

Radiative transfer calculations of the diffuse ionized gas in disc galaxies with cosmic ray feedback

Bert Vandenbroucke,¹★ Kenneth Wood,¹ Philipp Girichidis,² Alex S. Hill^{3,4,5}
and Thomas Peters⁶

¹*SUPA, School of Physics and Astronomy, University of St Andrews, North Haugh, St Andrews KY16 9SS, UK*

²*Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany*

³*Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia V6T 1Z1, Canada*

⁴*Space Science Institute, Boulder, CO, USA*

⁵*National Research Council Canada, Herzberg Program in Astronomy and Astrophysics, Dominion Radio Astrophysical Observatory, PO Box 248, Penticton, British Columbia V2A 6J9, Canada*

⁶*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany*

Accepted 2018 February 16. Received 2018 February 6; in original form 2017 November 30

ABSTRACT

The large vertical scale heights of the diffuse ionized gas (DIG) in disc galaxies are challenging to model, as hydrodynamical models including only thermal feedback seem to be unable to support gas at these heights. In this paper, we use a three-dimensional Monte Carlo radiation transfer code to post-process disc simulations of the Simulating the Life-Cycle of Molecular Clouds project that include feedback by cosmic rays. We show that the more extended discs in simulations including cosmic ray feedback naturally lead to larger scale heights for the DIG which are more in line with observed scale heights. We also show that including a fiducial cosmic ray heating term in our model can help to increase the temperature as a function of disc scale height, but fails to reproduce observed DIG nitrogen and sulphur forbidden line intensities. We show that, to reproduce these line emissions, we require a heating mechanism that affects gas over a larger density range than is achieved by cosmic ray heating, which can be achieved by fine tuning the total luminosity of ionizing sources to get an appropriate ionizing spectrum as a function of scale height. This result sheds a new light on the relation between forbidden line emissions and temperature profiles for realistic DIG gas distributions.

Key words: radiative transfer – methods: numerical – cosmic rays – galaxies: ISM – galaxies: structure.

1 INTRODUCTION

Observations of the extended interstellar medium (ISM) in the Milky Way and other galaxies have shown the existence of a diffuse ionized component of the ISM with scale heights of more than a kpc above the star-forming disc (Haffner et al. 2009). Due to its temperatures of $\sim 10^4$ K, this component is usually referred to as the warm ionized medium. The optical emission line ratios and inferred temperature ($\sim 10^4$ K) of the gas imply that it is photoionized, so that it is also referred to as the diffuse ionized gas (DIG). The most obvious possible source of ionizing radiation is the O and B stars in the galactic disc itself (Reynolds 1990). Previous studies have shown that this radiation is sufficient to ionize the DIG at high altitudes, provided that the gas density distribution is not smooth (Wood et al. 2005, 2010; Barnes et al. 2014). A fractal or lognormal density

distribution is implied by observations and arises naturally from turbulence in the DIG (Elmegreen 1997; Berkhuijsen & Fletcher 2008; Hill et al. 2008). The importance of photoionization indicates that the gas is heated by UV radiation.

To reproduce the observed scale height of the DIG in 3D models of the galactic disc, a detailed model of the structure of the disc is needed. Barnes et al. (2014) used the models of Hill et al. (2012), which include thermal stellar feedback in a full magnetohydrodynamical (MHD) treatment of the ISM, and showed that they fail to reproduce the required densities at high altitude that could explain the observed scale height of the DIG. Furthermore, they also show the need for an additional source of heating that would explain the observed intensities of the [N II] 6584 Å and [S II] 6725 Å forbidden lines as a function of the H α intensity. These results indicate that a simple model including only thermal supernova feedback is unable to support an extended DIG and unable to heat the gas to high enough temperatures. This means we either need a more complete model of stellar feedback that also includes radiative feedback and

* E-mail: bv7@st-andrews.ac.uk

stellar winds that can change the structure of the ISM (Gatto et al. 2017; Peters et al. 2017), an additional form of heating that affects gas at low densities (Reynolds, Haffner & Tufte 1999), or both.

To address these issues, we repeat the analysis of Barnes et al. (2014) for a new sample of MHD simulations that include cosmic ray feedback (Girichidis et al. 2016b). Unlike thermal feedback, cosmic ray feedback does not couple directly to the local gas but can diffuse to higher altitudes through the magnetic field. This means it can heat the gas more efficiently and support a thicker disc. It also provides a non-local heating term that could provide the additional heating needed to explain observed line ratios.

The structure of this paper is as follows. In Section 2, we summarize the new ionization code we used, and give some more details about the simulation models. We then show that our models are converged (Section 3.1), and illustrate the effect of changing various model parameters (Section 3.2). We then move on to a full description of the time evolution of the models (Section 3.3), and of the emission line maps that we produce (Section 3.4). We end with our conclusions.

2 METHOD

2.1 Code

We use the new Monte Carlo radiative transfer code `CMAcIOnIZE` (Vandenbroucke & Wood 2018),¹ which is essentially a rewritten version of the ionization code of Wood, Mathis & Ercolano (2004). The code employs a basic model whereby ionizing radiation from one or more sources is propagated through a discrete density grid, and only absorption and re-emission by hydrogen and helium are taken into account. As Wood et al. (2010) showed, the effect of dust absorption is minimal and can hence be neglected. The ionizing part of the source spectrum is sampled using a number of discrete photon packets, which are emitted isotropically from the source location(s). Each photon packet is then followed while it traverses the simulation volume, until a randomly sampled optical depth is reached. At this point, the photon packet is absorbed, and is re-emitted at a randomly sampled frequency. We do not adopt the so-called on the spot approximation, whereby re-emitted photons are absorbed locally, and instead treat re-emitted photons in the same way as source photons. Photon packets whose frequency drops below the ionization threshold, or that leave the simulation box without scattering (through a non-periodic boundary), are removed from the system.

