

Fluorescent rings in star-free dark matter haloes

Calvin Sykes^{1,2*}, Michele Fumagalli^{1,2}, Ryan Cooke²,
Tom Theuns¹, Alejandro Benítez-Llambay¹

¹*Institute for Computational Cosmology, Durham University, Durham DH1 3LE, UK*

²*Centre for Extragalactic Astronomy, Durham University, Durham DH1 3LE, UK*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Photoheating of the gas in low-mass dark matter (DM) haloes prevents baryons from cooling, leaving the haloes free of stars. Gas in these ‘dark’ haloes remains exposed to the ultraviolet background (UVB), and so is expected to emit via fluorescent recombination lines. We present a set of radiative transfer simulations, which model dark haloes as spherical gas clouds in hydrostatic equilibrium with a DM halo potential, and in thermal equilibrium with the UVB at redshift $z = 0$. We use these simulations to predict surface brightnesses in $H\alpha$, which we show to have a characteristic ring-shaped morphology for haloes in a narrow mass range between $\approx 10^{9.5}$ and $10^{9.6} M_\odot$. We explore how this emission depends on physical parameters such as the DM density profile and the UVB spectrum. We predict the abundance of fluorescent haloes on the sky, and discuss possible strategies for their detection. We demonstrate how detailed observations of fluorescent rings can be used to infer the properties of the haloes which host them, such as their density profiles and the mass-concentration relation, as well as to directly measure the UVB amplitude.

Key words: radiative transfer – galaxies: dwarf – cosmology: dark matter

1 INTRODUCTION

A fundamental prediction of the cold dark matter (CDM) paradigm is the existence of a large number of dwarf haloes (those with mass $\lesssim 10^{10} M_\odot$) resulting from the hierarchical nature of structure formation. To achieve consistency with observational constraints, it is necessary that only a small fraction of these low-mass haloes host luminous galaxies (e.g. Klypin et al. 1999; Moore et al. 1999). This requirement can be fulfilled by the effects of baryonic feedback processes, as demonstrated by hydrodynamical simulations (e.g. Babul & Rees 1992; Efstathiou 1992). Specifically, cosmic reionisation heats baryonic matter to $\sim 10^4$ K (see Meiksin 2009 for a review), inhibiting baryons from condensing into low-mass haloes to form stars. Furthermore, gas which does accrete can later be expelled by external ram-pressure stripping as the halo moves through the intergalactic medium (Benítez-Llambay et al. 2013).

Detection of these star-free ‘dark’ haloes would therefore provide convincing evidence in support of CDM, and a number of methods for discovering them have been proposed. The presence of dark haloes may be inferred gravitationally, either via their dynamical influence on other objects (Erkal & Belokurov 2015; Feldmann & Spolyar 2015), or via gravitational lensing (Vegetti et al. 2010; Hezaveh et al. 2016). Alternatively, one may exploit the fact that the haloes do contain a reservoir of almost pristine gas originating from the limited accretion that occurred before reionisation. Observing this gas either in absorption against

a luminous background source (Rees 1986), or directly in emission, could allow the presence of the dark haloes to be inferred.

Previously, Sternberg et al. (2002, hereafter S02) considered the emission properties of the gas bound to dark haloes, in the context of comparing them with observed high-velocity clouds (HVCs). This work concluded that for HVCs to be dark halo candidates, they must be ‘circumgalactic’ objects located relatively close to the Milky Way, and pressure-confined by a hot galactic corona. More recently, Benítez-Llambay et al. (2017, hereafter BL17) used the APOSTLE suite of Local Group hydrodynamical simulations (Sawala et al. 2016) to model the state of gas within dark haloes. This work identified two populations: ‘COSWEBs’, haloes which are found to lose almost all their residual gas after reionisation due to ram pressure stripping; and ‘RELHICs’, generally more isolated systems which are not subject to this additional effect, and thus have relatively constant baryonic content since reionisation.¹ RELHICs with sufficiently high halo masses were found to develop an approximately kiloparsec-sized neutral core, with the surrounding gas remaining ionised by the diffuse ultraviolet background (UVB).

Whereas BL17 considered the prospect of detecting H I 21 cm emission from these objects, in this work we follow S02 by investigating a second potential source of emission: the fluorescent $H\alpha$ line, in order to reassess its expected properties in light of the ‘extragalactic’ environment for RELHICs favoured by APOSTLE.

* E-mail: calvin.v.sykes@durham.ac.uk

¹ These acronyms stand for “COSmic WEB Stripped systems” and “Reionisation-Limited H I Clouds” respectively.

Fluorescent emission occurs when atoms ionised by the UVB recombine to produce neutral atoms in excited states, which subsequently undergo radiative cascades to return to the ground state. Although the $n = 2 \rightarrow 1$ Ly α transition is intrinsically brighter, we consider the H α transition ($n = 3 \rightarrow 2$) due to its non-resonant nature and rest-frame optical wavelength, which are especially advantageous for the ground detection of low-redshift sources.

In this work, we perform photoionisation simulations of the gas within dark haloes in the local Universe. We calculate radial profiles of the H α surface brightness, and investigate its dependence on physical properties of the haloes and the UVB. These calculations use an adapted version of the code described by [Cooke & Pettini \(2016, hereafter CP16\)](#), which follows the procedure outlined in [S02](#). These calculations are qualitatively similar to those performed by the CLOUDY spectral synthesis code ([Ferland et al. 2017](#)), but with two important additional features. Firstly, we model gas in hydrostatic equilibrium with a dark matter halo to obtain realistic gas density profiles. Secondly, we perform the calculations in (projected) spherical coordinates, in order to capture the radial dependence of optical depths within the cloud.

This paper is organised as follows: in Section 2, we describe our numerical method. We use this method in Section 3 to compute grids of models with a variety of input parameters to determine the resulting radial surface brightness profiles. We then combine our results with the APOSTLE simulations in Section 4 to predict the abundance of H α -fluorescent haloes on the sky, and the prospects for their detection. Lastly, in Section 5 we present a discussion of our results and their significance for constraining the structure and environment of a fluorescent halo, based on its observed H α emission. Throughout, we assume a set of cosmological parameters ($H_0 = 67.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_\Lambda = 0.685$, $\Omega_M = 0.315$, $\Omega_B = 0.0491$) consistent with *Planck* measurements ([Planck Collaboration et al. 2014](#)).

2 SIMULATION DESIGN

We use a modified version of the spherically-symmetric ionisation balance code described in [CP16](#). A full description of this code is given therein; here we will provide a brief summary and describe the modifications we have made.²

2.1 Code outline

We consider clouds of primordial gas, consisting only of hydrogen and helium with abundance ratio by number $n_{\text{He}}/n_{\text{H}} = 0.083$ (corresponding to a primordial helium mass fraction $Y_{\text{p}} = 0.24$). This gas is embedded within a spherically-symmetric Navarro-Frenk-White (NFW; [Navarro et al. 1996b](#)) dark matter halo, with radial density profile given by

$$\rho(r) = \frac{\rho_s}{x(1+x)^2}; \quad x \equiv \frac{r}{r_s}, \quad (1)$$

where r_s and ρ_s are characteristic length and density scales of the halo. The gas is assumed to be in hydrostatic equilibrium with a potential $\varphi(r)$, such that $dP(r) = -\rho_g(r)d\varphi(r)$, where $P(r)$ and $\rho_g(r)$ are the gas pressure and density profiles respectively. For simplicity, we neglect the self-gravity of the gas such that $\varphi(r)$

is defined solely by the dark matter halo, which is a reasonable assumption for the DM-dominated systems we consider.

