the azimuthal median, and the shaded regions show the interquartile range. The dashed green line shows the 'Lyα Em' sub-sample from S2011.

to the fact that in fiducial model case there are only a few, isolated clumps, whereas here there is substructure over a large range of impact parameters. If every stacked galaxy exhibits this behaviour then the mean-stacked profile would not regress towards the median shown here but instead stay elevated.

Figure 4.6 shows the spectrum of all Lyα photons exiting the radiative transfer simulation box. There is a clear distinction between the two models. The fiducial model has a more pronounced redward peak, and a reduced blueward peak relative to the inefficient feedback model. This can be understood as a less extreme example of the effect shown in Figure 2.2 whereby an expanding medium boosts the redward peak and suppresses the blueward peak. Since the fiducial model has stronger feedback, and thus a stronger outflow than the inefficient feedback model, this is exactly the behaviour we expect to see in the spectrum.

We note also that the red peak has a larger velocity shift in the fiducial model than in the inefficient feedback model. Again, this is entirely consistent with the inefficient feedback model producing a weaker wind with a lower outflow velocity. As we discuss in § 4.4 the velocity offset of the peaks in both models are too low compared to observations, but aside from that the line shape of the fiducial model is in much better agreement with

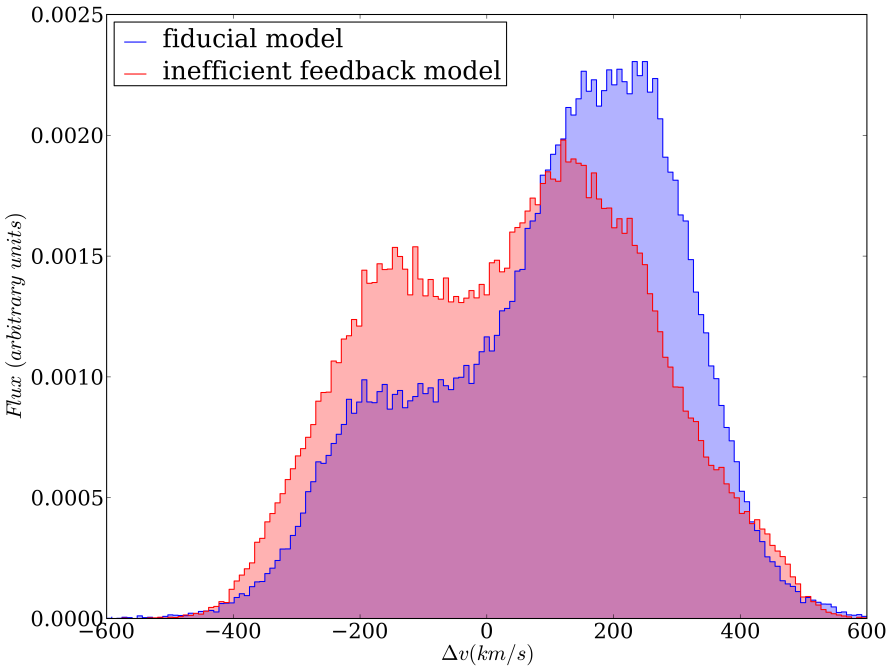

Figure 4.6: The spectrum of all Lyα photons exiting the simulation box for both the fiducial model and inefficient feedback models. The spectrum is in arbitrary units, and each has been normalised to unity.

the observational data.

4.3.3 Combined Absorption and Emission

Comparing our fiducial model with the inefficient feedback model we can state that judged by absorption alone the fiducial model better fits observations, while if we consider emission only the message is less clear. If we assume that for both models the median surface brightness profile is a good proxy for what would result from mean-stacking many more simulated galaxies, then neither model is a better fit to observations than the other. However, this assumption may not be true, and instead the mean profile of the inefficient feedback model could be a better proxy for the mean-stacked profile. If this is the case then the fiducial model again provides a better fit to observations.

Taking both absorption and emission into account it is clear that the fiducial model is favoured, with the caveat that this study is limited to a single simulated galaxy. At this stage, this preference comes primarily from the absorption result. As explained above, whether the emission actively supports this preference or merely does not contradict it hinges on the results of simulating more galaxies.

4.4 Discussion

We have demonstrated that our fiducial wind model can to a large extent reproduce the observables detailed in S2010 and S2011. In § 4.3.1 we noted that the simulated absorption equivalent widths within the inner $\sim$ 40kpc of the galaxy are lower than observations, suggesting that in this region either the simulated HI gas density is too low, the simulated velocity dispersion is too low, or a combination of both. In § 4.3.2 we saw that the simulated emission profile is too peaked compared to observations, i.e. the Lyα photons are arriving at the observer too directly (to see this intuitively consider the case of a point source in a vacuum - with no scattering the observer would see a 1-pixel wide step function emission profile). The emission surface brightness profile results confirm one of the options from the absorption result: a higher gas density would result in more Lyα photon scatterings, flattening the emission profile as required to match observations; a higher velocity dispersion would doppler shift the HI away from Lyα resonance, allowing the Lyα photons a more direct path to the observer and further increasing the peakiness of the emission profile. Thus the emission surface brightness profile results suggest that more successful wind models should have a higher central HI gas density. However, this alone cannot be the whole explanation. In both cases the emission spectra (Figure 4.6) show a red peak at a velocity offset less than that observed ($\sim$400km/s, see Steidel et al. 2010). This is indicative of the outflow velocity not being high enough. It appears then that a combination of both increased HI density in the inner regions and increased outflow velocities are needed to bring simulations into line with all of the available observations.

At first glance, the success of the outflow model is surprising since it is by necessity extremely simplified. In particular, it has the non-physical characteristic of temporary decoupling of the kicked wind particles, which is necessary to produce efficient feedback. In fact, Dalla Vecchia & Schaye (2008) suggest that the density threshold at which hydrodynamical coupling of the kicked particles to their neighbours is reinstated may be the dominant factor in determining the morphology of this class of outflow models.

While we need further simulations to confirm this, we speculate that since both of the tested models show similar absorption and emission profiles at $b > 50$kpc, the CGM signal in Lyα absorption and emission may be largely unaffected by outflows at these impact parameters. Indeed, Shen et al. (2013) show that outflows do not disrupt cold inflows. Thus if the signal at $b > 50$kpc is dominated by cold accretion flows it is reasonable to expect that our strong feedback model would not significantly affect the absorption and emission profiles in this region. However, we caution against interpretating the fact that the $b > 50$kpc signal does not appear to be coming from outflows to mean that it must necessarily be caused by inflows.

This limited radial influence of the outflow is perhaps a manifestation of the fact that a wind has some finite sphere of influence. Visually, we can see this by referring to the lower-left panel of Figure 4.1, where we can see that beyond $\sim$ 50kpc the amplitude of the velocity field drops off. The extent of this sphere will surely depend on the velocity of the stochastic kick given to wind particles but may also be affected by the aforementioned density threshold at which wind particles become re-coupled to hydrodynamics.

So far we only present the results of one simulated galaxy, and it would be premature to judge the fiducial outflow model on this one result. Since S2010 and S2011 both deal with averaged/stacked data the correct comparison to make is to similarly stacked simulated galaxies. Genel et al. (2012) simulate a suite of galaxies so in future this is feasible.

In fact the need for more simulations is deeper than may be immediately obvious. The two different feedback models we have presented produce galaxies with different properties from the same halo initial conditions. This, of course, is exactly why we are interested in various feedback models but as we saw in § 4.3 it makes comparison of the resulting galaxies difficult. Recall that the halo we chose to 're-simulate' with GADGET was chosen such that the galaxy which forms under the fiducial feedback model was similar to the mean of the S2010/S2011 samples. Using the inefficient feedback model the galaxy which formed was brighter and so we needed to rescale the luminosity of the galaxy to compare to the fiducial model and observations. This introduces an inconsistency: the stars which ionize/illuminate the CGM in Lyα do not have the same properties as those which generated the outflow velocity field and CGM properties. If we instead chose to compare two galaxies with the two models, based on the final formed galaxy properties (that is, choose galaxies from both models to match observations) we are faced with the problem that we would be comparing galaxies with different initial conditions. Since what we are really interested in is the impact of the different models on the evolution of the CGM we prefer the approach we have taken where we keep the initial conditions the same across the two models. Nevertheless, we acknowledge that this is not ideal and suggest that the real solution is to simulate a large statistical sample of galaxies for both models, from which a sample with properties matched to observations can be selected and compared.

Other outflow models such as those described in Dalla Vecchia & Schaye (2012) and Dalla Vecchia & Schaye (2008), avoid the unphysical temporary decoupling from hydrodynamics common to Springel & Hernquist (2003)-style winds. The resulting outflows are qualitatively different in spatial and velocity morphology from those produced by Springel & Hernquist (2003) and would thus make for an interesting comparison. Additionally, there are free parameters in the outflow model used by Genel et al. (2012), which should be explored in future.

Our simulations do not include spatially extended emission from the streams, either powered by recombination or collisional excitation. Previous works have shown that this may provide a large Lyα luminosity (e.g. Dijkstra & Loeb, 2009). Accounting for this emission may boost the amount of Lyα emission at large impact parameters, thus flattening the predicted surface brightness profile.

We note that the difference in the observed spectra from the two models we have presented may be observable with instruments such as MUSE. In particular a stronger outflow appears to diminish the 'twin peaks' profile of the Lyα line, emphasising the red peak and diminishing the blue peak.

This work has possible relevance to the enigmatic Lyα blobs reported by Steidel et al. (2000) and Matsuda et al. (2004). It is presently unclear what the physical mechanism powering these objects is. One possibility is that the blobs are powered by a central galaxy with the photons scattered off the surrounding medium, in a similar fashion to that

considered in this chapter. Hayes et al. (2011) presents observations supporting this model of Lyα blobs.

4.4.1 Comparison with Previous Work

As we have already pointed out, the absorption line data is better fit by our fiducial model but the inefficient feedback model also does a good job of reproducing the observations at higher impact parameters. The similarity of the profiles in Figure 4.3 goes some way towards explaining why previous analyses of Goerdt et al. (2012) (simulating a 138 physical kpc box at redshift 2.3) and Shen et al. (2013) (simulating a 263 physical kpc box at redshift 2.8) can reproduce the absorption line data without strong outflows, instead attributing most of the Lyα absorption to inflowing cold streams.

Goerdt et al. (2012) reproduces the Lyα absorption line data extremely well, even towards the centre of the galaxy. This is in contrast to our inefficient feedback model, which underpredicts the amount of Lyα absorption in the inner $\sim$ 40kpc. One reason for this may be that Goerdt et al. (2012) use a simple self-shielding criteria to calculate the ionisation state of the gas, and furthermore include no local sources. It is therefore reasonable that our simulations which contain local sources, mostly concentrated in the galaxy, and explicitly compute the ionisation state, should have a higher ionisation fraction in the central region close to the galaxy. This would naturally lead to a lower Lyα absorption equivalent width. In the outer regions, where local sources are less dominant, we expect better agreement with the Goerdt et al. (2012) result and indeed this is exactly what we see.

An alternate explanation for the discrepancy could be the way we calculate the initial conditions for our CRASH runs. Whereas we assume all gas is in photoionisation equilibrium with the UVB, Goerdt et al. (2012) use a self-shielding criteria. In the dense regions towards the centre of the galaxy this may cause us to overestimate the ionisation fraction of the gas. If this is happening then it is also happening in our fiducial model, which also lies under the observations at very small impact parameter, and a better treatment of the UVB in our calculations may yield better agreement with observations for our strong feedback model.

Shen et al. (2013) does include local sources, albeit in a simplified fashion, placing all sources at the centre of the galaxy. They have similar problems to our model in reproducing the central two absorption data points from S2010.