For each cell in the simulation volume, we keep track of the path-lengths traversed by photon packets that pass through it. After a sufficient number of photon packets have been evolved in this way, the path-length counters in each cell can be used to obtain a good approximation of the mean ionizing intensities in that cell. We then use these values to calculate the local ionization equilibrium. This will likely change the properties of the cell, so that we need to repeat the whole process until convergence is reached.

As the ionization equilibrium also depends on the temperature, we also need to keep track of heating and cooling processes that might affect the thermal equilibrium of the ISM. We consider heating caused by direct UV absorption by hydrogen and helium, heating due to indirect, ‘on the spot’ absorption of $\text{He I Ly}\alpha$, and heating caused by absorption by polycyclic aromatic hydrocarbons. As an optional extra heating term, we also consider heating

by cosmic rays, following Wiener, Zweibel & Oh (2013). Apart from cooling by recombination of ionized hydrogen and helium, and cooling by free–free emission (bremsstrahlung), we also consider a number of metals to obtain cooling rates: C, N, O, Ne, and S. To this end, intensity counters are also stored for different ions of these metals: C^+ , C^{++} , N^0 , N^+ , N^{++} , O^0 , O^+ , O^{++} , Ne^+ , Ne^{++} , S^+ , S^{++} , and S^{+++} .

Since the full combined ionization and temperature algorithm is more computationally expensive than a more approximate approach whereby the temperature is kept fixed, we first explore the available parameter space using a fixed temperature of 8000 K (which corresponds to the average equilibrium temperature of the observed DIG), and only use the combined algorithm to obtain line intensities for the models with the most realistic parameter values.

In runs that use the full version of the code, we can use the resulting ionization structure of the various coolants and the equilibrium temperature to compute forbidden line emission (Wood et al. 2004; Vandenbroucke & Wood 2018). This allows us to produce line emission maps and line ratios that can be directly compared with observations.

2.2 ISM density field

We use simulated density fields from the Simulating the Life-Cycle of Molecular Clouds (SILCC) project (Walch et al. 2015; Girichidis et al. 2016a), more specifically, the three different models described by Girichidis et al. (2016b): a model with only thermal stellar feedback, a model with only cosmic ray stellar feedback, and a model with both forms of feedback. Thermal stellar feedback consists of adding 10^{51} erg of energy to the gas surrounding a supernova event. This happens either as pure energy injection in regions where the Sedov–Taylor expansion of the supernova is resolved, or as a momentum injection in regions where it is not resolved. Cosmic ray feedback consists of adding 10^{50} erg of energy to the cosmic ray energy equation. This equation is evolved as an extra separate equation during the MHD integration, and couples to the hydrodynamics as an extra pressure term in the momentum and energy equations. The cosmic ray energy equation assumes a simplified transport equation based on an isotropic particle distribution function, and neglects the effect of cosmic ray streaming. None of the models we post-process includes a prescription for photoionization feedback.

The simulations themselves were run in a box of $2 \times 2 \times \pm 20$ kpc, with periodic x and y boundaries, and using an adaptive mesh with a resolution of 15.6 pc ($128 \times 128 \times 2560$ cells) in the high-resolution region. The three simulations respectively cover an evolution over 257.3, 263.7, and 256.9 Myr. In this work, we mainly focus on the results after 250 Myr of evolution, although we will also briefly discuss the time evolution of the ionization structure. It is worth noting that the Girichidis et al. (2016b) simulations do not include a self-consistent modelling of star formation using sink particles, nor early stellar feedback from radiation or stellar winds, like e.g. Peters et al. (2017).

To perform the post-processing with our radiative transfer code, we resample the density grids on a Cartesian grid of $128 \times 128 \times 256$ cells, in a box with dimensions $2 \times 2 \times \pm 2$ kpc, and with the same periodic boundaries as the simulations themselves. Extending the box vertically to larger sizes does not change our results, since the density at large vertical heights drops to negligible values. The adopted number of cells was found to provide an optimal trade-off between accuracy and computational efficiency. We tested the convergence of our results using a simulation with a higher resolution, which oversamples the original

¹ <https://github.com/bwvdbro/CMAcIOnize>

SILCC data. The average vertical column density of the gas within our box is $\sim 10 \text{ M}_{\odot} \text{ pc}^{-2}$ ($\sim 10^{21} \text{ atoms cm}^{-2}$) for all three models at 250 Myr, consistent with the average H I column density in the Milky Way (Dickey & Lockman 1990), and in line with the column densities found close to Milky Way spiral arms (Nakanishi & Sofue 2003).

The simulated density fields use a fixed solar metallicity, so we have to make some assumptions for the abundances of He and the coolants we track. We will use the same values as in Barnes et al. (2014): $\text{O}/\text{H} = 4.3 \times 10^{-4}$, $\text{N}/\text{H} = 6.5 \times 10^{-5}$ (Simpson et al. 2004; Jenkins 2009), $\text{S}/\text{H} = 1.4 \times 10^{-5}$ (Daflon et al. 2009), $\text{He}/\text{H} = 0.1$, and $\text{Ne}/\text{H} = 1.17 \times 10^{-4}$ (Mathis 2000).

2.3 Ionizing sources

The ionizing sources are luminous O and B stars. Garmany, Conti & Chiosi (1982) found an average stellar surface density of 24 kpc^{-2} for O stars in the solar neighbourhood, so we randomly sample 96 sources within our $2 \times 2 \text{ kpc}^2$ box. The x and y coordinates of the sources are uniformly sampled, while the z coordinate is sampled from a normal distribution. The scale height of the stellar disc is a parameter for our models. We use a value of 63 pc (Maíz-Apellániz 2001) for most of our models. The typical ionizing luminosity of an O star is $\sim 10^{49} \text{ s}^{-1}$ (Sternberg, Hoffmann & Pauldrach 2003); however, it is not a priori clear how much of this radiation makes it out of the dense, H II region surrounding the star and into the ISM (Wood et al. 2010). We therefore treat the fraction of the $\sim 10^{49} \text{ s}^{-1}$ luminosity that is effectively used to ionize the ISM as a second parameter of our source model.