Haloes are constructed by choosing a virial mass M_{200} , defined as the mass contained in a sphere with average density $200\rho_{\text{crit}}$, where ρ_{crit} is the critical density of the Universe. The radius of this sphere is the virial radius R_{200} , such that:

$$M_{200} = \frac{4\pi}{3} R_{200}^3 \times 200\rho_{\text{crit}}. \quad (2)$$

NFW haloes may be parameterised by a concentration parameter $c_{200} \equiv R_{200}/r_s$. In our fiducial models, we determine values of c_{200} as a function of M_{200} using the mass-concentration relation of [Ludlow et al. \(2016\)](#). We investigate alternative mass-concentration relations and density profiles in Section 3.2.2 in order to determine the dependence of our results on these assumptions. A total gas mass M_g is then assigned to the halo using the analytic model employed by [BL17](#), which determines the gas mass required for hydrostatic equilibrium as a function of M_{200} .

Given the assumed absence of local star formation, the gas is illuminated solely by photons from the diffuse UVB. We adopt the [Madau & Haardt \(2015, hereafter MH15\)](#) UVB in our fiducial model, although we explore the effect of varying the shape of this spectrum in Section 3.2.1. Under the assumption of spherical symmetry, the UVB irradiates the gas isotropically. However, the gas column density has an angular dependence which varies with position, meaning that the local background radiation intensity, and hence the local photoionisation rate, is a function of both depth within the cloud and direction of the incident radiation.

In addition to photoionisation, we consider primary and secondary collisional ionisation, ionisation resulting from helium recombination radiation, and charge transfer ionisation. We include radiative, dielectronic and charge transfer recombinations. Ionisation equilibrium for each species is enforced by equating the appropriate ionisation and recombination processes in each case. The rate coefficients assumed for these reactions are unmodified from those detailed in [CP16](#), with the exception of the rates for collisional ionisation, which we take from the CHIANTI atomic database ([Dere et al. 1997; Zanna et al. 2015](#)) in order to enable calculation of the rates at temperatures $\sim 10^3 \text{ K}$, which may be reached in the neutral core (the lower limit in the previously-used [Dere \(2007\)](#) rates was instead $\sim 10^4 \text{ K}$).

The version of the code described in [CP16](#) additionally assumed strict thermodynamic equilibrium, such that the temperature at each radial coordinate is that for which the computed heating and cooling rates are equal. The heating rate includes both photoionisation heating and secondary heating by primary photoelectrons, while the cooling rate incorporates contributions from collisional excitation/ionisation cooling, recombination cooling, Bremsstrahlung cooling and Compton cooling/heating. We take the formulae used to determine each of these contributions from [Cen \(1992\)](#). At low gas densities ($n_{\text{H}} \lesssim 10^{-4.8} \text{ cm}^{-3}$), cooling becomes inefficient and the gas temperature will instead be that which results in a heating timescale equal to the Hubble time. We incorporate this additional regime by switching from calculating the equilibrium temperature to using the tabulated temperature-density relations given in [Wiersma et al. \(2009\)](#) when the gas density falls below this transition point.

The form of the calculations follows that described by [S02](#): assuming the gas is initially isothermal and fully-ionised, we compute the gas pressure profile required to maintain hydrostatic equilibrium. A boundary condition is required to determine the pressure profile; in common with [BL17](#), we use the assumption that at sufficient dis-

² The code is made available here:

https://github.com/calvin-sykes/spherical_cloudy

tance from the centre of the halo ($\approx 100 R_{200}$), the density should approach the cosmic mean baryon (number) density $\bar{n} \approx 10^{-6.7} \text{ cm}^{-3}$ at $z = 0$.

From the computed pressure profile, we determine the radial density distribution, and hence the intensity of the radiation field within the cloud. We then solve for ionisation equilibrium to calculate the neutral fractions of each species as a function of radius. The temperature profile is then derived as described above, and used to recompute the pressure profile. This process is iterated until a convergence criterion is met, requiring that the relative change in the computed neutral fractions between successive iterations is no more than 0.1% at any radial coordinate.

Once a converged ionisation structure has been found, the $H\alpha$ surface brightness $\Sigma_{H\alpha}$ may be computed. At a radial position within the halo r , $H\alpha$ photons are produced at a rate given by the volume emissivity $\epsilon_{H\alpha} = n_{HII}(r)n_e(r)\alpha_{H\alpha}(r)$, where n_{HII} and n_e are the densities of ionised hydrogen and electrons respectively. The effective recombination coefficient $\alpha_{H\alpha}$ expresses the rate per unit electron and ion densities at which the $n = 3 \rightarrow 2$ transition occurs. We obtain values for this coefficient, as a function of temperature, from [Osterbrock & Ferland \(2006\)](#).

For a line of sight intersecting the halo with impact parameter b , the projected surface brightness is found as the line integral of the emissivity along the line of sight

$$\Sigma_{H\alpha}(b) = \frac{h\nu_{H\alpha}}{2\pi} \int_b^{R_{200}} \frac{r}{\sqrt{r^2 - b^2}} n_{HII}(r)n_e(r)\alpha_{H\alpha}(r) dr, \quad (3)$$

where h is Planck's constant and $\nu_{H\alpha} = 4.57 \times 10^{14} \text{ Hz}$ is the frequency of an $H\alpha$ photon.

2.2 Tests of numerical accuracy

To validate our code, we model a series of plane-parallel gas slabs, varying the volume and total column densities of the slab. For each slab, we initially specify a constant temperature $T_{\text{gas}} = 10^4 \text{ K}$. In this case the ionisation structure of the slab may be predicted analytically; in particular, the surface brightness is expected to approach a constant value, in a similar manner to the result for $\text{Ly}\alpha$ fluorescence reported by [Gould & Weinberg \(1996\)](#); see Appendix A). The use of a plane-parallel geometry also enables us to compare our results with those obtained from the spectral synthesis code CLOUDY ([Ferland et al. 2017](#)).

Fig. 1 demonstrates that in the high-column density limit, our code reproduces the analytic prediction well, with the only deviation occurring at high volume densities where collisional ionisation becomes significant. Note that our code makes a number of simplifications as compared to CLOUDY: we assume case B recombination, rather than self-consistently computing the diffuse radiation field produced by recombinations within the cloud. Also, our treatment of collisional processes is less sophisticated than the one in CLOUDY. We further explore the effects of these approximations in Appendix B.

In addition to verifying the ionisation balance portion of our code, we compare the physical properties of our haloes to the analytic model of [BL17](#). In Fig. 2, we compare radial density and temperature profiles from this model, which is calibrated for the [Haardt & Madau \(2012\)](#) UVB model, and from our code run with the same spectrum. We obtain good agreement for low-mass haloes and the outer regions of higher-mass haloes, where the gas remains fully ionised. For haloes which are sufficiently massive to develop an ionisation front, our code predicts higher densities and lower temperatures within the neutral region. This occurs because in the [BL17](#)




Figure 1. $H\alpha$ surface brightness as a function of total H column density for isothermal, H -only plane parallel models. Dotted curves are obtained from the code we describe here, solid ones from CLOUDY; each same-coloured pair corresponds to a particular constant gas density n_H . The horizontal dashed line indicates the analytic result for $H\alpha$ surface brightness in the photoionised, optically-thick limit.

model, a constant value of the mean mass per particle $\mu = 0.62$ is used, as appropriate for fully-ionised primordial gas. In this work, we explicitly calculate the ionisation state of the gas as a function of radius, and so can determine μ as

$$\mu(r) = \frac{\sum_i X_i(r) A_i(r)}{\sum_i X_i(r) + X_e(r)}, \quad (4)$$

where X_i is the number abundance of species i relative to H , A_i is the mass number of species i , and the e subscript denotes electrons. For fully neutral gas, Eq. 4 yields $\mu = 1.23$, resulting in higher pressures which require higher gas densities to maintain hydrostatic equilibrium. Since the cooling rate is an increasing function of density, the equilibrium temperature for neutral gas is consequently lower.