With respect to emission, S2011 presents a simple analytic model to explain their observations. They consider a spherically symmetric outflowing HI CGM, modulated by a covering fraction which is a power law function of galactocentric radius. Radiative transfer is treated with an extremely basic prescription. Nevertheless this model provides a good fit to their observations. The success of this model provided motivation for us to test whether the underlying assumption of a central source scattering of an outflowing CGM could stand up to the scrutiny applied by our hydrodynamic simulations and a full treatment of radiative transfer.

Laursen et al. (2009) perform a similar radiative transfer treatment to this work. However, their work differs to that presented here in some important ways. The galaxies they

simulate are not targeted to match the sample of galaxies from S2011, as they simulate galaxies at a different redshift (z=3.6) and with no strong feedback. Nevertheless, their results are comparable with a very similar surface brightness profile shape.

4.5 Summary

We have introduced a new test for galactic outflow models, which self-consistently combines hydrodynamical simulations and Lyα radiative transfer. Crucially we use constraints from both Lyα absorption and emission to test our models. The fiducial outflow model which we have presented in this chapter broadly reproduces both absorption and emission observations, although the Lyα emission spectrum remains problematic.

We also showed that there are differences in these two diagnostics when a different feedback model is used. Furthermore comparison of our results for the two feedback models hint that galactic outflows may predominantly affect the inner $\sim$50kpc of the CGM. This suggests that future Lyα observations of the inner CGM may be key to gaining a better understanding of the galactic outflows which appear to be important to galaxy formation and evolution.

In the next chapter we follow up on these findings by testing other outflow models with the same techniques.

Chapter 5

Model comparisons

In the previous chapter we discussed in depth the procedure we used to simulate the Lyα absorption and emission observables for a specific model, and compared our fiducial model (Genel et al., 2012) to one with inefficient stellar feedback (Springel & Hernquist, 2003).

In this chapter we apply the same procedure to a wider selection of wind models, and discuss the Lyα absorption and emission results that arise from them.

5.1 Outflow models

In this chapter we present the results of re-simulating the same halo/initial conditions (specifically halo *s396* from Genel et al. 2012 as per the preceding chapter) with different outflow models. Here we briefly summarise the properties of the tested outflow models:

- **fiducial**

  This is the fiducial feedback model (Genel et al., 2012) from Chapter 4. This simulation uses the star-formation model of Schaye & Dalla Vecchia (2008) together with the Oppenheimer & Davé (2008) wind prescription, modified to give higher wind velocities (i.e. a stronger/more efficient galactic wind). Note that the Oppenheimer & Davé (2008) wind prescription which this model extends is itself built on the feedback model from Springel & Hernquist (2003) (model *n71* immediately below). As in Springel & Hernquist (2003) wind particles are explicitly given a kick out of the galaxy.

- **n71**

  Uses the feedback model from Springel & Hernquist (2003). This is the 'inefficient feedback model' from the previous chapter. Again we briefly describe the model here but refer the interested reader to the original paper for comprehensive details. This model has a subresolution multi-phase treatment of the ISM where the hot phase is heated by supernovae, and the cold phase is in turn heated by radiative cooling of the hot phase gas. Star formation is assumed to proceed with some given efficiency

and timescale, converting cold gas into stars. Particles are explicitly designated as wind particles and given a velocity kick out of the galaxy.

- **n00**

  Uses the feedback model from Scannapieco et al. (2008). For a detailed discussion of the model see the aforementioned paper, in addition to Scannapieco et al. (2005) and Scannapieco et al. (2006). In brief, this model includes a decoupled multi-phase ISM model so that dissimilar particles cannot be SPH neighbours. Furthermore when a supernova occurs the energy is distributed between hot and cold ISM phases (the fraction going to one is parameterised, the other follows naturally). The hot phase particles receive energy immediately from supernovae. The cold phase particles keep track of the energy they accumulate from supernovae until they have the same entropy as hot neighbour particles, at which point the accumulated energy is injected into the cold phase. Gas particles are assumed to have star formation with a specified efficiency if they exceed a threshold density and are in a converging flow. In this model the wind arises naturally from supernova feedback, and the treatment of the multi-phase model. Once cold phase particles accumulate enough energy, they become hot phase particles, and the resulting thermal expansion can break out of the galaxy thus forming a wind.

- **n88**

  Uses the feedback model described in Aumer et al. (2013). This is a significant update to the (Scannapieco et al., 2005, 2006, 2008) model described above (*n00*). Here, supernova feedback injects energy into the ISM in both thermal and kinetic forms. Improvements have been made to the treatment of metal production and cooling rates. Turbulent diffusion of metals, and radiation pressure from young stars have also been added.

- **n89**

  Same feedback model as n88, but with a newer version of GADGET.

- **n85**

  Uses the feedback model described in Aumer et al. (2013) but without radiation pressure from young stars.

5.2 Results

Figure 5.1 shows the absorption equivalent width profiles from the above models. The y-axis is the absorption equivalent width in ångströms and the x-axis is the galacto-centric impact parameter at which the absorption profile is averaged. These are all presented in the same format as in the previous chapter (cf. Figure 4.3 in § 4.3.1).

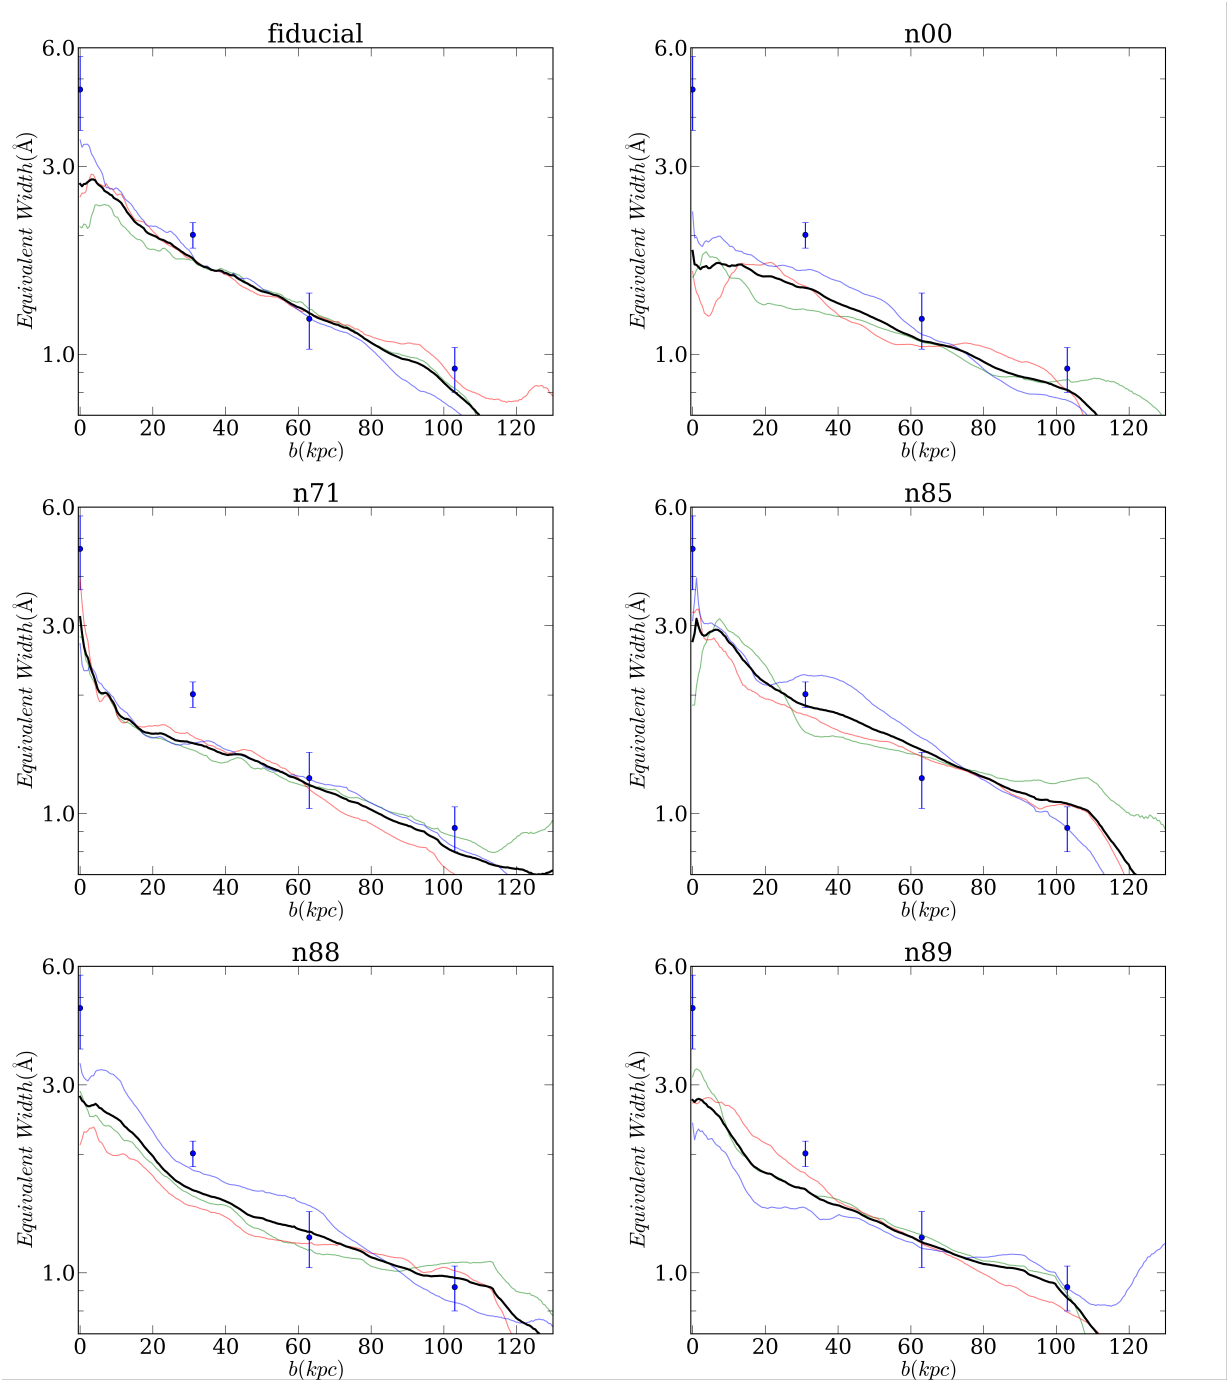

Figure 5.1: Lyα absorption equivalent width profiles for the feedback models discussed in § 5.1. In each panel the red, green, and blue lines show the mean radial profile obtained along the three primary axes. The black line is the radial profile averaged over the three axes. Blue points and error bars are observations from Steidel et al. (2010).

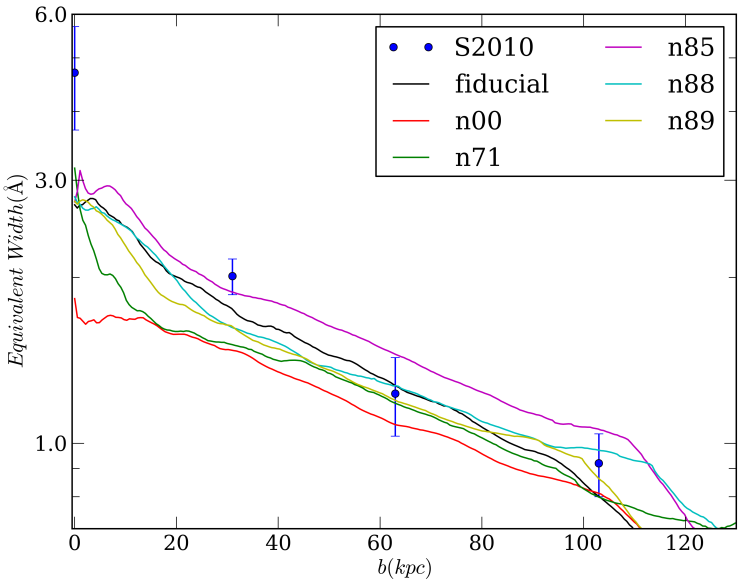

Figure 5.2: Comparison of the mean Lyα absorption equivalent width profiles of the above models. Blue points and error bars are observations from Steidel et al. (2010).