The spectrum of the ionizing sources is based on the stellar models of Pauldrach, Hoffmann & Lennon (2001); we use the data tables compiled by Sternberg et al. (2003) to obtain the spectrum of a 40 000 K star with a surface gravity of $\log(g/(\text{cm s}^{-1})) = 3.40$, on which we linearly interpolate.

The density fields we use contain asymmetries that can potentially propagate into our radiative transfer calculation, especially when our source distribution is sampled randomly and independent of the density field. We address this issue by testing the robustness of our results against a change in the seed for the random generator used to sample our source distribution.

2.4 Simulations

In total, 52 radiative transfer simulations were run for each of the three feedback models: one reference model, 10 models to test the convergence of our model, six to test the effect of parameter changes in the ionizing source model, 10 more advanced models used to calculate emission lines, and 25 models to show the time evolution of the ionized disc. The latter used SILCC snapshots at different times throughout the simulation, while all the other simulations were run using the snapshot at 250 Myr. The different parameters and our adopted naming convention for the first two groups of models are shown in Table 1. Our reference model uses 20 iterations for a $128 \times 128 \times 256$ cell grid with 10^7 photon packets, with a source scale height of 63 pc and a per source luminosity of $4.26 \times 10^{49} \text{ s}^{-1}$, and fixed source positions (set by the fixed random seed 42). Note that simulations Ci20g128p7 and Ir42149s063 both refer to this reference model.

The more advanced models use the same parameters, but with 10^8 photon packets, to ensure converged coolant fractions.

3 RESULTS AND DISCUSSION

3.1 Convergence

Before we can test convergence of our results, we have to make clear what convergence means in our case. We are interested in the vertical scale of the neutral and ionized disc, so we require vertical disc profiles which are sufficiently converged. To quantify convergence, we compare profiles for the average ionized and neutral gas densities. These are defined as, respectively, the spatially averaged ionized and neutral gas densities in $2 \times 2 \text{ kpc}$ planar slices with constant height z above or below the disc. Examples of density profiles are shown in the top panel of Fig. 1.

For each parameter of interest, we compare profiles for different parameter values with a reference model by computing the relative difference

$$\text{RD}(X, z) = 2 \left| \frac{n(z) - n_X(z)}{n(z) + n_X(z)} \right|, \quad (1)$$

with X the parameter value that identifies the reference model. A model is considered sufficiently converged if the relative difference is of the level of 1 per cent or less for all values of z .

We start with the number of iterations of our algorithm that is needed to obtain converged densities. The bottom row of Fig. 1 shows the relative difference between the density profiles after 5, 10, 15, and 20 iterations, and the reference result after 25 iterations. The largest differences between the low iteration number models and the reference model are located near the central disc. The ionized disc between 0.5 and 2 kpc is converged to 1 per cent accuracy after 20 iterations, which is the number we will use for all our simulations.

A very important parameter in our algorithm is the number of photon packets used to discretize the stellar radiation field. If this number is too low, discretization error will dominate our simulation, and the radiation field will be unable to spread throughout the simulation box. The top rows of Fig. 1 show the average densities for different numbers of photon packets, and for a fixed grid resolution of $128 \times 128 \times 256$ cells. It is immediately clear that the results for the simulations with only 10^5 photon packets are qualitatively different from those of the other simulations, illustrating how the radiation field is unable to efficiently ionize the box if too few photon packets are used. Just, as in the case of the number of iterations, using a low number of photons also affects the scale height of the ionized disc. The relative difference drops below 1 per cent for the simulations with 10^7 photon packets, and this is the number that we will use in consecutive simulations.

A last important parameter for the convergence of our models is the grid resolution. Unlike the number of iterations or photon packets, this parameter is not immediately linked to our photoionization algorithm. This means that we can obtain a converged photoionization solution for any given grid resolution by using enough photon packets and iterations. If the grid samples the density field very badly, then this solution will not be converged to the solution imposed by the underlying density field, but from the point of view of the photoionization algorithm, it will be converged none the less. Conversely, parameters that lead to a converged result for one grid resolution will not necessarily work for another resolution. The central row of Fig. 1 shows how using a $256 \times 256 \times 512$ grid with only 10^7 photon packets seems to change our result, indicating that we need at least this resolution to obtain a converged result. However, if we rerun the same model with 10 times more photon packets, then the solution looks much more similar to the $128 \times 128 \times 256$

Table 1. Naming convention for the convergence tests and source model parameter exploration.

Simulation name	Number of iterations	Grid resolution	Number of photon packets
Ci20g128p7	20	$128 \times 128 \times 256$	10^7
Ci05g128p7	5	$128 \times 128 \times 256$	10^7
Ci10g128p7	10	$128 \times 128 \times 256$	10^7
Ci15g128p7	15	$128 \times 128 \times 256$	10^7
Ci25g128p7	25	$128 \times 128 \times 256$	10^7
Ci20g064p7	20	$64 \times 64 \times 128$	10^7
Ci20g256p7	20	$256 \times 256 \times 512$	10^7
Ci20g256p8	20	$256 \times 256 \times 512$	10^8
Ci20g128p5	20	$128 \times 128 \times 256$	10^5
Ci20g128p6	20	$128 \times 128 \times 256$	10^6
Ci20g128p8	20	$128 \times 128 \times 256$	10^8
Simulation name	Random seed	Ionizing luminosity ($\text{s}^{-1} \text{ source}^{-1}$)	Scale height (pc)
Ir42l49s063	42	4.26×10^{49}	63.0
Ir19l49s063	19	4.26×10^{49}	63.0
Ir55l49s063	55	4.26×10^{49}	63.0
Ir42l48s063	42	4.26×10^{48}	63.0
Ir42l50s063	42	4.26×10^{50}	63.0
Ir42l49s032	42	4.26×10^{49}	31.5
Ir42l49s126	42	4.26×10^{49}	126.0

result, showing that the latter resolution is sufficient for convergence. This is to be expected, as this is the resolution used by Girichidis et al. (2016b), so we are effectively overresolving the input density field with a $256 \times 256 \times 512$ grid. As pointed out by de Avillez & Breitschwerdt (2004) and Hill et al. (2012), this resolution does not guarantee converged density fields in the mid-plane of the SILCC simulations, and might cause us to overestimate the ionizing luminosity necessary to ionize the extended disc.