2.3 Star formation threshold

In the following section, we will show that more massive dark matter haloes will retain a more massive reservoir of gas; correspondingly, they will have larger neutral cores and higher peak values of $\Sigma_{H\alpha}$. However, these cores will be denser and colder, increasing the likelihood of star formation occurring. UV radiation emitted locally from young stars within the halo will readily exceed the intensity of the UVB, while their supernovae would expel gas from the halo. Hence, we require that the total amount of star formation is small, which imposes an upper limit on the halo masses for which fluorescence is expected.

Star formation primarily occurs within dense molecular clouds, and so the presence of H_2 may be used as a proxy for star formation. Ideally, we would self-consistently determine the molecular fraction within our code; however, the complexity of doing this puts it beyond the scope of this work. Consequently, we instead consider the formation of H_2 as a post-processing step. Analytic models, such as that described by [Krumholz et al. \(2009a\)](#), permit this by prescribing the surface molecular fraction $f_{H_2} \equiv \Sigma_{H_2}/(\Sigma_{H I} + \Sigma_{H_2})$ as a function of the column density of $H I$ and an estimate of the




Figure 2. Radial density (top) and temperature (bottom) profiles from [BL17](#) model (dashed lines) and our code (solid lines). The models agree well, except for the inner regions of high-mass haloes, where the transition to neutral gas changes the mean mass per particle μ , an effect which is not included in the [BL17](#) model.

H_2 -dissociating Lyman-Werner flux.³ We note that this model is primarily intended for disk galaxies with active star formation and higher metallicity than the gas we consider (which we assume to have primordial composition, although for the purposes of computing f_{H_2} we must assume a non-zero metallicity $Z = 0.001Z_\odot$). While the exact values of f_{H_2} we obtain may not be fully accurate, we expect the column density threshold for H_2 formation to be reliable. This threshold is predicted by the model to occur at $N_{\text{H I}} \sim 10^{23} \text{ cm}^{-2}$ which, when the absence of metals or dust is taken into consideration, is in good agreement with the observed paucity of damped Ly α (DLA) systems above a similar column density threshold, believed to be due to the onset of H_2 formation (e.g. [Schaye 2001](#); [Krumholz et al. 2009b](#)). Additionally, the upper bound on halo mass of $M_{200} \lesssim 10^{9.6} M_\odot$ implied by this threshold is consistent with the fraction of ‘dark’ haloes found in the APOSTLE hydrodynamical simulations, which exceeds 50% for halo masses below this value ([Benítez-Llambay et al. 2017](#)). Accordingly, excluding all haloes for which $N_{\text{H I}} \geq 10^{23} \text{ cm}^{-2}$ provides a threshold below which we may discount star formation.

An additional upper limit can arise from the constraint that the total gas fraction M_{gas}/M_{200} remains small, such that we are justified in neglecting the gas self-gravity. At the halo masses required for this constraint to become relevant, we find that column densities already exceed the limit for molecular gas formation defined above. However, the steep increase in gas density within the ionisation front raises the possibility of gas becoming locally self-

³ More advanced models for f_{H_2} exist, e.g. [Krumholz \(2013\)](#). These model the formation of H_2 at low densities more accurately, but give the same results in the high-density regime that is relevant to our purposes.




Figure 3. Radial profiles for hydrogen neutral fraction (top) and $\text{H}\alpha$ surface brightness (bottom) for haloes in our fiducial model, as a function of halo mass M_{200} . Note the sudden transition in the ionisation state of the gas at $M_{200} = 10^{9.5} M_\odot$. This corresponds to the formation of the limb-brightened ring in emission, shown here as a peak in the $\Sigma_{\text{H}\alpha}$ profile.

gravitating within the neutral core. Adopting the stricter condition $M_{\text{gas}}(\leq r)/M_{\text{DM}}(\leq r) \ll 1 \forall r < R_{200}$ to avoid this possibility would result in a small ($\lesssim 0.005$ dex) reduction in our upper threshold on halo mass. Nevertheless, in the absence of further work to model the chemical state of this gas, especially during early times when the greater mean density of the Universe would have yielded even steeper gas density profiles, our upper thresholds on halo mass should themselves be interpreted as upper limits.

3 PROPERTIES OF FLUORESCENT HALOES

3.1 Fiducial model

We first present results from our fiducial model, in which the UVB is given by the [MH15](#) spectrum and an NFW density profile is assumed for the dark matter halo, along with the [Ludlow et al. \(2016\)](#) mass-concentration relation and a redshift $z = 0$.

We compute models for haloes with masses $10^{9.38} M_\odot$ and above, and show the resulting hydrogen neutral fractions $X_{\text{H I}} \equiv n_{\text{H I}}/n_{\text{H}}$ and $\text{H}\alpha$ surface brightnesses $\Sigma_{\text{H}\alpha}$ in Fig. 3. At the lower end of this range, hydrostatic equilibrium is achieved for a bound gas mass within R_{200} of $M_{\text{gas}} \gtrsim 10^7 M_\odot$. With projected column densities $N_{\text{H I}} \lesssim 10^{15} \text{ cm}^{-2}$, this gas is optically thin to the UVB, and hence is highly ionised throughout. The gas remains rarefied, with peak densities $n_{\text{H}} \sim 10^{-3} \text{ cm}^{-3}$, and therefore $\Sigma_{\text{H}\alpha}$ is low at all impact parameters. As the halo mass increases, the central gas density also increases. Consequently, at small impact parameters the projected column density and surface brightness gradually increase, while remaining largely unchanged along lines of sight avoiding




Figure 4. Projected $H\alpha$ surface brightness for haloes in our fiducial model. The range of halo masses shown is chosen such that in the top-left panel, a neutral core first begins to form. As the halo mass increases, the neutral core enlarges and the limb-brightened peak in $\Sigma_{H\alpha}$ becomes more prominent until we reach the most massive halo allowed by our star formation criterion (bottom-right). Red dashed lines indicate 10% of the virial radius.

the high-density central region. For $M_{200} = 10^{9.47} M_{\odot}$, the central column density reaches 10^{18}cm^{-2} , corresponding to an optical depth at the Lyman limit $\tau_{912} \sim 1$. A sharp transition in the gas properties occurs as this threshold is exceeded: the gas becomes self-shielded from the UVB and develops a neutral core. By $M_{200} = 10^{9.49} M_{\odot}$ the central neutral fraction approaches 1.

Fig. 4 shows the projected surface brightness distributions for

haloes at this mass and above, clearly illustrating their characteristic ring-shaped surface brightness morphology. This occurs because the $H\alpha$ emissivity is sharply peaked at the ionisation front: at larger radii, the gas density is reduced, while inside the ionisation front few $H\text{II}$ recombinations occur. Hence, lines of sight tangent to the surface of the neutral core pass through a greater fraction of high-emissivity gas than those passing through the centre of the neutral

region. This ‘limb brightening’ effect produces surface brightness profiles which peak at the projected radius of the ionisation front. As the halo mass increases further, the core grows in size (see also Fig. 3), and rapidly becomes cooler and denser, leading to the formation of molecular H_2 , and our previously-discussed star formation criterion is exceeded for haloes above $10^{9.55} M_\odot$. For a halo of this mass, the surface brightness profile peaks at $r = 0.70$ kpc with a value of $\Sigma_{H\alpha, \max} = 6.58 \times 10^{-20} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$.