Figure 5.2 combines the mean absorption equivalent width profiles from the above models into one figure to facilitate a comparison.

Comparing the absorption equivalent width mean profiles for the various models (Figure 5.2) we see that all models result in absorption profiles that have a broadly similar shape. All of the absorption profiles are more-or-less consistent with the outer two observed data points at 63kpc and 103kpc from Steidel et al. (2010). All but *n85* are inconsistent with the 2nd data point, and all fail to rise sharply enough in the inner ∼30 kpc to match the observed data point.

Judged purely from the absorption results, model *n85* is the preferred model. Meanwhile the fiducial, *n88*, and *n89* models give very similar absorption profiles to each other, and come second to the *n85* model when judged by the observational data. Models *n00* and *n71* are the least favoured, lying farthest from the observed data point at an impact parameter of 31kpc. Model *n00* in particular exhibits a pronounced flattening of the profile towards the centre of the galaxy, in contrast to the other models and the observational data.

Nevertheless the differences between all of the models are surprisingly small, and this again re-inforces the notion that we need to use all of the available data (i.e. not just absorption profiles but also emission surface brightness profiles and spectra) when trying to distinguish between models.

Therefore we now examine the Lyα surface brightness emission profiles. These are shown in Figure 5.3. Again, they are presented in the same format as the previous chapter (cf. Figure 4.5 in § 4.3.2). Figures 5.4 and 5.5 combine the median and mean surface brightness profiles respectively in order to facilitate comparison of the models.

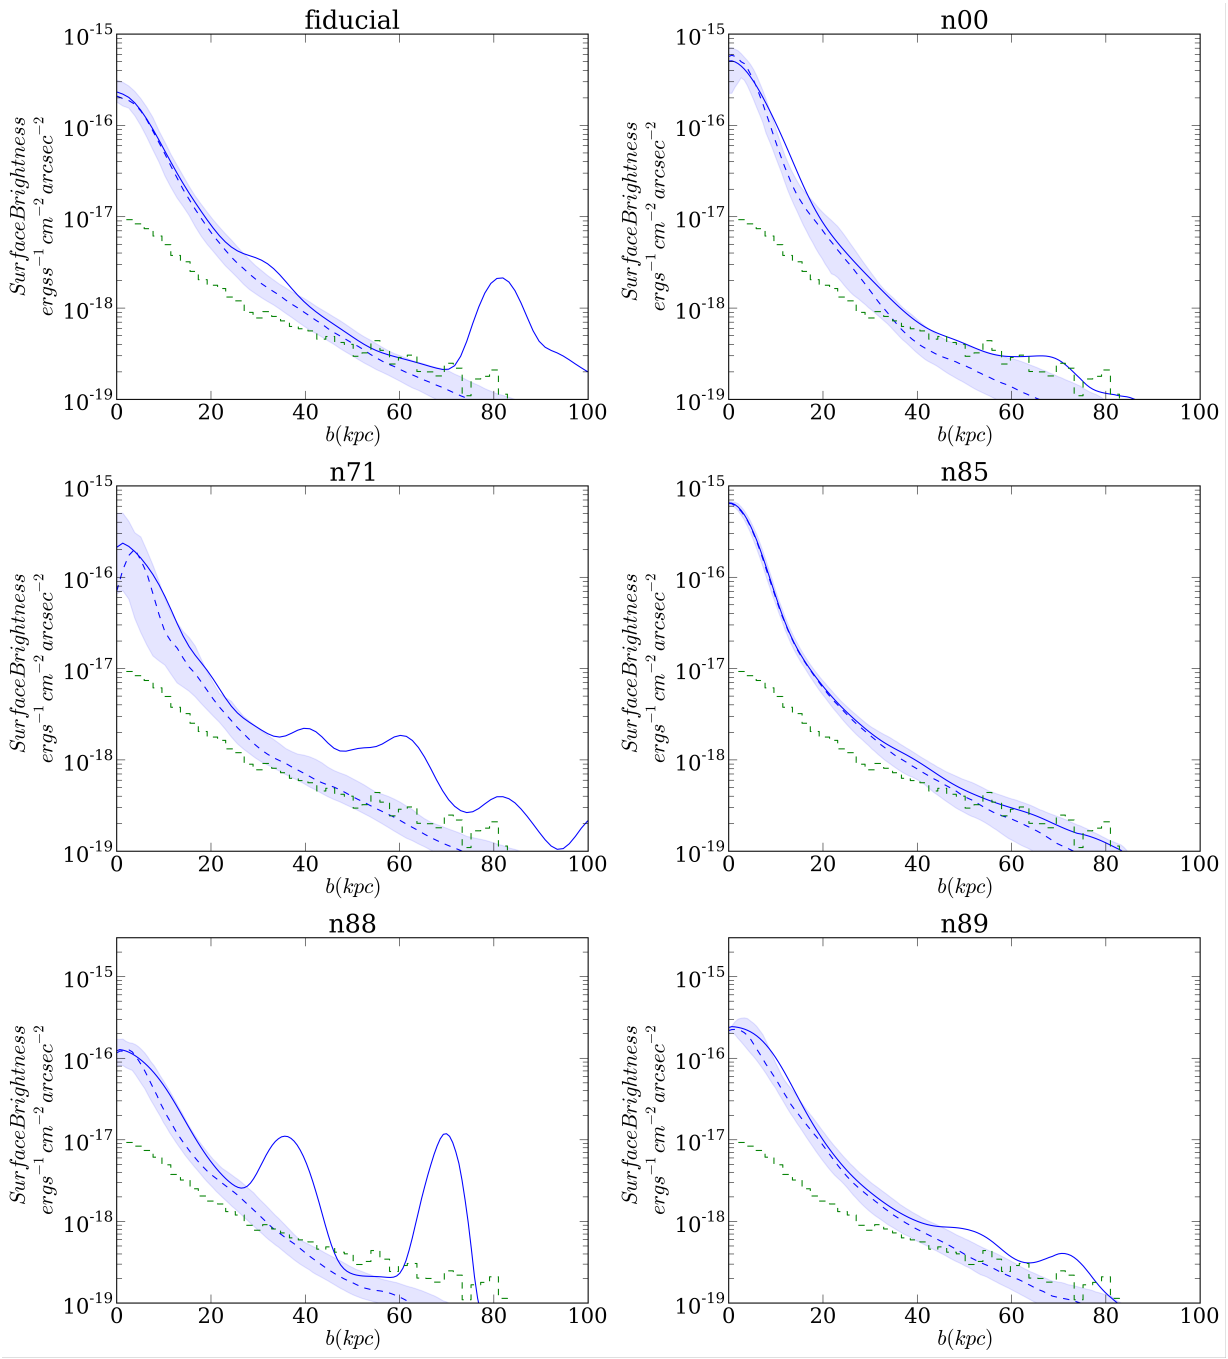

Figure 5.3: Lyα surface brightness profiles for the feedback models discussed in § 5.1. In these panels the profiles are azimuthally averaged and also averaged across 6 viewing orientations corresponding to the 3 primary axes. The solid blue line is the mean-averaged profile, whereas the dashed blue line is the median-averaged profile. The blue shaded region represents the 25^{th} and 75^{th} percentiles. The dashed green line shows the observations from Steidel et al. (2011).

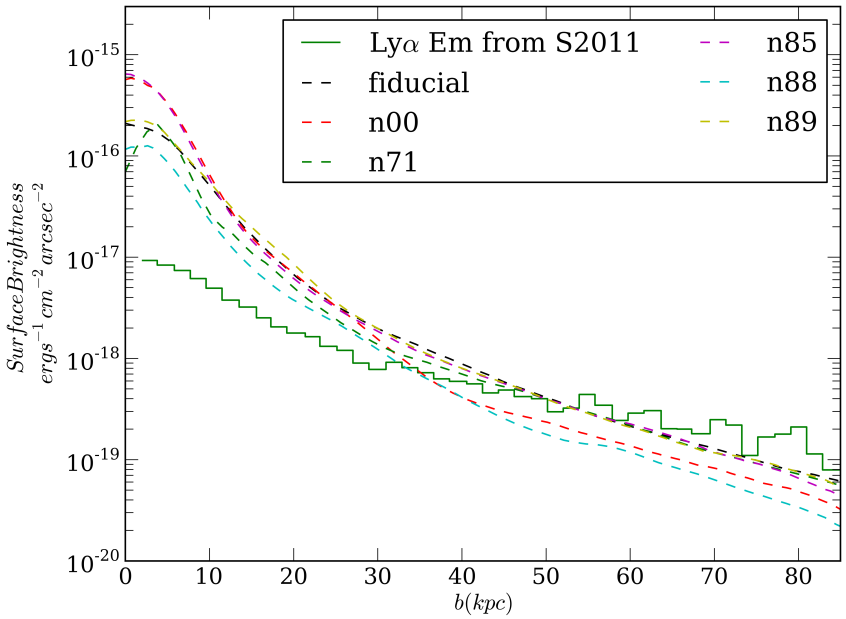

Figure 5.4: Comparison of the *median* Lyα emission profiles for the discussed outflow models. The dashed green line shows the observations from Steidel et al. (2011).

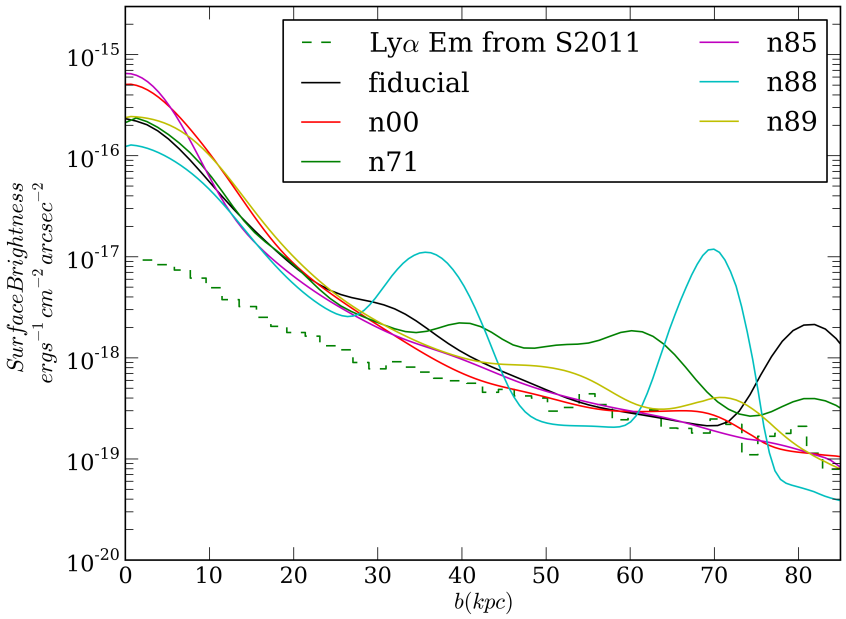

Figure 5.5: Comparison of the *mean* Lyα emission profiles for the previous models in addition to the fiducial model from chapter 4. The dashed green line shows the observations from Steidel et al. (2011).