We conclude that 20 iterations of our algorithm, using a grid of $128 \times 128 \times 256$ cells and 10^7 photon packets, are sufficient to obtain converged vertical density profiles for the ionized disc.

3.2 Ionizing source model

Before we can discuss the effect of changing the source luminosity or scale height of the source distribution on our results, we need to quantify the scatter introduced by our sampling of the source distribution. To this end, we rerun each of our models with two different seeds for the random generator that generates the positions of our discrete sources.

The central row of Fig. 2 shows the relative difference in densities for models with different random seeds (and hence different positions for the 96 ionizing sources). Apart from a small difference in the width of the central density peak, the relative difference between the different models is of the order of 10 per cent in the neutral gas density. There is a clear asymmetry in the neutral gas disc for the SILCC model without cosmic ray feedback, which shows that the exact positions of the sources can enhance asymmetries present in the density field. However, for the ionized gas density, and especially the extended disc which is of most interest to us, the relative differences between different source models are small. We conclude that the overall structure of our solution is independent of the numerical details of our ionizing source model, so that we can use a single realization of the source model to study the structure of the ionized disc.

We already mentioned above that there is some uncertainty on the ionizing luminosity of the individual sources in our model, since we do not know how much radiation is absorbed in the local envelope surrounding each source, which we do not resolve. To address this uncertainty, we run models with different ionizing luminosities per ionizing source. The top rows of Fig. 2 show the density profiles for three different values of the luminosity: the generic value of $4.26 \times 10^{49} \text{ s}^{-1}$, a value that is 10 times lower, and a value that is 10 times higher. The latter is not really physical, but it is instructive to see how it affects our results.

It is immediately clear from the figure that a low value of the luminosity can lead to results that are qualitatively completely different: if the luminosity is too low, then there is insufficient radiation to ionize the entire box, and we end up with an extended neutral disc instead of an ionized disc. However, when the luminosity is high enough, an ionized disc is created, the structure of which is relatively independent of the total ionizing luminosity. Just as with the different random seeds above, we see that a different luminosity has some effect on the width of the central density profile, but does not really affect the extended disc (provided the luminosity is high enough to ionize it). The neutral density profile is affected, but only in the regions where the neutral gas fraction is already quite low.

The bottom row of Fig. 2 shows the relative differences for models with different values for the ionizing source distribution scale height: the generic value of 63 pc, a value which is half of that, and a value which is twice the generic value. It is clear that the scale height of the sources has a small effect on the value of the neutral fraction, as it is easier for hard ionizing radiation to escape from the dense inner disc if the source distribution is more extended.

We conclude that the extended ionized disc in our models is robust against changes in our ionizing source model, provided that the ionizing luminosity of the combined sources is sufficient to ionize the disc. These results are in agreement with previous results of Wood et al. (2010).




Figure 1. *Top row:* average neutral (full line) and ionized (dotted line) density for a given number of photon packets used in our algorithm. The shaded regions represent the scatter within slices of equal height above the disc. For clarity, the region close to the mid-plane has been plotted on a linear scale, while the outer regions are plotted on a logarithmic scale. The dashed vertical lines indicate the transition from the linear to the logarithmic region. Note that the large scatter near the mid-plane causes some of the shaded regions to have negative lower bounds, which cannot be represented on the logarithmic scale. *Other rows:* relative difference between the average densities of the indicated models and the reference model for different model parameters: *top:* the number of photon packets, *middle:* grid resolution, *bottom:* the number of iterations. The dashed horizontal lines represent the target 1 per cent convergence limit.

3.3 Time evolution

Fig. 3 shows the build-up of the density profiles over time, by post-processing snapshots at a 10 Myr interval from time 0 Myr to the final snapshot at 250 Myr we extensively discussed before, for each of the three feedback models. For the model with only thermal feedback, the profiles quickly reach a stable configuration in less than 50 Myr, for the models with cosmic ray feedback, the build-up is slower and it takes approximately 100 Myr to reach a stable disc.

The figure also shows how all three models at first have very low disc densities, which slowly increase over time. For the model with only thermal feedback, this increase is limited, while for the other two models the disc growth is more extended.

It is important to note that these results were obtained in post-processing, so there is no dynamical effect of the ionizing radiation on the evolution of the simulation model (other than the feedback model employed by the SILCC simulations themselves).




Figure 2. *Top row:* average neutral (full line) and ionized (dotted line) density for a given ionizing luminosity per ionizing source. The shaded regions represent the scatter within slices of equal height above the disc. For clarity, the region close to the mid-plane has been plotted on a linear scale, while the outer regions are plotted on a logarithmic scale. The dashed vertical lines indicate the transition from the linear to the logarithmic region. Note that the large scatter near the mid-plane causes some of the shaded regions to have negative lower bounds, which cannot be represented on the logarithmic scale. *Bottom rows:* relative difference between the average densities of the indicated models and the reference model for different model parameters: *top:* ionizing luminosity per source, *middle:* random seed used to generate source positions, *bottom:* scale height of the ionizing disc.

Furthermore, all our results used the same source model, independent of the evolution of the underlying SILCC model.