Ring-shaped $H\alpha$ surface brightness morphologies were also predicted by S02, with a comparable peak value of $\Sigma_{H\alpha, \max} \approx 4 \times 10^{-20} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$. However, this is attained for a UVB spectrum with hydrogen photoionisation rate $\approx 20\%$ lower than the MH15 spectrum (see Section 3.2.1 below). Additionally, S02 adopt an external bounding pressure, which acts in conjunction with the DM potential to confine the gas toward the centre of the halo, allowing a dense, self-shielding core to form at a much lower halo mass $M_{200} \approx 10^8 M_\odot$.

3.2 Model variations

3.2.1 UVB spectrum

Discrete sources of ionising radiation, such as a nearby AGN or bright galaxy, can produce local inhomogeneities in the UVB. Furthermore, uncertainties present in the modelling of the UVB result in variations in the intensity and spectral shape of the background computed by different models at a fixed redshift. Accordingly, we explore the sensitivity of our results to such changes in the UVB. The intensity of the UVB may be parameterised by the hydrogen photoionisation rate, which is given by

$$\Gamma = 4\pi \int_{\nu_{\text{th}}}^{\infty} \frac{J(\nu)\sigma(\nu)}{h\nu} d\nu, \quad (5)$$

where $J(\nu)$ is the mean intensity, $\sigma(\nu)$ is the cross-section for photoionisation of H^0 , and $\nu_{\text{th}} = 13.6 \text{ eV}/h$. In Fig. 5, we show stacked radial surface brightness profiles as a function of halo mass, with colour encoding the value of $\Sigma_{H\alpha}$. The panel in the centre of the upper horizontal row corresponds to our fiducial model, in which the photoionisation rate is $\Gamma_{\text{fid}} = 5.47 \times 10^{-14} \text{ s}^{-1}$. Panels to the left and right show results obtained with Γ respectively decreased and increased by a factor of two from its fiducial value. We find that increasing the photoionisation rate shifts the halo mass scale at which fluorescent rings are predicted upward, while the surface brightness of these rings also increases. Conversely, a reduced photoionisation rate allows dimmer rings to form at lower halo masses.

We interpret these trends as being due to the ability of a more intense UVB to penetrate a greater total column density of gas (i.e. N_H) before becoming attenuated. In the situation we consider, the gas volume density is a function of radius, which itself becomes steeper with increasing halo mass. This results in the ionisation front forming at a larger characteristic gas number density, which is only reached for more massive haloes. In turn, the greater density of the gas at the ionisation front boosts its emissivity, and hence the peak surface brightness. We further explore the reasons for surface brightness variations between models in Section 5.

To consider variations in the shape of the UVB spectrum, we adopt the parameterisation from Crighton et al. (2015), which introduces the UVB shape parameter α_{UV} . This parameter is used to modify the ionising spectrum as follows:

$$J_\nu(E) = \begin{cases} N_\Gamma \times H(E) & E \leq E_0 \\ N_\Gamma \times H(E) \times (E/E_0)^{\alpha_{\text{UV}}} & E_0 < E \leq E_1 \\ N_\Gamma \times H(E) \times (E_1/E_0)^{\alpha_{\text{UV}}} & E > E_1, \end{cases} \quad (6)$$

where $H(E) \equiv J_{\nu, \text{MH15}}(E)$, $E_0 = 1 \text{ Ryd}$ and $E_1 = 10 \text{ Ryd}$. We additionally scale these spectra by a factor $N_\Gamma \equiv \Gamma_{\text{fid}}/\Gamma$, such that the photoionisation rate is unchanged from its fiducial value, in order to examine the effect of modifying the spectral shape in isolation. The resulting spectra for $\alpha_{\text{UV}} = 0$ (corresponding to our fiducial model), $+1$, -1 and -2 are plotted in Fig. 6.

In Fig. 5, the central column shows $H\alpha$ surface brightness profiles as a function of M_{200} for haloes illuminated by the same four spectra. We find that the halo mass required to form a fluorescent ring is reduced for more negative values of α_{UV} , corresponding to softer spectra with a greater proportion of ionising photons having energies close to the Lyman limit (since the spectra have been renormalised to maintain a constant value of Γ). These photons are more likely to be absorbed, due to the $\sim \nu^{-3}$ dependence of the photoionisation cross-section, and so will on average be absorbed within a lower total gas column density. As described previously, the coupling between column and volume densities means that this corresponds to the ionisation front forming in lower-density gas, which maintains hydrostatic equilibrium with a shallower dark matter potential. In addition to the reduction in emissivity due to the lower gas density, an additional effect is caused by the reduced flux of photons with energies close to 24.6 eV , at which helium is singly ionised. This produces a larger He I fraction at the H ionisation front, further reducing the emissivity by lowering the electron density. Opposite arguments apply to the harder $\alpha_{\text{UV}} = 1$ spectrum, which produces larger peak surface brightnesses but inhibits neutral cores from forming until higher halo masses are reached.

3.2.2 Halo density profile

In our fiducial model, we have used an NFW density profile to model dark matter haloes, with the concentration parameter c_{200} determined using the Ludlow et al. (2016) relation, which is obtained from haloes in the Copernicus Complexio (CoCo) N-body simulations (Bose et al. 2016; Hellwing et al. 2016). The CoCo suite also includes runs in which a warm dark matter (WDM) cosmology is used. In contrast to CDM cosmologies, where c_{200} is a monotonically increasing function of M_{200} , the WDM mass-concentration relation turns over at a specific mass scale (Bose et al. 2017). $H\alpha$ surface brightness profiles obtained by using this relation are shown in the bottom-left corner of Fig. 5. At the halo masses at which we predict fluorescent ring emission, the WDM mass-concentration relation results in haloes which are approximately 50% less concentrated at fixed M_{200} than the equivalent halo following the L16 relation, and so have systematically lower central dark matter densities. Consequently, hydrostatic equilibrium is maintained with a shallower gas density profile, which produces lower central gas densities. Hence, the column density required to form a neutral core is first reached at higher halo mass, and once neutral cores form they extend to larger projected radii.

We also consider the possibility of varying the density profile assumed for the dark matter, by computing surface brightness profiles for haloes with ‘cored’ density profiles, where the central dark matter density approaches a finite value rather than diverging to infinity, as occurs for ‘cuspy’ profiles such as NFW. Observations of nearby dwarf galaxies, which imply central density gradients shallower than an NFW profile, have been interpreted either as a manifestation of the self-interacting nature of dark matter (e.g. Spergel & Steinhardt 2000), or as a result of baryonic interactions (Navarro et al. 1996a; Pontzen & Governato 2012, and references therein). The latter scenario generally requires supernova feedback from stars forming within the halo (e.g. Benítez-Llambay et al.




Figure 5. $\Sigma_{\text{H}\alpha}$ as a function of impact parameter b and M_{200} for different model variations. In each panel, the y-axis limits are set by the minimum mass required for the formation of a neutral core, and the maximum mass allowed by the star formation criterion. Results from our fiducial model are shown on the upper-central panel, with panels to the left and right illustrating the effect of varying the photoionisation rate, and panels above and below that of changing the spectral index α_{UV} . Panels in the bottom left and right corners show results for models with a different halo mass-concentration relation and halo density profile (note the wider x-axis limits in the latter).




Figure 6. UVB spectra in the range $1 \text{ Ryd} < E < 10 \text{ Ryd}$ for different values of the slope parameter α_{UV} , where $\alpha_{UV} = 0$ corresponds to the fiducial **MH15** spectrum. Vertical dashed lines indicate the H I, He I and He II ionisation edges.