If we focus first on the median emission profiles (Figure 5.4) we see that, reminiscent of our findings in Chapter 4, the profiles for the various outflow models are almost indistinguishable. This becomes more evident in light of the interquartile ranges shown in Figure 5.3. In fact, as in Chapter 4, the normalisation of the stellar mass in the galaxy in all of these simulations is somewhat arbitrarily set by the stellar mass of the galaxy in our fiducial simulation. The spread in the profiles would be even less if we attempted to match the normalisation of the individual simulations to the observations rather than to the fiducial model (for example by allowing the Lyα escape fraction to be a free parameter varying across simulations).

Turning to the mean surface brightness emission profiles (Figure 5.5) we see more differences between models. In particular the *fiducial*, *n00*, and *n85* models have fairly smooth profiles, whereas models *n71* and *n88* have 'bumpier' profiles, with *n89* lying somewhere in the spectrum between these two sets. The smooth profile models all produce similar mean profiles from about 20kpc outwards. In the inner 20kpc the profiles rise at different rates, giving central surface brightnesses varying by a factor of ~ 3. The bumps present in the bumpier profiles start to become pronounced around 35kpc from the centre. As noted in § 4.3.2 these bumps (and the deviation of the mean from the median at these impact parameters) are associated with substructure in the simulations, as confirmed by a visual examination of the density fields produced by the simulations. Thus, if we make the assumption that the 'bumpiness' of the profiles (equivalently the 'bittiness' of the surface brightness maps) is a consistent property of a given model (we re-visit this assumption in § 5.3), then given the mean profile of a galaxy we can distinguish between at least some of the models using 'bumpiness' as the metric.

Finally, we examine the emission spectra. Figure 5.6 shows the emission spectrum of the various models, analagous to Figure 4.6 in § 4.3.2. Figure 5.7 overlays the spectra to allow direct comparison. In these spectra the sign convention is such that a positive velocity offset corresponds to a lower frequency. All spectra show the Lyα line rising above the continuum between -400km/s and 400km/s. Echoing what we found in § 4.4 the red peak (when it exists) in these models is too close to the systemic velocity compared to observations, i.e. the velocity offset is too low. (Steidel et al., 2010) find the main Lyα peak at around +400km/s, whereas all of our simulated models have a red peak around 150km/s - 200km/s. This is indicative of the outflow velocities of our models being too low. In spite of this, several of the profiles show similarities to observations. In particular, the *fiducial*, *n88*, and *n89* models show a strongly enhanced red peak and suppressed blue peak. The *fiducial* model has a suppressed but still identifiable blue peak, as do the observations (albeit the velocity offset is different as previously mentioned). Additionally the ratio of the red to blue peak (~ 2.5) is similar to the observed peak strength ratio.

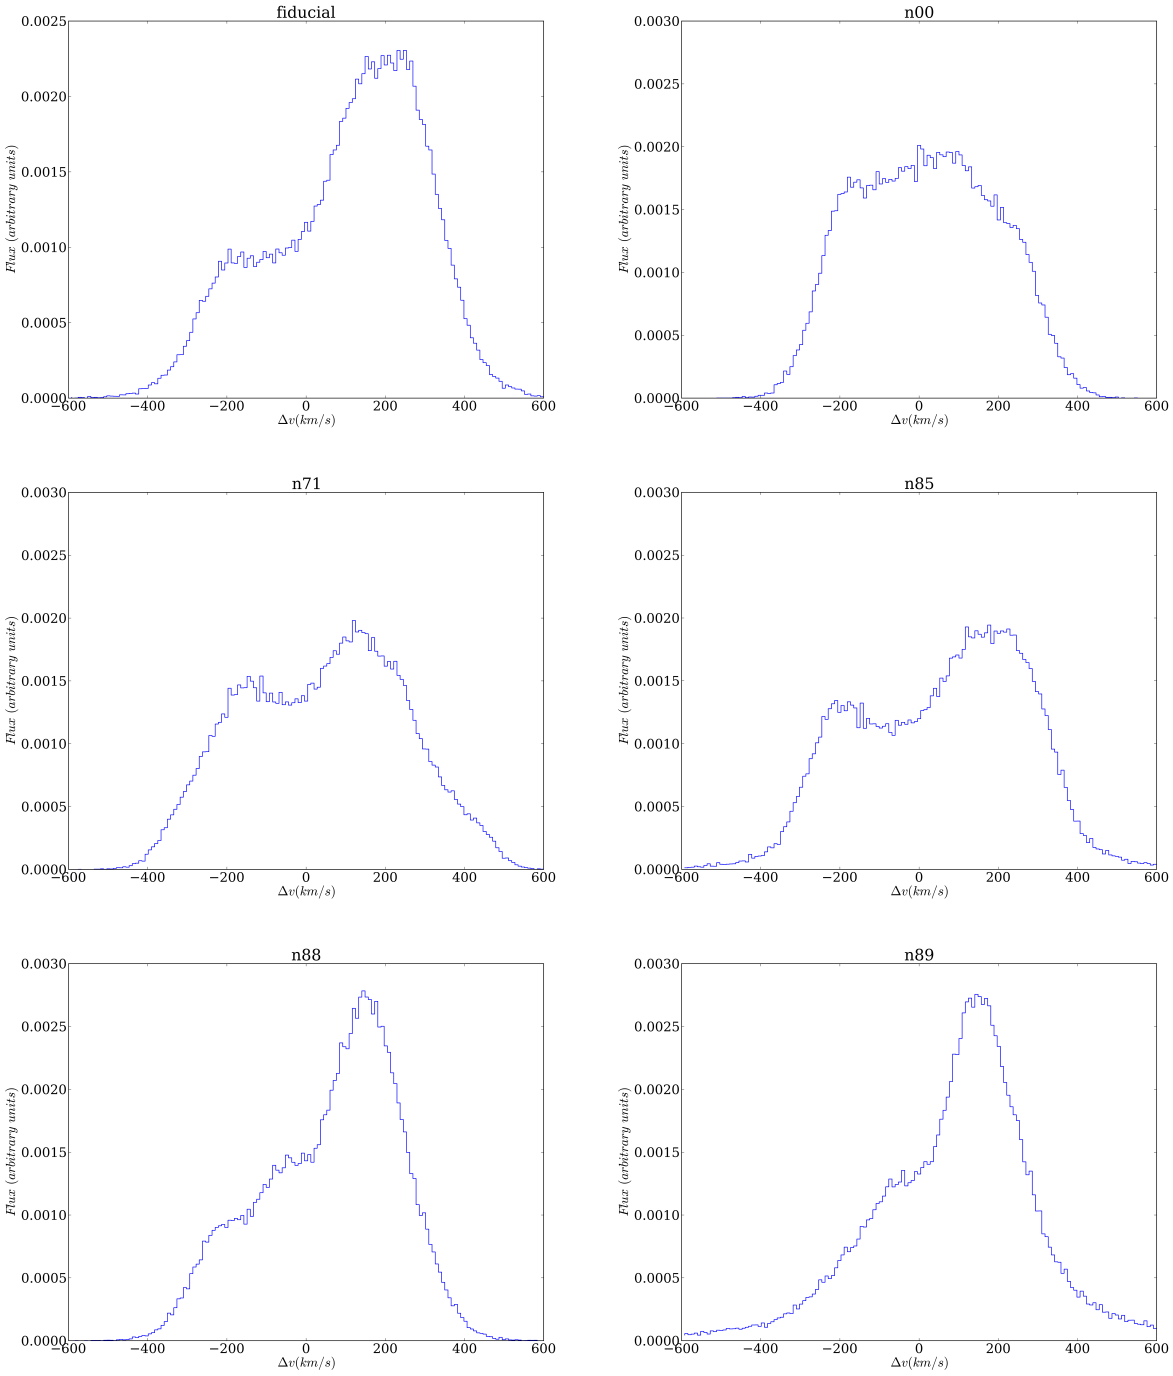

Figure 5.6: Lyα emission spectrum for all photons leaving the simulations run for the outflow models in § 5.1. The y-axis is flux in arbitary units, where the spectrum is normalised to unity. The x-axis is the velocity offset from line centre. In these spectra the sign convention is such that a positive velocity offset corresponds to a lower frequency.

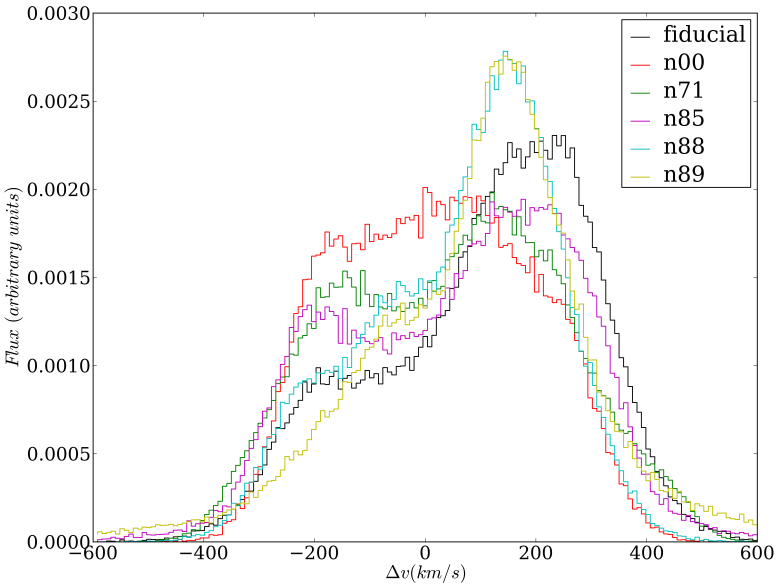

Figure 5.7: Comparison of the Lyα emission spectra from the outflow models described in § 5.1

5.3 Discussion

The results in § 5.2 highlight the importance of combining all of the available observational data, and simulating the same observables when trying to distinguish between feedback models. Taken on their own each of the three tests gives us some insight into how well our models are reflecting reality. However it is only when they are combined that a clearer picture of our models' successes, failures, and future directions becomes more apparent; we elaborate on this in point 4 below.

We just saw in the previous section that the 'bumpiness' of the Lyα surface brightness profile can be useful in distinguishing between outflow models. This is a nice result, but we would do well to ask what, if any, practical use it has. The reason we need to ask this is that the current observational surface brightness threshold of 1-2 $\times 10^{-18}$ erg s^{-1}cm^{-2}arcsec^{-2} dictates that we need to stack galaxies in order to see the SB profile beyond the central $\sim$20kpc. Therefore when we compare the simulations of a single galaxy to observations we are in fact comparing to stacked data (52 galaxies in the case of the data from Steidel et al. 2011). This topic has already been discussed in § 4.3.2, but we re-examine the situation in light of the results presented in this § 5.2. For cases such as the *fiducial* model, *n00*, *n85*, and *n89* we expect that if we were to simulate a suite of similar galaxies the small amount of substructure present would be smoothed out by the stacking procedure. This occurs because the occurance of substructure can be at various projected impact parameters depending on physical separation from the centre of the galaxy, and

orientation angle. Thus for smooth models, most realisations will not contain substructure at a given impact parameter – therefore mean-stacking many realisations washes out the substructure. We would expect then that in the case of e.g. model *n88* (Figure 5.3) this would result in the mean-stacking of many realisations producing a profile more similar to the median profile shown. On the other hand, 'bumpy' profiles such as the one produced by model *n71* (Figure 5.3) have a lot of substructure. If we again assume that this 'bumpiness' is a generic property of galaxies evolved with the *n71* wind prescription, then we would expect that by stacking many simulations the mean profile would not be washed out as in the case of the smooth models. That is, the *mean* profile for 'bumpy' models like *n71* may be comparable to stacked observations. If this is the case then we can use observations to distinguish between 'bumpy' models like *n71* and smooth models.