3.4 Emission lines

To calculate emission lines, we used the full version of CMACIONIZE, which treats the temperature of the ISM as a variable rather than a constant, and uses heating and cooling terms to obtain a converged temperature and ionization state for each cell in the grid. Since the abundances of most coolants are significantly lower than those of

hydrogen and helium, we increased the number of photons for these simulations with a factor of 10 to get better signal-to-noise ratios for the coolant ionic fractions. The convergence of the resulting H α profiles and forbidden line emission ratios is shown in Fig. 4 for the model with both thermal and cosmic ray feedback. We see that 10^8 photon packets are sufficient to get converged H α profiles and emission line ratios.

Fig. 5 shows simulated maps of the H α intensity for our different models. It is clear that the H α emission is significantly more extended in the simulations including cosmic ray feedback. Note that




Figure 3. Average neutral (full line) and ionized (dotted line) density profile at different times during the simulations. For clarity, the region close to the mid-plane has been plotted on a linear scale, while the outer regions are plotted on a logarithmic scale. The dashed vertical lines indicate the transition from the linear to the logarithmic region.




Figure 4. *Top:* $H\alpha$ intensity profiles as a function of height and relative difference for simulation with three different values of the photon packet number: 10^7 (p7), 10^8 (p8), and 10^9 (p9). *Bottom:* line ratios for the same simulations.

the $H\alpha$ intensities we find here are significantly higher than those found by Barnes et al. (2014) (see Section 3.4.2 below).

Fig. 6 shows horizontally averaged $H\alpha$ profiles. It is clear that the simulation with only thermal feedback has very low intensities at even moderate heights above the mid-plane, which indicates that this model does not have an extended ionized disc. By fitting an exponential profile of the form $H\alpha(z) = H\alpha_0 \exp(-|z|/z_s)$ at heights above and below 500 pc for the models with cosmic ray feedback, we can determine the $H\alpha$ scale height z_s for those models. This yields values of 0.643 and 0.659 kpc for the model with only cosmic ray feedback, and 0.674 and 0.723 kpc for the model with both thermal and cosmic ray feedback. These values are in line with observed $H\alpha$ scale heights in the Milky Way, which range from ~ 100 pc in the inner Galaxy (Madsen & Reynolds 2005), over ~ 400 pc in the near spiral arms (Haffner, Reynolds & Tuft 1999; Hill et al. 2014; Krishnarao et al. 2017) and values of ~ 700 pc in the local solar neighbourhood (Gaensler et al. 2008; Savage & Wakker 2009), to values of more than 1 kpc in the far Carina arm (Krishnarao et al. 2017). The values we find are also significantly higher than the ~ 200 pc scale heights found in Barnes et al. (2014).

Gaensler et al. (2008) propose to measure the filling factor of ionized gas by comparing $H\alpha$ scale heights with electron scale heights obtained from pulsar dispersion measurements. Assuming an ionizing gas filling factor which is constant as a function of height, the electron scale height z_{s,n_e} (which traces the electron density n_e) and $H\alpha$ scale height z_{s,n_e^2} (which scales as n_e^2) should be linked by $z_{s,n_e}/z_{s,n_e^2} = 2$. However, if these scale heights are inferred from integrated quantities, then the presence of neutral regions in the integral could skew this ratio to higher values for low altitudes. We test this for our simulations by fitting scale heights to $n_e(z)$ and




Figure 5. $H\alpha$ intensity maps for the three models at $t = 250$ Myr. The $H\alpha$ intensity is given in Rayleighs (R); 1 R corresponds to a column emissivity of 10^6 photons $\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$.




Figure 6. $H\alpha$ intensity profiles as a function of height for the three feedback models. The shaded regions show the scatter within planar regions of equal scale height. The dashed lines show exponential fits to the part of the curve with $|z| > 500$ pc (we fitted the parts below and above the disc separately). The scale heights derived from these fits are indicated in the legend. The black dashed line indicates the sensitivity limit for the Wisconsin $H\alpha$ Mapper (WHAM), a representative instrument used to observe emission lines in the Milky Way disc. Note that the large scatter near the mid-plane causes some of the shaded regions to have negative lower bounds, which cannot be represented on the logarithmic scale.

$n_e^2(z)$, as shown in Fig. 7. As can be seen, the z_{s,n_e^2} values we find are in good agreement with the integrated $H\alpha$ values. Furthermore, the z_{s,n_e} values we find are consistent with a ratio of 2.

When we look at emission line ratios, the correspondence between model and observations is much less favourable however.




Figure 7. Average electron density n_e (top) and squared electron density n_e^2 (bottom) as a function of height above the disc. The red line shows the average value within slices of equal height, the shaded region shows the scatter within each slice. The dashed lines indicate exponential fits to the part of the curve with $|z| > 500$ pc (we fitted the parts below and above the disc separately). The scale heights derived from these fits are indicated in the legend.




Figure 8. Line ratios for the simulations with and without cosmic ray heating. The different lines correspond to different values of the cosmic ray heating factors C_{cr} or C'_{cr} , as indicated in the legend. The blue line corresponds to the line ratios of the reference model without cosmic ray heating. All values have been binned in 50 logarithmically spaced bins. The shaded regions show the scatter within the bins. The black dashed line is the WHAM sensitivity limit, while the blue dots represent observational data from the WHAM survey (Haffner et al. 1999). Note that the large scatter for some $H\alpha$ intensities causes the shaded regions to have negative lower bounds, which cannot be represented on the logarithmic scale.