2018), and so evidence of a core in a halo hosting a fluorescent ring would be suggestive of self-interacting dark matter (SIDM). We choose to model Burkert density profiles, for which (Burkert 1995):

$$\rho(r) = \frac{\rho_0 r_0^3}{(r + r_0)(r^2 + r_0^2)}, \quad (7)$$

where r_0 is the core radius, which we choose to be a function of M_{200} such that $R_{200}/r_0 = c_{200}$, in analogy to the NFW scale radius r_s . This choice results in core radii $r_0 \sim 2 - 3 \text{ kpc}$, somewhat larger than the typical size $r_0 \sim 1 \text{ kpc}$ predicted for SIDM haloes of the mass we consider, but within the range allowed by constraints on the self-interaction (see e.g. Tulin & Yu 2018). We caution that these assumptions will influence the resulting fluorescent emission; accordingly, the results we present here, shown in the bottom-right panel of Fig. 5, should be treated as illustrative only, rather than a specific prediction of a particular SIDM model. As with the Bose et al. (2017) mass-concentration relation, the reduction in central dark matter density caused by the dark matter core results in larger, more extended neutral gas cores developing at higher halo mass. However, the more extreme modification to the dark matter density profile results in a much more substantial change in the size of the neutral gas core. In both these models and those using the WDM mass-concentration relation, we find the magnitude of $\Sigma_{H\alpha, \text{max}}$ to be slightly increased with respect to our fiducial model. This occurs because of the larger physical size of the emitting region results in a longer path length over which $\epsilon_{H\alpha}$ is integrated to obtain the surface brightness, but is a relatively minor effect when compared to the impact of varying the properties of the UVB.

4 DETECTABILITY OF FLUORESCENT HALOES

We now make predictions for the abundance with which dark haloes exhibiting H α fluorescence appear on the sky, and hence consider the prospects for their detection, using results from the APOSTLE suite of high-resolution hydrodynamical simulations (Sawala et al. 2016). These simulations follow twelve galaxy groups selected to resemble the Local Group (LG) on the basis of their dynamical properties; we consider the five volumes simulated at the highest (‘L1’) resolution corresponding to a gas (DM) particle mass




Figure 7. Number density distribution of H α -fluorescent haloes in APOSTLE. The median values are shown by thick solid lines; shaded regions indicate 10th – 90th percentiles. From left to right, we show the number density of all RELHICs with H I column densities below our star formation threshold, those with surface brightnesses exceeding three different limiting values, and those with discernible ring-shaped surface brightness profiles. Here, $\bar{N}_{\text{H I}} \equiv N_{\text{H I}}/[\text{cm}^2]$ and $\bar{\Sigma}_{\text{H}\alpha} \equiv \Sigma_{\text{H}\alpha}/[\text{erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}]$. The right-hand axis indicates the corresponding number counts per APOSTLE volume.

$1.0 (5.0) \times 10^4 M_{\odot}$. From the $z = 0$ outputs of these simulations, we select all haloes with $M_{200} > 10^8 M_{\odot}$ located within 3.5 Mpc of the simulated LG’s barycentre, before discarding all ‘luminous’ haloes (i.e. those containing one or more star particles). The remaining ‘dark’ haloes are then divided into the two populations (COSWEBs and RELHICs) introduced in Section 1, where haloes with negligible gas content at $z = 0$ are placed into the former category and discounted.⁴ From this classification, we obtain a total of 263 RELHICs across the five volumes, with median halo mass $M_{200} = 10^{9.12} M_{\odot}$. We match each of these haloes with a corresponding calculation from Section 3, deselecting 23 objects due to their exceeding our star formation criterion (Section 2.3); for the remaining 240 we obtain a prediction of their peak H α surface brightness under our fiducial model. In Fig. 7, we show the resulting number density of all fluorescent objects, and of subsets with surface brightnesses exceeding chosen threshold values.

Fig. 7 indicates that high surface brightness fluorescence is rare, as are ring-shaped structures, with a median value of $3^{+2.6}_{-2.0}$ objects per volume, where the sub- and superscripts indicate the 10th and 90th percentile limits. This is not unexpected, due to the very narrow range in halo mass for which we expect bright ring fluorescence.⁵ Consequently, detecting fluorescent rings in the surroundings of the LG is impractical, since a prohibitively deep, all-sky H α survey would be required. However, searching for fluorescent haloes around nearby galaxies, albeit still challenging, becomes more feasible, since surface brightness remains independent of distance for cosmologically-nearby sources, but the angular extent of a volume equivalent to that simulated in APOSTLE decreases substantially. In

⁴ The group finder SUBFIND (Springel et al. 2001) is used to quantify the bound gas mass.

⁵ We caution that ‘bright’ is used here in a relative sense only; in absolute terms a surface brightness $\sim 10^{-20} \text{ erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}$ is still observationally challenging. Nevertheless, a combination of deep observations, high spectral resolution, and pixel-based stacking techniques has been successfully used to image surface brightnesses approaching this threshold (e.g. Fumagalli et al. 2017; Wisotzki et al. 2018).




Figure 8. Redshift evolution of the apparent surface brightness (solid curve; left scale) and the angular separation corresponding to a comoving kpc (dashed curve; right scale).

this situation, ultra-deep observations along single pointings have the potential to uncover one or more fluorescent rings.

Additionally, the requirement for an expensive optical search may be circumvented by exploiting the large H I column densities that will be associated with strongly-fluorescent systems. These imply that haloes hosting fluorescent rings should also be bright 21 cm H I emitters, and so existing H I sources may be used to provide a set of targets for subsequent H α observations, once deep (but affordable) broad-band imaging has been used to rule out associations with galaxies. Of particular interest is the catalogue of ultra-compact high velocity clouds (UCHVCs) detected in surveys like ALFALFA (Adams et al. 2013), which BL17 demonstrate to have H I emission properties which overlap at least partially with corresponding predictions for APOSTLE RELHICs. This is in contrast to the (non ‘ultra-compact’) HVCs, given the ‘extragalactic’ environment for RELHICs that APOSTLE favours. S02 found that the latter scenario would require HVCs to be underconcentrated and unrealistically spatially extended in order to reproduce their observed angular sizes.

Throughout the preceding discussion, we have considered fluorescence from haloes at redshift $z = 0$. At earlier times, the UVB has a greater intensity, resulting in an increased photoionisation rate Γ . As detailed in Section 3.2.1, this will correspond to greater intrinsic surface brightnesses. However, the observed surface brightness is dimmed by a factor $1/(1+z)^4$ relative to its intrinsic value. In Fig. 8, we illustrate that the quantity $\Gamma(z)/(1+z)^4$, which is directly proportional to the observed surface brightness, is maximised at $z \sim 0.2$. Also shown in Fig. 8 is the variation with redshift of the angular separation corresponding to a proper kpc, which decreases rapidly with redshift. In addition, halo concentration increases with redshift, which we expect to result in more compact fluorescent rings. Hence, resolving fluorescent objects at $z \gg 0$ is likely to be challenging. As we noted in Section 3.2.1, a boost in the UVB photoionisation rate can be provided even at $z \simeq 0$ by proximity to a highly-luminous source, such as an AGN (Cantalupo et al. 2005). To obtain detailed predictions of the expected surface brightness in this case, further modelling is required to determine the precise effects of the resulting change in spectral shape on the properties of fluorescent rings. Additionally, hydrodynamic modelling is needed to establish the conditions (e.g. separation from the source) needed to avoid stripping of gas from a candidate halo.