The results shown in § 5.2 extend our earlier finding in § 4.4 that different feedback models produce identifiably different emission line profiles. As discussed in § 4.3.2 the main difference is the degree of enhancement [suppression] of the red [blue] peak. This is related to the strength of the produced outflow. The clear differences which the various models exhibit (compare e.g. *n00* and *fiducial*) add another valuable piece of data to help us distinguish between models.

Ultimately it is true that different wind prescriptions lead to different Lyα observables although the differences are usually not radical. What we have presented in this chapter and the preceding chapter is a methodology to test wind models, and a demonstration that different models can produce results with some differences, and some striking similarities. In order to move beyond this proof-of-concept phase several things are needed:

1. We mentioned in § 5.2 that we assume that the 'bumpiness' of the resulting galaxy is a general result of a given feedback model. Critically we need to run more simulations to test whether it is in fact true that certain models generically produce smooth and bumpy galaxies as seen in Lyα . If this proves to be the case then perhaps the best hope for using Lyα as a diagnostic is to derive heuristics from many simulations that allow models to be identified. Observations can then be compared to these heuristics.

2. We also need to properly match the selection criteria of observations and simulations. Here, in order to make use of an existing suite of simulations, we merely select a halo (and correspondingly a set of initial conditions) based on the properties of the galaxy which evolved with a fiducial wind prescription. We then re-simulate that same halo with different wind prescriptions, rescaling sources to match the luminosity of the fiducial galaxy. We justify this by noting that: i) we are only demonstrating a proof-of-concept in this comparison, ii) we are interested in the structure of the CGM not the actual galaxy - although we acknowledge that the two are intimately related, and iii) the alternative of selecting a different halo based on the galaxy which evolves under a different wind prescription brings its own complication of comparing galaxies with different initial conditions. To treat the selection properly a wide range of galaxies with various outflow prescriptions should be simulated, and then from this a statistical sample of galaxies should be selected by matching the properties of

the evolved galaxies under different models to the properties of the observed galaxies. This brings us back to the need for a larger suite of simulations (i.e. point 1, above).

3. From the observational side, there is also scope for progress. Improved sensitivity might lead to imaging of individual systems. If there are clear trends in the appearance of galaxies which result from different models then improved constraints would increase the discriminating power of these tests. Observing more galaxy-galaxy pairs would also improve the constraints for the absorption test, and would help to provide more impact parameter data points. Looking at Figure 5.1 we see that several models are only marginally compatible with the data at 63kpc and 103kpc. Shrinking the observational error bars here would be of great help, with the caveat that we still need to consider point 2, above.

4. We need to consider more wind models. *None* of the models shown here accurately reproduce the central ∼25kpc of galaxies – neither in Lyα absorption nor Lyα emission. Additionally *none* of the models shown here correctly predict the observed Lyα emission spectrum. Therefore, if we are to find a model which matches observations we need some other models which need to be quite different to those tested here.

   From the absorption results we can see that all models shown here fail to rise fast enough towards the centre of the galaxy. This implies that either there is not enough HI in the central region, or that the line-of-sight velocity dispersion is too low towards the centre of the simulated galaxies, or both. A deficiency of central HI could be due to too much photoionization by local sources, too much star formation (obviously these two causes are related), too strong winds, or a combination of these factors.

   The emission surface brightness profile results all rise too quickly towards the centre of the simulated galaxies. That is, the Lyα from the galaxy is reaching the simulated observer too directly. This suggests that there is not enough HI in the central region, or that it has too large a bulk motion relative to the sources such that it is blue- or redshifted out of resonance with the incident Lyα photons.

   The emission spectrum results indicate that the outflow velocities in our simulations are too low.

   The explanation that is consistent with all of the absorption and emission results is that our outflows are too weak, while at the same time there is not enough HI in the central region. Indeed, the absorption profile from our inefficient wind model (*n71*) does show a sharper rise in the central region than the other models (Figure 5.2), suggesting that increasing the HI in the central region may alleviate the absorption profile issue. It is, however, not clear exactly how the central HI density can be increased while the outflow velocity is simultaneously increased. We note, however, that the central HI density needs to be increased in the central region, whereas the outflowing gas affecting the emission spectrum need not be centrally located.

   Interestingly *n00*, the only model here to use purely thermal supernova energy injection, produces the most different result from the others in absorption (where the

absorption profile flattens off too much in the central region) and in the emission spectrum (where the characteristic double-peaked Lyα profile is strikingly absent). This does perhaps hint that some forms of energy injection lead to better results than others. Finally, we note that the models we have presented here are all quite similar and often build off one another. To some extent this explains the similarity of the results.

Perhaps the clearest message from the results presented in this chapter is that hydrodynamical simulations are extremely hard to get right, and the fact that we can match a great many observables by implementing subresolution models should not lead us into a false belief that we are able to simulate galaxies successfully. There is still a great deal of work to be done in the simulation of galaxy formation and evolution. No doubt some progress will come as a direct result of the continuing march of technology (see chapter 1.2), whilst the rest will be down to improved models.

Observations are also hard at this redshift, but as ever will continue to improve with new techniques and instrumentation. What we have done in this work is provided the final piece of the puzzle which will allows us to test outflow models in detail. Perhaps this is not fully viable today but with this work we have shown that we are ready to exploit improved simulations and better observations, whenever they should arrive.

Chapter 6

Summary

Efforts to understand the processes which govern galaxy formation and evolution are at the frontier of modern cosmology. A thorough understanding of the formation and evolution of galaxies would build a bridge between the 'dark sector' and the baryonic and radiative processes that are accessible to direct observation. In particular galactic outflows remain poorly understood yet play a critical role in galactic evolution and the baryon cycle - the exchange of baryons between galaxies, the circumgalactic medium, and the intergalactic medium. In fact, until now extremely simplified shell models have been a popular description of galactic outflows. This popularity is fuelled by the simplicity of the models, but the fact that they are also surprisingly good at reproducing much of the observed phenomenology cannot be neglected.

In this thesis we have attempted to further our understanding of galactic outflows. We have done this by pursuing two complimentary approaches. On the one hand, recent observations have yielded galactic spectra with 'blue bumps' that pose a challenge to 'traditional' outflow shell models. Therefore, in Chapter 3 we attempt to make the minimal necessary modification to traditional outflow shell models that allows them to remain consistent with observations. On the other hand, we know that at some point spherical expanding shell models must break down - the reality of galaxy evolution is not so clean and simple. As such, in Chapters 4 and 5 we try to start from cosmological initial conditions, simulate the many processes governing galaxy formation and evolution with hydrodynamics, and then perform radiative transfer to compare our simulations with observations. Concordant with their scope, these projects have met varying degrees of success.

In Chapter 3 we find that an extremely simple modification has important implications for traditional shell models. Our modification adds a small static shell of HI outside, and directly adjacent to, the outflowing spherical shell. With this modification we find that our modified shell model can produce Lyα line profiles similar to the 'blue bumps' in the observed phenomenology. This is a feat hard to achieve with standard shell models and demonstrates that Fermi-like acceleration of Lyα photons across shocks may be important for the observed spectra of young galaxies. As a consequence the observed Lyα blue bumps may be a valuable tool for extracting knowledge of galaxy evolution from observations.

In Chapters 4 and 5 we have presented an ambitious attempt to simulate the Lyα

observables from high-redshift Lyα emitters. Yet the hubris which allowed us to even consider embarking on such an endeavour was tempered by a healthy skepticism of what could, at this time, be achieved. Indeed while we may have hoped at the outset that different feedback models should, would, nay *must* reveal vastly different Lyα observables, we were not surprised to find that the reality of the situation is more nuanced.

Simulation of emission surface brightness profiles, which at first seemed like a promising new test for galactic outflow models, in the end yielded only relatively small (though still potentially useful) differences between the tested models. The absorption equivalent width and emission spectrum tests yielded more discriminating power, although both were inconsistent with the observed data. In light of these results it would be easy to be disheartened with this approach. Yet such thinking misses a larger point.

From our vantage point today we often look back and perceive the progress of science as a staccato of milestone discoveries: Newton's Laws, the discovery of oxygen, relativity, the discovery of DNA, the realisation that 'island universes' were galaxies, the realisation that we ourselves live in such a galaxy. And yet, in reality a great deal of progress, and much of the body of knowledge which allows such leaps to be made, is obtained via evolutionary progression. We make models, we push them until they break, and then we use the knowledge of how they break to refine them or make new ones. Scientific progress is an iterative process, and having a test which our models fail is crucial to closing the iterative loop.

Judged in this light, the work described in Chapters 4 and 5 is a resounding success. The new tests which we present may not show a great deal of variation between models, but one result is very clear: *none* of the models hold up to *either* of the three observational tests. Of course, this is not entirely surprising - we are dealing with extremely simplified models which require sub-grid recipes (star formation, interaction between the multiple phases of the ISM, energy deposition from supernovae), and in which we know that we have introduced unphysical aspects (the most egregious of which are the adding of winds 'by hand' and the temporary hydrodynamical decoupling of the particles used to implement the wind). Nevertheless, models which when judged by other criteria have seemed very successful (see § 4.2 for the successes of our fiducial model), have been found wanting. In fact we have presented 3 tests which, since in principle they can be passed independently, combined give a strong constraint on future outflow models.

In closing, the outlook for further progress via hydrodynamical and radiative transfer simulations is bright. As discussed in § 1.2, even pessimistic predictions suggest that we have close to a decade left for Moore's Law to run. This corresponds to an increase in computing power of ~ 2 orders of magnitude. Needless to say such an increase in available power would be extremely welcome. By combining hydrodynamical simulations with numerical radiative transfer we are able to predict the observable signatures of outflow models. As the available computing power increases and numerical hydrodynamics techniques advance, our confidence in our ability to predict the observables of a given model will only increase. Combined with the promise of more, and better, observations we expect such modelling to be a powerful asset in deciphering the complexities of galactic outflows, and understanding their place in moulding the universe that we call home.

Appendix A

Fermi-like Acceleration Appendix

A.1 Origin of $P(\mu) \propto \mu$

The goal of this appendix is to clarify why the distribution of angles at which Lyα photons cross the shock front scales as $P(\mu) \propto \mu$ (where $\mu \equiv \cos\theta$), rather than isotropically.

Photons that cross the shock front will on average have traversed an optical depth $\tau = 1$. Depending on μ, $\tau = 1$ corresponds to a different (frequency dependent) physical depth, $d(\nu)$, away from the shock front. It is easy to see that $d \propto \mu$, so that at fixed frequency the volume of gas that the photon is likely to have last scattered in is $V_{\rm scat} \propto \mu$. If we further assume that the density of scattering events is homogeneously distributed in the gas, then this leads directly to the relation $I \propto \mu$, where I is the intensity of photons. Therefore $P(\mu) \propto \mu$, which we showed in Figure 3.2, and which was found previously by Ahn et al. (2001) and Garavito-Camargo et al. (2014).

Figure A.1 shows the average angle at which photons cross the shock front, $\langle\mu\rangle = \int_0^1 d\mu\, \mu P(\mu)$, as a function of crossing number l. As we mentioned in the paper, odd [even] number of shock crossings correspond to crossings from left-to-right [right-to-left]. This Figure clearly shows that $\langle\mu\rangle \sim 2/3$ for even shock-crossings, while it approaches $\langle\mu\rangle \sim 0.5$ for odd-shock crossings when $l > 3$.