Fig. 8 shows the $[N \text{ II}]/H\alpha$ and $[S \text{ II}]/H\alpha$ ratios as a function of the $H\alpha$ intensity. We notice a clear upwards trend for both line ratios for increasing $H\alpha$ intensity. This trend is completely opposite to the downwards trend found in observations (Haffner et al. 1999), and agrees with earlier results found by Barnes et al. (2014). Since $[N \text{ II}]/H\alpha$ and $[S \text{ II}]/H\alpha$ trace warm ionized gas, the lack of high line ratios for low $H\alpha$ intensities indicates that the temperature at high altitudes above the disc (where the $H\alpha$ intensity is low) is too low in our models. This can also be seen in Fig. 9, where the equilibrium gas temperature is shown as a function of altitude. The equilibrium temperature in our models only increases very slowly across the ionized disc. If we want our models to line up with the observations, we hence need an extra heating mechanism that can increase the temperature at high altitudes.

3.4.1 Cosmic ray heating

One possible source of heating would be heating by cosmic rays. The cosmic ray model employed by the SILCC simulations introduces an extra pressure term and associated adiabatic heating and cooling due to cosmic rays, but does not constitute a direct heating term (Girichidis et al. 2016b). This means we cannot directly use the simulation snapshots to provide us with the required additional heating. We know however that cosmic rays can directly transfer energy to the gas by the generation of dampened Alfvén waves

(Wentzel 1971), a process which scales with $n_e^{-1/2}$, where n_e is the electron density in the gas (Wiener et al. 2013).

To test if cosmic ray heating can explain the observed line ratios, we run additional photoionization simulations with the advanced model that treats both the ionization balance and the temperature as a variable, and includes an extra heating term that scales with $n_e^{-1/2}$:

$$H_{\text{cr}} = (1.2 \times 10^{-27} \text{ erg cm}^{-9/2} \text{ s}^{-1}) C_{\text{cr}} n_e^{-1/2} e^{-|z|/h_{\text{cr}}}, \quad (2)$$

where C_{cr} is a constant parameter of our model, z is the height above the disc, h_{cr} is the scale height of the cosmic ray heating, and the numerical factor was based on the factor given in Wiener et al. (2013). As in Wiener et al. (2013), we choose a value $h_{\text{cr}} = 1.333 \text{ kpc}$. We will consider different values of the parameter C_{cr} .

We also implemented a second cosmic ray heating scheme that is based on the more exact expression for the cosmic ray heating (Wiener et al. 2013):

$$H_{\text{cr}} = v_A \cdot \nabla P_{\text{cr}}, \quad (3)$$

where v_A is the Alfvén speed ($v_A = \frac{B}{\sqrt{\mu_0 m_p n_e}}$, with B the magnetic field, μ_0 the vacuum permeability, and m_p the proton mass), and P_{cr} is the cosmic ray pressure [$P_{\text{cr}} = (\gamma_{\text{cr}} - 1)e_{\text{cr}}$, with $\gamma_{\text{cr}} = 4/3$ the cosmic ray polytropic index, and e_{cr} the cosmic ray energy density]. B and e_{cr} are present in the SILCC outputs, so our alternative heating




Figure 9. *Top:* average temperature profiles for the simulations with and without cosmic ray feedback. The coloured lines are the results for different values of the cosmic ray heating factors C_{cr} or C'_{cr} , as indicated in the legend. The blue line $C_{\text{cr}} = 0$ is the reference model without cosmic ray heating. The shaded regions show the scatter within slices of equal scale height. *Bottom:* temperature as a function of ionized density for the same simulations. The results have been binned in 50 logarithmically spaced bins; the shaded regions represent the scatter within the bins. The vertical dashed lines represent the density threshold below which the gas will likely be shock heated, an effect which we do not take into account in our treatment.

term then can be formulated as

$$H_{\text{cr}} = (1.2 \times 10^{-27} \text{ erg cm}^{-9/2} \text{ s}^{-1}) C'_{\text{cr}} \frac{\mathbf{B} \cdot \nabla e_{\text{cr}}}{\sqrt{n_e}}, \quad (4)$$

with C'_{cr} a constant parameter. This heating term will be trivially zero for the SILCC model without cosmic ray feedback.

Since the heating terms described above diverge for very low values of the electron density, we have to be careful when adding them to our temperature calculation, as low electron densities do not only occur in ionized cells with low densities, but also in denser, neutral cells, where the electron density is low due to the high neutral fraction of the gas. We therefore, before entering the complex temperature iteration for a cell, compute the neutral fraction at a fixed temperature of 8000 K for that cell. If the cell has a neutral fraction that is higher than a cut-off value of 0.75, we assume the cell is fully neutral and set its temperature to 500 K. In this case, we do not enter the detailed cooling and heating routine, and avoid problems with finding the neutral fraction in cells with high densities and low electron densities.

Similarly, we also need to take care when heating the cells to very high temperatures, as we only included the relevant heating




Figure 10. Hot gas filling factors for the simulations with cosmic ray heating. The different lines correspond to different values of the cosmic ray heating factor C_{cr} , as indicated in the legend. The blue line corresponds to the hot gas filling factor found in the SILCC snapshots themselves.

and cooling processes up to temperatures of about 30 000 K. We therefore manually reset the temperature to 30 000 K if cosmic ray heating pushes it to higher values.

Fig. 9 shows the temperature profiles for various strengths of the cosmic ray heating. It is clear that cosmic ray heating is able to push the average temperature to higher values. However, it is also worth noting that a significant fraction of the gas ends up at 30 000 K, the imposed temperature limit, especially for cells at higher altitudes above the disc. These cells will have very low $\text{H}\alpha$ intensities, so they do not contribute significantly to the observed line ratios. Fig. 10 shows the hot gas filling factor in each horizontal slice of the simulation, defined as the fraction of the 128×128 cells in the slice with a temperature of 30 000 K. For a comparison, we also show the hot gas filling factor in the actual SILCC snapshots. It is clear that the models with high cosmic ray feedback contain a lot of cells with very high temperatures. To reproduce the hot gas filling factor of the SILCC simulations in the models with cosmic ray feedback, we would need a cosmic ray heating parameter with about 10 percent of the value advocated in Wiener et al. (2013), or the more advanced cosmic ray heating model with the highest parameter value.