Figure 9. Summary of properties of ring fluorescence. The top panel shows the halo mass range for which ring fluorescence is predicted, for all models considered in Section 3. Our fiducial model is on the far left, followed by the following variations: three different values of the UVB spectral slope α_{UV} , UVB photoionisation rates Γ a factor of 2 below and above the fiducial value, the WDM (Bose et al. 2017) mass-concentration relation, and the cored dark matter density profile given by Burkert (1995). Subsequent panels show, for the same set of haloes in each model, values of the peak surface brightness, the radius at which this peak is reached, and the full width at half maximum of the radial surface brightness.

5 PHYSICS GOVERNING THE PROPERTIES OF FLUORESCENT HALOES

In Section 3, we presented surface brightness profiles resulting from varying input parameters to our photoionisation code. In the event of a successful detection of an object with ring-shaped fluorescent H α emission, we would like to determine the physical properties of the halo hosting the ring along with the incident radiation field. Hence, in this section we consider the causes of the trends displayed in our results in more detail, and discuss how they may be used to

infer the properties of a fluorescent ring and its environment from its appearance.

We summarise our results in Fig. 9, which shows the range of masses for which ring fluorescence is predicted in the different model variations we consider, along with the corresponding ranges in the peak value surface brightness, the projected radius at which this peak occurs, and the full width at half maximum of the surface brightness profile. Depending on its properties, purely detecting a fluorescent ring may provide constraints on its environment. For example, detection of a ring with $\Sigma_{\text{H}\alpha, \text{max}} \gtrsim 10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ would strongly suggest a UVB intensity enhanced with respect to the MH15 spectrum. From the bottom two panels of Fig. 9, we also see that with ‘non-standard’ dark matter (i.e. WDM or cored density profiles), fluorescent rings are expected to reach their peak surface brightness at larger projected radii than in the fiducial cuspy CDM case. However, the surface brightness profiles for the cored haloes we model are also broader, meaning that combined measurements of r_{max} and Δr could potentially be used to distinguish WDM from CDM, and cored from cuspy density profiles. We again note that these properties depend on the exact parameterisation of the density profile. For example, choosing a smaller core radius, corresponding to a more concentrated halo, would translate to reductions in r_{max} and Δr .

We now turn to considering why different input parameters affect the properties of fluorescent rings in the ways shown in Fig. 9. An important trait universal to all our models is that, despite significant variation in the sizes and peak surface brightnesses of the fluorescent rings they predict, the column density of the gas at r_{max} has a characteristic value $N_{\text{HI}} = N_{\text{crit}} \approx 10^{19} \text{ cm}^{-2}$. We illustrate this in Fig 10, where we also show that this corresponds to the column density at which the neutral fraction is $X_{\text{HI}} = 0.5$. While the association of the ionisation front with a column density of this magnitude is a standard result of radiative transfer calculations, it is nevertheless significant that we obtain self-similar X_{HI} and $\Sigma_{\text{H}\alpha}$ profiles despite the substantial variation in the shapes of physical density profiles between our different models.

One notable exception to this self-similarity is the $\alpha_{\text{UV}} = 1$ model, in which the surface brightness profile dips after the peak at N_{crit} , but then rises again toward higher N_{HI} . This occurs due to the increased flux of high-energy photons in the harder UVB spectrum assumed by this model, which cause ionisations at a larger typical value of column density (corresponding to smaller radii) due to their reduced photoionisation cross-section. Additionally, the energetic photoelectrons they liberate can go on to ionise further atoms via secondary collisional processes. Hence, the residual ionised fraction X_{HII} within the predominantly-neutral core is increased. In combination with the increase in n_{H} toward smaller radii, this produces the observed doubly-peaked surface brightness profile.

With the ‘fixed point’ at N_{crit} identified, the changes in ring brightness and radius between models may be described in terms of changing $N_{\text{HI}} - n_{\text{H}}$ and $N_{\text{HI}} - r$ relations respectively. The former follows since the peak emissivity is given by $\varepsilon_{\text{H}\alpha, \text{max}} = X_{\text{HII}} n_{\text{H}} n_e \alpha_{\text{H}\alpha}(T)|_{r_{\text{max}}}$. Since r_{max} corresponds to the ionisation front where $X_{\text{HII}} = 0.5$, and so $n_e \approx n_{\text{H}}$ while also $T \approx 10^4 \text{ K}$, the value of $\varepsilon_{\text{H}\alpha, \text{max}}$ is set by $n_{\text{H}}(r_{\text{max}})$. Hence, a change in the local density of gas where N_{HI} reaches N_{crit} will produce a change in the peak surface brightness. By the same token, a change in r_{max} implies a change in the position of the ionisation front, and therefore a change in the radius at which $N_{\text{HI}} = N_{\text{crit}}$. With these ‘scaling relations’ identified, it is then possible to take the properties of an observed fluorescent ring and work backwards to determine the physical properties of the host halo. By comparing these inferred val-




Figure 10. Normalised H α surface brightness profiles for the most massive halo satisfying our star formation criterion in each model we consider (black), shown with corresponding H neutral fractions X_{HI} (dashed blue).




Figure 11. Surface brightness as a function of H I column density for different photoionisation rates. In the inset panel, we show that the three $N_{\text{HI}} - \Sigma_{\text{H}\alpha}$ relations are self-similar when the surface brightness is scaled by the ratio $f_{\Gamma} = \Gamma / \Gamma_{\text{fid}}$.

ues with those predicted by our simulations, quantitative constraints may then be placed on which models are applicable to observed objects, as we discuss next.

We first consider trends in the maximum surface brightness $\Sigma_{\text{H}\alpha, \text{max}}$; as described previously, this is primarily set by the gas number density at the ionisation front. The physical size of the neutral region provides an additional influence, by altering the path length of sightlines passing through gas with emissivity close to maximal. However, the most extreme change in ring geometry with respect to our fiducial model occurs when a cored density profile is assumed, but the resulting values of $\Sigma_{\text{H}\alpha, \text{max}}$ differ by a much smaller margin. Hence, we infer that the ring brightness is set primarily by properties of the UVB. In particular, the models in which we consider a doubling or halving of Γ predict changes in surface brightness by the same amount. This is shown explicitly by plotting $\Sigma_{\text{H}\alpha, \text{max}}$ as a function of N_{HI} (Fig. 11), where we demonstrate in the inset panel that $\Sigma_{\text{H}\alpha} / f_{\Gamma}$ curves, where f_{Γ} is the multiplicative factor applied to the photoionisation rate, are self-similar. Hence, assuming the other properties of the UVB (i.e. α_{UV}) are similar to the MH15 spectrum, measuring the peak surface brightness of an




Figure 12. H I column density profiles (blue) for models with different dark matter density profiles, with corresponding H α surface brightness profiles (red). Profiles are shown as a function of impact parameter b , both in physical units (top panel) and normalised by r_{\max} (bottom panel). For each model, profiles are shown for the most massive halo which satisfies our star formation criterion.

observed ring and inferring a value for f_{Γ} would allow a direct determination of the UVB photoionisation rate (see also Fumagalli et al. 2017).

In contrast, the projected radius r_{\max} at which fluorescent rings are maximally bright is essentially independent of the properties of the UVB, changing only when the mass-concentration relation or density profile of the underlying dark matter halo is altered. As the column density at r_{\max} remains equal to N_{crit} , it is the $N_{\text{HI}}-r$ relation that is changing in this case, as we illustrate in the upper panel of Fig. 12. Identifying this trend reinforces our previous observation that measuring r_{\max} in combination with Δr permits cuspy CDM, cored DM (e.g. SIDM) and WDM to be differentiated. However, doing so requires knowledge of the distance to the ring in order to determine physical values of r_{\max} and Δr . Without this, the H α and H I profiles are only known as a function of b/r_{\max} , a situation which results in them becoming highly self-similar. We illustrate this in the lower panel of Fig. 12. However, if we instead assume a fixed family of models and vary the halo mass, we find a systematic trend in the shape of the H I profile, which becomes steeper for $N_{\text{HI}} > N_{\text{crit}}$ with increasing halo mass (Fig. 13). Changes also occur in the H α profiles, which become more sharply peaked at $b = r_{\max}$ for higher halo masses. Using complementary observations of H α and H I in this way may permit better constraints to be made, although the relatively subtle change in e.g. r_{\max} produced by the WDM models, which correspond to a significant 50% reduction in halo concentration c_{200} , imply that obtaining the necessary accuracy from observations would be challenging.