An explanation for this is offered in Figure A.2, which shows the mean distance from the shock front (scattering depth) of the last scattering before crossing the shock as a function of the shock crossing number l. This distance follows an inverse exponential distribution. As the photon becomes more blueshifted with increasing l, the mean free path of the photon also increases, so that for each successive l the photon is coming from deeper within the slab, explaining the increasing trend in Figure A.2. However, as the mean free path increases, the photons which propagate the farthest actually propagate all the way through the gas slab and exit the simulation. Therefore, the distribution of scattering depth is truncated for the next scattering count, resulting in a flattening in the growth of scattering depth as a function of l. This is precisely the behaviour we observe for odd-l in Figure A.2. The scattering depth distribution is truncated, but the photon blueshift is unaffected. Therefore, the gas between the next scattering event and the shock

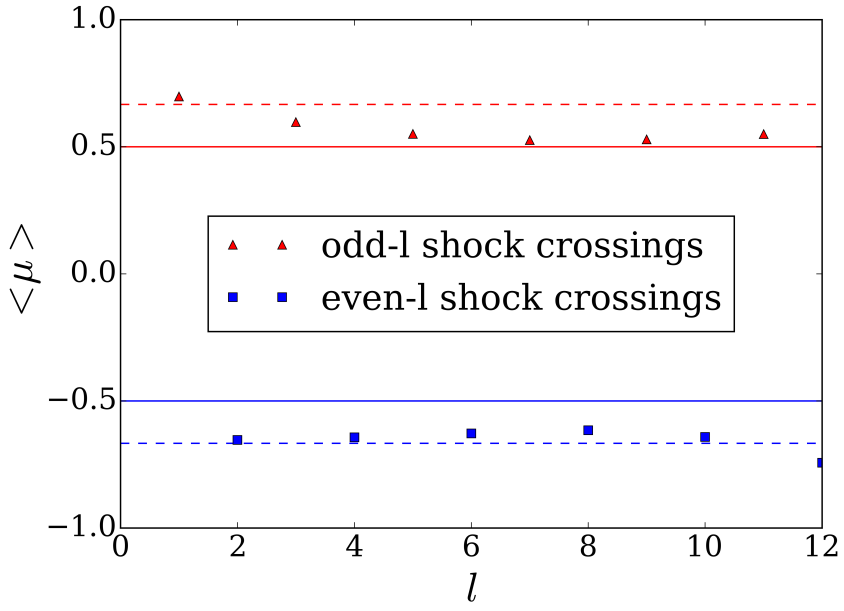

Figure A.1: Mean projection as a function of shock crossing number, l. The solid lines represent the expected values for isotropic scattering, while the dashed lines show $<\mu>= \pm 2/3$. Red triangular points are for odd-l shock crossings, blue square points are for even-l shock crossings.

front becomes effectively optically thin to the photon, and we expect photons to cross the shock front isotropically, which is why $\langle\mu\rangle$ approaches 0.5 in Figure A.1.

For even-l we do not observe this transition to isotropic crossing. This is because $N_1 \ll N_c$, so that at a given l and blueshift the mean free path of a photon is much lower in the right (static) slab and the scattering depth distribution for even-l (i.e. crossings where the previous scattering was in the right, static slab) has not yet become truncated. As discussed above it is the truncation of the scattering depth distribution that causes the transition to isotropic shock crossing, and so in its absence for even-l the Lyα photons still cross the shock front following $P(\mu) \propto \mu$.

Finally, we note that the highest shock crossing counts that occur in our simulation are rare events, and the resultant poor statistics are responsible for the turnaround in Figure A.2 (recall that we are sampling from an inverse exponential distribution).

A.2 Interstellar dust

In dusty media the increase in path length caused by scattering can allow the dust to significantly attenuate the Lyα flux that escapes. Because photons of different frequencies have different scattering cross-sections the average number of scatterings they undergo will be different. Therefore they will have different changes in path-length, and subsequently undergo different levels of dust attenuation.

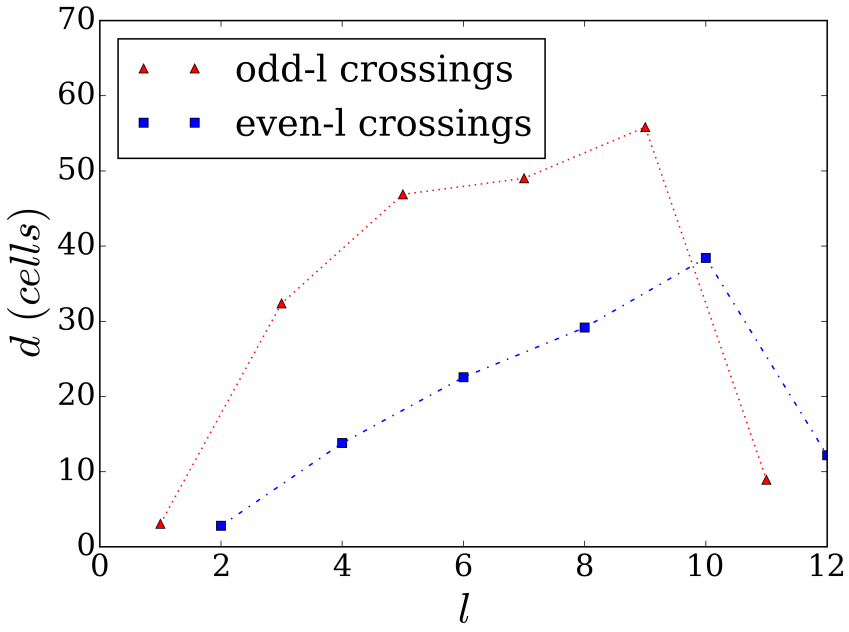

Figure A.2: Mean distance between the shock and the location of the last scattering before crossing the shock vs shock crossing number, l. d is shown in simulation units, i.e. cells (each slab is 128 cells wide). Red triangular points and the red dotted line show the odd-l crossings, while blue square points and the blue dot-dash line show the even-l crossings.

We take a direct approach to address this issue and perform numerical experiments where grey dust is added to either the inner, outer, or both shells. The results are shown in Figure A.3. It is clear that the differences in the emergent spectrum caused by the different dust prescriptions are very small.

The addition of dust slightly enhances the blue bump. This occurs because once a photon is Fermi accelerated the optical depth of the HI shells to the photon is vastly reduced and the photon easily escapes, avoiding or reducing the number of further scatterings. Thus, a Fermi acclerated photon has a shorter total path-length so that the presence of dust affects it less than an unaccelerated photon. This results in a slightly enhanced blue peak in the normalised spectra shown in Figure A.3.

A.3 Varying the HI Column Density in the Static Shell

Figure A.4 shows that the blue bump remains visible even if we further reduce $N_{\rm HI}^{\rm stat}$ by order(s) of magnitude.

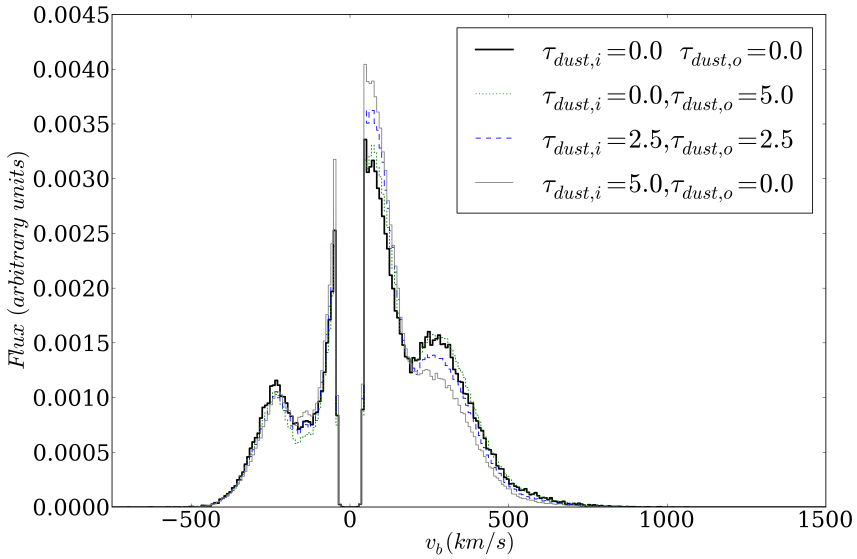

Figure A.3: Emergent spectrum for a shell-model with $(N_{\rm HI}^{\rm shell}, N_{\rm HI}^{\rm stat}, v_{\rm shell}) = (10^{19}\,{\rm cm}^{-2}, 10^{18}\,{\rm cm}^{-2}, 200{\rm km\,s}^{-1})$ including three different dust prescriptions. For each histogram $\tau_{dust,i}$ refers to the optical depth of the inner shell, $\tau_{dust,o}$ refers to the optical depth of the outer shell. The reference dust-free spectrum is shown with a heavy black line.

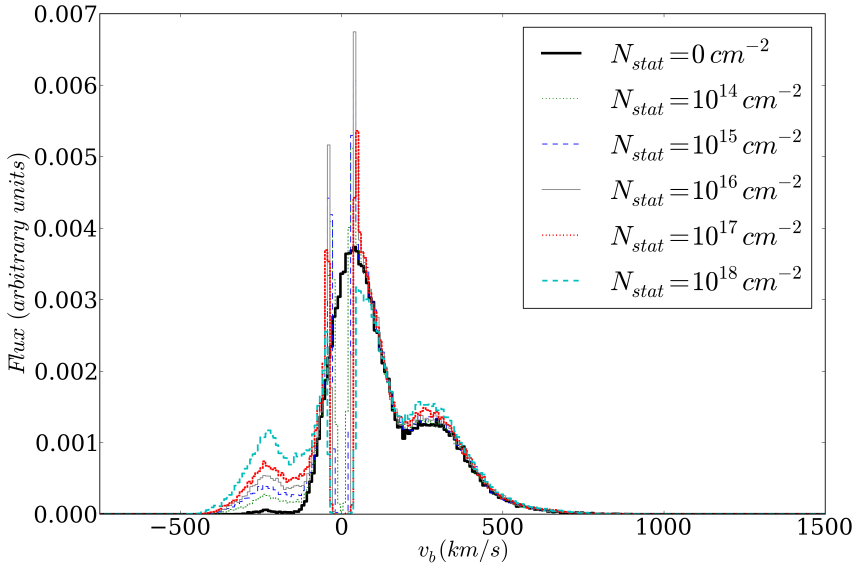

Figure A.4: Same as Figure A.3, but here we vary $N_{\rm HI}^{\rm stat}$ while keeping the other parameters fixed. The prominence of the blue peak reduces with HI column density, but a blue bump clearly remains even when $N_{\rm HI}^{\rm stat} = 10^{15}$ cm^{-2}.

Appendix B

Derivation of absorption equivalent width vs impact parameter relation

First we start with the standard definition of equivalent width (albeit in frequency-space):

$$\mathrm{EW}_\nu(b) = \int_{\nu_{min}}^{\nu_{max}} \frac{[1 - I(\nu, b)]}{I_0} d\nu \tag{B.1}$$

where ν is the frequency, b is the impact parameter, I_0 is the initial intensity, and I is the intensity at the detector, which is dependent on the temperature, T:

$$I(\nu, b) = I_0 e^{-\tau(\nu,T,b)} \tag{B.2}$$

In all of our calculations we set $I_0 = 1$, whence:

$$\mathrm{EW}_\nu(b) = \int_{\nu_{min}}^{\nu_{max}} [1 - e^{-\tau(\nu,T,b)}] d\nu \tag{B.3}$$

The optical depth of the HI gas that a sightline passes through is dependent on the physical column of gas through which the sightline passes, and the cross section of the gas at the specified frequency:

$$\tau(\nu, T, b) = n\sigma(\nu, T)s(b) \tag{B.4}$$

where n is the HI density, T is the temperature, $\sigma(\nu, T)$ is the frequency and temperature-dependent cross-section, and $s(b)$ is the physical length of HI traversed by a sightline at impact parameter, b.