Note that both cosmic ray heating models produce similar results. As can be seen from the bottom panel of Fig. 9, both mechanisms predominantly heat low-density ionized gas, and cause a sharp increase in temperature at a density threshold set by the value of the heating parameter.

Also note that gas with densities lower than $\sim 10^{-2} \text{ cm}^{-3}$ will most likely be shock heated and ionized to temperatures $> 10^5 \text{ K}$, an effect that we did not incorporate in our post-processing treatment. Therefore, we find that all but the strongest cosmic ray heating realizations only heat gas that would already be heated by other physical mechanisms.

Finally, Fig. 8 shows the line ratios for the different cosmic ray heating models. It is clear that the cosmic ray heating does not affect the high $\text{H}\alpha$ intensity end of the curve, which corresponds to radiation emitted by the central disc. Cosmic ray heating does affect the low-intensity line ratios in the diffuse ionized disc, and effectively changes the sign of the slope of the curve, leading to a trend that is more in line with the observed trend, although the effect is only noticeable for high values of the cosmic ray heating factor. This seems to indicate that the cosmic ray heating does not




Figure 11. Average temperature and $H\alpha$ intensity as a function of height above the disc, temperature as a function of ionized density, and $[\text{NII}]/H\alpha$ and $[\text{SII}]/H\alpha$ line strengths as a function of $H\alpha$ intensity for the models with different values for the total ionizing luminosity (Q_{ref} represents the reference model with an ionizing luminosity of $4.26 \times 10^{49} \text{ s}^{-1}$ per ionizing source). The full lines represent the binned average values as indicated on the legend, the shaded regions represent the scatter within the bins. The dashed lines in the middle row represent the density threshold below which the gas will likely be shock heated, an effect we do not take into account in our treatment. The dashed lines in the bottom two rows represent the sensitivity limit of the WHAM instrument, while the blue dots represent observational data from the WHAM survey (Haffner et al. 1999). Note that the large scatter for some scale heights and $H\alpha$ intensities causes the shaded regions to have negative lower bounds, which cannot be represented on the logarithmic scale.

Table 2. Scale height of the H α disc for the simulations with lower total ionizing luminosities.

Ionizing luminosity ($\text{s}^{-1} \text{ source}^{-1}$)	Scale height $z < 0$ (kpc)	Scale height $z > 0$ (kpc)
TH1_CR0		
4.26×10^{49}	0.931	2.064
1.065×10^{49}	0.964	2.065
4.26×10^{48}	0.980	2.077
TH0_CR1		
4.26×10^{49}	0.659	0.643
1.065×10^{49}	0.662	0.641
4.26×10^{48}	0.352	0.051
TH1_CR1		
4.26×10^{49}	0.723	0.674
1.065×10^{49}	0.204	0.585
4.26×10^{48}	0.055	0.056

affect the temperature of the intermediate density ionized gas which is responsible for the observed H α emission.

3.4.2 Lower total luminosity

We require a relatively high luminosity to ionize out to large scale heights, compared to the values that were used by Barnes et al. (2014). This leads to H α luminosities that are high compared to observed H α luminosities. It is instructive to see how the temperature and line ratios change if a lower luminosity is used.

To this end, we run a set of simulations with the full version of the code, using lower luminosities for the ionizing sources: a version with an ionizing luminosity of $4.26 \times 10^{48} \text{ s}^{-1}$ per ionizing source (10 per cent of the generic value), and a version with a luminosity of $1.065 \times 10^{49} \text{ s}^{-1}$ per ionizing source (25 per cent of the generic value).

The results of these simulations are shown in Fig. 11. As expected, the lower luminosities lead to more noisy temperature and H α profiles, as the sources are no longer able to ionize the entire low-density disc. However, for the SILCC models including cosmic ray feedback, the H α scale height for the simulation with 25 per cent of the total ionizing luminosity is still comparable to the full luminosity result (as can be seen from Table 2), and hence in line with observed scale heights in the Milky Way. Furthermore, the runs with lower intensity show a clear increase in temperature for intermediate density ionized gas, which in turn translates into line ratios that increase with decreasing H α intensity.

We hence find that a decrease in total ionizing luminosity can actually lead to higher local temperatures and can hence explain the observed line ratios. This counter-intuitive result can be explained by the hardening of the ionizing spectrum for large scale heights (Wood & Mathis 2004): low-frequency photons are preferentially absorbed at low heights, so that the spectrum that reaches higher heights contains relatively more high-frequency hydrogen ionizing photons, which heat the gas to higher temperatures. If the total ionizing luminosity is high, the fraction of photons absorbed at low heights will be low, and the hardening of the spectrum will happen over a large scale. However, if the luminosity is just enough to ionize out to the boundary of the simulation box, we see the full hardening of the spectrum and resulting rise in temperature, while still keeping the gas ionized. There is however some fine tuning required: decreasing the total luminosity further to 10 per cent of

the generic value significantly decreases the H α scale length, and leads to results that are no longer in line with observations.

4 CONCLUSION

In this work, we post-processed star-forming disc simulations of the SILCC project with the new Monte Carlo radiative transfer code CMACIONIZE. We showed that the more extended galactic disc present in MHD simulations that include cosmic ray feedback naturally leads to larger scale heights for the H α emission from the DIG in these discs, provided that enough ionizing radiation is present to ionize out to large scale heights. This more extended disc is a direct consequence of the specific way cosmic rays couple to the hydrodynamics of the ISM, which makes it possible to more efficiently heat the gas in a larger region compared to thermal feedback.