Figure 13. H I column density profiles (solid lines), with corresponding H α surface brightness profiles (red). Profiles are shown as a function of b/r_{\max} for haloes of different mass in our fiducial model.

6 SUMMARY AND CONCLUSIONS

We have reported the results from radiative transfer calculations which allow the simulation of gas bound to a gravitational potential such as that due to a dark matter halo and illuminated by an isotropic ionising background. By comparison with the well-known CLOUDY software, as well as previous analytic modelling, we have established that for the intended usage, despite making a number of simplifications, the code we use is robust and its results accurate. We have used this code to model the gas clouds bound to low-mass dark matter haloes, and obtained surface brightness profiles for the H α line, which is produced via fluorescent emission due to the ionising background. Our key findings are:

- (i) We confirm the results of previous studies, finding that haloes which develop a self-shielded neutral core possess surface brightness profiles with a characteristic ring-shaped morphology.
- (ii) This occurs for only a narrow range of halo masses, bounded from below by the requirement of a sufficiently deep potential, and from above by the requirement that star formation remains suppressed.
- (iii) Using numerical simulations, we predict a small population of $3^{+2.6}_{-2.0}$ (10-90% c. i.) fluorescent rings in a spherical 3 Mpc volume centred on the Local Group. Detecting these objects via a blind H α survey would be impractical, but pre-selection techniques, e.g. via 21 cm H I imaging, can provide catalogues of promising target sources, particularly beyond the LG where the sky coverage required to map a region of equivalent volume is reduced.
- (iv) We quantify how variations of the intensity (i.e. photoionisation rate) and spectral shape of the ionising background map directly to changes in the intensity of the fluorescent emission. Modifications to the dark matter potential instead alter the apparent size and shape of the fluorescent ring.
- (v) We show that reverse-engineering these relations would enable a direct determination of the UVB photoionisation rate from the intrinsic brightness of an observed ring; equally, constraints on the nature and distribution of dark matter may be inferred from the ring's geometry.

The existence of a large number of low-mass dark matter haloes which do not host a luminous galaxy is an unavoidable prediction of the CDM paradigm. While challenging, observing the distinctive ring-shaped fluorescent H α emission from the gas bound to these

haloes would provide a probe of this cosmological model on an as-yet untested scale. In the longer term, a detailed analysis of this emission would unlock further insights, both into the nature of dark matter and the properties of the extragalactic ionising background.

ACKNOWLEDGEMENTS

CS acknowledges support by a Science and Technology Facilities Council studentship [grant number ST/R504725/1]. MF & TT acknowledge support by the Science and Technology Facilities Council [grant number ST/P000541/1]. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 757535). During this work, RJC was supported by a Royal Society University Research Fellowship. We thank the members of the APOSTLE project for providing access to data. This work has benefited from the public Python packages NUMPY, SCIPY, ASTROPY and MATPLOTLIB. This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, and STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure.

REFERENCES

- Adams E. A. K., Giovanelli R., Haynes M. P., 2013, *ApJ*, 768, 77
 Babul A., Rees M. J., 1992, *MNRAS*, 255, 346
 Benítez-Llambay A., Navarro J. F., Abadi M. G., Gottlöber S., Yepes G., Yehuda Hoffman Steinmetz M., 2013, *ApJL*, 763, L41
 Benítez-Llambay A., et al., 2017, *MNRAS*, 465, 3913
 Benítez-Llambay A., Frenk C. S., Ludlow A. D., Navarro J. F., 2018, arXiv:1810.04186 [astro-ph]
 Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, *MNRAS*, 455, 318
 Bose S., et al., 2017, *MNRAS*, 464, 4520
 Burkert A., 1995, *ApJ*, 447, L25
 Cantalupo S., Porciani C., Lilly S. J., Miniati F., 2005, *ApJ*, 628, 61
 Cen R., 1992, *ApJS*, 78, 341
 Cooke R., Pettini M., 2016, *MNRAS*, 455, 1512
 Crighton N. H. M., Hennawi J. F., Simcoe R. A., Cooksey K. L., Murphy M. T., Fumagalli M., Prochaska J. X., Shanks T., 2015, *MNRAS*, 446, 18
 Dere K. P., 2007, *A&A*, 466, 771
 Dere K. P., Landi E., Mason H. E., Fossi B. C. M., Young P. R., 1997, *A&AS*, 125, 149
 Efsthathiou G., 1992, *MNRAS*, 256, 43P
 Erkal D., Belokurov V., 2015, *MNRAS*, 454, 3542
 Feldmann R., Spolyar D., 2015, *MNRAS*, 446, 1000
 Ferland G. J., et al., 2017, arXiv:1705.10877 [astro-ph]
 Fumagalli M., Haardt F., Theuns T., Morris S. L., Cantalupo S., Madau P., Fossati M., 2017, *MNRAS*, 467, 4802
 Gould A., Weinberg D. H., 1996, *ApJ*, 468, 462
 Haardt F., Madau P., 2012, *ApJ*, 746, 125
 Hellwing W. A., Frenk C. S., Cautun M., Bose S., Helly J., Jenkins A., Sawala T., Cytowski M., 2016, *MNRAS*, 457, 3492
 Hezaveh Y. D., et al., 2016, *ApJ*, 823, 37
 Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, *ApJ*, 522, 82
 Krumholz M. R., 2013, *MNRAS*, 436, 2747
 Krumholz M. R., McKee C. F., Tumlinson J., 2009a, *ApJ*, 699, 850
 Krumholz M. R., Ellison S. L., Prochaska J. X., Tumlinson J., 2009b, *ApJ*, 701, L12

- Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214
 Madau P., Haardt F., 2015, *ApJL*, 813, L8
 Meiksin A. A., 2009, *Rev. Mod. Phys.*, 81, 1405
 Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *ApJ*, 524, L19
 Navarro J. F., Eke V. R., Frenk C. S., 1996a, *MNRAS*, 283, L72
 Navarro J. F., Frenk C. S., White S. D. M., 1996b, *ApJ*, 462, 563
 Osterbrock D. E., Ferland G. J., 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*, 2 edn. University Science Books, Sausalito, CA
 Planck Collaboration et al., 2014, *A&A*, 571, A16
 Pontzen A., Governato F., 2012, *MNRAS*, 421, 3464
 Rees M. J., 1986, *MNRAS*, 218, 25P
 Sawala T., et al., 2016, *MNRAS*, 457, 1931
 Schaye J., 2001, *ApJ*, 562, L95
 Spergel D. N., Steinhardt P. J., 2000, *PRL*, 84, 3760
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
 Sternberg A., McKee C. F., Wolfire M. G., 2002, *ApJS*, 143, 419
 Tulin S., Yu H.-B., 2018, *Phys. Rep.*, 730, 1
 Vegetti S., Koopmans L. V. E., Bolton A., Treu T., Gavazzi R., 2010, *MNRAS*, 408, 1969
 Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, 393, 99
 Wisotzki L., et al., 2018, *Nature*, 562, 229
 Zanna G. D., Dere K. P., Young P. R., Landi E., Mason H. E., 2015, *A&A*, 582, A56