In order to simplify our final expression[1] we use the approximation that $\sigma_\nu \approx \sigma_{\nu,doppler}$, where $\sigma_{\nu,doppler}$ is the absorption cross-section calculated using the thermal doppler broadening line profile, $\phi_{doppler}$.

$$\sigma_{\nu,doppler}(\nu, T) = \frac{g_2}{g_1} \left(\frac{c^2}{8\pi\nu^2} \right) A_{21} \phi_{doppler}(\nu, T) \tag{B.5}$$

[1] neither LAF nor SLAF make this simplification, and instead use the Voigt profile

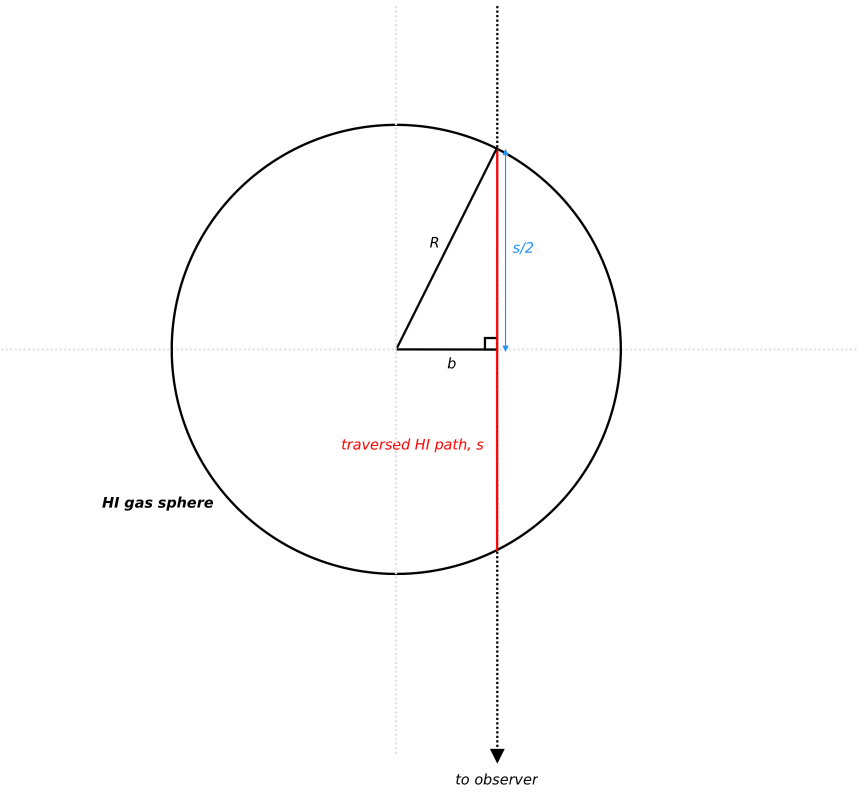

Figure B.1: Schematic showing the geometry of the test case. R is the radius of the gas sphere, b is the impact parameter. The line segment in red is s, the segment of the path which intersects the gas sphere.

$$\phi_{doppler}(\nu, T) = \frac{e^{-\frac{(\nu-\nu_0)^2}{\Delta\nu_D(T)^2}}}{\sqrt{\pi}\Delta\nu_D(T)} \tag{B.6}$$

where $g_2 = 8$ and $g_1 = 2$ are the statistical weights of the second and first energy levels, A_{21} is the Einstein A-coefficient for the Lyα transition, ν_0 is the line-centre frequency, $\Delta\nu_D(T) = \frac{\nu_0}{c}\sqrt{2k_bT/m_H}$ is the Doppler width, k_b is the Boltzmann constant, and m_H is the mass of a hydrogen atom.

Figure B.1 shows a schematic of the test case. It is apparent from simple geometry that $(s/2)^2 + b^2 = R^2$, yielding:

$$s(b) = 2\sqrt{R^2 - b^2} \tag{B.7}$$

Finally, combining Equations B.3, B.4, and B.7 we obtain:

$$\mathrm{EW}_\nu(b) = \int_{\nu_{min}}^{\nu_{max}} [1 - e^{-2n\sqrt{R^2-b^2}\sigma(\nu,T)}]d\nu \tag{B.8}$$

the same result as combining Equations 2.2 and 2.3.

Bibliography

Adams, T. F. 1972, ApJ, 174, 439

Adams J. J., Hill G. J., MacQueen P. J., 2009, ApJ, 694, 314

Ahn, S.-H., Lee, H.-W., & Lee, H. M. 2001, ApJ, 554, 604

Ahn, S.-H., Lee, H.-W., Lee, H. M., 2002, ApJ, 567, 922

Aumer, M., White, S. D. M., Naab, T., Scannapieco, C., MNRAS, 434, 3142

Barkana, R. & Loeb, A., 2001, Physics Reports, 349, 125

Barnes, L. A., Haehnelt, M. G., Tescari, E., & Viel, M. 2011, MNRAS, 416, 1723

Behrens, C., & Braun, H. 2014, A&A, 572, A74

Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N.,, Limon, M., Mather, J., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G., Wilkinson, D. T., Wollack, E., Wright, E. L., Bulletin of the American Astronomical Society, 29, 1353

Benson, A. J., Bower, R. G., Frenk, C. S., Lacey, C. G., Baugh, C. M., Cole, S., 2003, ApJ, 599, 38

Cantalupo, S., Porciani, C., Lilly, S. J., Miniati, F., 2005, ApJ, 628, 61

Cantalupo, S., Lilly, S. J. & Porciani, C., 2007, ApJ, 657, 135

Chabrier, G., 2003, PASP, 115, 763

Chung, A. S., Dijkstra, M., Ciardi, B., Gronke, M., 2016, MNRAS, 455, 884

Ciardi, B. and Ferrara, A., 2005, Space Sci. Rev., 16, 625

Ciardi, B. and Ferrara, A. and Marri, S. and Raimondo, G., 2001, MNRAS, 324, 381

Chonis, T. S., Blanc, G. A., Hill, G. J., et al. 2013, ApJ, 775, 99

Cole, S., Lacey, C. G., Baugh, C. M., Frenk, C. S., 2000, MNRAS, 319, 168

Dalla Vecchia, C. and Schaye, J., 2008, MNRAS, 387, 1431

Dalla Vecchia, C. & Schaye, J., 2012, MNRAS, 426, 140

Dekel, A., & Birnboim, Y. 2006, MNRAS, 368, 2

Dekel, A., Birnboim, Y., Engel, G., Freundlich, J., Goerdt, T., Mumcuoglu, M., Neistein, E., Pichon, C., Teyssier, R., Zinger, E., 2009, *Nature*, 457, 451

Dijkstra, M., Haiman, Z., & Spaans, M. 2006a, ApJ, 649, 14

Dijkstra, M., Haiman, Z., & Spaans, M. 2006b, ApJ, 649, 37

Dijkstra, M., Lidz, A., Wyithe, J. S. B., 2007, MNRAS, 377, 1175

Dijkstra, M. & Loeb, A., 2009, MNRAS, 400, 1109

Dijkstra, M. & Wyithe, J. S. B., 2010, MNRAS, 408, 352

Dijkstra, M. and Kramer, R. 2012, MNRAS, 424, 1672

Dijkstra, M. 2014, PASA, 31, e040

Djorgovski S., 1988, in Kron R. G., Renzini A., eds, Towards Understanding Galaxies at Large Redshift Vol. 141 of Astrophysics and Space Science Library, Lyman-alpha galaxies at large redshifts., 259

Fan, X., Carilli, C. L., Keating, B., 2006, ARA&A, 44, 415

Faucher-Giguère, C.-A., Kereš, D., Dijkstra, M., Hernquist, L., Zaldarriaga, M., 2010, ApJ, 725, 633

Feldmeier, J. J., Hagen, A., Ciardullo, R., Gronwall, C., Gawiser, E., Guaita, L., Hagen, L. M. Z., Bond, N. A., Acquaviva, V., Blanc, G. A., Orsi, A., Kurczynski, P., 2013, ApJ, 776, 75

Fontana, A., Vanzella, E., Pentericci, L., Castellano, M., Giavalisco, M., Grazian, A., Boutsia, K., Cristiani, S., Dickinson, M., Giallongo, E., Maiolino, R., Moorwood, A., Santini, P., 2010, ApJL, 725, L205

Frenk, C. S. and White, S. D. M., 2012, Annalen der Physik, 524, 507

Fryxell, B., Olson, K., Ricker, P., Timmes, F. X., Zingale, M., Lamb, D. Q., MacNeice, P., Rosner, R., Truran, J. W., Tufo, H., 2000, ApJS, 131, 273

Fukugita, M. & Peebles, P. J. E., 2004, ApJ, 616, 643

Garavito-Camargo, J. N., Forero-Romero, J. E., & Dijkstra, M. 2014, ApJ, 795, 120

Garel, T., Blaizot, J., Guiderdoni, B., Schaerer, D., Verhamme, A., Hayes, M., 2012, MNRAS, 422, 310

Genel, S., Naab, T., Genzel, R., Förster Schreiber, N. M., Sternberg, A., Oser, L., Johansson, P. H., Davé, R., Oppenheimer, B. D., Burkert, A., 2012, ApJ, 745, 11

Gnedin, N. Y., Kravtsov, A. V., Chen, H.-W., 2008, ApJ, 672, 765

Goerdt, T., Dekel, A., Sternberg, A., Gnat, O., Ceverino, D., 2012, MNRAS, 424, 2292

Graziani, L., Maselli, A. & Ciardi, B., 2013, MNRAS, 431, 722

Hatton, S., Devriendt, J. E. G., Ninin, S., Bouchet, F. R., Guiderdoni, B., Vibert, D., 2003, MNRAS, 343, 75

Haardt, F. and Madau, P., 2001, Clusters of Galaxies and the High Redshift Universe Observed in X-rays

Haardt, F. and Madau, P., 2012, ApJ, 746, 125

Haiman, Z. and Rees, M. J., 2001, ApJ, 556, 87

Hayes, M., Scarlata, C., Siana, B., 2011, *Nature*, 476, 304

Harrington, J. P. 1973, MNRAS, 162, 43

Hashimoto, T., Verhamme, A., Ouchi, M., et al. 2015, arXiv:1504.03693

Hopkins, P. F., Quataert, E. & Murray, N., 2011, MNRAS, 417, 950

Hopkins, P. F., Quataert, E. & Murray, N., 2012, MNRAS, 421, 3522

Iye, M., Ota, K., Kashikawa, N., Furusawa, H., Hashimoto, T., Hattori, T., Matsuda, Y., Morokuma, T., Ouchi, M., Shimasaku, K., 2006, *Nature*, 443, 186

Jeeson-Daniel, A., Ciardi, B., Maio, U., Pierleoni, M., Dijkstra, M., Maselli, A., 2012, MNRAS, 424, 2193

, 1984, ApJL, 284, L9

Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, MNRAS, 363, 2

Kollmeier, J. A., Zheng, Z., Davé, R., Gould, A., Katz, N., Miralda-Escudé, J., Weinberg, D. H., 2010, ApJ, 708, 1048

Kulas, K. R., Shapley, A. E., Kollmeier, J. A., et al. 2012, ApJ, 745, 33

Kunth, D., Mas-Hesse, J. M., Terlevich, E., et al. 1998, A&A, 334, 11

Lacey, C. and Cole, S., 1993, MNRAS, 262, 627

Lake, E., Zheng, Z., Cen, R., et al. 2015, ApJ, 806, 46

Laursen, P., Razoumov, A. O., & Sommer-Larsen, J.,2009, ApJ, 696, 853

Laursen, P., Sommer-Larsen, J., Razoumov, A. O., 2011, ApJ, 728, 52

Lehnert, M. D., Nesvadba, N. P. H., Cuby, J.-G., Swinbank, A. M., Morris, S., Clément, B., Evans, C. J., Bremer, M. N., Basa, S., 2010, *Nature*, 467, 940