However, most of our simulations are unable to reproduce observed nitrogen and sulphur line ratios, even when additional heating due to cosmic rays is included in the model, since these models only affect the temperature of low-density ionized gas that does not contribute to observed line ratios. Only if we use a total ionizing luminosity that is marginally sufficient to ionize the DIG, we are able to increase the temperature of intermediate density ionized gas and reproduce observed Milky Way line ratios. This fine tuning can have several explanations:

(i) The ionizing radiation itself is partially responsible for setting the scale height of the ionized gas. This would naturally lead to a correlation between the total ionizing luminosity and the scale height. However, we cannot study this effect in this work, as there is no direct dynamical link between the SILCC simulations and our post-processing tool. We plan to repeat our analysis for the Peters et al. (2017) simulations that do include photoionization feedback.

(ii) There is another physical heating mechanism that is responsible for the extra heating that is necessary to explain observed line ratios, and that affects intermediate rather than low-density ionized gas. In this case, we would not need to fine tune the ionizing luminosity.

(iii) The dense gas surrounding young O stars absorbs about 75 per cent of the ionizing radiation from the star, so that only 25 per cent is left to ionize the more diffuse surrounding gas. As we do not resolve the densest gas, this is an effect that is likely. But to find more accurate values of the ionizing escape fraction, more detailed simulations of star-forming clouds are necessary.

We hence have no satisfactory explanation for the observed line ratios in the DIG of the Milky Way, and leave this for future work. We do find that the observed decreasing line ratios of nitrogen and sulphur with increasing H α intensity are not necessarily explained by an increasing average ISM temperature with increasing scale height, as is often assumed. Instead, these line ratios could also be consistent with an increase in temperature dispersion with increasing scale height, with an overall constant or decreasing average temperature.

ACKNOWLEDGEMENTS

We want to thank the anonymous referee for constructive and insightful remarks that significantly improved the quality of this work. BV and KW acknowledge support from STFC grant ST/M001296/1. PG acknowledges support from the DFG Priority Program 1573 Physics of the Interstellar Medium as well as funding from the European Research Council under ERC-CoG grant CRAGSMAN-646955. ASH is supported by NASA through

grant number HST-AR-14297 from Space Telescope Science Institute, which is operated by AURA, Inc. under NASA contract NAS 5-26555, to Haverford College and by NASA ATP grant NNX17AH80G to Space Science Institute.

REFERENCES

- Barnes J. E., Wood K., Hill A. S., Haffner L. M., 2014, *MNRAS*, 440, 3027
- Berkhuijsen E. M., Fletcher A., 2008, *MNRAS*, 390, L19
- Daflon S., Cunha K., de la Reze R., Holtzman J., Chiappini C., 2009, *AJ*, 138, 1577
- de Avillez M. A., Breitschwerdt D., 2004, *A&A*, 425, 899
- Dickey J. M., Lockman F. J., 1990, *ARA&A*, 28, 215
- Elmegreen B. G., 1997, *ApJ*, 477, 196
- Gaensler B. M., Madsen G. J., Chatterjee S., Mao S. A., 2008, *Publ. Astron. Soc. Aust.*, 25, 184
- Garman C. D., Conti P. S., Chiosi C., 1982, *ApJ*, 263, 777
- Gatto A. et al., 2017, *MNRAS*, 466, 1903
- Girichidis P. et al., 2016a, *MNRAS*, 456, 3432
- Girichidis P. et al., 2016b, *ApJ*, 816, L19
- Haffner L. M., Reynolds R. J., Tufte S. L., 1999, *ApJ*, 523, 223
- Haffner L. M. et al., 2009, *Rev. Mod. Phys.*, 81, 969
- Hill A. S., Benjamin R. A., Kowal G., Reynolds R. J., Haffner L. M., Lazarian A., 2008, *ApJ*, 686, 363
- Hill A. S., Joung M. R., Mac Low M.-M., Benjamin R. A., Haffner L. M., Klingenberg C., Waagan K., 2012, *ApJ*, 750, 104
- Hill A. S., Benjamin R. A., Haffner L. M., Gostisha M. C., Barger K. A., 2014, *ApJ*, 787, 106
- Jenkins E. B., 2009, *ApJ*, 700, 1299
- Krishnarao D., Haffner L. M., Benjamin R. A., Hill A. S., Barger K. A., 2017, *ApJ*, 838, 43
- Madsen G. J., Reynolds R. J., 2005, *ApJ*, 630, 925
- Maíz-Apellániz J., 2001, *AJ*, 121, 2737
- Mathis J. S., 2000, *ApJ*, 544, 347
- Nakanishi H., Sofue Y., 2003, *PASJ*, 55, 191
- Pauldrach A. W. A., Hoffmann T. L., Lennon M., 2001, *A&A*, 375, 161
- Peters T. et al., 2017, *MNRAS*, 466, 3239
- Reynolds R. J., 1990, *ApJ*, 348, 153
- Reynolds R. J., Haffner L. M., Tufte S. L., 1999, *ApJ*, 525, L21
- Savage B. D., Wakker B. P., 2009, *ApJ*, 702, 1472
- Simpson J. P., Rubin R. H., Colgan S. W. J., Erickson E. F., Haas M. R., 2004, *ApJ*, 611, 338
- Sternberg A., Hoffmann T. L., Pauldrach A. W. A., 2003, *ApJ*, 599, 1333
- Vandenbroucke B., Wood K., 2018, *Astron. Comput.*, 23, 40.
- Walch S. et al., 2015, *MNRAS*, 454, 238
- Wentzel D. G., 1971, *ApJ*, 163, 503
- Wiener J., Zweibel E. G., Oh S. P., 2013, *ApJ*, 767, 87
- Wood K., Mathis J. S., 2004, *MNRAS*, 353, 1126
- Wood K., Mathis J. S., Ercolano B., 2004, *MNRAS*, 348, 1337
- Wood K., Haffner L. M., Reynolds R. J., Mathis J. S., Madsen G., 2005, *ApJ*, 633, 295
- Wood K., Hill A. S., Joung M. R., Mac Low M.-M., Benjamin R. A., Haffner L. M., Reynolds R. J., Madsen G. J., 2010, *ApJ*, 721, 1397

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.