APPENDIX A: ANALYTIC ESTIMATE OF $H\alpha$ SURFACE BRIGHTNESS IN OPTICALLY-THICK LIMIT

In this Appendix we present a derivation of an analytic result for the $H\alpha$ surface brightness in the optically-thick limit (see Fig. 1), following the method discussed by Gould & Weinberg (1996). For a constant-density, plane-parallel slab of pure hydrogen composition, photoionisation equilibrium gives

$$\alpha(1 - X_{H1})^2 n_H^2 = \Gamma X_{H1} n_H \quad (A1)$$

where $X_{H1} = n_{H1}/n_H$, and α is the recombination coefficient for H^+ . At a depth ℓ inside the cloud, the photoionisation rate is given by its value at the illuminated face of the cloud Γ_0 , attenuated by the total optical depth $\tau(\ell)$, i.e.

$$\Gamma(\ell) = \Gamma_0 \exp[-\tau(\ell)] \quad (A2)$$

The Ly α surface brightness in units of photons $\text{cm}^{-2} \text{s}^{-1}$ is given by

$$\Sigma_{Ly\alpha} = f_{Ly\alpha} \int_0^\ell \alpha(1 - X_{H1})^2 n_H^2 d\ell' \quad (A3)$$

where $f_{Ly\alpha} \sim 0.6$ is the fraction of recombinations that lead to emission of a Ly α photon (Osterbrock & Ferland 2006), and α is the recombination coefficient. Substituting Eqs. A1 and A2 into the above then gives

$$\Sigma_{Ly\alpha} = f_{Ly\alpha} \int_0^\ell X_{H1} n_H \Gamma_0 \exp(-\tau) d\ell' \quad (A4)$$

We next define a photon number-weighted average of the photoionisation cross section $\langle\sigma\rangle$ as

$$\langle\sigma\rangle \equiv \frac{\int_{\nu_{th}}^\infty \sigma(\nu) \phi(\nu) d\nu}{\int_{\nu_{th}}^\infty \phi(\nu) d\nu} \quad (A5)$$

where $\phi(\nu) \equiv J(\nu)/h\nu$ gives the number of photons per second with frequency ν . With this definition, the total optical depth becomes

$$\tau \equiv \int_0^\ell X_{H1} \langle\sigma\rangle n_H d\ell' \quad \text{and so} \quad d\tau = X_{H1} \langle\sigma\rangle n_H d\ell' \quad (A6)$$

Inserting these results into Eq. A4 gives

$$\Sigma_{\text{Ly}\alpha} = f_{\text{Ly}\alpha} \frac{\Gamma_0}{\langle \sigma \rangle} \int_0^{\tau'=\tau} \exp(-\tau') d\tau' \quad (\text{A7})$$

In the optically thick limit, $\tau \rightarrow \infty$ and so $\Sigma_{\text{Ly}\alpha} = f_{\text{Ly}\alpha} \Gamma_0 / \langle \sigma \rangle$. Rewriting this result using Eq. A5 and the definition of Γ_0 previously given in Eq. 5 yields

$$\Sigma_{\text{Ly}\alpha} = 4\pi f_{\text{Ly}\alpha} \int_{\nu_{\text{th}}}^{\infty} \phi(\nu) d\nu \quad (\text{A8})$$

which is the Gould & Weinberg (1996) result: in the optically thick limit, all photons are absorbed and therefore the Ly α surface brightness tends to a constant fraction of the illuminating flux set by $f_{\text{Ly}\alpha}$. To obtain an equivalent result for H α , we multiply by the case B flux ratio $F_{\text{Ly}\alpha}/F_{\text{H}\alpha} = 8.5$ for $T = 10^4$ K (Osterbrock & Ferland 2006). The resulting prediction, $\Sigma_{\text{H}\alpha, \text{GW}} = 2.15 \times 10^{-20} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$, provides a useful check of our code.

APPENDIX B: ADDITIONAL TESTS OF IONISATION BALANCE

In this Appendix, we provide details of a number of additional comparisons with CLOUDY that were performed to verify the operation of the code described in Section 2.1.

Firstly, we consider the effects of our use of the case B recombination assumption, by comparing the hydrogen neutral fractions $X_{\text{H I}} \equiv n_{\text{H I}}/n_{\text{H}}$ that we obtain with corresponding predictions from CLOUDY, both using its default self-consistent treatment of recombination radiation, and overriding this behaviour to also assume the case B limit. The upper panel of Fig. B1 indicates that this approximation results in a slight underestimate of the neutral fraction at column densities $N_{\text{H}} < 10^{18} \text{ cm}^{-2}$.

The volume emissivities plotted in the bottom panel of Fig. B1 appear to show that this optically thin gas contributes the majority of the H α emission; however, this is an effect of the artificial constant-density assumption used in the slab models. For the physically-motivated models we consider, the gas density profile is dictated by hydrostatic equilibrium so that n_{H} rapidly decreases with increasing radius and the H α emissivity peaks near the ionisation front. This is located at $N_{\text{H}} \gtrsim 10^{18} \text{ cm}^{-2}$, where the case B limit is applicable; accordingly, we do not expect the absence of self-consistent diffuse radiation in our code to significantly affect our results. Fig. B1 also indicates emissivities below the CLOUDY values at high column and volume densities. This is likely caused by the exclusion of collisional excitation processes in our code, as disabling these processes in CLOUDY produces the dashed curves in Fig. B1, with which our results are a better match. Since the emissivity of the predominantly neutral gas found at high N_{H} is at least two orders of magnitude lower than that of gas near the ionisation front, we expect the net impact of this discrepancy on the total surface brightness to be minimal.

Finally, in Fig. B2 we show results for a set of plane parallel models run at a higher gas temperature of 20,000 K. At this temperature, collisional ionisation dominates throughout the slab, resulting in neutral fractions $\lesssim 0.05$. There is a systematic offset of $\sim 10\%$ in the neutral fractions predicted by our code and by CLOUDY, which is due to a difference of similar magnitude in the collisional ionisation rate coefficients assumed by the two codes. Rescaling the CHIANTI rates we employ to match those reported by CLOUDY produces the




Figure B1. Upper panel: H I neutral fraction as a function of total H column density for isothermal, H-only plane parallel models. Results from our code are shown by dotted lines; over much of the plotted N_{H} range, these are indistinguishable from the solid lines, which show CLOUDY results with the Case B assumption switched on. Relaxing this assumption results in the curves shown by dashed lines. Lower panel: H α volume emissivity for the same models. Our models are again shown with dotted lines, while solid and dashed lines correspond to Case B CLOUDY models with collisional excitation processes enabled and disabled respectively.




Figure B2. H I neutral fraction as a function of total H column density for isothermal, H-only plane parallel models with $T_{\text{gas}} = 20,000$ K. Solid and dotted curves correspond to output generated by CLOUDY and our code respectively, while the dot-dashed curves show our predictions if the collisional ionisation rates we use are rescaled to match those in CLOUDY.

dot-dashed curves in Fig. B2, reducing the offset in the neutral fractions to below 1%. Hence, excluding differences in the input data used, we find excellent agreement between our code and CLOUDY in the collisionally-ionised limit.

In summary, while the simplicity of the code we present here as compared to CLOUDY results in some differences in the predicted values of $X_{\text{H I}}$ and $\varepsilon_{\text{H}\alpha}$, these occur for regimes that are either absent in our physically-motivated models, or contribute little to the observed surface brightness. Therefore, our code's predictions appear to be robust in its intended domain of applicability.