Leitherer, C., Schaerer, D., Goldader, J. D., González Delgado, R. M., Robert, C., Kune, D. F., de Mello, D. F., Devost, D., Heckman, T. M., 1999, ApJS, 123, 3

Lovell, M. R., Eke, V., Frenk, C. S., Gao, L., Jenkins, A., Theuns, T., Wang, J., White, S. D. M., Boyarsky, A., Ruchayskiy, O., 2012, MNRAS, 420, 2318

Maselli, A., Ferrara, A., & Ciardi, B., 2003, MNRAS, 345, 379

Maselli, A., Ciardi, B., & Kanekar, A., 2009, MNRAS, 393, 171

Mathis, J. S., Mezger, P. G., & Panagia, N., 1983, A&A, 128, 212

Matsuda, Y., Yamada, T., Hayashino, T., Tamura, H., Yamauchi, R., Ajiki, M., Fujita, S. S., Murayama, T., Nagao, T., Ohta, K., Okamura, S., Ouchi, M., Shimasaku, K., Shioya, Y., Taniguchi, Y., 2004, AJ, 128, 569

Matsuda, Y., Yamada, T., Hayashino, T., Yamauchi, R., Nakamura, Y., Morimoto, N., Ouchi, M., Ono, Y., Umemura, M., Mori, M., 2012, MNRAS, 425, 878

Momose, R., Ouchi, M., Nakajima, K., Ono, Y., Shibuya, T., Shimasaku, K., Yuma, S., Mori, M., Umemura, M., 2014, MNRAS, 442, 110

Moore, G.E., 1965, Electronics, 38.8, 82

Mori, M. and Umemura, M. and Ferrara, A., 2004, ApJ, 613, L97

Mori, M. and Umemura, M., 2006, *Nature*, 440, 644

Neufeld D. A., McKee C. F., 1988, ApJL, 331, L87

Neufeld, D. A. 1990, ApJ, 350, 216

Oppenheimer, B. D. and Davé, R., 2006,MNRAS, 373, 1265

Oppenheimer, B. D. and Davé, R., 2008, MNRAS, 387, 577

Orsi, A., Lacey, C. G. & Baugh, C. M., 2012, MNRAS, 425, 87

Osterbrock, D. E. 1962, ApJ, 135, 195

Peebles, P. J. E. & Dicke, R. H., 1968, ApJ, 154, 891

Peebles, P. J. E., 1970, AJ, 75, 13

Peebles, P. J. E., 1973, PASJ, 25, 291

Peebles, P. J. E., 1974, ApJL, 189, L51

Peebles, P. J. E., 1980, The large-scale structure of the universe

Peebles, P. J. E., 1998, arXiv:astro-ph/9806201

Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., Hook, I. M., Kim, A. G., Kim, M. Y., Lee, J. C., Nunes, N. J., Pain, R., Pennypacker, C. R., Quimby, R., Lidman, C., Ellis, R. S., Irwin, M., McMahon, R. G., Ruiz-Lapuente, P., Walton, N., Schaefer, B., Boyle, B. J., Filippenko, A. V., Matheson, T., Fruchter, A. S., Panagia, N., Newberg, H. J. M., Couch, W. J., Project, T. S. C. 1999, ApJ, 517, 565

Pierleoni, M., Maselli, A., Ciardi, B., 2009, MNRAS, 393, 872

Planck Collaboration, Ade, P. A. R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A. J., Barreiro, R. B., Bartlett, J. G., et al., 2015, arXiv:1502.01589

Prescott, M. K. M., Dey, A., Brodwin, M., Chaffee, F. H., Desai, V., Eisenhardt, P., Le Floc'h, E., Jannuzi, B. T., Kashikawa, N., Matsuda, Y., Soifer, B. T., 2012, ApJ, 752, 86

Rakic, O., Schaye, J., Steidel, C. C., Rudie, G. C., 2012, ApJ, 751, 94

Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., Leibundgut, B., Phillips, M. M., Reiss, D., Schmidt, B. P., Schommer, R. A., Smith, R. C., Spyromilio, J., Stubbs, C., Suntzeff, N. B., Tonry, J., 1998, AJ, 116, 1009

Rosdahl, J. and Blaizot, J., 2012, MNRAS, 423, 344

Rubin, V. C., Thonnard, N., Ford, Jr., W. K., 1978, ApJL, 225, L107

Rudie, G. C., Steidel, C. C., Trainor, R. F., Rakic, O., Bogosavljević, M., Pettini, M., Reddy, N., Shapley, A. E., Erb, D. K., Law, D. R., 2012, ApJ, 750, 67

Scannapieco, C., Tissera, P. B., White, S. D. M., Springel, V., 2005, MNRAS, 364, 552

Scannapieco, C., Tissera, P. B., White, S. D. M., Springel, V., 2006, MNRAS, 371, 1125

Scannapieco, C., Tissera, P. B., White, S. D. M., Springel, V., 2008, MNRAS, 389, 1137

Scannapieco, C., Wadepuhl, M., Parry, O. H., Navarro, J. F., Jenkins, A., Springel, V., Teyssier, R., Carlson, E., Couchman, H. M. P., Crain, R. A., Dalla Vecchia, C., Frenk, C. S., Kobayashi, C., Monaco, P., Murante, G., Okamoto, T., Quinn, T., Schaye, J., Stinson, G. S., Theuns, T., Wadsley, J., White, S. D. M., Woods, R., 2012, MNRAS, 423, 1726

Schaerer, D., Hayes, M., Verhamme, A., & Teyssier, R. 2011, A&A, 531, A12

Schaye, J. & Dalla Vecchia, C., MNRAS, 383, 1210

Shen, S., Madau, P., Guedes, J., Mayer, L., Prochaska, J. X., Wadsley, J., 2013, ApJ, 765, 89

Shibuya, T., Kashikawa, N., Ota, K., Iye, M., Ouchi, M., Furusawa, H., Shimasaku, K., Hattori, T., 2012, ApJ, 752, 114

Smith, B. M., Malhotra, S., Rhoads, J., Finkelstein, S., Zheng, Z., Wang, J., 2012, American Astronomical Society Meeting Abstracts, 219, 340.06

Springel, V. and Hernquist, L., 2002, MNRAS, 333, 649

Springel, V. and Hernquist, L., 2003, MNRAS, 339, 289

Springel, V., 2005, MNRAS, 364, 1105

Springel, V., White, S. D. M., Jenkins, A., Frenk, C. S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., Peacock, J. A., Cole, S., Thomas, P., Couchman, H., Evrard, A., Colberg, J., Pearce, F., 2005, *Nature*, 435, 629

Springel, V., 2011, Computational Star Formation, 270, 203

Steidel, C. C., Adelberger, K. L., Shapley, A. E., Pettini, M., Dickinson, M., Giavalisco, M., 2000, ApJ, 532, 170

Steidel, C. C., Erb, D. K., Shapley, A. E., Pettini, M., Reddy, N., Bogosavljević, M., Rudie, G. C., Rakic, O., 2010, ApJ, 717, 289

Steidel, C. C., Bogosavljević, M., Shapley, A. E., Kollmeier, J. A., Reddy, N. A., Erb, D. K., Pettini, M., 2011, ApJ, 736, 160

Strauss M. A., McCarthy P. J., Spinrad H., Djorgovski S., van Breugel W., Liebert J., Dickinson M., 1987, in Bergeron J., Kunth D., Rocca-Volmerange B., Tran Thanh Van J., eds, High Redshift and Primeval Galaxies 3C 326.1 - A forming galaxy at a redshift of 1.82?, 85

Tasitsiomi, A., 2006, ApJ, 645, 792

Tauber, J. A., Mandolesi, N., Puget, J.-L., Banos, T., Bersanelli, M., Bouchet, F. R., Butler, R. C., Charra, J., Crone, G., Dodsworth, J., et al., 2010, A&A, 520, A1

Teyssier, R., 2002, A&A, 385, 337

van de Voort, F., Schaye, J., Altay, G., Theuns, T., 2012, MNRAS, 421, 2809

Verhamme, A., Schaerer, D., & Maselli, A. 2006, A&A, 460, 397

Verhamme, A., Schaerer, D., Atek, H., & Tapken, C. 2008, A&A, 491, 89

Verhamme, A., Dubois, Y., Blaizot, J., et al. 2012, A&A, 546, A111

,Wadsley, J. W., Stadel, J., Quinn, T., 2004, New Astronomy, 9, 137

White, S. D. M., Frenk, C. S., Davis, M., 1983, ApJL, 274, L1

White, S. D. M., Davis, M., Frenk, C. S., 1984, MNRAS, 209, 27P

White, S. D. M. and Frenk, C. S., 1991, ApJ, 379, 52

Yajima, H., Li, Y., Zhu, Q., Abel, T.,2012,MNRAS, 424, 884

Yajima, H., Li, Y., Zhu, Q., Abel, T.,2015,ApJ, 801, 52

Yusef-Zadeh, F. and Morris, M. and White, R. L., 1984, ApJ, 278, 186

Zheng, Z., & Miralda-Escudé, J. 2002, ApJ, 578, 33

Zwicky, F., 1933, Helvetica Physica Acta, 6, 110

Acknowledgements

Firstly, I would like to thank Benedetta Ciardi for advising my thesis. There have been many times when I have given her cause for great concern, and I am immensely grateful for her patience, and for trusting me enough to allow me to find my own way. My PhD experience has been much richer, and all the more valuable for it.

I must also extend my greatest thanks to Mark Dijkstra. It is no exaggeration to say that this thesis would not have been possible without his advising, guidance, sharing of ideas, and encouragement. Thank you for your faith in me, however misplaced it may have been at times.

Thanks to Marat Gilfanov and Simon White for serving on my PhD committee. I'll be honest: those meetings weren't always fun, but they served a purpose – we got there in the end, folks!

A huge thanks to Simon White, Gerhard Börner, Martin Faessler, Jochen Weller, Dieter Braun, and Bernadett Weinzierl for agreeing to form my defence committee.

Thanks to Koki Kakiichi and Luca Graziani for countless fruitful discussions during this work. Thanks also to Akila Jeeson-Daniel for providing the CRASH pipeline, and in particular for all the help at the beginning of my PhD.

I would also like to thank Kathryn, Holger, and Christopher Kreckel for their support, encouragement, Feuerzangenbowle, and persistence in telling me to 'just finish it'. Christopher, I know that's what those sounds meant. Holger, many thanks for translating my abstract. I trust that you haven't pranked me. If you have, I take the thanks back.

Further thanks to Thomas Ertl for proofreading my German abstract. If you haven't told me that Holger pranked me, that's even worse.

Warmest thanks to all my MPA/MPE friends and colleagues who accompanied me through this journey and kept it fun. You know who you are.

A word of thanks also to all of the secretaries at MPA who made life in Germany as smooth as possible for this silly Ausländer.

Apologies to Rashid Sunyaev, whose quotation has been taken horrendously out of context to close my thesis. Sorry, but it was too good a chance to pass up.

I am thankful to all of my friends who constantly supported me in this endeavour but special mention must go to a couple of people very dear to my heart.

Frau Dr. KV, thank you for always, always supporting me, helping me through the tough times, and most of all for your companionship. It took me a while to catch up with you but in the end we both made it to the summit.

IC, thank you so much for the final push over the finish line. You must have felt like Sisyphus, but you achieved what he never could.

Finally, I would like to thank my family for their love and support. I am especially grateful to my parents for their unquestioning support despite my endless tilting at windmills.

It's a pity you're so slow.
— *Rashid Sunyaev*
Private communication to the author
MPA Coffee Corner, 14.01.2015